COMBINATORICA
Bolyai Society – Springer-Verlag

COMBINATORICA **17** (1) (1997) 109–123

BLOW-UP LEMMA

JÁNOS KOMLÓS, GÁBOR N. SÁRKÖZY, ENDRE SZEMERÉDI

Received August 16, 1995

Regular pairs behave like complete bipartite graphs from the point of view of bounded degree subgraphs.

1. Introduction

The Regularity Lemma [16] is a powerful tool in Graph Theory and its applications. It basically says that every graph can be well approximated by the union of a constant number of random-looking bipartite graphs — called regular pairs (see the definitions below).

These bipartite graphs share many local properties with random bipartite graphs, e.g. most degrees are about the same, most pairs of vertices have about as many common neighbours as is expected in a random graph, and so on. These particular local properties imply other ones. Most importantly, they imply that every fixed bipartite graph H can be found as a subgraph in any large enough regular pair G (with a given positive density).

This however is not an exclusive property of regular pairs, or that of random graphs, since classical Extremal Graph Theory tells us that any given bipartite graph H is contained as a subgraph in all large enough dense graphs (Kővári–T. Sós-Turán [8] and Erdős-Stone [9]). The power of using the Regularity Lemma becomes apparent only when extended to much larger subgraphs H. Two examples are the theorem of Chvátal–Rödl–Szemerédi–Trotter [6] stating that all bounded degree graphs have linear Ramsey numbers, and the Alon-Yuster theorem [4] providing almost perfect coverings of a large G with copies of a small H.

Mathematics Subject Classification (1991): 05 C 35

Some new papers deal with embedding problems where H and G have the same order (*spanning subgraphs*). Classical Ramsey-Turán theory tells us that this is hopeless (the Ramsey and Turán numbers are huge) unless H is small in some sense. Two natural classes of such small H are trees (forests) and bounded degree graphs. An example for the first one is the proof of the Bollobás conjecture on spanning trees [11]. Examples for the second one are papers of the same authors on the Pósa–Seymour conjecture [12, 13] and on the Alon-Yuster conjecture [14]. These papers use a randomized version of the greedy algorithm to embed most vertices of H and then finish the embedding by using a König–Hall argument. The proofs were further refined to lead to the tool called Blow-up Lemma in this paper. It basically says that regular pairs behave like complete bipartite graphs from the point of view of bounded degree subgraphs.

1.1. Notations

All graphs are simple, that is, they have no loops or multiple edges. $v(G)$ is the number of vertices in G (order), $e(G)$ is the number of edges in G (size). $\deg(v)$ (or $\deg_G(v)$) is the degree of vertex v (within the graph G), and $\deg(v,Y)$ (or $\deg_G(v,Y)$) is the number of neighbours of v in Y. $\delta(G), \Delta(G)$ and $t(G)$ are the minimum degree, maximum degree and average degree of G. $N(x)$ (or $N_G(x)$) is the set of neighbours of the vertex x, and $e(X,Y)$ is the number of edges between X and Y. A bipartite graph G with color-classes A and B and edges E will sometimes be written as $G=(A,B,E)$. For disjoint X,Y, we define the *density*

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

The density of a bipartite graph $G=(A,B,E)$ is the number

$$d(G) = d(A,B) = \frac{|E|}{|A| \cdot |B|}.$$

For two disjoint subsets A,B of $V(G)$, the bipartite graph with vertex set $A \cup B$ which has all the edges of G with one endpoint in A and the other in B is called the pair (A,B).

A pair (A,B) is *ε-regular* if for every $X \subset A$ and $Y \subset B$ satisfying

$$|X| > \varepsilon|A| \quad \text{and} \quad |Y| > \varepsilon|B|$$

we have

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

A pair (A,B) is *(ε,δ)-super-regular* if for every $X \subset A$ and $Y \subset B$ satisfying

$$|X| > \varepsilon|A| \quad \text{and} \quad |Y| > \varepsilon|B|$$

we have

$$e(X,Y) > \delta|X||Y|,$$

and furthermore,

$$\deg(a) > \delta|B| \quad \text{for all} \quad a \in A, \quad \text{and} \quad \deg(b) > \delta|A| \quad \text{for all} \quad b \in B.$$

H is embeddable into G if G has a subgraph isomorphic to H, that is, if there is a one-to-one map (injection) $\varphi: V(H) \to V(G)$ such that $\{x,y\} \in E(H)$ implies $\{\varphi(x), \varphi(y)\} \in E(G)$. The cardinality of a set S will mostly be denoted by $|S|$, but sometimes we write $\#S$.

