1 Introduction

Recall that given two graphs K, H, their *Ramsey number* $r(K, H)$ is defined as the minimum N such that whenever we color the edges of K_N red and blue, we either see a red copy of K or a blue copy of H. The fact that these numbers exist (i.e. are finite) is (a simple corollary of) Ramsey's theorem, and the main question in graph Ramsey theory is to obtain good estimates on the number $r(K, H)$ for K, H in various natural classes of graphs.

One of the simplest ways to lower-bound $r(K, H)$ is the following construction, usually attributed to Chvátal and Harary. Suppose that H is connected. We take $\chi(K) - 1$ blue cliques, each of size $\mathsf{v}(H) - 1$, and connect all these blue cliques by red edges. In other words, the blue graph is a disjoint union of $\chi(K) - 1$ copies of $K_{\mathsf{v}(H)-1}$, and the red graph is a complete $(\chi(K) - 1)$-partite graph. Note that the blue graph cannot contain any copy of H, since H is connected and has $\mathsf{v}(H)$ vertices, and thus cannot fit in any connected component of the blue graph. Moreover, the red graph can contain no copy of K, by the definition of $\chi(K)$. This yields a coloring on $(\chi(K) - 1)(\mathsf{v}(H) - 1)$ vertices with no red K or blue H, which implies that

$$r(K, H) \geq (\chi(K) - 1)(\mathsf{v}(H) - 1) + 1. \tag{1}$$

In general, this bound is hopelessly bad. For instance, if $K = H = K_k$, it yields a bound of $r(K_k, K_k) \geq (k - 1)^2 + 1$, while the correct behavior is known to be exponential in k. More generally, if K and H are any graphs of density at least ε on k vertices, then a random coloring shows that $r(K, H) \geq 2^{c_\varepsilon k}$ for some $c_\varepsilon > 0$, which is far larger than the at-most-quadratic bound of (1).

However, for sparse graphs, the bound (1) isn't so bad. For instance, it is a famous result of Chvátal, Rödl, Szemerédi, and Trotter that if H is a graph with maximum degree at most Δ, then $r(H, H) \leq C_\Delta \mathsf{v}(H)$ for some constant C_Δ independent of H. Since we will also have $\chi(H) \leq \Delta + 1$, we see that this behavior of $r(H, H)$ is similar to the lower bound (1), in that it is linear in $\mathsf{v}(H)$ with a constant independent of H. A major generalization of the Chvátal–Rödl–Szemerédi–Trotter theorem was a recent breakthrough by Lee, who proved that $r(H, H) \leq C_d \mathsf{v}(H)$ for every d-degenerate graph H. Recall that H is said to be d-degenerate if every subgraph of H contains a vertex of degree at most d. Thus, being d-degenerate is a much more general property than having maximum degree at most Δ, and Lee's result essentially says that for all "sparse" H, the Ramsey number $r(H, H)$ is linear in $\mathsf{v}(H)$.

However, one can ask a more refined question: when is (1) *exactly* tight? While one can ask this for arbitrary choices of K and H, the most well-studied version is when we take K to be a fixed clique K_k. In this case, Burr defined H to be *k-good* if $r(K_k, H) = (k - 1)(\mathsf{v}(H) - 1) + 1$. As it turns out, many sufficiently large (in terms of k) sparse graphs are k-good. For instance, the earliest result in this direction is due to Chvátal, who showed that every tree is k-good for all k. Burr and Erdős began systematically studying k-good graphs, and proved a number of results in this direction, e.g. showing that all sufficiently long cycles are k-good for all k. Simiarly to the above, they conjectured that all sufficiently large H with bounded maximum degree are k-good for all k. However, this conjecture was disproven by Brandt.

