MASTER'S THESIS

Master of Science in Advanced Mathematics and Mathematical Engineering

Title: A spectral approach to Szemerédi's Regularity Lemma

Author: Marta Altarriba Fatsini

Advisor: Oriol Serra

Department: Mathematics

Academic year: 2020/2021

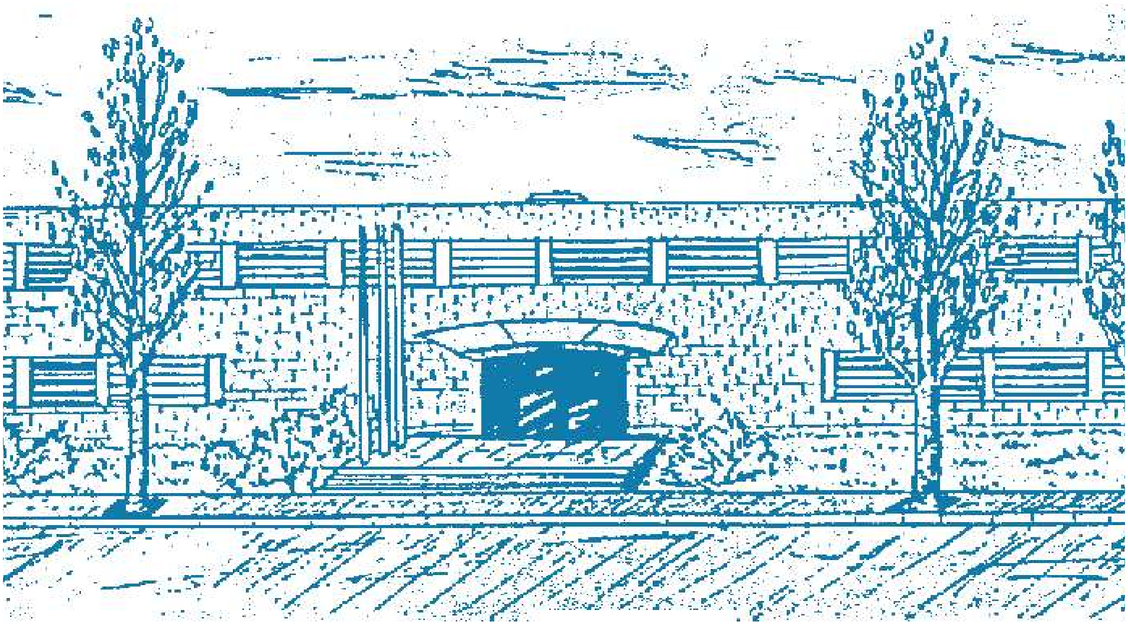

UNIVERSITAT POLITÈCNICA DE CATALUNYA
BARCELONATECH
Facultat de Matemàtiques i Estadística

Universitat Politècnica de Catalunya
Facultat de Matemàtiques i Estadística

Master in Advanced Mathematics and Mathematical Engineering
Master degree thesis

A spectral approach to Szemerédi's Regularity Lemma

Marta Altarriba Fatsini

Supervised by Oriol Serra

June 2021

Thanks to my advisor Oriol Serra for his support and supervision, to my friends Roser, Andreu and Roger for sharing their time with me, and to Martí for his patience and love.

Abstract

Szemerédi's Regularity Lemma was introduced by Endré Szemerédi [15] in his proof of the Erdős conjecture on arbitrary long arithmetic progressions in dense subsets of the integers. It has become one of the most powerful tools on Mathematics, especially in Graph Theory and Combinatorics. Roughly speaking, the lemma says that for any graph there is a partition of the vertices into a bounded number of parts such that edges between most different parts behave almost randomly. Recently, Tao gave a spectral version of the Regularity Lemma which originated on work of Frieze and Kannan which applies to self–adjoint operators. Its application to adjacency matrices provides a spectral proof of Szemerédi's Regularity Lemma.

This thesis has two main purposes. The first one is to discuss in detail the spectral proof and the decomposition of the adjacency matrix used to describe the partition. The second one is to study the natural extension of the notion of regularity and the Regularity Lemma itself for self–adjoint matrices. The associated Counting and Removal Lemmas are also discussed.

Keywords

Szemerédi's Regularity Lemma, Spectral Graph Theory.

Contents

Introduction

Szemerédi's Regularity Lemma was introduced by Endré Szemerédi in his proof of his theorem on arbitrary long arithmetic progressions in dense subsets of the integers. It has become one of the most powerful tools on Mathematics, especially in Graph Theory and Combinatorics, see for instance the survey of Komlós, Shokoufandeh, Simonovits and Szemerédi [11].

To properly state the theorem, we require some background on regularity for graphs. Let G be a graph on n vertices and fix $\epsilon > 0$. We say that a pair of subsets of vertices $X, Y \subset V$ is *ϵ–regular* if for all subsets $A \subset X$ and $B \subset Y$ such that $|A| \geq \epsilon |X|$ and $|B| \geq \epsilon |Y|$, we have that

$$|d(A, B) - d(X, Y)| \leq \epsilon,$$

where

$$d(X, Y) := \frac{e(X, Y)}{|X| |Y|}$$

is the *edge–density* of the pair (X, Y), namely the number of edges with one end in X and the other in Y over the maximum possible number of edges between them. A partition of the vertices $\mathcal{P} = \{V_1, \dots, V_k\}$ is *ϵ–regular* if

$$\sum_{\substack{(i,j)\in[k]^2 \\ (V_i,V_j) \text{ not } \epsilon\text{–regular}}} |V_i| |V_j| \leq \epsilon |V|^2 ,$$

that is to say, almost all pairs are ϵ–regular.

Roughly speaking, an ϵ–regular partition of the graph G is such that the edges between almost all parts are well–distributed, or in other words, the edges between most different parts behave random–like. The main theorem can be stated as follows.

Theorem (Szemerédi's Regularity Lemma)**.** For every $\epsilon > 0$, there exists a constant $M = M(\epsilon)$ such that every graph has an ϵ–regular partition into at most M parts.

The power of the above theorem relies in the fact that the number of parts of the ϵ–regular partition only depends on the constant ϵ and as a consequence, is the same reglardless of the size of the graph we are considering. Moreover, the original proof given by Szemerédi in [15] is quite straightforward and does not use any sophisticated mathematical techniques. The bound on M given by such proof is considerably large, and there have been several improvements on the optimalisation of M such as Gowers (see [9]), which however shows that it can be as large as a power tower of two's of height proportional to $\log(\epsilon^{-1})$.

Frieze and Kannan introduced an initial version of what will be the spectral interpretation of the Regularity Lemma in [8]. Their main motivation was on the algorithmic side as regular partitions can be applied to tackle many hard problems in Graph Theory such as the Max Cut or the Graph Bisection problems. In that exposition, they come up with a polynomial algorithm to obtain regular partitions based on a spectral decomposition of the adjacency matrix of the graph. This approach indeed provides approximate algorithms to several such problems and has had a huge impact in this area.

The spectral version of the Regularity Lemma appeared as an entry from 2012 in the blog of Tao and no paper version is available. The entry is given in the usual expository style of Tao and leaves room for a number of technical details. A more precise description appears in the preprints by Cioaba and Martin [7] and by Robertson [12] that were posted in Arxiv with no intention of being published in a refereed journal. A published paper by Szegedy [14] gives explicitely the statement of the Spectral Regularity Lemma, focusing in the much more abstract setting of the so–called kernel operators in probability spaces and its applications to the theory of graph limits.

The potential applications of the spectral version of the Regularity Lemma, besides the ones by Szegedy to graph limits, are not developed in either of the above mentioned references. These include the standard applications to the Graph Counting and Removal Lemmas. It is worth noticing here the recent development of removal lemmas for matrices by Alon and Ben–Eliezer [2] which do not use the spectral approach and focuses on the appearance of a given submatrix, a casa that can not be captured by spectral techniques which are invariant by similarity relations. Another line of applications concerns the description of regular partitions for specific families of graphs for which the spectra and orthogonal basis of eigenvectors are known.

The main purpose of this Master Thesis is to understand the spectral proof of Szemerédi's Regularity Lemma suggested by Tao and to study the additional information that the spectral partition provides compared to the one given by the classical proof. The structure of the present dissertation is the following.

In the introductory chapter, we introduce the notion of ϵ–regularity and then we give a first proof of Szemerédi's Regularity Lemma. To appreciate the importance of this tool, we give a few examples of a practical usage of the Regularity Lemma: the Graph Counting and Removal Lemmas, which are motivated by Roth's Theorem on 3–term arithmetic progressions, and the Erdős–Stone–Simonovits Theorem, a main result in Extremal Graph Theory.

In Chapter 2, a classical context in Spectral Graph Theory is given in order to provide the fundamental concepts that will be used throughout the thesis. Among other results, we bring attention to Courant–Fischer inequalities and the Mixing Lemma, and remark how are they related to pseudorandomness and regularity.

In Chapter 3, we prove the main results of the thesis. Although we are interested in a spectral proof of Szemerédi's Regularity Lemma for graphs, we begin with a few regularity theorems for self–adjoint matrices, in particular:

- The first theorem we prove takes $\epsilon > 0$ and a self–adjoint $n \times n$ matrix T with coefficients in $\mathbb{C}$ such that $\mathrm{Tr}(T^2) \leq n^2$, and using the spectrum of T and a basis of eigenvectors, provides a decomposition of T into three self–adjoint matrices $T = T_1 + T_2 + T_3$, a partition $\{V_1, \ldots, V_M\}$ of $[n]$ and an exceptional set of pairs Σ in $[M]^2$. This partition enjoys some particular properties for pairs (V_s, V_t): an upper bound for the sum of products $|V_s|\,|V_t|$ for $(s,t) \in \Sigma$, and bounds for each matrix T_1, T_2 and T_3 for pairs $(s,t) \notin \Sigma$.

- The following result takes the decomposition, the partition and the set of pairs Σ from the above theorem and reformulates the properties that they satisfy into what will be defined later as ϵ–regularity for matrices. This statement will be refered to as the Spectral Regularity Theorem, since it is due to spectral arguments on the matrix T and is the result from which Szemerédi's Regularity Lemma directly follows.

- The two last matricial regularity results apply to sets of matrices. The first one takes r self–adjoint $n \times n$ matrices such that their sum T satisfies $\mathrm{Tr}(T^2) \leq n^2$ and allows us to find a partition which is ϵ–regular for each of them at the same time. The second one takes again a set of r matrices and also results in an ϵ–regular partition for all of them simultaneously but in this case we impose the condition on the trace to each of the matrices.

Also in Chapter 3, we have the reformulation of those matricial results to graph statements. In particular, we have the spectral proof of Szemerédi's Regularity Theorem, edge–coloring and edge–multicoloring regularity theorems given by the results on sets of matrices mentioned before, and two weaker results for multigraphs and directed graphs. At the end of the chapter we find a few examples of the spectral partition obtained when applying those theorems to widely–studied graphs such as the complete graph (for which the result is trivial) or Cayley graphs.

In the last chapter, the Graph Removal Lemma from Chapter 1 is recovered and extended to a multigraph version using one of the regularity theorems from the third chapter. These applications were one of the main motivations of the present work.

Chapter 1

Szemerédi's Regularity Lemma

Szemerédi's Regularity Lemma is a powerful tool in the study of large dense graphs which allows us to partition the set of vertices of any graph into a bounded number of parts so that the edges between most different parts behave random–like. This first chapter is split into two sections: the statement and proof of Szemerédi's Regularity Lemma together with the essential background on ϵ–regular partitions, and then classic applications of that result, the Graph Removal Lemma and the Erdős–Stone–Simonovits Theorem. The chapter follows the explanations in the notes of Zhao, see [19].

1.1 Statement and proof

Let $G = (V, E)$ be a simple graph and let $X, Y \subset V$. Denote by $E(X, Y) \subset E$ the set of edges with one end in X and the other in Y, and let $e(X, Y) = |E(X, Y)|$.

Definition 1.1. The *edge–density* of X and Y is the number of edges between X and Y in G over the maximum number of edges possible, that is

$$d(X, Y) := \frac{e(X, Y)}{|X|\,|Y| - |X \cap Y|}.$$

This quantity is a real number between $[0, 1]$. In the case $X = Y$, we can simplify notation with $E(X) = E(X, X)$ and $e(X) = |E(X)|$, and compute the density $d(X)$ as

$$d(X) = d(X, X) = \frac{e(X)}{|X|\,(|X| - 1)}.$$

When $X = V$, we refer to the quantity $d(G) = d(V)$ as the *density* of the graph G.

Definition 1.2. The pair (X, Y) is *ϵ–regular* if for all subsets $A \subset X$ and $B \subset Y$ with $|A| \geq \epsilon\,|X|$ and $|B| \geq \epsilon\,|Y|$, we have

$$|d(A, B) - d(X, Y)| \leq \epsilon.$$

That is, for all subsets of X and Y large enough, their edge–density is essentially $d(X, Y)$. On the other hand, if a pair (X, Y) is not ϵ–regular, then there exist $A \subset X$ and $B \subset Y$ with $|A| \geq \epsilon\,|X|$ and $|B| \geq \epsilon\,|Y|$ such that $|d(A, B) - d(X, Y)| > \epsilon$. In this situation, we say that the pair (A, B) witnesses the irregularity of (X, Y).

Definition 1.3. A partition $\mathcal{P} = \{V_1, \dots, V_k\}$ of the vertex set V of a graph G is *ϵ–regular* if

$$\sum_{\substack{(i,j)\in[k]^2 \\ (V_i,V_j) \text{ not } \epsilon\text{–regular}}} |V_i|\,|V_j| \le \epsilon\,|V|^2 .$$

Informally speaking, if we have an ϵ–regular partition of a graph then the edges are evenly distributed among ϵ–regular pairs, or equivalently, the edges between most different parts behave almost randomly. Szemerédi's Regularity Lemma states that for every sufficiently large graph, we can find an ϵ–regular partition with an upper bounded number of parts. The power of this result relies in the fact that such bound is independent of the size of the graph, because for large dense graphs the size of some of the parts are also considerably large. In that case, the random–like performance between ϵ–regular pairs reminds us of the one expected in a truly random bipartite graph where edges are chosen with some fixed probability.

Theorem 1.4 (Szemerédi's Regularity Lemma)**.** *For every $\epsilon > 0$, there exists a constant $M = M(\epsilon)$ such that every graph has an ϵ–regular partition into at most M parts.*

The proof of the above theorem that we are about to see was given in 1978 by Szemerédi [15]. It consists in taking a trivial initial partition and while the partition is not ϵ–regular, find irregular pairs and refine considering the subsets which are precisely witnesses of the irregularity. Before going into the proof, we require a few lemmas which will guarantee that we indeed arrive to an ϵ–regular partition in a finite number of iterations. To do so, we use what is called an "energy increment" argument: we define an upper bounded quantity associated to a partition, called energy, and see that it increases by a fixed constant every time we refine the partition.

Let $U, W \subset V$ be sets of vertices of G and let $\mathcal{P}_U = \{U_1, \dots, U_s\}$ and $\mathcal{P}_W = \{W_1, \dots, W_t\}$ be partitions of U and W respectively.

Definition 1.5. Let

$$q(U,W) := \frac{|U|\,|W|}{|V|^2} d(U,W)^2$$

and

$$q(\mathcal{P}_U, \mathcal{P}_W) := \sum_{i=1}^{s} \sum_{j=1}^{t} q(U_i, W_j).$$

The *energy* of a partition $\mathcal{P} = \{V_1, \dots, V_k\}$ of the vertices of a graph G is

$$q(\mathcal{P}) := q(\mathcal{P}, \mathcal{P}) = \sum_{i=1}^{k} \sum_{j=1}^{k} q(V_i, V_j) = \sum_{i=1}^{k} \sum_{j=1}^{k} \frac{|V_i|\,|V_j|}{|V|^2} d(V_i, V_j)^2.$$

Note that it is a real number in $[0,1]$, since

$$q(\mathcal{P}) = \sum_{i=1}^{k} \sum_{j=1}^{k} \frac{|V_i|\,|V_j|}{|V|^2} d(V_i, V_j)^2 \le \sum_{i=1}^{k} \sum_{j=1}^{k} \frac{|V_i|\,|V_j|}{|V|^2} = \frac{1}{|V|^2} \sum_{i=1}^{k} |V_i| \sum_{j=1}^{k} |V_j| = 1.$$

Lemma 1.6. $q(\mathcal{P}_U, \mathcal{P}_W) \geq q(U, W)$.

Proof. Choose vertices $u \in U$ and $w \in W$ independent and uniformly. Define the random variable Z to be the edge–density of parts U_i and W_j containing u and w. Clearly, the probability of choosing $u \in U_i$ and $w \in W_j$ is $\frac{|U_i||W_j|}{|U||W|}$. Hence, the expectation of Z is

$$\mathbb{E}[Z] = \sum_{i=1}^{s}\sum_{j=1}^{t} \frac{|U_i|\,|W_j|}{|U|\,|W|} d(U_i, W_j) = \frac{1}{|U||W|}\sum_{i=1}^{s}\sum_{j=1}^{t} e(U_i, W_j) = \frac{e(U, W)}{|U||W|} = d(U, W),$$

and

$$\mathbb{E}[Z]^2 = d(U, W)^2 = \frac{|V|^2}{|U|\,|W|} q(U, W).$$

Similarly, the expectation of Z^2 is

$$\mathbb{E}[Z^2] = \sum_{i=1}^{s}\sum_{j=1}^{t} \frac{|U_i|\,|W_j|}{|U|\,|W|} d(U_i, W_j)^2 = \frac{|V|^2}{|U|\,|W|} q(\mathcal{P}_U, \mathcal{P}_W).$$

By properties of the expectation, $\mathbb{E}[Z^2] \geq \mathbb{E}[Z]^2$ and the lemma holds. □

Lemma 1.7. *If $\mathcal{Q}$ is a refinement of $\mathcal{P}$, then $q(\mathcal{Q}) \geq q(\mathcal{P})$.*

Proof. Let $\mathcal{P} = \{V_1, \ldots, V_k\}$ and $\mathcal{Q} = \{\mathcal{Q}_{V_1}, \ldots, \mathcal{Q}_{V_k}\}$, where $\mathcal{Q}_{V_i}$ is the refinement of V_i in $\mathcal{Q}$. Then, applying Lemma 1.6 to each pair (V_i, V_j) we obtain

$$q(\mathcal{Q}) = \sum_{i=1}^{k}\sum_{j=1}^{k} q(\mathcal{Q}_{V_i}, \mathcal{Q}_{V_j}) \geq \sum_{i=1}^{k}\sum_{j=1}^{k} q(V_i, V_j) = q(\mathcal{P}),$$

where we expressed the energy of $\mathcal{Q}$ in terms of its subpartitions. □

Lemma 1.8. *If the pair (U, W) is not ϵ–regular as witnessed by $U_1 \subset U$ and $W_1 \subset W$, then*

$$q\left(\{U_1, U \setminus U_1\}, \{W_1, W \setminus W_1\}\right) > q(U, W) + \epsilon^4 \frac{|U|\,|W|}{|V|^2}.$$

Proof. Define a random variable Z as in the proof of Lemma 1.6, taking $\{U_1, U \setminus U_1\}$ and $\{W_1, W \setminus W_1\}$ as partitions. The variance of Z is

$$\mathrm{Var}(Z) = \mathbb{E}[Z^2] - \mathbb{E}[Z]^2 = \frac{|V|^2}{|U|\,|W|}\left(q\left(\{U_1, U \setminus U_1\}, \{W_1, W \setminus W_1\}\right) - q(U, W)\right).$$

Observe that $|Z - \mathbb{E}[Z]| = |d(U_1, W_1) - d(U, W)|$ when $u \in U_1$ and $w \in W_1$, which happens with probability $\frac{|U_1||W_1|}{|U||W|}$. Then,

$$\mathrm{Var}(Z) = \mathbb{E}\Big[\,(Z - \mathbb{E}[Z])^2\,\Big] \geq \frac{|U_1|\,|W_1|}{|U|\,|W|}\left(d(U_1, W_1) - d(U, W)\right)^2 > \epsilon^4$$

because U_1 and W_1 witnessed the irregularity of the pair (U, W), which means $|U_1| \geq \epsilon |U|$, $|W_1| \geq \epsilon |W|$ and

$$|d(U_1, W_1) - d(U, W)| > \epsilon.$$

$\square$

Lemma 1.9. *If a partition $\mathcal{P} = \{V_1, \ldots, V_k\}$ is not ϵ–regular, then there exists a refinement $\mathcal{Q}$ of $\mathcal{P}$ where every V_i is partitioned into at most 2^k parts such that*

$$q(\mathcal{Q}) \geq q(\mathcal{P}) + \epsilon^5.$$

Proof. For all irregular pairs (V_i, V_j), find simultaneously the subsets $A_{ij} \subset V_i$ and $A_{ji} \subset V_j$ that witness irregularity. Let $\mathcal{Q}$ be a common refinement of $\mathcal{P}$ by $\{A_{ij}\}_{(i,j)\in[k]^2}$. Observe that each V_i is partitioned into at most 2^k parts as we wanted. If we compute the energy of $\mathcal{Q}$, we get

$$q(\mathcal{Q}) = \sum_{(i,j)\in[k]^2} q(\mathcal{Q}_{V_i}, \mathcal{Q}_{V_j}) = \sum_{\substack{(i,j)\in[k]^2 \\ (V_i,V_j)\ \epsilon\text{–regular}}} q(\mathcal{Q}_{V_i}, \mathcal{Q}_{V_j}) + \sum_{\substack{(i,j)\in[k]^2 \\ (V_i,V_j)\ \text{not}\ \epsilon\text{–regular}}} q(\mathcal{Q}_{V_i}, \mathcal{Q}_{V_j}),$$

where $\mathcal{Q}_{V_i}$ is the partition of V_i in the refinement $\mathcal{Q}$. By Lemma 1.7, for each term of the second sum we have

$$q(\mathcal{Q}_{V_i}, \mathcal{Q}_{V_j}) \geq q(\{A_{ij}, V_1 \setminus A_{ij}\}, \{A_{ji}, V_j \setminus A_{ji}\}).$$

Now, by Lemma 1.8,

$$\sum_{\substack{(i,j)\in[k]^2 \\ (V_i,V_j)\ \text{not}\ \epsilon\text{–regular}}} q\left(\{A_{ij}, V_i \setminus A_{ij}\}, \{A_{ji}, V_j \setminus A_{ji}\}\right) \geq \sum_{\substack{(i,j)\in[k]^2 \\ (V_i,V_j)\ \text{not}\ \epsilon\text{–regular}}} q(V_i, V_j) + \epsilon^4 \frac{|V_i|\,|V_j|}{|V|^2}.$$

For the first sum of $q(\mathcal{Q})$, again by Lemma 1.6 we have $q(\mathcal{Q}_{V_i}, \mathcal{Q}_{V_j}) \geq q(V_i, V_j)$ for each term. Putting all this bounds together, we obtain the inequality

$$q(\mathcal{Q}) \geq \sum_{(i,j)\in[k]^2} q(V_i, V_j) + \sum_{\substack{(i,j)\in[k]^2 \\ (V_i,V_j)\ \text{not}\ \epsilon\text{–regular}}} \epsilon^4 \frac{|V_i|\,|V_j|}{|V|^2} \geq q(\mathcal{P}) + \epsilon^5,$$

where the last step is given by irregularity of the initial partition $\mathcal{P}$. $\square$

Finally, Szemerédi's Regularity Lemma follows from the above lemmas.

