Martin Marciniszyn, Jozef Skokan, Reto Spöhel and Angelika Steger

Threshold functions for asymmetric Ramsey properties involving cliques

Book Section (Refereed)

Original citation:
Originally published in Diaz, J; Jansen, K; Rolim, J; Zwick, U; *Approximation, randomization, and combinatorial optimization. Algorithms and techniques*. Lecture notes in computer science 4110, Berlin and Heidelberg, Germany, Springer-Verlag, 2006, pp. 462-474
DOI: 10.1007/11830924_42

Threshold Functions for Asymmetric Ramsey Properties Involving Cliques

Martin Marciniszyn[1,⋆], Jozef Skokan[2,⋆⋆], Reto Spöhel[1 ⋆⋆⋆], and Angelika Steger[1]

[1] Institute of Theoretical Computer Science, ETH Zurich, 8092 Zurich, Switzerland
{mmarcini,rspoehel,steger}@inf.ethz.ch
[2] Instituto de Matemática e Estatística, Universidade de São Paulo, 05508-090 São Paulo, SP, Brazil
jozef@member.ams.org

Abstract. Consider the following problem: For given graphs G and $F_1, \dots, F_k$, find a coloring of the edges of G with k colors such that G does not contain F_i in color i. For example, if every F_i is the path P_3 on 3 vertices, then we are looking for a proper k-edge-coloring of G, i.e., a coloring of the edges of G with no pair of edges of the same color incident to the same vertex.

Rödl and Ruciński studied this problem for the random graph $G_{n,p}$ in the symmetric case when k is fixed and $F_1 = \cdots = F_k = F$. They proved that such a coloring exists asymptotically almost surely (a.a.s.) provided that $p \leq bn^{-\beta}$ for some constants $b = b(F,k)$ and $\beta = \beta(F)$. Their proof was, however, non-constructive. This result is essentially best possible because for $p \geq Bn^{-\beta}$, where $B = B(F,k)$ is a large constant, such an edge-coloring does not exist. For this reason we refer to $n^{-\beta}$ as a *threshold function*.

In this paper we address the case when $F_1, \dots, F_k$ are cliques of different sizes and propose an algorithm that a.a.s. finds a valid k-edge-coloring of $G_{n,p}$ with $p \leq bn^{-\beta}$ for some constants $b = b(F_1, \dots, F_k, k)$ and $\beta = \beta(F_1, \dots, F_k)$. Kohayakawa and Kreuter conjectured that $n^{-\beta(F_1,\dots,F_k)}$ is a threshold function in this case. This algorithm can be also adjusted to produce a valid k-coloring in the symmetric case.

1 Introduction

The edge-chromatic number $\chi'(G)$ is one of the classical and well studied graph parameters. It is defined as the minimum number of colors k such that G allows for a k-edge-coloring with no pair of adjacent edges of the same color. Viewed

⋆ The author was supported by FAPESP (Proj. Temático–ProNEx Proc. FAPESP 2003/09925–5).

⋆⋆ The author was supported by NSF grant INT-0305793, by NSA grant H98230-04-1-0035, and by FAPESP (Proj. Temático–ProNEx Proc. FAPESP 2003/09925–5 and Proc. FAPESP 2004/15397–4).

⋆⋆⋆ The author was supported by SNF grant 200021-108158

from a slightly different perspective, one can equivalently define $\chi'(G)$ as the minimum number of colors k such that G admits a k-edge-coloring avoiding monochromatic paths of length 2. This definition has led to a fruitful and well-studied area in deterministic graph theory. For given graphs G and F, is there an edge-coloring with k colors of G that avoids a monochromatic copy of F?

It follows from Ramsey's celebrated result [1] that *every* k-coloring of the edges of the complete graph on n vertices contains a monochromatic copy of F if n is sufficiently large. While this seems to rely on the fact that K_n is a very dense graph, Folkman [2] and, in a more general setting, Nešetřil and Rödl [3] showed that there also exist locally sparse graphs $G = G(F)$ with the property that every k-coloring of the edges of G contains a monochromatic copy of F. By transferring the problem into a random setting, Rödl and Ruciński [4] showed that in fact such graphs G are quite frequent. More precisely, they proved the following result. Let

$$G \to (F)^e_k$$

denote the property that *every* edge-coloring of G with k colors contains a monochromatic copy of F. Recall that in the binomial random graph $G_{n,p}$ on n vertices, every edge is present with probability $0 \leq p = p(n) \leq 1$ independently of all other edges. Then the theorem of Rödl and Ruciński reads as follows.

