

Ca' Foscari University of Venice

DEPARTMENT OF ENVIRONMENTAL
SCIENCES, INFORMATICS AND STATISTICS



Master's Degree programme in
Computer Science, Data Management and Analytics
Second Cycle (D.M. 270/2004)

M.Sc. Thesis

Graph Compression Using The Regularity Method

SUPERVISOR:
Prof. Marcello Pelillo

ASSISTANT SUPERVISOR:
Dott. Marco Fiorucci

GRADUAND:
Francesco Pelosin
Matr. 839220

Academic Year
2016 - 2017

Abstract

We are living in a world which is getting more and more interconnected and, as physiological effect, the interaction between the entities produces more and more information. This high throughput generation calls for techniques able to reduce the volume of the data, but still able to preserve the carried knowledge. Data compression and summarization techniques are one of the possible approaches to face such problems. The aim of this thesis is to devise a new pipeline for compressing and decompressing a graph by exploiting Szemerédi's Regularity Lemma. In particular, it has been developed a procedure called CoDec (Compression-Decompression) which is based on Alon et al's constructive version of the Regularity Lemma. We provide an extensive experimental evaluation to measure how robust is the framework as we both corrupt the structures carried by the graph and add noisy edges among them. The experimental results make us confident that our method can be effectively used as a graph compression technique able to preserve meaningful patterns of the original graph.

Acknowledgements

I would like to express my gratitude to my supervisor Marcello Pelillo for accepting me as a candidate and assigning me such stimulating research topic for my thesis.

I also have to express my gratitude to Marco Fiorucci as he played a crucial role between me and my supervisor. He guided me during the creation of this thesis and he was always kind enough to discuss all the problems faced and to improve all the ideas that I had.

Since this thesis ends an important period of my life, I have to thank both my old and new friends that I've met, for all the happy moments and beautiful discussions that we had.

I have to thank my family for the continuous moral and financial support during these two years. In particular, my brother and his family to always have kind words for me in both happy and difficult times. Last but not least, I have to thank my parents who always constituted a source of love for me and that allowed me to pursue this degree.

Contents

1	Introduction	3
1.1	Contributions	4
1.2	Structure of the work	4
2	Context and Background	6
2.1	Extremal Graph Theory and Ramsey Theory	6
2.2	Szemerédi’s Regularity Lemma	7
2.3	Alon et al.’s Algorithm	10
3	Description of the devised technique	13
3.1	Overview of the CoDec Procedure	13
3.2	A New Refinement	15
3.3	Graph Decompression	17
3.4	Post-Decompression filtering	19
4	Experimental Results	20
4.1	Evaluation measures	20
4.2	Synthetic Graphs Experiments	22
4.3	Real Network Experiments	26
4.4	Analysis and considerations	29
5	Conclusions	34
5.1	Future Works	34

Chapter 1

Introduction

We are living in a world which is getting more and more interconnected and, as physiological effect, the interaction between the entities produces more and more information. A structure to store, represent and manage such information is a graph. Graphs are mathematical objects consisting in a vertex set V which represents the entities and an edge set E representing the links between them. Graphs are ubiquitous, representing a variety of natural processes as diverse as friendships between people, communication patterns and interactions between neurons in the brain. Often the chosen data structure to represent a graph is its adjacency matrix A , which can be binary, corresponding to whether there exists an interaction between two vertices, or numerical, corresponding to the strength of the connection.

This high throughput generation calls for techniques able to reduce the volume of the data but still able to preserve the knowledge that carry. The data mining community has taken a strong interest in the task since devising techniques that are able to cope with the intrinsic complexity and dimension of data imply several advantages. The trivial advantage is the reduction of data volume and storage but we can also store only meaningful patterns by filtering the noise which data naturally have. Data compression and summarization techniques are one of the possible approaches to face such problems. The survey of Liu et al [1] provides a comprehensive overview of the state-of-the-art methods for summarizing graph data. The essence of such techniques is to transform the input graph to a more concise representation by preserving the essential patterns which carry the most relevant information to describe the phenomenon under study.

In particular, we exploit the power of Szemerédi's Regularity Lemma to devise a well-found framework for compressing large graphs. The Regularity Lemma was introduced by the Abel laureate Endre Szemerédi in the mid 70's as a tool to prove his famous theorem on arithmetic progressions. This tool roughly states that we can approximate every large graph by partitioning it into a small number of parts such that the bipartite subgraph between almost every pair of parts is random-like. The Lemma guarantees that the partitioning, which in fact represent the compression of the graph, when properly decompressed, preserves some topological properties of the input graph.

The first formulation of the Regularity Lemma is an existence predicate, Szemerédi proved the existence of a regular partition but he did not provide a method to obtain it. In the last decades several mathematicians devised constructive versions of the Regularity Lemma. However, all these algorithms hide a tower constant that makes the implementation unfeasible.

Sperotto and Pelillo [2] reported arguably the first practical application of the Regularity Lemma, the original motivation was to study how to take advantage of the partitions in a pairwise clustering context. In particular, they proposed the first approximate version of the exact algorithm introduced by Alon et al. [3].

In the past few years different approximate algorithms, explicitly inspired by the Regularity Lemma, have been applied in pattern recognition, bioinformatics and social network analysis. The reader can refer to [4] for a survey of these emerging techniques. The last improved version based on the Alon et al. algorithm has been introduced by Fiorucci et al. [5] where they improved the work proposed by Sperotto and Pelillo by introducing new heuristics. This thesis tries to further improve and investigate the conditions of application of the proposed algorithm.

1.1 Contributions

This thesis aims to improve and further investigate the conditions of application of Fiorucci et al.'s approximate version [5] of Alon et al.'s algorithm [3]. In particular, the first contribution of this thesis relies on the introduction of a new heuristic for the refinement step which exploits the internal properties of the classes composing a regular partition. The heuristic idea relies on intuitive motivation which has its roots on the combinatorial nature of the lemma. It tries to mitigate the required tower bound proved by Gowers [6] on the number of the classes composing a regular partition, making, as we will see, the algorithmic implementation able to find partitions even in medium graphs.

The second and even more important contribution is the definition of a pipeline that we called CoDec. The pipeline first compresses a graph using our new version of Alon et al.'s algorithm, and then decompresses the obtained structure into another graph. After the decompression it exploits a new post-decompression step which aims to improve the structural preservation quality.

The work then points out the limits and possible strengths of the technique through an extensive batch of experiments both conducted over synthetic graphs and real networks. We will then analyze the results and draw some considerations on the devised technique.

Last but not least, the outline of the thesis provides, as much as it could, a progressive and concise dive into such deep topic. Several figures and simple conceptual explanations of the mathematical formulas which constitutes the nature of this powerful lemma, have been reported to offer an accessible reading to whom approaches this topic for the first time. The work deliberately does not report all the proofs of the theorems and lemmas cited since we can find them in the original papers, a reference will point us to the correct resource.

1.2 Structure of the work

This thesis has been divided in chapters each of which tries to give a full description of the topic discussed by giving the essential informations needed to understand the core idea. In particular:

- In Chapter 2 we will briefly see the two main mathematical areas where Szemerédi's Regularity Lemma originated. It is important to understand the mathematical nature of the lemma to fully grasp its details. Then, we will progressively dive into its specifications until we will be able to fully state it and complete it with the correct details. After understanding the origins and the power of Szemerédi's Regularity Lemma, we will study Alon et al.'s algorithm which basically is the reformulation of the lemma in algorithmic terms, that is the algorithm that we exploited and tweaked to constitute the compression step of the CoDec procedure.
- In Chapter 3 we will describe the CoDec procedure in all its steps and specifications. In particular, we will give a brief overview of the main steps accompanied with several sketch figures to clarify the logic behind. Then we will describe all its steps that are Compression (we will only focus on the new heuristic developed since the algorithm is explained in Section 2.3), Decompression and Post-Decompression Filtering.

- In Chapter 4 we will discuss the experimental results. Firstly we define the measures used to validate the experiments, then we will move on to describe how the synthetic data is created and where real data is taken from in order to test the framework. We will provide an exhaustive description of the experiments and we will report them through several images and tables. Then we perform some analysis and make some considerations of the results obtained.
- In Chapter 5 we will draw some conclusions that emerged from the experimental results. Then we will point out further investigations that could improve the framework.

Chapter 2

Context and Background

In this section we will see the background needed to understand the nature of the technique devised during this work, in particular we will briefly describe and illustrate some examples of the problem treated in the mathematical areas from which Szemerédi's Regularity Lemma comes from. It is very important to point out the the origins of the method exploited in our work since it gives the right philosophy in order to think in the correct terms.

2.1 Extremal Graph Theory and Ramsey Theory

Szemerédi's Regularity Lemma is one of the most powerful tools in Extremal Graph Theory and it is widely used by mathematicians to prove theorems which require to demonstrate the existence of some substructure in a graph. In fact these kind of questions are the questions mainly posed by these two mathematical areas: Extremal Graph Theory and Ramsey Theory.

Extremal Graph Theory

Extremal Graph Theory is a sub-area of Combinatorics which focuses on the study of any discrete or finite system, in particular, the mathematical objects treated are finite Sets, Colourings and Graphs. It is maybe one of the oldest subareas of Combinatorics and an example of the problems that treats is:

Question. *What is the maximum number of edges in a triangle-free graph of n vertices?*

A complete bipartite graph (with part sizes as equal as possible) has $\lfloor \frac{n}{2} \rfloor \lceil \frac{n}{2} \rceil \approx \frac{n^2}{4}$ edges and no triangle, this result comes from Mantel's theorem which has been proven in 1907 and later generalized by Turán in 1943. The proof is pretty simple and a visual representation is given by Figure 2.1. If we add a single edge in a complete bipartite graph we are, in fact, creating a triangle, so the maximum number of edges of a triangle free graph is the one given by Mantel. More generally the objective of Extremal Graph Theory is to suppose the existence (or not) of a certain sub-structure and then discover how large the graph should be for that sub-structure to exists (or not).

Ramsey Theory

A closely related field of Extremal Graph theory is Ramsey Theory. What it basically aims to do is finding "order in chaos". It is slightly more general than Extremal Graph Theory and the questions which asks treat, not only graphs, but also different mathematical objects.

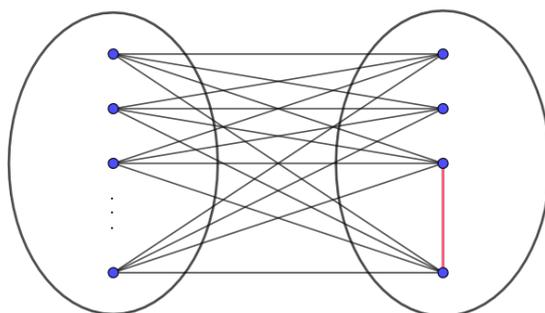


Figure 2.1: A complete bipartite graph, it can be seen that if we add an edge (red) we are creating a triangle.

We could roughly describe Ramsey Theory as the collection of problems where you have some Coloring of some object and you're looking for some very structured subset. If a given object is very very large and colored in an arbitrary way however it is colored you will find the the sub-object you're looking for. Since the Regularity Lemma has been used as a tool to prove the main Szemerédi Theorem which treats arithmetic progressions, we will illustrate a simple example of the problems studied in Ramsey Theory using arithmetic progressions, but before seeing it let's give two definitions:

Definition. An arithmetic progression is a sequence of l natural numbers of constant difference d with an initial value i .

For example the progression with $l = 5, d = 7, i = 23$ is 23, 30, 37, 44, 51.

Definition. A colouring of a non-empty set S with $r \in \mathbb{N}$ colours is just a colourful way of denoting a map $c : S \rightarrow [r]$. The number of $\{1, \dots, r\}$ are usually the colours, and $s \in S$ is painted with colour $i \in [r]$ if $c(s) = i$

Example Now we are ready to illustrate the example: By colouring the natural numbers using only two colours, say red and blue, it is easy to see that 9 consequent numbers e.g. $1, 2, \dots, 9$, are needed to ensure that there is an arithmetic progression of length 3. Why? Let us try to prove the opposite, so suppose the sequence $1, 2, \dots, 9$ does not contain arithmetic progressions of length 3. For this reason, number 1, 5 and 9 cannot be all equally coloured. So assume first that 1 and 5 are red, and 9 is blue. Since 1 and 5 are red, 3 has to be blue. But 9 is also blue, so 6 has to be red. Now 5 and 6 are red, forcing 4 and 7 to be blue. Number 8 has to be red since 7 and 9 are blue, and number 2 has to be red since 3 and 4 are blue. But then 2, 5 and 8 are all red, which is a contradiction. The case where 1 and 9 are red and 5 is blue is treated similarly. On the other hand, the sequence of length 8, given by RBRB-BRBR has no arithmetic progressions of length 3. Thus 9 is a sharp bound for this property. This is also called van der Waerden number $W(2,3)=9$.

