THE REGULARITY LEMMA AND GRAPH THEORY

ARIEL HAFFTKA

Abstract. The Szemerédi Regularity Lemma states that any sufficiently large graph G can be partitioned into a bounded (independent of the size of the graph) number of regular, or "random-looking," components. The resulting partition can be viewed as a regularity graph R. The Key Lemma shows that under certain conditions, the existence of a subgraph H in R implies its existence in G. We prove the Regularity Lemma and the Key Lemma.

Contents

1. Introduction

In 1975 Endre Szemerédi proved a 40 year old Ramsey Theory type conjecture of Erdös and Turán: every set of natural numbers of positive upper density contains arithmetic progressions of arbitrary length. Szemerédi began his proof with a "lemma on bipartite graphs" [1]. In 1976, he generalized this result to simple graphs [2], a result which became known as the Regularity Lemma. The Regularity Lemma states that any graph can almost entirely be partitioned into a bounded number of components, such that the edge density between any two of these components is nearly random, or ϵ-regular, as will be defined in Section 2. Notably, for a given lower bound m and ϵ, the upper bound $M = M(m, \epsilon)$ given by the Regularity Lemma depends only on m and ϵ and holds for arbitrarily large graphs.

One application of the Regularity Lemma is in proving that certain conditions on a large graph G (such high enough edge density) can force the existence of a given subgraph H. The partition resulting from the Regularity Lemma gives rise to a regularity graph R whose vertices are the components of the partition and where two components are connected by an edge if and only if the edge density between the two components in G is sufficiently high. The Key Lemma states that under the right conditions, the existence of a bounded degree subgraph H in a certain generalized regularity graph (see Section 4) implies its existence in G.

Date: August 22, 2008.

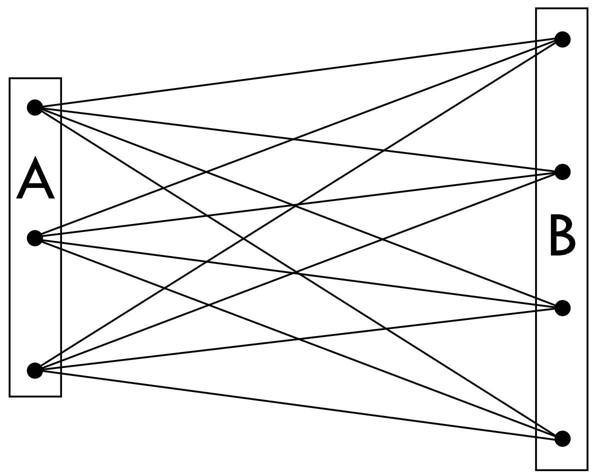

FIGURE 1. A complete $(3,4)$-bipartite graph.

In this paper we prove the Regularity Lemma and the Key Lemma. We shall draw main ideas, lemma statements, and some notations from [3]. The illustrations, examples, and most of the explanations are original.

2. Definitions, Examples, and Statement of the Regularity Lemma

Let $G = (V, E)$ be a simple graph on n vertices, i.e., a connected, undirected graph with no loops.

Definition 2.1. (Edge density) For disjoint $A, B \subset V$, let

$$d(A,B) = \frac{e(A,B)}{|A||B|},$$

where $e(A,B)$ is the number of edges with a vertex in A and a vertex in B.

First observe some obvious properties of edge density. The quantity $d(A,B)$ is the ratio of the number of edges between A and B to the maximum possible number of such edges. Thus if the subgraph on A and B is a complete bipartite graph, the density is 1 (as in Figure 1). If no edges connect A and B, the density is 0. Edge density is always a rational number in $[0,1]$. Here are two slightly less trivial examples: In Figure 2, we have $d(A,B) = 15/16$ and $d(C,D) = 5/8$.

We shall now introduce ϵ-regular pairs. Consider two disjoint vertex sets, A and B. Intuitively, the pair of sets $\{A, B\}$ is said to be ϵ-regular if the edges between A and B are distributed fairly homogeneously.

Definition 2.2. (ϵ-Regular Pair) Let $\epsilon > 0$. Finite nonempty disjoint sets $A, B \subset V$ are said to be ϵ-regular if for all $A' \subset A$ and $B' \subset B$, we have

$$|A'| \geq \epsilon|A| \text{ and } |B'| \geq \epsilon|B| \implies |d(A,B) - d(A',B')| \leq \epsilon.$$

First observe that any $\{A, B\}$ is trivially ϵ-regular for all $\epsilon \geq 1$, since it is impossible for $|d(A,B) - d(A',B')|$ to be greater than 1. Thus ϵ-regularity is only meaningful for $\epsilon < 1$.

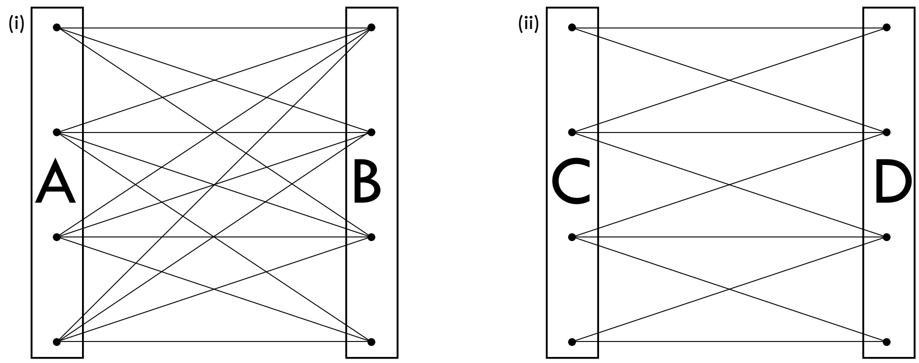

FIGURE 2. (*i*) The pair $\{A, B\}$ is (1/3)-regular. (*ii*) The pair $\{C, D\}$ is not (1/3)-regular.

The justification for the above definition is by comparison to large random graphs. Suppose, for example, that A and B are countably infinite. Let $p \in [0, 1]$ and for each $a \in A$ and $b \in B$ let $\{a, b\}$ be an edge with probability p. Then for large finite subsets $A' \subset A$ and $B' \subset B$ (taking the limit as $|A'|, |B'| \longrightarrow \infty$), the expected number of edges would be $p|A'||B'|$. So the expected density would be $\frac{p|A'||B'|}{|A'||B'|} = p$. But of course, if we were to pick out a single vertex $a \in A$ and a single vertex $b \in B$, then $d(\{a\}, \{b\}) = 0$ or 1. Thus we can only be reasonably certain that the edge density will be very close to p for certain large subsets of A and B. We use ϵ to specify an upper bound for the size of sets which do not necessarily behave randomly.

Here are some examples of ϵ-regular pairs. The pair $\{A, B\}$ in Figure 1 is ϵ-regular for all ϵ, since $d(A, B) = 1$ and since for any subsets $A' \subset A$ and $B' \subset B$, we also have $d(A', B') = 1$. More generally, if the subgraph induced by A and B is a complete bipartite graph, then $\{A, B\}$ is ϵ-regular for all ϵ. Any pair $\{\{a\}, \{b\}\}$ of singletons is ϵ-regular for all ϵ, since the only nonempty subset of a singleton is itself.

In Figure 2, the pair $\{A, B\}$ is (1/3)-regular while the pair $\{C, D\}$ is not. Let us see why. We already observed that $d(A, B) = 15/16$ and $d(C, D) = 5/8$. To check (1/3)-regularity, we need to look at subsets of size at least $\lceil (1/3)4 \rceil = 2$. First we show that $\{A, B\}$ is (1/3)-regular. Suppose not. there are subsets $A' \subset A$ and $B' \subset B$ of size at least 2 with $|d(A', B') - 15/16| > 1/3$. Thus $d(A', B') < 29/48 < 3/4$. But since the bipartite graph between A and B is complete except for one edge and since $|A'| \geq 2$ and $|B'| \geq 2$, we also have $d(A', B') \geq 3/4$, a contradiction.