Proof of Theorem 1.4. Start with a trivial partition. While the partition is not ϵ–regular, apply repeatedly Lemma 1.9 to refine it. Recall that the energy is upper bounded by 1, but at each iteration it increases by at least ϵ^5. Therefore, we obtain an ϵ–regular partition in at most ϵ^{-5} steps. $\square$

At each iteration, if the partition $\mathcal{P}$ has k parts we obtain a refinement of at most 2^{2^k} parts. In the end, the algorithm produces a partition of at most a tower of twos of height ϵ^{-5}. Gowers proved in [9] that the bound is a tower of twos yet the height of a lower bound can be improved to $\log(\epsilon^{-1})$.

1.2 The Graph Removal Lemma

A classic application of Szemerédi's Regularity Lemma is the well–known Graph Removal Lemma. This result states that a graph G which contains a few copies of a given subgraph H can be made H–free by removing a small number of edges, and gives the relation between the number of initial copies of H and the maximum number of edges removed.

Among the results that use the Graph Removal Lemma we remark Roth's Theorem, which is a significant example of how Szemerédi's Regularity Lemma can be used to solve other than Graph Theory problems. Roth's Theorem states that every subset of the integers with positive upper density contains a 3–term arithmetic progression. The idea behind the proof is to wisely define a graph G and then apply the Triangle Removal Lemma, which is the particular case of the Removal Lemma taking a triangle as the subgraph H (see Zhao [19], Theorem 3.19). Another important result that can be proved using the same strategy as in the Removal Lemma is the Erdős–Stone–Simonovits Theorem, for which we will see the proof later.

The structure of the proof of the Graph Removal Lemma can be divided into three steps:

1. Partition: apply the Regularity Lemma to obtain a regular partition of the vertices of G.

2. Clean: remove edges between irregular pairs, low–density pairs and pairs for which one part is too small.

3. Count: if the cleaned graph still has a copy of H, then the number of copies in the original graph must have been larger.

For the last step we will use the Graph Counting Lemma, which as the name suggests, counts the number of instances of H that we can find in a graph G when it meets some specific ϵ–regularity conditions. Intuitively, the Counting Lemma says that if there are subsets of vertices of G indexed by the vertices of H such that a pair of them is ϵ–regular whenever there is an edge in H between their corresponding vertices, we can expect a random–like behaviour between such pairs and hence the number of copies of H can be approximated using the edge–densities of those regular pairs.

Theorem 1.10 (Graph Counting Lemma)**.** *Let H be a graph on k vertices and let $\epsilon > 0$. Let G be a graph on n vertices with vertex subsets $V_1, \ldots, V_k \subseteq V(G)$ such that (V_i, V_j) is ϵ–regular whenever $\{i, j\} \in E(H)$. Then, the number of tuples $(v_1, \ldots, v_k) \in V_1 \times \cdots \times V_k$ such that $\{v_i, v_j\} \in E(G)$ whenever $\{i, j\} \in E(H)$ is within $\epsilon \cdot e(H) \cdot |V_1| \cdots |V_k|$ of*

$$\left(\prod_{\{i,j\} \in E(H)} d(V_i, V_j) \right) \left(\prod_{i=1}^{k} |V_i| \right).$$

Proof. Let us prove a probabilistic version of the theorem: if we have vertices $v_1 \in V_1, \dots, v_k \in V_k$ chosen uniformly and independently at random, then

$$\left| \Pr\left(\{v_i, v_j\} \in E(G) \;\; \forall \{i,j\} \in E(H) \right) - \prod_{\{i,j\} \in E(H)} d(V_i, V_j) \right| \leq \epsilon \cdot e(H). \tag{1}$$

That is, the probability that the tuple $(v_1, \dots, v_k)$ is indeed a copy of H in G is within $\epsilon \cdot e(H)$ of

$$\prod_{\{i,j\} \in E(H)} d(V_i, V_j).$$

Define $\mathbb{P}_H := \Pr\left(\{v_i, v_j\} \in E(G) \;\; \forall \{i,j\} \in E(H) \right)$. Relabelling if necessary, assume $\{1,2\} \in E(H)$. If H' denotes the graph obtained by removing the edge $\{1,2\}$ from H, we will show first that

$$|\mathbb{P}_H - d(V_1, V_2)\mathbb{P}_{H'}| \leq \epsilon. \tag{2}$$

Let us fix $v_3, \dots, v_k$ and choose randomly only v_1, v_2. Observe that if last inequality is satisfied under this extra constraint, it also holds for the weaker case where all k vertices are chosen randomly. Let

$$A_1 := \{v_1 \in V_1 : \{v_1, v_i\} \in E(G) \text{ whenever } \{1, i\} \in E(H')\},$$
$$A_2 := \{v_2 \in V_2 : \{v_2, v_i\} \in E(G) \text{ whenever } \{2, i\} \in E(H')\},$$

be the possible choices for v_1 and v_2 which, with $v_3, \dots, v_k$, give an embedding of H' in G. If we reformulate the probabilistic statement (2) with $v_3, \dots, v_k$ fixed, we have to prove that

$$\left| \frac{e(A_1, A_2)}{|V_1|\,|V_2|} - d(V_1, V_2) \frac{|A_1|\,|A_2|}{|V_1|\,|V_2|} \right| \leq \epsilon.$$

In the case $|A_1| \leq \epsilon\,|V_1|$ or $|A_2| \leq \epsilon\,|V_2|$, both terms are at most ϵ, so the inequality holds. Otherwise, if $|A_1| > \epsilon\,|V_1|$ and $|A_2| > \epsilon\,|V_2|$, by ϵ–regularity of (V_1, V_2) we have

$$\left| \frac{e(A_1, A_2)}{|V_1|\,|V_2|} - d(V_1, V_2) \frac{|A_1|\,|A_2|}{|V_1|\,|V_2|} \right| = |d(A_1, A_2) - d(V_1, V_2)| \frac{|A_1|\,|A_2|}{|V_1|\,|V_2|} \leq \epsilon,$$

which proves (2).

We are going to complete the proof by induction on the number of edges of H. Since $e(H') = e(H) - 1$, the inequality (1) holds when replacing H for H'. Therefore,

$$\begin{aligned}
\left| \mathbb{P}_H - \prod_{\{i,j\} \in E(H)} d(V_i, V_j) \right| &\leq |\mathbb{P}_H - d(V_1, V_2)\mathbb{P}_{H'}| + d(V_1, V_2) \left| \mathbb{P}_{H'} - \prod_{\{i,j\} \in E(H')} d(V_i, V_j) \right| \\
&\leq \epsilon + d(V_1, V_2) \cdot \epsilon \cdot e(H') \\
&\leq \epsilon \left(1 + e(H')\right) \\
&= \epsilon \cdot e(H).
\end{aligned}$$

□

Now we have all the necessary tools to prove the Graph Removal Lemma following the above procedure.

Theorem 1.11 (Graph Removal Lemma)**.** *For each graph H and each $\epsilon > 0$ there exists $\delta > 0$ such that every graph on n vertices with fewer than $\delta n^{|V(H)|}$ copies of H can be made H–free by removing no more than ϵn^2 edges.*

Proof. Let G be a graph on n vertices. First, apply Szemerédi's Regularity Lemma 1.4 to obtain an $\epsilon/4$–regular partition $\mathcal{P} = \{V_1, \dots, V_M\}$ of $V(G)$. Secondly, consider the graph G' obtained from G by removing all the edges between V_i and V_j whenever

(a) (V_i, V_j) is not $\epsilon/4$–regular,

(b) $d(V_i, V_j) < \epsilon/2$, or

(c) V_i or V_j has less than $\frac{\epsilon n}{4M}$ vertices.

The number of removed edges in (a) is, by $\epsilon/4$–regularity of the partition,

$$\sum_{(V_i,V_j) \text{ not } \epsilon/4\text{–regular}} e(V_i, V_j) \le \sum_{(V_i,V_j) \text{ not } \epsilon/4\text{–regular}} |V_i|\,|V_j| \le \frac{\epsilon}{4} n^2.$$

For low density pairs in (b), we have

$$\sum_{d(V_i,V_j)<\epsilon/2} e(V_i, V_j) = \sum_{d(V_i,V_j)<\epsilon/2} d(V_i, V_j)\,|V_i|\,|V_j| \le \frac{\epsilon}{2} \sum_{(i,j)\in[M]^2} |V_i|\,|V_j| = \frac{\epsilon}{2} n^2.$$

Finally, for small subsets in (c), the number of removed edges is

$$\sum_{|V_i| \text{ or } |V_j| < \frac{\epsilon n}{4M}} e(V_i, V_j) \le n \cdot \frac{\epsilon n}{4M} \cdot M = \frac{\epsilon}{4} n^2,$$

because each vertex in G could be adjacent to at most $\frac{\epsilon n}{4M}$ vertices of each of the at most M small subsets of the partition, which gives a total number of removed edges of at most ϵn^2.

Let $V(H) = [k]$ and suppose that there is a tuple $(v_1, \dots, v_k) \in V_{s_1} \times \cdots \times V_{s_k}$ for some $s_1, \dots, s_k \in [M]$ such that $\{v_i, v_j\} \in E(G')$ whenever $\{i, j\} \in E(H)$, i.e. there is a copy of H in G' realised by $(v_1, \dots, v_k)$. Due to the cleaning step, all pairs (V_{s_i}, V_{s_j}) are ϵ–regular in G', have edge–density at least $\epsilon/2$ and $|V_{s_i}|, |V_{s_j}| \ge \frac{\epsilon n}{4M}$. Making use of this bounds and the Counting Lemma 1.10, the number of copies of H in G' is at least

$$\begin{aligned}
\frac{1}{|\mathrm{Aut}(H)|} &\left(\left(\prod_{\{i,j\}\in E(H)} d(V_{s_i}, V_{s_j}) \right) \left(\prod_{i=1}^{k} |V_{s_i}| \right) - \epsilon \cdot e(H) \prod_{i=1}^{k} |V_{s_i}| \right) \\
&\ge \frac{1}{|\mathrm{Aut}(H)|} \left(\left(\frac{\epsilon}{2}\right)^{e(H)} \left(\frac{\epsilon n}{4M}\right)^k - \epsilon \cdot e(H) \left(\frac{\epsilon n}{4M}\right)^k \right) \\
&= \frac{1}{|\mathrm{Aut}(H)|} \left(\frac{\epsilon}{4M}\right)^k \left(\left(\frac{\epsilon}{2}\right)^{e(H)} - \epsilon \cdot e(H) \right) n^k,
\end{aligned}$$

where the constant $1/\left|\mathrm{Aut}(H)\right|$ fixes the overcounting in the case $s_1 = \cdots = s_k$.

By choosing a sufficiently small δ, namely

$$\delta < \frac{1}{|\mathrm{Aut}(H)|}\left(\frac{\epsilon}{4M}\right)^k \left(\left(\frac{\epsilon}{2}\right)^{e(H)} - \epsilon \cdot e(H)\right),$$

we get a contradiction: if there is some copy of H left in the cleaned graph G', the Counting Lemma implies that we can actually find more than δn^k copies. However, the original graph G had at most δn^k copies, so we conclude that the graph G' obtained from G by removing at most ϵn^2 edges is H–free. $\square$

The Graph Removal Lemma can be easily translated to asymptotic notation: if a graph on n vertices contains $o\left(n^{|V(H)|}\right)$ copies of a graph H then we can make it H–free by removing $o\left(n^2\right)$ edges.

An interesting result of Extremal Graph Theory that can be proved using a similar strategy is the Erdős–Stone–Simonovits Theorem. It gives the expression of the extremal number $\mathrm{ex}(n, H)$, i.e. the minimum number such that every graph on n vertices and $\mathrm{ex}(n, H)$ edges contains an instance of H, in terms of the chromatic number of a graph H.

Theorem 1.12 (Erdős–Stone–Simonovits). *For every fixed graph H, its extremal number is*

$$\mathrm{ex}(n, H) = \left(1 - \frac{1}{\chi(H) - 1} + o(1)\right)\frac{n^2}{2}.$$

Proof. Let $\chi(H) = r + 1$ and fix $\epsilon > 0$. Let G be any graph on n vertices with at least $(1 - \frac{1}{r} + \epsilon)\frac{n^2}{2}$ edges. We are going to prove that, if $n = n(\epsilon, H)$ is sufficiently large, then G contains a copy of H. First, we apply Szemerédi's Regularity Lemma 1.4 to G to obtain a δ–regular partition $\{V_1, \ldots, V_M\}$ of its vertices, where

$$\delta := \frac{1}{2e(H)}\left(\frac{\epsilon}{8}\right)^{e(H)}.$$

Let G' be the graph obtained from G by removing all the edges between V_i and V_j whenever

(a) (V_i, V_j) is not δ–regular,

(b) $d(V_i, V_j) < \epsilon/8$, or

(c) $|V_i|$ or $|V_j|$ has less than $\frac{\epsilon n}{8M}$ vertices.

The number of removed edges in (a) is, by δ–regularity of the partition,

$$\sum_{(V_i,V_j) \text{ not } \delta\text{–regular}} e(V_i, V_j) \le \sum_{(V_i,V_j) \text{ not } \delta\text{–regular}} |V_i||V_j| \le \frac{\delta}{n}^2.$$

For low denisty pairs in (b), we have

$$\sum_{d(V_i,V_j)<\epsilon/8} e(V_i,V_j) = \sum_{d(V_i,V_j)<\epsilon/8} d(V_i,V_j)|V_i||V_j| \leq \frac{\epsilon}{8} \sum_{(i,j)\in[M]^2} |V_i||V_j| = \frac{\epsilon}{8}n^2.$$

Finally, for small subsets in (c), the number of removed edges is

$$\sum_{|V_i| \text{ or } |V_j|<\frac{\epsilon n}{4M}} e(V_i,V_j) \leq n \cdot \frac{\epsilon n}{8M} \cdot M = \frac{\epsilon}{8}n^2,$$

because each vertex in G can be adjacent to at most $\frac{\epsilon n}{8M}$ vertices of each of the at most M small subsets of the partition. Hence, the total number of removed edges satisfies

$$\delta n^2 + \frac{\epsilon}{8}n^2 + \frac{\epsilon}{8}n^2 \leq \frac{3}{8}\epsilon n^2$$

and the cleaned graph G' has at least $\left(1 - \frac{1}{r} + \frac{\epsilon}{4}\right)\frac{n^2}{2}$ edges.

By Turán's Theorem (see [1]) we know that G' contains a copy of K_{r+1}, and assume that it is given by vertices $[r+1]$, relabelling if necessary. Let $s_1, \ldots, s_{r+1} \in [M]$ (possibly repeated) be such that $(1, \ldots, r+1) \in V_{s_1} \times \cdots \times V_{s_{r+1}}$. Due to the cleaning step, we know that all pairs (V_{s_i}, V_{s_j}) are δ–regular in G', have edge–density at least $\epsilon/8$ and $|V_{s_i}|, |V_{s_j}| \geq \frac{\epsilon n}{8M}$.

Let $V(H) = [k]$ and consider a coloring $c : [k] \to [r+1]$, which exists because $\chi(H) = r+1$. Therefore, subsets $V_{s_{c(1)}}, \ldots, V_{s_{c(k)}}$ meet the hypothesis of the Counting Lemma 1.10 and thus we obtain that the number of homomorphisms of H in G' is at least

$$\begin{aligned}\left(\prod_{\{i,j\}\in E(H)} d\left(V_{s_{c(i)}}, V_{s_{c(j)}}\right)\right)\left(\prod_{i=1}^{k}\left|V_{s_{c(i)}}\right|\right) - \delta \cdot e(H) \cdot \left|V_{s_{c(1)}}\right| \cdots \left|V_{s_{c(k)}}\right| \\ \geq \left(\left(\frac{\epsilon}{8}\right)^{e(H)} - \delta \cdot e(H)\right)\left(\frac{\epsilon}{8M}\right)^k n^k.\end{aligned}$$

Observe that we are counting all homomorphisms, although there are some which result in a subgraph of H instead of H. However, the number of non–injective maps $[k] \to [n]$ is

$$n^k - \frac{n!}{(n-k)!} = O(n^{k-1}),$$

yet the order of total the number of homomorphisms is k for our choice of δ. In conclusion, for a sufficiently large n there is an injective map $[k] \to [n]$ so there is indeed a copy of H in G'. $\square$

Chapter 2

Spectral Graph Theory

The *spectrum* of the graph $G = (V, E)$ is defined as the spectrum of its adjacency matrix $A(G)$. The purpose of Spectral Graph Theory is to study structural properties of a graph which can be derived from its spectrum. In this chapter we will see some basic results which will provide the background on this subject needed to prove the main results of Chapter 3. The proofs of the following statements can be found in any standard Spectral Graph Theory book, for example Brouwer–Haemers [4].

2.1 Introduction and examples

First of all, note that $A = A(G)$ is a symmetric 0–1 matrix, and therefore all its eigenvalues are real. Throughout the chapter, n will be the number of vertices of G and m the size of the edge set. The spectrum will be denoted by $(\lambda_1, \dots, \lambda_n)$ where the eigenvalues are labeled according to absolute value in nonincreasing order, that is $|\lambda_1| \geq \cdots \geq |\lambda_n|$. Recall that the entry $(a_{ij}^{(k)})$ of A^k counts the number of walks of length k from vertex i to j.

Proposition 2.1. *Let G be a graph on n vertices and let A be its adjacency matrix. Then,*

(i) $\sum_{i=1}^n \lambda_i = \mathrm{Tr}(A) = 0$.

(ii) $\sum_{i=1}^n \lambda_i^2 = \mathrm{Tr}(A^2) = 2m \leq 2\binom{n}{2} \leq n^2$.

(iii) $\sum_{i=1}^n \lambda_i^3 = \mathrm{Tr}(A^3) = 6T$, *where T is the number of triangles in G.*

Another elementary spectral result is the value of the largest eigenvalue of an r–regular graph, that is a graph such that the degree of all vertices is r.

Proposition 2.2. *Let G be an r–regular connected graph. Then $\lambda_1 = r$ and $|\lambda_i| < r$ for all $i > 1$.*

Let us now review the spectrum of some remarkable families of graphs.

Proposition 2.3 (Circulant graphs)**.** *Let G be a graph with adjacency matrix $A = circ(a_1, \dots, a_n)$, where $(a_1, \dots, a_n)$ is the first row and row k is a cyclic shift of row $k-1$. Then the j–th eigenvalue of G is*

$$\lambda_j = \sum_{k=1}^{n} a_k \omega^{jk},$$

where $\omega = e^{2\pi i/n}$ is a primitive n–th root of unity.
Moreover,

$$\mathbf{u}_j = \frac{1}{\sqrt{n}}(\omega^j, \omega^{2j}, \ldots, \omega^{(n-1)j}, 1), \quad j = 1, \ldots, n,$$

is an orthonormal basis of eigenvectors of G. In particular,

(i) for the complete graph K_n we have

$$\mathrm{Spec}(K_n) = (n-1, -1, \ldots, -1),$$

and $\mathbf{u}_1 = \frac{1}{\sqrt{n}}\mathbf{1}$, $\mathbf{u}_j = \frac{1}{\sqrt{2}}(\mathbf{e}_1 - \mathbf{e}_j)$ for $j > 1$ (where $\mathbf{e}_j$ represents the vector with a 1 in position j and a 0 everywhere else) is an orthonormal basis of eigenvectors.

