THE REGULARITY LEMMA

SGT 2013[1]

1. Density and ϵ-regular pairs

This section aims to introduce notions of density and regularity, as well as bringing together some of the simple results surrounding them.

Definition 1.1. *(Density) Let G be a graph, and let $X, Y \subseteq V(G)$ be disjoint. In particular G could be bipartite with vertex classes X and Y. We define the density, $d(X,Y)$, of the pair (X,Y) as: $d(X,Y) := \frac{e(X,Y)}{|X||Y|}$.*

Definition 1.2. *(ϵ-regularity) Let $\epsilon > 0$. Given a graph G and disjoint vertex sets $A, B \subseteq V(G)$ we say the pair (A,B) is ϵ-regular if for every $X \subseteq A$ and $Y \subseteq B$ such that $|X| > \epsilon|A|$ and $|Y| > \epsilon|B|$ we have $|d(X,Y) - d(A,B)| < \epsilon$.*

The definition formalises the concept of a pair of vertex classes having the edges between them distributed fairly uniformly.

Lemma 1.3. *Let (A,B) be an ϵ-regular pair of density d and $Y \subseteq B$ such that $|Y| > \epsilon|B|$. Then all but at most $\epsilon|A|$ vertices in A have more than $(d-\epsilon)|Y|$ neighbours in Y.*

Proof. Let $X := \{x \in A | \ d_Y(x) \leq (d-\epsilon)|Y|\} \subseteq A$. Assume $|X| > \epsilon|A|$. Then as (A,B) is an ϵ-regular pair and $|Y| > \epsilon|B|$, we have that $d(X,Y) > d - \epsilon$.

However, $e(X,Y) \leq |X|(d-\epsilon)|Y|$ by definition of X. Thus, $d(X,Y) \leq d - \epsilon$, a contradiction. So our assumption that $|X| > \epsilon|A|$ was false. $\square$

The next lemma tells us that reasonable size subgraphs of regular pairs are also regular.

Lemma 1.4. *(Slicing Lemma) Let $\alpha > \epsilon > 0$ and $\epsilon' := max\{\epsilon/\alpha, 2\epsilon\}$. Let (A,B) be an ϵ-regular pair with density d. Suppose $A' \subseteq A$ such that $|A'| \geq \alpha|A|$, and $B' \subseteq B$ such that $|B'| \geq \alpha|B|$. Then (A', B') is an ϵ'-regular pair with density $d' = d \pm \epsilon$.*

Proof. Firstly, $|A'| \geq \alpha|A| > \epsilon|A|$ and $|B'| \geq \alpha|B| > \epsilon|B|$ since $\alpha > \epsilon$. Thus, as (A,B) is ϵ-regular we know $|d' - d| < \epsilon$. Consider $X \subseteq A'$ and $Y \subseteq B'$ such that $|X| > \epsilon'|A'|$ and $|Y| > \epsilon'|B'|$. Then as $\epsilon' \geq \epsilon/\alpha$ and $|A'| \geq \alpha|A|$ we have

$$|X| > \epsilon'|A'| \geq \frac{\epsilon}{\alpha}|A'| \geq \epsilon|A|.$$

Similarly we obtain $|Y| > \epsilon|B|$.

Therefore, as (A,B) is ϵ-regular, we have $|d(X,Y) - d| < \epsilon$. Thus, as $\epsilon \leq \epsilon'/2$ and by the triangle inequality,

$$|d(X,Y) - d'| = |(d(X,Y) - d) + (d - d')| \leq |d(X,Y) - d| + |d' - d| < \epsilon + \epsilon \leq \epsilon'.$$

[1]These notes have been prepared by D. Kühn and D. Osthus for the participants of the SGT 2013 summer school in Oléron. The first five sections are based closely on the M.Sci thesis of A. Treglown, and we are grateful for his permission to use his .tex file. We are also grateful to Katherine Staden and Tim Townsend for help with preparing the document.

So, by definition, (A', B') is an ϵ'-regular pair. $\square$

The Slicing Lemma tells us that not too small subgraphs of an ϵ-regular pair are also regular with density close to that of the original pair. To get an idea as to why it is useful suppose that we are in a situation, where, for whatever reason, we only consider some of the vertices in an ϵ-regular pair. Then it seems good to know that all the properties of the original pair do not just disappear. But for now we will be content with an application of the Slicing Lemma which links the notion of regularity to that of super-regularity.