1.2. The Blow-up Lemma

Theorem 1 (Blow-up Lemma). *Given a graph R of order r and positive parameters δ, Δ, there exists a positive $\varepsilon = \varepsilon(\delta, \Delta, r)$ such that the following holds. Let n_1, n_2, ..., n_r be arbitrary positive integers and let us replace the vertices v_1, v_2, ..., v_r of R with pairwise disjoint sets V_1, V_2, ..., V_r of sizes n_1, n_2, ..., n_r (blowing up). We construct two graphs on the same vertex-set $V = \cup V_i$. The first graph* **R** *is obtained by replacing each edge $\{v_i, v_j\}$ of R with the complete bipartite graph between the corresponding vertex-sets V_i and V_j. A sparser graph G is constructed by replacing each edge $\{v_i, v_j\}$ arbitrarily with an (ε, δ)-super-regular pair between V_i and V_j. If a graph H with $\Delta(H) \le \Delta$ is embeddable into* **R** *then it is already embeddable into G.*

1.3. Sketch of proof

The classical proofs using the Regularity Lemma for embedding a graph H into a graph G typically follow a greedy embedding algorithm like this: For the location of the first vertex of H select a typical vertex v_1 of G — that is, any vertex with a typical degree. After having embedded k vertices of H into $v_1, \ldots, v_k$ of G, select v_{k+1} from the appropriate color class of G to be any vertex with a large degree into the common neighbourhood of the already selected vertices in the same color class. This selection guarantees that all selected vertices in the same color class have a large common neighbourhood, and thus not only do we end up with a copy of the desired H, but we may obtain many such copies (proportional to $order(G)^{order(H)}$). This can easily be extended from embedding bipartite graphs to embedding graphs with a higher but fixed chromatic number. We will use a randomized version of this greedy algorithm.

Let $n = \sum n_i$. We will assume that $|V(H)| = |V(G)| = n$. We embed the vertices of H one-by-one by following a randomized greedy algorithm (**RGA**), which works smoothly until there is only a small proportion of $V(H)$ left, and then it may get stuck hopelessly. To avoid this, we will set aside a positive proportion of special vertices of H called "buffer vertices".

In Phase 1 of the **RGA** we will embed the vertices of H one-by-one into G until all non-buffer vertices are embedded. For each x not embedded yet (including the buffer vertices) we keep track of an ever shrinking *host-set* $V_{t,x}$ that the image of x is confined to at time t. These sets are defined in the following natural way. If x is to be embedded into V_i, then for each t we define an auxiliary set $C_{t,x}$ to be the intersection of V_i with the common neighbourhood (in G) of the images of the neighbours of x (in H) already embedded by time t. Clearly x must be embedded into the set of vertices in $C_{t,x}$ still unused by time t. This set is the set $V_{t,x}$. Now if x is embedded at time t then we select the image of x randomly (uniformly) from $V_{t-1,x}$. While the sets $C_{t,x}$ are all large by the regularity condition, $V_{t,x}$ — the unused part of $C_{t,x}$ — may get very small or even empty if we use a greedy algorithm. This is the reason we use a randomized algorithm; the randomization ensures that we do not delete vertices from the host-sets $V_{t,x}$ in a worse case manner, but only proportionately to the time t. Some exceptions to this proportionality may be accidentally created, and then we immediately deal with these exceptional vertices by putting them into a first-in first-out queue (FIFO) $q(t)$ whose vertices have a priority at the selection of the next vertices to be embedded.

This embedding is continued in Phase 1 until all non-buffered vertices are embedded, unless the queue gets too long, in which case we abort the algorithm. T will denote the (random) last moment of Phase 1. In Phase 2, we embed the leftover $n-T$ vertices by using a König-Hall type argument.

2. Proof

We are given the parameters $r = |V(R)|$, and ε, δ of the super-regularity condition. We define several new parameters: $\delta', \delta'', \delta'''$ are used by the RGA algorithm and κ by the König-Hall argument of Section 2.7, while n_0 and γ are auxiliary parameters. The parameters have the following order:

$$\varepsilon \ll \kappa \ll \gamma \ll \delta''' \ll \delta'' \ll \delta' \ll \delta \tag{1}$$

where $\alpha \ll \beta$ means that α is small enough in terms of β. One possible choice is the following:

$$\delta' = \frac{\delta^{\Delta+1}}{4r^2\Delta^2}, \quad \delta'' = (\delta')^2, \quad \delta''' = (\delta')^3, \quad \gamma = (\delta')^4, \quad \kappa = (\delta')^{3\Delta}.$$

We fix a threshold $n_0 \geq 1/\delta'$ and will always assume that $\varepsilon \leq \varepsilon_0(\delta, \Delta, r, n_0)$. For easier reading, we mostly use the letter x for vertices of H, and the letter v for vertices of the host graph G.

Given an embedding of H into $\mathbf{R}$, it defines an *assignment*

$$\psi : V(H) \to \{V_1, V_2, \ldots, V_r\},$$

and we want to find an *embedding* φ of H into G compatible with ψ, that is, a map

$$\varphi : V(H) \to V(G), \quad \varphi \text{ is one-to-one}$$

such that $\varphi(x) \in \psi(x)$ for all $x \in V(H)$, and $\{x,y\} \in E(H)$ implies $\{\varphi(x), \varphi(y)\} \in E(G)$. We will call the sets V_i *clusters*, and write $X_i = \psi^{-1}(V_i)$ for $i = 1, 2, \dots, r$.

