Monochromatic triangles in three-coloured graphs

James Cummings[a,1], Daniel Král'[b,2], Florian Pfender[c], Konrad Sperfeld[d], Andrew Treglown[e,2], Michael Young[f,3]

[a] *Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh PA 15213, USA*
[b] *Mathematics Institute, DIMAP and Department of Computer Science, University of Warwick, Coventry CV4 7AL, UK.*
Previous affiliation: Computer Science Institute, Faculty of Mathematics and Physics, Charles University, Malostranské náměstí 25, 118 00 Prague, Czech Republic
[c] *University of Colorado Denver, Mathematical and Statistical Sciences, Denver, CO, 80202, USA*
[d] *Universität Rostock, Institut für Mathematik, 18057 Rostock, Germany*
[e] *School of Mathematical Sciences, Queen Mary, University of London, London, E1 4NS, UK.*
Previous affiliation: Computer Science Institute, Faculty of Mathematics and Physics, Charles University, Malostranské náměstí 25, 118 00 Prague, Czech Republic
[f] *Department of Mathematics, Iowa State University, Ames, IA 50011, USA*

Abstract

In 1959, Goodman [9] determined the minimum number of monochromatic triangles in a complete graph whose edge set is 2-coloured. Goodman [10] also raised the question of proving analogous results for complete graphs whose edge sets are coloured with more than two colours. In this paper, for n sufficiently large, we determine the minimum number of monochromatic triangles in a 3-coloured copy of K_n. Moreover, we characterise those 3-coloured copies of K_n that contain the minimum number of monochromatic triangles.

1. Introduction

The *Ramsey number* $R_k(G)$ of a graph G is the minimum $n \in \mathbb{N}$ such that every k-colouring of K_n contains a monochromatic copy of G. (In this paper we say a graph K is *k-coloured* if we have coloured the *edge set* of K using k colours. Note that the edge colouring need not be proper.) A famous theorem of Ramsey [17] asserts that $R_k(G)$ exists for all graphs G and all $k \in \mathbb{N}$.

In light of this, it is also natural to consider the so-called *Ramsey multiplicity* of a graph: Let $k, n \in \mathbb{N}$ and let G be a graph. The *Ramsey multiplicity* $M_k(G, n)$ of G is the minimum number of monochromatic copies of G over all k-colourings of K_n. (Here, we are counting unlabelled copies of G in the sense that we count the number of distinct monochromatic subgraphs of K_n that are isomorphic to G.) In the case when $k = 2$ we simply write $M(G, n)$. The following classical result of Goodman [9] from 1959 gives the precise value of $M(K_3, n)$.

[1] This author's research was partially supported by National Science Foundation grant DMS-1101156.
[2] The work leading to this invention has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013)/ERC grant agreement no. 259385.
[3] This author's research was supported by NSF grant DMS 0946431.

Theorem 1 (Goodman [9])**.** *Let $n \in \mathbb{N}$. Then*

$$M(K_3, n) = \begin{cases} n(n-2)(n-4)/24 & \text{if } n \text{ is even;} \\ n(n-1)(n-5)/24 & \text{if } n \equiv 1 \mod 4; \\ (n+1)(n-3)(n-4)/24 & \text{if } n \equiv 3 \mod 4. \end{cases}$$

A graph G is *k-common* if $M_k(G,n)$ asymptotically equals, as n tends to infinity, the expected number of monochromatic copies of G in a random k-colouring of K_n. Erdős [6] conjectured that K_r is 2-common for every $r \in \mathbb{N}$. Note that Theorem 1 implies that this conjecture is true for $r = 3$. However, Thomason [22, 23] disproved the conjecture in the case when $r = 4$. Further, Jagger, Šťovíček, and Thomason [13] proved that any graph G that contains K_4 is not 2-common. Recently, Cummings and Young [4] proved that graphs G that contain K_3 are not 3-common. The introductions of [4] and [12] give more detailed overviews of k-common graphs.

The best known general lower bound on $M(K_r, n)$ was proved by Conlon [3]. Some general bounds on $M_k(K_r, n)$ are given in [7]. See [2] for a (somewhat outdated) survey on Ramsey multiplicities.

The problem of obtaining a 3-coloured analogue of Goodman's theorem also has a long history. In fact, it is not entirely clear when this problem was first raised. In 1985, Goodman [10] simply refers to it as "an old and difficult problem". Prior to this, Giraud [8] proved that, for sufficiently large n, $M_3(K_3, n) > 4\binom{n}{3}/115$. Wallis [24] showed that $M_3(K_3, 17) \leq 5$ and then, together with Sane [20], proved that $M_3(K_3, 17) = 5$. (Greenwood and Gleason [11] proved that $R_3(K_3) = 17$, therefore, $M_3(K_3, 16) = 0$.)

The focus of this paper is to give the exact value of $M_3(K_3, n)$ for sufficiently large n, thereby yielding a 3-coloured analogue of Goodman's theorem. Moreover, we characterise those 3-coloured copies of K_n that contain exactly $M_3(K_3, n)$ monochromatic triangles.

Given $n \in \mathbb{N}$ we define a special collection of 3-coloured complete graphs on n vertices, $\mathcal{G}_n$ as follows:

- Consider the (unique) 2-coloured copy K of K_5 on $[5]$ without a monochromatic triangle. Replace the vertices of K with disjoint vertex classes $V_1, \dots, V_5$ such that $||V_i| - |V_j|| \leq 1$ for all $1 \leq i, j \leq 5$ and $|V_1| + \cdots + |V_5| = n$. For all $1 \leq i \neq j \leq 5$, add all possible edges between V_i and V_j using the colour of ij in K. For each $1 \leq i \leq 5$, add all possible edges inside V_i in a third colour. Denote the resulting complete 3-coloured graph by $G_{ex}(n)$ (see Figure 1).

- $\mathcal{G}_n$ consists of $G_{ex}(n)$ together with all graphs obtained from $G_{ex}(n)$ by recolouring a (possibly empty) matching $M_{i,j}$ in $G_{ex}(n)[V_i, V_j]$ with the third colour for all $1 \leq i \neq j \leq 5$, such that the recolouring does not introduce any new monochromatic triangles (see Figure 1).

Notice that the graphs in $\mathcal{G}_n$ only contain monochromatic triangles of one colour. The following is our main result.

Theorem 2. *There exists an $n_0 \in \mathbb{N}$ such that the following holds. Suppose G is a complete 3-coloured graph on $n \geq n_0$ vertices which contains the smallest number of monochromatic triangles amongst all complete 3-coloured graphs on n vertices. Then G is a graph from $\mathcal{G}_n$.*

Corollary 3. *There exists an $n_0 \in \mathbb{N}$ such that the following holds. Suppose $n \geq n_0$ and write $n = 5m + r$ where $m, r \in \mathbb{N}$ such that $0 \leq r \leq 4$. Then*

$$M_3(K_3, n) = r\binom{m+1}{3} + (5-r)\binom{m}{3}.$$

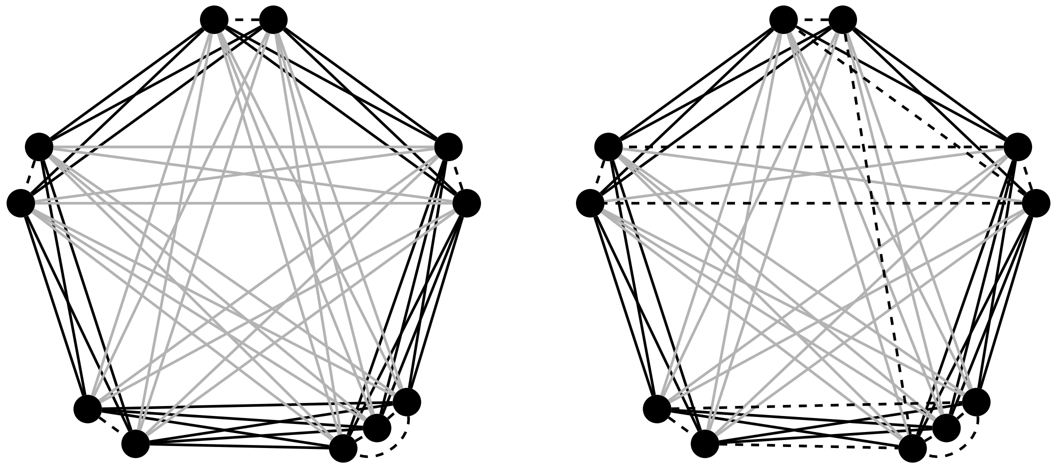

Figure 1: $G_{ex}(11)$ and another element of $\mathcal{G}_{11}$

The proof of Theorem 2 uses Razborov's method of flag algebras [18] together with a probabilistic argument.