Theorem 1.1 (Brandt 1996). *For any $k \geq 3$ and any Δ sufficiently large, a random Δ-regular graph is with high probability not k-good.*

In fact, Brandt showed something even more general, namely that any sufficiently good expander is not k-good. Since a random regular graph is a good expander, this implies the result above. To prove this result in the case $k = 3$, consider a coloring where the red graph is a blow-up of C_5 where each part has $(\frac{1}{2} - \delta)n$ vertices. The red graph is clearly triangle-free. Moreover, if H is a subgraph of the blue graph with n vertices, then H has at least δn vertices in two consecutive parts of the C_5. Since the edges between these parts are all red, this implies that H has two vertex subsets of size δn with no edges between them, contradicting its expansion properties. This shows that $r(K_3, H) \geq (\frac{5}{2} - 5\delta)n > 2(n-1) + 1$ for a sufficiently good expander H on n vertices.

The main result that I will be discussing is a theorem of Nikiforov and Rousseau, which can be seen as a sort of converse to Brandt's theorem; roughly speaking, it says that all large graphs that are *not* good expanders (in an appropriate sense) are k-good for any k. To state it, we will need the following definition.

Definition 1.2. Let H be an n-vertex graph, and let $\gamma, \eta > 0$ be real numbers. We say that H is *(γ, η)-splittable* if there is a set $S \subseteq V(H)$ with $|S| \leq n^{1-\gamma}$ such that every connected component of $H \setminus S$ has at most ηn vertices.

Thus, H is splittable if it has a small separator, namely a small vertex subset whose removal splits the graph into many small components. In particular, note that being splittable is roughly the same as being a bad expander, because in a good expander all large vertex sets will have edges between them, and in particular it is impossible to disconnect the graph by removing a small number of vertices.

Theorem 1.3 (Nikiforov–Rousseau 2007). *For every $k \geq 2, d \geq 1$, and $0 < \gamma < 1$, there exist $\eta > 0$ and $n_0 \in \mathbb{N}$ such that the following holds. If H is a d-degenerate (γ, η)-splittable graph on $n \geq n_0$ vertices, then H is k-good.*

This result is enormously general, and in fact I am only stating a fairly special case of Nikiforov and Rousseau's actual theorem. As simple corollaries of this theorem, Nikiforov and Rousseau were able to resolve all but one[1] open question about Ramsey goodness that had been asked by Burr and Erdős. It is hard to overstate just how powerful this theorem is.

Example. One simple example of a graph H which satisfies the assumptions of Theorem 1.3 is K'_m, the subdivision of K_m. This graph has $m + \binom{m}{2}$ vertices, and is 2-degenerate because it is a bipartite graph where one side is 2-regular. Moreover, if we delete the m

[1] This last question, namely that hypercubes are k-good for all k, was ultimately resolved by Fiz Pontiveros, Griffiths, Morris, Saxton, and Skokan.

original vertices of K_m, then K'_m becomes a bunch of isolated vertices, implying that it has a separator of size $m = O(\sqrt{\mathsf{v}(K'_m)})$. Burr and Erdős had conjectured that K'_m is k-good for all sufficiently large m, and Nikiforov and Rousseau's theorem confirms this conjecture.

An extremely general class of splittable degenerate graphs is the class of planar graphs. Indeed, Euler's identity implies that every planar graph is 5-degenerate, and the famous Lipton–Tarjan separator theorem implies that every sufficiently large planar graph is $(1/2, \eta)$-splittable for all η. Thus, every sufficiently large planar graph is k-good for all k. Even more generally, the Alon–Seymour–Thomas separator theorem implies that any sufficiently large graph with some forbidden minor is $(1/2, \eta)$-splittable. Moreover, a result of Mader says that such graphs also have bounded degeneracy. Putting these together, we see that in any non-trivial minor-closed family, all sufficiently large graphs are k-good for all k.

2 Proof sketch and preliminaries

To prove Theorem 1.3, we need to show that if $N = (k-1)(n-1)+1$, then every red/blue edge-coloring of K_N will contain either a red K_k or a blue copy of H, where n is sufficiently large and H is d-degenerate and (γ, η)-splittable. The proof proceeds roughly as follows.