Now we show that $\{C, D\}$ is not (1/3)-regular. Let C' be the two vertices on the top left and D' be the two vertices on the bottom right, in the right side of Figure 2. Then $d(C', D') = 1/4$, so $|d(C', D') - d(C, D)| = |1/4 - 5/8| = 3/8 > 1/3$, which contradicts (1/3)-regularity.

Now we introduce the partitions we will use. In fact, usually we will not use partitions, but rather collections of disjoint sets covering most, but not all, vertices of a graph. The phrase "collection of disjoint sets" will be used with the implied meaning of "almost partition", or a partition of all of a graph except for a few

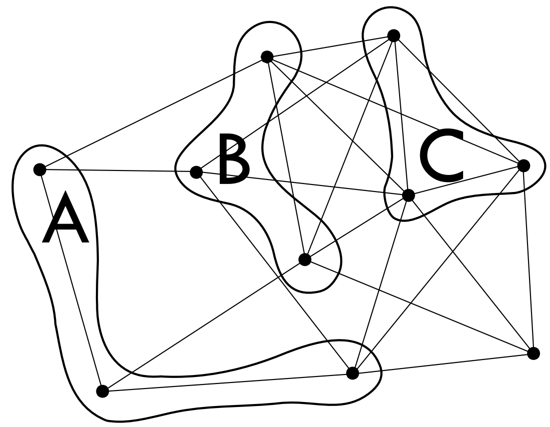

FIGURE 3. A (2/3)-regular collection.

leftover vertices. We will only use the word "partition" when its precise meaning is intended, that is to refer to a collection of disjoint sets which cover a set.

Given a large dense graph, we would like to group its vertices into components, in such a way that "most" choices of two components form ϵ-regular pairs and such that all components are the same size. We also would like to group most of the vertices, although not necessarily all, into these components.

Definition 2.3. (ϵ-Regular Collection) A collection of pairwise disjoint vertex sets $\mathcal{A} = \{A_1, \dots, A_k\}$ is said to be ϵ-regular if

(1) (regular pairs) At most ϵk^2 pairs $\{A_i, A_j\}$ are not ϵ-regular, for $1 \leq i < j \leq k$.
(2) (equal size) $|A_1| = \cdots = |A_k|$.
(3) (almost a partition) $|V \setminus \bigcup_{A \in \mathcal{A}}| \leq \epsilon |V|$.

An example of a (2/3)-regular collection is Figure 3. It is easy to check condition (1) of (2/3)-regularity: that the pairs $\{A, B\}, \{B, C\}$, and $\{A, C\}$ are (2/3)-regular. As an example, we will prove that $\{A, C\}$ is $(2, 3)$-regular. First observe that $d(A, C) = 2/9$. To check ϵ-regularity, we must show that for any subsets $A' \subset A$ and $C' \subset C$ of size at least $(2/3)(3) = 2$, we have $|d(A', C') - 2/9| \leq 2/3$. Suppose this were false. Then there exist A' and C' such that $d(A', C') > 2/3 + 2/9 = 8/9$. But since $d(A', B')$ is of the form $i/9$ for some i, we have $d(A', C') = 1$, which implies the graph between A' and C' is a complete $(2, 2)$ bipartite graph. This contradicts the picture.

Now we check the two other conditions of (2/3)-regularity. Condition (2) is satisfied because $|A| = |B| = |C|$, while condition (3) is satisfied because only 1 vertex is excluded from the collection, and $1 \leq 6 = (2/3)3^2 = \epsilon k^2$.

Since every pair of singletons is ϵ-regular, every graph on n vertices trivially admits an ϵ-regular collection of n disjoint singletons. Thus the importance of the Regularity Lemma is in the claim that we can put an upper (and lower) bound on the number of sets needed to form an ϵ-regular collection, independent of how large a graph is.

Theorem 2.4. *(Regularity Lemma) For every $\epsilon > 0$ and $m \in \mathbb{N}$, there exists an integer $M = M(\epsilon, m)$ such that every graph G on m or more vertices admits an ϵ-regular collection $\mathcal{A} = \{A_1, \ldots, A_k\}$ with $m \le k \le M$.*

3. Proof of the Regularity Lemma

Our proof will follow [3]. The main idea is as follows. We define an index τ of a pair of disjoint vertex sets A and B. We extend this definition so τ becomes a function from the set $\{$collections of disjoint subsets of $V\} \longrightarrow \mathbb{R}$. In Lemma 3.2 we show that τ is bounded above. Then we describe an algorithm that takes an ϵ-irregular collection and produces a refinement with index increased by the positive constant $\epsilon^5(1-\epsilon)^2$ (assuming $\epsilon < 1$) depending only on ϵ. Thus by the boundedness of τ, the algorithm will eventually result in an ϵ-regular collection.

The algorithm involves two steps. In Lemma 3.6 we show that if a pair of disjoint sets A, B is not ϵ-regular, then there exist partitions of $\mathcal{A}$ of A and $\mathcal{B}$ of B such that the index of $\{\mathcal{A}, \mathcal{B}\}$ is greater than the index of $\{A, B\}$ by at least $\frac{\epsilon^4|A||B|}{|V|^2}$. In Lemma 3.11, we use the previous result to show that if a collection $\mathcal{A}$ of disjoint vertex sets is not ϵ-regular, then a refinement exists with index increased by at least the nonzero constant $\epsilon^5(1-\epsilon)^2$, depending only on ϵ.

We begin by defining the index τ, first for a pair of disjoint vertex sets; then for a finite collection $\{A_i\}_{i\in I}$ of disjoint vertex sets; then for a pair of collections $\{A_i\}_{i\in I}$ of disjoint sets in A and $\{B_j\}_{j\in J}$ of disjoint sets in B; and finally we define an extended version τ^* which will only be used in Lemma 3.11.

Definition 3.1. (Index τ)

(1) For disjoint $A, B \subset V$, let
$$\tau(A, B) = \frac{e(A, B)^2}{|A||B||V|^2}.$$

(2) For a collection $\mathcal{A} = \{A_i\}_{i\in I}$ of pairwise disjoint vertex sets, let
$$\tau(\mathcal{A}) = \sum_{\{i,j\}\subset I} \tau(A_i, A_j).$$

(3) Let A and B be disjoint vertex sets. For collections $\mathcal{A} = \{A_i\}_{i\in I}$ of pairwise disjoint subsets of A and $\mathcal{B} = \{B_j\}_{j\in J}$ of pairwise disjoint subsets of B, let
$$\tau(\mathcal{A}, \mathcal{B}) = \sum_{i\in I, j\in J} \tau(A_i, B_j).$$

(4) Furthermore, let
$$\tau^*(\mathcal{A}) = \tau\left(\mathcal{A} \cup \left\{\{x\} \middle| x \in V \setminus \bigcup_{A\in\mathcal{A}} A\right\}\right).$$

We now show that the index τ is bounded above.

Lemma 3.2. *Let $\mathcal{A} = \{A_i\}_{i\in I}$ be a collection of pairwise disjoint vertex sets. Then $\tau(\mathcal{A}) \le 1$.*

Proof.

$$\begin{aligned}
\tau(\mathcal{A}) &= \sum_{\{i,j\}\subset I} \tau(A_i, A_j) \\
&= \sum_{\{i,j\}\subset I} \frac{d(A_i, A_j)^2 |A_i||A_j|}{|V|^2} \\
&\leq \frac{1}{|V|^2} \sum_{\{i,j\}\subset I} |A_i||A_j| \\
&\leq \frac{1}{|V|^2} \left(\sum_{i\in I} |A_i| \right) \left(\sum_{j\in J} |A_j| \right) \\
&\leq \frac{1}{|V|^2} |V||V| \\
&= 1.
\end{aligned}$$

□

The Cauchy-Schwarz inequality

$$\sum_{i=1}^{n} \frac{\alpha_i^2}{\beta_i} \geq \frac{(\sum_i^n \alpha_i)^2}{\sum_{i=1}^n \beta_i}. \tag{3.3}$$

for $\alpha_i \in \mathbb{R}$ and $\beta_i \in \mathbb{R}^+$, will be used in Lemmas 3.4 and 3.6.