Theorem 1 ([4], [5], [6]). *Let $k \geq 2$ and F be a non-empty graph that is not a forest. Then there exist constants $b, B > 0$ such that*

$$\lim_{n\to\infty} \mathbb{P}\left[G_{n,p} \to (F)^e_k\right] = \begin{cases} 0 & \text{if } p \leq bn^{-1/m_2(F)} \\ 1 & \text{if } p \geq Bn^{-1/m_2(F)} \end{cases} ,$$

where

$$m_2(F) := \max\left\{ \frac{|E(H)|-1}{|V(H)|-2} : H \subseteq F \wedge |V(H)| \geq 3 \right\} .$$

A function $p_0 = p_0(n)$ like the function $n^{-1/m_2(F)}$ in Theorem 1 is called threshold or threshold function. In Theorem 1, this function can be motivated as follows. For the sake of simplicity, suppose that $m_2(F) = (|E(F)|-1)/(|V(F)|-2)$. Then, for $p = cn^{-1/m_2(F)}$, the expected number of copies of F containing a given edge of $G_{n,p}$ is a constant depending on c. If this constant is close to zero, the copies of F in $G_{n,p}$ are loosely scattered and a valid coloring should thus exist. On the other hand, if this constant is large, the copies of F in $G_{n,p}$ highly intersect with each other, and the existence of a valid coloring becomes unlikely.

In Theorem 1 the same graph F is forbidden in every color class. We can generalize this setup by allowing for k different forbidden graphs, one per color. Within classical Ramsey theory the study of these so-called asymmetric Ramsey properties led to many interesting questions (see e.g. [7]) and results, most notably the celebrated paper of Kim [8] where he established an asymptotically sharp bound on the Ramsey number $R(3,t)$, that is, the minimum number n such that every 2-edge-coloring of the complete graph on n vertices contains either a red triangle or a blue clique of size t.

Within the random setting only very little is known about asymmetric Ramsey properties. Let

$$G \to (F_1, \dots, F_k)^e$$

denote the property that in *every* edge-coloring of G with k colors, there exists a color i such that F_i is contained in the subgraph of G spanned by the edges which are assigned to i. In [9] Kohayakawa and Kreuter proved the following result for cycles C_ℓ of fixed length ℓ.

Theorem 2 ([9]). *Let $k \geq 2$ and $3 \leq \ell_1 \leq \dots \leq \ell_k$ be integers. Then there exist constants $b, B > 0$ such that*

$$\lim_{n\to\infty} \mathbb{P}\left[G_{n,p} \to (C_{\ell_1}, \dots, C_{\ell_k})^e\right] = \begin{cases} 0 & \text{if } p \leq bn^{-1/m_2(C_{\ell_2}, C_{\ell_1})} \\ 1 & \text{if } p \geq Bn^{-1/m_2(C_{\ell_2}, C_{\ell_1})} \end{cases} ,$$

where

$$m_2(C_{\ell_2}, C_{\ell_1}) := \frac{\ell_1}{\ell_1 - 2 + 1/m_2(C_{\ell_2})} .$$

On the basis of their results in [9], Kohayakawa and Kreuter formulated the following conjecture.

Conjecture 3 ([9]). Let F_1, F_2 be graphs with $1 < m_2(F_1) \leq m_2(F_2)$. Then there exists a constant $b > 0$ such that for all $\varepsilon > 0$, we have

$$\lim_{n\to\infty} \mathbb{P}\left[G_{n,p} \to (F_1, F_2)^e\right] = \begin{cases} 0 & \text{if } p \leq (1-\varepsilon)bn^{-1/m_2(F_1,F_2)} \\ 1 & \text{if } p \geq (1+\varepsilon)bn^{-1/m_2(F_1,F_2)} \end{cases} ,$$

where

$$m_2(F_1, F_2) := \max\left\{ \frac{|E(H)|}{|V(H)| - 2 + 1/m_2(F_1)} : H \subseteq F_2 \wedge |V(H)| \geq 2 \right\} .$$

The threshold function from Conjecture 3 is supported by the following observation. The expected number of copies of F_2 in $G_{n,p}$ with $p = \Theta\big(n^{-1/m_2(F_1,F_2)}\big)$ is

$$\Theta\left(n^{|V(F_2)|} p^{|E(F_2)|}\right) = \Omega\left(n^{2-1/m_2(F_1)}\right) .$$

Since every edge-coloring of $G_{n,p}$ must avoid monochromatic copies of F_2 in color 2, there is at least one edge of color 1 in every subgraph of $G_{n,p}$ isomorphic to F_2. Select one such edge from each copy of F_2 arbitrarily. It is plausible that these edges span a graph G' with edge density $\Omega\big(n^{-1/m_2(F_1)}\big)$ that satisfies certain pseudo-random properties. As it turns out, that seems just about the right density in order to embed a copy of F_1 into G', no matter which edges were selected from the original graph.

