ON POSSIBLE UNIFORM TURÁN DENSITIES

DYLAN KING, SIMÓN PIGA, MARCELO SALES, AND BJARNE SCHÜLKE

Abstract. Given a family of 3-graphs $\mathcal{F}$, the *uniform Turán density* $\pi_{\therefore}(\mathcal{F})$ is defined as the infimum $d \in [0,1]$ for which any sufficiently large uniformly d-dense 3-graph — that is, a 3-graph which has edge-density at least d on all linearly sized subsets — contains a copy of some $F \in \mathcal{F}$. Let $\Pi_{\therefore,\mathrm{fin}}$ denote the set of all possible uniform Turán densities of finite families. Erdős, Hajnal, and Rödl introduced a family of constructions for lower bounds on uniform Turán densities called palette constructions. We show that $\Pi_{\therefore,\mathrm{fin}}$ contains every d that is obtained as the uniform density of an optimized palette construction. A corollary of this is that $\Pi_{\therefore,\mathrm{fin}}$ contains the set of Lagrangians of 3-graphs and includes irrational numbers. Our work complements a recent result of Lamaison, which states that every value in $\Pi_{\therefore,\mathrm{fin}}$ can be approximated by uniform densities of palette constructions.

§1 Introduction

For $n \in \mathbb{N}$ and a family $\mathcal{F}$ of k-uniform hypergraphs (or k-graphs), let the *extremal number* $\mathrm{ex}(n,\mathcal{F})$ be the maximum number of edges in a k-graph G on n vertices that does not contain a copy of any $F \in \mathcal{F}$. Such a k-graph G is called $\mathcal{F}$*-free*. It is well known that the quantity $\mathrm{ex}(n,\mathcal{F})/\binom{n}{k}$ is decreasing [19], and therefore one may define the *Turán density* of a family $\mathcal{F}$ as

$$\pi(\mathcal{F}) := \lim_{n\to\infty} \frac{\mathrm{ex}(n,\mathcal{F})}{\binom{n}{k}}.$$

When the family $\mathcal{F} = \{F\}$ is a single k-graph, we usually denote $\pi(\mathcal{F})$ by $\pi(F)$. Let $\Pi_{\infty}^{(k)}$ be the set of all possible Turán densities of families of k-graphs and let $\Pi_{\mathrm{fin}}^{(k)}$ be the set of Turán densities of finite families $\mathcal{F}$.

The study of Turán densities was initiated by Turán [36], who determined $\mathrm{ex}(n,F)$ when F is the complete (2-)graph. Erdős, Stone, and Simonovits [9, 11], generalised this by establishing that

$$\pi(F) \geqslant \frac{\chi(F)-2}{\chi(F)-1}\,,$$

The third author is supported by US Air force grant FA9550-23-1-0298.

The fourth author is supported by the Young Scientist Fellowship IBS-R029-Y7.

The work on this article was supported by the SQuaRE “Variants of the hypergraph Turán problem” at the American Institute of Mathematics.

where $\chi(F)$ is the chromatic number of F. Their proof also gives that

$$\Pi_{\infty}^{(2)} = \Pi_{\text{fin}}^{(2)} = \left\{ \frac{k}{k+1} : k \in \mathbb{N}_{\geqslant 0} \right\}.$$

For higher uniformities, the problem becomes considerably harder and, despite much effort, remains wide open even for 3-graphs. Determining the Turán density even for seemingly "simple" 3-graphs is notoriously difficult. Perhaps the most famous open problem is finding the Turán density of the complete 3-graph on four vertices, $K_4^{(3)}$. In fact, even $\pi(K_4^{(3)-})$, the Turán density of the 3-graph with four vertices and three edges, is unknown.

A different line of research investigates properties of the set of Turán densities. Disproving a 1000$ conjecture by Erdős, Frankl and Rödl [15] showed that for $k \geqslant 3$ the set $\Pi_{\infty}^{(k)}$ is not well-ordered, i.e., there exists some $\alpha \in [0,1)$ such that for every $\varepsilon > 0$, the set $(\alpha, \alpha+\varepsilon) \cap \Pi_{\infty}^{(k)}$ is not empty. This indicates how much more difficult the hypergraph Turán problem is for k-graphs with $k \geqslant 3$. Recently, Pikhurko [30] proved a series of results concerning $\Pi_{\infty}^{(k)}$ and $\Pi_{\text{fin}}^{(k)}$. In particular, he showed that $\Pi_{\infty}^{(k)} \neq \Pi_{\text{fin}}^{(k)}$, $\Pi_{\infty}^{(k)}$ is uncountable, and, using a result by Brown and Simonovits [7], that $\Pi_{\infty}^{(k)} = \overline{\Pi_{\text{fin}}^{(k)}}$. However, the full description of the sets $\Pi_{\infty}^{(k)}$ and $\Pi_{\text{fin}}^{(k)}$ remains open. For more on the hypergraph Turán problem, we refer to the survey by Keevash [20].