2.1. Preprocessing

Before we start the **RGA** algorithm, we order the vertices of H into a sequence $S = (x_1, x_2, \dots, x_n)$ which is more or less, but not exactly, the order in which the vertices will be embedded. Let $m_i = \lfloor \delta' n_i \rfloor$ and $m = \sum m_i$. For each i, choose a set of m_i vertices in X_i such that any two of these altogether m vertices are at a distance at least four in H. The only restriction we make is that no vertex of X_i should be chosen that is a neighbour of any vertex in any other X_j of size less than $n_i/(2r\Delta)$. (This only excludes less than half the n_i vertices of X_i, and its relevance can only be seen later.) Selecting such a subset of m_i vertices is possible, for H is a bounded degree graph.

These vertices $b_1, \dots, b_m$ will be called the *buffer vertices*, and they will be the last vertices in S. The sequence S starts with the neighbourhoods $N_H(b_1), N_H(b_2), \dots, N_H(b_m)$. The length of this initial segment of S will be denoted by T_0. Thus $T_0 = \sum_{i=1}^{m} |N_H(b_i)| \le \Delta m$. If M_i is the number of vertices of X_i in this initial segment then it is easy to see that $M_i < 2r^2 \Delta^2 \delta' n_i$.

The middle of S is an arbitrary ordering of the remaining vertices of H.

(When certain images are *a priori* restricted as in Remark 13 in Section 3, we also list the restricted vertices at the beginning of S right after the neighbours of the buffer vertices.)

2.2. The RGA algorithm

At time 0, set $q(0) = Q(0) = \emptyset$, and $C_{0,x} = V_{0,x} = \psi(x)$ for all $x \in V(H)$.

Phase 1.

For $t \ge 1$, **repeat** the following two steps.

Step 1 — Extending the embedding: Select the vertex $s(t)$ of H to be embedded at time t in the following way.

1. If $q(t-1) \ne \emptyset$ then $s(t)$ is the first element of $q(t-1)$.
2. If $q(t-1) = \emptyset$ then $s(t)$ is the first unembedded vertex in the sequence S.

Now assume that $s(t)=x$ was selected. Consider those vertices $v\in V_{t-1,x}$ for which

$$\deg_G(v, V_{t-1,y}) > \delta|V_{t-1,y}| \tag{2}$$

holds for all y for which $\{x,y\}\in E(H)$ and y is still unembedded. Pick one such v randomly (uniformly) as the image $\varphi(x)$.

Step 2 — Updating: For each unembedded vertex y, set

$$C_{t,y} = \begin{cases} C_{t-1,y}\cap N_G(\varphi(x)) & \text{if } \{x,y\}\in E(H) \\ C_{t-1,y} & \text{otherwise,} \end{cases}$$

and

$$V_{t,y} = \begin{cases} V_{t-1,y}\cap N_G(\varphi(x)) & \text{if } \{x,y\}\in E(H) \\ V_{t-1,y}\setminus\{\varphi(x)\} & \text{otherwise.} \end{cases}$$

(That is, $V_{t,y}=C_{t,y}\setminus\{\varphi(s(u)):u\le t\}$.) Define the leftover sets

$$L_t(V_i) = V_i\setminus\{\varphi(s(u)) : u\le t\}.$$

To obtain $q(t)$, remove from $q(t-1)$ the just embedded x (if $q(t-1)\neq\emptyset$), and append to $q(t-1)$ all unembedded vertices $y\notin q(t-1)$ such that $|V_{t,y}|\le\delta''|L_t(\psi(y))|$. Set $Q(t)=Q(t-1)\cup q(t)$, $q_i(t)=q(t)\cap X_i$, and $Q_i(t)=Q(t)\cap X_i$.

If for some i, $|Q_i(t)|>\delta'''m_i$, then set $T=t$ and **halt with failure.**

Else, if there are no more unembedded non-buffer vertices left, then set $T=t$ and go to Phase 2, **otherwise** set $t\leftarrow t+1$ and go back to Step 1.

End of Phase 1

Phase 2.

Find a system of distinct representatives for the sets $V_{T,y}$ for all unembedded y and **halt with success.** If there is no such set of distinct representatives then **halt with failure.**

End of RGA

We claim that this algorithm finds a good embedding of H into G with high probability. Recall that we fixed an n_0 (arbitrarily) and assumed that $\varepsilon\le\varepsilon_0(\delta,\Delta,r,n_0)$.

Claim 2. *There are positive constants K,λ, depending on δ,Δ,r but not on ε and the n_i, such that Phase 1 succeeds with a probability exceeding $1-\sum_i(K\varepsilon)^{\lambda n_i}$, and Phase 2 succeeds with a probability exceeding $1-\left(\sum_{n_i\ge n_0}e^{-\lambda n_i}+K\varepsilon\right)$.*

Claim 2 clearly proves the Blow-up Lemma, since by choosing a large enough n_0 and then a small enough ε guarantees that the failure probability for the RGA algorithm is very small. In particular, the algorithm succeeds with a positive probability.

2.3. Some simple facts

We start with the most important property of super-regular (and regular) pairs.