In this paper, we consider the threshold function p_0 for cliques $K_{\ell_1}, \dots, K_{\ell_k}$. A threshold phenomenon consists of two separate statements, the so-called 0-statement and the 1-statement, which are usually proved in entirely different ways. In our setting, the two statements are as follows. For the 1-statement one has to show that if $p \geq Bp_0$, a random graph $G_{n,p}$ asymptotically almost surely (a.a.s.) satisfies $G_{n,p} \to (K_{\ell_1}, \ldots, K_{\ell_k})^e$, i.e., *every* k-edge-coloring of $G_{n,p}$ contains at least one of the forbidden monochromatic cliques. For the 0-statement we suppose that $p \leq bp_0$ for some sufficiently small constant $b > 0$ and need to provide a k-edge-coloring of a random graph $G_{n,p}$ that avoids every forbidden clique K_{ℓ_i}, $1 \leq i \leq k$, in the corresponding color class i.

A standard way of attacking the 1-statement, which was also pursued in [9], is via the sparse version of Szemerédi's regularity lemma, which was independently developed by Kohayakawa [10] and Rödl (unpublished). Using properties of regularity, one can find a monochromatic copy of a forbidden subgraph in the colored graph $G_{n,p}$. Unfortunately, generalizing this argument from cycles to cliques requires a proof of Conjecture 23 in [11] of Kohayakawa, Łuczak, and Rödl. This so-called KŁR-Conjecture formulates a probabilistic version of the classical embedding lemma for dense graphs. It implies many interesting extremal results for random graphs. In their monograph on random graphs [12], Janson, Łuczak, and Ruciński consider the verification of this conjecture one of the most important open questions in the theory of random graphs. Despite recent progress [13], the conjecture is, in its full generality, still wide open. However, assuming that it is true, a proof of the 1-statement is routinely obtained. We omit the proof in this extended abstract due to space restrictions.

From an algorithmic or constructive point of view, the 0-statement is much more interesting. The way of proving it that was pursued in [5] and [9] is by contradiction. This approach shows the existence of a coloring, but provides no efficient way of obtaining the coloring from the proof. Our approach is constructive. We provide a (polynomial-time) algorithm that computes a valid coloring for graphs that satisfy certain properties. We employ techniques similar to those in [5] and [9] in order to prove that these properties a.a.s. hold in $G_{n,p}$ with p sufficiently small. Indeed, the results in [5] yield that our algorithm also computes valid colorings of $G_{n,p}$ in the symmetric case, unless the forbidden graph is one of a few special cases, e.g., a triangle. In fact, the symmetric case of triangles was solved in [6] by different methods.

We prove the threshold from Conjecture 3 for cliques. As in Theorems 1 and 2, the threshold function is slightly weaker than conjectured, allowing for distinct constants in the 0- and the 1-statement.

Theorem 4 (Main Result). *Let $k \geq 2$ and $\ell_1 \geq \cdots \geq \ell_k \geq 3$ be integers. Then there exist constants $b, B > 0$ such that*

$$\lim_{n\to\infty} \mathbb{P}\left[G_{n,p} \to (K_{\ell_1}, \ldots, K_{\ell_k})^e\right] = \begin{cases} 0 & \text{if } p \leq bn^{-1/m_2(K_{\ell_2}, K_{\ell_1})} \\ 1 & \text{if } p \geq Bn^{-1/m_2(K_{\ell_2}, K_{\ell_1})} \end{cases} ,$$

where

$$m_2(K_{\ell_2}, K_{\ell_1}) := \frac{\binom{\ell_1}{2}}{\ell_1 - 2 + 1/m_2(K_{\ell_2})} ,$$

and the 1-statement holds provided Conjecture 23 in [11] is true for K_{ℓ_2}.

In this extended abstract, we will outline the proof of the 0-statement of Theorem 4 under the additional assumption that $\ell_2 > 3$. For $\ell_2 = 3$, we face additional difficulties. Due to space restrictions, we focus on the main case and sketch how to deal with triangles in Section 3.

1.1 Notation

Our notation is mostly adopted from [12]. All graphs are simple and undirected. We abbreviate the number of vertices $|V(G)|$ of a graph G by $v(G)$ and similarly the number of edges $|E(G)|$ by $e(G)$. We say that a property $\mathcal{P}$ holds in $G_{n,p}$ *asymptotically almost surely* (a.a.s.) if we have

$$\lim_{n\to\infty} \mathbb{P}\left[G_{n,p} \text{ satisfies } \mathcal{P}\right] = 1 \ .$$

2 An algorithm for computing valid edge colorings

Suppose $G = G_{n,p}$ with $p \leq bn^{-1/m_2(K_{\ell_2},K_{\ell_1})}$ is given. In order to provide a valid coloring of G, it suffices to compute a 2-coloring of $E(G)$ such that there is no copy of K_{ℓ_1} in color 1 and no copy of K_{ℓ_2} in color 2. That implies the 0-statement of Theorem 4 also for k-colorings. Hence, we focus on 2-colorings and abbreviate ℓ_1 by r and ℓ_2 by ℓ in the following. For the rest of this section, $r > \ell > 3$ shall remain fixed.