Suppose we are given such a coloring, and we assume that it contains no red K_k. Using Szemerédi's regularity lemma, we partition K_N into a bounded number of vertex sets such that most pairs of parts are ε-regular, for some appropriate ε chosen later. We can also arrange it so that each part is ε-regular with itself. We first argue that each part must be almost monochromatic in blue, for otherwise some part would contain a red K_k by the regularity. Moreover, we can also ensure that the red graph between the parts looks almost $(k-1)$-partite, for otherwise we would find a red K_k spanning some of the parts. We can also argue that this red graph must be quite dense between its "almost color classes".

At this point, the situation already looks pretty similar to the construction we used to derive the lower bound (1), namely we roughly have $k-1$ blocks that are nearly monochromatic blue, while the edges between these parts are almost entirely red. Recall that in deriving (1), we used the fact that H had n vertices and each blue clique had only $n-1$ vertices. So to finish the proof, we must use the fact that $N = (k-1)(n-1)+1$ to be able to fully embed H in the blue graph. If all the blue blocks are really blue cliques, this is fine, since by pigeonhole one of them must have at least n vertices. We can not assume this, but it turns out that one of the parts has n vertices and is almost a blue clique, and then H can be found in it through a careful embedding algorithm.

2.1 Regularity review

Given two vertex sets X, Y in a graph, let $e(X, Y)$ denote the number of pairs in $X \times Y$ that are edges, and let $d(X, Y) = e(X, Y)/|X||Y|$ denote the edge density between X and Y. Given a parameter $\varepsilon > 0$, we say that (X, Y) is *ε-regular* if $|d(X', Y') - d(X, Y)| \leq \varepsilon$ for all $X' \subset X, Y' \subset Y$ with $|X'| \geq \varepsilon|X|, |Y'| \geq |Y|$. Note that in this definition we don't require X and Y to be disjoint, and in particular if the pair (X, X) is ε-regular, then we'll just say that X is ε-regular (or that it's ε-regular with itself).

The version of Szemerédi's regularity lemma that we'll use is the following.

Lemma 2.1 (Regularity lemma). *For every $\varepsilon > 0$, there is an $M = M(\varepsilon) \in \mathbb{N}$ such that every N-vertex graph G has a vertex partition $V(G) = V_0 \sqcup V_1 \cdots \sqcup V_k$ with $m \leq M$ and the following properties.*

- *$|V_0| < \varepsilon N$, and $|V_1| = \cdots = |V_m|$.*
- *Each V_i for $i \geq 1$ is ε-regular with itself.*
- *For each $i \geq 1$, the number of j such that (V_i, V_j) is not ε-regular is at most $(1 - \varepsilon)m$.*

To prove[2] this, we first apply the ordinary version of Szemerédi's regularity lemma with some appropriately chosen $\varepsilon' < \varepsilon$. If the irregular pairs aren't well-distributed (in the sense that some part participates in many irregular pairs), then there can't be too many such bad parts. We can arbitrarily cut them up and distribute their vertices among the other parts, which only makes the regularity slightly worse. Next, we can use a lemma of Conlon and Fox to find an ε'-regular subset of each part. By pulling out such subsets repeatedly, we can almost partition each part into ε'-regular subparts, and then we can again arbitrarily redistribute the remaining vertices without making the regularity much worse.

The main property we'll need of regular pairs is the counting lemma; here is a simple version that will suffice for our purposes.

Lemma 2.2 (Counting lemma). *For every $\delta > 0$, there exists some $\varepsilon > 0$ such that the following holds. If $X_1, \ldots, X_k$ are (not necessarily disjoint) vertex sets with (X_i, X_j) ε-regular and with $d(X_i, X_j) \geq \delta$ for all $i \neq j$, then there is a copy of K_k with one vertex in each X_i.*

2.2 Other results we'll need

Here are a few other relatively standard results that will be used in the proof. The first is a convenient version of the stability result for Turán's theorem.

Theorem 2.3 (Andrásfai–Erdős–Sós, 1974). *For every $k \geq 3$, there exists some $\tau = \tau(k) > 0$ such that if G is a K_k-free graph on n vertices with minimum degree at least $(1 - \frac{1}{k-1} - \tau)n$, then G is $(k-1)$-partite.*

The most well-known version of this is when $k = 3$, which says that a triangle-free graph with minimum degree more than $\frac{2}{5}n$ must be bipartite. This is tight, as shown by a blowup of C_5, which is $(\frac{2}{5}n)$-regular, triangle-free, and not bipartite.