Definition 1.5. *(Super-regularity) Given a graph G and disjoint vertex sets $A, B \subseteq V(G)$, we say the pair (A, B) is (ϵ, δ)-super-regular if it is ε-regular and $d_B(a) \geq \delta|B|$ for all $a \in A$, and $d_A(b) \geq \delta|A|$ for all $b \in B$.*

Next we see that given a regular pair we can approximate it by a super-regular pair.

Lemma 1.6. *If (A, B) is an ϵ-regular pair with density d in a graph G (where $0 < \epsilon < 1/3$), then there exists $A' \subseteq A$ and $B' \subseteq B$ with $|A'| \geq (1-\epsilon)|A|$ and $|B'| \geq (1-\epsilon)|B|$, such (A', B') is a $(2\epsilon, d')$-super-regular pair, where $d' = d \pm 3\epsilon$.*

Proof. Let A' be the set of all vertices $x \in A$ such that $d_B(x) \geq (d-\epsilon)|B|$. Notice Lemma 1.3 implies $|A'| \geq (1-\epsilon)|A|$. Similarly, let B' be the set of all vertices $y \in B$ such that $d_A(y) \geq (d-\epsilon)|A|$. So again, $|B'| \geq (1-\epsilon)|B|$. Now let $\alpha := 1/2 > \epsilon$. We have $|A'| \geq (1-\epsilon)|A| > \alpha|A|$ and $|B'| \geq (1-\epsilon)|B| > \alpha|B|$. So, by the Slicing Lemma, (A', B') is a 2ϵ-regular pair with density d', where $d' = d \pm \epsilon$. Further, if $x \in A'$, $d_B(x) > (d-\epsilon)|B|$ and if $y \in B'$, $d_A(y) > (d-\epsilon)|A|$. Since $|A'| \geq (1-\epsilon)|A|$ and $|B'| \geq (1-\epsilon)|B|$ this tells us

$$d_{B'}(x) > (d-\epsilon)|B| - \epsilon|B| > (d-3\epsilon)|B'|$$

and similarly for A'. $\square$

2. The Regularity Lemma

Next we introduce Szemerédi's Regularity Lemma. We will not state the original version of the lemma, but a 'cleaner' form of it. Later we will see that the theorem below is just one of a number of forms of the lemma.

Theorem 2.1. *(Szemerédi's Regularity Lemma [8]) For every $\epsilon > 0$ and every $M' \in \mathbb{N}$ there exists an integer $M(\epsilon, M')$ such that every graph G of order $n \geq M'$ admits a partition $\{V_0, V_1, \ldots, V_k\}$ of $V(G)$ such that:*

(i) *$M' \leq k \leq M$,*
(ii) *$0 \leq |V_0| \leq \epsilon|G|$,*
(iii) *$|V_1| = \cdots = |V_k| = m$,*
(iv) *all but at most ϵk^2 of the pairs (V_i, V_j) with $1 \leq i < j \leq k$ are ϵ-regular.*

We call the set V_0 an exceptional set as, unlike usual partitions, V_0 may be empty. Further, the classes V_i are known as *clusters* and the partition described is called an *ϵ-regular partition.*

What the result essentially says is that all sufficiently large and dense graphs can be approximated by a 'random' graph. We can disregard the vertices in V_0 and the edges that lie inside some V_i or in a pair (V_i, V_j) which is not ϵ-regular. Then we have a subgraph such that all edges between two of our clusters are distributed fairly uniformly, as we would expect in a random graph. So V_0 acts like a bin, one which cannot get too full by condition (ii). We can choose M' to be large so that the clusters are not too large and hence, most edges go between different clusters. Thus, together with condition (iv) this ensures not too many edges are disregarded when considering the random-like approximation we mentioned. The upper bound M, however, means that large graphs will have clusters that are large too.

3. The degree form of the Regularity Lemma and the Key Lemma

The next result, the so-called degree form of the Regularity Lemma, is a very useful and applicable form of the Regularity Lemma.