Fact 3.. *Let (A,B) be an (ε,δ)-super-regular pair. Then for any $Y\subset B$, $|Y|>\varepsilon|B|$ we have*

$$\#\{x\in A:\deg(x,Y)\le\delta|Y|\}\le\varepsilon|A|.$$

Now let $T(x)$ denote the time in which x is embedded in Phase 1, if x is not embedded in Phase 1 then write $T(x)=T+1$. While Claim 2 says that the (combined) queues $Q_i(T)$ at the end of Phase 1 are small (that is, Phase 1 concludes with success) *with probability close to 1*, the following trivial fact, that follows from the halting rule of the algorithm, says that the (combined) queues one step earlier are all *necessarily* small.

Fact 4. *For all i,*

$$|Q_i(T-1)|\le\delta'''m_i.$$

This, together with the fact that the only buffer vertices that may ever get used up in Phase 1 are the ones that ever get into the queue, implies the next fact. Let us write A_i for the set of vertices of V_i unembedded in Phase 1, and A_+ for the set of all unembedded vertices:

$$A_i=L_T(V_i),\quad\text{and}\quad A_+=\bigcup_i A_i. \tag{3}$$

Fact 5. *For all i and $t\le T$,*

$$|L_t(V_i)|\ge|A_i|\ge m_i-|Q_i(T-1)|\ge(1-\delta''')m_i. \tag{4}$$

Hence, if $m_i\ge 1$ then $|A_i|>\delta' n_i/3$.

Fact 6. *If $x\in Q(T)$ then $|A_+\cap C_{T(x)-1,x}|\le\delta''|\psi(x)|$.*

Indeed, $x\in Q(T)$ means that $x\in q(t)$ for some $t\le T(x)-1$. Hence $|V_{T(x)-1,x}|\le\delta''|\psi(x)|$, and then $|A_+\cap V_{T(x)-1,x}|\le\delta''|\psi(x)|$. It remains to note that $A_+\cap V_{T(x)-1,x}=A_+\cap C_{T(x)-1,x}$ since $V_{t,x}=C_{t,x}\cap L_t(\psi(x))$ for all $t\le T(x)$.

The following two observations will be used repeatedly. They trivially follow from the description of the RGA algorithm.

Fact 7. *For all x and all $t\le T(x)-1$,*

$$|C_{t,x}|>\delta|C_{t-1,x}|\quad\textit{and}\quad|V_{t,x}|>\delta|V_{t-1,x}|.$$

Let us write $\nu(t,x)$ for the number of neighbours (in H) of x embedded by time t:

$$\nu(t,x)=\#\{u\le t:\{x,s(u)\}\in E(H)\}. \tag{5}$$

Fact 8. *Let $x \in V(H)$ and $t \le T(x) - 1$ be arbitrary. Then*

$$|C_{t,x}| > \delta^{\nu(t,x)}|\psi(x)| \ge \delta^{\Delta}|\psi(x)|.$$

The next important lemma says that during Phase 1 the host-sets $V_{t,x}$ cannot get too small. (This is just a formal property of the algorithm, and it does not reflect successful performance. Poor performance of the algorithm is manifested by a premature termination of Phase 1.)

Lemma 9. *For all i and $x \in X_i$,*

$$|V_{T(x)-1,x}| > \delta^{\Delta}\delta''|A_i| - |Q_i(T-1)|.$$

Hence, whenever the RGA algorithm selects a vertex from a cluster V_i of size $n_i \ge 1/\delta'$ then it selects randomly from a set of size greater than γn_i.

Proof. Write $t = T(x) - 1$ and assume first that $x \notin q(t)$. Hence $x \notin Q_i(t)$. Thus $|V_{t,x}| > \delta''|L_t(V_i)|$. If, on the other hand, $x \in q(t)$ then let $t' = \min\{u : x \in q(u)\}$ be the time when x got into the queue. By the definition of the queue,

$$|V_{t'-1,x}| > \delta''|L_{t'-1}(V_i)| \ge \delta''|A_i|.$$

If $t' = t$ then by Fact 7 we have

$$|V_{t,x}| > \delta|V_{t-1,x}| = \delta|V_{t'-1,x}| > \delta\delta''|A_i|.$$

If $t' < t$ then $t' \le T - 1$, and we have

$$|V_{t,x}| > \delta^{\Delta}|V_{t'-1,x}| - |V_i \cap \{\varphi(s(u)) : t' \le u \le t\}| \tag{6}$$

$$> \delta^{\Delta}\delta''|A_i| - |q_i(t')| \ge \delta^{\Delta}\delta''|A_i| - |Q_i(T-1)|, \tag{7}$$

since x got into the queue at time t', is still in the queue at time t, and as long as the queue is not totally empty only the elements of $q(t')$ can be selected for embedding.

This proves the inequality. To prove the last claim of the lemma, note that at most $\Delta\varepsilon n_i$ vertices violate (2), and hence the algorithm selects from a set of size greater than $\delta^{\Delta}\delta''|A_i| - |Q_i(T-1)| - \Delta\varepsilon n_i > (\delta^{\Delta}\delta''\delta'/3 - \delta'''\delta' - \Delta\varepsilon)n_i \ge \gamma n_i$ (provided $m_i \ge 1$ and ε is small enough). ∎

2.4. The main lemma

The sets $C_{t,x}$ are easy to handle (see Fact 8), but the sets $V_{t,x}$ are hard to control. We estimate the probability that too many host-sets $V_{t,x} \subset V_i$ get small by observing that in this case the "trace" of these x's on the eventual leftover set of vertices — that is, the intersections of the sets $C_{t,x}$ with $L_T(V_i)$ — would be small, which would have a negligible probability if we could somehow fix the leftover set $L_T(V_i)$. Thus, the following lemma will play a key role in the proof. Its proof is postponed to the end of Section 2.