Here we consider a variant of the Turán density suggested by Erdős and Sós [10, 12]. Throughout the rest of the paper, we focus on 3-graphs. For $d \in [0,1]$ and $\eta > 0$, we say that a 3-graph H on n vertices is (d,η)-*dense* if for all $X \subseteq V(H)$, we have

$$e(X) \geqslant d\binom{|X|}{3} - \eta n^3 .$$

The *uniform Turán density* $\pi_{\therefore}$ of a family $\mathcal{F}$ of 3-graphs is defined as

$$\begin{aligned}\pi_{\therefore}(\mathcal{F}) = \sup\{d \in [0,1] : &\text{ for every } \eta > 0 \text{ and } n \in \mathbb{N}, \text{ there exists} \\ &\text{an } \mathcal{F}\text{-free, } (d,\eta)\text{-dense 3-graph } H \text{ with } |V(H)| \geqslant n\}.\end{aligned}$$

In other words, $\pi_{\therefore}(\mathcal{F})$ is the smallest $d \in [0,1]$ such that there is some $\eta > 0$ such that every sufficiently large 3-graph H on n vertices that is $(d+o(1),\eta)$-dense contains a copy of some $F \in \mathcal{F}$.

Erdős and Sós specifically asked to determine $\pi_{\therefore}(K_4^{(3)})$ and $\pi_{\therefore}(K_4^{(3)-})$. Similarly as with the original Turán density, these problems turned out to be very difficult. Only recently, Glebov, Král', and Volec [18] and Reiher, Rödl, and Schacht [34] independently solved the latter, showing that $\pi_{\therefore}(K_4^{(3)-}) = 1/4$, which confirmed a conjecture by Erdős and Sós. We refer to Reiher's survey [31] for a full description of the landscape of extremal problems in uniformly dense hypergraphs.

Similarly as for the original Turán density, let $\Pi_{\therefore,\infty}$ be the set of all possible uniform Turán densities of families and $\Pi_{\therefore,\mathrm{fin}}$ the set of all possible uniform Turán densities of finite families. In order to state and discuss the main result of this paper, we need to introduce the concept of Lagrange polynomials, which go back to the work of Motzkin and Straus [27]. Given $n \in \mathbb{N}$ and a subset of ordered triples (which will later be called a "palette") $P \subseteq [n]^3$, we define the *Lagrange polynomial of P* as

$$\lambda_P(x_1, \ldots, x_n) = \sum_{(i,j,k) \in P} x_i x_j x_k \,,$$

and the *Lagrangian of P*, denoted as Λ_P, as the maximum of $\lambda_P(x_1, \ldots, x_n)$ subject to $x_1 + \ldots + x_n = 1$ and $x_i \in [0, 1]$ for all $i \in [n]$. Let Λ_{pal} be the set of all possible Λ_P. In [21], it was shown that $\Lambda_{\mathrm{pal}} \subseteq \Pi_{\therefore,\infty}$ and, as a corollary, that the set $\Pi_{\therefore,\infty}$ is not well-ordered and contains irrational numbers. However, the families constructed in the proof were all infinite, which does not shed light on the possible values of $\Pi_{\therefore,\mathrm{fin}}$. Here we extend the result in [21] by showing that $\Lambda_{\mathrm{pal}} \subseteq \Pi_{\therefore,\mathrm{fin}}$.

Theorem 1.1. *For all $\lambda \in \Lambda_{\mathrm{pal}}$, there is a finite family $\mathcal{F}$ of 3-graphs with $\pi_{\therefore}(\mathcal{F}) = \lambda$.*

Until recently, the only known members of $\Pi_{\therefore,\mathrm{fin}}$ were $0, 1/27, 4/27, 1/4$, and $8/27$ [16–18, 32, 34]. Recently, two infinite families of uniform Turán densities were obtained: one converging to $1/2$ [24] and another being the uniform Turán densities of large stars [26]. All of these densities are rational numbers. A corollary of Theorem 1.1 is that there exist irrational uniform Turán densities of finite families (e.g., see Observation 6.1, [21]).