Lemma 3.4. *(τ does not decrease when taking refinements.) Let $A, B \subset V$ be disjoint, $\mathcal{A} = \{A_i\}_{i\in I}$ be a partition of A, and $\mathcal{B} = \{B_j\}_{j\in J}$ be a partition of B. Then*

$$\tau(\mathcal{A}, \mathcal{B}) \geq \tau(A, B).$$

Furthermore, if $\mathcal{A}'$ is a partition of A that refines $\mathcal{A}$, then

$$\tau(\mathcal{A}') \geq \tau(\mathcal{A}).$$

Proof. (1) First we show that $\tau(\mathcal{A}, \mathcal{B}) \geq \tau(A, B)$.

$$\begin{aligned}
\tau(\mathcal{A},\mathcal{B}) &= \sum_{i\in I, j\in J} \tau(A_i, B_j) \\
&= \sum_{i\in I, j\in J} \frac{e(A_i,B_j)^2}{|A_i||B_j||V|^2} \\
&\geq \frac{\left(\sum_{i\in I, j\in J} e(A_i,B_j)\right)^2}{|V|^2 \sum_{i\in I, j\in J} |A_i||B_j|}, \quad \text{(by Cauchy-Schwarz 3.3)}. \\
&= \frac{e(A,B)^2}{|V|^2(\sum_{i\in I}|A_i|)(\sum_{j\in J}|B_j|)}, \quad \left(\text{since } \sum_{i\in I, j\in J} e(A_i,B_j) \text{ counts all the} \right. \\
&\quad \left. \text{edges from } A \text{ to } B\right). \qquad (3.5) \\
&= \frac{e(A,B)^2}{|V|^2|A||B|} \\
&= \tau(A,B).
\end{aligned}$$

(2) The second claim follows from the first. For each $X \in \mathcal{A}$, let

$$\mathcal{A}_X = \{A' \in \mathcal{A}' | A' \subset X\},$$

so that $\mathcal{A}_X$ is the partition of X induced by $\mathcal{A}'$. Then

$$\begin{aligned}
\tau(\mathcal{A}') &= \sum_{\{A',A''\}\subset\mathcal{A}'} \tau(A',A'') \\
&\geq \sum_{\{X,X'\}\subset\mathcal{A}} \; \sum_{A'\in\mathcal{A}_X, A''\in\mathcal{A}_{X'}} \tau(A',A'') \quad \text{(since } \mathcal{A}' \text{ refines } \mathcal{A}). \\
&= \sum_{\{X,X'\}\subset\mathcal{A}} \tau(\mathcal{A}_X,\mathcal{A}_{X'}) \\
&\geq \sum_{\{X,X'\}\subset\mathcal{A}} \tau(X,X'), \quad \text{(by part 1)}. \\
&= \tau(\mathcal{A}).
\end{aligned}$$

□

The next lemma is the first of two crucial parts of our proof. We look at a single pair of disjoint vertex sets $\{A,B\}$ and show that if it is not ϵ-regular, then there exist partitions of A and of B, each into two sets, with index increased by at least $\frac{\epsilon^4|A||B|}{|V|^2}$. The main idea is the following: Since $\{A,B\}$ is not ϵ-regular, there exist subsets of A and of B that contradict the definition of ϵ-regularity. We use these sets to construct partitions of A and B. The rest of the proof involves computation and an application of Cauchy-Schwarz to prove the inequality.

Lemma 3.6. *Let $A, B \subset V$ be disjoint. If the pair $\{A,B\}$ is not ϵ-regular, there exist partitions $\mathcal{A}$ of A and $\mathcal{B}$ of B such that $|\mathcal{A}| \leq 2$, $|\mathcal{B}| \leq 2$, and*

$$\tau(\mathcal{A},\mathcal{B}) \geq \tau(A,B) + \frac{\epsilon^4|A||B|}{|V|^2}.$$

Proof. 1. *Constructing the partitions.*

Since $\{A, B\}$ is not ϵ-regular, there exist $A_1 \subset A$ and $B_1 \subset B$ such that $|A_1| \geq \epsilon|A|, |B_1| \geq \epsilon|B|$, and

$$|d(A_1, B_1) - d(A, B)| \geq \epsilon. \tag{3.7}$$

Let $A_2 = A \backslash A_1$ and $B_2 = B \backslash B_1$, so that $\mathcal{A} := \{A_1, A_2\}$ is a partition of A and $\mathcal{B} := \{B_1, B_2\}$ is a partition of B. At least one of A_2 and B_2 is nonempty, since otherwise (3.7) would be false. Without loss of generality, assume A_2 is nonempty. Then we can write

$$\mathcal{A} = \{A_i\}_{i \in I}; \qquad \mathcal{B} = \{B_j\}_{j \in J};$$

where $I = \{1, 2\}$ and $J = \{1, 2\}$ if B_2 is nonempty and $J = \{1\}$ if B_2 is empty.

2. *Proving the inequality.*

We claim that $\tau(\mathcal{A}, \mathcal{B})$ is sufficiently larger than $\tau(A, B)$. For brevity, let

$$\begin{aligned}
\lambda :=& d(A_1, B_1) - d(A, B);\\
a :=& |A|;\ b := |B|,\ e := e(A, B);\\
a_i :=& |A_i|;\ b_j := |B_j|;\ e_{ij} = e(A_i, B_j);\\
&\text{with } b_2 = 0 = e_{i2} \text{ if } B_2 \text{ is empty.}
\end{aligned}$$

We are borrowing some notation from [3]. The goal is to express the difference between $\tau(\mathcal{A}, \mathcal{B})$ and $\tau(A, B)$ in terms of λ, and then the claim will follow from $|\lambda| \geq \epsilon$.

Since by definition, $\lambda = \frac{e_{11}}{a_1 b_1} - \frac{e}{ab}$, we have $e_{11} = \lambda a_1 b_1 + \frac{e a_1 b_1}{ab}$. Thus we have two simple identities.

First,

$$\begin{aligned}
e_{11}^2 &= \left(\lambda a_1 b_1 + \frac{e a_1 b_1}{ab}\right)^2\\
&= \lambda^2 a_1^2 b_1^2 + \frac{2\lambda e a_1^2 b_1^2}{ab} + \frac{e^2 a_1^2 b_1^2}{a^2 b^2}.
\end{aligned} \tag{3.8}$$

Second,

$$\begin{aligned}
(e - e_{11})^2 &= \left(e\left(\frac{ab - a_1 b_1}{ab}\right) - \lambda a_1 b_1\right)^2\\
&= e^2\left(\frac{ab - a_1 b_1}{ab}\right)^2 - 2\lambda a_1 b_1 e\left(\frac{ab - a_1 b_1}{ab}\right) + \lambda^2 a_1^2 b_1^2.
\end{aligned} \tag{3.9}$$