Theorem 3.1. *(Degree form of the Regularity Lemma) For every $\epsilon > 0$ and every M' there are integers M and n_0 such that if $d \in [0, 1]$ and G is a graph with $|G| \geq n_0$, then there exists a partition of $V(G)$ into $k+1$ clusters $V_0, V_1, ..., V_k$, and there is a spanning subgraph $G' \subseteq G$ with the following properties:*

- $M' \leq k \leq M$,
- $|V_0| \leq \epsilon |G|$,
- *all clusters V_i for $i \geq 1$ are of the same size $m \leq \lceil \epsilon |G| \rceil$,*
- $d_{G'}(v) > d_G(v) - (d + \epsilon)|G|$ *for all* $v \in V(G)$,
- $e(G'[V_i]) = 0$ *for all* $i \geq 1$,
- *all pairs $G'[V_i, V_j]$ $(1 \leq i < j \leq k)$ are ϵ-regular, each with density either 0 or greater than d.*

Sketch of proof of Theorem 3.1 To obtain a partition as in Theorem 3.1, apply Theorem 2.1 with parameters d, ε', M'' satisfying $1/M'', \varepsilon' \ll \varepsilon, d, 1/M'$ to obtain clusters $V_1', \ldots, V_{k'}'$ and an exceptional set V_0'. (Here $a \ll b < 1$ means that there is an increasing function f such that all the calculations in the argument work as long as $a \leq f(b)$.) Let $m' := |V_1'| = \cdots = |V_{k'}'|$. Now delete all edges between pairs of clusters which are not ε'-regular and move any vertices into V_0' which were incident to at least $\varepsilon n/10$ (say) of these deleted edges. Secondly, delete all (remaining) edges between pairs of clusters whose density is at most $d + \varepsilon'$. Consider such a pair (V_i', V_j') of clusters. For every vertex $x \in V_i'$ which has more than $(d + 2\varepsilon')m'$ neighbours in V_j' mark all but $(d + 2\varepsilon')m'$ edges between x and V_j'. Do the same for the vertices in V_j' and more generally for all pairs of clusters of density at most $d + \varepsilon'$. It is easy to check that in total this yields at most $\varepsilon' n^2$ marked edges. Move all vertices into V_0' which are incident to at least $\varepsilon n/10$ of the marked edges. Thirdly, delete any edges within the clusters. Finally, we need to make sure that the clusters have equal size again (as we may have lost this property during the deletion process). This can be done by splitting up the clusters into smaller subclusters (which contain almost all the vertices and have equal size) and moving a small number of further vertices into V_0'. A straightforward calculation shows that the new exceptional set V_0 has size at most εn as required. □

G' can be thought of as obtained from G by tidying up. The main point of Theorem 3.1 is that this can be done in such a way that the degree of each vertex is only reduced slightly. We can 'clean up' G' further by defining the *pure graph* of G to be $G'' = G' - V_0$.

We now introduce a very important type of graph which is vital in applying the Regularity Lemma.

Definition 3.2. *(Reduced graph) Let G be a graph and $\{V_1, \dots, V_k\}$ a partition of $V(G)$. Given two parameters $\epsilon > 0$ and $d \in [0, 1)$ we define the reduced graph R of G as follows: its vertices are the clusters $V_1, \dots, V_k$ and there exists an edge between V_i and V_j precisely when (V_i, V_j) is ϵ-regular with density more than d.*

Most proofs involving the Regularity Lemma use reduced graphs. Often we will look at graph G and apply Theorem 3.1 to get the pure graph G''. We then take the reduced graph R of G''. In this case we also say that R is the reduced graph of G. Notice R provides an overview of the layout of the graph G'': it shows us when there exists a reasonable number of edges between two given clusters in G''. We can also define a reduced graph from the ϵ-regular partition obtained when applying Theorem 2.1. If we have a property of G such as an edge or degree condition, this will often give us a similar condition for R. An example of this is the lemma below.