Lemma 10 (Main Lemma). *We are given integers p and i, a set $A \subset V_i$, $|A| > \delta^{-\Delta}\varepsilon n_i$, and a set $X \subset X_i$ of p vertices any two at a distance (in H) greater than two. Then,*

$$P\left(|A \cap C_{T(x)-1,x}| < \delta^{\Delta}|A| \quad \text{for all} \;\; x \in X\right) \le (\Delta\varepsilon/\gamma)^p.$$

2.5. Small clusters

Assume $\varepsilon < 1/n_0$, and let V_i be a "small " cluster, that is, a cluster of size $n_i < n_0$. The RGA algorithm gets stuck on V_i if either $|Q_i(T)| > \delta''' m_i$ (stuck in Phase 1), or we cannot find a system of distinct representatives from the leftover sets $V_{T,x}$, $x \in X_i \setminus \{s(u) : u \le T\}$ (stuck in Phase 2). Neither one of these two possibilities will occur if $C_{t,x} = V_i$ for all $x \in X_i$ and all $t \le T(x) - 1$ (which implies $Q_i(T) = \emptyset$). This, in turn, will certainly happen if each element in the neighbourhood $N^{\cup}(X_i) = \bigcup_{x \in X_i} N_H(x)$ is selected from the common neighbourhood $N^{\cap}(V_i) = \bigcap_{v \in V_i} N_G(v)$. Hence the probability of failure on V_i is upper bounded, according to Lemma 9, by

$$\sum_{y \in N^{\cup}(X_i)} \frac{|\psi(y) \setminus N^{\cap}(V_i)|}{\gamma|\psi(y)|}.$$

(This was not a perfect deduction, for Lemma 9 makes the assumption $n_j \ge 1/\delta'$ about the sizes of the neighbour-sets $V_j = \psi(y)$. This, however, does not make the obtained upper bound invalid, for the contribution to the above probability of a neighbour-cluster V_j of size less than $1/\delta'$ is 0, since the conditions $n_i\varepsilon < 1$, $n_j\varepsilon < 1$ imply that the super-regular pair $\{V_i, V_j\}$ is in fact a complete bipartite graph.)

Now the regularity condition implies $\deg(v, V_j) \ge (1-\varepsilon)|V_j|$ for all $v \in V_i$ and all neighbour clusters V_j. Hence

$$\frac{|\psi(y) \setminus N^{\cap}(V_i)|}{\gamma|\psi(y)|} \le \frac{\Delta\varepsilon}{\gamma}.$$

Since also $|N^{\cup}(X_i)| \le \Delta |X_i| < \Delta n_0$, the probability of failure on V_i is less than $\Delta^2 n_0 \varepsilon/\gamma$, which can be made arbitrarily small by choosing a small enough ε. We proved the following lemma.

Lemma 11. *Let $n_i = |V_i| < n_0$, and assume $\varepsilon < 1/n_0 \le \delta'$. The probability that the RGA algorithm fails on V_i is less than $\Delta^2 n_0 \varepsilon/\gamma$.*

2.6. Analysis of Phase 1

The following lemma immediately implies that Phase 1 succeeds with probability $1 - o(1)$.

Lemma 12. *Let $m_i \ge 1$ (that is, $n_i \ge 1/\delta'$). Then, for any integer p,*

$$P(|Q_i(T)| \ge p) < 4^{n_i} (\Delta\varepsilon/\gamma)^{p/(\Delta^2+1)}.$$

Indeed, let us apply Lemma 12 with $p = \lceil \delta''' m_i \rceil$. We get the estimate (using Lemma 11)

P(Phase 1 of the RGA algorithm ends with failure)

$$= P(|Q_i(T)| > \delta''' m_i \ \text{ for some } i)$$

$$< \Delta^2 n_0 \varepsilon/\gamma \sum_{n_i < 1/\delta'} 1 + \sum_{n_i \ge 1/\delta'} 4^{n_i} (\Delta\varepsilon/\gamma)^{p/(\Delta^2+1)} < \sum_i (K\varepsilon)^{\lambda n_i}. \tag{8}$$

For the analysis of Phase 1, it remains to prove Lemma 12.