2.2 Szemerédi's Regularity Lemma

Szemerédi's Regularity Lemma arises from the composition of the two above areas, as we said it is a powerful lemma which Szemerédi used as a tool to prove his famous theorem on the arithmetic progressions, which let him win the prestigious Abel Prize in 2012. The theorem is a generalization of a much older problem posed by Erdős and Turán in 1936.

By defining upper density of a set $A \subset \mathbb{N}$ as:

$$\tilde{d}(A) = \limsup_{n \rightarrow \infty} \frac{|A \cap [n]|}{n}$$

and by denoting $[n] = \{1, \dots, n\}$ the set of all positive integers from 1 to n , the conjecture says:

Conjecture (Erdős and Turán, 1936). *If $\tilde{d}(A) > 0$, then A contains arbitrarily long arithmetic progressions.*

In 1953, Klaus Friedrich Roth proved that any subset of the integers with positive upper density contains an arithmetic progression of length 3. In 1969, Endre Szemerédi proved that the subset must contain an arithmetic progression of length 4, and then in 1975 proved that any subset with positive upper density must contain arithmetic progressions of arbitrary length, known as Szemerédi’s Theorem [7]. The theorem has been proven in several different ways for example Furstenberg [8] proved it through Ergodic Theory, Gowers [9] using Harmonic Analysis and many other, according to Terry Tao there are at least 16 different ways to prove it. Each of these proves really opened a new area of mathematics, so this conjecture is sometimes described as: “*the greatest triumph of Hungarian school of mathematics*”, which says that if you ask the right question, and try to solve it, the connections between different areas will show themselves and you will discover the theory anyway instead of building it.

In order to prove his famous theorem, which is very intricate, Szemerédi had to introduce several other mathematical lemmas and tools, the Regularity Lemma is one of these. The Lemma may be thought of as an answer to the following vague question:

Vague Version *Every graph G can be “well approximated” by a bounded number of “quasi-random graphs”*

Surely any graph does not look like a random graph, but maybe any graph looks like a combination of random graphs, that’s what he claimed and proved. The interesting thing is that both the approximation of the graph and the amount of randomness of the random graphs only depend by one parameter ϵ . The number of graphs needed is $k(\epsilon)$ a function of ϵ . In order to understand it better we introduce another less vague version.

Less Vague Version *Given any graph G , we can partition the vertex set $V(G)$ into a bounded number of pieces (at most k , say), such that for “almost every” pair (A, B) is approximated well by a quasi-random graph.*

To quantify the “almost every” statement we can say that there could be some pairs which are not “quasi-random” and which we can say nothing about, but the fraction of those will only be $\epsilon \binom{k}{2}$ as well. Another important thing is that each classes pair has its own density, Figure 2.2 visualize the concept

Technical Version Now we can enter into the technical details of the Lemma. First we introduce the notion of density of a classes pair. Let $G = (V, E)$ be an undirected graph with no self-loops, and let $X, Y \subseteq V$ be two disjoint subsets of vertices of G .

Definition. (Pair Density) *We define the edge density of the pair (X, Y) as:*

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

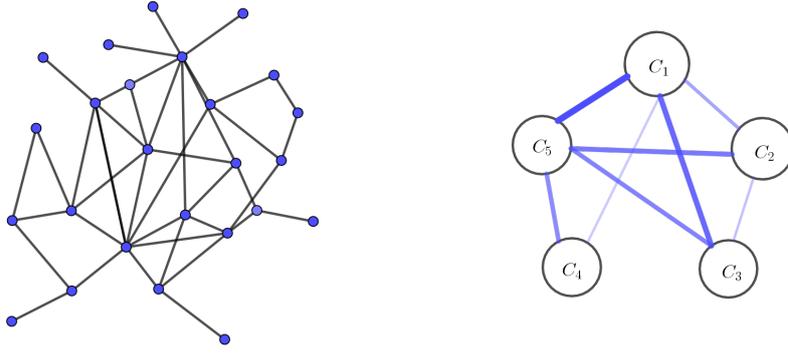


Figure 2.2: Toy visualization of the Regularity Lemma, at left we have the graph and at right we have the approximation, with eight ϵ -regular classes pair, the density of the pair is expressed with the color of the edges. Note: there could be some classes which are not regular (no edge).

This is just the amount of edges shared by the two sets divided by the maximum number of edges that there could be between the two sets, naturally the quantity lies between 0 and 1.

Now we are going to define what it means to a pair to be “quasi-random”, in particular we will call the pair ϵ -regular and the property is defined as follows:

Definition. (ϵ -regularity) A pair (A, B) of sets in $V(G)$ are ϵ -regular if $\forall X \subset A, Y \subset B$ such that $|X| \geq \epsilon|A|, |Y| \geq \epsilon|B|$ we have:

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

Which translated says: given a pair of vertices sets A, B , we say they are ϵ -regular if by taking two sufficiently small subsets of the two the density between them is roughly the same as the density of A, B . This indicates some sort of homogeneity of the structure of the pairs.

In order to define Szemerédi’s Regularity Lemma we last need another tool which is the Embedding Lemma (also known as Key Lemma) introduced by Komlos et al. [10] [11], that actually represents the crucial property of ϵ -regularity. Before stating the lemma we need to clarify some notation: Given a graph H and an integer m , let $H(m)$ denote the graph obtained by “blowing up” each vertex of H to size m , i.e., each vertex $j \in V(H)$ is replaced by a set A_j of size m . Thus $H(m)$ has vertex set $\bigcup_j A_j$, and edge set $\{uv : u \in A_i \text{ for some } ij \in E(H)\}$.

Given a graph H , an integer m and $\delta > \epsilon > 0$, let $\mathcal{G}(H, m, \epsilon, \delta)$ denote the family of graphs G such that $V(G) = V(H(m)), G \subset H(m)$, and $G[A_i, A_j]$ is ϵ -regular and has density at least δ whenever $ij \in E(H)$.

Lemma. (Embedding Lemma) For every $\epsilon > 0, \delta > 0$ there exist a n_0 such that if $V(G) = A \cup B \cup C$, Let $\Delta \in \mathbb{N}$, let $\delta > 0$, and let $\epsilon_0 = \delta^\Delta / (\Delta + 2)$. Let R be a graph, let $m, t \in \mathbb{N}$ with $t \leq \epsilon_0 m$, and let $H \subset R(t)$ with maximum degree at most Δ . If $\epsilon_0 > \epsilon > 0$ and $G \in \mathcal{G}(R, m, \epsilon, \delta + \epsilon)$, then G contains at least $(\epsilon_0 m)^{V(H)}$ copies of H .

Which translated says that for the purpose of finding small subgraphs in a graph G , we can treat sufficiently dense ϵ -regular pairs like complete bipartite graphs. A consequence of the latter fact is that the graph that we want to approximate should have a fairly high density, it is very important to keep this in mind because the lemma does not apply for very sparse graphs. Ultimately it states that the topological structure of the graph which approximates the original one, is guaranteed to be present inside the original one at least a certain number of times. This lemma guarantees the fact that we are able to compress/approximate a graph with a regular partition and still being able to preserve topological structures. Now we are ready to fully states the regularity lemma.

Lemma. (*Szemerédi's Regularity Lemma*) Let $\epsilon > 0, m \in \mathbb{N}$. There exists $k_0 = k_0(\epsilon, m) \in \mathbb{N}$ such that for any graph G there exists a partition¹ $V(G) = C_0 \cup C_1 \cup \dots \cup C_k$ where $m \leq k \leq k_0$ such that:

- $|C_0| \leq \epsilon|V(G)|$
- $|C_i| = \dots = |C_k|$
- (C_i, C_j) is ϵ -regular, for all but $\epsilon \binom{k}{2}$ pairs i, j with $1 \leq i < j \leq k$

Notes All the classes must have the same cardinality except for the set C_0 that can have less nodes inside. As already pointed out there could be some irregular pairs, the fraction admitted can not exceed $\epsilon \binom{k}{2}$. Another very important requirement, has been already mentioned: the analyzed graph should be sufficiently dense, it is a direct consequence of the Embedding Lemma. We also absolutely need to report this important theorem from Gowers [6]:

Theorem. (*Gowers 1997*) For any $\epsilon > 0$, there is a graph so that any application of the Regularity Lemma requires that the number of clusters is at least a number which is a tower of twos of height proportional to $\log \frac{1}{\epsilon}$.

It says that the smaller the ϵ parameter (the approximation factor) the more clusters (classes) we need to have in a ϵ -regular partition, in fact we should have a huge number of classes since there is a tower constant to be respected. This implies that the original graph size is astronomically big and this constitutes the major factor which makes the implementation infeasible.

Said that and with these assumptions in mind, we now move to the algorithmic formulation.

2.3 Alon et al.'s Algorithm

In their original paper Alon et al. [3] showed that the complexity of deciding if a given partition is ϵ -regular is more difficult than actually building one then they devised an algorithm to build one. In fact they showed that creating an ϵ -regular partition can be done with a feasible temporal complexity $O(n^2)$. Let's dig into the algorithm: we will report only the main theorem and some definitions which are crucial to understand the structure of the algorithm.

The following theorem is the core of Alon et al.'s algorithm, take note that we are not going to prove any theorem since all the proofs can be read in the original paper and are all quite long to handle.

Theorem. (*Alon et al., 1994*) For every $\epsilon > 0$ and every positive integer t there is an integer $Q = Q(\epsilon, t)$ such that every graph with $n > Q$ vertices has an ϵ -regular partition into $k + 1$ classes, where $t \leq k \leq Q$. For every fixed $\epsilon > 0$ and $t \geq 1$ such a partition can be found in $O(M(n))$ sequential time, where $M(n) = O(n^{2.376})$ is the time for multiplying two $n \times n$ matrices with 0,1 entries over the integers. It can also be found in time $O(\log n)$ on an EREW PRAM with a polynomial number of parallel processors.

As we already said, the theorem simply states that it exists an algorithm to build such partition and it has a feasible temporal complexity (but keep in mind that we still have the tower constant from Gowers on the number of classes required by a partition). The next question now is: How can we do this? Before moving on we need several other definitions.

¹sometimes also called equitable partition

Let H be a bipartite graph with equal color classes $|A| = |B| = n$, from now on it will represent one of the pair of the final partition. We define the average degree of H as:

$$\bar{d}(A, B) = \frac{1}{2n} \sum_{i \in A \cup B} \deg(i)$$

where $\deg(i)$ is the degree of vertex i .

For two distinct vertices $y_1, y_2 \in B$ define the neighbourhood deviation of y_1 and y_2 with:

$$\sigma(y_1, y_2) = |N(y_1) \cap N(y_2)| - \frac{\bar{d}^2}{n}$$

where $N(x)$ is the set of neighbours of vertex x . For a subset $Y \subseteq B$ the deviation of Y is defined as:

$$\sigma(Y) = \frac{\sum_{y_1, y_2 \in Y} \sigma(y_1, y_2)}{|Y|^2}$$

Let $0 < \epsilon < 1/16$. It can be shown that if there exists $Y \subseteq B, |Y| > \epsilon n$ such that $\sigma(Y) \geq \epsilon^3 \frac{n}{2}$, then at least one of the following cases occurs:

1. $\bar{d} < \epsilon^3 n$ (H is ϵ -regular)
2. there exists in B a set of more than $\frac{1}{8} \epsilon^4 n$ vertices whose degree deviate from \bar{d} by at least $\epsilon^4 n$
3. there are subsets $A' \subset A, B' \subset B, |A'| \geq \frac{\epsilon^4}{n} n, |B'| \geq \frac{\epsilon^4}{n} n$ and $|\bar{d}(A', B') - \bar{d}(A, B)| \geq \epsilon^4$

In case the first two points does not hold one must find A', B' sets as described in point 3. Now the proof of the latter Lemma provides an intuitive way to find the given subsets but, as we said, it is not reported here. Anyway, the main idea is that the procedure selects a subset of the nodes whose degree “deviate” the most from the average. More formally: for each $y_0 \in B$ with $|\deg(y_0) - \bar{d}| < \epsilon^4 n$ we find the set of vertices $B_{y_0} = \{y \in B | \sigma(y_0, y) \geq 2\epsilon^4 n\}$. The proof given in the original paper guarantees the existence of at least one such y_0 for which $|B_{y_0}| \geq \frac{\epsilon^4}{4} n$. The subsets $B' = B_{y_0}$ and $A' = N(y_0)$ are the required ones. These two subsets represent the collection of nodes that contribute more to the irregularity of the pair (A, B) , in fact they play an important role and we will address them as certificates. We will call complements all the other nodes of the class which are not part of the certificates.