Putting it all together,

$$\begin{aligned}
|V|^2\tau(\mathcal{A},\mathcal{B}) =& |V|^2 \sum_{i\in I, j\in J} \tau(A_i, B_j) \\
=& \sum_{i\in I, j\in J} \frac{e_{ij}^2}{a_i b_j} \\
=& \frac{e_{11}^2}{a_1 b_1} + \sum_{i\in I, j\in J, (i,j)\neq(1,1)} \frac{e_{ij}^2}{a_i b_j} \\
\geq& \frac{e_{11}^2}{a_1 b_1} + \frac{\left(\sum_{i\in I, j\in J, (i,j)\neq(1,1)} e_{ij}\right)^2}{\sum_{i\in I, j\in J, (i,j)\neq(1,1)} a_i b_j}, \quad \text{(by Cauchy-Schwarz 3.3)}. \\
=& \frac{e_{11}^2}{a_1 b_1} + \frac{(e-e_{11})^2}{ab - a_1 b_1}, \quad \left(\text{since } a = \sum_{i\in I} a_i,\ b = \sum_{j\in J} b_j, \text{ and } e = \sum_{i\in I, j\in J} e_{ij}\right). \\
=& \frac{\lambda^2 a_1^2 b_1^2 + \frac{2\lambda e a_1^2 b_1^2}{ab} + \frac{e^2 a_1^2 b_1^2}{a^2 b^2}}{a_1 b_1} + \frac{e^2\left(\frac{ab - a_1 b_1}{ab}\right)^2 - 2\lambda a_1 b_1 e \left(\frac{ab - a_1 b_1}{ab}\right) + \lambda^2 a_1^2 b_1^2}{ab - a_1 b_1}, \\
& \text{(by 3.8 and 3.9)}. \\
=& \left(\lambda^2 a_1 b_1 + \frac{2\lambda e a_1 b_1}{ab} + \frac{e^2 a_1 b_1}{a^2 b^2}\right) + \frac{e^2(ab - a_1 b_1)}{a^2 b^2} - \frac{2\lambda a_1 b_1 e}{ab} + \frac{\lambda^2 a_1^2 b_1^2}{ab - a_1 b_1} \\
=& \lambda^2 a_1 b_1 + \frac{e^2 ab}{a^2 b^2} + \frac{\lambda^2 a_1^2 b_1^2}{ab - a_1 b_1} \\
\geq& \epsilon^4 ab + \frac{e^2}{ab}, \quad \Big(\text{since } |\lambda| \geq \epsilon;\ a_1 \geq \epsilon a;\ b_1 \geq \epsilon b; \text{ and } a_1 < a \text{ [since } A_2 \text{ is} \\
& \text{nonempty]}\Big).
\end{aligned}$$

Thus $\tau(\mathcal{A},\mathcal{B}) \geq \tau(A,B) + \frac{\epsilon^4 |A||B|}{|V|^2}$. □

The next lemma is the second crucial part of our proof. In it we show that if a collection $\mathcal{P} = \{A_1, \dots, A_k\}$ of pairwise disjoint, equal-sized vertex sets covering most of V is not ϵ-regular, then there exists a refinement $\mathcal{P}'$ of $\mathcal{P}$, covering almost as much as $\mathcal{P}$, with index increased by a positive constant depending only on ϵ. We use the following broad definition of refinement.

Definition 3.10. (Refinement) Let $\mathcal{P}$ and $\mathcal{P}'$ be collections of pairwise disjoint subsets of V. We say that $\mathcal{P}'$ refines $\mathcal{P}$ if each set in $\mathcal{P}'$ is a subset of some set in $\mathcal{P}$.

Trivially, refinement is a partial ordering on the set {collections of disjoint vertex sets}, with $\mathcal{P}' \leq \mathcal{P}$ if $\mathcal{P}'$ refines $\mathcal{P}$. In the proof of the following, we will use the term *maximal* with respect to this partial ordering and a given condition.

Lemma 3.11. *Let* $0 < \epsilon < 1$; $\mathcal{P} := \{A_1, \dots, A_k\}$ *be a collection of pairwise disjoint subsets of* V; $|A_1| = \cdots = |A_k| := c$; *and* $|V \setminus \bigcup_{i=1}^k A_i| \leq \epsilon|V|$. *If* $\mathcal{P}$ *is not* ϵ*-regular, then there exists a collection* $\mathcal{P}' = \{A_1', \dots, A_l'\}$ *of pairwise disjoint equal-sized subsets of* V *refining* $\mathcal{P}$ *such that*

(1) *(The collection $\mathcal{P}'$ covers almost as much as $\mathcal{P}$):*

$$|V\setminus\bigcup_{i=1}^{l} A_i'| \le |V\setminus\bigcup_{i=1}^{k} A_i| + \frac{|V|}{2^{k+1}}.$$

(2) *(The collection $\mathcal{P}'$ has extended index τ^* increased by a positive constant depending only on ϵ):*

$$\tau^*(\mathcal{P}') \ge \tau^*(\mathcal{P}) + \epsilon^5(1-\epsilon)^2.$$

(3) *(The collection $\mathcal{P}'$ is bounded):* $k \le l \le k4^k$

Proof. The proof consists of two parts: first constructing the partitions and second proving the inequalities.

1. *Constructing the partitions.*

The main idea of the construction is as follows: Since the initial collection $\mathcal{P}$ is not ϵ-regular, there exist more than ϵk^2 ϵ-irregular pairs. We will apply Lemma 3.6 to obtain, for each such pair $\{A_i, A_j\}$, a partition $\mathcal{A}_{ij}$ of A_i and a partition $\mathcal{A}_{ji}$ of A_j with index increased by at least $\frac{\epsilon^4|A||B|}{|V|^2}$.

The next step will be to use the partitions $\mathcal{A}_{ij}$ to create the desired collection $\mathcal{P}'$. First we will take the refinement $\mathcal{A}$ of $\mathcal{P}$ induced by all the $\mathcal{A}_{ij}$ and $\mathcal{A}_{ji}$. Second we will slice the sets of $\mathcal{A}$ into sets of equal size, to produce the desired collection $\mathcal{P}'$. See Figure 4 for a diagram of the construction.

We begin the construction. For $1 \le i < j \le k$, choose partitions $\mathcal{A}_{ij}$ of A_i and $\mathcal{A}_{ji}$ of A_j as follows.

(1) If $\{A_i, A_j\}$ is ϵ-regular, let $\mathcal{A}_{ij} := \{A_i\}$ and $\mathcal{A}_{ji} := \{A_j\}$.
(2) If $\{A_i, A_j\}$ is not ϵ-regular, then by Lemma 3.6 there exist partitions $\mathcal{A}_{ij}$ of A_i and $\mathcal{A}_{ji}$ of A_j such that $|\mathcal{A}_{ij}|, |\mathcal{A}_{ji}| \le 2$ and

$$\tau(\mathcal{A}_{ij}, \mathcal{A}_{ji}) \ge \tau(A_i, A_j) + \epsilon^4\frac{c^2}{|V|^2}. \tag{3.12}$$

For each i, let $\mathcal{A}_i$ be the unique partition of A_i induced by all partitions $\mathcal{A}_{ij}$ with $j \in [1,k]\setminus\{i\}$. In terms of topology, $\mathcal{A}_i$ is the set of minimal nonempty open sets, or "atoms," in the finite topology generated on A_i by the sets in $\mathcal{A}_{ij}$ with $j \in [1,k]\setminus\{i\}$. Thus each $\mathcal{A}_i$ contains at most 2^{k-1} sets, since each starting set A_i has been cut into two pieces at most $k-1$ times.

Let

$$\mathcal{A} := \bigcup_{i=1}^{k} \mathcal{A}_i. \tag{3.13}$$

It follows from $|\mathcal{A}_i| \le 2^{k-1}$ that

$$k \le |\mathcal{A}| \le k2^{k-1},$$

satisfying part (3) of the claim.

Now we further refine $\mathcal{A}$ to obtain a collection with all sets of the same size. Let $\mathcal{P}' := \{A_1', \dots A_l'\}$ be a maximal collection of disjoint sets of size $d := \lfloor\frac{c}{4^k}\rfloor$ that refines $\mathcal{A}$. A trivial example exists: Since each $\mathcal{A}_i$ covers the set A_i, $|A_i| = c$, and $\mathcal{A}_i$ has at most 2^{k-1} sets, it follows that at least one component of $\mathcal{A}_i$ has size at least $\frac{c}{2^{k-1}} > d$. Thus we could pick a single set of d vertices from at least one component of each $\mathcal{A}_i$. Since V is finite, a maximal collection certainly exists. We claim $\mathcal{P}'$ has the desired properties.