(ii) $\mathrm{Spec}(C_n) = (2, 2\cos(2\pi/n), 2\cos(4\pi/n), \ldots, 2\cos(2(n-1)\pi/n)$.

Proposition 2.4 (Complement of a graph)**.** *Let G be an r–regular graph and let $\overline{G}$ be the complement of G. If* $\mathrm{Spec}(G) = (r, \lambda_2, \ldots, \lambda_n)$ *then*

$$\mathrm{Spec}(\overline{G}) = (n-r-1, -\lambda_n - 1, \ldots, -\lambda_2 - 1).$$

Moreover, the eigenvectors of G are also eigenvectors of $\overline{G}$.

Proposition 2.5 (Line graphs)**.** *Let $G = (V, E)$ be an r–regular graph on n vertices and let $L(G)$ denote the line graph of G, which has E as vertex set and two vertices are adjacent if the corresponding edges are incident in G. Then, the eigenvalues of the line graph are $\lambda_i - 2 + r$ for $i = 1, \ldots, n$ and -2 with multiplicity $m - n$.*

Another interesting family are Kneser graphs. The *Kneser graph $K(n, k)$* has the k–subsets of $[n]$ as vertex set, and two subsets are adjacent if they are disjoint. One famous example is $K(5, 2)$, also known as the Petersen graph.

Proposition 2.6 (Kneser graphs)**.** *The eigenvalues of the Kneser graph $K(n, k)$, $n \geq 2k$, are*

$$(-1)^t \binom{n-k-t}{k-t}, \quad t = 0, \ldots, k.$$

The eigenvectors of $(-1)^t\binom{n-k-t}{k-t}$ are the vectors

$$\mathbf{x} = (x_K : K \in \binom{[n]}{k}),$$

where each coordinate x_K is

$$x_K = \sum_{T \in \binom{K}{t}} y_T$$

with the y_T satisfying

$$\sum_{T \cap K = \emptyset} y_T = (-1)^t \sum_{T \subset K} y_T.$$

Proof. The Kneser graph is r–regular with $r = \binom{n-k}{k}$, and by Proposition 2.2 this is the largest eigenvalue λ_1. We are going to determine the remaining eigenvalues by finding plausible eigenvectors. Suppose that $\mathbf{x} = (x_K : K \in \binom{[n]}{k})$ is an eigenvector of the adjacency matrix of the Kneser graph $K(n, k)$. Then, for each k–subset K of $[n]$, we must have

$$\sum_{K' \cap K = \emptyset} x_{K'} = \lambda x_K, \tag{1}$$

where λ is the corresponding eigenvalue, because if $A = (a_{KK'})$ is the adjacency matrix indexed by k–subsets of $[n]$, then by definition $a_{KK'} = 1$ if and only if $K \cap K' = \emptyset$. One way to obtain x_K for every K satisfying the above equality (1) for some λ and all $K \in \binom{[n]}{k}$ is as follows. Fix $t < k$ and consider numbers $\{y_T : T \in \binom{[n]}{t}\}$ to be specified later and define, for each k–subset K,

$$x_K = \sum_{T \in \binom{K}{t}} y_T.$$

Observe that every t–subset T not meeting K apears in $\binom{n-k-t}{k-t}$ k–subsets disjoint from K: one T is fixed we must complete it with $k - t$ elements not in $K \cup T$ in order to obtain a k–subset K' disjoint from K and containing T. Therefore, equation (1) reads

$$\sum_{K' \cap K = \emptyset} x_{K'} = \sum_{T \cap K = \emptyset} \binom{n-k-t}{k-t} y_T = \lambda \sum_{T \subset K} y_T.$$

We would be done if we can find real numbers y_T such that, for each $K \in \binom{[n]}{k}$, we have

$$\sum_{T \cap K = \emptyset} y_T = (-1)^t \sum_{T \subset K} y_T, \tag{2}$$

which would give the eigenvector corresponding to eigenvalue $(-1)^t \binom{n-k-t}{k-t}$.

Let us show that there are in fact $\binom{n}{t} - \binom{n}{t-1}$ independent vectors $(y_T : T \in \binom{[n]}{T})$ satisfying equation (2). Actually it suffices to see that, for each $U \in \binom{[n]}{t-1}$, we have

$$\sum_{T \supset U} y_T = 0. \tag{3}$$

In this case, for each $i = 0, \ldots, t$ we have

$$0 = \sum_{U : |U \cap K| = i} \sum_{T : T \supset U} y_T = \sum_{T : |T \cap K| = i} \sum_{U : U \subset T} y_T + \sum_{T : |T \cap K| = i+1} \sum_{U : U \subset T} y_T, \tag{4}$$

because whenever $|U \cap K| = i$, sets T containing U satisfy $i \leq |T \cap K| \leq i + 1$ and the second equality in the above equation follows by exchanging the order of summation. Now, for every set T with $|T \cap K| = i + 1$ we have precisely $(t - i)$ sets U contained in T with $|U \cap K| = i$, while if $|T \cap K| = i$ then T contains $(i+1)$ sets U contained in T with $|U \cap K| = i$. Hence equation (4) gives

$$0 = (t-i) \sum_{T:|T\cap K|=i} y_T + (i+1) \sum_{T:|T\cap K|=i+1} y_T.$$

The above equation gives a recurrence in i which leads to

$$\sum_{T:|T\cap K|=i+1} y_T = -\frac{t-i}{i+1} \sum_{T:|T\cap K|=i} y_T = \cdots = (-1)^{i+1} \binom{t}{i+1} \sum_{T:|T\cap K|=0} y_T$$

and, for $i+1 = t$,

$$\sum_{T:T\subset K} y_T = (-1)^t \sum_{T:|T\cap K|=0} y_T,$$

which is precisely (2). Moreover, (3) gives a linear system with $\binom{n}{t}$ variables and $\binom{n}{t-1}$ equations, so the vector space of solutions has dimension $\binom{n}{t} - \binom{n}{t-1}$. Those vectors are clearly independent for distinct values of t (because they correspond to distinct eigenvalues) and therefore we have

$$1 + \sum_{t=1}^{k-1} \binom{n}{t} - \binom{n}{t-1} = \binom{n}{k}$$

linearly independent eigenvectors. □

2.2 Main results and connection with regularity

A central tool in Spectral Graph Theory are the Courant–Fischer inequalities. Observe that the theorem is not restricted to adjacency matrices of graphs, it holds for any real symmetric matrix.

Theorem 2.7 (Courant–Fischer inequalities)**.** *Let A be a real symmetric matrix. Then*

$$\lambda_k = \max_{V\in\mathcal{V}_k} \min_{\mathbf{v}\in V, \|\mathbf{v}\|_2=1} \mathbf{v}^T A\mathbf{v},$$

where $\mathcal{V}_k$ denotes the family of all k–subspaces of $\mathbb{R}^n$. Similarly,

$$\lambda_k = \min_{V\in\mathcal{V}_{n-k+1}} \max_{\mathbf{v}\in V, \|\mathbf{v}\|_2=1} \mathbf{v}^T A\mathbf{v}.$$

Proof. See Brouwer–Haemers [4], Theorem 2.4.1.

One of the consequences is the relation of the average and maximum degrees of a graph G with the first eigenvalue of the adjacency matrix A.

Corollary 2.8. *Let G be a connected graph. If $d(G)$ and $\Delta(G)$ denote the average and maximum degrees of G, then*

$$d(G) \le \lambda_1 \le \Delta(G).$$

The spectrum of the Laplacian of a graph G is also related to some properties of G such as connectivity. Recall that the *Laplacian* $L(G) = (l_{ij})$ is the $n \times n$ defined as

$$l_{ij} = \begin{cases} -1 & \text{if } \{i,j\} \in E, \\ d(i) & \text{if } i = j, \\ 0 & \text{otherwise,} \end{cases}$$

that is to say, $L(G) = D - A$ where D is the diagonal matrix with entries the degrees of the vertices.

Proposition 2.9. *Let $\mu_1 \leq \cdots \leq \mu_n$ be the spectrum of the Laplacian of a graph G. Then,*

(i) $\mu_1 = 0$ and has eigenvector $\mathbf{1} = (1, \ldots, 1)$.

(ii) If G is connected then $\mu_2 > 0$.

(iii) If G is r–regular then $\mu_i = r - \lambda_i$, where λ_i is the i–th eigenvalue of A.

It follows from the above proposition that $L = L(G)$ is semipositive definite. Interpreting L as a bilinear operator, for every vector $\mathbf{v} = (v_1, \ldots, v_n) \in \mathbb{R}^n$ we have

$$\mathbf{v}^T L \mathbf{v} = \sum_{\{i,j\} \in E} (v_i - v_j)^2.$$

The above equality is the basis of many applications of spectral graph theory via the Courant–Fischer inequalities. For instance, given that $\mu_1 = 0$ and has eigenvector $\mathbf{1}$, then for the second smallest eigenvalue μ_2 and for every unit vector $\mathbf{v}$ orthogonal to $\mathbf{1}$ we have

$$\mathbf{v}^T L \mathbf{v} \geq \mu_2 \left\| \mathbf{v} \right\|_2^2 = \mu_2.$$

By choosing $\mathbf{v} = \mathbf{1}_U$ to be the indicator function of a subset $U \subset V$ of vertices of G, and using the fact that the bilinear operator L is invariant by translations by the vector $\mathbf{1}$ (again because $\mathbf{1}$ is an eigenvector of eigenvalue $\mu_1 = 0$), we have

$$e(U, V \setminus U) = \mathbf{1}_U^T L \mathbf{1}_U = (\mathbf{1}_U - c\mathbf{1})^T L (\mathbf{1}_U - c\mathbf{1}) \geq \mu_2 \left\| \mathbf{1}_U - c\mathbf{1} \right\|_2,$$

where $\mathbf{1}_U - c\mathbf{1}$ is the projection of $\mathbf{1}_U$ onto the space orthogonal to $\langle \mathbf{1} \rangle$ for some $c \in \mathbb{R}$.

Thus μ_2 is related to the expansion properties of the graph G which refer to the number of edges leaving a set with respect to the size of the set. One basic example of this relation are the Cheeger inequalities. These are bounds for the *isoperimetric number* of a graph G, also called the *Cheeger constant*, that is

$$i(G) = \min \left\{ \frac{e(U, V \setminus U)}{|U|} : U \subset V, |U| \leq \frac{n}{2} \right\}.$$

Theorem 2.10 (Cheeger inequalities)**.** *Let G be a connected graph with n vertices. Then*

$$\frac{\mu_2}{2} \leq i(G)$$

and if $G \neq K_n$,

$$i(G) \leq \sqrt{2\Delta\mu_2}.$$

Proof. See Brouwer–Haemers [4], Proposition 4.5.2.

The above theorem shows that the *spectral gap*, the distance between the two smallest eigenvalues of the Laplacian matrix, gives a spectral description of the expanding properties of a graph.

A second consequence of the spectral approach, which is relevant for the Regularity Lemma, is the so–called Mixing Lemma. We will say that a graph G is an (n, r, λ)*–expander* if it has n vertices, it is r–regular and has second adjacency eigenvalue in absolute value at most λ, namely $|\lambda_2| \leq \lambda$.

Theorem 2.11 (Mixing Lemma)**.** *Let G be an (n, r, λ)–expander. For every pair U, V of disjoint sets of vertices,*

$$\left| e(U, V) - \frac{r}{n} |U| \, |V| \right| \leq \lambda \, |U| \, |V| \, .$$

Proof. See Brouwer–Haemers [4], Proposition 4.3.2.

An interpretation of the above inequality is that the number of edges between two sets is close to the number one would find if the edges were placed randomly, and the deviation is smaller as the value of λ is smaller. One more time, the spectral gap describes a structural property of the graph in terms of the random–like distribution of the edges. This leads to the notion of pseudorandomness, a notion introduced by Thomason in [18]. A graph G is said to be (p, α)*–jumbled* if, for every subset U of vertices, the inequality

$$\left| e(U) - p\binom{|U|}{2} \right| \leq \alpha |U|$$

holds. Chung, Graham and Wilson showed that this notion indeed captures the nature of randomness in the sense that a (p, α)–jumbled graph enjoys a number of properties which happen almost surely in the binomial random graph where every edge is independently chosen with probability p, which is an informal definition of what we call a *pseudorandom graph* (see [6]). Among these properties, we find that the largest eigenvalue of the adjacency matrix of G satisfies

$$|\lambda_1 - pn| \leq \epsilon n$$

while $\lambda = \max\{|\lambda_2|, |\lambda_n|\} \leq \epsilon n$.

The proof of the Mixing Lemma uses the fact that, for disjoint sets U, V of vertices,

$$\mathbf{1}_U^T A \mathbf{1}_V = e(U, V).$$

The condition of regularity of the graph can be disposed of in several ways. One of them is using the *cut norm* of a matrix A, which is defined as

$$\|A\|_\square = \sup\{\left|x^T A y\right| : \|x\|_\infty, \|y\|_\infty \leq 1\}.$$

This norm was introduced in Frieze–Kannan [8] and was used in the definition of graph limits to define an appropriate metric for the space of graphs and graphons. The above is one among several equivalent definitions of the cut norm. We observe that

$$\|A\|_\square \leq \text{rad}(A),$$

where $\text{rad}(A)$ denotes the *spectral radius* of A, the largest eigenvalue of A in absolute value. Thus for every pair of disjoint sets U and V we have

$$e(U, V) = \mathbf{1}_U^T A \mathbf{1}_V \leq \|A\|_\square \leq \text{rad}(A),$$

which means that a graph with an adjacency matrix with small spectral radius implies that $e(U, V)$ is small.

In addition to the Courant–Fischer inequalities, a common tool in spectral graph theory is the use of the Spectral Theorem. If $\mathbf{u}_1, \ldots, \mathbf{u}_n$ is an orthogonal basis of eigenvectors, listed in the order of nonincreasing eigenvalues, then

$$A = \sum_{i=1}^{n} \lambda_i \mathbf{u}_i \mathbf{u}_i^T.$$

Among the many applications of this spectral decomposition in spectral graph theory there is the following theorem given by Chung (see [5]).

Theorem 2.12 (Spectral bound of diameter)**.** *Let G be an r–regular graph with second largest eigenvalue in absolute value λ. The diameter $D(G)$ satisfies*

$$D(G) \leq \left\lceil \frac{\log n}{\log(r/\lambda)} \right\rceil.$$

Proof. See Brouwer–Haemers [4], Proposition 4.7.1.

Once more, the smaller the value of λ the closer is the diameter to $\log n$, the smallest possible in order of magnitude among the possible values of the diameter in a regular graph.

Chapter 3

Spectral regularity theorems

This central chapter contains the spectral proof of Szemerédi's Regularity Lemma. Furthermore, the same spectral method is used to prove regularity theorems for graphs with edge–colorings and multicolorings, and weak versions of regularity for multigraphs and directed graphs.

In order to do so, we prove first more general results for matrices. We start with the Spectral Regularity Theorem for self–adjoint matrices which provides the definition ϵ–regular partition of a matrix. Then, anticipating the results for graphs that we want to prove, we prove generalisations of the matrix which give partitions that are ϵ–regular for a set of matrices simultaneously. Then, we translate the matricial versions to the desired graph statements.

To illustrate the results, there is a final section with some examples of the partitions we obtain using spectral tools.

3.1 Spectral regularity for matrices

This section is a step by step review of the first parts of the generalisation and explanation of Cioaba and Martin (see [7]) of the spectral proof given by Tao in [16]. In addition, there are some observations of interest concerning the interpretation of the obtained partition.

First of all, we are going to prove a theorem which provides a descomposition $T = T_1+T_2+T_3$ of a self–adjoint $n \times n$ matrix T together with a partition $\{V_0, \dots, V_M\}$ of $[n]$ and a set of "irregular" pairs $\Sigma \subset \{0, \dots, M\}^2$. Matrices T_1, T_2 and T_3 of the decomposition are also self–adjoint, and they have the following properties:

- T_1 is ϵ–constant on each submatrix on $V_s \times V_t$ for a pair $(s,t) \notin \Sigma$,
- $\|T_2\|_2^2 \le \epsilon^3 n^2$, and
- if $\mathrm{rad}(T_3)$ denotes the spectral radius of T_3, then

$$\mathrm{rad}(T_3) \le \frac{\epsilon}{n} |V_s| \, |V_t| \, .$$

The decomposition can be interpreted as a structured part T_1 close to a step function with respect to the given partition, T_3 a random noise on the same partition, involving fluctuations bounded by its largest eigenvalue which is small, and a negligible part T_2.

Throughout the section, we will consider self–adjoint $n \times n$ matrices T with coefficients in $\mathbb{C}$, their spectrum $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ ordered according to absolute value in nonincreasing order and an orthonormal basis of eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{C}^n$. We denote the $n \times n$ all–ones matrix by J_n.

Theorem 3.1 (Matrix Regularity Theorem)**.** *Let T be a self–adjoint $n \times n$ matrix such that $\mathrm{Tr}(T^2) \leq n^2$. For every $\epsilon > 0$, there are constants $M = M(\epsilon)$ and $N = N(\epsilon)$ such that if $n \geq N$, there is a partition $\mathcal{P} = \{V_0, \dots, V_M\}$ of $[n]$, a decomposition $T = T_1 + T_2 + T_3$ where T_1, T_2 and T_3 are self–adjoint and a set of pairs $\Sigma \subset \{0, \dots, M\}^2$ such that*

- *for all $(s,t) \in [M]^2$, there exists d_{st} such that for all $a \in V_s$ and $b \in V_t$,*

$$|(T_1)_{ab} - d_{st}| \leq \epsilon,$$

- *for all $(s,t) \in \{0, \dots, M\}^2 \setminus \Sigma$,*

$$\sum_{a \in V_s} \sum_{b \in V_t} |(T_2)_{ab}|^2 \leq \epsilon^2 |V_s||V_t|,$$

- *for all $(s,t) \in \{0, \dots, M\}^2 \setminus \Sigma$,*

$$n \cdot \mathrm{rad}(T_3) \leq \epsilon |V_s||V_t|,$$

- *and*

$$\sum_{(s,t) \in \Sigma} |V_s||V_t| \leq \epsilon n^2.$$

Proof. Since T is self–adjoint, we can consider the spectral decomposition

$$T = \sum_{j=1}^{n} \lambda_j \mathbf{u}_j \mathbf{u}_j^*.$$

Observe that for each $j \in [n]$, we have that

$$j \lambda_j^2 \leq \lambda_1^2 + \cdots + \lambda_j^2 \leq \mathrm{Tr}(T^2) \leq n^2,$$

and as a consequence,

$$|\lambda_j| \leq \frac{n}{\sqrt{j}}. \tag{1}$$

To find the spectral decomposition of T into three matrices, we are going to use a function $F : \mathbb{N} \to \mathbb{N}$ to be defined later which depends on ϵ that satisfies $F(x) > x$ for all $x \in \mathbb{N}$. If $F^{(k)}$ denotes the composition of F k times, consider the following partition of $[n]$ into $\lfloor 1/\epsilon^3 \rfloor$ intervals:

- $I_k := \left[F^{(k-1)}(1), F^{(k)}(1) - 1 \right]$ for $k = 1, \dots, \lfloor 1/\epsilon^3 \rfloor - 1$, and

- $I_{\lfloor 1/\epsilon^3 \rfloor} := \left[F^{(\lfloor 1/\epsilon^3 \rfloor - 1)}(1), n\right]$.

Choose $K \in \left[\lfloor 1/\epsilon^3 \rfloor\right]$ satisfying

$$\sum_{j \in I_K} \lambda_j^2 \leq \epsilon^3 n^2, \tag{2}$$

which indeed exists, otherwise we would have

$$\sum_{k=1}^{\lfloor 1/\epsilon^3 \rfloor} \sum_{j \in I_k} \lambda_j^2 = \mathrm{Tr}(T^2) > \frac{1}{\epsilon^3} \epsilon^3 n^2 = n^2.$$

Consider the decomposition $T = T_1 + T_2 + T_3$ where

$$\begin{aligned} T_1 &:= \sum_{k=1}^{K-1} \sum_{j \in I_k} \lambda_j \mathbf{u}_j \mathbf{u}_j^*, \\ T_2 &:= \sum_{j \in I_K} \lambda_j \mathbf{u}_j \mathbf{u}_j^*, \\ T_3 &:= \sum_{k=K+1}^{\lfloor 1/\epsilon^3 \rfloor} \sum_{j \in I_k} \lambda_j \mathbf{u}_j \mathbf{u}_j^*, \end{aligned}$$

which are also self–adjoint by definition.