Alon et al.’s algorithm allows us to find the certificates which contribute more to the irregularity of a pair, before stating the algorithm we will report an important result of Szemerédi which combined with Alon et el.’s algorithm allows us to build such partitions. Given an ϵ -regular partition \mathcal{P} of a graph $G = (V, E)$ into C_0, C_1, \dots, C_k classes Szemerédi [12] defines a measure called index of partition:

$$sze_ind(\mathcal{P}) = \frac{1}{k^2} \sum_{s=1}^k \sum_{t=s+1}^k d(C_s, C_t)^2$$

Since $0 \leq d(C_s, C_t) \leq 1, 1 \leq s, t \leq k$, it can be seen that $sze_ind(\mathcal{P}) \leq \frac{1}{2}$

In his main proof, Szemerédi state that if a partition \mathcal{P} violates the regularity condition, then it can be refined by a new partition and, in this case, the latter measure should increase.

Alon et al.’s Alogirhtm Now we are ready to see the main algorithm:

1. Create the initial partition: Arbitrarily divide the vertices of G into an equitable partition \mathcal{P}_1 with classes C_0, C_1, \dots, C_b where $|C_i| = \lfloor \frac{n}{b} \rfloor$ hence $|C_0| < b$

2. Check Regularity: For every pair (C_r, C_s) of \mathcal{P}_i , verify if it is ϵ -regular or find $X \subseteq C_r, Y \subseteq C_s, |X| \geq \frac{\epsilon^4}{16}C_1, |Y| \geq \frac{\epsilon^4}{16}C_1$ such that $|\bar{d}(X, Y) - \bar{d}(C_r, C_s)| \geq \epsilon^4$
3. Count regular pairs: if there are at most $\epsilon \binom{k_i}{2}$ pairs that are not ϵ -regular then, stop. \mathcal{P}_i is an ϵ -regular partition
4. Refine: Apply the refinement algorithm and obtain a partition \mathcal{P}' with $1 + k_i 4^{k_i}$ classes
5. Reiterate

Besides all the notation and mathematics here we have an algorithm which guarantees the possibility to reach an ϵ -regular partition that respects the strong regularity conditions. Several other algorithms have been created most of them relaxed the conditions of ϵ -regularity to a weak notion, in particular the Frieze-Kannan algorithm [13] is the most famous one, which is based on an intriguing relation between the regularity conditions and the singular values of matrices.

Chapter 3

Description of the devised technique

3.1 Overview of the CoDec Procedure

Thanks to the previous sections we now have a basic context and idea of the domain of the devised technique. In particular, now we have a grasp on how Alon et al’s algorithm works, therefore we can describe the CoDec Procedure (Co-mpression Dec-ompression) which is the core of the work done during this thesis. We will split the algorithm in different steps as Figure 3.1 shows, but before describing the work, let’s just clarify some aspects that are valid throughout all the description.

Notes

- Since we are entering to a more algorithmic description of the work, from now on we will use matrices as data structures to represent any undirected graph with no self loops. In particular the matrices are symmetric with an empty diagonal (no self loops).
- From now on with the words “reduced graph” or “compression” we will refer to a ϵ -regular partition \mathcal{P} (of size k) of an undirected graph with no self loops achieved by running the approximated version of Alon et al.’s algorithm in addition with the new heuristic. We sometimes address it with the symbol \mathcal{P} or in its matrix representation with *RED* (from REDuced graph).

Description The procedure takes as input an undirected (weighted or unweighted) graph of n nodes, expressed as a symmetric $n \times n$ matrix G , then:

1. It searches and collect different ϵ -regular partitions $\mathcal{P}_1, \mathcal{P}_2 \dots$ by running several approximated Alon et al.’s algorithm instances in conjunction with the new refinement heuristic, over a set of ϵ candidates. In particular, the procedure uniformly sample the ϵ parameter space in the interval $[0, 0.5]$.
- 1b. We select the best partition \mathcal{P}^* from the collected ones. The partition is the one which both maximizes the partition cardinality $|\mathcal{P}|$ (therefore maximizes the Szemerédi Index according to its formulation) and minimizes the value of the parameter ϵ (the resulted graph should be a better approximation of the graph according to the theory). We then build a $|\mathcal{P}^*| \times |\mathcal{P}^*|$ matrix to contain the compression of the matrix G that we will call *RED*.
2. The procedure decompress *RED* by exploiting the Embedding Lemma (Section 3.3), giving us a weighted $n \times n$ matrix that we will call *SZE* and which should carry the structural properties of the graph.

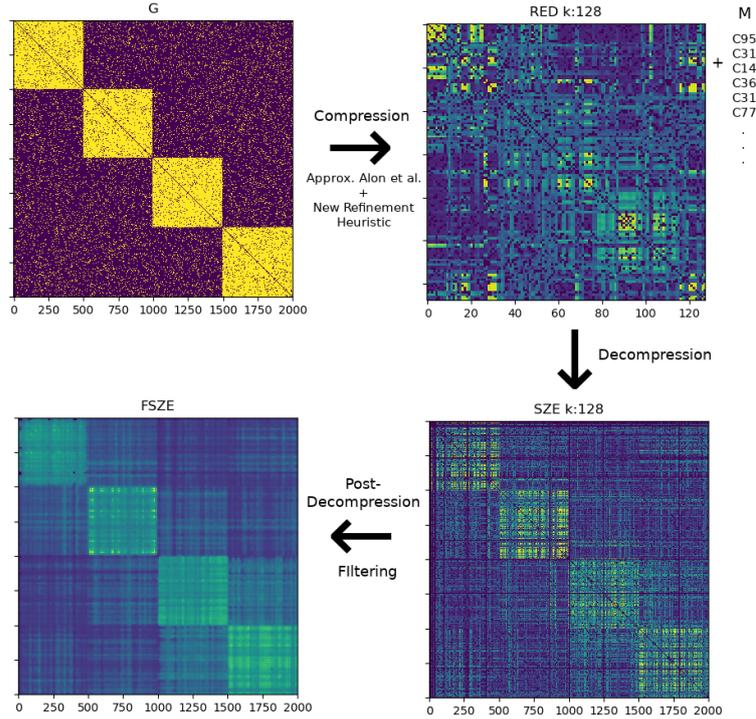


Figure 3.1: Schematic representation of the CoDec Procedure. The input graph is a noisy graph with four balanced clusters representing the structure carrying information. The compression step produces a partition of 128 classes coded in the RED matrix plus a vector of membership M . The second phase exploits the Embedding Lemma to reconstruct the compressed structure. The third and last step Post-Decompression applies a filter to highlight the structural preservation.

3. We then apply a filtering $f(\cdot)$ technique to the matrix to highlight the patterns in the SZE matrix giving us the final $n \times n$ matrix $f(SZE)$ or $FSZE$. This procedure is described in Section 3.4.

Algorithm 1 CoDec Procedure. Note that for simplicity we consider the first 6 steps as compression

```

1: procedure CODEC( $G$ )
2:   for  $\epsilon$  in  $(0, 5)$  do
3:      $\mathcal{P}_i = \text{APPROX\_ALON}(G, \epsilon)$ 
4:     Store  $\mathcal{P}_i$  in a vector  $V$ 
5:    $\mathcal{P}^* = \text{BEST\_PARTITION}(V)$ 
6:    $RED, M = \text{COMPRESS}(\mathcal{P}^*)$ 
7:    $SZE = \text{DECOMPRESSION}(RED, M)$ 
8:    $FSZE = \text{POST-DECOMPRESSION}(SZE)$ 

```

Observations In point 1 a smarter approach has been considered but not tested yet, it consists in a binary search with a parametrized step size over the ϵ parameter space $[0, 0.5]$. This improvement would significantly reduce the number of runs and speed up the overall running time, although the required time for a single run in a medium-size graph ($n < 10000$) is acceptable. This is possible due to the fact that, experimentally, it has been noted that the cardinality of the partitions found grows as ϵ grows until we are no more able to find any useful partition.

The compression in point 1 is expressed with a mapping from the nodes to their membership to a particular class through a $1 \times n$ vector, which encodes the nodes class membership M , plus a $|\mathcal{P}^*| \times |\mathcal{P}^*|$

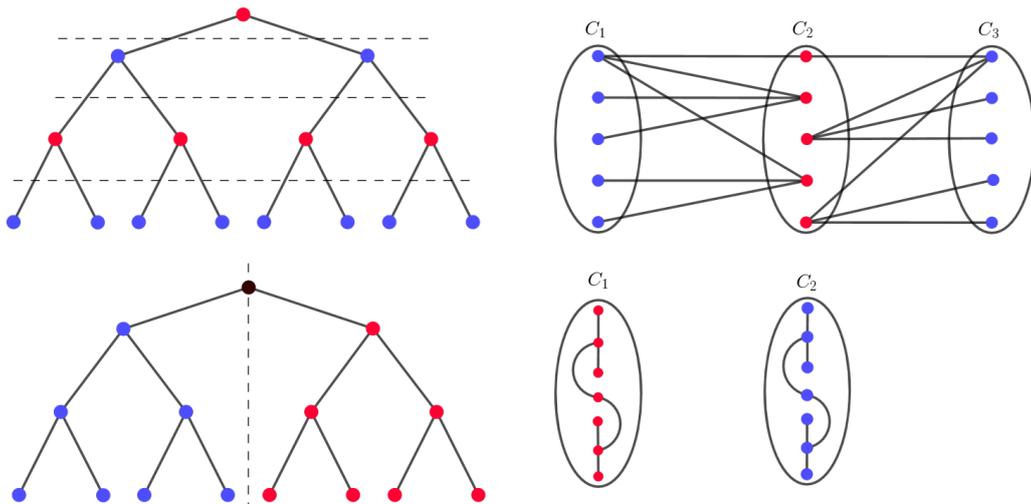


Figure 3.2: A toy example: it can be seen that by achieving different indegree average values, we converge to a particular coloring of the graph. In the upper image we have an internal average degree of 0 while in the case below we have an internal average degree of 1.81 and in fact they represent two different colorings.

matrix RED to encode the densities between the classes¹.

3.2 A New Refinement

In this Section we see the main idea that this work proposes as heuristic for the refinement procedure of Alon et al.'s approximated algorithm.

As we said, in the original paper of Alon et al. they show that the Regularity Lemma can be formulated in algorithmic terms, but they do not provide a constructive version of how to reach the desired partition through the refinement step. Other works tried to implement some kind of heuristics [14] but the choices are pretty simple and only take in consideration properties of the nodes. With this thesis we introduce a new heuristic which exploits more grounded intuitions.

Idea The key idea of this refinement is that we now try to reach the final partition by exploiting properties of the classes instead of the nodes. In particular it can be drawn a correspondence between the average internal degree of the nodes in a class and a particular graph coloring. As can be seen in Fig. 3.2 when we drive the internal degree² of the nodes inside a class, we are coloring the graph by highlighting some regular patterns. We then assume that the information to be preserved is represented in the repetition of some sub-patterns and we drive the refinement to keep these substructures between the classes. By trusting the Alon algorithm intuition we then aim to converge to a meaningful partition which highlights the structure of the input graph. Obviously the analyzed graph must have some intrinsic patterns in order for this idea to work.

Details We need to specify some details and decisions made to fully understand the heuristic:

- At the very beginning we need to provide a starting partition for the approximated Alon et al.'s algorithm, the choice is to split the graph in 4 classes and distribute the nodes in a random way.

¹We use only half of the matrix achieving a $O(|\mathcal{P}|^2)$ space complexity with $\frac{1}{2}$ constant factor

²By internal degree we mean the degree of a node with respect to the nodes in the same class

- In the original definition the number of classes at each iteration should grow astronomically with a tower function as lower bound. Since this detail is computationally infeasible we decided that the number of classes always doubles at each iteration.