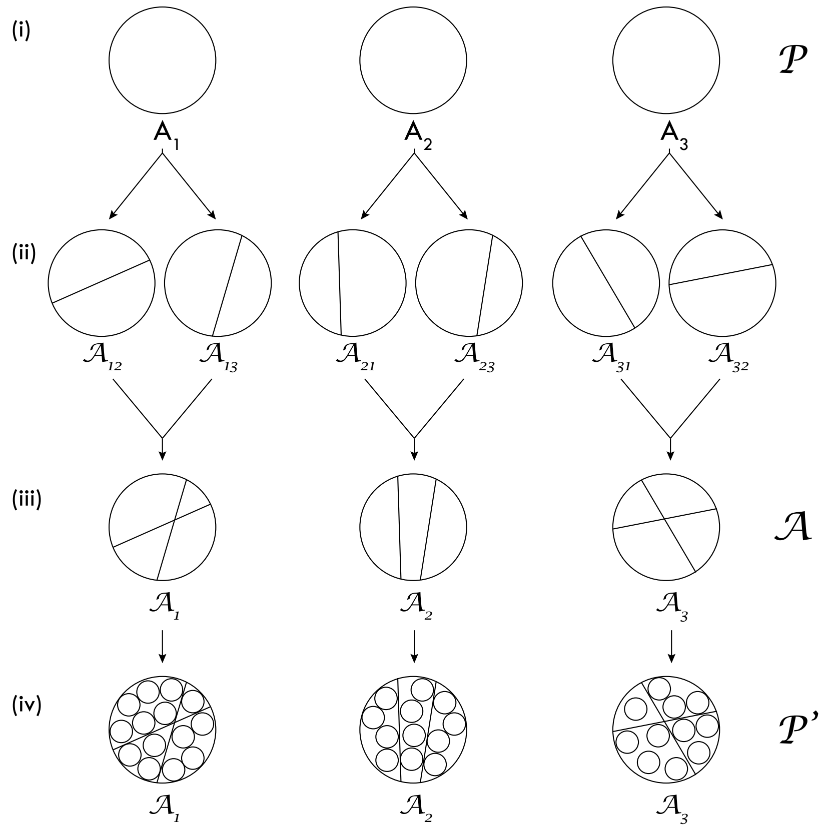

FIGURE 4. The construction of $\mathcal{P}'$, for $k = 3$ and assuming all pairs $\{A_i, A_j\}$ start ϵ-irregular. (*i*) We begin with the collection $\mathcal{P}$. (*ii*) Using Lemma 3.6, each set is partitioned in two different ways, with respect to the other sets. For example, the partition $\mathcal{A}_{12}$ is a partition of A_1 with respect to A_2. (*iii*) The induced partition $\mathcal{A}_i$ is taken on each set A_i. (*iv*) The sets are cut into pieces of size d.

2. *Proving the inequalities.*

Since each component of $\mathcal{P}'$ is of size $d = \lfloor \frac{c}{4^k} \rfloor$ and is contained in at least one set A_i of size c, we have $k \leq l \leq k4^k$.

Since $\mathcal{P}'$ excludes less than d vertices from each component of $\mathcal{A}$,

$$\begin{aligned}
|V\setminus \bigcup_{i=1}^{l} A_i'| \leq & |V\setminus \bigcup_{i=1}^{k} A_i| + d|\mathcal{A}| \\
\leq & |V\setminus \bigcup_{i=1}^{k} A_i| + \left(\frac{c}{4^k}\right) k2^{k-1} \\
= & |V\setminus \bigcup_{i=1}^{k} A_i| + \frac{ck}{2^{k+1}} \\
\leq & |V\setminus \bigcup_{i=1}^{k} A_i| + \frac{|V|}{2^{k+1}}, \quad (\text{since } \mathcal{P} \text{ is a partition of } |V| \text{ or fewer vertices into } k \text{ sets of size } c)\,.
\end{aligned}$$

We claim that $\tau^*(\mathcal{P}') \geq \tau^*(\mathcal{P}) + \epsilon^5(1-\epsilon)^2$. Let

$$\mathcal{A}_0' = \left\{\{x\}|x \in V\setminus \bigcup_{i=1}^{l} A_i'\right\}; \qquad \mathcal{A}_0 = \left\{\{x\}|x \in V\setminus \bigcup_{i=1}^{k} A_i\right\}.$$

The goal is to use Lemma 3.4 to show that $\tau^*(\mathcal{P}') \geq \tau^*(\mathcal{P})$. We must use τ^*, not τ, for the following reason. Although the collection $\mathcal{P}'$ refines $\mathcal{P}$, it may cover fewer vertices than $\mathcal{P}$. Lemma 3.4, however, only applies if $\mathcal{P}'$ refines $\mathcal{P}$ and covers the same vertices. The extended index τ^* avoids this problem by considering all the vertices not covered by a partition to be singletons in the partition. By definition, $\tau^*(\mathcal{P}') = \tau\left(\mathcal{P}' \cup \mathcal{A}_0'\right)$. Now we can make two observations: First, $\mathcal{P}' \cup \mathcal{A}_0'$ refines $\mathcal{A} \cup \mathcal{A}_0$ (since $\mathcal{P}'$ refines $\mathcal{A}$ and each set in $\mathcal{A}_0'$ is a singleton). Second, $\mathcal{P}' \cup \mathcal{A}_0'$ and $\mathcal{A} \cup \mathcal{A}_0$ both cover all of V. Thus $\mathcal{P}' \cup \mathcal{A}_0'$ and $\mathcal{A} \cup \mathcal{A}_0$ satisfy the conditions of Lemma 3.4.

$$\begin{aligned}
\tau^*(\mathcal{P}') = & \tau\left(\mathcal{P}' \cup \mathcal{A}_0'\right) \\
\geq & \tau\left(\mathcal{A} \cup \mathcal{A}_0\right), \Big(\text{ by Lemma 3.4, since } \mathcal{P}' \cup \mathcal{A}_0' \text{ refines } \mathcal{A} \cup \mathcal{A}_0 \text{ and by the above} \\
& \text{discussion}\Big). \\
= & \tau\left(\bigcup_{i=0}^{k} \mathcal{A}_i\right), \text{by (3.13).}
\end{aligned}$$

$$= \sum_{1\le i<j\le k} \tau(\mathcal{A}_i, \mathcal{A}_j) + \sum_{i=0}^{k} \tau(\mathcal{A}_i) + \sum_{i=1}^{k} \tau(\mathcal{A}_0, \mathcal{A}_i),$$

$$\ge \sum_{1\le i<j\le k} \tau(\mathcal{A}_{ij}, \mathcal{A}_{ji}) + \tau(\mathcal{A}_0) + \sum_{i=1}^{k} \tau(\mathcal{A}_0, \{A_i\}), \quad \Big(\text{by Lemma 3.4, since}$$

in the first term $\mathcal{A}_i$ refines $\mathcal{A}_{ij}$ and $\mathcal{A}_j$ refines $\mathcal{A}_{ji}$; in the second term the summands are nonnegative; and in the third term $\mathcal{A}_i$ refines $\{A_i\}\Big)$.

$$\ge \sum_{1\le i<j\le k} \tau(A_i, A_j) + \epsilon k^2 \frac{\epsilon^4 c^2}{|V|^2} + \tau(\mathcal{A}_0) + \sum_{i=1}^{k} \tau(\mathcal{A}_0, \{A_i\}), \quad \Big(\text{since at least } \epsilon k^2$$

irregular pairs $\{A_i, A_j\}$ satisfy (3.12)$\Big)$.

$$= \tau^*(\mathcal{P}) + \epsilon^5 \frac{k^2 c^2}{|V|^2}, \quad (\text{by definition of } \tau^*)\,.$$

$$\ge \tau^*(\mathcal{P}) + \epsilon^5 (1-\epsilon)^2, \quad \Bigg(\text{since } |V| - ck = |A_0| \le \epsilon |V| \text{ and } 0 < \epsilon < 1 \text{ imply}$$

$$\left(\frac{ck}{|V|}\right)^2 \ge (1-\epsilon)^2\Bigg).$$

□

The Regularity Lemma now follows by repeated application of Lemma 3.11 and by the boundedness of τ^*.

Proof. (Proof of the Regularity Lemma)

Let $\epsilon > 0$ and $m \in \mathbb{N}$. Lemma 3.2 shows that τ, and consequently τ^*, is bounded above by 1. Thus

$$b := \left\lceil \frac{1}{e^5 (1-\epsilon)^2} \right\rceil$$

is an upper bound on the number of times Lemma 3.11 can be applied before a partition is forced to be ϵ-regular. Choose $k_0 \in \mathbb{N}$ large enough that

$$2^{k_0} \ge b/\epsilon.$$

Let $\phi : \mathbb{N} \longrightarrow \mathbb{N}$ be given by $\phi(x) = x4^x$. Let

$$M = \max\left\{\phi^b(k_0), \frac{2k_0}{\epsilon}\right\},$$

where the exponent denotes iteration.