Proof of Lemma 12. Write $p' = \lceil p/(\Delta^2+1) \rceil$. Given a set $X \subset X_i$, $|X| = p$, let us estimate the probability $P(X \subset Q(T))$. Use the greedy method to select a subset $X' \subset X$ of p' vertices such that any two vertices of X' are at a distance more than two. Using (4) and Fact 6 we see that the probability $P(X' \subset Q(T))$ is upper bounded by the probability of the event that there exists an $A \subset V_i$, $|A| > m_i/2 > \delta^{-\Delta} \varepsilon n_i$ such that, for all $x \in X'$,

$$|A \cap C_{T(x)-1,x}| \le \delta'' n_i < \frac{2\delta''}{\delta'} m_i \le \frac{\delta^{\Delta}}{2} m_i < \delta^{\Delta} |A|.$$

For a given A this probability is at most $(\Delta\varepsilon/\gamma)^{p'}$ (by the Main Lemma). The number of choices for X and A is less than 4^{n_i}; hence Lemma 12. ∎

2.7. Analysis of Phase 2

Let B be the set of unembedded vertices at the beginning of Phase 2, and let $B_i = B \cap X_i$. Note that all these vertices are buffer vertices. (On the other hand, while most buffer vertices are unembedded since they are at the end of the sequence S, not necessarily are all unembedded, for some may have got into the queue and got embedded.) For such a vertex x, write $V(x) = V_{T,x}$ for the host-set at the beginning of Phase 2. Hence $V(x) \subset A_i$ for all $x \in B_i$, and $|B_i| = |A_i|$ (the sets A_i were defined in (3)). Let $\kappa \leq \gamma$.

For each i, we would like to select a system of distinct representatives from the sets $\{V(x) : x \in B_i\}$. The König-Hall condition for the existence of such representatives clearly follows from the following three conditions. For all i,

(9) $|V(x)| > \kappa|A_i|$ for all $x \in B_i$,

(10) $\left|\bigcup_{x\in X} V(x)\right| > (1-\kappa)|A_i|$ for all subsets $X \subset B_i$, $|X| > \kappa|A_i|$,

(11) $\bigcup_{x\in X} V(x) = A_i$ for all subsets $X \subset B_i$, $|X| > (1-\kappa)|A_i|$.

Because of Lemma 11, we can restrict our attention to clusters V_i of size $n_i \geq n_0 (\geq 1/\delta')$.

Proof of (9). For all $x \in B_i$ we have $V(x) = V_{T,x} = V_{T(x)-1,x}$, and thus Lemma 9 implies that $|V(x)| > \gamma n_i \geq \kappa n_i$ for all $x \in B_i$, proving (9). Note that this always holds, not only with large probability.

Proof of (10). To prove that (10) holds with a large probability, let i be fixed and estimate, for fixed sets $X \subset B_i$, $|X| = p > \kappa|A_i|$, and $A' \subset A_i$, $|A'| \geq \kappa|A_i|$, the probability of the event that

$$A' \cap V(x) = \emptyset \text{ for all } x \in X.$$

Since $|A_i| > \delta' n_i/3$ by Fact 5, we have (for small enough ε) $|A'| > \delta^{-\Delta}\varepsilon n_i$ and we can apply the Main Lemma. Hence the probability in question (given successful completion of Phase 1) is at most $(\Delta\varepsilon/\gamma)^p < (\Delta\varepsilon/\gamma)^{\kappa\delta' n_i/3}$. (Note that we used here again the identity $A_+ \cap V_{T,x} = A_+ \cap C_{T,x}$ for all $x \in B$.) Thus, the probability that (10) fails for some i is at most

$$(12) \qquad \Delta^2 n_0 \varepsilon/\gamma \sum_{n_i < 1/\delta'} 1 + \sum_{n_i \geq 1/\delta'} 4^{n_i} (\Delta\varepsilon/\gamma)^{\kappa\delta' n_i/3} < \sum_i (K\varepsilon)^{\lambda n_i}.$$

Proof of (11). We want to prove that, with a large probability, every vertex $a \in A_i$ belongs to many sets $V(x)$, $x \in B_i$. In fact, we will show this for every vertex $a \in V_i$. For this purpose, let us fix i, a set $X \subset B_i$, $|X| \geq (1-\kappa)|B_i|$, and a vertex $a \in V_i$, and let us estimate the probability that for all $x \in X$, $a \notin V(x)$.

The event $a \in V(x)$ is equivalent to $\varphi(N_H(x)) \subset N_G(a)$. Let us fix $x \in X$, and denote its neighbourhood by $Z = \{z_1, \ldots, z_\ell\}$. Since a is a buffer vertex, its neighbours, the z's, are all embedded in the first T_0 steps of the **RGA** algorithm. If a vertex $z \in Z \cap X_j$ is still unembedded at time t, (that is, $t \le T(z)-1$) then, since the total number of vertices embedded into V_j in the first T_0 steps is $M_j < 2r^2\Delta^2\delta' n_j$, we have $|V_{t,z}| \ge |C_{t,z}| - M_j$, whence, by Fact 8,

$$|V_{t,z}| > \delta^\Delta n_j - 2r^2\Delta^2\delta' n_j \ge \delta^\Delta n_j/2.$$

Consequently, all queues are empty during this first stage of the algorithm, and so *the first T_0 points are embedded in their original order.*

As we embed the vertices $z_1, z_2, \ldots$, we would like to see if they get into $W = N_G(a)$. But there may be some edges between them, so it is a good idea to make sure that their degrees into W are also large. The easiest way to ensure this would be to run the RGA algorithm with initial sets $W_z = \psi(z) \cap W$ rather than the full $\psi(z)$ (restriction of the process to W). While such a modification of the RGA algorithm may work for a specific a, the number of a is n_i (even for a fixed i), and we cannot modify the RGA to work simultaneously for all $a \in V_i$. However we will show that this happens with a *positive probability.* In other words, we show that, with a positive probability, the original RGA algorithm has the following three additional properties for all $z \in Z$.

$$\varphi(z) \in W. \tag{13}$$

$$\deg_G(\varphi(z), V_{T(z)-1,y} \cap W) > \delta |V_{t-1,y} \cap W| \tag{14}$$

for all $y \in Z$ embedded after z (that is, (2) holds with the sets restricted to W).