We describe an algorithm that finds a valid edge-coloring of G a.a.s. The basic idea of the algorithm is to remove edges from the graph successively. An edge e is deleted from G if there are no two cliques of size ℓ and r respectively that intersect exactly on e. Assuming that all edges of G can be removed in this way, it is easy to create a valid coloring by inserting them in the reverse order one by one, always assigning a valid color instantly. The actual algorithm is more complex since sometimes one has to *forget* about the existence of certain small cliques in order to remove really all edges from G. As we shall see, we can easily deal with those cliques later.

In order to simplify notation, we define, for any graph G, the families

$$\mathcal{L}_G := \{L \subseteq G : L \cong K_\ell\} \quad \text{and} \quad \mathcal{R}_G := \{R \subseteq G : R \cong K_r\}$$

of all ℓ-cliques and r-cliques in G respectively. Furthermore, we introduce the family

$$\mathcal{L}^*_G := \big\{L \in \mathcal{L}_G : \forall e \in E(L) \ \exists R \in \mathcal{R}_G \text{ s.t. } E(L) \cap E(R) = \{e\}\big\} \ \subseteq \ \mathcal{L}_G \ .$$

The algorithm is given in Figure 1. Note that edges are removed from and inserted into a working copy $G' = (V, E')$ of G. The local variable $\mathcal{L}$ contains the same elements as $\mathcal{L}_{G'}$ up to the first execution of lines 12-13. In general, we have $\mathcal{L} \subseteq \mathcal{L}_{G'}$.

Lemma 5. *If algorithm* Asym-Edge-Col *terminates without error, then it has indeed found a valid coloring of* G*.*

```
ASYM-EDGE-COL(G = (V, E))
 1  s ← EMPTY-STACK()
 2  E′ ← E
 3  ℒ ← ℒ_G
 4  while E′ ≠ ∅
 5  do if ∃ e ∈ E′ s.t. ∄(L, R) ∈ ℒ × ℛ_{G′=(V,E′)} : E(L) ∩ E(R) = {e}
 6        then for  each L ∈ ℒ : e ∈ E(L)
 7                do s.PUSH(L)
 8                   ℒ.REMOVE(L)
 9                s.PUSH(e)
10                E′.REMOVE(e)
11        else  if ∃L ∈ ℒ \ ℒ*_{G′=(V,E′)}
12                then s.PUSH(L)
13                     ℒ.REMOVE(L)
14        else  error "stuck"
15  while s ≠ ∅
16  do if s.TOP() is an edge
17        then e ← s.POP()
18                e.SET-COLOR( blue )
19                E′.ADD(e)
20        else  L ← s.POP()
21                if L is entirely blue
22                   then f ←  any e ∈ E(L) s.t. ∄R ∈ ℛ_{G′=(V,E′)} : E(L) ∩ E(R) = {e}
23                        f.SET-COLOR( red )
```

Fig. 1. The implementation of algorithm ASYM-EDGE-COL.

Proof. First, we argue that the algorithm never creates a blue copy of K_ℓ. Observe that *every* copy of K_ℓ that exists in G' is pushed on the stack in the first loop. Therefore, in the execution of the second loop, the algorithm must check the coloring of every such copy. Due to the order of the elements on the stack, each check is performed only after all edges of the corresponding clique were inserted and colored. For every blue copy of K_ℓ, one particular edge is recolored to red. Since red edges are never flipped back to blue, no blue copy of K_ℓ can occur in the coloring found by the algorithm.

It remains to prove that the assignment of color red to some edge by the algorithm can never create an entirely red copy of K_r. By the condition on f in line 22 of the algorithm, at the very moment there exists no copy of K_r in G' that intersects with L exactly in f. So there is either no K_r containing f at all, or every such copy contains also another edge from L. In the latter case, those copies cannot become entirely red since L is entirely blue.

Our last step is to show that the edge f in line 22 always exists. Since the second loop inserts edges into G' in the reverse order in which they were deleted during the first loop, when we select f in line 22, G' has the same structure as at the moment when L was pushed on the stack. This could have happened either in line 7, when there exists no r-clique in G' that intersects with L on some particular edge $e \in E(L)$, or in line 12, when L satisfies the condition of the if-clause in line 11. In both cases we have $L \notin \mathcal{L}^*_{G'}$, and, therefore, there exists an edge $e \in E(L)$ such that all currently existing copies of K_r do not intersect with L exactly in e. □

It remains to prove the following lemma.