Let G be a graph on $n \ge m$ vertices. We aim to show that G admits an ϵ-regular collection containing between m and M sets. If $n \le M$, G admits an ϵ-regular collection into singletons. We can thus assume $n > M$.

Let $\mathcal{A} = \{A_1, \ldots, A_{k_0}\}$ be a collection of disjoint subsets of V, each of size $\lfloor \frac{n}{k_0} \rfloor$ and let $\mathcal{A}^s = \{A_1^s, \ldots, A_{k_s}^s\}$ denote a result of iterating Lemma 3.11 s times.

We need to check that after each iteration, $\mathcal{A}^s$ either is ϵ-regular or satisfies the hypotheses of Lemma 3.11. The sets of $\mathcal{A}^s$ are of the same size, so all there is to check is that $\mathcal{A}^s$ covers all but at most ϵn vertices. For $0 \le s \le b$, we have

$$\begin{aligned}
|V \setminus \bigcup_{i=1}^{k_s} A_i^s| =& |V \setminus \bigcup_{i=1}^{k_0} A_i| + \frac{sn}{2^{k_0+1}}, \quad \text{(by Lemma 3.11)}. \\
\le& k_0 + \frac{bn}{2^{k_0+1}}, \quad \big(\text{since } \mathcal{A} \text{ covers all but at most } k_0 \text{ vertices} \\
& \text{and since } s \le b\big). \\
\le& k_0 + \frac{\epsilon n}{2}, \quad \left(\text{since } 2^{k_0} \ge \frac{b}{\epsilon} \text{ implies } \frac{b}{2^{k_0}} \le \epsilon\right). \\
\le& \frac{\epsilon n}{2} + \frac{\epsilon n}{2}, \quad \left(\text{since } n \ge M \ge \frac{2k_0}{\epsilon} \text{ implies } k_0 \le \frac{\epsilon n}{2}\right). \\
=& \epsilon n.
\end{aligned}$$

Thus Lemma 3.11 can be iterated until $\mathcal{A}$ becomes ϵ-regular. The final ϵ-regular collection will have at most $\phi^b(k) = M$ components. □

4. The Key Lemma

We have taken the Key Lemma, regularity graphs, and some notations from [3].

The Regularity Lemma allows one to approximate the structure of a large dense graph G with a smaller graph R, called a regularity graph. The main idea of the Key Lemma is that the regularity graph can be used to describe sugraphs of the original graph.

Suppose for example that G is a large dense graph and that the Regularity Lemma has provided an ϵ-regular collection $\mathcal{A} = \{A_1, \dots, A_k\}$. We let R be the graph whose vertex set is $\mathcal{A}$. We let $\{A_i, A_j\}$ be an edge of R if the pair is ϵ-regular and if the edge density $d(A_i, A_j)$ is greater than some fixed choice of minimum density $d \in (0, 1]$ (we will make this precise in Definition 4.1).

For each edge $\{A_i, A_j\}$ of the regularity graph, we can put a lower bound on the number of edges of G with one vertex in A_i and one vertex in A_j. In fact, we can show that almost all vertices of A_i are connected to at least a $d - \epsilon$ fraction of the vertices of A_j (Lemma 4.4). Thus, if for example the regularity graph contains a triangle, it would be reasonable to hope that a triangle also exists in G.

Definition 4.1. (Regularity Graph) Let G be a graph, $\epsilon > 0$, $d \in (0, 1]$, and $\mathcal{A} = \{A_1, \dots, A_k\}$ be an ϵ-regular collection with $|A_1| = \cdots = |A_k| := l$. The regularity graph of G with respect to $\mathcal{A}$ and with parameters ϵ, d, and l is the graph R with vertex set $V(R) := \mathcal{A}$ and edge set

$$E(R) := \Big\{\{A_i, A_j\} \Big| \{A_i, A_j\} \text{ is an } \epsilon\text{-regular pair of density } d(A_i, A_j) \ge d\Big\}.$$

The regularity graph can be generalized as follows.

Definition 4.2. (s-fold Regularity Graph) Let $s \in \mathbb{N}$ and let R be a regularity graph, as above. The s-fold regularity graph R^s is obtained by replacing each vertex A_i of R with an s-set $A_i^s := \{v_i^1, \dots, v_i^s\}$ of s distinct vertices and by replacing each edge $\{A_i, A_j\}$ with a complete bipartite graph between A_i^s and A_j^s.

The Key Lemma asserts that under the right conditions, the existence of a graph in R^s implies its existence in G.

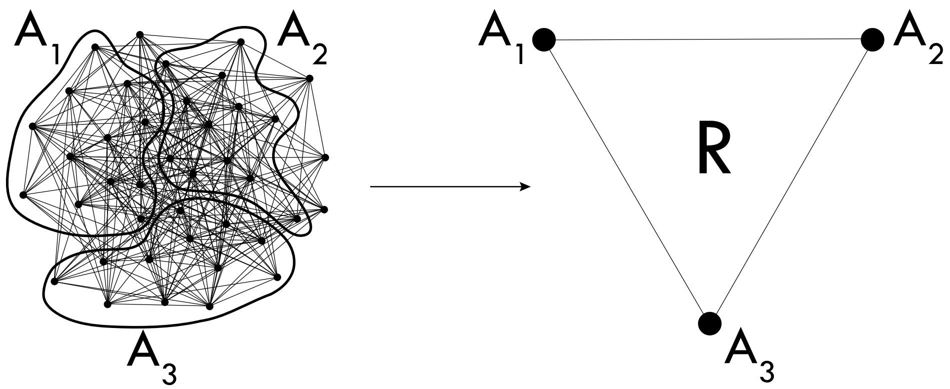

FIGURE 5. An example of a regularity graph produced by an ϵ-regular collection.

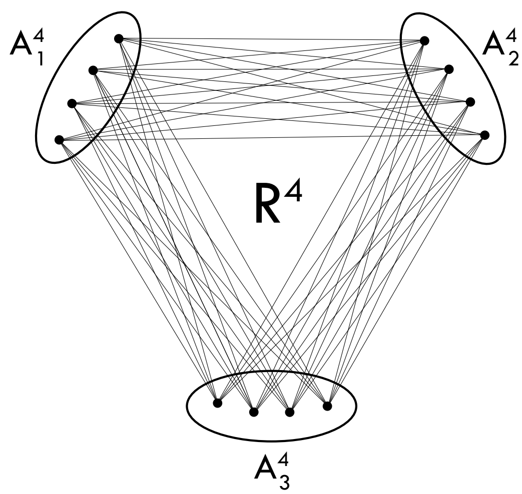

FIGURE 6. The 4-fold regularity graph produced by R from Figure 5.

Theorem 4.3. *(The Key Lemma) For all $d \in (0,1]$ and any integer $\Delta \geq 0$, there exists $\epsilon_0 > 0$ such that the following holds: if G is any graph; s is any natural number; R^s is any s-fold regularity graph of G with parameters $\epsilon \leq \epsilon_0$, d, and $l \geq s/\epsilon_0$, and H is any graph with maximum degree $\Delta(H) \leq \Delta$; then*

$$H \subset R^s \implies H \subset G.$$

Before proving the Key Lemma, we need the following trivial lemma.

Lemma 4.4. *Let $\{A, B\}$ be an ϵ-regular pair of density d; $Y \subset B$ with $|Y| \geq \epsilon|B|$; and*

$$X = \{x \in A | \deg(x, Y) < (d - \epsilon)|Y|\}.$$

Then $|X| < \epsilon|A|$.