And finally, if $z \in Z$ is unembedded at time $t \le T_0$, then

$$|C_{t,z} \cap W_z| > \delta^{\nu(t,z)} |W_z| \ge \delta^\Delta |W_z| \ge \delta^{\Delta+1} |\psi(z)|, \tag{15}$$

where ν was defined in (5). (That is, Fact 8 holds with the sets restricted to W.) Note that (15) and $z \in X_j$ would imply

$$|V_{t,z} \cap W| \ge \delta^{\Delta+1} n_j - M_j > (\delta^{\Delta+1} - 2r^2\Delta^2\delta') n_j > \delta^{\Delta+1} n_j/2.$$

Indeed, we have these properties preserved inductively if only at time t when a $z \in Z$ is embedded we happen to choose z from the set $V_{t-1,z} \cap W$ of size greater than $\delta^{\Delta+1}|\psi(z)|/2$ — unless we choose such a z in a way that the degree condition (14) is violated. Given an arbitrary past, the (conditional) probability that z is indeed selected from this small set is greater than $\delta^{\Delta+1}/2$, and the probability that the selection violates the degree condition is, by the regularity condition, less than $\Delta\varepsilon/\gamma$. Hence the (conditional) probability that conditions (13), (14) and (15) will all be satisfied for all $z \in Z$ (conditioned under an arbitrary past) is greater than $\mu = (\delta^{\Delta+1}/2 - \Delta\varepsilon/\gamma)^\Delta$, since conditional probabilities always multiply.

Also, the neighbourhoods $N(x)$ are disjoint for different $x \in X$, for the points $x \in X$ were at least four apart in H. Hence the above lower bounds on the conditional probabilities of success are valid even when conditioned under the full past (when at time t we start embedding the neighbourhood $Z = N(x)$ for a new x, we start fresh with $C_{t-1,z} = \psi(z)$ for all $z \in Z$). Hence we can multiply the upper bounds $1-\mu$ on the conditional failure probabilities and get that the probability that none of the events $\varphi(N_H(x)) \subset N_G(a)$ will occur is less than

$$(1-\mu)^{|X|} < (1-\mu)^{\delta' n_i/4}.$$

Finally, writing $h(x) = -[x \log x + (1-x)\log(1-x)]$, the probability that (11) fails on V_i (with $|V_i| \geq 1/\delta'$) is less than

$$n_i \binom{n_i}{\kappa n_i} (1-\mu)^{\delta' n_i/4} < n_i \exp\{(h(\kappa) - \mu\delta'/4)n_i\} \leq n_i \exp\{-\mu\delta' n_i/8\}.$$

Thus, the probability that (11) fails on **some** V_i is less than

$$\sum_{n_i \geq n_0} e^{-\lambda n_i} + K\varepsilon. \qquad \blacksquare$$

2.8. Proof of the Main Lemma

Let $E(x)$ denote the event $\{|A \cap C_{T(x)-1,x}| < \delta^{\Delta}|A|\}$, and write $E(X) = \cap_{x\in X} E(x)$. We want to get the estimate $P(E(X)) \leq (\Delta\varepsilon/\gamma)^p$ (recall that $p = |X|$). Let $E(t,x)$ be the event $\{|A \cap C_{t,x}| < \delta^{\nu(t,x)}|A|\}$, and define the random instance

$$\tau(x) \doteq \inf\{t : E(t,x)\}.$$

If there is such an instance then write $k(x) = s(\tau(x))$ for the vertex that "killed" x. Note that $k(x)$ is a neighbour of x in H and the map $k(x)$ is one-to-one (and hence so is τ).

Now let us fix a map $k_0 : X \to V(H)$, $k_0(x) \in N_H(x)$ for all $x \in X$. We will show that

$$P(E(X) \cap \{k = k_0 \text{ on } X\}) \leq (\varepsilon/\gamma)^p.$$

This will prove the Main Lemma, for the number of such maps k_0 is at most Δ^p.

We embed the vertex $y = k_0(x)$ at time $\tau(x)$, but we already know at time $t = \tau(x) - 1$ that y will be embedded in the next step. At that time the event $E(t,x)$ is still false, that is, we still have $\{|A \cap C_{t,x}| \geq \delta^{\nu(t,x)}|A|\}$, and hence the set $Y = A \cap C_{t,x}$ is of size greater than $\varepsilon|\psi(x)|$. In the next step y will be embedded into a vertex v with $\deg_G(v, Y) < \delta|Y|$. But there are at most $\varepsilon|\psi(y)|$ such vertices out of the more than $\gamma|\psi(y)|$ vertices v is to be selected from. Thus the (conditional) probability (given the full past until time $\tau(x)-1$) that such a v is selected is less than ε/γ. Hence, the probability that this happens for every $x \in X$ is less than $(\varepsilon/\gamma)^p$ as claimed. ∎

Remark. A formally simpler way to present the above proof would be to use induction on n to prove the inequality

$$P\left(|A \cap C_{T(x)-1,x}| < \delta^{\Delta}|A| \quad \text{for all } \ x \in X\right) \leq \prod_{x \in X} (\deg(x)\varepsilon/\gamma).$$

(Simply check if $s(1)$ is the neighbour of any vertex in X or not, and apply induction.)