Let $J := F^{(K-1)}(1)$. We are going to define a partition of $[n]$ using the eigenvectors of T_1, namely $\mathbf{u}_1, \ldots, \mathbf{u}_{J-1}$. Consider the square of side $2\sqrt{J/(\epsilon n)}$ centered at the origin of the complex plane and divide it into $4J^4/\epsilon^4$ subsquares of side $\sqrt{\epsilon^3/(J^3 n)}$. For each $\mathbf{u}_j = (u_j(1), \ldots, u_j(n))$, define $\mathcal{P}(\mathbf{u}_j)$ to be the partition of $[n]$ by coloring $a \in [n]$ by the square where $u_j(a)$ belongs to. All values outside the square, precisely the $a \in [a]$ for which either the imaginary or the real part of $u_j(a)$ is larger than $\sqrt{J/(\epsilon n)}$, are contained in an exceptional part Σ_j. If a value lies in the border of a square, assign the index to any of the parts corresponding to an adjacent square.
Let $\mathcal{P} = \{V_0, \ldots, V_M\}$ be the partition of $[n]$ where

$$V_0 := \bigcup_{j=1}^{J-1} \Sigma_j$$

and $\{V_1, \ldots, V_M\}$ is the common refinement of $[n] \setminus V_0$ of partitions $\mathcal{P}(\mathbf{u}_j)$ for $j = 1, \ldots, J-1$.

For $j \in [J-1]$, the partition $\mathcal{P}$ has the following properties:

(*i*) The size of each Σ_j is at most $\epsilon n/J$ because $\|\mathbf{u}_j\|_2 = 1$. Therefore,

$$|V_0| \leq (J-1)\frac{\epsilon n}{J} < \epsilon n.$$

(ii) For every $a \in V_1 \sqcup \cdots \sqcup V_M$ we have that $u_j(a)$ lies inside the big square so its magnitude is upper bounded by

$$|u_j(a)| \leq \sqrt{\frac{2J}{\epsilon n}}.$$

(iii) For each $s > 0$ and $a, b \in V_s$, we have

$$|u_j(a) - u_j(b)| \leq \sqrt{\frac{2\epsilon^3}{J^3 n}}.$$

(iv) The number of parts M is upper bounded by the total number of small squares of all the partitions, namely

$$M \leq \left(\frac{2\sqrt{J/(\epsilon n)}}{\sqrt{\epsilon^3/(J^3 n)}} \right)^{2J} = \left(\frac{2J^2}{\epsilon^2} \right)^{2J} = \left(\frac{4J^4}{\epsilon^4} \right)^{J}. \tag{3}$$

To complete the proof, we need to find a set of pairs $\Sigma \subset \{0, \ldots, M\}^2$ such that

$$\sum_{(s,t) \in \Sigma} |V_s| \, |V_t| \leq \epsilon n^2,$$

and check that the partition $\mathcal{P}$ satisfies the conditions on T_1, T_2 and T_3.

Define $\sigma \subset \{0, \ldots, M\}^2$ as the set of pairs (s, t) such that

$$\sum_{a \in V_s} \sum_{b \in V_t} |(T_2)_{ab}|^2 > \epsilon^2 \, |V_s| \, |V_t| \, .$$

Therefore,

$$\epsilon^2 \sum_{(s,t) \in \sigma} |V_s| \, |V_t| < \sum_{(s,t) \in \sigma} \sum_{a \in V_s} \sum_{b \in V_t} |(T_2)_{ab}|^2 \leq \sum_{a,b \in [n]} |(T_2)_{ab}|^2 \leq \epsilon^3 n^2,$$

where last inequality is a consequence of equation (2) expressed in terms of T_2 and also the symmetry of T_2, namely

$$\sum_{j \in I_K} \lambda_j^2 = \mathrm{Tr}(T_2^2) = \sum_{a,b \in [n]} |(T_2)_{ab}|^2 \leq \epsilon^3 n^2.$$

Let Σ be the set of pairs $(s, t) \in \{0, \ldots, M\}^2$ such that $(s, t) \in \sigma$ or either $s = 0$, $t = 0$ or $\min(|V_s| \, , |V_t|) < \epsilon n / M$. Therefore,

$$\sum_{(s,t) \in \Sigma} |V_s| \, |V_t| \leq \sum_{(s,t) \in \sigma} |V_s| \, |V_t| + 2n \, |V_0| + 2 \sum_{|V_s| < \epsilon n / M} n \, |V_s| \leq \epsilon n^2 + 2\epsilon n^2 + 2M \frac{\epsilon n}{M} n = 5\epsilon n^2 \tag{4}$$

as we wanted.

By definition of σ, for pairs $(s,t) \notin \sigma$ (and consequently, for $(s,t) \notin \Sigma$) we have

$$\sum_{a \in V_s} \sum_{b \in V_t} |(T_2)_{ab}|^2 \leq \epsilon^2 |V_s| |V_t| . \tag{5}$$

Thus, we have proved the desired properties of T_2.

To prove the rest of the theorem, we are going to show first that for any pair $(s,t) \in [M]^2$, the entries of T_1 over the block $V_s \times V_t$ differ among themselves by at most 4ϵ. For any $a, c \in V_s$ and $b, d \in V_t$,

$$\begin{aligned}
|(T_1)_{ab} - (T_1)_{cd}| &= \left| \sum_{j<J} \lambda_j u_j(a) u_j(b) - \lambda_j u_j(c) u_j(d) \right| \\
&\leq \sum_{j<J} |\lambda_j| \, |u_j(a)u_j(b) - u_j(c)u_j(b) + u_j(c)u_j(b) - u_j(c)u_j(d)| \\
&\leq \sum_{j<J} n \left(|u_j(b)| \, |u_j(a) - u_j(c)| + |u_j(c)| \, |u_j(b) - u_j(d)| \right) \\
&\leq Jn \left(2 \cdot \sqrt{\frac{2J}{\epsilon n}} \cdot \sqrt{\frac{2\epsilon^3}{J^3 n}} \right) \\
&= 4\epsilon,
\end{aligned}$$

where we used that $|\lambda_j| \leq n/\sqrt{j} \leq n$.
Now, if d_{st} is the mean of the entries of T_1 in block $V_s \times V_t$, we have

$$\begin{aligned}
|(T_1)_{ab} - d_{st}| &= \left| (T_1)_{ab} - \frac{\sum_{c \in V_s, d \in V_t} (T_1)_{cd}}{|V_s| \, |V_t|} \right| \\
&\leq \frac{1}{|V_s| \, |V_t|} \left| |V_s| \, |V_t| \, (T_1)_{ab} - \sum_{(c,d) \in V_s \times V_t} (T_1)_{cd} \right| \\
&\leq \frac{1}{|V_s| \, |V_t|} \sum_{(c,d) \in V_s \times V_t} |(T_1)_{ab} - (T_1)_{cd}| \\
&= \frac{1}{|V_s| \, |V_t|} \sum_{\substack{(c,d) \in V_s \times V_t \\ (c,d) \neq (a,b)}} |(T_1)_{ab} - (T_1)_{cd}| \\
&\leq \frac{|V_s| \, |V_t| - 1}{|V_s| \, |V_t|} \cdot 4\epsilon \\
&< 4\epsilon.
\end{aligned} \tag{6}$$

Thirdly, for the condition on T_3, note that all of its eigenvalues satisfy inequality (1). In particular, for the spectral radius of T_3 we have

$$n \cdot \mathrm{rad}(T_3) = n \left| \lambda_{F(J)} \right| \leq \frac{n^2}{\sqrt{F(J)}}.$$

It only remains to check that $n^2/\sqrt{F(J)} \leq \epsilon |V_s| |V_t|$ for all pairs $(s,t) \notin \Sigma$. Since $|V_s|, |V_t| \geq \epsilon n/M$ for such pairs, it suffices to choose F so that $F(J) \geq M^4/\epsilon^6$. We already have seen in (3) that $M \leq (4J^4/\epsilon^4)^J$, so if F satisfies

$$F(x) \geq \frac{1}{\epsilon^6} \left(\frac{4x}{\epsilon^4}\right)^{4x},$$

we get the condition on T_3, namely

$$n \cdot \mathrm{rad}(T_3) \leq \frac{n^2}{\sqrt{F(J)}} \leq \frac{\epsilon^3 n^2}{M^2} \leq \epsilon |V_s| |V_t|. \tag{7}$$

To sum up, if we choose $\epsilon/5$ instead of ϵ, inequalities (4), (6), (5) and (7) become

$$\begin{aligned} \sum_{(s,t)\in\Sigma} |V_s| |V_t| &\leq \frac{5}{5}\epsilon n^2 = \epsilon n^2, \\ |(T_1)_{ab} - d_{st}| &\leq \frac{4}{5}\epsilon < \epsilon, \\ \sum_{a\in V_s} \sum_{b\in V_t} |(T_2)_{ab}|^2 &\leq \frac{\epsilon^2}{5^2} |V_s| |V_t| < \epsilon^2 |V_s| |V_t|, \\ n \cdot \mathrm{rad}(T_3) &\leq \frac{\epsilon}{5} |V_s| |V_t| < \epsilon |V_s| |V_t|, \end{aligned}$$

which complete the proof. □

Observation 3.2. Note that, for ϵ and F fixed, the maximum value K that may be chosen is $\lfloor 1/\epsilon^3 \rfloor$, and hence $J_{\max} = F^{(\lfloor 1/\epsilon^3 \rfloor - 1)}(1)$. Therefore, a bound for the number of parts M valid for any self–adjoint matrix T of size $n \geq N(\epsilon)$ is

$$M \leq \left(\frac{4J_{\max}^4}{\epsilon^4}\right)^{J_{\max}}.$$

However, if there exists another $J < J_{\max}$ for which the sum of squares of eigenvalues with indexes in the interval $\left[J, F(J)-1\right]$ is small enough, then the bound (3) is significantly smaller. For that reason, if we choose $J = F^{(K-1)}(1)$ for the first $K \in \left[\lfloor 1/\epsilon^3 \rfloor\right]$ for which

$$\sum_{j\in I_K} \lambda_j^2 \leq \epsilon^3 n^2$$

holds, we are obtaining a partition with the minimum number of parts (using this particular spectral method).

For instance, consider a matrix T with $\mathrm{Tr}(T^2) \leq (1-\epsilon^3)n^2$ and let I_K be an interval for which (2) holds. Then

$$\mathrm{Tr}(T^2) > \sum_{k=1}^{K-1} \sum_{j\in I_k} \lambda_j^2 + \sum_{k=K+1}^{\lfloor 1/\epsilon^3 \rfloor} \sum_{j\in I_k} \lambda_j^2 > \left(\frac{1}{\epsilon^3} - 1\right) \epsilon^3 n^2 = (1-\epsilon^3)n^2,$$

so there must be another interval I_k for which (2) also holds, and therefore the same algorithm provides two distinct ϵ–regular partitions.

Considering this observation, the constants depending on ϵ of the statement are clearly determined:

- $N(\epsilon) := F^{(\lfloor 1/\epsilon^3 \rfloor - 1)}(1)$, otherwise the last part of the partition is empty, and
- $M(\epsilon) := \left(\frac{4J_{\max}^4}{\epsilon^4}\right)^{J_{\max}} + 1$.

Observation 3.3 (Basis of real eigenvectors)**.** Note that in the special case when $\mathbf{u}_1, \dots, \mathbf{u}_{J-1} \in \mathbb{R}^n$, then for all $a \in [n]$ and $j \in [J-1]$, $u_j(a)$ lies in $\mathbb{R}$, which is the border of $4J^2/\epsilon^2$ squares. Therefore, the indexes can be distributed into at most $1 + 2J^2/\epsilon^2$ parts, which is a significantly better bound for the total number of parts.

Actually, in the case of a real symmetric matrix T, we can always find an orthonormal basis of eigenvectors in $\mathbb{R}^n$. Let $\mathbf{u} \in \mathbb{C}^n$ be an eigenvector of $\lambda \neq 0$, we have that

$$(T - \lambda I)\mathbf{u} = (T - \lambda I)\Re(\mathbf{u}) + (T - \lambda I)\Im(\mathbf{u})i = 0$$

where $\Re(\mathbf{u})$ and $\Im(\mathbf{u})$ represent the real and imaginary parts of the vector $\mathbf{u}$ respectively. Thus $\Re(\mathbf{u})$ and $\Im(\mathbf{u})$ are both real eigenvectors for the eigenvalue λ.

Let $\mathbf{u}_1, \dots, \mathbf{u}_n \in \mathbb{C}^n$ be a basis of eigenvectors. If we take their real and imaginary parts we still have a set which spans the whole space of eigenvectors, and a maximal linearly independent subset provides us with a basis of real eigenvectors. It might not be orthogonal for eigenvectors of the same eigenvalue, but we may use the Gram–Schmidt process to find an orthonormal basis.

Observation 3.4. The hypothesis $\mathrm{Tr}(T^2) \leq n^2$ is needed to make sure that

(i) the bound on the eigenvalues (1) holds, and

(ii) there exists a K for which the inequality (2) holds,

so we could replace them by these two conditions without changing the result. For instance, take a self–adjoint matrix T with spectrum $\lambda_1, \dots, \lambda_n$ such that $\mathrm{Tr}(T^2) \leq n^2$ and assume there exists $K > 1$ for which (2) holds, and consider an orthonormal basis of eigenvectors $\mathbf{u}_1 \dots, \mathbf{u}_n$. We can construct a self–adjoint $n \times n$ matrix $\widetilde{T}$ that satisfies (i) and (ii) but $\mathrm{Tr}(\widetilde{T}^2) > n^2$ as follows:

$$\widetilde{T} = n\mathbf{u}_1\mathbf{u}_1^* + \sum_{j>1}^{n} \lambda_j \mathbf{u}_j \mathbf{u}_j^*,$$

a matrix with eigenvalues $\widetilde{\lambda}_1 = n$ and $\widetilde{\lambda}_j = \lambda_j$ for $j > 1$ which satisfies

$$\mathrm{Tr}(\widetilde{T}^2) = \sum_{j=1}^{n} \widetilde{\lambda}_j^2 = n^2 + \sum_{j=2}^{n} \widetilde{\lambda}_j^2 > n^2,$$

and both (i) and (ii) clearly hold by definition of $\widetilde{T}$.

Observation 3.5 (Refinement of the spectral partition)**.** Let $\mathcal{Q} = \{W_0, \dots, W_{\widetilde{M}}\}$ be a refinement of $\mathcal{P} = \{V_0, \dots, V_M\}$, where $\widetilde{M} = M + \Delta$ for some natural number $\Delta > 0$ and the first m parts of $\mathcal{Q}$ are the refinement of V_0, namely

$$V_0 = \bigcup_{s=0}^{m} W_s.$$

Define $\widetilde{\sigma} \subset \{0, \dots, \widetilde{M}\}^2$ as the set of pairs (s,t) such that

$$\sum_{a \in W_s} \sum_{b \in W_t} \left|(T_2)_{ab}\right|^2 > \epsilon^2 \left|W_s\right| \left|W_s\right|.$$

Let $\widetilde{\Sigma}$ be the set of pairs $(s,t) \in \{0, \dots, \widetilde{M}\}^2$ such that $(s,t) \in \widetilde{\sigma}$ or either $s \in \{0, \dots, m\}$, $t \in \{0, \dots, m\}$ or $\min(\left|W_s\right|, \left|W_t\right|) < \epsilon n / \widetilde{M}$.

- For all $(s,t) \in \{m+1, \dots, \widetilde{M}\}^2$, there exists $\widetilde{d_{st}} \in \mathbb{C}$ such that for all $a \in W_s$ and $b \in W_t$ we have
$$\left|(T_1)_{ab} - \widetilde{d_{st}}\right| \leq \epsilon,$$
because in particular, $a \in W_s \in V_{s'}$ and $b \in W_t \in V_{t'}$ for some $(s', t') \in [M]^2$. As a consequence, we can choose either $\widetilde{d_{st}} = d_{s't'}$ or the mean of the entries of T_1 on block $W_s \times W_t$.

- For all $(s,t) \in \{0, \dots, \widetilde{M}\}^2 \setminus \widetilde{\Sigma}$, by definition of $\widetilde{\sigma}$ we have
$$\sum_{a \in W_s} \sum_{b \in W_t} \left|(T_2)_{ab}\right|^2 \leq \epsilon^2 \left|W_s\right| \left|W_t\right|.$$

- For all $(s,t) \in \{0, \dots, \widetilde{M}\}^2 \setminus \widetilde{\Sigma}$, we have
$$n \cdot \mathrm{rad}(T_3) = n \left|\lambda_{F(J)}\right| \leq \frac{n^2}{\sqrt{F(J)}},$$
but we want $n^2/\sqrt{F(J)} \leq \epsilon \left|W_s\right| \left|W_t\right|$ for all pairs $(s,t) \notin \widetilde{\Sigma}$. Since $\left|W_s\right|, \left|W_t\right| \geq \epsilon n / \widetilde{M}$ for such pairs, if we chose F so that $F(J) \geq \widetilde{M}^4/\epsilon^6 = (M + \Delta)^4/\epsilon^6$, for instance
$$F(x) \geq \frac{1}{\epsilon^6} \left(\left(\frac{4x^4}{\epsilon^4} \right)^x + \Delta \right)^4,$$
then the refinement $\mathcal{Q}$ satisfies
$$n \cdot \mathrm{rad}(T_3) \leq \epsilon \left|W_s\right| \left|W_t\right|.$$

- Analogously to (4), by choosing an appropriate ϵ we have
$$\sum_{(s,t) \in \widetilde{\Sigma}} \left|W_s\right| \left|W_t\right| \leq \epsilon n^2.$$

In conclusion, if we apply Theorem 3.1 to a matrix T and $F(J) \geq \widetilde{M}^4/\epsilon^6$ for some $\widetilde{M}$, then we can refine the obtained partition $\mathcal{P}$ into at most $\widetilde{M}$ parts and find a set $\widetilde{\Sigma} \subset \{0, \ldots, \widetilde{M}\}^2$ such that fulfill the required properties on T_1, T_2 and T_3.

In other words, not any refinement $\mathcal{Q}$ of $\mathcal{P}$ is a suitable partition for T: for the property on T_3 we need that the F used in Theorem 3.1 is such that $F(J) \geq |\mathcal{Q}|^4/\epsilon^6$.

Although we could use last theorem to prove Szemerédi's Regularity Lemma with spectral techniques, we are going to prove an intermediate statement that brings the above result a little bit closer to the graph ϵ–regularity defined in Chapter 1.

Theorem 3.6 (Spectral Regularity Theorem)**.** *For every $\epsilon > 0$, there are constants $M = M(\epsilon)$ and $N = N(\epsilon)$ such that every $n \times n$ self–adjoint matrix T with $n \geq N$ such that $\mathrm{Tr}(T^2) \leq n^2$ has a partition of $[n]$, $\{V_0, \ldots, V_M\}$, and an exceptional set of pairs $\Sigma \subset \{0, \ldots, M\}^2$ such that*

$$\sum_{(s,t)\in\Sigma} |V_s||V_t| \leq \epsilon n^2,$$

and for all pairs $(s,t) \notin \Sigma$ and vectors $\mathbf{v}_A, \mathbf{v}_B \in \mathbb{C}^n$ such that $\mathrm{supp}(\mathbf{v}_A) \subset V_s$, $\mathrm{supp}(\mathbf{v}_B) \subset V_t$, $\|\mathbf{v}_A\|_2^2 \leq |V_s|$ and $\|\mathbf{v}_B\|_2^2 \leq |V_t|$, we have

$$|\mathbf{v}_B^*(T - d(V_s, V_t)J_n)\mathbf{v}_A| \leq \epsilon\,|V_s|\,|V_t|\,,$$

where $d(V_s, V_t)$ is the mean of the entries in the block matrix $V_s \times V_t$.

Proof. Apply Theorem 3.1 to matrix T. Let $\mathcal{P}$ be the partition obtained and Σ the set of irregular pairs. Observe that

$$\sum_{(s,t)\in\Sigma} |V_s|\,|V_t| \leq \epsilon n^2,$$

is a direct consequence of Theorem 3.1. To complete the proof, we will see that for pairs $(s,t) \notin \Sigma$ and vectors $\mathbf{v}_A, \mathbf{v}_B \in \mathbb{C}^n$ as in the statement, we have

$$\begin{aligned}|\mathbf{v}_B^*\left(T - d(V_s, V_t)J_n\right)\mathbf{v}_A| &\leq |\mathbf{v}_B^*\left(T_1 - d_{st}J_n\right)\mathbf{v}_A| + |\mathbf{v}_B^* T_2 \mathbf{v}_A| + |\mathbf{v}_B^* T_3 \mathbf{v}_A| + |\mathbf{v}_B^*\left(d_{st} - d(V_s, V_t)\right)J_n\mathbf{v}_A| \\ &\leq \epsilon\,|V_s|\,|V_t|\end{aligned}$$

for some number d_{st}. The proof will consist in finding a bound for each summand.