Notations We need to clarify the notation since there are several entities to be considered:

- X_i, X_j : certificates of the irregular pair C_i, C_j
- \bar{X}_i, \bar{X}_j : complements of the irregular pair C_i, C_j . Just to be clear: $X_i \cup \bar{X}_i = C_i$
- $d(C_i, C_j)$: density between the two classes as specified in Section 2.2. If $i = j$ then we consider the internal density (or in-density) defined as

$$d(C_i, C_i) = \frac{e(C_i, C_i)}{|C_i|^2}$$

- C_i^1, C_i^2 : notation to express the two new classes created from an irregular pair ³
- indegree : degree of a node calculated with respect to the nodes of the same class
- in-hubs : hub nodes with respect to the indegree i.e. nodes with highest indegree value

Description We are now ready to define the new refinement procedure. Given a partition $\mathcal{P}_i = C_0, C_1, \dots, C_k$ at some iteration i of approximated Alon et al.'s algorithm:

1. Randomly pick a class C_i
2. If C_i is irregular with other classes then choose the candidate C_j which maximizes the following score:

$$S = d(C_i, C_j) + (1 - |d(C_j, C_j) - d(C_i, C_i)|) \rightarrow [0, 2]$$

Now we have a classes pair candidate (C_i, C_j) ready to be refined⁴. We then process the two certificates of the pair, in particular we have two possible scenarios for each certificate. We consider only X_i for the sake of illustration, but the exact same steps will be applied for X_j :

- (a) Case $d(X_i, X_i) \geq 0.5$ (Densification)
 - i. Sort X_i nodes by indegree to reveal in-hubs $\{h_1, h_2, \dots, h_{|X_i|}\}$
 - ii. Unzip hubs to new classes : $C_i^1 : \{h_1, h_3, h_5, \dots, h_{|X_i|-1}\}$ and $C_i^2 : \{h_2, h_4, h_6, \dots, h_{|X_i|}\}$
 - iii. Fill C_i^1 and C_i^2 up to $|C_i|/2$ by taking the “most”⁵ connected nodes from the union of the complements $\bar{X}_i \cup \bar{X}_j$
- (b) Case $d(X_i, X_i) < 0.5$ (Sparsification)
 - i. Random sampling $|X_i|/2$ nodes into C_i^1 and C_i^2
 - ii. Fill C_i^1 and C_i^2 up to $|C_i|/2$ by taking the “least” connected nodes from the union of the complements $\bar{X}_i \cup \bar{X}_j$

3. If class C_i is ϵ -regular with all the other classes

- (a) Sort C_i nodes by indegree to reveal in-hubs $\{h_1, h_2, \dots, h_{|C_i|}\}$

³just for clarity there must be C^{j+1}, C^{j+2} too, but for brevity we just consider the class i over the whole description

⁴each class will be split in two, giving birth to 4 new classes $C_i^1, C_i^2, C_j^1, C_j^2$

⁵Weighted case: the nodes with the highest value. Unweighted case: just the nodes which share the most connections with the considered set

(b) Unzip nodes to new classes : $C_i^1 : \{h_1, h_3, h_5, \dots, h_{|C_i|-1}\}$ and $C_i^2 : \{h_2, h_4, h_6, \dots, h_{|C_i|}\}$

4. Repeat the above points until there are no more classes left.

* Since the new refinement split each class in two, it could happen that we have an odd cardinality of a class, in this case we put the exceeding node in the “trash-set” C_0 . When the cardinality of this set exceeds ϵn and we have enough nodes, we equally distribute them between the classes, otherwise the procedure declares the partition irregular.

Algorithm 2 Refinement step performed at the i -th iteration of the approximated Alon et. al’s algorithm

```

1: procedure REFINEMENT( $\mathcal{P}^i$ )
2:   for each class  $C_i$  in  $\mathcal{P}^i$  do
3:     if  $C_i$  is  $\epsilon$ -regular with all the other classes then
4:        $C_i = \text{SORT\_BY\_INDEGREE}(C_i)$ 
5:        $C_i^1, C_i^2 = \text{UNZIP}(C_i)$   $\triangleright$  This statement may add a node to  $C_0$  if  $|C_i|$  is odd
6:     else
7:       Select  $C_j$  with most similar internal structure
8:       Get certificates  $(X_i, X_j)$  and complements  $(\bar{X}_i, \bar{X}_j)$  of the pair
9:       if  $d(X_i, X_j) < 0.5$  then
10:         $C_i^1, C_i^2 = \text{SPARSIFICATION}(X_i, \bar{X}_i \cup \bar{X}_j)$   $\triangleright$  This statement may add a node to  $C_0$ 
11:      else
12:         $C_i^1, C_i^2 = \text{DENSIFICATION}(X_i, \bar{X}_i \cup \bar{X}_j)$   $\triangleright$  This statement may add a node to  $C_0$ 
13:      Perform step 9,10,11,12 for  $X_j$ 
14:   if  $|C_0| > \epsilon n$  and  $|C_0| > |\mathcal{P}^{i+1}|$  then
15:     Uniformly distribute nodes of  $C_0$  between all the classes
16:   else
17:     return  $(\mathcal{P}^{i+1}, \text{irregular})$ 
18:   return  $(\mathcal{P}^{i+1}, \text{regular})$ 

```

Observations Since this new refinement is a heuristic we don’t have any formal proof of the correctness, but we can give some intuitive explanations of the procedure.

The point 2 defines a score in order to select the irregular pair C_i, C_j that we will refine (split in C_i^1, C_i^2 and C_j^1, C_j^2) in the next step. In particular this score aims to select pairs who share the most similar internal structure, this is to follow the intuition of the heuristic which has been illustrated with Figure 3.2. By selecting the most similar classes pair we aim to select a particular topological substructure, at least this should be the connection. After that we proceed to purify the two classes with step 2a and 2b in other words we aim to converge to a pattern.

The Densification branch 2a also contribute to maximize the Szemerédi index since we are unzipping in-hub nodes to the new two classes, therefore this should result in a higher density between the new classes and a higher Szemerédi index (recall: the Szemerédi Index grows if we have a very dense graph or if we have a huge amount of classes).

3.3 Graph Decompression

In this Section we describe the decompression phase where we take the structure that represents the best partition and we decompress it to have a weighted matrix of the same size of the original one. In this particular phase we fully exploit the Embedding Lemma that, as we said, is the reason why the Regularity Lemma is useful. At this point we have selected the best partition \mathcal{P}^* coded through

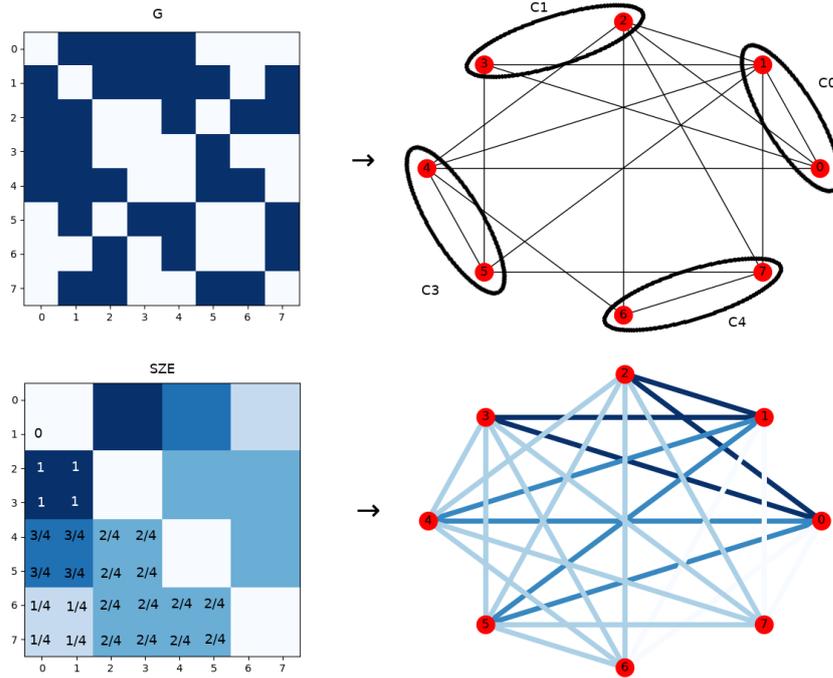


Figure 3.3: Consider the partition pair (C_3, C_4) its density is $\frac{1}{2}$ since there could be 4 edges between them but in G , which is our corrupted graph, we have only 2 edges between them. In the *SZE* reconstruction we will fully connect the nodes inside the classes by setting the edges weights equal to the density between the pair. From this toy example we could also see how the Szemerédi’s Regularity Lemma can be applied as a tool to approximate the corrupted graph G and to reveal the structural pattern of the original graph (let’s assume that the original graph was a fully connected graph).

a $1 \times n$ vector which represents a mapping of the nodes to their respective class, plus a $|\mathcal{P}^*| \times |\mathcal{P}^*|$ matrix *RED* to encode the densities between the classes.

The procedure decompresses the matrix by fully connecting the nodes of each ϵ -regular classes pair. It is mandatory to fully connect the nodes because that’s a requirement of the Embedding Lemma as we already described in Section 2.2. In particular given an ϵ -regular pair with density d we will fully connect the nodes by assigning a weight equal to d . This means that the decompression phase generates a weighted undirected graph even if the input graph is unweighted. The resulting matrix encodes the structural patterns of the original graph, ideally higher the density between the pair the more structure we captured. The latter consideration is a consequence of the fact that higher density pairs contribute the most to the partition index, therefore it means that a pair with high density should contain more “structural information”. We will call the reconstructed/decompressed $n \times n$ weighted matrix *SZE*.

The reconstruction used in the experiments is very simple and it is label-preserving, this is a very important feature because allows us to visualize the results and compute some error measures. In order to follow the theory, the reconstruction phase does not put edges between nodes of the same class, or between the ϵ -irregular pair, although the framework provides such possibility. In particular if we also connect the irregular pairs the result would be a fully connected approximated graph, experimentally it has been noted that this feature does not significantly improve the results.

Another decompression feature has been taken in consideration: we can try to keep track also of the internal densities of the classes. The idea is to exploit this information when we blow up the reduced matrix. In particular we could create an Erdős-Rényi graph with probability equal to the internal density of the class in order to approximate its internal structure. Unfortunately this feature



Figure 3.4: An example of median filtering, taken from Szeliski [15] book.

has not been extensively studied, so the results has been omitted.

3.4 Post-Decompression filtering

In this phase we abstract the notion of the graph and we treat its data representation i.e. the matrix purely as an image in order to perform some morphological operations. This turns out to be a good trick to mitigate the approximation of our algorithm and to let the structural pattern emerge from the decompressed matrix.

In particular after the Alon compression and reconstruction, we filter the matrix SZE with a fixed size median kernel giving birth to the $FSZE$ matrix. Median filtering is a nonlinear morphological operation mainly used in Computer Vision since it is very effective at removing “salt and pepper” noise in images Figure 3.4 illustrate such application. It works by moving a kernel through the image pixel by pixel, replacing each value with the median value of neighbouring pixels.

The pattern of neighbours is called the “kernel” (or window) and it slides over the entire image. The median is calculated by first sorting all the pixel values from the window into numerical order, and then replacing the pixel being considered with the middle (median) pixel value (50 percentile). The filtering phase is essential to clean up the reconstructed matrix since in the reconstruction phase we fully connects the classes and obviously we could add some extra edges. In particular, as we mentioned in the decompression section, the edges with higher weight should carry more structural information, to filter the noise of the low density weights we run a median kernel that should highlight the matrix structural properties.

Experimentally we decided to use a fixed window size, but in principle one can study the optimal size in function of some graph invariant although this possibility has not been taken in consideration since a fixed size provided a good solution.

When we apply any filter in images there is a problem when taking in account the pixels in the border of the image, in our case we adopted a reflection padding technique to handle matrix borders. There is not a strong motivation to use a reflection padding, it is purely an implementation detail since it does not severely affect the result.

Chapter 4

Experimental Results

In this section we report the experimental setup, the results obtained and we will draw some considerations. First of all let's clarify that two different experiments have been conducted. The first one reported in Section 4.2 studies the CoDec Procedure applied to various Synthetic Graphs, the objective is to do a deep study of the devised technique. The second type of experiments reported in Section 4.3 uses Real Networks, the aim is to verify the procedure over real cases and to assert the approximation feature of the framework.

Almost all the experiments involving medium size graphs ($n < 5000$ nodes), have been conducted with my personal machine: an Intel Core i5 @ 2.60GHz HP Pavilion 15 Notebook with 8GB of RAM (DDR3 Synchronous 1600 MHz) running Arch-Linux with kernel version 4.14.4-1, but the experiments involving graphs with $n \geq 5000$ have been conducted in the [SCSCF](#) which is the scientific computational cluster made of 7 nodes of the Ca' Foscari university. The reason for the latter choice is for both time (since we performed a lot of repetitions) and resources reasons (my personal Notebook we are able to process single runs of graphs up to 17000 nodes, but it couldn't work all day for several days without causing damage to the hardware).

4.1 Evaluation measures

In order to validate the results, in the experimental evaluation we have used three main measures. The first two are somehow a "direct" estimation of dissimilarity between any two matrices, to evaluate the approximation quality of the method. As we will discuss we could argue if these two provide a good estimation. In fact we will discuss it later on the consideration section. The last measure estimates the structural preservation of the technique by comparing two clustering vectors.

Direct Measures

Two direct measures have been used to compute the dissimilarity between the original graph and the result of our procedure. To illustrate the measures let's define two matrices $\mathbf{A} = (a_{ij})$ and $\mathbf{B} = (b_{ij})$ both of size $n \times n$:

- Normalized l_1 -distance

$$l_1(\mathbf{A}, \mathbf{B}) = \frac{\sum_{i=1}^n \sum_{j=1}^n |a_{ij} - b_{ij}|}{n^2} \rightarrow [0, 1]$$

- Normalized l_2 -distance

$$l_2(\mathbf{A}, \mathbf{B}) = \frac{\sqrt{\sum_{i=1}^n \sum_{j=1}^n (a_{ij} - b_{ij})^2}}{n^2} \rightarrow [0, 1]$$

A high value of one of the two measure means that the two matrices are very “distant” or dissimilar, while lower values means similar matrices. We would like to have low values.

Structural Preservation Measure

In order to measure the preservation of the structure in cluster-like graphs, it has been implemented a column-wise Knn Voting System clustering technique (we will call it KVS). Given a graph, the KVS produces a clustering vector, to validate the correctness of the grouping, then we calculate the Adjusted Random Index (*ARI*) between the graph’s original clustering L and the predicted one C .