3. Concluding remarks

Remark 13. *When using the Blow-up Lemma, we usually need the following strengthened version: Given $c>0$, there are positive numbers $\varepsilon=\varepsilon(\delta,\Delta,r,c)$ and $\alpha=\alpha(\delta,\Delta,r,c)$ such that the Blow-up Lemma in the equal size case (all $|V_i|$ are the same) remains true if for every i there are certain vertices x to be embedded into V_i whose images are a priori restricted to certain sets $C_x \subset V_i$ provided that*

(i) *each C_x within a V_i is of size at least $c|V_i|$,*

(ii) *the number of such restrictions within a V_i is not more than $\alpha|V_i|$.*

References

[1] N. Alon, R. Duke, H. Leffman, V. Rödl, and R. Yuster: Algorithmic aspects of the regularity lemma, *FOCS*, **33** (1993), 479–481; Journal of Algorithms, **16** (1994), 80–109.

[2] N. Alon and E. Fischer: 2-factors in dense graphs, Discrete Math., to appear.

[3] N. Alon, S. Friedland, G. Kalai: Regular subgraphs of almost regular graphs, *J. Combinatorial Theory*, B **37** (1984), 79–91. See also N. Alon, S. Friedland, G. Kalai: Every 4-regular graph plus an edge contains a 3-regular subgraph, *J. Combinatorial Theory*, B **37** (1984), 91–92.

[4] N. Alon and R. Yuster: Almost H-factors in dense graphs, *Graphs and Combinatorics*, **8** (1992), 95–102.

[5] N. Alon and R. Yuster: H-factors in dense graphs, *J. Combinatorial Theory*, Ser. B, to appear.

[6] V. Chvátal, V. Rödl, E. Szemerédi, and W. T. Trotter Jr.: The Ramsey number of a graph with bounded maximum degree, *Journal of Combinatorial Theory*, B **34** (1983), 239–243.

[7] K. Corrádi and A. Hajnal: On the maximal number of independent circuits in a graph, *Acta Math. Acad. Sci. Hung.*, **14** (1963), 423–439.

[8] T. Kővári, Vera T. Sós, and P. Turán: On a problem of Zarankiewicz, *Colloq. Math.*, **3** (1954), 50–57.

[9] P. Erdős and A. H. Stone: On the structure of linear graphs, *Bull. Amer. Math. Soc.*, **52** (1946), 1089–1091.

[10] A. Hajnal and E. Szemerédi: Proof of a conjecture of Erdős, *Combinatorial Theory and its Applications* vol. II (P. Erdős, A. Rényi and V. T. Sós eds.), Colloq. Math. Soc. J. Bolyai 4, North-Holland, Amsterdam (1970), 601–623.

[11] J. Komlós, G. N. Sárközy, and E. Szemerédi: Proof of a packing conjecture of Bollobás, AMS Conference on Discrete Mathematics, DeKalb, Illinois (1993), *Combinatorics, Probability and Computing*, **4** (1995), 241–255.

[12] J. Komlós, G. N. Sárközy, and E. Szemerédi: On the Pósa-Seymour conjecture, submitted to the *Journal of Graph Theory.*

[13] J. Komlós, G. N. Sárközy, and E. Szemerédi: On the square of a Hamiltonian cycle in dense graphs, Proceedings of Atlanta'95, *Random Structures and Algorithms*, **9** (1996), 193–211.

[14] J. Komlós, G. N. Sárközy, and E. Szemerédi: Proof of the Alon-Yuster conjecture, in preparation.

[15] J. Komlós and M. Simonovits: Szemerédi's Regularity Lemma and its applications in graph theory, Bolyai Society Mathematical Studies 2, *Combinatorics, Paul Erdős is Eighty* (Volume 2), (D. Miklós, V. T. Sós, T. Szőnyi eds.), Keszthely (Hungary) (1993), Budapest (1996), 295–352.

[16] E. Szemerédi: Regular partitions of graphs, Colloques Internationaux C.N.R.S. N$^{\underline{o}}$ 260 — *Problèmes Combinatoires et Théorie des Graphes*, Orsay (1976), 399–401.

János Komlós

Department of Mathematics
Rutgers University
New Brunswick, NJ 08903
`komlos@math.rutgers.edu`

Gábor N. Sárközy

Department of Computer Science
Rutgers University
New Brunswick, NJ 08903
`gsarkozy@cs.rutgers.edu`

Endre Szemerédi

Hungarian Academy of Sciences
Budapest, Hungary
`szemered@cs.rutgers.edu`