The next useful result will be the following form of the dependent random choice lemma. Dependent random choice is an extremely useful proof technique, which allows one to find "popular" subsets of a graph with many common neighbors, as follows.

[2]This is just a proof sketch, but see [arXiv:2001.00407, Lemma 1] for a full proof (of a nearly identical result).

Theorem 2.4. *For every $s \geq 2, \beta > 0, \lambda > 0$, there exist $\mu > 0$ and $n_0 \in \mathbb{N}$ such that the following holds. Suppose G is a bipartite graph with parts A, B of sizes at least n_0, and assume that $d(A, B) \geq \beta$. Then there is a subset $A' \subset A$ with $|A'| \geq |A|^{1-\lambda}$ such that every s-tuple of vertices in A' has at least $\mu|B|$ common neighbors in B.*

To prove this, we pick a small random subset of B, and define A' to be the common neighborhood of B. Then intuitively, we'd expect subsets of A' to have many common neighbors, basically because the distribution of A' is biased towards sets with many common neighbors. One can show, just by linearity of expectation, that with positive probability we will indeed get the desired result.

3 Proof of Theorem 1.3

Recall that we set $N = (k-1)(n-1)+1$ and we are given a two-coloring of $E(K_N)$, namely a partition of the edges into two graphs R and B. We suppose that R is K_k-free, and wish to find a copy of H in B, where H is a d-degenerate (γ, η)-splittable graph. Throughout this proof, I will be vague about the order various parameters need to be picked and their relative sizes. It is important (and not obvious) that the parameters can all be picked in a consistent order, but I will ignore this for the sake of clarity. In particular, parameters like δ and ε may suddenly appear without being defined or chosen; in every such case, you should trust me (or check!) that one can in fact pick them so that the argument works.

We begin by applying the regularity lemma to the graph R, with some appropriately small parameter ε. We obtain a partition of the vertices into $V_0 \sqcup V_1 \sqcup \cdots \sqcup V_m$, where m is bounded, $|V_0| \leq \varepsilon N$, and $|V_1| = \cdots = |V_m| = t$, for some integer t. We also have that each V_i for $i > 0$ is ε-regular with itself, and that each V_i participates in at most εm irregular pairs.

We first claim that each part V_i with $i \in [m]$ has blue density at least $1-\delta$. For if this were not the case, then some V_i would have red density at least δ, and would be ε-regular with itself. By the counting lemma, this implies that this part V_i contains a red K_k, contradicting our assumption. So $d(B[V_i]) \geq 1 - \delta$ for all $i \in [m]$.

Now we define reduced graphs R^* and B^*, as follows. Both of them will have vertex set $[m]$. The edges of B^* consist of those pairs (i, j) for which $d_B(V_i, V_j) > 1-\delta$, while the edges of R^* consist of those pairs for which (V_i, V_j) is ε-regular and $d_R(V_i, V_j) \geq \delta$. Note that R^* and B^* are edge-disjoint. Moreover, if (i, j) is not an edge in R^* or B^*, then we must have that (V_i, V_j) is not ε-regular. In particular, this means that every vertex in $[m]$ has degree at least $(1-\varepsilon)m$ in $R^* \cup B^*$.

Lemma 3.1. *If some vertex in B^* has degree at least $(1+\alpha)\frac{m}{k-1}$, then there is a copy of H in B.*

Proof sketch. Suppose without loss of generality that this vertex is $1 \in [m] = V(B^*)$. This means that $d_B(V_1, V_j) > 1-\delta$ for at least $(1+\alpha)\frac{m}{k-1}$ choices of j. Note that these $(1+\alpha)\frac{m}{k-1}$ sets collectively contain $(1+\alpha)\frac{m}{k-1}t \geq (1+\alpha')N/(k-1) > (1+\alpha'')n$ vertices, since $N = (k-1)(n-1)+1$. Moreover, by the above, each part V_i has blue density at least $1-\delta$.