Lemma 3.3. *Let G be a graph such that $\delta(G) \geq c|G|$ where c is a constant. Suppose we have applied the degree form of the Regularity Lemma to G and have defined from this the reduced graph R with parameters ϵ and d such that $2\epsilon \leq d$. If $d < c/2$ then $\delta(R) \geq (c - 2d)|R|$.*

Proof. Suppose not. Then there exists some $V_i \in V(R)$ such that $d(V_i) < (c-2d)|R|$. (Here we let $\{V_0, V_1, \dots, V_k\}$ be the ϵ-regular partition of $V(G)$, such that V_0 is the exceptional set and $|V_j| = m$ for all $j \in \{1, \dots, k\}$.) Let G'' be the pure graph of G.

By assumption, fewer than $(c - 2d)|R|$ of our clusters V_j form an ϵ-regular pair of density more than d with V_i in G''. At most m^2 edges go between such a pair. So the total number of edges coming out of V_i in G'' is less than

$$(c - 2d)|R|m^2 \leq (c - 2d)|G|m$$

since $|R|m \leq |G|$.

Given any vertex $x \in V(G'')$ we know that $d_{G''}(x) > d_G(x) - (d + 2\epsilon)|G|$. So for all $x \in V_i$, we have $d_{G''}(x) > (c - d - 2\epsilon)|G|$ since $\delta(G) \geq c|G|$. G'' contains no edges between vertices in V_i. So the number of edges coming out of V_i in G'' is just $\sum_{x \in V_i} d_{G''}(x)$. Thus,

$$(c - d - 2\epsilon)|G|m < \sum_{x \in V_i} d_{G''}(x) < (c - 2d)|G|m.$$

This implies $d < 2\epsilon$, a contradiction to our hypothesis. So our assumption is false, proving the claim. □

The lemma tells us that if we have a reduced graph R with parameters ϵ and d sufficiently small, then the minimum fraction of vertices a vertex in R is adjacent to is close to the corresponding fraction for the vertices in G. The fact that we can tell something about the minimum degree of the reduced graph of G from the minimum degree of G is very useful.

On the other hand, we will see that under certain conditions, properties of a reduced graph R of G are inherited by G. This is realised in the Key Lemma below. Before we state the Key Lemma we need some more notation. Given a graph R and positive integer t, let $R(t)$ be the graph obtained from R by replacing every vertex $x \in V(R)$ by a set U_x of t independent vertices, and joining $u \in U_x$ to $v \in U_y$ precisely when xy is an edge in R. That is we replace the edges of R by copies of $K_{t,t}$. We will refer to $R(t)$ as a 'blown-up' copy of R.

Theorem 3.4. *(Key Lemma) Given $d > \epsilon > 0$, a graph R, and a positive integer m, let us construct a graph G by replacing every vertex of R by m vertices, and replacing the edges of R with ϵ-regular pairs of density at least d. Let H be a subgraph of $R(t)$ with h vertices and maximum degree $\Delta > 0$, and let $\delta := d - \epsilon$ and $\epsilon_0 := \frac{\delta^\Delta}{2+\Delta}$. If $\epsilon \leq \epsilon_0$ and $t - 1 \leq \epsilon_0 m$, then $H \subseteq G$.*

The proof will actually guarantee not just one but $\Omega(m^h)$ copies of H in G.

Proof. We let $V(R) =: \{V_1, \dots, V_k\}$. By definition of G we will also use the convention that V_i is the vertex set in G that corresponds to the vertex $V_i \in V(R)$. We let U_i^t denote the vertex set of size t in $R(t)$ corresponding to V_i. Now H is a subgraph of $R(t)$ with vertices $u_1, \dots, u_h$ say. Each vertex u_i lies in one of the sets U_j^t. This defines a map $\sigma : [h] \mapsto [k]$. Our aim is to embed many labelled copies of H in G. We will define embeddings of the form $u_i \mapsto v_i \in V_{\sigma(i)}$. Thus, $v_1, \dots, v_h$ will be distinct.