KVS Description KVS-clustering technique works as follows:

1. Given a weighted matrix, for each row of the matrix we take the column positions of the k highest values
2. We check, in the these positions, which are the true clustering labels by looking at the respective positions in L , giving us k labels predictions
3. We set the current row label prediction to the label which is the most repeated in the k predictions

Since the KVS-clustering requires a parameter k to run, we try different values and then we pick the one which maximizes the ARI, specifically we tried with three different values of $k = 5, 7, 9$ we must be sure that k is an odd number to always have point 3 satisfied.

ARI Description The KVS clustering technique outputs a $1 \times n$ array of predicted labels C , when we calculate the Adjusted Random Index with respect to the labeling vector L , the similarity score is between -1 and 1. A random prediction clustering has an ARI close to 0 while 1 stands for perfect match, -1 means completely opposite labeling.

This estimation has been introduced by Hubert et al. [16] and it is a measure to compare two clusterings. If L is a ground truth clustering vector and C is the clustering that we want to evaluate, let us define a and b as:

- a : the number of pairs of elements that are in the same set in L and in the same set in C
- b : the number of pairs of elements that are in different sets in L and in different sets in C

The raw (unadjusted) Random Index (RI) is then given by:

$$RI = \frac{a + b}{L_2^{n_{samples}}}$$

Where $L_2^{n_{samples}}$ is the total number of possible pairs in the dataset (without ordering). However the RI score does not guarantee that random label assignments will get a value close to zero (especially if the number of clusters is in the same order of magnitude as the number of samples). To counter this effect we can discount the expected Random Index $E[RI]$ of random labelings by defining the adjusted Rand index as follows:

$$ARI = \frac{RI - E[RI]}{\max(RI) - E[RI]}$$

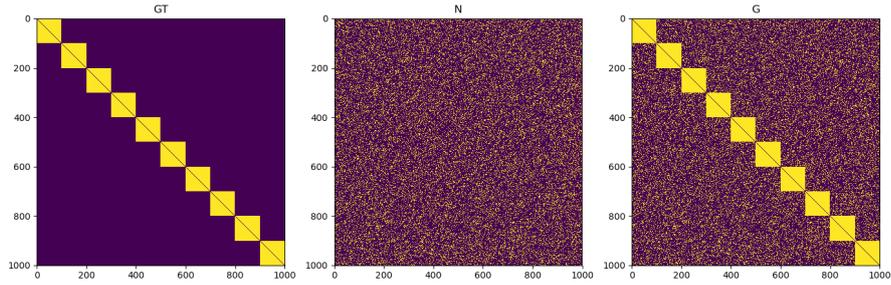


Figure 4.1: A Synthetic Graph G composed by 1000 nodes with internoise of 0.2 and 10 clusters. The labeling vector will be a 1×1000 vector where the first 10 elements will be set to 1 (first cluster label) then the next 10 elements will be set to 2 and so on to 10.

4.2 Synthetic Graphs Experiments

Synthetic Graph Generation

In this subsection we will see how the synthetic graphs, used for the experiments, are generated. As already said, we are interested to undirected graphs that can be seen as a composition of “structure” which carries the knowledge of the phenomenon, plus some “noise”.

That said, we can formalize a bit the notion of the graphs that we will analyze, in particular we can see them as:

$$G = GT + N$$

Where G represents the graph, GT (namely Ground Truth) represents the structure and N represents the noise.

Given a graph G of n nodes, we model the noise with Erdős-Renyi Graph (ERG) of size n with probability of connection between two nodes p (from now on we will call this parameter inter-cluster noise level, or inter-noise). The structural part GT is a graph of n nodes with C disconnected groups (cluster) of nodes of variable size.

For the evaluation of the experimental results, we also need a vector of labels L that represents the correct clustering of the nodes according to the ground truth. The labeling vector is a $1 \times n$ vector which indicates the membership of a node to the respective cluster. Figure 4.1 exemplifies the creation of a synthetic graph.

Basic Version Now we describe the basic procedure to generate unweighted Synthetic Graphs:

- Generate a $n \times n$ adjacency matrix A with A_{ij} equal to a realization of a $U(0, 1)$ r.v..
- If $A_{ij} \leq$ inter-noise level then put an edge (we’re creating the ERG i.e. N)
- Create C disconnected cliques (clusters) on the diagonal, each one with size $k \times k$ where $k = n/C$ (we’re creating a GT made of balanced cliques).

This procedure outputs an unweighted adjacency matrix G which represent an undirected unweighted Synthetic Graph.

Extension Now we can extend this generation procedure to produce a wide spectrum of graphs, in particular let’s consider these possibilities:

- Adding weights both to noise and to the structure.

- Imbalance the clusters by varying the size of each one.
- Introducing noise inside the clusters (we will refer to this noise as intra-noise or corrosion).

In this work not all the combination of the extensions have been analyzed since the spectrum is almost infinite, but the generation tool provides such possibility.

Experimental Setup

In this section we will describe the setup of the experiments conducted and we will introduce the naming conventions used to report in the results.

Experiment 1

- For the first experiment several dimensions of the graphs have been considered. In particular we tried graphs with size $n = 1000, 2000, 5000, 10000, 25000$.
- In particular, fixing a size n , we then generate 5 different graphs with inter-noise = 0.2, 0.4, 0.5, 0.6, 0.8.
- We run the CoDec Procedure and calculate the measures specified in Section 4.1 to validate the results.
- We performed a thresholding to unweight the FSZE matrix by minimizing this quantity:

$$\min_{0 < t < 1} l_2(FSZE_t, GT)$$

Which translated means that we tried to threshold the weights of the $FSZE$ matrix by minimizing $l_2(FSZE, GT)$ measure. This operation gives birth to the $UFSZE$ (Unweighted Filtered SZEmeredi matrix) (fourth column of the visual representation in Figure 4.4). This step has been inserted in order to study some possible relation between the values of the matrix produced by the CoDec Procedure and the density of the original graph.

- The experiment has been repeated a certain number of times r to generate some statistics. With medium size graphs the repetitions are 20 while with big graphs 5 or 3.
- We then visualize a single run

Experiment 2

- For the second experiment we considered a fixed dimension for the synthetic graphs $n = 4000$.
- Given a fixed size n , we generated imbalanced and balanced cluster-like graphs, by varying both internoise and intranoise, exploring all the combinations of these values: 0.2, 0.4, 0.5, 0.6, 0.8.
- We run the CoDec Procedure and calculate the Adjusted Random Index measure specified in Section 4.1 to validate the preservation of the structural properties.
- The experiment has been conducted one time, therefore there is no repetition.
- We then visualize a single run of a single graph.

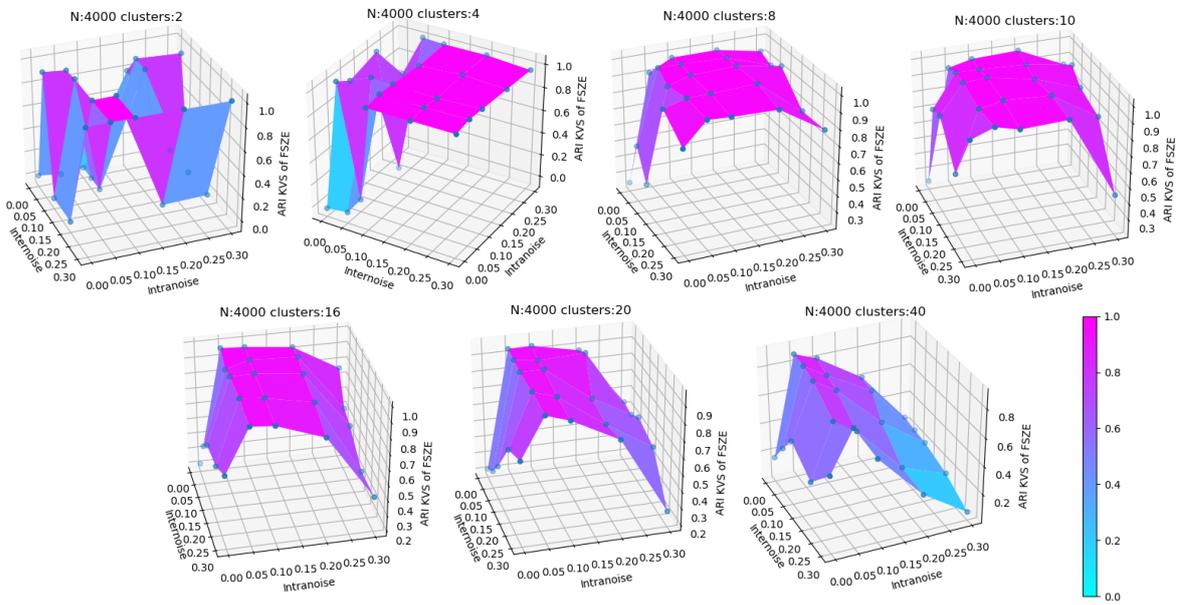


Figure 4.2: Experiment 2: Internoise and Intra-noise analysis over noisy graphs with imbalanced clusters as structure

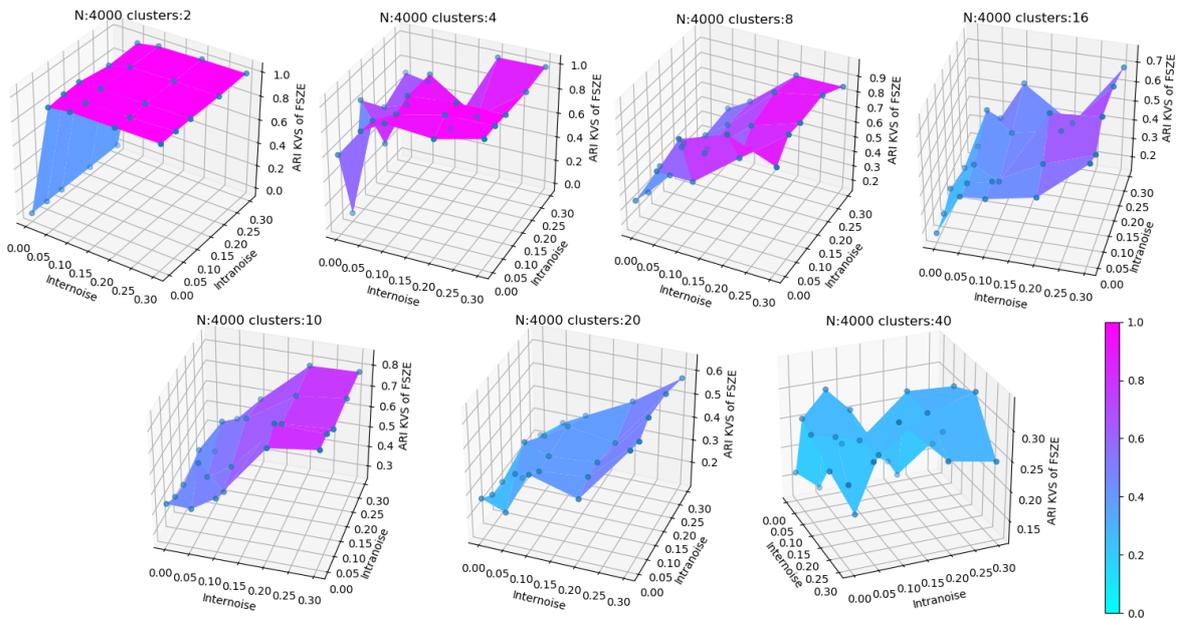


Figure 4.3: Experiment 2: Internoise and Intra-noise analysis over noisy graphs with balanced clusters as structure.