Goodman [10] also raised the question of establishing k-coloured analogues of Theorem 1 for $k \geq 4$. Let $k \geq 3$ and $n \in \mathbb{N}$. Fox [7] gave an upper bound on $M_k(K_3, n)$ by considering the following graphs: Set $m := R_{k-1}(K_3) - 1$. Consider a $(k-1)$-coloured copy K of K_m on $[m]$ without a monochromatic triangle. Replace the vertices of K with disjoint vertex classes $V_1, \dots, V_m$ such that $||V_i| - |V_j|| \leq 1$ for all $1 \leq i, j \leq m$ and $|V_1| + \cdots + |V_m| = n$. For all $1 \leq i \neq j \leq m$, add all possible edges between V_i and V_j using the colour of ij in K. For each $1 \leq i \leq m$, add all possible edges to V_i using a kth colour. Denote the resulting complete k-coloured graph by $G_{ex}(n, k)$. (Thus, $G_{ex}(n) = G_{ex}(n, 3)$.)

Question 4. Let $k \geq 4$ and $n \in \mathbb{N}$ be sufficiently large. Is $M_k(K_3, n)$ equal to the number of monochromatic triangles in $G_{ex}(n, k)$?

2. Notation

We will make the convention that the set of colours used in a k-colouring of the edges of a graph is $[k]$. In the case of a 3-colouring we will generally refer to the colours 1, 2 and 3 as "red", "blue" and "green". When H and H' are two k-coloured graphs, an *isomorphism* between them is a function $f : V(H) \to V(H')$ which is a graph isomorphism and respects the colouring. Two k-coloured graphs H and H' are *isomorphic* ($H \cong H'$) if and only if there is an isomorphism between them.

Given $r \in \mathbb{N}$, we denote the complete graph on r vertices by K_r and define $R(r, r) := R_2(K_r)$. Given k and $c \in [k]$, we define K_c^r to be the k-coloured complete graph in which every edge of K_r is given the colour c. We define $\mathcal{K}^r$ to be $\{K_c^r : c \in [k]\}$, that is to say the set of monochromatic K_r's. Suppose G is a k-coloured graph and let $v \in V(G)$ and $i \in [k]$. Then we will use $N_i(v)$ to denote the set of vertices in G that receive an edge of colour i from v.

For a graph G and a vertex set $V \subseteq V(G)$, we denote by $G[V]$ the subgraph of G induced by V. Given $v_1, \dots, v_m \in V(G)$ we write $G[v_1, \dots, v_m]$ for $G[\{v_1, \dots, v_m\}]$, and for disjoint subsets V

and W of $V(G)$ we denote by $G[V, W]$ the bipartite graph with vertex classes V and W whose edge set consists of those edges between V and W in G. When G is a k-coloured graph, we view $G[V]$ as a k-coloured graph with the edge colouring inherited from G, and do likewise for $G[v_1, \dots, v_m]$ and for $G[V, W]$.

Throughout the paper, we write, for example, $0 < \nu \ll \tau \ll \eta$ to mean that we can choose the constants ν, τ, η from right to left. More precisely, there are increasing functions f and g such that, given η, whenever we choose some $\tau \leq f(\eta)$ and $\nu \leq g(\tau)$, all calculations needed in our proof are valid. Hierarchies with more constants are defined in the obvious way. Finally, the set of all k-subsets of a set A is denoted by $[A]^k$.

In the proof of Theorem 2 we will omit floors and ceilings whenever this does not affect the argument.

3. Graph densities

From this point on we are exclusively concerned with 3-colourings, mostly colourings of complete graphs. Suppose H and G are 3-coloured complete graphs where $|H| \leq |G|$. Let $d(H, G)$ denote the number of sets $V \in [V(G)]^{|H|}$ such that $G[V] \cong H$, and define the *density of H in G* as

$$p(H, G) := \frac{d(H, G)}{\binom{|G|}{|H|}}.$$

This quantity has a natural probabilistic interpretation, namely it is the probability that if we choose a set $\mathbf{V} \in [V(G)]^{|H|}$ uniformly at random then $\mathbf{V}$ induces an isomorphic copy of H.

When $\mathcal{H}$ is a family of 3-coloured complete graphs H of some fixed size k with $k \leq |G|$, we define

$$p(\mathcal{H}, G) := \sum_{H \in \mathcal{H}} p(H, G),$$

that is to say the probability that a random $\mathbf{V} \in [V(G)]^k$ induces a coloured graph isomorphic to an element of $\mathcal{H}$. In the sequel we generally write "H' is an $\mathcal{H}$" as an abbreviation for "H' is isomorphic to some $H \in \mathcal{H}$", "G contains an $\mathcal{H}$" as an abbreviation for "G contains an induced isomorphic copy of an element of $\mathcal{H}$", and "an $\mathcal{H}$ in G" for "an induced copy of some element of $\mathcal{H}$ in G".

For $n \geq |H|$ we let $p^{\min}(H, n)$ be the minimum value of $p(H, G)$ over all 3-coloured complete graphs G on n vertices. When $\mathcal{H}$ is a family of 3-coloured complete graphs H of some fixed size $k \leq n$, we let $p^{\min}(\mathcal{H}, n)$ be the minimum value of $p(\mathcal{H}, G)$ over all 3-coloured complete graphs G on n vertices.

We now define a certain class $\mathcal{H}$ of "bad" 3-coloured complete graphs on 4 vertices. As motivation, we note that we are defining a set of 3-coloured graphs H such that $\max_{G \in \mathcal{G}_n} p(H, G) = 0$.

Let $\mathcal{H}(i, j, k)$ be the class of 3-coloured complete graphs on 4 vertices with a monochromatic triangle, i extra edges of that same colour, and j and k edges of the other colours, respectively (with $i + j + k = 3$, $j \geq k$). Define $\mathcal{H} := \mathcal{H}(2, 1, 0) \cup \mathcal{H}(0, 2, 1)$.

The following result about graph densities will be used in the proof of Theorem 2. It provides an (asymptotically) optimal lower bound on the density of monochromatic triangles, and also asserts that copies of colourings from the class $\mathcal{H}$ are rare in any colouring that comes close enough to achieving this bound. The proof is given in Section 4.

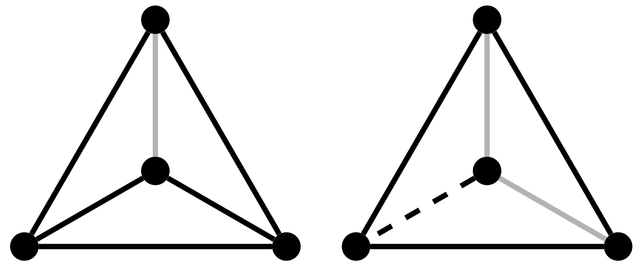

Figure 2: Representative elements of $\mathcal{H}(2,1,0)$ and $\mathcal{H}(0,2,1)$.

Proposition 5. *For all $\varepsilon > 0$ there is n_0 such that for all 3-coloured complete graphs G on at least n_0 vertices:*

1. $p(\mathcal{K}^3, G) \geq 0.04 - \varepsilon$.
2. *If* $p(\mathcal{K}^3, G) \leq 0.04$, *then* $p(\mathcal{H}, G) < \varepsilon$.

4. Flag algebras

In this section we use the method of *flag algebras* due to Razborov [18] to prove Proposition 5. The flag machinery described in subsections 4.1 and 4.2 is due to Razborov, as is the idea of using semidefinite programming for search of valid inequalities using this framework.

4.1. Some background

We start by describing how the main concepts of the general theory of flag algebras look in the case of 3-coloured complete graphs. Let $\mathcal{M}_l$ be the set of isomorphism classes of 3-coloured complete graphs on l vertices. It is helpful to know $|\mathcal{M}_l|$ for small values of l; computing this value is a classical enumeration problem [21], in particular $|\mathcal{M}_l| = 1, 1, 3, 10, 66, 792$ for $l = 0, 1, 2, 3, 4, 5$.

A *type* σ is a 3-coloured complete graph whose underlying set is of the form $[k] = \{1, 2, \ldots, k\}$ for some k, where we write $|\sigma| = k$. A σ*-flag* is a 3-coloured complete graph which contains a labelled copy of σ, or more formally a pair (M, θ) where M is a 3-coloured complete graph and θ is an injective map from $[k]$ to $V(M)$ that respects the edge-colouring of σ. Two σ-flags are *isomorphic* if there is an isomorphism that respects the labelling. More formally, f is a *flag isomorphism* from (M_1, θ_1) to (M_2, θ_2) if $f : V(M_1) \to V(M_2)$ is an isomorphism of coloured graphs and $f \circ \theta_1 = \theta_2$.