Lemma 6. *There exists a positive constant $b = b(\ell, r)$ such that the algorithm* Asym-Edge-Col *a.a.s. terminates on $G_{n,p}$ with $p \leq bn^{-1/m_2(K_\ell, K_r)}$ without error.*

2.1 Proof of Lemma 6

We prove Lemma 6 by providing an algorithm Grow that, if Asym-Edge-Col fails on an arbitrary graph G, explicitly computes a subgraph $F \subseteq G$ which is either too large or too dense to appear in $G_{n,p}$ with p as in the lemma. More precisely, we shall show that for any graph F that Grow may return, the probability that F appears in $G_{n,p}$ is small compared to the size of $\mathcal{F}$, the class of all graphs that Grow may return. It follows that $G_{n,p}$ a.a.s. does not contain any of these graphs, which implies Lemma 6 by contradiction. Note that we employ algorithm Grow only for proving the lemma. It does not contribute to the running time of algorithm Asym-Edge-Col.

In order to formulate algorithm Grow, we need some definitions. Let

$$\gamma = \gamma(\ell, r) := 1/m_2(K_\ell, K_r) - 2/\left(\ell + r - 3\right) \ .$$

Note that for $r > \ell > 3$, we have

$$\gamma(\ell, r) = \frac{2\big((\ell^2 - 3\ell - 2)r - 2\ell(\ell - 3)\big)}{r(r-1)(\ell+1)(\ell+r-3)} > 0 \ .$$

Remark 7. Observe that $\gamma(3, r)$ is negative for $r \geq 3$. This is why we have to modify our proof for the case $\ell = 3$, see Section 3. The proof we present here also covers the symmetric case for $\ell = r \geq 5$ since then $\gamma(\ell, \ell) > 0$.

For any graph F, let

$$\lambda(F) := v(F) - e(F)/m_2(K_\ell, K_r) \ .$$

The definition of $\lambda(F)$ is motivated by the fact that the number of copies of F in $G_{n,p}$ with $p = bn^{-1/m_2(K_\ell, K_r)}$ has order of magnitude

$$n^{v(F)} p^{e(F)} = b^{e(F)} n^{\lambda(F)} \ .$$

For any graph F, we call an edge $e \in E(F)$ *eligible for extension* if it satisfies

$$\nexists (L, R) \in \mathcal{L}_F \times \mathcal{R}_F \text{ s.t. } E(F) \cap E(F) = \{e\} \ .$$

GROW($G' = (V, E)$)
$i \leftarrow 0$
$F_0 \leftarrow$ any $R \in \mathcal{R}_{G'}$
while $i < \log(n) \ \wedge \ \lambda(F_i) > -\gamma$
do if $\exists R \in \mathcal{R}_{G'} \setminus \mathcal{R}_{F_i}$ s.t. $|V(R) \cap V(F_i)| \geq 2$
  then $F_{i+1} \leftarrow F_i \cup R$
  else $e \leftarrow$ ELIGIBLE-EDGE(F_i)
    $F_{i+1} \leftarrow$ EXTEND-L(F_i, e, G')
  $i \leftarrow i + 1$
return F_i

EXTEND-L(F, e, G')
$L \leftarrow$ any $L \in \mathcal{L}^*_{G'} : e \in E(L)$
$F' \leftarrow F \cup L$
for each e' **in** $E(L) \setminus E(F)$
do $R_{e'} \leftarrow$ any $R \in \mathcal{R}_{G'} : E(L) \cap E(R) = \{e'\}$
  $F' \leftarrow F' \cup R_{e'}$
return F'

Fig. 2. The implementation of algorithm GROW.

The implementation of algorithm GROW is shown in Figure 2. The intended input is the graph $G' \subseteq G$ after ASYM-EDGE-COL got stuck. It proceeds as follows: the seed F_0 is any copy of K_r in G'. In every iteration i, it extends F_i to F_{i+1} by adding new vertices and edges to it. As long as there are copies of K_r in G' that intersect with F_i in at least two vertices but not in all edges, it greedily adds those to F_i. If there are no such copies, it calls a subroutine ELIGIBLE-EDGE that takes F_i as input and returns an edge $e \in E(F_i)$ eligible for extension that is *unique up to isomorphism of* F_i, i.e., in such a way that for any two isomorphic graphs F and F', there exists an isomorphism φ with $\varphi(F) = F'$ such that

$$\text{ELIGIBLE-EDGE}(F') = \varphi(\text{ELIGIBLE-EDGE}(F)) \ .$$

Note that this implies in particular that e depends only on the graph F_i and not on the surrounding graph G'. Clearly, one way to implement this procedure would be keeping a large table of representatives for all isomorphism classes of graphs with up to n vertices that maps to each entry one particular edge eligible for extension. Since we only want to show the existence of a certain structure in G' and do not care about complexity issues here, the actual implementation of that procedure is irrelevant. Procedure EXTEND-L then adds a graph $L \in \mathcal{L}^*_{G'}$ that contains the edge e returned by ELIGIBLE-EDGE to F_i. It glues to each new edge $e' \in E(L) \setminus E(F_i)$ a graph $R_{e'} \in \mathcal{R}_{G'}$ that intersects with L only in e'. The algorithm stops and returns $F_i \subseteq G' \subseteq G$ as soon as $\lambda(F_i) \leq -\gamma$ or $i \geq \log(n)$.