So we've found at least n vertices among the parts $V_1, \{V_j\}$, and almost all the edges within these parts and between V_1 and the remaining parts are blue. In particular, it will be very easy to now embed H in B. Specifically, we embed the small separator of H inside V_1, and then can greedily embed the remaining vertices in the parts V_j, using the splittability property to only care about connecting these vertices to each other and to V_1, and using the d-degeneracy assumption to greedily embed the vertices in an appropriate order. Crucially, the fact that essentially all edges we care about are blue means that this is easy—the only thing that could go wrong is not having enough room to fit all of H in, and our assumption says that we have $(1+\alpha'')n$ vertices to work with, which is plenty of wiggle room. $\square$

Therefore, we may assume that $\Delta(B^*) < (1+\alpha)\frac{m}{k-1}$. Since every vertex in $R^* \cup B^*$ has degree at least $(1-\varepsilon)m$, this implies that

$$\delta(R^*) > (1-\varepsilon)m - (1+\alpha)\frac{m}{k-1} > \left(1 - \frac{1}{k-1} - \tau\right)m.$$

Moreover, we claim that R^* is K_k-free. Indeed, if R^* had a K_k, this would yield k parts that are ε-regular between them with red density at least δ, which gives a red K_k in R by the counting lemma. Therefore, by the Andrásfai–Erdős–Sós theorem, R^* is $(k-1)$-partite. Let $Z_1, \ldots, Z_{k-1}$ be the color classes of R^*. Additionally, let $U_1, \ldots, U_{k-1}$ denote the partition of $V(K_N) \setminus V_0$ induced by $Z_1, \ldots, Z_{k-1}$; namely, for each vertex $v \in V(K_N)$, if it lies in some part $V_i \subset Z_\ell$, then we declare $v \in U_\ell$.

Note that each part U_ℓ is dense in blue, because it consists of parts V_i each of which has $d_B > 1-\delta$, and since each Z_ℓ is an independent set in R^*, we also have that $d_B(V_i, V_j) > 1-\delta$ if $i, j \in Z_\ell$. Moreover, we have that $|Z_\ell| \le \Delta(B^*) + 1 \le (1+\alpha)\frac{m}{k-1}$. Moreover, since this holds for all Z_ℓ, and because they form a partition of $[m]$, we can get a corresponding lower bound $|Z_\ell| > (1-k\alpha)\frac{m}{k-1}$. Thus, the partition $U_1, \ldots, U_{k-1}$ is an almost equitable partition of almost all of $V(K_N)$ into $k-1$ parts that are extremely dense in blue; this is essentially the structure we are looking for, and now we simply need to refine this structure some more to eventually find a copy of H.

The first step of this refinement is proving that there are very few blue edges between distinct $U_\ell, U_{\ell'}$.

Lemma 3.2. *Suppose we had that*

$$\sum_{1 \le \ell < \ell' \le k-1} e_B(U_\ell, U_{\ell'}) \ge \beta \binom{N}{2}.$$

Then we can find a blue copy of H.

Proof. By averaging first over the pairs (ℓ, ℓ'), and then over parts $V_i \subset U_\ell$, we can find that there exist V_i and $U_{\ell'}$ such that

$$e_B(V_i, U_{\ell'}) \ge \beta |V_i||U_{\ell'}|.$$

Additionally, recall that V_i is a subset of some U_ℓ, and that $d_B(V_i, U_\ell \setminus V_i) > 1-\delta$. We would really like to mimic the proof of Lemma 3.1, where we embed the separator of H into V_1 and the remainder into the rest of U_ℓ. However, we are in trouble, because

$$|U_\ell| = t|Z_\ell| > (1-k\alpha)\frac{mt}{k-1} \approx (1-k\alpha)\frac{N}{k-1} \approx (1-k\alpha)n,$$

which is close to but strictly smaller than n. To fix this, we will embed a small number of vertices of H in $U_{\ell'}$. However, the issue we run into is that our assumption that $d(V_i, U_{\ell'}) > \beta$ is rather weak, and so we can't do the simple greedy embedding we did earlier. Instead, we will need to use the dependent random choice lemma, Theorem 2.4, in two different ways.