For the first one, if d_{st} is the mean in the entries of T_1 in block $V_s \times V_t$, observe that

$$|\mathbf{v}_B^*(T_1 - d_{st}J_n)\mathbf{v}_A| \leq \sum_{a\in V_s}\sum_{b\in V_t} |(T_1)_{ab} - d_{st}|\,|v_A(a)|\,|v_B(b)| < \epsilon\,\|\mathbf{v}_A\|_1\,\|\mathbf{v}_B\|_1 \leq \epsilon\,|V_s|\,|V_t|\,, \quad (8)$$

where we applied the triangle inequality, the condition on T_1 given by Theorem 3.1 and the norm inequality given by Cauchy–Schwarz, namely

$$\begin{aligned}
\|\mathbf{v}\|_1 &= \sum_{j\in\operatorname{supp}(\mathbf{v})} |v_j| \\
&= \sum_{j\in\operatorname{supp}(\mathbf{v})} |v_j|\cdot 1 \\
&\leq \left(\sum_{j\in\operatorname{supp}(\mathbf{v})} \sqrt{|v_j|^2}\right)^{1/2} \left(\sum_{j\in\operatorname{supp}(\mathbf{v})} 1^2\right)^{1/2} \\
&= \sqrt{|\operatorname{supp}(\mathbf{v})|}\,\|\mathbf{v}\|_2\,,
\end{aligned}$$

which applied to $\mathbf{v}_A$ and $\mathbf{v}_B$ is precisely

$$\begin{aligned}
\|\mathbf{v}_A\|_1 &\leq \sqrt{V_s}\,\|\mathbf{v}_A\|_2 \leq |V_s|\,, \\
\|\mathbf{v}_B\|_1 &\leq \sqrt{V_t}\,\|\mathbf{v}_B\|_2 \leq |V_t|\,.
\end{aligned}$$

Secondly, for any $(s,t)\notin\Sigma$, we apply Cauchy–Schwarz and the condition on T_2 given by Theorem 3.1 and obtain

$$\begin{aligned}
|\ \mathbf{v}_B^* T_2 \mathbf{v}_A|^2 &= \left|\sum_{a\in V_s}\sum_{b\in V_t}(T_2)_{ab} v_A(a)\overline{v_B(b)}\right|^2 \\
&\leq \left(\sum_{a\in V_s}\sum_{b\in V_t}|(T_2)_{ab}|^2\right)\left(\sum_{a\in V_s}\sum_{b\in V_t}|v_A(a)|^2|v_B(b)|^2\right) \\
&\leq \epsilon^2\,|V_s|\,|V_t|\,\|\mathbf{v}_A\|_2\,\|\mathbf{v}_B\|_2 \\
&\leq \epsilon^2\,|V_s|^2\,|V_t|^2\,,
\end{aligned}$$

which implies

$$|\mathbf{v}_B^* T_2 \mathbf{v}_A| \leq \epsilon\,|V_s|\,|V_t|\,. \tag{9}$$

For the third term, if $\operatorname{rad}(T_3)$ is the spectral radius of T_3, we have

$$|\mathbf{v}_B^* T_3 \mathbf{v}_A| \leq \operatorname{rad}(T_3)\,\|\mathbf{v}_A\|_2\,\|\mathbf{v}_B\|_2 \leq n\cdot\operatorname{rad}(T_3) \leq \epsilon\,|V_s|\,|V_t|\,, \tag{10}$$

where the first inequality is a consequence of Cauchy–Schwarz and the second one is given by Theorem 3.1.

Finally, observe that by the triangle inequality,

$$\begin{aligned}
|d_{st} - d(V_s,V_t)| &= \frac{1}{|V_s|\,|V_t|}\left|\mathbf{1}_{V_s}^T(T_1 - T)\mathbf{1}_{V_t}\right| \\
&= \frac{1}{|V_s|\,|V_t|}\left|\mathbf{1}_{V_s}^T(T_2 + T_3)\mathbf{1}_{V_t}\right| \\
&\leq \frac{1}{|V_s|\,|V_t|}\left(\left|\mathbf{1}_{V_s}^T T_2 \mathbf{1}_{V_t}\right| + \left|\mathbf{1}_{V_s}^T T_3 \mathbf{1}_{V_t}\right|\right) \\
&\leq 2\epsilon.
\end{aligned} \tag{11}$$

Making use of the bounds (8), (9), (10), (11), we have that for all pairs $(s,t) \notin \Sigma$, and for any $\mathbf{v}_A$, $\mathbf{v}_B \in \mathbb{C}^n$ as before we have

$$\begin{aligned} \left|\mathbf{v}_B^* \left(T - d(V_s,V_t)J_n\right)\mathbf{v}_A\right| &\leq \left|\mathbf{v}_B^* \left(T_1 - d_{st}J_n\right)\mathbf{v}_A\right| + \left|\mathbf{v}_B^* T_2 \mathbf{v}_A\right| + \left|\mathbf{v}_B^* T_3 \mathbf{v}_A\right| + \left|\mathbf{v}_B^* \left(d_{st} - d(V_s,V_t)\right) J_n \mathbf{v}_A\right| \\ &\leq \left(\epsilon + \epsilon + \epsilon + 2\epsilon\right) |V_s|\,|V_t| \\ &= 5\epsilon\,|V_s|\,|V_t|\,. \end{aligned}$$

Choosing $\epsilon/5$ instead of ϵ, we obtain the desired result. $\square$

Observation 3.7 (Definition of ϵ–regular partition of a matrix)**.** If we choose ϵ^3 instead of ϵ in Theorem 3.6, we obtain

$$\sum_{(s,t)\in\Sigma} |V_s||V_t| \leq \epsilon^3 n^2 \leq \epsilon n^2,$$

for irregular pairs and

$$|\mathbf{v}_B^*(T - d(V_s,V_t)J_n)\mathbf{v}_A| \leq \epsilon^3\,|V_s|\,|V_t|\,,$$

for regular pairs. Whenever $\|\mathbf{v}_A\|_2^2 \geq \epsilon\,|V_s|$ and $\|\mathbf{v}_B\|_2^2 \geq \epsilon\,|V_t|$, if we divide both sides of the inequality by the product of the square of the norms of $\mathbf{v}_A$ and $\mathbf{v}_B$ we have that

$$\frac{|\mathbf{v}_B^*(T - d(V_s,V_t)J_n)\mathbf{v}_A|}{\|\mathbf{v}_A\|_2^2\,\|\mathbf{v}_B\|_2^2} \leq \epsilon^3 \frac{|V_s|\,|V_t|}{\|\mathbf{v}_A\|_2^2\,\|\mathbf{v}_B\|_2^2} \leq \epsilon.$$

From now on, we will call such partition an ϵ–regular partition for a self–adjoint matrix T.

Observation 3.8. Let $J = F^{(K-1)}(1) \in \mathbb{N}$ and vectors $\mathbf{u}_1, \ldots, \mathbf{u}_{J-1} \in \mathbb{C}^n$ be fixed. The spectral proof gives the same partition for all matrices T with spectrum $\lambda_1, \ldots, \lambda_n$ for which

(i) $\sum\limits_{j\in I_K} \lambda_j^2 \leq \epsilon^3 n^2$, and

(ii) $\mathbf{u}_1, \ldots, \mathbf{u}_{J-1}$ are the first $J-1$ vectors of an orthonormal basis of eigenvectors of T.

Let us now study how the partition changes when the spectrum, the J or the basis of eigenvectors vary. Let T and $\widetilde{T}$ be two self-adjoint $n\times n$ matrices such that $\mathrm{Tr}(T^2), \mathrm{Tr}(\widetilde{T}^2) \leq n^2$. Let $\lambda_1, \ldots, \lambda_n$ and $\widetilde{\lambda}_1, \ldots, \widetilde{\lambda}_n$ be their spectrums and $\mathbf{u}_1, \ldots, \mathbf{u}_n$ and $\widetilde{\mathbf{u}}_1, \ldots, \widetilde{\mathbf{u}}_n$ orthonormal basis of eigenvectors of T and $\widetilde{T}$ respectively. Let $J = F^{(K-1)}(1)$ and $\widetilde{J} = F^{(\widetilde{K}-1)}(1)$ be such that

$$\sum_{j\in I_K} \lambda_j^2 \leq \epsilon^3 n^2$$

and

$$\sum_{j\in I_{\widetilde{K}}} \widetilde{\lambda}_j^2 \leq \epsilon^3 n^2.$$

We remark the following situations, where the partitions obtained from the spectral proof are somehow related:

(a) $J = \widetilde{J}$ and $\{\mathbf{u}_1, \dots, \mathbf{u}_{J-1}\} = \{\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{J-1}\}$: the ϵ–regular spectral partitions obtained are the same.

(b) $J = \widetilde{J}$ and $\{\mathbf{u}_1, \dots, \mathbf{u}_{J-1}\} \neq \{\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{J-1}\}$: the bound (3) for the number of parts is the same, because the complex plane is split likewise. The coordinates of $\{\mathbf{u}_1, \dots, \mathbf{u}_{J-1}\}$ and $\{\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{J-1}\}$ may be different, so the ϵ–regular partitions may vary for the indexes for which the value of the corresponding coordinate lies in different parts of the complex plane.

(c) $J < \widetilde{J}$ and $\{\mathbf{u}_1, \dots, \mathbf{u}_{J-1}\} = \{\tilde{\mathbf{u}}_1, \dots, \tilde{\mathbf{u}}_{J-1}\}$: the complex plane is split into different parts, so the partition may not be suitable for both matrices. In the very specific case when $\widetilde{J} = c^{2/3}J$ for some constant $c \in \mathbb{N}$, then each square of the complex plane defined by J is divided into c^2 subsquares in the partition for $\widetilde{J}$. Then, the ϵ–regular partition for $\widetilde{T}$ is a refinement of the partition for T.

Having in mind that our goal is to prove regularity theorems for graphs with a coloring or multicoloring of the edges, let us give the proof of the analogous versions for matrices. The result follows the proof by Robertson [12].

Theorem 3.9. *For every $\epsilon > 0$, there are constants $M = M(\epsilon, r)$ and $N = N(\epsilon, r)$ such that for every set of self–adjoint $n \times n$ matrices $\{T^{[1]}, \dots, T^{[r]}\}$ with $n \geq N$ such that $\sum_{i=1}^{r} \operatorname{Tr}\left((T^{[i]})^2\right) \leq n^2$ there exists a partition of $[n]$ which is ϵ–regular for $T^{[1]}, \dots, T^{[r]}$ simultaneously.*

Proof. For $i = 1, \dots, r$, let $\lambda_1^{[i]}, \dots, \lambda_n^{[i]} \in \mathbb{R}$ be the eigenvalues of $T^{[i]}$ and consider an orthonormal basis of eigenvectors $\mathbf{u}_1^{[i]}, \dots, \mathbf{u}_n^{[i]} \in \mathbb{C}^n$.

We will prove the theorem in three steps:

1. Find an appropriate $J \in \mathbb{N}$.

2. Apply Theorem 3.6 to each $T^{[i]}$ using the J from last step and refine to obtain a partition $\mathcal{P}$.

3. Check that $\mathcal{P}$ is ϵ–regular for $T^{[i]}$ for $i = 1, \dots, r$.

Let us begin with the definition of a proper $J \in \mathbb{N}$. Let $F : \mathbb{N} \to \mathbb{N}$ be a function to be defined later which depends on ϵ and r that satisfies $F(x) > x$ for all $x \in \mathbb{N}$. Consider the partition of $[n]$ into intervals $I_1, \dots I_{\lfloor 1/\epsilon^3 \rfloor}$ as in Theorem 3.1. Define $J := F^{(K-1)}(1)$ for some $K \in \left[\lfloor 1/\epsilon^3 \rfloor\right]$ satisfying

$$\sum_{i=1}^{r} \sum_{j \in I_K} (\lambda_j^{[i]})^2 \leq \epsilon^3 n^2.$$

Note that this K exists, otherwise we would have

$$\sum_{k=1}^{\lfloor 1/\epsilon^3 \rfloor} \sum_{i=1}^{r} \sum_{j \in I_k} (\lambda_j^{[i]})^2 = \sum_{i=1}^{r} \operatorname{Tr}\left((T^{[i]})^2\right) > \frac{1}{\epsilon^3}\epsilon^3 n^2 = n^2.$$

For each $i = 1, \ldots, r$, apply Theorem 3.6 to $T^{[i]}$ and let $\mathcal{P}^{[i]} := \{V_0^{[i]}, \ldots, V_{M^{[i]}}^{[i]}\}$ be the partitions we obtain. Let $\mathcal{P} := \{V_0, \ldots, V_M\}$ be the partition where

$$V_0 := \bigcup_{i=1}^{r} V_0^{[i]},$$

and $\{V_1, \ldots, V_M\}$ is a refinement of partitions $\mathcal{P}^{[i]}$ for indexes in $[n] \setminus V_0$. Observe that

$$|V_0| \leq \sum_{i=1}^{r} |V_0^{[i]}| \leq \sum_{i=1}^{r} (J-1)\frac{\epsilon n}{J} < r\epsilon n.$$

The total number of parts obtained in the refinement is

$$M \leq \prod_{i=1}^{r} M^{[i]} \leq \left(\frac{4J^4}{\epsilon^4}\right)^{rJ}. \tag{12}$$

Consider the decomposition $T^{[i]} = T_1^{[i]} + T_2^{[i]} + T_3^{[i]}$ as in Theorem 3.1. To complete the proof, we will check the ϵ–regularity of the partition $\mathcal{P}$ for each $T^{[i]}$. The proof is analogous to Theorems 3.1 and 3.6 but with partition $\mathcal{P}$ instead of $\mathcal{P}^{[i]}$.

We are going to check first the condition on irregular pairs. Observe that

$$\sum_{j \in I_K} {\lambda_j^{[i]}}^2 = \operatorname{Tr}\left((T_2^{[i]})^2\right) = \sum_{a,b \in [n]} |(T_2^{[i]})_{ab}|^2 \leq \epsilon^3 n^2.$$

Define $\sigma^{[i]} \subset [M]^2$ so that for all pairs $(s,t) \notin \sigma^{[i]}$,

$$\sum_{a \in V_s} \sum_{b \in V_t} |(T_2^{[i]})_{ab}|^2 \leq \epsilon^2 |V_s| |V_t|. \tag{13}$$

Let $\Sigma^{[i]}$ be the set of all irregular pairs for $T^{[i]}$: $(s,t) \in \{0, \ldots, M\}^2$ such that $(s,t) \in \sigma^{[i]}$, $s = 0$, $t = 0$ or $\min(|V_s|, |V_t|) < \epsilon n/M$. Therefore,

$$\begin{aligned} \sum_{(s,t) \in \Sigma^{[i]}} |V_s| |V_t| &\leq \sum_{(s,t) \in \sigma^{[i]}} |V_s| |V_t| + 2n |V_0| + 2 \sum_{|V_s| < \epsilon n/M} n |V_s| \\ &\leq \epsilon n^2 + 2r\epsilon n^2 + 2M\frac{\epsilon n}{M} \\ &= (3 + 2r)\epsilon n^2. \end{aligned} \tag{14}$$

For the ϵ–regularity of the pairs $(s,t) \notin \Sigma^{[i]}$, we are going to check that for any vectors $\mathbf{v}_A, \mathbf{v}_B \in \mathbb{C}^n$ as in the statement,

$$\begin{aligned}&\left|\mathbf{v}_B^* \left(T^{[i]} - d^{[i]}(V_s,V_t)J_n\right)\mathbf{v}_A\right| \\ &\quad \le \left|\mathbf{v}_B^* \left(T_1^{[i]} - d_{st}^{[i]}J_n\right)\mathbf{v}_A\right| + \left|\mathbf{v}_B^* T_2^{[i]}\mathbf{v}_A\right| + \left|\mathbf{v}_B^* T_3^{[i]}\mathbf{v}_A\right| + \left|\mathbf{v}_B^* \left(d_{st}^{[i]} - d^{[i]}(V_s,V_t)\right)J_n\mathbf{v}_A\right|,\end{aligned}$$

for real numbers $d_{st}^{[i]}$ (the mean of the entries of $T_1^{[i]}$ in block $V_s \times V_t$).

Clearly, if the bound (8) for the term of $T_1^{[i]}$ holds for $V_{s'}^{[i]}, V_{t'}^{[i]} \in \mathcal{P}^{[i]}$ for some pair (s',t'), it also holds for the refinement $\mathcal{P}$ because in particular, $\mathrm{supp}(v_A) \subset V_s \subseteq V_{s'}^{[i]}$ and $\mathrm{supp}(\mathbf{v}_B) \subset V_t \subseteq V_{t'}^{[i]}$. Therefore, we have

$$\left|\mathbf{v}_B^*(T_1^{[i]} - d_{st}^{[i]}J_n)\mathbf{v}_A\right| \le \epsilon\,|V_s|\,|V_t|\,. \tag{15}$$

Similarly, it is easy to see that the bound (9) for $T_2^{[i]}$ also holds for $V_s, V_t \in \mathcal{P}$ since we have specifically defined $\sigma^{[i]}$ to satisfy (13). For that reason, we have

$$\left|\mathbf{v}_B^* T_2^{[i]}\mathbf{v}_A\right| \le \epsilon\,|V_s|\,|V_t|\,. \tag{16}$$

For the term of $T_3^{[i]}$, the F defined in Theorem 3.1 is not valid because in this case the bound for M is (12), which is weaker than (3). Then, to make sure that

$$\left|\mathbf{v}_B^* T_3^{[i]}\mathbf{v}_A\right| \le \epsilon\,|V_s|\,|V_t|\,,$$

it suffices to choose an F satisfying

$$F(x) \ge \frac{1}{\epsilon^6}\left(\frac{4x^4}{\epsilon^4}\right)^{4rx},$$

because then, since $|V_s|\,,|V_t| \ge \epsilon n/M$ for all pairs $(s,t) \notin \Sigma^{[i]}$, we have

$$\left|\mathbf{v}_B^* T_3^{[i]}\mathbf{v}_A\right| \le n\cdot \mathrm{rad}(T_3^{[i]}) \le \frac{n^2}{\sqrt{F(J)}} \le \frac{\epsilon^3 n^2}{M^2} \le \epsilon\,|V_s|\,|V_t|\,. \tag{17}$$

Finally, observe that the bound (11) is also satisfied by the refinement partition $\mathcal{P}$, namely

$$\left|d_{st}^{[i]} - d^{[i]}(V_s,V_t)\right| \le 2\epsilon. \tag{18}$$

Making use of the bounds (15), (16), (17) and (18), we have that for all pairs $(s,t) \notin \Sigma^{[i]}$, and for any vectors $\mathbf{v}_A, \mathbf{v}_B \in \mathbb{C}^n$ as before, we have

$$\left|\mathbf{v}_B^* \left(T^{[i]} - d^{[i]}(V_s,V_t)J_n\right)\mathbf{v}_A\right| \le 5\epsilon\,|V_s|\,|V_t|\,.$$

To conclude, if we do all the computations choosing $\frac{\epsilon^3}{3+2r}$ instead of ϵ, last inequality becomes

$$\left|\mathbf{v}_B^*\left(T^{[i]} - d^{[i]}(V_s, V_t)J_n\right)\mathbf{v}_A\right| \leq \frac{5}{3+2r}\epsilon^3 \left|V_s\right| \left|V_t\right| \leq \epsilon^3 \left|V_s\right| \left|V_t\right|,$$

and the conditon on irregular pairs (14) becomes

$$\sum_{(s,t)\in\Sigma^{[i]}} |V_s|\,|V_t| \leq \frac{3+2r}{3+2r}\epsilon^3 n^2 < \epsilon n^2.$$

Therefore, $\mathcal{P}$ is an ϵ–regular partition for $T^{[1]}, \ldots, T^{[r]}$ simultaneously. $\square$

In next section we will see that matrices of monochromatic subgraphs of a graph with an edge–coloring are a particular case of last theorem. In order to obtain a regularity theorem for edge–multicolorings, we need the following statement which is very similar to the previous one, but with a weaker condition on the traces of $T^{[i]}$. In this case, we require $\mathrm{Tr}(T^{[i]}) \leq n^2$ for all $i = 1, \ldots, r$ instead of $\sum_{i=1}^{r} \mathrm{Tr}\left((T^{[i]})^2\right) \leq n^2$. Although the proof is almost identical, we are going to review it to remark the differences on a few constants.

Theorem 3.10. *For every $\epsilon > 0$, there are constants $M = M(\epsilon, r)$ and $N = N(\epsilon, r)$ such that for every set of self–adjoint $n\times n$ matrices $\{T^{[1]}, \ldots, T^{[r]}\}$ with $n \geq N$ such that $\mathrm{Tr}\left((T^{[i]})^2\right) \leq n^2$ there exists a partition of $[n]$ which is ϵ–regular for $T^{[1]}, \ldots, T^{[r]}$ simultaneously.*

Proof. For $i = 1, \ldots, r$, let $\lambda_1^{[i]}, \ldots, \lambda_n^{[i]} \in \mathbb{R}$ be the eigenvalues of $T^{[i]}$ and consider an orthonormal basis of eigenvectors $\mathbf{u}_1^{[i]}, \ldots, \mathbf{u}_n^{[i]} \in \mathbb{C}^n$.

The structure of the proof is exactly the one in Theorem 3.9. However, there are differences in some of the inequalities:

- the condition for the definition of J,
- the bound for trace of $(T_2^{[i]})^2$ (and consequently, the bound for pairs in $\Sigma^{[i]}$), and
- the appropriate value of ϵ that completes the proof.