We argue that GROW terminates without error, i.e., that ELIGIBLE-EDGE always finds an edge eligible for extension, and that EXTEND-L always finds suitable graphs L and $R_{e'}$, $e' \in E(L)$. Let us consider the properties of G' when

ASYM-EDGE-COL gets stuck. As the condition in line 5 of ASYM-EDGE-COL fails, G' is in the family

$$\mathcal{C}(\ell, r) := \big\{G = (V, E) : \forall e \in E(G) \ \exists (L, R) \in \mathcal{L}_G \times \mathcal{R}_G \text{ s.t. } E(L) \cap E(R) = \{e\}\big\} \ .$$

In fact, every edge of G' is contained in a copy $L \in \mathcal{L}$, and as the condition in line 11 fails as well, G' is even in the smaller family

$$\mathcal{C}^*(\ell, r) := \big\{G = (V, E) : \forall e \in E(G) \ \exists L \in \mathcal{L}^*_G \text{ s.t. } e \in E(L)\big\} \ \subseteq \ \mathcal{C}(\ell, r) \ .$$

Claim 8. *Algorithm* GROW *terminates without error on any nonempty input graph $G' \in \mathcal{C}^*(\ell, r)$. Moreover, every iteration of the while-loop adds at least one edge to F.*

Proof. Suppose there is no edge in F_i that is eligible for extension. Then we have $F_i \in \mathcal{C}(\ell, r)$ by definition. This implies that every vertex $v \in V(F_i)$ has degree at least $(\ell - 1) + (r - 1) - 1 = \ell + r - 3$, i.e., $e(F_i)/v(F_i) \geq (\ell + r - 3)/2$. It follows that

$$\lambda(F_i) \leq e(F_i)\left(\frac{2}{\ell + r - 3} - \frac{1}{m_2(K_\ell, K_r)}\right) = -e(F_i)\gamma \leq -\gamma \ ,$$

where we used that $\gamma = \gamma(\ell, r)$ is positive. Consequently, GROW terminates in line 3 without calling ELIGIBLE-EDGE. Hence, ELIGIBLE-EDGE always returns an edge eligible for extension when called from GROW.

Property $\mathcal{C}^*(\ell, r)$ of G' guarantees the existence of suitable graphs L and $R_{e'}$, $e' \in E(L)$, when EXTEND-L is called. Moreover, by definition of $\mathcal{L}^*_{G'}$, there exists, in particular, $R_e \in \mathcal{R}_{G'}$ such that e is the intersection of R_e and L. When EXTEND-L(F, e, G') is called, R_e has already been added to F during a previous iteration in lines 4 and 5 of GROW. Hence, the L returned in line 1 of EXTEND-L is not contained in F, as otherwise e would not be eligible for extension. On the other hand, it is clear that an R found in line 4 adds at least one new edge to F. Together this proves that every iteration adds at least one edge to F. □

Now, we will consider the evolution of F in more detail. We say that iteration i of the while-loop in procedure GROW is *non-degenerate* if we have the following assertions:

- The condition in line 4 evaluates to false and, hence, EXTEND-L is called.
- In line 2 of EXTEND-L, we have $V(F) \cap V(L) = e$.
- In every execution of line 5 of EXTEND-L, we have $V(F') \cap V(R_{e'}) = e'$.

Otherwise, we call iteration i *degenerate*. In non-degenerate iterations, F_{i+1} is uniquely defined up to isomorphism for a given F_i, depending only on the implementation of subroutine ELIGIBLE-EDGE, which determines the position where to attach the next K_ℓ. An easy calculation yields the next claim.

Claim 9. *If iteration i of the while-loop in procedure* GROW *is non-degenerate, we have*

$$\lambda(F_{i+1}) = \lambda(F_i) \ .$$

In a degenerate iteration i, the structure of F_{i+1} does not only depend on F_i, but varies with the structure of G'. Suppose that F_i is extended with an r-clique in line 5. This R can intersect with F_i in virtually every possible way. Moreover, there may be several copies of K_r which satisfy the condition in line 4. The same is true for the graphs added in lines 2 and 5 of Extend-L. Thus, degenerate iterations cause difficulties since they enlarge the family of graphs that algorithm Grow may potentially return. However, we will show that at most a *constant* number of degenerate iterations can occur before the algorithm terminates. This allows us to control the number of non-isomorphic graphs that can be the output of Grow. The key to proving this is the next claim.