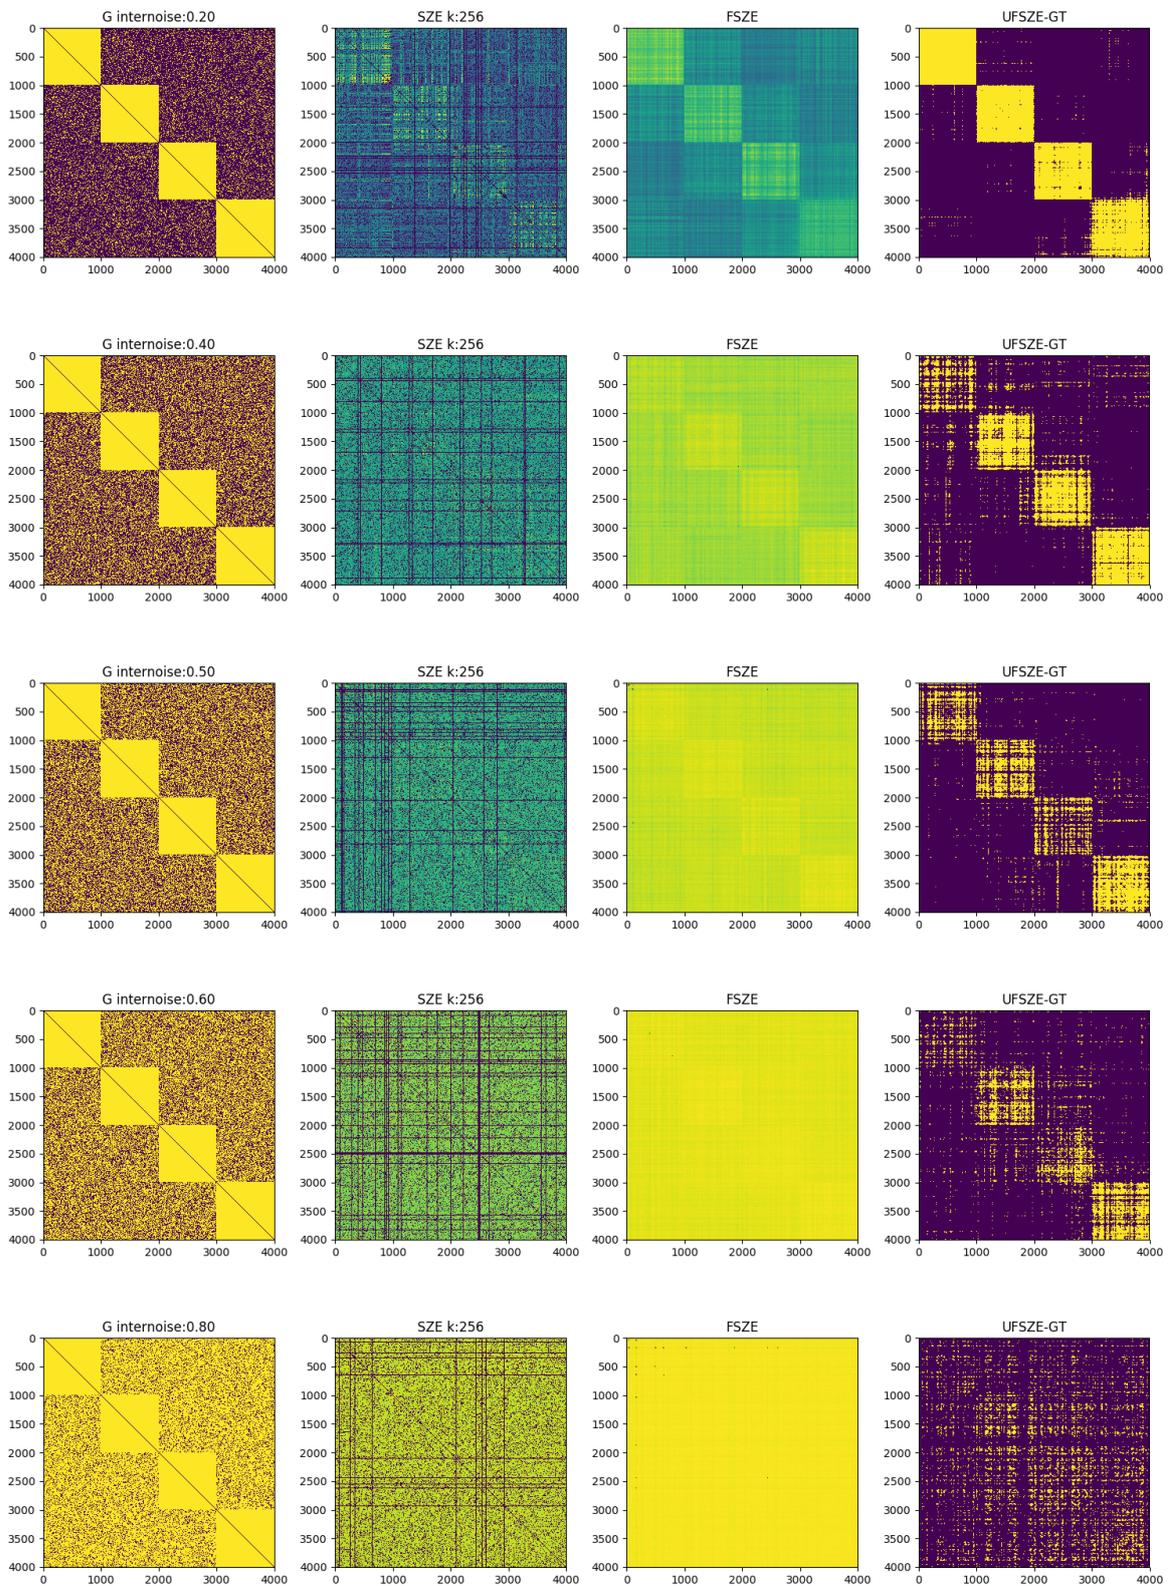


Figure 4.4: Visualization of a single repetition of Experiment 1

$n = 1000 \quad r = 20$			
Inter-noise	Avg ϵ	Avg k	Avg size_idx
0.2	0.275	64	0.0836
0.4	0.275	64	0.1498
0.5	0.275	64	0.1921
0.6	0.275	64	0.2418
0.8	0.275	64	0.3562

Inter-noise	Avg ARI Kvs FSZE \pm std	$l_2(FSZE, GT) \pm$ std	$l_1(FSZE, GT) \pm$ std
0.2	0.9069 \pm 0.01591	0.4119 \pm 0.0007	0.3928 \pm 0.0004
0.4	0.9466 \pm 0.00451	0.5016	0.5009 \pm 0.0009
0.5	0.8068	0.5543	0.5467
0.6	0.7587	0.6071	0.5862
0.8	0.0669	0.7299	0.669

Table 4.1: Experiment 1: graphs with $n = 1000$ and 20 repetitions. The standard deviation has been omitted in the cases where its below 10^{-5}

$n = 2000 \quad r = 20$			
Inter-noise	Avg ϵ	Avg k	Avg size_idx
0.2	0.2975	128	0.0912
0.4	0.275	128	0.1515
0.5	0.275	128	0.1942
0.6	0.275	128	0.2434
0.8	0.275	128	0.3586

Inter-noise	Avg ARI Kvs FSZE \pm std	$l_2(FSZE, GT) \pm$ std	$l_1(FSZE, GT) \pm$ std
0.2	0.997 \pm 0.0012	0.3764 \pm 0.0006	0.3549 \pm 0.0006
0.4	0.9671	0.4972	0.4966
0.5	0.8228	0.5543	0.5468
0.6	0.7947	0.6086	0.5876
0.8	0.0842	0.7299	0.6686

Table 4.2: Experiment 1: graphs with $n = 2000$ and 20 repetitions

4.3 Real Network Experiments

The majority of the experiments have been conducted with synthetic graphs because we have a ground truth to validate the structural preservation of the CoDec procedure. Since our framework is able to both approximate the given graph and highlight the underlying structural pattern we opted to test the first property with real world data. Although one of the requirements of the algorithm is to have a dense graph, we tested the devised procedure with real world datasets which, instead, are very sparse. The problem is that we don't have the ground truth and we can only calculate direct dissimilarity measures between the output matrix and the original one. The datasets have been taken from two famous repositories of real networks: the Stanford Large Network Dataset Collection SNAP [17] and from the Konect repository of the University Koblenz-Landau Konect [18].

The setup of the experiments is very simple and we describe it as follows:

- All the networks used are unweighted and undirected.
- We read the network in a matrix data structure and we apply the CoDec Procedure.
- Since we don't have the ground truth to test the structural preservation we calculate the $l_2(G, FSZE)$ and the $l_1(G, FSZE)$ metrics to measure the dissimilarity of the two matrices.

$n = 5000 \quad r = 5$			
Inter-noise	Avg ϵ	Avg k	Avg <i>sze_idx</i>
0.2	0.2435	256	0.0936
0.4	0.2525	256	0.1529
0.5	0.2525	256	0.1958
0.6	0.23	256	0.2445
0.8	0.23	256	0.3599

Inter-noise	Avg ARI Kvs FSZE \pm std	$l_2(FSZE, GT) \pm$ std	$l_1(FSZE, GT) \pm$ std
0.2	0.9981 \pm 0.0028	0.3695 \pm 0.0025	0.3525 \pm 0.0045
0.4	0.9935 \pm 0.0024	0.5027 \pm 0.0021	0.5016 \pm 0.0021
0.5	0.9857	0.5559	0.547
0.6	0.6684	0.6164	0.5944
0.8	0.2057	0.7362	0.6714

Table 4.3: Experiment 1: graphs with $n = 5000$ and 5 repetitions

$n = 10000 \quad r = 5$			
Inter-noise	Avg ϵ	Avg k	Avg <i>sze_idx</i>
0.2	0.275	512	0.0948
0.4	0.2525	512	0.1533
0.5	0.2525	512	0.1962
0.6	0.2525	512	0.2449
0.8	0.23	512	0.3607

Inter-noise	Avg ARI Kvs FSZE \pm std	$l_2(FSZE, GT) \pm$ std	$l_1(FSZE, GT) \pm$ std
0.2	0.9993 \pm 0.0003	0.3575 \pm 0.0038	0.3369 \pm 0.0035
0.4	0.9907	0.5007	0.5019
0.5	0.9918	0.558	0.5476
0.6	0.7772	0.6228	0.5942
0.8	0.194	0.7357	0.6714

Table 4.4: Experiment 1: graphs with $n = 10000$ and 5 repetitions

- Every network has been processed with 20 repetitions (just to show that the results does not significantly vary).

The table reports some statistics for each dataset, here is the naming convention:

- k : the cardinality of the best partition, namely $|\mathcal{P}^*|$.
- ϵ : the corresponding epsilon of the best partition found.
- nirr: number of irregular pair of the final partition.
- *sze_idx*: the Szemerédi Index described in Section 2.3.
- tcompress: time to compute the best partition and express it though *RED* (step 1 and 2 of the CoDec Procedure described in Section 3.1). It is expressed in seconds.
- tdecompress: time to reconstruct the matrix and decompress the *RED* matrix to the *SZE* (step 3 of the CoDec Procedure described in Section 3.1). It is expressed in seconds.
- tfiltering: time to apply the median filter on the *SZE* matrix giving *FSZE* (step 4 of the CoDec Procedure described in Section 3.1).

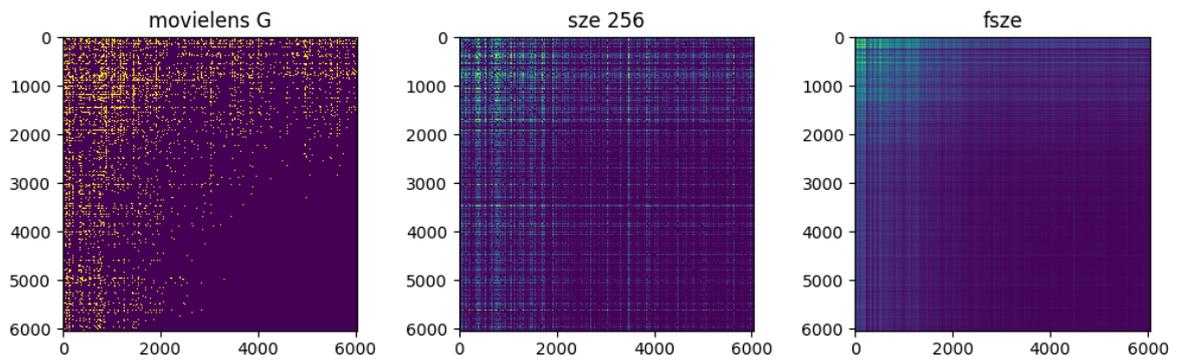


Figure 4.5: MovieLens [19] network.

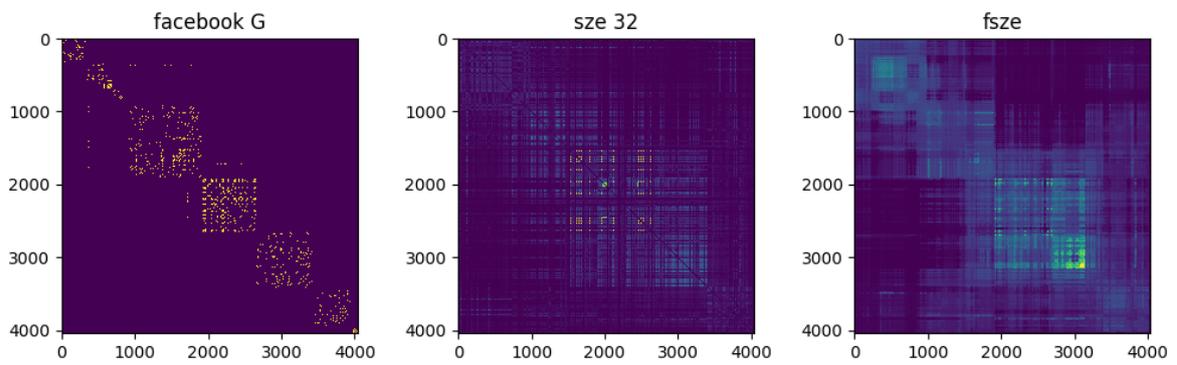


Figure 4.6: Facebook [20] network.