Let us begin with the definition of a proper $J \in \mathbb{N}$. Let $F : \mathbb{N} \to \mathbb{N}$ be a function to be defined later which depends on ϵ and r that satisfies $F(x) > x$ for all $x \in \mathbb{N}$. Consider the partition of $[n]$ into intervals $I_1, \ldots, I_{\lfloor 1/\epsilon^3 \rfloor}$ as in Theorem 3.1. Define $J := F^{(K-1)}(1)$ for some $K \in \left[\lfloor 1/\epsilon^3 \rfloor\right]$ satisfying

$$\sum_{i=1}^{r} \sum_{j\in I_K} (\lambda_j^{[i]})^2 \leq r\epsilon^3 n^2.$$

Note that this K exists, otherwise we would have

$$\sum_{k=1}^{\lfloor 1/\epsilon^3 \rfloor} \sum_{i=1}^{r} \sum_{j\in I_k} (\lambda_j^{[i]})^2 = \sum_{i=1}^{r} \mathrm{Tr}\left((T^{[i]})^2\right) > \frac{1}{\epsilon^3} r\epsilon^3 n^2 = rn^2.$$

For each $i = 1, \dots, r$, apply Theorem 3.6 to $T^{[i]}$. Define the partition $\mathcal{P} := \{V_0, \dots, V_M\}$ and the decomposition as in Theorem 3.9. We will check the ϵ–regularity of the partition $\mathcal{P}$ for each $T^{[i]}$.

We are going to check first the condition on irregular pairs. Observe that

$$\sum_{j \in I_K} (\lambda_j^{[i]})^2 = \mathrm{Tr}\left((T_2^{[i]})^2\right) = \sum_{a,b \in [n]} |(T_2^{[i]})_{ab}|^2 \leq r\epsilon^3 n^2.$$

Define $\sigma^{[i]} \subset [M]^2$ as in Theorem 3.9 and let $\Sigma^{[i]}$ be the set of all irregular pairs for $T^{[i]}$: $(s,t) \in \{0, \dots, M\}^2$ such that $(s,t) \in \sigma^{[i]}$, $s = 0$, $t = 0$ or $\min(|V_s|, |V_t|) < \epsilon n/M$. Therefore,

$$\begin{aligned} \sum_{(s,t) \in \Sigma^{[i]}} |V_s| |V_t| &\leq \sum_{(s,t) \in \sigma^{[i]}} |V_s| |V_t| + 2n |V_0| + 2 \sum_{|V_s| < \epsilon n/M} n |V_s| \\ &\leq r\epsilon n^2 + 2r\epsilon n^2 + 2M \frac{\epsilon n}{M} \\ &= (3r+2)\epsilon n^2. \end{aligned} \tag{19}$$

The ϵ-regularity of pairs $(s,t) \notin \Sigma^{[i]}$ can be checked in the exact same way as in theorem 3.9. Therefore, we have that for all pairs $(s,t) \notin \Sigma^{[i]}$, and for any $A \subset V_s$ and $B \subset V_t$, we have

$$\left| \mathbf{1}_B^T \left(T^{[i]} - d^{[i]}(V_s, V_t) J_n \right) \mathbf{1}_A \right| \leq 5\epsilon |V_s| |V_t|.$$

To conclude, if we do all the computations choosing $\frac{\epsilon^3}{3r+2}$ instead of ϵ, last inequality becomes

$$\left| \mathbf{1}_B^T \left(T^{[i]} - d^{[i]}(V_s, V_t) J_n \right) \mathbf{1}_A \right| \leq \frac{5}{3r+2} \epsilon^3 |V_s| |V_t| \leq \epsilon^3 |V_s| |V_t|,$$

and the conditon on irregular pairs (19) becomes

$$\sum_{(s,t) \in \Sigma^{[i]}} |V_s| |V_t| \leq \frac{3r+2}{3r+2} \epsilon^3 n^2 < \epsilon n^2.$$

Therefore, $\mathcal{P}$ is an ϵ–regular partition for $T^{[1]}, \dots, T^{[r]}$ simultaneously. □

3.2 Spectral regularity for graphs

In this section we are going to see how the regularity of the partition of the adjacency matrix T of a graph G provided by the Spectral Regularity Theorem 3.6 can be translated to a regular partition of the graph itself. The following lemma shows that graph regularity is indeed a direct consequence of matrix regularity of its adjacency matrix.

Lemma 3.11. *Consider an ϵ–regular partition $\mathcal{P}$ of the adjacency matrix of a graph G for some $\epsilon > 0$. Then, $\mathcal{P}$ is an ϵ–regular partition of G.*

Proof. Let G be a graph on n vertices. Let $\mathcal{P} = \{V_0, \dots, V_M\}$ be an ϵ–regular partition for its adjacency matrix T and let $\Sigma \subset \{0, \dots, M\}^2$ be the set of exceptional pairs.

Consider a pair $(s,t) \in \{0, \dots, M\}^2 \setminus \Sigma$ and subsets $A \subset V_s$ and $B \subset V_t$. Let $\mathbf{1}_A$ and $\mathbf{1}_B$ be their characteristic vectors respectively. By definition of characteristic vector, we have $\operatorname{supp}(\mathbf{1}_A) = A$, $\operatorname{supp}(\mathbf{1}_B) = B$, $\|\mathbf{1}_A\|_2^2 = |A|$ and $\|\mathbf{1}_B\|_2^2 = |B|$. Thus, by ϵ–regularity of the partition for T,

$$\frac{\left|\mathbf{1}_B^T(T - d(V_s,V_t)J_n)\mathbf{1}_A\right|}{|A|\,|B|} = |d(A,B) - d(V_s,V_t)| \leq \epsilon,$$

which is precisely the definition of ϵ–regularity of the pair (s,t) for the graph G. For irregular pairs, it is clear that

$$\sum_{(s,t)\in\Sigma} |V_s||V_t| \leq \epsilon n^2.$$

Therefore, we can conclude that $\mathcal{P}$ is an ϵ–regular partition for the graph G. $\square$

Let us now revisit Szemerédi's Regularity Lemma. Again, the proof we are going to see is based on the explanation of Cioaba and Martin in [7]. It is a pretty straightforward consequence of Theorem 3.6 and the above lemma.

Theorem 3.12 (Szemerédi's Regularity Lemma)**.** *For every $\epsilon > 0$, there are constants $M = M(\epsilon)$ and $N = N(\epsilon)$ such that every graph on $n \geq N$ vertices has an ϵ–regular partition into at most M parts.*

Proof. Let G be a graph on n vertices and let T be its adjacency matrix. Observe that the entry (a,b) of the matrix T^k counts the number of walks of length k between vertices a and b. Taking $k = 2$, we have

$$\sum_{j=1}^{n} \lambda_j^2 = \operatorname{Tr}(T^2) = 2e(G) \leq 2\binom{n}{2} \leq n^2.$$

The matrix T satisfies the hypothesis of Theorem 3.6. Let $\mathcal{P} = \{V_0, \dots, V_M\}$ be the ϵ–regular partition obtained from the theorem and let $\Sigma \subset \{0, \dots, M\}^2$ be the set of exceptional pairs. By Lemma 3.11, $\mathcal{P}$ is an ϵ–regular partition of the vertices of G. $\square$

This spectral proof also provides an ϵ–regular partition of G as in Theorem 1.4, but there are remarkable differences. One example is the bound on the number of parts: recall that for the partition obtained in 1.4 (the energy proof), the bound was a power tower of ϵ^{-5} twos, and in the spectral case, as we observed in 3.2 it may be a lot larger, namely

$$M \leq \left(\frac{4(F^{(\lfloor 1/\epsilon^3\rfloor - 1)}(1))^4}{\epsilon^4}\right)^{F^{(\lfloor 1/\epsilon^3\rfloor - 1)}(1)}.$$

Therefore, we can conclude that the interest on the spectral version is not in the optimisation of the size of the ϵ–regular partition. In [8], Frieze and Kannan came up with the spectral method (in fact, they obtained a slightly different algorithm that then then derivated to the proof by Tao in [16]) in order to find a more efficient way to compute partitions.

As we have commented in the previous section, another interesting situation to study is the case of a graph with an edge–coloring. We define the ϵ–regularity of a partition for a graph with an edge–coloring to be ϵ–regularity of the partition for each of the monochromatic subgraphs. Therefore, the Edge–coloring Regularity Theorem will provide a partition of the vertices which is ϵ–regular simultaneously on the set of monochromatic subgraphs.

Theorem 3.13 (Edge–coloring Regularity Theorem)**.** *For every $\epsilon > 0$ and $r \in \mathbb{N}$, there are constants $M = M(\epsilon, r)$ and $N = N(\epsilon, r)$ such that for every graph on $n \geq N$ vertices and every r–coloring of the edges there is a partition of the vertices into at most M parts which is ϵ–regular in the monochromatic subgraphs.*

Proof. Let $G = ([n], E^{[1]} \sqcup \cdots \sqcup E^{[r]})$ be a graph on n vertices and an r–edge–coloring. Let T be its adjacency matrix and $\lambda_1, \ldots, \lambda_n \in \mathbb{R}$ its spectrum. For $i = 1, \ldots, r$, define $G^{[i]} = ([n], E^{[i]})$ and let $T^{[i]}$ be its adjacency matrix. Denote the spectrum of $T^{[i]}$ by $\lambda_1^{[i]}, \ldots, \lambda_n^{[i]} \in \mathbb{R}$. Note that

$$\begin{aligned}
\sum_{j=1}^{n} \lambda_j^2 &= \operatorname{Tr}\left(T^2\right) \\
&= \operatorname{Tr}\left((T^{[1]} + \cdots + T^{[r]})^2\right) \\
&= \sum_{i=1}^{r} \operatorname{Tr}\left((T^{[i]})^2\right) + \sum_{\substack{1 \leq i_1, i_2 \leq r \\ i_1 \neq i_2}} \operatorname{Tr}\left(T^{[i_1]} T^{[i_2]}\right) \\
&= \sum_{i=1}^{r} \sum_{j=1}^{n} (\lambda_j^{[i]})^2
\end{aligned}$$

because the diagonal entries of $T^{[i_1]}T^{[i_2]}$ are zero since $E^{[i_1]} \cap E^{[i_2]} = \emptyset$ for all $i_1 \neq i_2$. As a consequence,

$$\sum_{i=1}^{r} \operatorname{Tr}\left((T^{[i]})^2\right) = \operatorname{Tr}(T^2) \leq n^2.$$

Apply Theorem 3.9 to matrices $T^{[1]}, \ldots, T^{[r]}$ and let $\mathcal{P}$ be the ϵ–regular partition obtained. By Lemma 3.11, $\mathcal{P}$ is an ϵ–regular partition for $G^{[1]}, \ldots, G^{[r]}$ simultaneously. □

The natural generalisation of last theorem is the case when G has an r–multicoloring of the edges: edge-sets $E^{[1]}, \ldots, E^{[r]}$ are not disjoint, which means we can assign to each edge up to r colors.

Theorem 3.14 (Edge–multicoloring Regularity Theorem)**.** *For every $\epsilon > 0$ and $r \in \mathbb{N}$, there are constants $M = M(\epsilon, r)$ and $N = N(\epsilon, r)$ such that for every graph on $n \geq N$ vertices and every r–multicoloring of the edges there is a partition of the vertices into at most M parts which is ϵ–regular in the monochromatic subgraphs.*

Proof. Let $G = ([n], E^{[1]} \cup \cdots \cup E^{[r]})$ be a graph on n vertices and an r–edge–multicoloring. For $i = 1, \ldots, r$, define $G^{[i]} = ([n], E^{[i]})$ and let $T^{[i]}$ be its adjacency matrix.

Apply Theorem 3.10 to matrices $T^{[1]}, \ldots, T^{[r]}$ and let $\mathcal{P}$ be the ϵ–regular partition obtained. By Lemma 3.11, $\mathcal{P}$ is an ϵ–regular partition for $G^{[1]}, \ldots, G^{[r]}$ simultaneously. $\square$

It is clear that the case of one single graph and the edge–coloring versions are particular choices of G in last theorem:

(i) Szemerédi's Regularity Lemma 3.12 corresponds to the case $r = 1$, and

(ii) the Edge-coloring Regularity Theorem 3.13 corresponds to the case $\bigcap\limits_{i=1}^{r} E^{[i]} = \emptyset$.

Thanks to the generality of the matrix of the Spectral Regularity Theorem 3.6, we are able to state regularity theorems for other types of graphs as long as the trace of the square of its adjacency matrix is properly bounded. We are going to see first the case of undirected multigraphs. We are including the proof for the sake of completeness although it is exactly the same as the one in Theorem 3.12. The weakness of this theorem compared to the previous ones is precisely that this result does not hold for any multigraph, because not any adjacency matrix of a multigraph satisfies the bound on the trace.

Theorem 3.15 (Multigraph Weak Regularity Theorem)**.** *For every $\epsilon > 0$, there are constants $M = M(\epsilon)$ and $N = N(\epsilon)$ such that every multigraph on $n \geq N$ vertices with adjacency matrix T satisfying $\mathrm{Tr}(T^2) \leq n^2$ has an ϵ–regular partition into at most M parts.*

Proof. Let G be a multigraph on n vertices and let T be its adjacency matrix, which satisfies the hypothesis of Theorem 3.6. Let $\mathcal{P} = \{V_0, \ldots, V_M\}$ be the ϵ–regular partition obtained and let $\Sigma \subset \{0, \ldots, M\}^2$ be the set of exceptional pairs. By Lemma 3.11, $\mathcal{P}$ is an ϵ–regular partition of the vertices of G. $\square$

The last theorem we are going to see is a weaker regularity theorem for directed graphs which will provide an ϵ–regular partition of a bounded number of parts for a graph of any size satisfying an extra condition that will be determined later. The definition of ϵ–regularity for digraphs that we are going to consider is the one from Alon–Shapira [3].

Let $G = (V, E)$ be a directed graph and let $X, Y \subseteq V$.

Definition 3.16. Define

$$\begin{aligned}
\overrightarrow{E}(X,Y) &:= \{(x,y) \in E : x \in X, y \in Y\},\\
\overleftarrow{E}(X,Y) &:= \{(y,x) \in E : y \in Y, x \in X\},\\
\overline{E}(X,Y) &:= \{(x,y) \in E : (y,x) \in E, x \in X, y \in Y\},
\end{aligned}$$

and let $\overrightarrow{e}(X, Y)$, $\overleftarrow{e}(X, Y)$ and $\overline{e}(X, Y)$ be the cardinals of those sets respectively.

The *directed edge–densities* of X and Y are

$$\begin{aligned}\overrightarrow{d}(X, Y) &:= \frac{\overrightarrow{e}(X, Y)}{|X||Y|}, \\ \overleftarrow{d}(X, Y) &:= \frac{\overleftarrow{e}(X, Y)}{|X||Y|}, \\ \overline{d}(X, Y) &:= \frac{\overline{e}(X, Y)}{|X||Y|}.\end{aligned}$$

Definition 3.17. The pair (X, Y) is *ϵ–regular* if for all $A \subset X$ and $B \subset Y$ with $|A| \geq \epsilon|X|$ and $|B| \geq \epsilon|Y|$, we have

$$\begin{aligned}|\overrightarrow{d}(A, B) - \overrightarrow{d}(X, Y)| &\leq \epsilon, \\ |\overleftarrow{d}(A, B) - \overleftarrow{d}(X, Y)| &\leq \epsilon, \\ |\overline{d}(A, B) - \overline{d}(X, Y)| &\leq \epsilon,\end{aligned}$$

the three at the same time.

Definition 3.18. A partition $\mathcal{P} = \{V_1, \ldots, V_k\}$ of V is *ϵ–regular* if

$$\sum_{\substack{(s,t)\in[k]^2 \\ (V_s,V_t) \text{ not } \epsilon\text{–regular}}} |V_s||V_t| \leq \epsilon|V|^2.$$

Alon and Shapira gave a proof using an energy increment argument analogous to the proof of Szemerédi's Regularity Lemma 1.4 in the first chapter (see [3]). The idea is, given any initial partition $\mathcal{Q} = \{W_1, \ldots, W_k\}$, to divide the directed graph into three undirected graphs $\overrightarrow{G}$, $\overleftarrow{G}$ and $\overline{G}$ with edge sets

$$\begin{aligned}\overrightarrow{E}(\mathcal{Q}) &= \bigcup_{s<t} \overrightarrow{E}(W_s, W_t), \\ \overleftarrow{E}(\mathcal{Q}) &= \bigcup_{s<t} \overleftarrow{E}(W_s, W_t), \\ \overline{E}(\mathcal{Q}) &= \bigcup_{s\neq t} \overline{E}(W_s, W_t)\end{aligned}$$

respectively. Then, start refining and updating the undirected graphs by adding the edges between subparts until the partition is ϵ–regular. The spectral approach inspired by this proof is the following.

Start with a random partition $\mathcal{Q}$. Define three undirected graphs $\overrightarrow{G}$, $\overleftarrow{G}$ and $\overline{G}$ as before. Then, we could apply Theorem 3.10 and obtain a partition which is ϵ–regular in the directed sense for the graph $\widetilde{G}$ with edge set $\widetilde{E} = \overrightarrow{E} \cup \overleftarrow{E} \cup \overline{E}$, with $\widetilde{E} \subset E$. The problem is that $\widetilde{E}$ may be a proper subset, which means that there are some edges of G that have not been taken into account, specifically the single inner edges of parts of $\mathcal{Q}$. In this case, the refinement step of Theorem 3.10 may not give a truly ϵ–regular partition for G in the directed sense.

One way to avoid this situation is to start with an initial partition for such $\widetilde{E} = E$, for instance, a partition into independent sets. In that setting, all edges are assigned a direction and consequently the spectral algorithm provides a suitable ϵ–regular partition for G. In fact, as we do not have this problem for double edges, a partition into independent sets of the graph with edge-set $E \setminus \overline{E}(V, V)$ would be enough. This method will give an ϵ–regular partition such that the bound on the number of parts depends on the size of that initial partition, which is precisely the chromatic number of the graph $(V, E \setminus \overline{E}(V, V))$.

In this case, the weakness of the following statement is that the bound not only depends on ϵ but also on the chromatic number of the graph specified above.

Theorem 3.19 (Directed Weak Regularity Theorem)**.** *For every $\epsilon > 0$ and $\chi \in \mathbb{N}$, there is $M = M(\epsilon, \chi)$ and $N = N(\epsilon)$ such that every graph on $n \geq N$ vertices such that the subgraph induced by single edges has chromatic number χ has an ϵ–regular partition into at most M parts.*

Proof. Let $G = ([n], E)$ be a directed graph and let $\mathcal{Q} = \{W_1, \dots, W_\chi\}$ be a χ–vertex–coloring of the simple graph $([n], E \setminus \overline{E}([n], [n]))$. Consider the multicoloring of the edges given by $E = \overrightarrow{E}(\mathcal{Q}) \cup \overleftarrow{E}(\mathcal{Q}) \cup \overline{E}(\mathcal{Q})$. Let $\overrightarrow{T}$, $\overleftarrow{T}$ and $\overline{T}$ be the adjacency matrices of graphs $\overrightarrow{G}$, $\overleftarrow{G}$ and $\overline{G}$ with edge–sets $\overrightarrow{E}(\mathcal{Q})$, $\overleftarrow{E}(\mathcal{Q})$ and $\overline{E}(\mathcal{Q})$ respectively. Apply Theorem 3.14 and let $\mathcal{P} = \{V_0, \dots, V_M\}$ be the refinement of $\mathcal{Q}$ and the partition given by the theorem. The size of $\mathcal{P}$ is bounded by

$$M \leq \chi \left(\frac{4J^4}{\epsilon^4} \right)^{3J}$$

and by choosing an appropriate F, for instance

$$F(x) \geq \frac{\chi^4}{\epsilon^6} \left(\frac{4x^4}{\epsilon^4} \right)^{12J},$$

we can conclude that $\mathcal{P}$ is ϵ–regular for $\overrightarrow{G}$, $\overleftarrow{G}$ and $\overline{G}$ simultaneously, or in other words, $\mathcal{P}$ is an ϵ–regular partition of G in the directed sense. $\square$

3.3 Examples

The first graph we are going to study is K_n, the complete graph on n vertices. Although any partition of the vertices is ϵ–regular, because the edge–density between two subsets of vertices is always 1, we are going to apply the spectral method to see which partition do we obtain. Recall that $\lambda_1 = n - 1$ and $\lambda_j = -1$ for $j > 1$, and $\mathbf{u}_1 = 1/\sqrt{n}\mathbf{1}$, $\mathbf{u}_j = 1/\sqrt{2}(\mathbf{e}_1 - \mathbf{e}_j)$ for $j > 1$

is an orthonormal basis of eigenvectors. Let us apply Szemerédi's Regularity Lemma 3.12 to K_n. Define the intervals $I_1, \ldots, I_{\lfloor 1/\epsilon^3 \rfloor}$ as in Theorem 3.1 and let $K \in \lfloor 1/\epsilon^3 \rfloor$ be the first integer such that

$$\sum_{j \in I_K} \lambda_j^2 \leq \epsilon^3 n^2.$$

We have two possibilities:

- If $K = 1$, then the following should be satisfied,

$$\begin{aligned}\sum_{j \in I_1} \lambda_j^2 &= (n-1)^2 + \sum_{j=2}^{F(1)-1} (-1)^2 \\ &= n^2 - 2n + 1 + F(1) - 1 \\ &= n^2 + F(1) - 2n - 1 \\ &\leq \epsilon^3 n^2,\end{aligned}$$

  which does not hold for example for $\epsilon = 1/2$ and $F(x) = 1/\epsilon^6 (4x/\epsilon^4)^{4x}$.