We denote by $\mathcal{F}_l^\sigma$ the set of isomorphism classes of σ-flags with l vertices. Note that if 0 is the empty type then $\mathcal{F}_l^0 = \mathcal{M}_l$. The flags of most interest to us are the elements of $\mathcal{F}_4^\sigma$ for various σ with $|\sigma| = 3$; it is easy to see that if $|\sigma| = 3$ then $|\mathcal{F}_4^\sigma| = 27$.

The notion of graph density described in the preceding section extends to σ-flags in a straightforward way. Given σ-flags $F \in \mathcal{F}_l^\sigma$ and $G \in \mathcal{F}_m^\sigma$ for $m \geq l$, we define $p(F, G)$ to be the density of isomorphic copies of F in G. More formally let $G = (M, \theta)$, choose uniformly at random a set $\mathbf{V} \in [V(M)]^l$ such that $\mathbf{V}$ contains $\mathrm{im}(\theta)$, and define $p(F, G)$ to be the probability that $(M[\mathbf{V}], \theta)$ is isomorphic (as a σ-flag) to F. By convention we will set $p(F, G) = 0$ in case $m < l$.

It is routine to see that if $l \leq m \leq n$, $F \in \mathcal{F}_l^\sigma$ and $H \in \mathcal{F}_n^\sigma$ then

$$p(F, H) = \sum_{G \in \mathcal{F}_m^\sigma} p(F, G) p(G, H). \tag{1}$$

This *chain rule* plays a central role in the theory.

More generally, given flags $F_i \in \mathcal{F}^\sigma_{l_i}$ for $1 \le i \le n$ and $G = (M, \theta) \in \mathcal{F}^\sigma_m$ where $m \ge \sum_i l_i - (n-1)|\sigma|$, we define a "joint density" $p(F_1, \dots, F_n; G)$. This is the probability that if we choose an n-tuple $(\mathbf{V}_1, \dots, \mathbf{V}_n)$ of subsets of $V(M)$ uniformly at random, subject to the conditions $\mathbf{V}_i \in [V(M)]^{l_i}$ and $\mathbf{V}_i \cap \mathbf{V}_j = \mathrm{im}(\theta)$ for $i \neq j$, then $(M[\mathbf{V}_i], \theta)$ is isomorphic to F_i for all i.

A sequence (G_n) of σ-flags is said to be *increasing* if the number of vertices in G_n tends to infinity, and *convergent* if the sequence of densities $(p(F, G_n))$ converges for every σ-flag F. A routine argument along the lines of the Bolzano-Weierstrass theorem shows that every increasing sequence has a convergent subsequence. If (G_n) is convergent, then we can define a map Φ on σ-flags by setting $\Phi(F) = \lim_{n\to\infty} p(F, G_n)$. We note that when $F \in \mathcal{F}^\sigma_l$ and $l \le m$, it follows readily from equation (1) that

$$\Phi(F) = \sum_{G \in \mathcal{F}^\sigma_m} p(F, G)\Phi(G). \tag{2}$$

Equation (2) suggests that in some sense "$F = \sum_{G\in\mathcal{F}^\sigma_m} p(F,G)G$", and the definition of the *flag algebra* $\mathcal{A}^\sigma$ makes this precise. We define $\mathcal{F}^\sigma_\infty = \bigcup_n \mathcal{F}^\sigma_n$, let $\mathbb{R}\mathcal{F}^\sigma_\infty$ be the real vector space consisting of finite formal linear combinations of elements of $\mathcal{F}^\sigma_\infty$, and then define $\mathcal{A}^\sigma$ to be the quotient of $\mathbb{R}\mathcal{F}^\sigma_\infty$ by the subspace $\mathcal{K}^\sigma$ generated by all elements of the form $F - \sum_{G\in\mathcal{F}^\sigma_m} p(F,G)G$. We will not be distinguishing between a flag F, its isomorphism class $[F] \in \mathcal{F}^\sigma$, the element $1[F] \in \mathbb{R}\mathcal{F}^\sigma_\infty$ and the element $1[F] + \mathcal{K}^\sigma \in \mathcal{A}^\sigma$.

If Φ is the map on σ-flags induced by a convergent sequence as above, then Φ extends by linearity to a map $\Phi : \mathbb{R}\mathcal{F}^\sigma_\infty \to \mathbb{R}$. The linear map Φ vanishes on $\mathcal{K}^\sigma$ by equation (2), and hence induces a linear map $\Phi : \mathcal{A}^\sigma \to \mathbb{R}$. So far $\mathcal{A}^\sigma$ is only a real vector space; we make it into an $\mathbb{R}$-algebra by defining a product as follows. Let $F_1 \in \mathcal{F}^\sigma_{l_1}$, $F_2 \in \mathcal{F}^\sigma_{l_2}$, let $m \ge l_1 + l_2 - |\sigma|$, and define

$$F_1 \cdot F_2 := \sum_{G\in\mathcal{F}^\sigma_m} p(F_1, F_2; G)G.$$

This can be shown [18, Lemma 2.4] to give a well-defined multiplication operation on $\mathcal{A}^\sigma$ independent of the choice of m, and it can also be shown [18, Theorem 3.3 part a] that if $\Phi : \mathcal{A}^\sigma \to \mathbb{R}$ is induced by a convergent sequence then $\Phi(F_1 \cdot F_2) = \Phi(F_1)\Phi(F_2)$, that is Φ is an algebra homomorphism from $\mathcal{A}^\sigma$ to $\mathbb{R}$. The converse is also true [18, Theorem 3.3 part b]: if Φ is such a homomorphism and $\Phi(F) \ge 0$ for all σ-flags F, then there exists an increasing and convergent sequence (G_n) such that $\Phi(F) = \lim_n p(F, G_n)$ for all flags F.

Following Razborov we let $\mathrm{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})$ be the set of homomorphisms induced by convergent sequences of σ-flags, and define a preordering on $\mathcal{A}^\sigma$ by stipulating that $A \le B$ if and only if $\Phi(A) \le \Phi(B)$ for all $\Phi \in \mathrm{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})$.

4.2. Averaging and lower bounds

The algebra $\mathcal{A}^\sigma$ has an identity element $1_\sigma = (\sigma, id_\sigma)$, and it is easy to see that $\Phi(1_\sigma) = 1$ for all $\Phi \in \mathrm{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})$. Accordingly we will identify the real number r and the element $r1_\sigma$. With this convention, the task of finding asymptotic lower bounds for quantities like the density of monochromatic triangles amounts to proving inequalities of the form "$F \ge r$ in $\mathcal{A}^0$" for some sum of 0-flags F and real number r. We will prove that

$$K^3_{red} + K^3_{blue} + K^3_{green} \ge 0.04.$$

Given a σ-flag $F = (M, \theta)$, we let $F|_0 = M$. We define $[\![F]\!]_\sigma = q_\sigma(F)M$, where $q_\sigma(F)$ is the probability that a random injective function $\boldsymbol{\theta}$ from $[|\sigma|]$ to $V(M)$ gives a σ-flag $(M, \boldsymbol{\theta})$ and this flag is isomorphic to F. This map on σ-flags extends to a linear map from $\mathcal{A}^\sigma$ to $\mathcal{A}^0$.

A key fact is that for every type σ and every $A \in \mathcal{A}^\sigma$, we have the inequality

$$[\![A^2]\!]_\sigma \geq 0, \tag{3}$$

where $A^2 = A \cdot A$. We will ultimately prove our desired lower bound by adding many inequalities of this form for various types σ and elements A of $\mathcal{A}^\sigma$.

Inequality (3) can be proved by elementary means; roughly speaking we average the square of the number of copies of F containing a particular copy of σ over all such copies and discard terms of low order. It can also be proved [18, Theorem 3.14] using the notion of *random homomorphism* discussed below in subsection 4.4.

We will prove that $K^3_{red} + K^3_{blue} + K^3_{green} \geq 0.04$ by proving an equation of the form

$$K^3_{red} + K^3_{blue} + K^3_{green} - 0.04 - \sum_i [\![L_i^2]\!]_{\sigma_i} = \sum \lambda_k M_k, \tag{4}$$

where the σ_i's are types, $L_i \in \mathcal{A}^{\sigma_i}$, the M_k's are 3-coloured complete graphs and $\lambda_k \geq 0$ for all k. Equation (4) clearly implies that $K^3_{red} + K^3_{blue} + K^3_{green} \geq 0.04$, which is the translation of claim 1 in Proposition 5 into the flag language.

Since there are increasing sequences of 3-coloured complete graphs in which the density of monochromatic triangles approaches 0.04, there are $\Phi \in \mathrm{Hom}^+(\mathcal{A}^0, \mathbb{R})$ such that $\Phi(K^3_{red} + K^3_{blue} + K^3_{green}) = 0.04$. For any such Φ we must have

(i) $\Phi([\![L_i^2]\!]_{\sigma_i}) = 0$ for all i,

(ii) $\lambda_k = 0$ for all k such that $\Phi(M_k) > 0$, and

(iii) $\Phi(M_k) = 0$ for all k such that $\lambda_k > 0$.