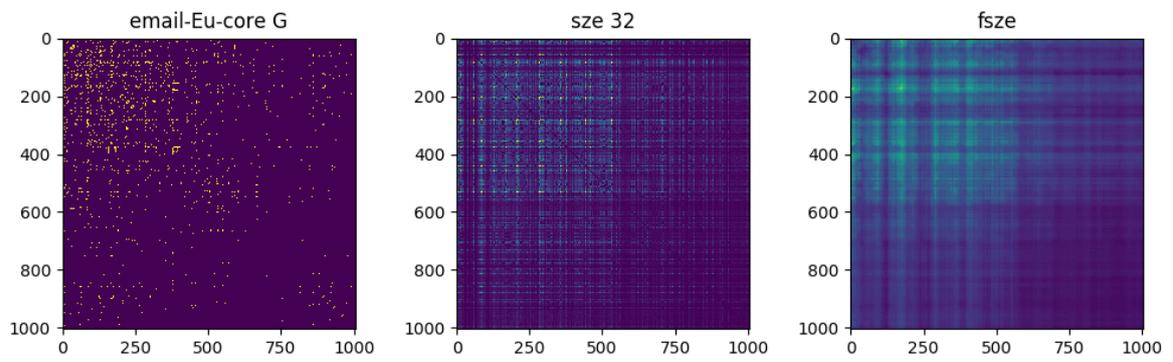


Figure 4.7: Email-Eu-core [21][22] network.

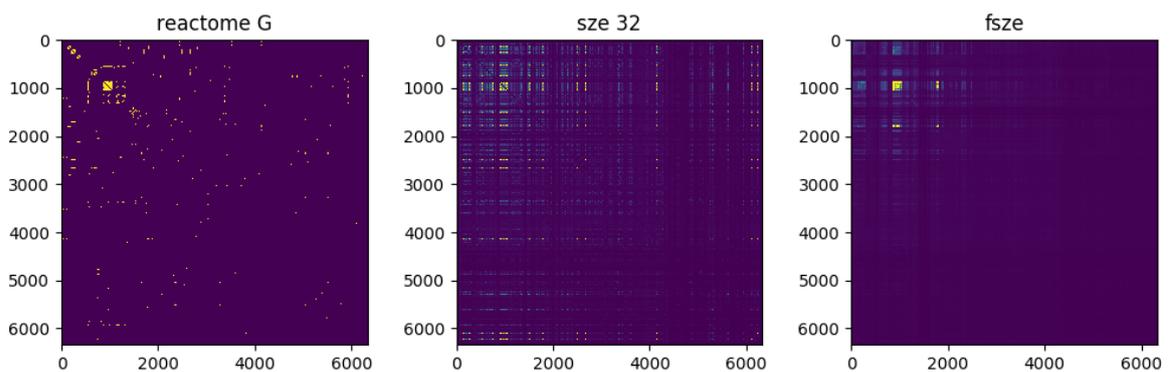


Figure 4.8: Reactome [23] network.

$n = 25000 \ r = 3$			
Inter-noise	Avg ϵ	Avg k	Avg <u>size_idx</u>
0.2	0.32	2048	0.0933
0.4	0.2975	2048	0.1555
0.5	0.275	2048	0.197
0.6	0.275	2048	0.2459
0.8	0.2975	2048	0.3615

Inter-noise	Avg ARI Kvs FSZE \pm std	$l_2(FSZE, GT) \pm$ std	$l_1(FSZE, GT) \pm$ std
0.2	0.998 ± 0.0007	0.3599 ± 0.009	0.3517 ± 0.0062
0.4	0.9995	0.4795	0.481
0.5	0.9885	0.5755	0.55
0.6	0.9482	0.6166	0.5928
0.8	0.6862	0.6758	0.6665

Table 4.5: Experiment 1: graphs with $n = 25900$ and 3 repetitions

Movielens $n = 6041 \ d = 0.0541 \ r = 20$				
	k	ϵ	<u>size_idx</u>	nirr
min	128	0.18200	0.00350	599
avg	217.6	0.23003	0.00386	2457.2
sd	60.2	0.02775	0.00015	1348.7
max	256	0.25000	0.00410	3839
	$l_2(SZE, G)$	$l_2(FSZE, G)$	$l_1(SZE, G)$	$l_1(FSZE, G)$
min	0.218	0.223	0.082	0.071
avg	0.219	0.224	0.084	0.074
sd	0.001	0.001	0.003	0.003
max	0.220	0.225	0.088	0.078
	tcompress/s	tdecompress/s	tfiltering/s	
min	252.33	253.61	473.55	
avg	284.52	286.91	511.04	
sd	23.13	23.37	29.67	
max	321.99	325.05	556.79	

Table 4.6: Movielens real network

4.4 Analysis and considerations

Analysis of Synthetic Graphs Experiments

The whole framework can be seen as a graph compression technique, beside the characteristic that allows us to store large graphs with less requirements, the procedure seems to preserve the underlying structural patterns which hold the information.

As can be seen from all the tables, our compression algorithm based on an approximated version of Alon et al.’s algorithm plus the new devised heuristic, finds best partitions with epsilon value varying between 0.13 to 0.27. We recall that in order to have a faithful approximation of the original graph, we should be able to find partitions with very low values of epsilon. This is very difficult to achieve since one of the theoretical requirements is the astronomical size of the input graph (due to the tower constant on the cardinality of the partition). Of course even with a framework able to handle very big graphs we would still need a perfect implementation of Alon et al.’s algorithm. Said that, and keeping in mind that the epsilon parameter controls the approximation factor, we could say that it is a good result. In fact we are able to process medium/big size graphs and achieve a very good compression with preservation of the structure.

As far as regards the compression property if we take in consideration Table 4.5, we can calculate

Facebook $n = 4039$ $d = 0.0108$ $r = 20$

	k	ϵ	sz idx	nirr
min	32	0.123	0.00010	12
avg	34.90	0.145	0.00015	34.09
sd	9.41	0.017	0.00007	20.52
max	64	0.191	0.00030	100
	$l_2(SZE, G)$	$l_2(FSZE, G)$	$l_1(SZE, G)$	$l_1(FSZE, G)$
min	0.101	0.103	0.018	0.015
avg	0.102	0.103	0.019	0.016
sd	0.001	0.000	0.000	0.001
max	0.103	0.103	0.020	0.017
	tcompress/s	tdecompress/s	tfiltering/s	
min	43.07	43.27	128.22	
avg	46.55	46.78	135.12	
sd	2.23	2.24	4.65	
max	50.35	50.67	145.48	

Table 4.7: Facebook real network

Email-Eu-core $n = 1005$ $d = 0.00257$ $r = 20$

	k	ϵ	sz idx	nirr
min	16	0.131	0.00040	3
avg	30.4	0.219	0.00076	29.1
sd	4.9	0.027	0.00014	10.7
max	32	0.250	0.00100	42
	$l_2(SZE, G)$	$l_2(FSZE, G)$	$l_1(SZE, G)$	$l_1(FSZE, G)$
min	0.156	0.158	0.040	0.035
avg	0.157	0.158	0.041	0.036
sd	0.001	0.000	0.001	0.001
max	0.159	0.158	0.043	0.039
	tcompress/s	tdecompress/s	tfiltering/s	
min	2.47	2.49	8.07	
avg	2.78	2.82	8.50	
sd	0.21	0.21	0.25	
max	3.18	3.22	8.92	

Table 4.8: Email-Eu-core real network

the compression rate with graphs of 25000 nodes. In particular a 25000×25000 matrix has been compressed in a 2048×2048 matrix plus a 25000 vector. If we weight each entry of a matrix and each entry of a vector as one byte (for simplicity) we have:

- Storage requirements for the graph: $(25000 \cdot 25000)/(2 \cdot 1024^2) \approx 298$ MB
- Storage requirements for the compressed graph: $(2048 \cdot 2048)/(2 \cdot 1024^2) = 2$ MB
- Storage requirements for the membership vector: $25000/1024 = 3.05$ KB

This constitutes a compression ratio of 147.25 and we save 99.32% of the space required. This is surely a very good result especially if we want to send the structure through a channel. This should suspect us since a compression rate such high it's uncommon, but this is a lossy compression and when the structure is decompressed it's meant to loose all the information which do not carry structural patterns, namely we only preserve useful informations.

From the visual inspection reported in Figure 4.4 the CoDec Procedure first property is that it preserves the underlying structure of the graph G . This property should be stronger as the graph dimension grows, this point is also verified from the experimental results conducted with Experiment

Reactome $n = 6328$ $d = 0.00073$ $r = 20$

	k	ϵ	size	idx	nirr
min	16	0.097	0.00010		4
avg	28.8	0.123	0.00011		17.05
sd	6.6	0.015	0.00002		7.6
max	32	0.165	0.00020		31
	$l_2(SZE, G)$	$l_2(FSZE, G)$	$l_1(SZE, G)$		$l_1(FSZE, G)$
min	0.082	0.083	0.011		0.009
avg	0.084	0.084	0.012		0.010
sd	0.001	0.001	0.000		0.000
max	0.085	0.085	0.013		0.011
	tcompress/s	tdecompress/s	tfiltering/s		
min	97.84	98.28	285.84		
avg	105.49	105.96	314.12		
sd	3.82	3.84	12.87		
max	110.89	111.38	330.39		

Table 4.9: Reactome real network

1, in particular from the Adjusted Random Index of the KVS-clustering reported in Tables 4.1, 4.2, 4.3, 4.4, 4.5. We can clearly see that we achieve better results as the dimension of the graph increases, in fact with $n = 25000$ we are still able to filter an inter-noise of 0.6 with an Adjusted Random Index of 0.94, while with $n = 1000$ it is not the case. This confirms that as the dimension of the graph grows we are in fact capturing and preserving important information, therefore we can derive the fact that our approximated implementation of Alon et al.’s algorithm plus the new heuristic is indeed behaving according to the original formulation.

Anther point is that Alon et al.’s algorithm does not tell us how to reach a regular partition at each iteration of the algorithm. In other words it does not fully specify the refinement step. Our new refinement heuristic has an underlying motivation which follows the combinatorial nature of the lemma and aims to mitigate the effect of the tower constant. In fact we find partitions with relatively small cardinality in medium/big graphs. The good side of this, is that we are able to compress the structure to very high rates and we are able to process small network too. There is also a bad side: since we have a low cardinality of the best partition it means that we will loose information in the decompression phase. That is because the internal structure of the partition classes is unknown and will be lost in the decompression phase. As we discussed in Section 3.3 we could keep track of the internal densities and later expand them through a random graph with probability of connection of two nodes equal to the internal density, but this has not been extensively tested and worth be further explored.

We then tried to study a possible relation between the density of the input graph and the optimal thresholding value of the $FSZE$ matrix which minimize the dissimilarity with the GT . The study have been reported in Figure 4.9. There seems to be a pattern for high densities, but as it decrease the distribution of the optimal threshold value it’s much wider. Maybe with more data and further investigations we can draw an empirical rule to optimally threshold the output of the CoDec procedure to produce an unweighted graph and to completely filter the noise by preserving only the structural part.

In Experiment 2 Figure 4.3 we tested the CoDec procedure robustness when both intranoise and internoise are added to a graph with balanced clusters as structure. We decided to use every combination of the values of the noise ranging from 0 to 0.3 since it represents the range where we achieve the best results and because of the computational time required. As said, we introduced both noisy edges between the clusters and corruption of the structure of the clusters. We can see that as the number

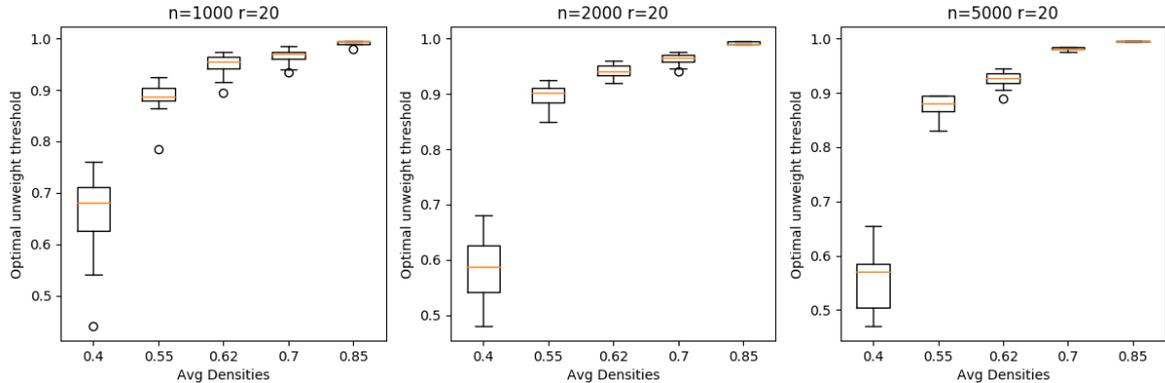


Figure 4.9: Relation between graph density and optimal unweighting threshold, the graphs considered were the graphs used for the Experiment 1.

of clusters increase the preservation is more difficult. This result can be explained by the fact that as the number of cliques increases the density of the original graph decreases, therefore we are getting far from the ideal conditions for the lemma to apply. Another important consideration is that the procedure is not able to preserve structure in total absence of noise, this is interesting since highlights the fact that the framework needs it to work and reinforce the denoisation property of the method.