First, we can find a subset $X \subset V_i$ with $|X| > \sigma|V_i|$, such that every vertex in X has at least $\beta'|U_{\ell'}|$ blue neighbors in $U_{\ell'}$ and at least $(1-\delta')|U_\ell \setminus V_i|$ blue neighbors in $U_\ell \setminus V_i$. We find this X by just deleting the vertices that don't satisfy these properties, and applying Markov's inequality.

Next, we apply Theorem 2.4 $k-2$ times, once each to the red graph between X and each U_a with $a \neq \ell$. By doing so repeatedly, we can find a set $Y \subset X$ with $|Y| \geq |X|^{1-k\lambda}$ such that every pair of vertices in Y has at least $\mu|U_a|$ common red neighbors in U_a, for each a. Now, if Y contains any red edge, then its two endpoints will have many common red neighbors in each U_a. In particular, there will be at least $\mu|V_a|$ common red neighbors inside some $V_a \subset U_a$, for each a. Now, the red graph between these U_a is ε-regular and has density at least δ, so by the counting lemma, we can find a red K_{k-2} among these common neighbors. Combining this with the red edge we started with, we find a red K_k, a contradiction. So there can be no red edge in Y, meaning that Y is a blue clique.

Now, we apply Theorem 2.4 again, this time to the blue graph between Y and $U_{\ell'}$. Recall that this graph has density at least β', since $Y \subset X$ and every vertex in X had at least $\beta'|U_{\ell'}|$ blue neighbors in $U_{\ell'}$. Now, this theorem allows us to find a subset $W \subset Y$ with $|W| > |V_1|^{1-\rho}$ such that every d-tuple of vertices in W has at least $\nu|U_{\ell'}|$ common blue neighbors in $U_{\ell'}$. Moreover, we still have that W is a blue clique, and that every vertex of W has at least $(1-\delta')|U_\ell \setminus V_i|$ blue neighbors in $U_\ell \setminus V_i$.

Now, we first embed the separator of H into W. This is easy because W is a blue clique, so all we need to check is that there is room, but there is because the separator has size at most $n^{1-\gamma}$, and $|W| > |V_1|^{1-\rho}$ for some arbitrarily small ρ. This is the reason we had to assume that H had a polynomially-small separator. We embed almost all the remaining vertices of H into $U_\ell \setminus V_i$, which is again easy because the blue density here is very high. However, since $|U_\ell|$ is slightly smaller than n, we finally need to embed the small number of remaining vertices into $U_{\ell'}$. This is again easy because of the degeneracy condition, and the fact that *every* d-tuple of vertices in W has at least $\nu|U_{\ell'}|$ common blue neighbors in $U_{\ell'}$. However, note that that we cannot embed too many vertices in $U_{\ell'}$, because ν is very small, and thus we can run into collisions if we try to embed more than roughly νn vertices in $U_{\ell'}$. Luckily, the parameters can be chosen so that this is ok, namely so that $U_\ell \setminus V_i$ has enough room to fit all but νn vertices of H. $\square$

Recall that all of the above was under the assumption that

$$\sum_{1 \le \ell < \ell' \le k-1} e_B(U_\ell, U_{\ell'}) \ge \beta \binom{N}{2}.$$

Since we can find a blue H in this case, we can now assume that

$$\sum_{1 \le \ell < \ell' \le k-1} e_B(U_\ell, U_{\ell'}) < \beta \binom{N}{2}. \tag{2}$$

Note that we've refined our structure further; now we have the dense blue parts $U_1, \ldots, U_{k-1}$, and we just showed that between them there are almost no blue edges. To refine it further, we can find a large induced subgraph where R itself is $(k-1)$-partite (rather than just the reduced graph R^* being $(k-1)$-partite). Indeed, (2) implies that R has edge density close to the Turán bound, since R is nearly complete between the $k-1$ parts $U_1, \ldots, U_{k-1}$. By repeatedly deleting the vertices of lowest degree, we can eventually get N' vertices, each of which has red degree at least $(1 - \frac{1}{k-1} - \tau)N'$. Since R is K_k-free, the Andrásfai–Erdős–Sós theorem implies that this induced subgraph is $(k-1)$-partite.