The last of these points is the key to proving the second claim in Proposition 5. We will verify that for all $H \in \mathcal{H}$, H is a linear combination of M_k's such that $\lambda_k > 0$. It follows that for all such H, $\Phi(H) = 0$ for any Φ with $\Phi(K^3_{red} + K^3_{blue} + K^3_{green}) = 0.04$. This assertion is exactly the translation into the flag language of claim 2 in Proposition 5.

4.3. Proof of Proposition 5

To prove Proposition 5 we need to specify ten types, several hundred flags, and ten 27×27 matrices. Rather than attempting to render the details of the proof in print, we have chosen to describe its structure here and make all the data available online, together with programs which can be used to verify them.

Let σ be a type and let $L_1, \ldots, L_t \in \mathcal{A}^\sigma$, where each L_i is a real linear combination of a fixed set of σ-flags $F_1, \ldots, F_n$. By standard facts in linear algebra,

$$L_1^2 + \ldots + L_t^2 = \sum_{ij} Q_{ij} F_i \cdot F_j$$

for some $n \times n$ positive semidefinite symmetric matrix Q, and conversely any expression of the form $\sum_{ij} Q_{ij} F_i \cdot F_j$ for a positive semidefinite Q is a sum of squares.

In our case we will have ten types τ_r for $1 \leq r \leq 10$, each with $|\tau_r| = 3$. The types are chosen to include representative elements of each isomorphism class of 3-coloured triangles.

For each type τ_r we will have a complete list $F_1^r, \ldots, F_{27}^r$ of the τ_r-flags on 4 vertices. In line with the discussion in subsection 4.2, we will specify for each r a 27×27 symmetric matrix Q^r and will actually prove an equation of the form

$$K_{red}^3 + K_{blue}^3 + K_{green}^3 - 0.04 - \sum_r [\![Q_r]\!]_{\sigma_r} = \sum \lambda_k M_k, \tag{5}$$

where $Q_r = \sum_{ij} Q_{ij}^r F_i^r F_j^r$, each matrix Q^r is positive semidefinite, and each coefficient λ_k is non-negative. The matrices Q^r will have rational entries, so the whole computation can be done exactly using rational arithmetic.

By the definition of flag multiplication, each product $F_i^r \cdot F_j^r$ can be written as a linear combination of elements of $\mathcal{F}_5^{\sigma_r}$, so each term $\sum_r [\![Q_r]\!]_{\sigma_r}$ is a linear combination of elements of $\mathcal{M}_5$. The 3-coloured complete graphs M_k appearing in equation (5) will be the 792 elements of $\mathcal{M}_5$. By extending $p(A, B)$ linearly from $\mathcal{F}_m^\sigma \times \mathcal{F}_m^\sigma$ to $\mathcal{F}_m^\sigma \times \mathbb{R}\mathcal{F}_m^\sigma$, $p(A, B)$ can stand for "the coefficient of A in B", $A \in \mathcal{F}_m^\sigma$ and $B \in \mathbb{R}\mathcal{F}_m^\sigma$ (we use this notation for $m = 5$ only).

Given the coefficient matrices Q^r, we must first verify that they are positive semidefinite (a routine calculation). We must then expand the left hand side of equation (5) in the form $\sum \lambda_k M_k$, and check that $\lambda_k \geq 0$ for all k. Clearly

$$\lambda_k = p(K_{red}^3, M_k) + p(K_{blue}^3, M_k) + p(K_{green}^3, M_k) - 0.04 - \sum_r p(M_k, [\![Q_r]\!]_{\sigma_r}),$$

and

$$p(M_k, [\![Q_r]\!]_{\sigma_r}) = \sum_{ij} Q_{ij}^r p(M_k, [\![F_i^r \cdot F_j^r]\!]_{\sigma_r}),$$

so the main computational task in verifying the proof is to compute the coefficients $p(M_k, [\![F_i^r \cdot F_j^r]\!]_{\sigma_r})$.

A useful lemma of Razborov gives a probabilistic interpretation of $p(M_k, [\![F_i^r \cdot F_j^r]\!]_{\sigma_r})$ which obviates the need to compute $F_i^r \cdot F_j^r$ and $[\![F_i^r \cdot F_j^r]\!]_{\sigma_r}$ before computing $p(M_k, [\![F_i^r \cdot F_j^r]\!]_{\sigma_r})$. The lemma states that for any type τ, any τ-flags K_1 and K_2 and any m which is large enough to express $[\![K_1 \cdot K_2]\!]_\tau$ as a linear combination of elements of $\mathcal{M}_m$, the coefficient $p(L, [\![K_1 \cdot K_2]\!]_\tau)$ of $L \in \mathcal{M}_m$ is the probability that choosing a random injection $\boldsymbol{\theta}$ from $V(\tau)$ to $V(L)$ and then random sets $\mathbf{X}$ and $\mathbf{Y}$ of the appropriate size with $\mathbf{X} \cap \mathbf{Y} = \mathrm{im}(\boldsymbol{\theta})$ gives flags $(L[\mathbf{X}], \theta)$ and $(L[\mathbf{Y}], \theta)$ such that $(L[\mathbf{X}], \theta)$ is isomorphic to K_1 and $(L[\mathbf{Y}], \theta)$ is isomorphic to K_2. The proof is straightforward.

To complete the proof of Proposition 5, we must now compute the coefficients λ_k and verify that for all k

(i) $\lambda_k \geq 0$;

(ii) For all $H \in \mathcal{H}$, if $p(H, M_k) > 0$ then $\lambda_k > 0$.

The data for the proof and a Maple worksheet which verifies it can be found online at the URL `http://www.math.cmu.edu/users/jcumming/ckpsty`. Further, the version of this paper on the arXiv [5] has an appendix with the data for the proof.

4.4. Semidefinite programming

The proof described in the preceding section was obtained using *semidefinite programming*. In our case we fixed the types σ_r and flags F_j^r, and set up a semidefinite programming problem where the unknowns are the matrices $(Q^1, \ldots, Q^{10})$ and the goal is to maximise a lower bound for $K^3_{red} + K^3_{blue} + K^3_{green}$. Using the CSDP and SDPA solvers, we produced lower bounds of the form $0.04 - \varepsilon$ where ε is very small (about 10^{-9} with CSDP, about 10^{-6} with standard precision SDPA, and about 10^{-17} with the high precision version SDPA-QD).

We now needed to perturb the coefficients in our matrices Q^r to achieve the optimal value 0.04 for the lower bound, i.e., we needed to find positive semidefinite matrices Q^r such that

$$p(K^3_{red}, M_k) + p(K^3_{blue}, M_k) + p(K^3_{green}, M_k) - \sum_r p(M_k, [\![Q_r]\!]_{\sigma_r}) - 0.04 \geq 0$$

for all k. A naive approach fails: if we make a small perturbation to satisfy these inequalities, then the resulting matrices Q^r typically have at least one small negative eigenvalue.

This issue is related to the theory of *random homomorphisms* from [18, Section 3.2], as explained in [19, Section 4]. If $\Phi \in \mathrm{Hom}^+(\mathcal{A}^0, \mathbb{R})$ and σ is a type such that (viewing σ as an element of $\mathcal{A}^0$) $\Phi(\sigma) > 0$, then we may use Φ to construct a certain probability measure on $\mathrm{Hom}^+(\mathcal{A}^\sigma, \mathbb{R})$, which we may view (using the probabilistic language) as a *random homomorphism* $\mathbf{\Phi}^\sigma$. One of the properties of $\mathbf{\Phi}^\sigma$ is that for any $F \in \mathcal{A}^\sigma$ the expected value of $\mathbf{\Phi}^\sigma(F)$ is given by the formula

$$\mathbb{E}(\mathbf{\Phi}^\sigma(F)) = \frac{\Phi([\![F]\!]_\sigma)}{\Phi([\![1]\!]_\sigma)}.$$

So, we can view the inequality $[\![F^2]\!]_\sigma \geq 0$ as an averaging argument analogous to the Cauchy-Schwartz theorem [18, Theorem 3.14].

Let Q_r be one of the quadratic forms appearing in a proof of the optimal bound and let $Q_r = \sum_k m_k^2$ where each m_k is a linear combination of the flags F_i^r. Recall from subsection 4.2 that if Φ is such that $\Phi(K^3_{red}+K^3_{blue}+K^3_{green}) = 0.04$, then $\Phi([\![Q_r]\!]_{\sigma_r}) = 0$ for all r. If $\Phi([\![Q_r]\!]_{\sigma_r}) = 0$, then it holds with probability one that $\mathbf{\Phi}^{\sigma_r}(m_k) = 0$. This yields that all eigenvectors of the matrix Q^r corresponding to non-zero eigenvalues must lie in a certain linear space, and the existence of this kind of relation explains the problem with perturbing an approximate solution to an exact one.