We will describe an algorithm below such that in the ith step we define v_i. We will need that there is a sufficient number of choices for v_i for each i in order to obtain the required number of labelled copies of H in G. Given some u_i we will have at each step a candidate set for v_i. At the jth step this will be called Y_i^j. Initially (i.e. in 'step' 0) we have $Y_i^0 := V_{\sigma(i)}$ for all i. In particular, $|Y_i^0| = m$ for all i. After every application of the algorithm we will want to update the candidate sets for all vertices v_j yet to be defined. That is, if $u_i u_j \in E(H)$ and v_i is defined, we can only consider vertices in V_j as candidates for v_j if they are adjacent to v_i in G. The algorithm at step $i \geq 1$ consists of two steps.

Step 1: Picking v_i. We pick $v_i \in Y_i^{i-1}$ such that

$$d_G(v_i, Y_j^{i-1}) > \delta |Y_j^{i-1}| \tag{1}$$

for all $j > i$ such that $v_i v_j \in E(H)$.

Step 2: Updating the candidate sets. We set $Y_j^i := Y_j^{i-1} \cap N(v_i)$ if $u_i u_j \in E(H)$, or $Y_j^i := Y_j^{i-1}$ otherwise (for $j > i$).

For $i < j$ we define $d_{ij} := |\{l \in [i] : u_l u_j \in E(H)\}|$. We claim that if $d_{ij} > 0$ then $|Y_j^i| > \delta^{d_{ij}} m$: Given u_j the initial candidate set is $V_{\sigma(j)} = Y_j^0$ which has size m. To obtain the candidate set Y_j^i, by definition of d_{ij} we have had to shrink Y_j^0 d_{ij} times. But given some $l \leq i$ such that $u_l u_j \in E(H)$ we have $|Y_j^l| = |Y_j^{l-1} \cap N(v_j)| > \delta |Y_j^{l-1}|$. So $|Y_j^i| > \delta^{d_{ij}} m$.

Notice that if $d_{ij} = 0$ then $|Y_j^i| = m$. Now for all $i < j$, $|Y_j^i| > \delta^\Delta \geq \epsilon m$ and $Y_j^i \subseteq V_{\sigma(j)}$. If u_j is a neighbour of u_i in H then since $u_i \in U_{\sigma(i)}^t$ and $u_j \in U_{\sigma(j)}^t$, we have that $V_{\sigma(i)}$ and $V_{\sigma(i)}$ are adjacent in R. So by definition of G we have that $(V_{\sigma(i)}, V_{\sigma(j)})$ is an ϵ-regular pair in G with density at least d. Thus, by Lemma 1.3 all but at most ϵm vertices of Y_{i-1}^i satisfy (1) for our specific j. Now u_i has at most Δ neighbours in H. So all but at most $\Delta \epsilon m$ vertices of Y_{i-1}^i satisfy (1) for all $j > i$. At most $t - 1$ vertices before u_i were given an image in $V_{\sigma(i)}$. Thus, we have at least

$$|Y_i^{i-1}| - \Delta \epsilon m - (t-1) > (\delta^\Delta - \Delta \epsilon) m - (t-1) \geq 0$$

free choices for each v_i. This shows we can apply the algorithm at each step, thus we obtain vertices $v_1, ..., v_h$ in G, and by the construction of our algorithm if $u_i u_j \in E(H)$ then $v_i v_j \in E(G)$. So $H \subseteq G$. ∎

Although the Key Lemma does not specify this, we usually think of R being a reduced graph of a graph G, and the way 'G' is specified in the Key Lemma, this refers to the pure graph G'' of G. Thus, the message one should get from the Key Lemma is the following: Given a graph G and the reduced graph R, G will contain a copy of any relatively sparse graph H for which we know that $H \subseteq R$ or $H \subseteq R(t)$. Note that all the copies of H in G which we found in the proof of the Key Lemma had the following property: If u_i is a vertex in a copy of H in $R(t)$ then it was embedded into $V_{\sigma(i)}$ where $u_i \in U^t_{\sigma(i)}$. Thus, we may be able to find another, disjoint, copy of H in G by applying the Key Lemma again.