Concretely, we can show that there exist sets $S_1, \ldots, S_{k-1}$ such that $|S_i| > (1-\theta)n$ for all i, such that each S_i is a blue clique, and such that every vertex in $S_1 \cup \cdots \cup S_{k-1}$ has at least $(1-\theta)n$ red neighbors in every other S_i. Let $S = S_1 \cup \cdots \cup S_{k-1}$.

Since each S_i is a blue clique, it is again easy to embed almost all of H in some S_i. The issue is fitting in the remaining θn vertices of H, and it is at this step were the very careful arithmetic comes into play. We will embed the remaining vertices of H in $T \coloneqq V(K_N) \setminus S$.

First, observe that if there is a vertex in T that has at least $k\theta|S_i|$ red neighbors in each S_i, we can greedily find a red K_k, since each vertex in S_i has at least $(1-\theta)|S_j|$ red neighbors in each other S_j. Therefore, we may assume that every vertex in T has at most $k\theta|S_i|$ red neighbors in some S_i. We partition T into $T_1 \cup \cdots \cup T_{k-1}$, where each vertex in T_i has at most $k\theta|S_i|$ red neighbors in S_i. Now, the sets $(S_i \cup T_i)_{i=1}^{k-1}$ partition $V(K_N)$, so there exists some i with $|S_i \cup T_i| \ge n$. Recall that S_i is a blue clique of size at least $(1-\theta)n$, and that $d_B(S_i, T_i) \ge 1 - k\theta$. So we may greedily embed H in $S_i \cup T_i$. One has to be a bit careful because there may be exactly n vertices in $S_i \cup T_i$, so there may be no room to spare. However, since the blue density is so high, and since the majority of this set is an actual blue clique, this embedding isn't so hard.

4 Further remarks

The full theorem of Nikiforov and Rousseau is substantially more general than what I presented above. Indeed, looking at this proof, we can see that at every step where we found a red K_k, we did so by applying the counting lemma (or something similar). Because of this, we should expect that this proof can find not only one red K_k, but many of them, and indeed this is the case. Nikiforov and Rousseau's full theorem says that in any coloring of $E(K_N)$ with $N = (k-1)(n-1)+1$, we can find either a blue copy of H as above, or cn^{k-2}

red K_k all of which share a common edge, for some constant $c > 0$. This configuration is called a *joint* of size cn^{k-2}.

As it turns out, once you can find a large joint, you can find many other k-partite structures as well. For instance, by applying an earlier result of Nikiforov, one can show that a joint of size cn^{k-2} contains a complete multipartite graph $K_{1,1,r,r,\ldots,r}$, where $r = \kappa \log n$ for some constant $\kappa > 0$. Thus, Nikiforov and Rousseau's theorem implies a goodness-type result, where in blue we are looking for a very sparse graph of size n, while in red we are looking for a dense k-partite graph of size $\Omega(\log n)$. Additionally, if one allows the red graph to be somewhat sparser, one can even make it have size $\Omega(n)$, as certain such graphs can also be found inside large joints.

The single Burr–Erdős question about Ramsey goodness that wasn't answered by Nikiforov and Rousseau concerned the hypercube graph Q_d. Burr and Erdős had conjectured that Q_d is s-good for all sufficiently large d. Since the family of hypercubes does not have uniformly bounded degeneracy (Q_d is d-degenerate), Nikiforov and Rousseau's theorem does not apply. However, this conjecture of Burr and Erdős is true, as was proved by Fiz Pontiveros, Griffiths, Morris, Saxton, and Skokan. Although this means that all the Burr–Erdős conjectures are resolved, it would still be very interesting to obtain a fuller classification of which graphs are k-good and which ones are not; in particular, it would be very interesting to prove a fuller converse to Brandt's result, essentially saying that all sparse bad expanders are k-good.