In our case we could not derive enough relations by the method of the preceding paragraph, and we guessed the necessary extra relations by inspecting numerical data computed to very high accuracy with the SDPA-QD software package. Oleg Pikhurkho [16] later offered us an explanation of the extra relations, which we give here with his permission.

Recall that we are seeking an equation of the form

$$K^3_{red} + K^3_{blue} + K^3_{green} - 0.04 - \sum_r [\![Q_r]\!]_{\sigma_r} = \sum_k \lambda_k M_k,$$

where each λ_k is non-negative and each Q_r is a positive semidefinite quadratic form in some set of σ_r-flags. As we already mentioned, if Φ is such that $\Phi(K^3_{red} + K^3_{blue} + K^3_{green}) = 0.04$, then $\Phi([\![Q_r]\!]_{\sigma_r}) = 0$ for all r. Suppose that we can find a perturbation Φ_ε of such a Φ for all ε in some interval $[0, \alpha)$, such that $\Phi_\varepsilon(K^3_{red} + K^3_{blue} + K^3_{green}) = 0.04 + \Theta(\varepsilon^3)$. In our particular problem, one such perturbation corresponds to extremal examples obtained from $G_{ex}(n)$ as follows: choose $\lfloor \varepsilon n/5 \rfloor$ vertices in two of the five monochromatic cliques and recolour edges between these two sets of vertices with the colour used inside the cliques.

It follows that $\Phi_\varepsilon([\![Q_r]\!]_{\sigma_r}) = O(\varepsilon^3)$ in the equation we seek. However, for certain choices of r and Φ, there exist choices of linear forms m in the σ_r-flags F_i^r such that $\Phi_\varepsilon([\![m^2]\!]_{\sigma_r}) = \Theta(\varepsilon^2)$. If $Q_r = \sum_k m_k^2$ as before, then by letting $\varepsilon \to 0$ we will obtain that every m_k must be orthogonal to all such m, which further restricts all eigenvectors of Q^r corresponding to non-zero eigenvalues. In our case we obtained additional constraints for those σ_r and Φ such that $\Phi(\sigma_r) = 0$, which were sufficient to complete the proof. However, this method can lead in general to discovering additional constraints for σ_r with $\Phi(\sigma_r) > 0$ as well.

5. Proof of Theorem 2

5.1. *Finding a standard subgraph of G*

Define constants $\varepsilon, \varepsilon_1, \varepsilon_2, \varepsilon_3, \varepsilon_4, \varepsilon_5, \varepsilon_6, \varepsilon_7, \varepsilon_8, \varepsilon_9, \varepsilon_{10}$ and integers n_0, n_1, n_2 such that n_0 and ε satisfy the assertion of Proposition 5 and

$$0 < 1/n_0 \ll \varepsilon \ll 1/n_1 \ll \varepsilon_1 \ll \varepsilon_2 \ll \varepsilon_3 \ll 1/n_2 \ll \varepsilon_4 \ll \varepsilon_5 \ll \varepsilon_6 \ll \varepsilon_7 \ll \varepsilon_8 \ll \varepsilon_9 \ll \varepsilon_{10} \ll 1. \quad (6)$$

Let G be a 3-coloured complete graph on $n \geq n_0$ vertices with $p(\mathcal{K}^3, G)$ minimised. We may assume the three colours used are red, green and blue. Note that, by the minimality of G, $p(\mathcal{K}^3, G) \leq p(\mathcal{K}^3, G_{ex}(n)) \leq 0.04$. Since $n \geq n_0$, Proposition 5 implies that $p(\mathcal{K}^3, G) \geq 0.04 - \varepsilon$ and $p(\mathcal{H}, G) < \varepsilon$.

Let us call an induced subgraph $G' \subseteq G$ ε_1-*standard* if

(i) $p(\mathcal{K}^3, G') \leq 0.04 + \varepsilon_1$;

(ii) $p(\mathcal{H}, G') = 0$.

Now we randomly pick n_1 vertices from G to induce a subgraph G'.

Claim 1. $\mathbb{P}(G'$ is ε_1-standard$) \geq 1 - \varepsilon_2$.

Proof. Since $1/n_1 \ll \varepsilon_1$, Proposition 5 implies that $p^{\min}(\mathcal{K}^3, n_1) > 0.04 - \varepsilon_1^2$. Thus, $Z := p(\mathcal{K}^3, G') - (0.04 - \varepsilon_1^2) > 0$. Note that $\mathbb{E}(Z) \leq \varepsilon_1^2$ since $\mathbb{E}(p(\mathcal{K}^3, G')) = p(\mathcal{K}^3, G) \leq 0.04$. Hence, by Markov's inequality,

$$\mathbb{P}(Z \geq \varepsilon_1) \leq \frac{\varepsilon_1^2}{\varepsilon_1} = \varepsilon_1$$

and therefore

$$\mathbb{P}(p(\mathcal{K}^3, G') \leq 0.04 + \varepsilon_1) \geq 1 - \varepsilon_1.$$

By Markov's inequality,

$$\mathbb{P}\left(p(\mathcal{H}, G') < \frac{2\varepsilon}{\varepsilon_2}\right) \geq 1 - \varepsilon_2/2.$$

Note that (6) implies that $2\varepsilon/\varepsilon_2 < 1/\binom{n_1}{4}$. Thus, the claim follows. □

In the next two subsections we will build up structure in our ε_1-standard subgraphs G', thereby obtaining that each such G' has 'similar' structure to $G_{ex}(n_1)$.

5.2. Properties of maximal monochromatic cliques in G'

Consider any ε_1-standard subgraph G' of G on n_1 vertices. Let $\mathcal{X}$ be the set of maximal monochromatic cliques of order at least 4 in G'. So a clique X_1 in $\mathcal{X}$ cannot strictly contain another clique $X_2 \in \mathcal{X}$. However, $\mathcal{X}$ may contain cliques that intersect each other. Since n_1 is sufficiently large, G' contains a $\mathcal{K}^4$ by Ramsey's theorem. Thus, $|\mathcal{X}| \geq 1$.

Claim 2. Let $X \in \mathcal{X}$ and $y \in V(G') \backslash V(X)$. All but one of the edges xy with $x \in V(X)$ have the same colour, which is different from the colour of X. The remaining edge is either of that same colour or of the colour of X.

Proof. Assume X is coloured red. By definition of $\mathcal{X}$, we cannot have that all edges between X and y are red. This implies that at most one such edge is red (else G' contains an $\mathcal{H}(2,1,0)$, a contradiction to (ii)). This in turn implies that there does not exist both green and blue edges between X and y (else G' contains an $\mathcal{H}(0,2,1)$). The claim now follows. □

Claim 3. Suppose $X_1, X_2 \in \mathcal{X}$ have different colours. Then X_1 and X_2 are vertex-disjoint.

Proof. Since X_1 and X_2 have different colours, $|V(X_1) \cap V(X_2)| \leq 1$. Suppose for a contradiction there exists a vertex $x \in V(X_1) \cap V(X_2)$. Suppose X_1 is red and X_2 is blue. For each $x_1 \in X_1 - x$, since x_1x is red, Claim 2 implies that all but at most one of the edges from x_1 to X_2 are red. Thus, there exists distinct $x', x'' \in X_1 - x$ and $x''' \in X_2 - x$ such that $x'x'''$ and $x''x'''$ are red. But since xx''' is blue, $G'[x, x', x'', x''']$ is an $\mathcal{H}(2,1,0)$, a contradiction to (ii). □

Claim 4.

(a) If $X_1, X_2 \in \mathcal{X}$ have different colours, then there is a vertex $v_1 \in V(X_1)$ and a vertex $v_2 \in V(X_2)$ such that all edges between $X_1 - v_1$ and $X_2 - v_2$ have the same colour, and this colour is different from the colours of X_1 and X_2.

(b) If $X_1, X_2 \in \mathcal{X}$ have the same colour, then either X_1 and X_2 share exactly one vertex v, and all edges between $X_1 - v$ and $X_2 - v$ have a common colour, or X_1 and X_2 are disjoint, there is a (possibly empty) matching of the colour of X_1 and X_2 between X_1 and X_2, and all other edges between X_1 and X_2 have the same colour, different from the colour of X_1 and X_2.