4. The Erdős-Stone Theorem

Definition 4.1. *Let $n \in \mathbb{N}$ and H be a graph. Then*

$$ex(n, H) := \max\{e(G) : |G| = n \text{ and } H \not\subseteq G\}.$$

The *Turán graph* $T_{r-1}(n)$ is the complete $(r-1)$-partite graph on n vertices such that the vertex classes are as equal as possible. We let $t_{r-1}(n) := e(T_{r-1}(n))$. It is not hard to see that

$$\lim_{n\to\infty} \frac{t_{r-1}(n)}{\binom{n}{2}} = 1 - \frac{1}{r-1}, \qquad \text{and for } n \geq r \qquad t_{r-1}(n) \leq \left(1 - \frac{1}{r-1}\right)\frac{n^2}{2}. \tag{2}$$

Theorem 4.2. *(Turán [9]) For $n, r \in \mathbb{N}$ with $r > 1$, we have $ex(n, K_r) = t_{r-1}(n)$. In particular if $|G| = n \geq r$ and $e(G) > (1 - \frac{1}{r-1})\frac{n^2}{2}$ then $K_r \subseteq G$.*

Notice that r is the chromatic number of K_r. Theorem 4.4 will show that this is important. Also, from our remarks before Theorem 4.2, we observe that the lower bound on the number of edges is asymptotically best possible.

Corollary 4.3. *For $r \in \mathbb{N}$ we have*

$$\lim_{n\to\infty} \frac{ex(n, K_r)}{\binom{n}{2}} = 1 - \frac{1}{r-1}.$$

The next result is an extension to this, giving a lower bound on the number of edges in a graph that guarantees some graph H as its subgraph.

Theorem 4.4. *(Erdős, Stone [3] and Erdős, Simonovits [2]) Given any $\epsilon > 0$ and any graph H there is an $N(H, \epsilon)$ such that if $n \geq N$ and G is a graph on n vertices with*

$$e(G) > \left(1 - \tfrac{1}{\chi(H)-1} + \epsilon\right)\tfrac{n^2}{2}$$

then $H \subseteq G$.

Later on we will use the Regularity Lemma to prove this result, though it can also be proven directly. From Theorem 4.4 we can deduce the following.

Theorem 4.5. *(Fundamental Theorem of Extremal Graph Theory)*

$$\lim_{n\to\infty} \frac{ex(n, H)}{\binom{n}{2}} = 1 - \frac{1}{\chi(H)-1}$$

What this tells us is that the chromatic number of H is the important value when it comes to forcing a copy H into another graph. We write $K_r^s := K_r(s)$ for the blow-up of K_r which is obtained from K_r by replacing every vertex by a set of s independent vertices.

Theorem 4.6. *For all integers $r \geq 2$ and $s \geq 1$, and for every $\epsilon > 0$, there exists an integer $n_0(r, s, \epsilon)$ such that if a graph G has $n \geq n_0$ vertices and*

$$e(G) \geq t_{r-1}(n) + \epsilon n^2,$$

then $K_r^s \subseteq G$.

Recall that $t_{r-1}(n)$ is the number of edges in the Turán graph $T_{r-1}(n)$. It is easy to see that Theorem 4.6 implies Theorem 4.4: Let $\epsilon > 0$ and H be any graph. We define $r := \chi(H)$ and $s := |H|$. Notice Theorem 4.4 is trivial for $\chi(H) = 1$ so we can assume $r \geq 2$. Let n_0 be the output of Theorem 4.6 under input r, s and $\frac{\epsilon}{2}$. Thus, n_0 depends only on ϵ and H. Further, if a graph G on $n \geq n_0$ vertices satisfies $e(G) > (1 - \frac{1}{\chi(H)-1} + \epsilon)\frac{n^2}{2}$ then

$$e(G) > \left(1 - \frac{1}{r-1}\right)\frac{n^2}{2} + \frac{\epsilon}{2}n^2 \overset{(2)}{\geq} t_{r-1}(n) + \frac{\epsilon}{2}n^2$$

and so Theorem 4.6 tells us that $K_r^s \subseteq G$. But $H \subseteq K_r^s$ so $H \subseteq G$ as required.

We now give a proof of Theorem 4.6.

Proof of Theorem 4.6. Suppose we have integers $r \geq 2$, $s \geq 1$ and $\gamma > 0$. Suppose that G is a graph with sufficiently large order n and $e(G) \geq t_{r-1}(n) + \gamma n^2$. We define $d := \gamma$ and $\epsilon > 0$ so that:

(1) $3\epsilon < \gamma$,
(2) $\epsilon \leq \frac{(\gamma-\epsilon)^{(r-1)s}}{2+(r-1)s}$.