Proof.

$$\begin{aligned} d(X,Y) &= \sum_{x\in X} \frac{\deg(x,Y)}{|X|\cdot|Y|} \\ &< \frac{(d-\epsilon)|X||Y|}{|X||Y|} \\ &= d-\epsilon. \end{aligned}$$

Therefore $|d(X,Y) - d(A,B)| > \epsilon$. Since $\{A,B\}$ is an ϵ-regular pair and since $|Y| \geq \epsilon$, this implies $|X| < \epsilon|A|$. □

Lemma 4.4 provides an upper bound on the number of vertices $x \in A$ with less than the expected number of neighbors in Y, provided Y is large enough. We will use this fact repeatedly in our proof of the Key Lemma. It will allow us to embed H into G by picking vertices one by one and, at each stage, using Lemma 4.4 to ensure that such a choice is possible. The proof involves keeping track, at each stage of the embedding, of which vertices in each component of the ϵ-regular collection are still valid choices and of which vertices can no longer be used because they contradict previously established adjacency relations. We will provide an algorithm and then prove by induction that the algorithm works.

While our proof borrows the main ideas and notation from [3], the algorithm is slightly based on that of [4] and the inductive part of the proof is largely original.

We will use the following notation in our proof:

Notation 4.5. (Neighbors) Let $G = (V,E)$ be a graph. For $v \in V$, let $N(v)$ be the set of vertices adjacent to v.

Proof. (Proof of the Key Lemma.) Choose $\epsilon_0 < d$ small such that

$$(\Delta+1)\epsilon_0 \leq (d-\epsilon_0)^{\Delta}. \tag{4.6}$$

Assume that H is a subgraph of R^s and let $V(H) := \{u_1,\dots,u_h\}$. Each vertex u_i of H is contained in one of the s-sets A^s_j for some $j \in [1,k]$. This defines a function $\sigma : [1,h] \longrightarrow [1,k]$ satisfying

$$u_i \in A^s_{\sigma(i)}$$

for each $i \in [1,h]$.

We aim to embed H into G by associating each vertex u_i of H with a unique vertex v_i in $A_{\sigma(i)}$. The goal is to show that $v_1 \in A_{\sigma(1)},\dots,v_h \in A_{\sigma(h)}$ can be chosen such that

1. (The v_i's are unique.) We have $u_i \neq u_j \implies v_i \neq v_j$.
2. (Edges are preserved.) If $\{u_i,u_j\} \in E(H)$, then $\{v_i,v_j\} \in E(G)$.

We will pick $v_1,\dots,v_h$ inductively. Immediately after having picked $v_1,\dots,v_n$, we denote by Y^n_i the set of vertices in $A_{\sigma(i)}$ that are still acceptable choices for v_i. At the start, $Y^0_1 = A_{\sigma(1)},\dots,Y^0_h = A_{\sigma(h)}$, since there are no established adjacency relations that might restrict the choices for $v_1,\dots,v_h$. At each stage, a vertex v_n is chosen in Y^{n-1}_n and the sets $Y^n_n,\dots,Y^n_h$ are created from the previous sets $Y^{n-1}_n,\dots,Y^{n-1}_h$ based on the adjacency relations introduced by the choice of v_n. The sets thus satisfy

$$A_{\sigma(i)} = Y^0_i \supset Y^1_i \supset \cdots \supset Y^{i-1}_i \supset Y^i_i = \{v_i\},$$

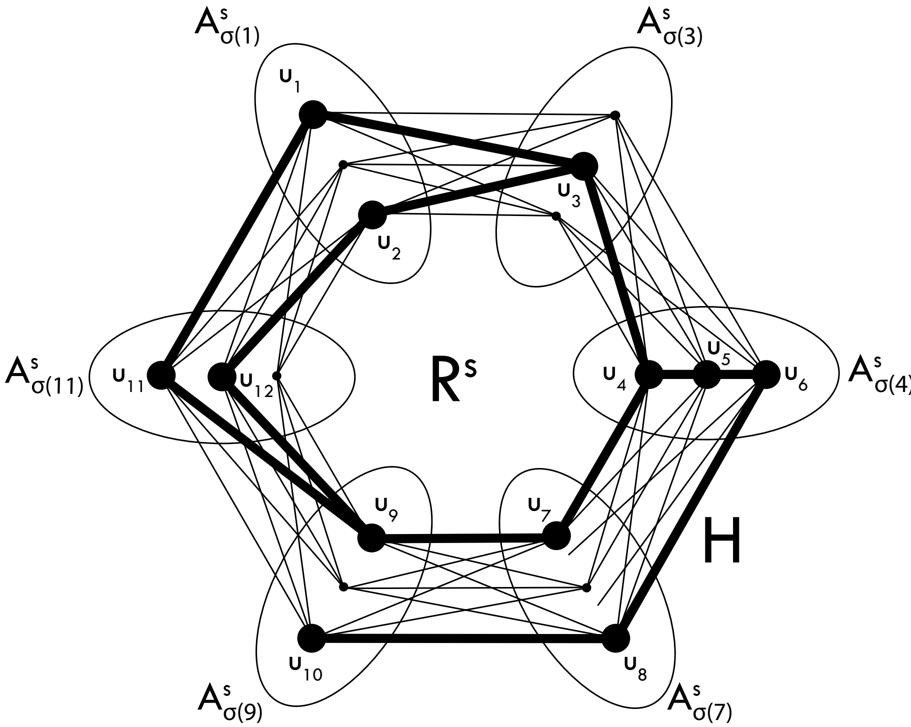

FIGURE 7. The graph H as a subgraph of R^s. In this example, R is a 6-cycle and $s = 3$.

for each $i \in [1, h]$.

The First Stage of the Algorithm

The first stage of the algorithm is to pick the vertex $v_1 \in Y_1^0 = A_{\sigma(1)}$ and to compute the sets $Y_1^1, \dots Y_h^1$. This stage is almost identical to the succeding stages. We explain it in detail to give the reader a picture of how the general algorithm will work, although this entire section is not really part of the proof. So the reader may want to skip to the section titled *Algorithm*.

The first step is to pick $v_1 \in A_{\sigma(1)}$. We need to pick v_1 in such a way that there is a good chance of a successful embedding. This can be done by looking at the set $N(u_1)$ of vertices in H adjacent to u_1. There are at most Δ such vertices. We label these vertices $N(u_1) = \{u_{\eta(1)}, \dots u_{\eta(d)}\}$, with $d := |N(u_1)| \leq \Delta$. Our goal can be accomplishing by picking $v_1 \in A_{\sigma(1)}$ such that $\deg(v_1, A_{\eta(i)})$ is large for each $u_{\eta(i)}$ adjacent to u_1. Lemma 4.4 will ensure that such a choice is possible. This will be made precise in the proof of proposition $P(n)$ later in the proof.

Once v_1 has been chosen, the next step is to update the sets $Y_1^0, \dots Y_h^0$ so that they accurately reflect the remaining acceptable choices for the to be determined vertices $v_2, \dots, v_h$. Since v_1 has been chosen, the set Y_1^0 collapses to a single point, and we have $Y_1^1 = \{v_1\}$.

Updating $Y_2^0, \dots, Y_h^0$ is slightly less trivial. In fact, we claim that

$$Y_i^1 := \begin{cases} Y_i^0 \cap N(v_1), & \text{if } \{u_i, u_1\} \in E(H) \\ Y_i^0, & \text{else} \end{cases}$$

for $i \in [2, h]$. To prove this, consider the two cases. First, assume that $\{u_i, u_1\} \in E(H)$. Then for any vertex $v \in Y_i^0$ we have:

$$\begin{aligned} v \text{ is an acceptable choice for } v_i &\iff \{v_1, v\} \in E(H), \quad \Big(\text{since } \{u_i, u_1\} \in E(H)\Big) \\ &\iff v \in N(v_1). \end{aligned}$$

For the second case, if $\{u_i, u_1\} \notin E(H)$, then the choice of v_1 has created no new restrictions on the acceptable choices for v_i, so $Y_i^1 = Y_i^0$, as claimed. The general algorithm is as follows:

Algorithm. For $n = 1, 2, \dots, h-1$,

(1) Pick $v_n \in Y_n^{n-1}$ such that $v_n \notin \{v_1, \dots, v_{n-1}\}$ and
$$\deg(v_n, Y_i^{n-1}) \geq (d - \epsilon)|Y_i^{n-1}|$$
holds for all $i \in [n+1, h]$ with $\{u_i, u_n\} \in E(H)$.