Claim 10. *There exists a constant $\kappa = \kappa(\ell, r) > 0$ such that if iteration i of the while-loop in procedure* Grow *is degenerate, we have*

$$\lambda(F_{i+1}) \leq \lambda(F_i) - \kappa \ .$$

The proof of Claim 10 is the main technical part of our work and beyond the scope of this extended abstract. In combination with Claim 9, it yields Claim 11, which in turn leads to a polylogarithmic bound on the number of non-isomorphic graphs that Grow can return. The proof of Claim 11 is omitted due to space restrictions.

Claim 11. *There exists a constant $m_0 = m_0(\ell, r)$ such that algorithm* Grow *performs at most m_0 degenerate iterations before it terminates, regardless of the input instance G'.*

Let $\mathcal{F}(\ell, r, n)$ denote a family of representatives for the isomorphism classes of all graphs that can be the output of Grow with parameters n and $\gamma(\ell, r)$ on *any* input instance G'.

Claim 12. *There exists $C = C(\ell, r)$ such that $|\mathcal{F}(\ell, r, n)| \leq \log(n)^C$.*

Proof. For $t \geq d \geq 0$, let $\mathcal{F}(t, d)$ denote a family of representatives for the isomorphism classes of all graphs F_t that algorithm Grow can generate after exactly t iterations if it performs exactly d degenerate iterations along the way, and let $f(t, d) := |\mathcal{F}(t, d)|$ denote its cardinality.

Observe that in every iteration, we add at most

$$K := \ell - 2 + \binom{\ell}{2}(r - 2)$$

new vertices to F, which is exactly the number of vertices added in a non-degenerate iteration. Hence, we have $v(F_t) \leq r + Kt$. It also follows that in every iteration, the new edges $E(F_{t+1}) \setminus E(F_t)$ span a graph from $\mathcal{G}_K$, where $\mathcal{G}_K$ denotes the set of all graphs on at most K vertices. F_{t+1} is uniquely defined if one specifies $G \in \mathcal{G}_K$, the number y of vertices in which G intersects F_t, and two ordered lists of vertices from G and F_t respectively of length y, which specify the mapping of the intersection vertices from G into F_t. Thus, the number of ways to extend F_t is bounded from above by

$$\sum_{G\in\mathcal{G}_K}\sum_{y=2}^{v(G)} v(G)^y(v(F_t))^y \le C_1(r+Kt)^K \le t^{C_2} \le \log(n)^{C_2} \ ,$$

where the constants C_1 and C_2 depend only on ℓ and r.

As the selection of the edge to be extended is unique up to isomorphism of F, the evolution of F is uniquely defined if there are no degenerate iterations along the way, regardless of the input instance G'. This implies in particular that $f(t,0)=1$ for all t, and more generally that for $t\ge d\ge 0$

$$f(t,d) \le \binom{t}{d}\big(\log(n)^{C_2}\big)^d \le \log(n)^{(C_2+1)d} \ .$$

Here the binomial coefficient corresponds to the choice of the d degenerate iterations. We conclude from Claim 11 that there exists a constant $C=C(\ell,r)>0$ such that

$$|\mathcal{F}(\ell,r,n)| \le \sum_{t=0}^{\log(n)}\sum_{d=0}^{m_0} f(t,d) \le (\log(n)+1)(m_0+1)\log(n)^{(C_2+1)m_0} \le \log(n)^C$$

for n sufficiently large. □

Claim 13. *There exists a constant $b>0$ such that for $p\le bn^{-1/m_2(K_\ell,K_r)}$, $G_{n,p}$ does not contain any graph from $\mathcal{F}(\ell,r,n)$ a.a.s.*

Proof. Let $\mathcal{F}_1$ and $\mathcal{F}_2$ denote the classes of graphs that algorithm Grow can output if it terminates due to the first or the second condition in line 3, respectively. Owing to Claim 12 we have a polylogarithmic bound on the cardinality of $\mathcal{F}=\mathcal{F}(\ell,r,n)=\mathcal{F}_1\cup\mathcal{F}_2$, and Claims 9 and 10 imply that $\lambda(F_i)$ is non-increasing. It follows that for $b:=\mathrm{e}^{-\lambda(F_0)-\gamma}$, the expected number of copies of graphs from $\mathcal{F}$ in $G_{n,p}$ with $p\le bn^{-1/m_2(K_\ell,K_r)}$ is bounded by