Proof. If $X_1, X_2 \in \mathcal{X}$ have different colours, then by Claim 3, X_1 and X_2 are vertex-disjoint. Suppose X_1 is red and X_2 is blue. Firstly, note that there does not exist distinct $x_1', x_1'' \in X_1$ and $x_2', x_2'' \in X_2$ such that both $x_1'x_2'$ and $x_1''x_2''$ are blue. Indeed, if such edges exist then by Claim 2, $x_1'x_2''$ and $x_1''x_2'$ are red. Again by Claim 2, this implies that every edge from x_2' to $X_1 - x_1''$ is blue and every edge from x_2'' to $X_1 - x_1'$ is blue. Let $a, b \in X_1 - \{x_1', x_1''\}$. Then $G'[a, b, x_2', x_2'']$ is an $\mathcal{H}(2,1,0)$, a contradiction.

An identical argument implies that there does not exist distinct $x_1', x_1'' \in X_1$ and $x_2', x_2'' \in X_2$ such that both $x_1'x_2'$ and $x_1''x_2''$ are red. By Claim 2 this implies that there exists at most one vertex $v_1 \in X_1$ such that v_1 sends at least one red edge to X_2 and there exists at most one vertex $v_2 \in X_2$ such that v_2 sends at least one blue edge to X_1. This implies that all the edges from $X_1 - v_1$ to $X_2 - v_2$ are green, and so (a) is satisfied.

Suppose $X_1, X_2 \in \mathcal{X}$ have the same colour, red say. Notice that $|V(X_1) \cap V(X_2)| \leq 1$, since otherwise a vertex in $V(X_1) \backslash V(X_2)$ would send at least two red edges to X_2, a contradiction to Claim 2. If $|V(X_1) \cap V(X_2)| = 1$ then it is easy to see that, by Claim 2, the first part of (b) holds.

If X_1 and X_2 are disjoint, then by Claim 2, no vertex in X_1 sends more than one red edge to X_2 and no vertex in X_2 sends more than one red edge to X_1. Thus, the red edges between X_1 and X_2 form a (possibly empty) matching. Applying Claim 2 again shows that the second part of (b) holds. $\square$

5.3. Properties of the clique graph

We now define a new 3-coloured complete graph F which we refer to as the *clique graph*. The vertex set of F consists of the elements of $\mathcal{X}$ together with the vertices in Y where $Y \subseteq V(G')$ is the set of vertices in G' not contained in any of the cliques in $\mathcal{X}$. If $x, y \in Y$ then, in F, we colour xy with the colour of xy in G. If $X_1, X_2 \in \mathcal{X}$ then, in F, we colour the edge X_1X_2 with the colour of the majority of the edges between X_1 and X_2 in G. (Note that this colour is well-defined by Claim 4.) Finally, given a vertex $y \in Y$ and $X \in \mathcal{X}$, in F we colour the edge yX with the colour of the majority of the edges between y and X in G. (This colour is well-defined by Claim 2.)

Claim 5. No $\mathcal{K}^3$ in F contains a vertex $X \in \mathcal{X}$. Moreover, F contains no $\mathcal{K}^4$.

Proof. The first part of the claim follows from Claim 2 since otherwise there would be an $\mathcal{H}(2,1,0)$ in G', a contradiction to (ii). The second part of the claim follows from the first part together with the definition of Y. $\square$

For every clique $X \in \mathcal{X}$, the edges in F leaving X must have different colours from X. Thus, we have $|\mathcal{X}| + |Y| \leq 35$. Indeed, otherwise each X in $\mathcal{X}$ is incident to 18 edges of the same colour in F. But then, since $R(4,4) = 18$, F contains a $\mathcal{K}^4$ or a $\mathcal{K}^3$ containing X, a contradiction to Claim 5. If $|\mathcal{X}| \leq 4$, then $p(\mathcal{K}^3, G') \geq 4\binom{\lfloor (n_1-34)/4 \rfloor}{3} / \binom{n_1}{3} > 0.04 + \varepsilon_1$, a contradiction to (i). Thus, $|\mathcal{X}| \geq 5$.

If there are three cliques in $\mathcal{X}$ of one colour, and another clique in $\mathcal{X}$ of a different colour, then it is easy to see by Claim 4 that there must be a monochromatic triangle between these four cliques, a contradiction to Claim 5. Similarly, we cannot have two cliques in $\mathcal{X}$ of one colour, and also cliques in $\mathcal{X}$ of the other two colours. Therefore, all cliques in $\mathcal{X}$ must have the same colour, say red.

Since $R(3,3) = 6$, if $|\mathcal{X}| \geq 6$, then again $F[\mathcal{X}]$ contains a $\mathcal{K}^3$, a contradiction. So $|\mathcal{X}| = 5$. Further, $Y = \emptyset$, since otherwise $F[\mathcal{X} \cup \{y\}]$ is 2-coloured and thus contains a $\mathcal{K}^3$ (for all $y \in Y$).

Claim 6. Let $\mathcal{X} = \{X_1, \ldots, X_5\}$. The following properties hold:

(α_1) $(1-\varepsilon_3)\frac{n_1}{5} \leq |X_i| \leq (1+\varepsilon_3)\frac{n_1}{5}$ for all $1 \leq i \leq 5$;

(α_2) $E(F)$ is 2-coloured with green and blue and consists of a green 5-cycle and a blue 5-cycle. We may assume that $X_1X_2X_3X_4X_5X_1$ is a green cycle and $X_1X_3X_5X_2X_4X_1$ is a blue cycle;

(α_3) Either the cliques in $\mathcal{X}$ are vertex-disjoint or there exists a unique vertex w that lies in each clique in $\mathcal{X}$ (and w is the only vertex which lies in more than one clique in $\mathcal{X}$).

Proof. Every clique in $\mathcal{X}$ contains at least $(1-\varepsilon_3)\frac{n_1}{5}$ vertices as otherwise

$$p(\mathcal{K}^3, G') \geq \left(\binom{(1-\varepsilon_3)\frac{n_1}{5}}{3} + 4\binom{(1+\varepsilon_3/4)\frac{n_1}{5}}{3} \right) \Big/ \binom{n_1}{3} \overset{(6)}{>} 0.04 + \varepsilon_1.$$

A similar calculation shows that every clique in $\mathcal{X}$ contains at most $(1+\varepsilon_3)\frac{n_1}{5}$ vertices. Every clique in $\mathcal{X}$ is red, thus $E(F)$ is 2-coloured with green and blue. Since F does not contain a monochromatic triangle, F must satisfy (α_2).

Suppose two of the cliques, say X_1 and X_2, share a vertex w. As X_1X_3 is blue and X_2X_3 is green, Claim 2 implies that, for every vertex $v \in X_3$ the edge $vw \in E(G)$ can be neither blue nor green, so it has to be red. But this implies that $w \in X_3$. By similar arguments, $w \in X_4 \cap X_5$. Thus, (α_3) holds. □

5.4. Obtaining structure in G from G′

Our next task is to find a special set $V' \subseteq V(G)$ such that $G[V']$ has 'similar' structure to $G_{ex}(n_2)$.

Claim 7. There exists a set $V' \subseteq V(G)$ such that the following properties hold:

(β_1) $|V'| = n_2$;

(β_2) V' has a partition into non-empty sets C_1, C_2, C_3, C_4, C_5 such that

- $\frac{|C_i|}{|C_j|} > 1 - \varepsilon_4$ for all $1 \leq i \neq j \leq 5$,
- all edges inside the C_i have the same colour, say red,
- all edges between C_i and C_{i+1} are green,
- all edges between C_i and C_{i+2} are blue (here indices are computed modulo 5);

(β_3) If we uniformly at random choose two vertices $u, v \in V(G)$, then with probability greater than $1 - \varepsilon_5$, the set $V' \cup \{u, v\}$ satisfies (β_2) as well.

Proof. Consider any ε_1-standard subgraph G' of G on n_1 vertices. Randomly select a set $W \subseteq V(G')$ of size n_2. Then with probability more than $1 - \varepsilon_4^3$, W satisfies (β_2). This follows from Claims 4 and 6. For example, by applying a Chernoff-type bound for the hypergeometric distribution (see e.g. [14, Theorem 2.10]), (α_1) implies that with probability greater than $1 - \varepsilon_4^4$, the first two conditions in (β_2) hold. Further, note that the probability that W contains the special vertex w from (α_3) (if it exists) is $n_2/n_1 \ll \varepsilon_4$ by (6).

Randomly select a set $W' \subseteq V(G)$ of size n_2. One can view this procedure as first randomly selecting a set $W'' \subseteq V(G)$ of size n_1, then randomly selecting a set $W' \subseteq W''$ of size n_2. By Claim 1, with probability at least $1 - \varepsilon_2$, $G[W'']$ is ε_1-standard.