(2) Set
$$Y_i^n := \begin{cases} Y_i^{n-1} \cap N(v_n), & \text{if } \{u_i, u_n\} \in E(H) \\ Y_i^{n-1}, & \text{else} \end{cases}$$
for all $i \in [n+1, h]$.

To prove that the algorithm works, we shall induct on the following statement.

Proposition $P(n)$

(1) Vertices $v_1, \dots, v_n$ have been chosen according to the algorithm.

(2) If $n < h$, then for each $i \in [n+1, h]$,
$$|Y_i^n| \geq (d - \epsilon)^{\delta(n,i)} l,$$

where
$$\delta(n, i) := \deg\big(u_i, \{u_1, \dots, u_n\}\big)$$
is the number of edges in H with one vertex equal to u_i and another vertex in the set $\{u_1, \dots, u_n\}$ and where $\delta(0, i) := 0$.

Proof of the base case.

Let $n = 0$. Then part (1) of $P(0)$ is true, since 0 vertices have been chosen. Part (2) is true because

$$|Y_i^0| = |A_{\sigma(i)}| = l = (d - \epsilon)^0 l = (d - \epsilon)^{\delta(0,i)} l.$$

Proof of the inductive step.

Let $n \in [1, h]$ and assume $P(n-1)$.

First we prove part (1) of $P(n)$. The goal is to show that Y_i^{n-1} is not too small, and then Lemma 4.4 will show that there are at least s acceptable choices for v_n. Then we use the Pigeonhole Principle to show that at least one of these s choices is distinct from $v_1, \dots, v_{n-1}$.

First we claim that for each $i \in [n, h]$, the set Y_i^{n-1} is not too small.

$$\begin{aligned}
|Y_i^{n-1}| &\geq (d-\epsilon)^{\delta_{(n,i)}} l, \quad (\text{ by inductive hypothesis.}) \\
&\geq (d-\epsilon)^{\Delta} l, \quad \big(\text{since } \delta_{(n,i)} \leq \Delta.\big) \\
&\geq (\Delta+1)\epsilon_0 l, \quad (\text{ by assumption.}) \\
&= \Delta\epsilon_0 l + \epsilon_0 l \\
&\geq \Delta\epsilon l + s, \quad (\text{ since } \epsilon \leq \epsilon_0 \text{ and } l \geq s/\epsilon_0)\,. \qquad (4.7)
\end{aligned}$$

Now we can apply Lemma 4.4. For each $i \in [n+1, h]$ such that $\{u_i, u_n\} \in E(H)$, the sets

$$A := A_{\sigma(n)}; \qquad B := A_{\sigma(i)};$$

form an ϵ-regular pair. And letting $Y := Y_i^{n-1} \subset B$, we have

$$|Y| = |Y_i^{n-1}| \geq \epsilon l = \epsilon|B|, \text{ by (4.7) and since } \Delta \geq 1.$$

Thus by Lemma 4.4,

$$|\{v \in A_{\sigma(n)} |\deg(v, Y) < (d-\epsilon)|Y|\}| < \epsilon l.$$

Applying the above result to the at most Δ values of i for which u_i is adjacent to u_n, it follows that at most $\Delta\epsilon l$ elements $v \in A_{\sigma(n)}$ are such that for some $i \in [n+1, h]$, both $\deg(v, Y_i^{n-1}) < (d-\epsilon)|Y_i^{n-1}|$ and $\{u_i, u_n\} \in E(H)$. Thus there are at least

$$\begin{aligned}
|Y_n^{n-1}| - \Delta\epsilon l &\geq \Delta\epsilon l + s - \Delta\epsilon l, \quad \Big(\text{ by 4.7.}\Big) \\
&= s. \qquad (4.8)
\end{aligned}$$

acceptable choices for v_n.

Now we must check that of the at least s acceptable choices for v_n in Y_n^{n-1} given by (4.8), at least one such choice will be such that v_n is distinct from all the previously chosen vertices $v_1, \ldots, v_{n-1}$. Suppose indirectly that all of the at least s acceptable choices for v_n are equal to v_i for some $i \in [1, n-1]$. Let

$$I = \left\{ i \in [1, n-1] \middle| v_i \text{ is one of the at least } s \text{ acceptable choices for } v_n \text{ in } Y_n^{n-1} \right\}.$$

Then by our assumption $|I| \geq s$. By the inductive hypothesis and since $I \subset [1, n-1]$, we have successfully embedded $u_i \mapsto v_i$ for all $i \in I$. Since $\{v_i\}_{i\in I} \subset Y_n^{n-1} \subset A_{\sigma(n)}$, we have $\{u_i\}_{i\in I} \subset A^s_{\sigma(n)}$. Since we also know $u_n \in A^s_{\sigma(n)}$ (as given), we have

$$\{u_i\}_{i\in I} \cup \{u_n\} \subset A^s_{\sigma(n)}, \qquad (4.9)$$

It is also given that $u_1, \ldots, u_n$ are pairwise distinct, thus

$$|\{u_i\}_{i\in I} \cup \{u_n\}| = |I| + 1 \geq s + 1. \qquad (4.10)$$

Lines (4.9) and (4.10) contradict the fact that $|A^s_{\sigma(n)}| = s$, by definition. We thus conclude that a distinct and acceptable choice for v_n exists in $Y_n^{n-1} \subset A_{\sigma(n)}$.

Next we prove part (2) of $P(n)$. Since the vertex v_n has been chosen, part (2) of the algorithm defines the sets Y_i^n for $i \in [n+1, h]$. We assert that $|Y_i^n| \geq (d-\epsilon)^{\delta_{(n,i)}}$ for each $i \in [n+1, h]$. Fix i. To prove the inequality, let

$$\delta := \begin{cases} 1, & \text{if } \{u_i, u_n\} \in E(H) \\ 0, & \text{else.} \end{cases}$$

to indicate whether or not $\{u_i, u_n\}$ is an edge. Then

$$\begin{aligned}|Y_i^n| &\geq (d-\epsilon)^{\delta}|Y_i^{n-1}|, \quad \text{(by the choice of } v_n \text{ and } Y_i^n \text{ in the algorithm).}\\ &\geq (d-\epsilon)^{\delta}(d-\epsilon)^{\delta_{(n-1,i)}}, \quad \text{(by inductive hypothesis).}\\ &= (d-\epsilon)^{\delta_{(n-1,i)}+\delta}\\ &= (d-\epsilon)^{\delta_{(n,i)}}, \quad \big(\text{since } \delta_{(n,i)} = \delta(n-1,i)+\delta, \text{ trivially}\big).\end{aligned}$$

This proves part (2) of $P(n)$. $P(h)$ follows by induction.

Finally, we observe that $P(h)$ implies the claim. Part (1) of $P(h)$ asserts that vertices $v_1, \dots, v_h$ have been chosen according to the algorithm. Part (2) of the algorithm ensures that our embedding $u_i \mapsto v_i$ is injective and edge-preserving. Thus $v_1, \dots, v_h$ is a copy of H in G.

□

5. Acknowledgements

I thank Asaf Hadari and Mike Shulman for reading my work and providing many corrections.

References

[1] E. Szemerédi On sets of integers containing no k elements in arithmetic progression. Acta Arithmetica. XXVII., pp. 199-245. 1975.
[2] E. Szemerédi Regular Partitions of Graphs In: Problémes Combinatoires et Théorie des Graphes, No. 260, p. 399. Proceedings July 9-13, 1976, in Orsay. Centre National de la Recherche Scientifique (CNRS). Paris, 1978.
[3] R. Diestel Graph Theory Springer-Verlag. 1997.
[4] J. Komlos, M. Simonovits Szemerédi's Regularity Lemma and Its Applications in Graph Theory In: Combinatorics, Paul Erdos is Eighty, vol. 2 Janos Bolyai Mathematical Society. Budapest, 1996.