Further we let $\delta := \gamma - 3\epsilon > 0$. Now (2.1) implies

$$\lim_{n\to\infty} \frac{t_{r-1}(n)}{n^2/2} = 1 - \frac{1}{r-1}.$$

In particular this implies that for any $\eta > 0$ provided n is large enough then $\frac{t_{r-1}(n)}{n^2/2} + \eta > 1 - \frac{1}{r-1}$. This will be useful later in the proof.

We apply the degree form of the Regularity Lemma to our given graph G with input ϵ and d. Thus, we obtain a partition $\{V_0, V_1, \dots, V_k\}$ of $V(G)$, where $m := |V_i| \leq \lceil \epsilon n \rceil$ (for all $i \in \{1, \dots, k\}$) and a pure graph G''. From this we obtain the reduced graph R with parameters ϵ and d.

Recall that $e(G'') > e(G) - (d + 3\epsilon)\frac{n^2}{2}$. Also, each edge in R corresponds to at most m^2 edges in G'' and every edge in G'' must correspond to such an edge in R. So $e(R)m^2 \geq e(G'')$. Hence, provided n is sufficiently large,

$$\begin{aligned} e(R) > \frac{e(G) - (d+3\epsilon)\frac{n^2}{2}}{m^2} &\geq \frac{1}{2}k^2\left(\frac{t_{r-1}(n) + \gamma n^2 - (d+3\epsilon)\frac{n^2}{2}}{\frac{1}{2}(mk)^2}\right) \\ &\geq \frac{1}{2}k^2\left(\frac{t_{r-1}(n) + \gamma n^2 - (d+3\epsilon)\frac{n^2}{2}}{\frac{1}{2}n^2}\right) = \frac{1}{2}k^2\left(\frac{t_{r-1}(n)}{\frac{1}{2}n^2} + 2\gamma - (d+3\epsilon)\right) \\ &= \frac{1}{2}k^2\left(\frac{t_{r-1}(n)}{\frac{1}{2}n^2} + \delta\right) > \frac{1}{2}k^2\left(1 - \frac{1}{r-1}\right). \end{aligned}$$

Therefore, by Turán's Theorem, $K_r \subseteq R$. So $K_r^s \subseteq R(s)$. If n is sufficiently large then the choice of ϵ enables us to apply the Key Lemma giving $K_r^s \subseteq G'' \subseteq G$, as required. $\square$

This proof is not the original one (and much more heavy-handed), but it nicely illustrates the regularity method. The idea of our proof is as follows: Choosing ϵ and d small enough, and n large we can ensure that the reduced graph R of G still has lots of edges, namely more than $t_{r-1}(k)$. Thus, we can apply Turán's Theorem to R which implies $K_r \subseteq R$. Hence, $K_r^s \subseteq R(s)$. But then provided n is large and ϵ is small, the Key Lemma tells us that $K_r^s \subseteq G$, as desired.

5. Applying the Regularity Lemma

In the last section we noted that the Key Lemma can be used as a way of embedding some graph H into a graph G. Thus, many applications of the Regularity Lemma are concerned with embedding and packing problems. In this section we give the framework of a typical application of the Regularity Lemma in these situations. Throughout this section H is a graph that we are trying to embed into G.

Step 1: Preparing G. We apply Theorem 3.1 to G. We can define the reduced graph R of G with suitable parameters ϵ and d.

Step 2: Finding structure in R. In an application of the Regularity Lemma we will have some condition on G. From this condition we can then obtain information about R. Typically we may have an edge or degree condition on G. From these conditions we can usually get a similar condition for R. An example of this is Lemma 3.3.

Thus, if we have a bound, involving a slack term, on the minimum degree of, or the number of edges in G, choosing ϵ and d small enough will often give us an equivalent property of R without the slack term. For example, to prove Theorem 4.4 we have to consider a graph G of order n such that $e(G) \geq t_{r-1}(n) + \gamma n^2$ where $\gamma > 0$. Choosing ϵ and d carefully we can deduce that $e(R) \geq t_{r-1}(k)$ where $|R| = k$.