- If $K > 1$, then

$$\sum_{j \in I_K} \lambda_j^2 = \sum_{j \in I_K} (-1)^2 = |I_K| \leq \epsilon^3 n^2.$$

Assume $K > 1$ and let $J = F^{(K-1)}(1)$. Consider the square of side $2\sqrt{J/(\epsilon n)}$ centered at the origin of the complex plane. The eigenvectors of K_n have values 0, $1/\sqrt{n}$ and $\pm 1/\sqrt{2}$. The entries 0 and $1/\sqrt{n}$ are inside the square of side $2\sqrt{J/(\epsilon n)}$, because $1/\sqrt{n} \leq \sqrt{J/(\epsilon n)}$. We will assume $n > 2J/\epsilon$, which implies that $1/\sqrt{2} > \sqrt{J/(\epsilon n)}$. The partition is defined as follows:

- $\mathcal{P}(\mathbf{u}_1) = \{[n]\}$ and $\Sigma_1 = \emptyset$, because $u_1(a) = 1/\sqrt{n}$ for all $a \in [n]$.
- $\mathcal{P}(\mathbf{u}_j) = \{[n] \setminus \Sigma_j, \Sigma_j\}$ where $\Sigma_j = \{1, j\}$ for $j > 1$.

The refinement of those partitions is $\mathcal{P} = \{V_0, V_1\}$ where

$$V_0 = \bigcup_{j=1} \Sigma_j = [J-1]$$

and $V_1 = [n] \setminus [J-1]$, which is trivially ϵ–regular.

For a general circulant graph C with adjacency matrix $T = circ(a_1, \ldots, a_n)$, we have that the j–th eigenvalue of C

$$\lambda_j = \sum_{k=1}^{n} a_k \omega^{jk},$$

and $\mathbf{u}_1, \ldots, \mathbf{u}_n \in \mathbb{C}^n$ where

$$\mathbf{u}_j = \frac{1}{\sqrt{n}} (\omega^j, \omega^{2j}, \ldots, \omega^{(n-1)j}, 1)$$

is an orthonormal basis of eigenvectors.

Let $K \in \lfloor 1/\epsilon^3 \rfloor$ be the first integer such that

$$\sum_{j\in I_K} \lambda_j^2 = \sum_{j\in I_K}\sum_{k=1}^{n} a_k\omega^{jk} = a_n + a_{n-1}\sum_{j\in I_K}\omega^{(n-1)j} + \cdots + a_2\sum_{j\in I_K}\omega^{2j} + a_1\sum_{j\in I_K}\omega^{j} \leq \epsilon^3 n^2.$$

Again, let $J = J = F^{(K-1)}(1)$ and consider the square of side $2\sqrt{J/(\epsilon n)}$ centered at the origin of the complex plane. The entries of the eigenvectors of G are n–th roots of unity multiplied by a $1/\sqrt{n}$ factor. Therefore, $|u_j(a)| = 1/\sqrt{n}$ for every $a, j \in [n]$. Similarly as before, all of these entries lie inside the square because $1/\sqrt{n} \leq \sqrt{J/(\epsilon n)}$.

The partitions in this case are not as clear as before: for each $j < J$, we have that $(u_j(1), \dots, u_j(n))$ is plotted conforming a regular polygon of n vertices and radius $1/\sqrt{n}$. For each j, we have a different permutation of the vertices of the polygon. This difficults us to give the explicit description of the partitions $\mathcal{P}(\mathbf{u}_j)$ for $j < J$ and in consequence the posterior refinement.

The interest on circulant graphs comes from the fact that they can be equivalently described as Cayley graphs of finite cyclic groups. Recall that a *Cayley graph* is a directed graph $\Gamma = \Gamma(G, S)$ where G is a group and S a generating set of of G. The graph Γ has vertex set indexed by the elements of G and there is an edge (g, h) whenever $gh^{-1} \in S$. If S is symmetric ($S = S^{-1}$) and does not contain the identity element of the group, the Cayley graph Γ can be represented as a simple undirected case, which is the case we were studying. For example, if $G = \mathbb{Z}_n$ and $S = \{g, g^{-1}\}$, then $\Gamma(G, S)$ is the cycle C_n.

This interpretation of circulant graphs leads us to the study of "ϵ–regularity" for groups. Analogously to the process of translating regularity for graphs to matrices that we have done previously in this chapter, we can reformulate the same concept but for groups.

Let G be a finite group and S a symmetric set of generators not containing the identity, and let $X, Y \subseteq G$. Define

$$d(X, Y) = \frac{|\{(g, h) \in X \times Y : gh^{-1} \in S\}|}{|X|\,|Y|}.$$

The pair (X, Y) is ϵ–regular if for all $A \subset X$ and $B \subset Y$ such that $|A| \geq \epsilon\,|X|$ and $|B| \geq \epsilon\,|Y|$, we have

$$|d(A, B) - d(X, Y)| \leq \epsilon.$$

This property can be interpreted as, if we take not too small subsets A and B of X and Y respectively, the elements $h \in B$ and $g \in A$ such that are related like $h = gs$ for some $s \in S$ are well distributed among A and B. If a partition $\mathcal{P} = \{V_0, \dots, V_M\}$ of the group G satisfies that

$$\sum_{\substack{(s,t)\in\{0,\dots,M\}^2 \\ (V_s,V_t)\text{ not }\epsilon\text{–regular}}} |V_s|\,|V_t| \leq \epsilon\,|G|^2,$$

we will say that $\mathcal{P}$ is ϵ–regular.

In terms of finite abelian groups, Green proved an Arithmetic Rregularity Lemma in [10]. It states that, when $G = \mathbb{F}_p^n$ for a fixed primer p, given any subset $S \subseteq G$ there exists a subspace H of bounded codimension such that S is Fourier–uniform with repect to almost all cosets of H. All the discussion by Green is done in terms of Fourier transforms, which are out of our area of expected knowledge. A more accessible interpretation of this result is given in Terry–Wolf [17].

We say that a graph (V, E) is *k–stable* if there are no $a_1, \ldots, a_k, b_1, \ldots, b_k \in V$ such that $a_i b_j \in E$ if and only if $i \le j$. A subset $S \subseteq G$ of a finite abelian group G is *k–stable* if the Cayley graph $\Gamma(G, S)$ if k–stable.

The result from Terry and Wolf is the following: for $\epsilon > 0$, $k \ge 2$ and a prime p, there is a constant $N = N(k, \epsilon, p)$ such that for all $n \ge N$, if $G = \mathbb{F}_p^n$ and $S \subseteq G$ is k–stable, then there is a subspace $H \leqslant G$ of codimension at most $O_k(\epsilon^{-O_k(1)})$ such that for any $g \in G$, either $|(S - g) \cap H| \le \epsilon |H|$ or $|H \setminus (S - g)| \le \epsilon |H|$.

The above result can be interpreted as follows. For a k–stable subset $S \subseteq G$, the density of S for each translate on the subspace H is either close to 1 or 0. Further work can be done in that direction by studying the spectral partition of the Cayley graph $\Gamma(G, S)$ for a k–stable set S, which we will not discuss in this dissertion, and more generally, studying how the spectral partition may be related to Green's Arithmetic Regularity Lemma.

Let us give a final example concerning multigraphs. An interesting example is the case of the multigraph correponding to the k–th power of a graph G^k. If G is a graph on n vertices, G^k has vertex set $[n]$ and has an edge between a and b for each walk from a to b in G.

Recall that if T is the adjacency matrix of G, the entry $(T^k)_{ab}$ is the number of walks from a to b. In other words, T^k is precisely the adjacency matrix of G^k. Therefore, if $\lambda_1, \ldots, \lambda_n \in \mathbb{R}$ are the eigenvalues of T and $\mathbf{u}_1, \ldots, \mathbf{u}_n \in \mathbb{C}^n$ is a basis of eigenvectors, then $\lambda_1^k, \ldots, \lambda_n^k$ are the eigenvalues of T^k and $\mathbf{u}_1, \ldots, \mathbf{u}_n$ is also a basis of eigenvectors.

If $\mathrm{Tr}(T^{2k}) \le n^2$, we can apply Theorem 3.15 and if the same K is valid for both G and G^k, that is

$$\sum_{j \in I_K} \lambda_j^2, \sum_{j \in I_K} \lambda_j^{2k} \le \epsilon^3 n^2,$$

then the same partition is ϵ–regular for vertices of G and G^k. A trivial example of this situation is when

$$\sum_{j \in I_K} \lambda_j^2 < 1 \le \epsilon^3 n^2,$$

because then $|\lambda_j| < 1$ and consequently $\lambda_j^{2k} \le \lambda_j^2$ for all $j \in I_K$.

Chapter 4

Removal Lemmas

The Removal Lemma is one of the key applications of the Regularity Lemma, as we have already seen in Chapter 1. As mentioned in the Introduction, no explicit development of a spectral removal lemma has been addressed in the literature. The purpose of this chapter is to accomplish this goal. To this end we start by giving a matrix version of the Counting and Removal Lemmas for graphs, which correspond to the Matrix Removal Lemma for 0–1 symmetric matrices with zero diagonal. This will motivate our statement of a spectral removal lemma. Once in this more general setting, a natural application is to obtain a removal lemma for multigraphs, which to our knowledge has not been formulated in the literature except for the case of the triangles. In this case, Shapira and Yuster [13] show that the Triangle Removal Lemma does not necessary hold for multigraphs and give some requirements for this removal version to hold. Our version for multigraphs provides general sufficient conditions for the Removal Lemma for multigraphs to hold which involves the maximum multiplicity of an edge.

4.1 The matricial version of the Graph Removal Lemma

In order to prove the matricial analogous of the Counting and Removal Lemmas in Chapter 1, we have to see first how the concepts of the statements for graphs are translated in terms of their adjacency matrices.

Definition 4.1. Let $A = (a_{ij})$ be a $n \times n$ symmetric 0–1 matrix and let $A_1, A_2 \subseteq [n]$. Define

$$w_A(A_1 \times A_2) := \left|\{a_{ij} > 0 : i \in A_1, j \in A_2\}\right|.$$

The *weight* of A is $w(A) := w_A([n] \times [n])$.

Let G be a graph on n vertices and let A be its adjacency matrix. The number of edges between two subsets of vertices $A_1, A_2 \subseteq [n]$ of G is $w_A(A_1 \times A_2)$ and the total number of edges of G is $w(A)/2$. Therefore, the edge-density between A_1 and A_2 corresponds to

$$d(A_1, A_2) = \frac{w_A(A_1 \times A_2)}{|A_1|\,|A_2|},$$

which is a real number between 0 and 1 (as expected).

For a k–tuple $Q = (v_1, \ldots, v_k) \in [n]^k$ with no repeated entries denote by A^Q the $k \times k$ matrix obtained obtained from A by choosing the rows and columns with indexes in Q ordered as in the k–tuple.

A graph H on k vertices with adjacency matrix B is a subgraph of G if and only if there is a k–tuple Q such that $B \leq A^Q$, where here the inequality is meant to be componentwise. We then say that G contains a copy of H, or equivalently, that A *contains a copy* of B. If there is no k–tuple with this property, we say that A is B–free.

Theorem 4.2 (Graph Counting Lemma, matricial version)**.** *Let $B = (b_{ij})$ be a $k \times k$ symmetric 0–1 matrix with zero diagonal and let $\epsilon > 0$. Let $A = (a_{ij})$ be a $n \times n$ symmetric 0–1 matrix with zero diagonal and let $V_1, \dots, V_k$ be subsets of $[n]$ such that (V_i, V_j) is ϵ–regular whenever $b_{ij} = 1$. Then, the number of copies of B in A is within $\epsilon \cdot \frac{w(B)}{2} \cdot |V_1| \cdots |V_k|$ of*

$$\left(\prod_{\substack{b_{ij}=1 \\ i<j}} d(V_i, V_j) \right) \left(\prod_{i=1}^{k} |V_i| \right).$$

Proof. Let us prove a probabilistic version of the theorem: if we have indexes $v_1 \in V_1, \dots, v_k \in V_k$ chosen uniformly and independendently at random, then the conclusion of the theorem is equivalent to

$$\left| \Pr\left(A^Q \geq B\right) - \prod_{\substack{b_{ij}=1 \\ i<j}} d(V_i, V_j) \right| \leq \epsilon \cdot \frac{w(B)}{2}. \tag{1}$$

That is, the probability that the tuple $(v_1, \dots, v_k)$ is indeed a copy of B in A is within $\epsilon \cdot w(B)/2$ of

$$\prod_{\substack{b_{ij}=1 \\ i<j}} d(V_i, V_j).$$

Suppose $w(B) > 0$. To simplify notation, define

$$\mathbb{P}_B := \Pr\left(A^Q \geq B\right).$$

Relabelling if necessary, assume $b_{12} = 1$. If $B' = (b'_{ij})$ denotes the $k \times k$ symmetric matrix with $b'_{12} = b'_{21} = 0$ and $b'_{ij} = b_{ij}$ for the rest of entries, we will show

$$|\mathbb{P}_B - d(V_1, V_2)\mathbb{P}_{B'}| \leq \epsilon.$$

Let us fix $v_3, \dots, v_k$ and choose randomly only v_1, v_2. Observe that if last inequality is satisfied under this extra constraint, it also holds for the weaker case where all k indexes are chosen randomly. Let

$$A_1 := \{v_1 \in V_1 : a_{v_1 v_i} = 1 \text{ whenever } b_{1i} = 1 \text{ for all } i \neq 2\},$$

$$A_2 := \{v_2 \in V_2 : a_{v_2 v_i} = 1 \text{ whenever } b_{2,i} = 1 \text{ for all } i \neq 1\},$$

be the possible choices for v_1 and v_2 which, with $v_3, \ldots, v_k$, give a copy of B' in A. The probabilistic statement with $v_3, \ldots, v_k$ fixed can be reformulated as

$$\left| \frac{w_A(A_1 \times A_2)}{|V_1|\,|V_2|} - d(V_1, V_2)\frac{|A_1|\,|A_2|}{|V_1|\,|V_2|} \right| \leq \epsilon.$$

Let us check it by cases. If $|A_1| \leq \epsilon\,|V_1|$ or $|A_2| \leq \epsilon\,|V_2|$, both terms are at most ϵ, so the inequality holds. Otherwise, if $|A_1| > \epsilon\,|V_1|$ and $|A_2| > \epsilon\,|V_2|$, by ϵ–regularity of (V_1, V_2) we have

$$\left| \frac{w(A_1 \times A_2)}{|V_1|\,|V_2|} - d(V_1, V_2)\frac{|A_1|\,|A_2|}{|V_1|\,|V_2|} \right| = |d(A_1, A_2) - d(V_1, V_2)|\,\frac{|A_1|\,|A_2|}{|V_1|\,|V_2|} \leq \epsilon.$$

We are going to complete the proof by induction on $w(B)/2$. Since $w(B)/2 = 1 + w(B')/2$, assume (1) holds when replacing B with B'. Therefore,

$$\begin{aligned}
\left| \mathbb{P}_B - \prod_{\substack{i<j \\ b_{ij}=1}} d(V_i, V_j) \right| &\leq |\mathbb{P}_B - d(V_1, V_2)\mathbb{P}_{B'}| + d(V_1, V_2)\left| \mathbb{P}_{B'} - \prod_{\substack{i<j \\ b'_{ij}=1}} d(V_i, V_j) \right| \\
&\leq \epsilon + d(V_1, V_2)\epsilon \cdot \frac{w(B')}{2} \\
&\leq \epsilon \left(1 + \frac{w(B')}{2} \right) \\
&= \epsilon \cdot \frac{w(B)}{2}.
\end{aligned}$$

□

Equipped with the above counting lemma we can formulate the following matricial version of the Removal Lemma. As before, we say that an $n \times n$ matrix contains a copy of a $k \times k$ matrix B if there is a subset $Q \subset [n]$ with $|Q| = k$ such that $B \leq A^Q$.

Theorem 4.3 (Graph Removal Lemma, matricial version)**.** *For each $k \times k$ symmetric 0–1 matrix $B = (b_{ij})$ with zero diagonal and each $\epsilon > 0$ there exists $\delta > 0$ such that every $n \times n$ symmetric 0–1 matrix $A = (a_{ij})$ with zero diagonal with fewer than $\delta n^{w(B)/2}$ copies of B can be made B–free by changing at most ϵn^2 pairs of entries (a_{ij}, a_{ji}) to zero.*

Proof. First, apply the Spectral Regularity Theorem 3.6 to matrix A in order to obtain an $\epsilon/4$–regular partition $\mathcal{P} = \{V_0, \ldots, V_M\}$ of $[n]$. Secondly, let A' be a 0–1 symmetric matrix that has zero entries on the blocks $V_i \times V_j$ and $V_j \times V_i$ whenever

(a) (V_i, V_j) is not $\epsilon/4$–regular,

(b) $d(V_i, V_j) < \epsilon/2$,

(c) V_i or V_j have size smaller than $\frac{\epsilon n}{4M}$.

For the rest of the entries, set $a'_{ij} = a_{ij}$.

The number of pairs (a_{ij}, a_{ji}) set to zero in (a) is, by $\epsilon/4$–regularity of the partition,

$$\sum_{(V_i,V_j) \text{ not } \epsilon/4\text{–regular}} w_A(V_i \times V_j) \leq \sum_{(V_i,V_j) \text{ not } \epsilon/4\text{–regular}} |V_i|\,|V_j| \leq \frac{\epsilon}{4} n^2.$$

For the low denisty pairs in (b), we have

$$\sum_{d(V_i,V_j)<\epsilon/2} d(V_i, V_j)\,|V_i|\,|V_j| \leq \frac{\epsilon}{2} \sum_{(i,j)\in\{0,\ldots,M\}^2} |V_i|\,|V_j| = \frac{\epsilon}{2} n^2.$$

And finally, for the small subsets in (c), the number of pairs of entries set to zero is at most

$$n \cdot \frac{\epsilon n}{4M} \cdot M = \frac{\epsilon}{4} n^2,$$

because for each index in $[n]$, there are at most $\frac{\epsilon n}{4M}$ entries for each of the at most M small subsets of the partition.

Suppose that there is a tuple $(v_1, \ldots, v_k) \in V_{s_1} \times \cdots \times V_{s_k}$ where $s_1, \ldots, s_k \in \{0, \ldots, M\}$ such that $a'_{v_i v_j} = 1$ whenever $b_{ij} = 1$. In other words, there is a copy of B in A' realised by $(v_1, \ldots, v_k)$. By the Counting Lemma 4.2, there are at least

$$\begin{aligned} &\frac{1}{|\mathrm{Aut}(B)|}\left(\left(\prod_{\substack{b_{ij}=1\\ i<j}} d(V_{s_i}, V_{s_j})\right)\left(\prod_{i=1}^{k} |V_{s_i}|\right) - \epsilon \cdot \frac{w(B)}{2} \prod_{i=1}^{k} |V_{s_i}|\right) \\ &\geq \frac{1}{|\mathrm{Aut}(B)|}\left(\left(\frac{\epsilon}{2}\right)^{\frac{w(B)}{2}} \left(\frac{\epsilon n}{4M}\right)^{k} - \epsilon \cdot \frac{w(B)}{2} \left(\frac{\epsilon n}{4M}\right)^{k}\right) \\ &= \frac{1}{|\mathrm{Aut}(B)|}\left(\frac{\epsilon}{4M}\right)^{k}\left(\left(\frac{\epsilon}{2}\right)^{\frac{w(B)}{2}} - \epsilon \cdot \frac{w(B)}{2}\right) n^k \end{aligned}$$

copies of B in A', where $\mathrm{Aut}(B)$ is the set of permutation matrices P such that $B = PBP^T$. If we choose

$$\delta < \frac{1}{|\mathrm{Aut}(B)|}\left(\frac{\epsilon}{4M}\right)^{k}\left(\left(\frac{\epsilon}{2}\right)^{\frac{w(B)}{2}} - \epsilon \cdot \frac{w(B)}{2}\right),$$

we get a contradiction: if in the matrix A' there is some copy of B left, we found that there are more than δn^k copies. However, the original matrix A had at most δn^k copies. In conclusion, the matrix obtained from A by changing at most ϵn^2 pairs (a_{ij}, a_{ji}) of entries to zero is B–free. $\square$

4.2 The Multigraph Removal Lemma

Let G be a multigraph on n vertices and let H be another multigraph on k vertices. Let $G(x, y)$ be a simple graph on n vertices with an edge between vertices i and j whenever

$$m_G(i, j) \geq m_H(x, y),$$

where $m_G(i, j)$ denote the number of edges in G between the pair of vertices $\{i, j\}$. We will also refer to this number as the *multiplicity* of the edge $\{i, j\}$ in G. Note that if r is the maximum edge multiplicity of H, then $\{G(x, y)\}_{(x,y)\in[k]^2}$ consists in of most r differents matrices, one for each possible multiplicity.