$$\begin{aligned}\sum_{F\in\mathcal{F}} n^{v(F)}p^{e(F)} = \sum_{F\in\mathcal{F}} b^{e(F)}n^{\lambda(F)} &\le \sum_{F\in\mathcal{F}_1} e^{(-\lambda(F_0)-\gamma)\log(n)}n^{\lambda(F)} + \sum_{F\in\mathcal{F}_2} n^{-\gamma}\\ &\le (\log(n))^C n^{-\gamma} = o(1) \ ,\end{aligned}$$

which implies the claim due to Markov's inequality. Here we used again that γ is positive. Note that crucially, for all $F\in\mathcal{F}_1$, we have $e(F)\ge\log(n)$ since F was generated in $\lceil\log(n)\rceil$ iterations, each of which introduces at least one new edge. □

Suppose now that algorithm Asym-Edge-Col applied to $G_{n,p}$ with p as claimed gets stuck, and consider $G'\subseteq G$ at this moment. The call to Grow(G') returns a copy of a graph $F\in\mathcal{F}(\ell,r,n)$ that is contained in G'. But we just proved that a.a.s. we have $F\not\subseteq G_{n,p}$, which contradicts our assumption. This proves that Asym-Edge-Col finds a valid coloring of $G_{n,p}$ with $p\le bn^{-1/m_2(K_\ell,K_r)}$ a.a.s.

3 Triangles

As stated in Remark 7, the proof presented in Section 2 does not cover the case $\ell = 3$ since $\gamma(3, r) = -1/(r^2 - r) < 0$. In particular, this implies that, for any $b > 0$, $G_{n,p}$ with $p = bn^{-1/m_2(K_3,K_r)}$ may contain copies of K_{r+1}. Since K_{r+1} is a member of the family $\mathcal{C}^*(3, r)$, Asym-Edge-Col will terminate with an error. Some rather technical work is required to show that, for $r \geq 6$, K_{r+1} is essentially the only graph in $\mathcal{C}^*(3, r)$ that is sparse enough to appear in $G_{n,p}$ and cause problems. Once this is established, it is not hard to prove that when Asym-Edge-Col gets stuck, G' is a.a.s. the union of edge-disjoint copies of K_{r+1} and can be colored easily. Some further complications arise in the cases $r = 4$ and $r = 5$, but the main line of argument is the same.

References

1. Ramsey, F.P.: On a problem of formal logic. Proceedings of the London Mathematical Society **30** (1930) 264–286
2. Folkman, J.: Graphs with monochromatic complete subgraphs in every edge coloring. SIAM J. Appl. Math. **18** (1970) 19–24
3. Nešetřil, J., Rödl, V.: The Ramsey property for graphs with forbidden complete subgraphs. J. Combinatorial Theory Ser. B **20** (1976) 243–249
4. Rödl, V., Ruciński, A.: Threshold functions for Ramsey properties. J. Amer. Math. Soc. **8** (1995) 917–942
5. Rödl, V., Ruciński, A.: Lower bounds on probability thresholds for Ramsey properties. In: Combinatorics, Paul Erdős is eighty, Vol. 1. Bolyai Soc. Math. Stud. János Bolyai Math. Soc., Budapest (1993) 317–346
6. Łuczak, T., Ruciński, A., Voigt, B.: Ramsey properties of random graphs. J. Combin. Theory Ser. B **56** (1992) 55–68
7. Chung, F., Graham, R.: Erdős on graphs. A K Peters Ltd., Wellesley, MA (1998) His legacy of unsolved problems.
8. Kim, J.H.: The Ramsey number $R(3, t)$ has order of magnitude $t^2/\log t$. Random Structures Algorithms **7** (1995) 173–207
9. Kohayakawa, Y., Kreuter, B.: Threshold functions for asymmetric Ramsey properties involving cycles. Random Structures Algorithms **11** (1997) 245–276
10. Kohayakawa, Y.: Szemerédi's regularity lemma for sparse graphs. In: Foundations of computational mathematics (Rio de Janeiro, 1997). Springer, Berlin (1997) 216–230
11. Kohayakawa, Y., Łuczak, T., Rödl, V.: On K^4-free subgraphs of random graphs. Combinatorica **17** (1997) 173–213
12. Janson, S., Łuczak, T., Rucinski, A.: Random graphs. Wiley-Interscience Series in Discrete Mathematics and Optimization. Wiley-Interscience, New York (2000)
13. Gerke, S., Marciniszyn, M., Steger, A.: A probabilistic counting lemma for complete graphs. In Felsner, S., ed.: 2005 European Conference on Combinatorics, Graph Theory and Applications (EuroComb '05). Volume AE of DMTCS Proceedings., Discrete Mathematics and Theoretical Computer Science (2005) 309–316