The same experiment has been conducted with imbalanced clusters, as reported in Figure 4.2, one can immediately see that we have better results with respect to the balanced cluster case. The procedure preserves structure in many more combinations of intra and inter noise. We can see that as the intranoise increases (corruption of structure) we achieve worse preservation that is, again, due to the fact that as we corrupt we are in fact decreasing the density of the graph. It is not the case in the balanced cluster case, maybe because the critical threshold is higher due to the different morphology of the graphs.

Analysis of Real Networks Experiments

Even if we assumed that the procedure needs a structural part of the graph made of cluster-like groups, the algorithm is still able to highlight some of the pattern of the graphs. We can not provide a measure of preservation since we don't have a ground truth, but it can be verified from the visual inspection of the matrices in Figures 4.5, 4.6, 4.7, 4.8. This point is fundamental since it highlights the fact that this procedure is able to preserve the patterns of the graphs. In fact we could see it as if the graph G is just a corroded version of a GT which is well represented by the output of the procedure.

Another important thing to notice is that even if we use graphs of medium size with a really low density, the framework is still able to find partitions and provide a visually appealing reconstruction with low l_2 dissimilarities.

Perhaps the most important result obtained is the one that we got when the CoDec procedure has been applied to the Facebook network. In fact we obtain better results than the algorithm proposed by Riondato et al. [24] which, according to this survey [1], is one of the algorithm at the state-of-the-art. The result reported by Riondato et al. is a l_2 dissimilarity measure which is slightly different from our measure since it is not normalized by n^2 but instead by n . In order to compare the results we just need to multiply our dissimilarity value by the size of the graph, namely $n = 4039$. The results of Riondato et al. have been reported in Figure 4.10 in particular we achieve an $l_2 = 0.103 \cdot 4039 = 416.017$ while their best average is 501 computed with 5 different runs. In order to fully compare the two techniques we should use the same graphs, unfortunately Riondato et al. tested their method with very big

Graph	k	ℓ_2 -rec.
		err. ($\times 10^2$)
Facebook	500	6.56
$ V = 4\,039$	750	6.18
$ E = 88\,234$	1000	5.81
	1250	5.42
	1500	5.01

Figure 4.10: Riondato et al. l_2 reconstruction error for the Facebook network. We achieve 415.61.

graphs that our developed framework can not yet handle. A workaround to test the two procedures would be to perform a sampling of the bigger graphs, but it requires an extensive study to select the best methodology to preserve the structural properties of the sampled graphs.

Another important fact is that the algorithms of Riondato et al. mainly use graphs with clusters, and they point out that there is a relation between a weakly-regular partitions and the minimization of the CutNorm Error measure that they defined. In particular each class of the partition represent a cluster in the original graph. Our technique abstract this notion and in fact it is able also to highlight different patterns as we already discussed.

Said that, we could strongly argue if the l_2 measure is a good proxy to measure how much structure has been preserved. In fact it is easy to see that as the graphs get more and more sparse we are closer to an empty graph which, unless our reconstruction adds a lot of noise, always gives low l_2 values. In Section 5.1 we discuss some other measures.

Chapter 5

Conclusions

In this thesis it has been devised a new pipeline for compressing and decompressing a graph by exploiting powerful Szemerédi’s Regularity Lemma. In particular, it has been developed a procedure called CoDec (Compression-Decompression) that uses an approximated algorithmic formulation of the strong Regularity Lemma given by Alon et al. as compression step, then after the decompression phase (which simply exploits the Embedding Lemma) it introduces a post-decompression step to better highlight the structural properties of the reconstructed graph.

We provided a new heuristic for the refinement step which exploits the internal structures of the classes of a regular partition. As the experiments show, the new heuristic mitigates the problematic tower constant on the number of classes which should compose a regular partition. In fact when prompted with sufficiently dense graphs, the framework always find regular partitions of different cardinalities.

We then provided an extensive experimental evaluation to measure how robust is the framework as we both corrupt the structures of the synthetic graphs, which carries the information, and add noisy edges between them. As we reported, the procedure is less robust to the corruption of the structure since it implies a decrement in the density of the graph. This is coherent to the requirements to the Regularity Lemma which finds its ideal usage in dense graphs. In fact, CoDec achieves good structural preservation properties when we add noise between the structures.

We also tested the pipeline as the graphs dimension increase. Turns out that our implementation seems to confirm the theoretical requirements needed by the strong Regularity Lemma, so by processing bigger graphs we achieve better results.

5.1 Future Works

In this Section we report some ideas that can be further explored in the future. In particular, we could study how the results of the procedure change when the threshold of the Sparsification and Densification branch of the new heuristic vary. This could be another experiment to better understand the properties of the new heuristic and could possibly open new insights.

Surely we should also perform several tests on weighted graphs. In particular, we could study if there is any relation between weighted internoise and intranoise of the input graph and the CoDec structural preservation ability. The introduction of weighted noise and the introduction of weighted structures could open new aspects of the algorithm both in positive and negative.

Another very important aspect would be to use different technologies to implement the framework. The whole pipeline has been coded using Python 3.6 (with heavy usage of the Numpy library). We have some limitations on the dimension of the graphs that we can process due to the choice of using

matrices as data structure representation for the graphs. We could use different data structures and we could of course make use of some dedicated libraries like [GraphTools](#) or [NetworkX](#) which are simply born for such applications. Another option would be to port the framework into another language like C++ to achieve a speedup on the running time, but the migration would not be immediate due to the complexity of the project and target language. A big improvement would be to redefine the framework in a distributed environment, this surely open new possibility to process bigger graphs and to explore different refinements.

As we said in the analysis of the result, the l_2 metric does not provide a good evaluation to compare the similarity of two graphs, this is always the case when we are evaluating techniques that offers an approximation of the graph. In particular, we would like a measure of how similar are the two structural part of the graphs considered. An idea would be to use an indirect measure based on a feature vector of some statistics of the graphs, like: average local clustering coefficient from Watts et al. [25], global clustering coefficient from Luce et al. [26], average degree and so on, and then measure the distance between the two feature vectors. This is an indirect measure of similarity, the problem of computing such feature vector is that we need more computational time and it may be a problem when dealing with huge graphs.

There is also the possibility to use the spectral norm that is defined as the square root of the maximum eigenvalue of the squared adjacency matrix. Since it combines the squared adjacency matrix (which counts the number of paths of length two), and eigenvalues (which are highly connected with spectral clustering), it could be a good proxy to evaluate the structural similarity of two (cluster like) graphs.

In this work we used the notion of strong regularity to compress an input graph while preserving some structural properties. However, we can relax the notion to produce weak regular partition exploiting Frieze and Kannan's [13] weak version of the Regularity Lemma. There would be some benefits to move from a strong to a weak regularity notion mainly because we are no longer required to work with very big graphs and because there are many papers continuously improving the tower bounds and providing new approximation algorithms like the last by Fox et al. [27].

Bibliography

- [1] Y. Liu, A. Dighe, T. Safavi, and D. Koutra, “Graph Summarization: A Survey,” *ACM Computing Surveys*, vol. V, no. N, 2017.
- [2] A. Sperotto and M. Pelillo, “Szemerédi’s Regularity Lemma and Its Applications to Pairwise Clustering and Segmentation,” in *Energy Minimization Methods in Computer Vision and Pattern Recognition* (A. L. Yuille, S.-C. Zhu, D. Cremers, and Y. Wang, eds.), (Berlin, Heidelberg), pp. 13–27, Springer Berlin Heidelberg, 2007.
- [3] N. Alon, R. A. Duke, H. Lefmann, V. Rödl, and R. Yuster, “The Algorithmic Aspects of the Regularity Lemma,” *J. Algorithms*, vol. 16, pp. 80–109, 1 1994.
- [4] M. Pelillo, I. Elezi, and M. Fiorucci, “Revealing Structure in Large Graphs,” *Pattern Recogn. Lett.*, vol. 87, pp. 4–11, 2 2017.
- [5] M. Fiorucci, A. Torcinovich, M. Curado, F. Escolano, and M. Pelillo, “On the Interplay Between Strong Regularity and Graph Densification,” in *Graph-Based Representations in Pattern Recognition* (P. Foggia, C.-L. Liu, and M. Vento, eds.), (Cham), pp. 165–174, Springer International Publishing, 2017.
- [6] W. T. Gowers, “Lower bounds of tower type for Szemerédi’s uniformity lemma,” *Geometric & Functional Analysis GAFA*, vol. 7, pp. 322–337, 5 1997.
- [7] E. Szemerédi, “On sets of integers containing no four elements in arithmetic progression,” *Acta Mathematica Academiae Scientiarum Hungaricae*, vol. 20, no. 1-2, pp. 89–104, 1969.
- [8] H. Furstenberg, Y. Katznelson, and D. Ornstein, “The Ergodic theoretical proof of Szemerédi’s theorem,” *Bulletin of the American Mathematical Society*, vol. 7, no. 3, pp. 527–552, 1982.
- [9] W. T. Gowers, “A New Proof of Szemerédi’s Theorem for Arithmetic Progressions of Length Four,” *Geometric & Functional Analysis GAFA*, vol. 8, pp. 529–551, 7 1998.
- [10] J. Komlos and M. Simonovits, “Szemerédi’s Regularity Lemma and its applications in graph theory,” *DIMACS Technical Report*, 1996.
- [11] J. Komlós, A. Shokoufandeh, M. Simonovits, and E. Szemerédi, “The Regularity Lemma and Its Applications in Graph Theory,” *Theoretical aspects of computer science (Tehran, 2000)*, vol. 2292, pp. 84–112, 2002.
- [12] E. Szemerédi, “Regular Partitions of Graphs,” tech. rep., Stanford, CA, USA, 1975.
- [13] A. Frieze and R. Kannan, “A simple algorithm for constructing Szemerédi’s Regularity Partition,” 1999.

- [14] K. O. Neil and S. L. Peters, “Experimental Improvements to Regularity Clustering,” tech. rep., 2014.
- [15] R. Szeliski, “Computer Vision : Algorithms and Applications,” *Computer*, vol. 5, p. 832, 2010.
- [16] L. Hubert and P. Arabie, “Comparing partitions,” *Journal of Classification*, vol. 2, pp. 193–218, 12 1985.
- [17] J. Leskovec and A. Krevl, “SNAP Datasets: Stanford large network dataset collection.” <http://snap.stanford.edu/data>, June 2014.
- [18] J. Kunegis, “KONECT – The Koblenz Network Collection,” in *Proc. Int. Conf. on World Wide Web Companion*, pp. 1343–1350, 2013.
- [19] GroupLens Research, “MovieLens data sets.” <http://www.grouplens.org/node/73>, October 2006.
- [20] J. Leskovec and J. J. Mcauley, “Learning to Discover Social Circles in Ego Networks,” in *Advances in Neural Information Processing Systems 25* (F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, eds.), pp. 539–547, Curran Associates, Inc., 2012.
- [21] H. Yin, A. R. Benson, J. Leskovec, and D. F. Gleich, “Local Higher-Order Graph Clustering,” in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining - KDD '17*, pp. 555–564, 2017.
- [22] J. Leskovec, J. Kleinberg, and C. Faloutsos, “Graph Evolution: Densification and Shrinking Diameters,” *ACM Trans. Knowl. Discov. Data*, vol. 1, 3 2007.
- [23] G. Joshi-Tope, M. Gillespie, I. Vastrik, P. D’Eustachio, E. Schmidt, B. de Bono, B. Jassal, G. R. Gopinath, G. R. Wu, L. Matthews, S. Lewis, E. Birney, and L. Stein, “Reactome: A knowledgebase of biological pathways,” *Nucleic Acids Research*, vol. 33, no. DATABASE ISS., 2005.
- [24] M. Riionato, D. Garcia-Soriano, and F. Bonchi, “Graph Summarization with Quality Guarantees,” *Data Min. Knowl. Discov.*, vol. 31, pp. 314–349, 3 2017.
- [25] D. J. Watts and S. H. Strogatz, “Collective dynamics of ‘small-world’ networks,” *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.
- [26] R. D. Luce and A. D. Perry, “A method of matrix analysis of group structure,” *Psychometrika*, vol. 14, no. 2, pp. 95–116, 1949.
- [27] J. Fox, L. M. Lovász, and Y. Zhao, “A fast new algorithm for weak graph regularity,” pp. 1–13, 2018.