In matricial terms, if $A = (a_{ij})$ and $B = (b_{ij})$ are the adjacency matrices of G and H, then $A(x, y) := (a(x, y)_{ij})$ is the adjacency matrix of the graph $G(x, y)$, which is defined by

$$a(x, y)_{ij} := \begin{cases} 1 \text{ if } a_{ij} \geq b_{xy} > 0, \\ 0 \text{ otherwise,} \end{cases}$$

Let $d_{xy} = d_{A(x,y)}$ be the edge–density computed on matrix $A(x, y)$.

Let us first prove the Multigraph Counting Lemma. Note that the ϵ–regularity conditions are imposed in matrices $A(i, j)$ instead of A as we could have imagined. We make use of this strategy because in a key step of the proof we need the density to be a real number between 0 and 1, and it may not be the case for any multigraph since the multiplicity of the edges can be arbitrarily large. We observe that this approach in the following statement is different from the naive one, consisting of decomposing the matrices A and B as a sum of 0–1 matrices and apply the edge–multicolored version of the Regularity Lemma, which may give a different and weaker statement.

Analogous to the above case, a multigraph H is a subgraph of G if and only if there is a k–tuple $Q = (v_1, \ldots, v_k)$ such that $B \leq A^Q$ in the sense that $a_{v_i v_j} \geq b_{ij}$ whenever $b_{ij} > 0$.

Theorem 4.4 (Multigraph Counting Lemma, matricial version)**.** *Let $B = (b_{ij})$ be a $k \times k$ symmetric matrix with $b_{ij} \in \mathbb{N}$ and zero diagonal and let $\epsilon > 0$. Let $A = (a_{ij})$ be a $n \times n$ matrix with $a_{ij} \in$ and zero diagonal and consider the matrices $A(i, j)$ determined by A and B. Let $V_1, \ldots, V_k \subseteq [n]$ be such that (V_i, V_j) is ϵ–regular for $A(i, j)$ whenever $b_{ij} > 0$. Then, the number copies of B in A is within $\epsilon \cdot \frac{w(B)}{2} \cdot |V_1| \cdots |V_k|$ of*

$$\left(\prod_{\substack{b_{ij}>0 \\ i<j}} d_{ij}(V_i, V_j) \right) \left(\prod_{i=1}^{k} |V_i| \right).$$

Proof. Let us prove a probabilistic version of the theorem: if we have indexes $v_1 \in V_1, \ldots, v_k \in$

V_k chosen uniformly and independendently at random, then

$$\left|\Pr\left(a_{v_iv_j} \geq b_{ij} \quad \forall b_{ij} > 0\right) - \prod_{\substack{b_{ij}>0 \\ i<j}} d_{ij}(V_i, V_j)\right| \leq \epsilon \cdot \frac{w(B)}{2}. \tag{2}$$

That is, the probability that the tuple $(v_1, \dots, v_k)$ is indeed a copy of B in A is within $\epsilon \cdot w(B)/2$ of

$$\prod_{\substack{b_{ij}>0 \\ i<j}} d_{ij}(V_i, V_j).$$

To simplify notation, define

$$\mathbb{P}_B := \Pr\left(a_{v_iv_j} \geq b_{ij} \quad \forall b_{ij} > 0\right).$$

Relabelling if necessary, assume $b_{12} > 0$. If $B' = (b'_{ij})$ denotes the $k \times k$ symmetric matrix with $b'_{12} = b'_{21} = 0$ and $b'_{ij} = b_{ij}$ for the rest of entries, we will show

$$|\mathbb{P}_B - d_{12}(V_1, V_2)\mathbb{P}_{B'}| \leq \epsilon.$$

Let us fix $v_3, \dots, v_k$ and choose randomly only v_1, v_2. Observe that if last inequality is satisfied under this extra constraint, it also holds for the weaker case where all k vertices are chosen randomly. Let

$$A_1 := \{v_1 \in V_1 : a_{v_1v_i} \geq b_{1i} \text{ whenever } b_{1i} > 0 \text{ for all } i \neq 2\},$$

$$A_2 := \{v_2 \in V_2 : a_{v_2v_i} \geq b_{2i} \text{ whenever } b_{2i} > 0 \text{ for all } i \neq 1\},$$

be the possible choices for v_1 and v_2 which, with $v_3, \dots, v_k$, give a copy of B' in A. The probabilistic statement with $v_3, \dots, v_k$ fixed can be reformulated as

$$\left|\frac{w_{A(1,2)}(A_1 \times A_2)}{|V_1|\,|V_2|} - d_{12}(V_1, V_2)\frac{|A_1|\,|A_2|}{|V_1|\,|V_2|}\right| \leq \epsilon.$$

Let us check it by cases. If $|A_1| \leq \epsilon\,|V_1|$ or $|A_2| \leq \epsilon\,|V_2|$, both terms are at most ϵ, so we have

$$\left|\frac{w_{A(1,2)}(A_1 \times A_2)}{|V_1|\,|V_2|} - d_{12}(V_1, V_2)\frac{|A_1|\,|A_2|}{|V_1|\,|V_2|}\right| \leq \epsilon.$$

Otherwise, if $|A_1| > \epsilon\,|V_1|$ and $|A_2| > \epsilon\,|V_2|$, by ϵ–regularity of (V_1, V_2) in $A(1,2)$ we have

$$\left|\frac{w_{A(1,2)}(A_1 \times A_2)}{|V_1|\,|V_2|} - d_{12}(V_1, V_2)\frac{|A_1|\,|A_2|}{|V_1|\,|V_2|}\right| = |d_{12}(A_1, A_2) - d_{12}(V_1, V_2)|\,\frac{|A_1|\,|A_2|}{|V_1|\,|V_2|} \leq \epsilon.$$

We are going to complete the proof by induction on the number of entries $b_{ij} > 0$. Since $w(B)/2 = 1 + w(B')/2$, assume (2) holds when replacing B with B'. Therefore,

$$\begin{aligned}\left|\mathbb{P}_B - \prod_{\substack{b_{ij}>0\\ i<j}} d_{ij}(V_i,V_j)\right| &\leq |\mathbb{P}_B - d_{12}(V_1,V_2)\mathbb{P}_{B'}| + d_{12}(V_1,V_2)\left|\mathbb{P}_{B'} - \prod_{\substack{b'_{ij}>0\\ i<j}} d_{ij}(V_i,V_j)\right| \\ &\leq \epsilon + d_{12}(V_1,V_2)\cdot\epsilon\cdot\frac{w(B')}{2} \\ &\leq \epsilon\left(1+\frac{w(B')}{2}\right) \\ &= \epsilon\cdot\frac{w(B)}{2}.\end{aligned}$$

□

Let us now use the above theorem to prove a matricial version of the Multigraph Removal Lemma. Since the Counting Lemma requires some kind of regularity for matrices $A(i,j)$ and not for A, we are not going to use the Multigraph Regularity Lemma 3.15. In spite of that, if $r = \max m_H(i,j)$ is the maximum edge–multiplicity of H, we are going to consider the r distinct matrices $A(i,j)$ and Theorem 3.10 will provide a suitable ϵ–regular partition.

Theorem 4.5 (Multigraph Removal Lemma, matricial version)**.** *For each $k\times k$ symmetric matrix $B=(b_{ij})$ with $b_{ij}\in\mathbb{N}$ and $r=\max b_{ij}$ and each $\epsilon>0$ there exists $\delta>0$ such that every $n\times n$ symmetric matrix $A=(a_{ij})$ with $a_{ij}\in\mathbb{N}$ with fewer than $\delta n^{w(B)/2}$ copies of B can be made B–free by changing $c(r)\epsilon n^2$ pairs of entries (a_{ij},a_{ji}) to zero for some constant $c(r)\in\mathbb{Q}$.*

Proof. If $r=\max b_{ij}$, then there are r distinct $A(i,j)$ matrices, each one corresponding to some $b_{ij}=k$ for $k=1,\dots,r$. Let $\{A_1,\dots,A_r\}$ be such matrices. First, apply Theorem 3.10 to matrices $\{A_1,\dots,A_r\}$ in order to obtain an $\epsilon/4$–regular partition $\mathcal{P}=\{V_0,\dots,V_M\}$ of $[n]$. Secondly, let A' be a symmetric matrix that has zero entries on the blocks $V_i\times V_j$ and $V_j\times V_i$ whenever

(a) (V_i,V_j) is not $\epsilon/4$–regular in some A_s,

(b) $d(V_i,V_j)<\epsilon/2$ in some A_s,

(c) V_i or V_j have size smaller than $\frac{\epsilon n}{4M}$.

For the rest of the entries, set $a'_{ij}=a_{ij}$.

The number of pairs (a_{ij},a_{ji}) set to zero in (a) is, by $\epsilon/4$–regularity of the partition,

$$\sum_{s=1}^{r}\sum_{\substack{(V_i,V_j)\text{ not}\\ \epsilon/4\text{–regular for }A_s}} w_{A_s}(V_i\times V_j)\leq\sum_{s=1}^{r}\sum_{\substack{(V_i,V_j)\text{ not}\\ \epsilon/4\text{–regular for }A_s}}|V_i|\,|V_j|\leq\frac{\epsilon}{4}rn^2.$$

If $d_s = d_{A_s}$ is the density in A_s, for the low density pairs in (b) we have

$$\sum_{s=1}^{r} \sum_{d_s(V_i,V_j)<\epsilon/2} d_s(V_i, V_j)\,|V_i|\,|V_j| \leq \frac{\epsilon}{2} r \sum_{(i,j)\in\{0,\dots,M\}^2} |V_i|\,|V_j| = \frac{\epsilon}{2} r n^2.$$

And finally, for the small subsets in (c), the number of pairs of entries set to zero is at most

$$n \cdot \frac{\epsilon n}{4M} \cdot M = \frac{\epsilon}{4} n^2,$$

because for each index in $[n]$, there are at most $\frac{\epsilon n}{4M}$ entries for each of the at most M small subsets of the partition.

Suppose that there is a tuple $(v_1, \dots, v_k) \in V_{s_1} \times \cdots \times V_{s_k}$ where $s_1, \dots, s_k \in \{0, \dots, M\}$ such that $a'_{v_i v_j} \geq b_{ij}$ whenever $b_{ij} > 0$. In other words, there is a copy of B in A' realised by $(v_1, \dots, v_k)$. By the Counting Lemma 4.4, there are at least

$$\begin{aligned}
&\frac{1}{|\mathrm{Aut}(B)|}\left(\left(\prod_{\substack{b_{ij}>0\\ i<j}} d_{ij}(V_{s_i}, V_{s_j})\right)\left(\prod_{i=1}^{k} |V_{s_i}|\right) - \epsilon \cdot \frac{w(B)}{2} \prod_{i=1}^{k} |V_{s_i}|\right) \\
&\quad \geq \frac{1}{|\mathrm{Aut}(B)|}\left(\left(\frac{\epsilon}{2}\right)^{\frac{w(B)}{2}} \left(\frac{\epsilon n}{4M}\right)^{k} - \epsilon \cdot \frac{w(B)}{2}\left(\frac{\epsilon n}{4M}\right)^{k}\right) \\
&\quad = \frac{1}{|\mathrm{Aut}(B)|}\left(\frac{\epsilon}{4M}\right)^{k}\left(\left(\frac{\epsilon}{2}\right)^{\frac{w(B)}{2}} - \epsilon \cdot \frac{w(B)}{2}\right) n^k
\end{aligned}$$

copies of B in A', where $\mathrm{Aut}(B)$ is the set of permutation matrices P such that $B = PBP^T$. If we choose

$$\delta < \frac{1}{|\mathrm{Aut}(B)|}\left(\frac{\epsilon}{4M}\right)^{k}\left(\left(\frac{\epsilon}{2}\right)^{\frac{w(B)}{2}} - \epsilon \cdot \frac{w(B)}{2}\right),$$

we get a contradiction: if in the matrix A' there is some copy of B left, we found that there are more than δn^k copies. However, the original matrix A had at most δn^k copies. In conclusion, the matrix obtained from A by changing at most $\frac{3r+1}{4}\epsilon n^2$ pairs (a_{ij}, a_{ji}) of entries to zero is B–free. □

We can omit the zero diagonal condition since in the definition of the k–tuple we already assume that entries are not repeated. We observe that in the context of the matrix versions the condition that the coefficients are in $\mathbb{N}$ can be relaxed to entries $\mathbb{Q}$, which may provide a more meaningful statement.

In the language of multigraphs the above two theorems can be rewritten as follows.

Theorem 4.6 (Multigraph Counting Lemma)**.** *Let H be a multigraph on k vertices and let $\epsilon > 0$. Let G be a multigraph on n vertices and consider the graphs G_{ij} determined by G and*

H. Let $V_1, \ldots, V_k \subseteq V(G)$ be such that (V_i, V_j) is ϵ–regular for G_{ij} whenever $\{i,j\} \in E(H)$. Then, the number of tuples $(v_1, \ldots, v_k) \in V_1 \times \cdots \times V_k$ such that $m_G(v_i, v_j) \geq m_H(i,j)$ whenever $\{i,j\} \in E(H)$ is within $\epsilon \cdot e(H) \cdot |V_1| \cdots |V_k|$ of

$$\left(\prod_{\{i,j\} \in E(H)} d_{ij}(V_i, V_j)\right)\left(\prod_{i=1}^{k} |V_i|\right).$$

Theorem 4.7 (Multigraph Removal Lemma)**.** *For each multigraph H with maximum edge-multiplicity r and each $\epsilon > 0$ there exists $\delta > 0$ such that every multigraph on n vertices with maximum edge–multiplicity R and fewer than $\delta n^{|V(H)|}$ copies of H can be made H–free by removing no more than $c(r)R\epsilon n^2$ edges for some constant $c(r) \in \mathbb{Q}$.*

Note that the number of edges removed depends on the maximum edge–multiplicities of both multigraphs. In fact, our results are much more innacurate with respect to the ones in Shapira–Yuster [13] in the case of triangles, because in our approach we allow the removal of a number of edges proportional to R which can be arbitrarily large. Let us remark some cases:

- When $r = R = 1$, we recover the Graph Counting Lemma 1.11.
- When $r = 1$, if $(v_1, \ldots, v_k)$ gives a copy of H in G, for some $\{i,j\} \in E(H)$ such that $m(v_i, v_j) \geq 1$, we have to remove all the edges between v_i and v_j to avoid that copy of H. In fact, $R > 1$ indicates that the dependency on r and R in the statement of Theorem 4.7 can not be ommitted.
- When $r > 1$ and $m > 1$ is the minimum multiplicity of edges of H then it is not necessary to set the corresponding entry of the adjacency matrix of G to zero to remove an instance of H, it suffices to set the value to $m - 1$. In that case, the number of removed edges can be improved to $(R - m + 1)c(r)\epsilon n^2$.

Conclusions and further work

The main purpose of this thesis is to study the spectral proof of Szemerédi's Regularity Lemma given by Tao [16], inspired by the article of Frieze and Kannan [8].

The most remarkable accomplishments in the present dissertation are the following:

- A complete and exhaustive review of the spectral proof of Szemerédi's Regularity Lemma by Tao, with an interpretation of the decomposition of the adjacency matrix.
- The spectral partition enjoys some noteworthy properties that the partition obtained from the classical proof may not satisfy. We remark the fact that since it is exclusively determined by the spectrum and a part of the eigenvectors, the same partition is valid for all graphs satisfying those spectral conditions.
- A generalisation of Szemerédi's Regularity Theorem to self-adjoint matrices fulfilling a condition on the trace, namely an upper bound on the trace of the second power of the matrix, following Cioaba–Martin [7]. We yielded a notion of ϵ–regularity for matrices which contain the familiar regularity for graphs for particular case of symmetric 0–1 matrices. We refer to this statement as the Spectral Regularity Theorem.
- The matricial version can be extended to a regularity theorem for a finite set of matrices. We studied two different situations: a first one where each of the matrices satisfy the condition on the trace required by the Spectral Regularity Theorem and another one where the trace bound must be realised by the sum of squares of those matrices.
- A spectral proof for a regularity theorem concerning graphs with an edge–(multi)coloring following Robertson [12]. In addition, we contribute with a reinterpretation of those matricial results which derive to weaker regularity statements for multigraphs and directed graphs.
- A practical use of the regularity theorem for a set of graphs in the proof of the Multigraph Removal Lemma, analogous to the classic example of the Graph Removal Lemma and Szemerédi's Regularity Lemma, together with the reformulation of the statements in terms of adjacency matrices.
- An explicit version of the Spectral Counting Lemma and the Spectral Removal Lemma for matrices which has not been developed int he literature.

Some interesting further work suggested by this thesis is the following:

- An important topic that we did not study is the algorithmic efficiency of the spectral method. We have briefly commented the non–optimality of the bound on the number of parts of the spectral partition. Nevertheless, the adjacency matrix decomposition actually originated in an algorithmic context in Frieze–Kannan [8], so a complete study on the spectral proof by Tao may be significant.

- Alhough there are some families of graphs for which we have expressions for the spectrum and a basis of eigenvectors, see some examples in Chapter 2, we have encountered some difficulties to give explicit ϵ–regular partitions for those graphs. In our case, a visible obstacle for the practical use of the spectral method is the exponential dependence on ϵ^{-1} of the bound for the number of parts and obviously the minimum size of the considered graphs. For example, in the case of Kneser graphs, although we know that the entries of the eigenvectors are the solution to linear systems with real coefficients, the magnitude of the graph makes it notably labourious. In conclusion, the ability to do some computations for large dense graphs may be useful to illustrate the particularities of the spectral partition with respect to the obtained in the classical proof of Szemerédi.

- Green proved an arithmetic reformulation of Szemerédi's Regularity Lemma which follow from Fourier analysis arguments [10]. An interesting study may result from connecting the spectral approach discussed with the arithmetic interpretation.

- The Removal Lemma has been extended to the general hypergraph setting with significant applications to the multidimensional Szemerédi Theorem. Tao mentions in his blog that it seems unlikely that spectral techniques could be used to handle this extension. Nevertheless, Frieze and Kannan do address this issue and consider regular decompositions of high dimensional matrices or tensors. It is likely that such decmmopositions may be used to obtain the removal lemma for hypergraphs, bypassing the use of spectral methods.

References

[1] M. Aigner, G. M. Ziegler. *Proofs from THE BOOK*, Springer, Berlin & Heidelberg, 1998, 169–172.

[2] N. Alon, O. Ben–Eliezer. *Removal Lemmas for Matrices*, arXiv:1609.04235, 2016.

[3] N. Alon, A.Shapira. *Testing subgraphs in directed graphs*, Proc. of the 35 ACM STOC, ACM Press, 2003, 700–709.

[4] A. E. Brouwer, W. H. Haemers. *Spectra of Graphs*, Springer, New York, NY, 2010.

[5] F. R. K. Chung. *Diameters and eigenvalues*, Journal of the American Mathematical Society, 1989, 187–196.

[6] F. R. K. Chung, R. L. Graham, R. M. Wilson. *Quasi–Random Graphs*, Combinatorica 9, 1989, 345–362.

[7] S. Cioaba, R. R. Martin. *Tao's spectral proof of the Szemerédi regularity lemma*, lecture notes, 2013.

[8] A. Frieze, R. Kannan. *Quick approximation to matrices and applications*, Combinatorica 19, 1999, 175–220.

[9] T. Gowers. *Lower bounds of tower type for Szemerédi's uniformity lemma*, Geometric And Functional Analysis 7, 1997, 322–337.

[10] B. Green. *A Szemerédi–type regularity lemma in abelian groups, with applications*, Geometric And Functional Analysis 15, 2005, 340–376.

[11] J. Komlós, A. Shokoufandeh, M. Simonovits, E. Szemerédi. *The Regularity Lemma and Its Applications in Graph Theory*, Lecture Notes in Computer Science vol. 2292, 2000, 84–112.

[12] S. Robertson. *Spectral proof of Szemerédi regularity with edge–coloring*, seminar notes, 2019.

[13] A. Shapira, R. Yuster. *Multigraphs (only) satisfy a weak triangle removal lemma*, The Electronical Journal of Combinatorics vol. 16, 2009.

[14] B. Szegedi. *Limits of kernel operators and the spectral regularity lemma*, European Journal of Combinatorics vol. 32, Issue 7, 2011, 1156–1167.

[15] E. Szemerédi. *Regular partitions of graphs*, Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, 1976), 1978, 399–401.

[16] T. Tao. *The spectral proof of the Szemerédi regularity lemma*, blog entry, 2012.

[17] C. Terry, J. Wolf. *Stable arithmetic regularity in the finite–field model*, Bulletin of the London Mathematical Society vol. 51, 2017.

[18] A. Thomason. *Pseudo–random graphs*, Annals of Discrete Math. 33, 307–331, 1987.

[19] Y. Zhao. *Graph Theory and Additive Combinatorics*, lecture notes, Massachusetts Institute of Technology, 2019, 49–76.