Together, this implies that with probability greater than $(1 - \varepsilon_2)(1 - \varepsilon_4^3) > 1 - \varepsilon_4^2$ a randomly chosen set $W' \subseteq V(G)$ of size n_2 satisfies (β_2). Similarly, with probability greater than $1 - \varepsilon_4^2$ a randomly chosen set $W' \subseteq V(G)$ of size $n_2 + 2$ satisfies (β_2).

Consider all pairs $(V', \{u, v\})$ such that $\{u, v\}, V' \subseteq V(G)$ and $|V'| = n_2$. (Note here we allow for $V' \cap \{u, v\} \neq \emptyset$.) With probability greater than $1 - 3\varepsilon_4^2$, a randomly selected such pair $(V', \{u, v\})$ has the property that both V' and $V' \cup \{u, v\}$ satisfy (β_2). Since $3\varepsilon_4^2 \ll \varepsilon_5$, this implies that there exists a set $V' \subseteq V(G)$ satisfying (β_1)–(β_3). □

Let V' be as in Claim 7. Set

$$E_0 := \{uv \in E(G) : V' \cup \{u, v\} \text{ does not satisfy } (\beta_2)\}.$$

Then $|E_0| < \varepsilon_5 n^2$ by (β_3). Let

$$V_0 := \{v \in V(G) : v \text{ is incident to at least } \varepsilon_6 n \text{ edges in } E_0\}.$$

Then $|V_0| < \varepsilon_6 n$ since $\varepsilon_5 \ll \varepsilon_6$. For each $1 \leq i \leq 5$ define

$$F_i := \{v \in V(G) \setminus V_0 : vw \text{ is red for all } w \in C_i\}.$$

Note that $V(G) = V_0 \cup F_1 \cup F_2 \cup F_3 \cup F_4 \cup F_5$. Further, notice that the F_i are disjoint. (Indeed, if there is a vertex $x \in F_i \cap F_j$ for some $i \neq j$ then all edges incident to x are in E_0. But then $x \in V_0$, a contradiction.)

Claim 8. For all $1 \leq i \leq 5$,

$$(1-\varepsilon_7)n/5 \leq |F_i| \leq (1+\varepsilon_7)n/5.$$

Proof. Suppose $|F_i| < (1-\varepsilon_7)n/5$ for some $1 \leq i \leq 5$. By definition of the F_j and (β_3), there are at most $\varepsilon_5 n^2$ edges in F_j that are not red (for each $1 \leq j \leq 5$). Thus, in each F_j, there are at most $\varepsilon_5 n^3$ triples that do not form a red triangle. Hence, there are at least

$$\begin{aligned}\binom{|F_i|}{3} + 4\binom{|V(G)\backslash(V_0 \cup F_i)|/4}{3} - 5\varepsilon_5 n^3 &\geq \binom{(1-\varepsilon_7)n/5}{3} + 4\binom{(4/5+\varepsilon_7/5-\varepsilon_6)n/4}{3} - 5\varepsilon_5 n^3 \\ &\overset{(6)}{>} (0.04+\varepsilon_1)\binom{n}{3}\end{aligned}$$

red triangles in G, a contradiction. The upper bound follows similarly. □

For each $v \in V(G)$ and $1 \leq i \leq 5$, let $r_i(v) := |N_{red}(v) \cap F_i|$, $b_i(v) := |N_{blue}(v) \cap F_i|$ and $g_i(v) := |N_{green}(v) \cap F_i|$. On the basis of these quantities, we define another partition of $V(G)$ as follows. For each $1 \leq i \leq 5$, set

$$V_i := \left\{ v \in V(G) : \begin{array}{rl} r_i(v) & \geq 0.199n, \\ g_{i+1}(v) & \geq 0.199n, \\ b_{i+2}(v) & \geq 0.199n, \\ b_{i+3}(v) & \geq 0.199n, \text{ and} \\ g_{i+4}(v) & \geq 0.199n \end{array} \right\}.$$

Claim 9. For each $1 \leq i \leq 5$, $F_i \subseteq V_i$.

Proof. Given any $v \in F_i$, v is incident to at most $\varepsilon_6 n$ edges in E_0. Thus, there are at most $\varepsilon_6 n$ vertices in F_i that v does not send a red edge to. Hence, Claim 8 implies that $r_i(v) \geq 0.199n$. Similar arguments give $g_{i+1}(v), b_{i+2}(v), b_{i+3}(v), g_{i+4}(v) \geq 0.199n$. □

Set $V^* := V(G) \backslash (V_1 \cup V_2 \cup V_3 \cup V_4 \cup V_5)$. Let E^* be the set of edges xy in $G[V_1 \cup V_2 \cup V_3 \cup V_4 \cup V_5]$ such that $x \in V_i$ and $y \in V_j$ for some $1 \leq i, j \leq 5$ and so that the colour of xy differs from that of the edges between C_i and C_j.

Claim 10. For each $1 \leq i \leq 5$, $G[V_i]$ is a red clique.

Proof. Claims 8 and 9 imply that $(1-\varepsilon_7)n/5 \leq |V_i| \leq (1+4\varepsilon_7)n/5$ for all $1 \leq i \leq 5$. Suppose for a contradiction that there is a blue edge vw with $v, w \in V_1$. Recolouring vw red creates at most

$$|V^*| + |V_1| + (0.004 + \varepsilon_7/5)n < 0.205n$$

new red triangles. (The $(0.004+\varepsilon_7/5)n$ term counts the maximum number of red edges a vertex in V_1 can send to $V_2 \cup V_3 \cup V_4 \cup V_5$.) On the other hand, the recolouring destroys at least

$$|F_3| + |F_4| - 2(0.002 + 2\varepsilon_7/5)n > 0.395n$$

blue triangles, contradicting the minimality of G. □

Claim 11. $E^* \subseteq E_0$.

Proof. Suppose $xy \in E^*$ where $x \in V_i$ and $y \in V_j$ for some $1 \leq i, j \leq 5$. The colour of xy differs from that of the edges between C_i and C_j. But Claim 10 implies that x only sends red edges to C_i and y only sends red edges to C_j. Thus, $xy \in E_0$. □

Claim 12. $V^* = \emptyset$.

Proof. Suppose that $v \in V^*$. We count the number of monochromatic triangles t_v containing v and two vertices from outside of V^*. First, if we were to recolour all edges from v to the smallest V_i red, from v to $V_{i+1} \cup V_{i+4}$ green, and from v to $V_{i+2} \cup V_{i+3}$ blue, then we would get at most

$$\binom{|V_i|}{2} + |E^*| \leq \binom{\lfloor n/5 \rfloor}{2} + |E_0| < (0.02 + \varepsilon_5)n^2$$

monochromatic triangles containing v and two vertices from outside of V^*, and at most $|V^*|n < \varepsilon_6 n^2$ new triangles containing v and another vertex from V^*. Thus, the minimality of G implies that

$$t_v < (0.02 + \varepsilon_5 + \varepsilon_6)n^2. \tag{7}$$

Recall our notation $r_i(v), g_i(v), b_i(v)$. Note that

$$\begin{aligned} t_v \geq\ & 0.5(r_1(v)^2 + r_2(v)^2 + r_3(v)^2 + r_4(v)^2 + r_5(v)^2) \\ & + g_1(v)g_2(v) + g_2(v)g_3(v) + g_3(v)g_4(v) + g_4(v)g_5(v) + g_5(v)g_1(v) \\ & + b_1(v)b_3(v) + b_2(v)b_4(v) + b_3(v)b_5(v) + b_4(v)b_1(v) + b_5(v)b_2(v) \\ & - 2\varepsilon_5 n^2. \end{aligned} \tag{8}$$

where the last term occurs since $\binom{r_i(v)}{2} \geq 0.5r_i^2 - n$ for each $1 \leq i \leq 5$ and as $|E^*| < \varepsilon_5 n^2$.

Our next task is to find a lower bound on

$$\begin{aligned} & 0.5(r_1(v)^2 + r_2(v)^2 + r_3(v)^2 + r_4(v)^2 + r_5(v)^2) + \gamma_1\gamma_2 + \gamma_2\gamma_3 + \gamma_3\gamma_4 + \gamma_4\gamma_5 + \gamma_5\gamma_1 \\ & + \beta_1\beta_3 + \beta_2\beta_4 + \beta_3\beta_5 + \beta_4\beta_1 + \beta_5\beta_2 \end{aligned} \tag{9}$$

under the assumptions that $\gamma_i, \beta_i \geq 0$ are integers and $|F_i| = r_i(v) + \gamma_i + \beta_i$ for all $1 \leq i \leq 5$. (Note that finding a lower bound on (9) gives us a lower bound on the right hand side of (8) and thus a lower bound on the value of t_v.) Notice that there is a choice of the values of the γ_i and β_i which minimise the value of (9) and which satisfy $\gamma_i = 0$ or $\beta_i = 0$ for all $1 \leq i \leq 5$. (For example, if there is a choice of the values of the γ_i and β_i which minimise the value of (9) but with $\gamma_1, \beta_1 > 0$ then this implies that $\gamma_2 + \gamma_5 = \beta_3 + \beta_4$. We can thus obtain another 'minimal' choice of the γ_i and β_i by resetting $\gamma_1 = 0$ and $\beta_1 = |F_1| - r_1(v)$.)