Once we have found a suitable property of R we can apply an embedding result to obtain structure in R. For example Turán's Theorem will be applied at this stage in the proof of Theorem 4.4.

Step 3: Applying the Key Lemma. In Step 2 we will have embedded graph(s) into R. We could be in a situation where we simply want to show that G contains a copy of a given graph H. Thus, in Step 2 we usually will have embedded H or some graph containing H into R. Provided ϵ is small and $|G|$ is large enough, the Key Lemma implies $H \subseteq G$ as required. Alternatively, we may have embedded a graph K into R which does not contain H but such that a 'blown-up' copy of K contains H.

At this stage of a proof involving the Regularity Lemma, we may wish to repeatedly apply the Key Lemma. We will do this if we wish to embed a graph K 'piece by piece' into G.

6. The Blow-up lemma

The following is a special case of the Blow-up lemma of Komlós, Sárközy and Szemerédi [4]. It implies that dense superregular pairs behave like complete bipartite graphs with respect to containing bounded degree graphs as subgraphs, i.e. if the superregular pair has vertex classes V_i and V_j then any bounded degree bipartite graph on these vertex classes is a subgraph of this superregular pair. The difference to the Key Lemma is that one can obtain a spanning copy of H.

Lemma 6.1 (Blow-up lemma, bipartite case)**.** *Given $d > 0$ and $\Delta \in \mathbb{N}$, there is a positive constant $\varepsilon_0 = \varepsilon_0(d, \Delta)$ such that the following holds for every $\varepsilon < \varepsilon_0$. Given*

$m \in \mathbb{N}$, let G^* be an (ε, d)-superregular bipartite graph with vertex classes of size m. Then G^* contains a copy of every subgraph H of $K_{m,m}$ with $\Delta(H) \leq \Delta$.

One of the earliest applications was by Komlós, Sárközy and Szemerédi [5] to prove the Pósa-Seymour conjecture for all large graphs: For every $k \geq 1$ there is an integer n_0 so that every graph G on $n \geq n_0$ vertices and with $\delta(G) \geq \frac{k}{k+1}n$ contains the kth power of a Hamilton cycle. There are many further applications, see e.g. [6].

For example, Kühn and Osthus [7] used it to determine, for all H, the minimum degree which forces a perfect H-matching in a dense graph G (here a perfect H-matching is a set of copies of H covering all vertices of G). Böttcher, Taraz and Schacht gave general conditions ensuring the existence of subgraphs of small bandwidth [1].

However, as we shall see, one does not always need the blow-up lemma to obtain spanning subgraphs.

References

[1] J. Böttcher, M. Schacht and A. Taraz, Proof of the bandwidth conjecture of Bollobás and Komlós, *Math. Ann.* **343** (2009), 175–205.

[2] P. Erdős and M. Simonovits, A limit theorem in graph theory, *Studia Sci. Math. Hungar.* **1** (1966) 51–57.

[3] P. Erdős and A.H. Stone, On the structure of linear graphs, *Bull. Amer. Math. Soc.* **52** (1946).

[4] J. Komlós, G.N. Sárközy and E. Szemerédi, Blow-up Lemma, *Combinatorica* **17** (1997) 109–123.

[5] J. Komlós, G. N. Sárközy and E. Szemerédi, Proof of the Seymour conjecture for large graphs, *Ann. Combin.* **2** (1998), 43–60.

[6] D. Kühn and D. Osthus, Embedding large subgraphs into dense graphs, *Surveys in Combinatorics 2009* (editors S. Huczynka, J. Mitchell, C. Roney-Dougal), Cambridge University Press, 2009, 137–167.

[7] D. Kühn and D. Osthus, The minimum degree threshold for perfect graph packings, *Combinatorica* **29** (2009) 65–107.

[8] E. Szemerédi, Regular partitions of graphs, *Colloques Internationaux C.N.R.S. No. 260 - Problèmes Combinatories et Thèorie des Graphes* (1976) 399-401.

[9] P. Turán, On an extremal problem in graph theory, *Matematikai és Fizikai Lapok* (1941) **48** 436–452.