Consider such a choice of the γ_i and β_i. So at least three of the γ_i equal 0 or at least three of the β_i equal 0. Assume that $\beta_1 = \beta_2 = 0$. Thus,

$$0.5r_1(v)^2 + 0.5r_2(v)^2 + \gamma_1\gamma_2 \geq (0.02 - \varepsilon_8)n^2 \tag{10}$$

since $r_1(v) + \gamma_1, r_2(v) + \gamma_2 \geq (1 - \varepsilon_7)n/5$. If $\gamma_3 = \gamma_5 = 0$, then similarly

$$0.5r_3(v)^2 + 0.5r_5(v)^2 + \beta_3\beta_5 \geq (0.02 - \varepsilon_8)n^2.$$

Together with (8) this implies that $t_v \geq (0.04 - 2\varepsilon_8 n^2 - 2\varepsilon_5)n^2$, a contradiction to (7). So $\beta_3 = 0$ or $\beta_5 = 0$. Assume that $\beta_3 = 0$. Thus, as before we have that

$$0.5r_2(v)^2 + 0.5r_3(v)^2 + \gamma_2\gamma_3 \geq (0.02 - \varepsilon_8)n^2. \tag{11}$$

Hence, (10) and (11) imply that (9) is bounded below by

$$(0.04 - 2\varepsilon_8)n^2 - 0.5r_2(v)^2.$$

In all other cases we obtain that (9) is bounded below by

$$(0.04 - 2\varepsilon_8)n^2 - 0.5r_{j'}(v)^2$$

for some $1 \leq j' \leq 5$. In particular, together with (8) this implies that

$$t_v \geq (0.04 - 2\varepsilon_8)n^2 - 0.5r_{j'}(v)^2 - 2\varepsilon_5 n^2$$

for some $1 \leq j' \leq 5$. Thus, (7) implies that $r_{j'}(v) \geq (0.2 - \varepsilon_9)n$ for some $1 \leq j' \leq 5$. This in turn implies that v lies in at least $\binom{(0.2-\varepsilon_9)n}{2} \geq (0.02 - \varepsilon_9)n^2$ red triangles in G. Together with (7), this also implies that $r_i(v) < \varepsilon_{10}n$ for all $i \in [5] \setminus \{j'\}$.

We may assume that $j' = 1$. Suppose that for some j, $g_j(v) \geq 0.0001n$ and $b_j(v) \geq 0.0001n$. Let $\{i_1, i_2, i_3\} = [5] \setminus \{1, j\}$. It is easy to see that this implies that there are at least

$$(0.0001n)^2 - |E^*|$$

green or blue monochromatic triangles containing v and vertices from V_j, V_{i_1}, V_{i_2} and V_{i_3}. Therefore, $t_v \geq (0.02 - \varepsilon_9)n^2 + (0.0001n)^2 - |E^*|$, a contradiction to (7).

Thus, for every $i \in \{2, 3, 4, 5\}$, either $g_i(v) < 0.0001n$ or $b_i(v) < 0.0001n$. If $b_2(v) \geq 0.0001n$ then it is easy to see that $b_4(v), b_5(v) < 0.0001n$ (else we get $(0.0001n)^2 - |E^*|$ blue triangles containing v, a contradiction). So $g_4(v), g_5(v) \geq 0.19n$. This implies that there are at least $(0.19n)^2 - |E^*|$ green triangles containing v, a contradiction. Thus, $b_2(v) < 0.0001n$. Similar arguments imply that $g_3(v), b_4(v), g_5(v) < 0.0001n$. This implies that $v \in V_1$, a contradiction. So indeed $V^* = \emptyset$, as desired. □

By Claims 10 and 12, $V(G)$ can be partitioned into 5 monochromatic cliques of the same colour. A straightforward calculation yields that the graphs in $\mathcal{G}_n$ are precisely those 3-coloured complete graphs on n vertices that minimise the number of monochromatic triangles among all 3-coloured complete graphs whose vertex set can be partitioned into 5 monochromatic cliques of the same colour. Thus, $G \in \mathcal{G}_n$ as desired.

Acknowledgements

We thank the referees for their helpful comments. Thanks also to Tom Bohman, Po-Shen Loh, John Mackey, Ryan Martin, Dhruv Mubayi, and Oleg Pikhurko for their interest and encouragement. The computational part was done using the Maple[4] computer algebra system [15] and the freely available CSDP [1] and SDPA [25] semidefinite programming solvers. In particular the robustness of CSDP and the availability of very high precision versions of SDPA were critical.

[1] B. Borchers. CSDP, A C Library for Semidefinite Programming. *Optimization Methods and Software* 11(1):613-623, 1999.

[2] S.A. Burr and V. Rosta, On the Ramsey multiplicity of graphs - problems and recent results, *J. Graph Theory* **4** (1980), 347–361.

[3] D. Conlon, On the Ramsey multiplicity of complete graphs, *Combinatorica* **32** (2012), 171–186.

[4] J. Cummings and M. Young, Graphs containing triangles are not 3-common, *J. Combin.* **2** (2011), 1–14.

[5] J. Cummings, D. Král', F. Pfender, K. Sperfeld, A. Treglown and M. Young, Monochromatic triangles in three-coloured graphs, arXiv:1206.1987v1.

[6] P. Erdős, On the number of complete subgraphs contained in certain graphs, *Magyar Tud. Akad. Mat. Kutató Int. Közl.* **7** (1962), 459–464.

[7] J. Fox, There exist graphs with super-exponential Ramsey multiplicity constant, *J. Graph Theory* **57** (2008), 89–98.

[8] G. Giraud, Sur les proportions respectives de triangles uni, bi- ou tricoloures dans un tricolouriage des arêtes du n-emble, *Discrete Math.* **16** (1976), 13–38.

[9] A.W. Goodman, On sets of acquaintances and strangers at any party, *Amer. Math. Monthly* **66** (1959), 778–783.

[10] A.W. Goodman, Triangles in a complete chromatic graph with three colours, *Discrete Math.* **57** (1985), 225–235.

[11] R.E. Greenwood and A.M. Gleason, Combinatorial relations and chromatic graphs, *Canadian J. Math.* **7** (1955), 1–7.

[12] H. Hatami, J. Hladký, D. Král', S. Norine and A. Razborov, Non-three-colorable common graphs exist, *Combin. Prob. Comput.* **21** (2012), 734–742.

[13] C. Jagger, P. Šťovíček, and A. Thomason, Multiplicities of subgraphs, *Combinatorica* **16** (1996), 123–141.

[14] S. Janson, T. Łuczak and A. Ruciński, *Random Graphs*, Wiley, 2000.

[15] Maple 14. Maplesoft, a division of Waterloo Maple Incorporated, Waterloo, Ontario.

[4] Maple is a trademark of Waterloo Maple Incorporated.

[16] O. Pikhurko, private communication, 2012.

[17] F.P. Ramsey, On a problem of formal logic, *Proc. London Math. Soc* **30** (1930), 264–286.

[18] A. Razborov, Flag algebras, *J. Symbol. Logic* **72** (2007), 1239–1282.

[19] A. Razborov, On 3-hypergraphs with forbidden 4-vertex configurations, *SIAM J. Discrete Math.* **24** (2010) 946–963.

[20] S.S. Sane and W.D. Wallis, Monochromatic triangles in three colours, *Bull. Austral. Math. Soc.* **37** (1988), 197–212.

[21] Sequence A004102. The On-Line Encyclopedia of Integer Sequences, http://oeis.org/A004102, 2012.

[22] A. Thomason, A disproof of a conjecture of Erdős in Ramsey theory. *J. London Math. Soc.* **39** (1989), 246–255.

[23] A. Thomason, Graph products and monochromatic multiplicities, *Combinatorica* **17** (1997), 125–134.

[24] W.D. Wallis, The number of monochromatic triangles in edge-colourings of a complete graph, *J. Combin. Inform. System. Sci.* **1** (1976), 17–20.

[25] M. Yamashita, K. Fujisawa, M. Fukuda, K. Nakata and M. Nakata. A high-performance software package for semidefinite programs: SDPA 7. Research Report B-463, Department of Mathematical and Computing Science, Tokyo Institute of Technology, Tokyo, Japan, 2010.