Interestingly, one of the core steps (Lemma 4.1) in our proof is about structural Ramsey theory. The proof of this key lemma relies on the partite construction method of Nešetřil and Rödl [28], and for this reason the bounds on the graphs in $\mathcal{F}$ are enormous. A generalization of the aforementioned lemma was obtained independently by Král, Kučerák, Lamaison, and Tardos [23].

§2 Palettes

In this section, we present the main technical result of this paper. To do so, we first define the notion of a palette.

Definition 2.1. *A* palette *is a pair $P = (C, E)$ consisting of a set of colors C and a set of patterns $E \subseteq C^3$.*

Although this definition is very similar to that of ordered 3-graphs, note that a palette may contain degenerate patterns, i.e., patterns that contain fewer than three colors. We denote the set of colors of a palette P by $C(P)$, while P should be understood as the set of patterns $E(P)$. Let $c(P)$ and $e(P)$ be the number of colors and the number of patterns of P, respectively, and let $d(P) := e(P)/c(P)^3$ be its density. Note that for $p \in P$, $\{p\}$ can be viewed as a palette itself and as usual we omit the parentheses when writing $C(p)$ etc. We say that P is *non-degenerate* if every pattern of P is non-degenerate, i.e., for every $p \in P$, it holds that $c(p) = 3$. Given a subset $U \subseteq C(P)$ of colors, let $P[U]$ be the induced subpalette on U. That is, the palette with $C(P[U]) = U$ and $P[U] := \{p \in P : C(p) \subseteq U\}$.

Given two palettes P and Q, a *(palette) homomorphism* from P to Q is a map $\psi : C(P) \to C(Q)$ such that for every pattern $p = (c_1, c_2, c_3) \in P$, we have $\psi(p) = (\psi(c_1), \psi(c_2), \psi(c_3)) \in Q$. As with hypergraphs, we usually do not distinguish between isomorphic palettes. If there is an injective homomorphism from P to Q, we also say that P is contained in Q (or that P is a subpalette of Q), denoted by $P \subseteq Q$. We say that a palette Q is a *blow-up* of a palette P if it can be obtained from P by replacing every color with some number of copies of itself. More formally, we say that Q is a blow-up of P with *partition structure* $C(Q) = \bigcup_{c \in C(P)} V_c$ for some pairwise disjoint sets V_c, $c \in C(P)$, if

$$Q = \{(x_1, x_2, x_3) : x_i \in V_{c_i} \text{ for all } i \in [3] \text{ and } (c_1, c_2, c_3) \in P\}\,.$$

Note that P is contained in a blow-up of Q if and only if there is a homomorphism from P to Q. In the case that there is not only a homomorphism from P to Q but an isomorphism, we denote this by $P \cong Q$.

Given a 3-graph $F = (V, E)$ on n vertices, we say that P *paints* F if there exists a total ordering $\dashv$ of V and a coloring $\chi : V^{(2)} \to C(P)$ such that for every edge $xyz \in E$ with $x \dashv y \dashv z$, we have

$$(\chi(xy), \chi(xz), \chi(yz)) \in P. \tag{2.1}$$

Sometimes we refer to such a tuple $(\dashv, \chi)$ as a painting of F (using P). If there is no painting of F using P, we say that P does not paint F, or alternatively, that P is F-deficient. We say that P does not paint a family $\mathcal{F}$, or is $\mathcal{F}$-deficient, if P does not paint F for every $F \in \mathcal{F}$.

Palettes were introduced in [8, 31, 35] in the context of describing a general lower bound construction for the uniform Turán density, called the *Palette construction*. Given a family $\mathcal{F}$ of 3-graphs and a palette $P \subseteq [t]^3$ on t colors such that P is $\mathcal{F}$-deficient, we construct an $\mathcal{F}$-free hypergraph H with vertex set $[n]$ as follows. Let $x_1, \ldots, x_t \in [0, 1]$ with $\sum_{i=1}^t x_i = 1$, and let $\chi : [n]^{(2)} \to [t]$ be an auxiliary coloring defined probabilistically by coloring each pair independently with

$$\mathbb{P}(\chi(ab) = i) = x_i, \quad \forall ab \in [n]^{(2)},\, i \in [t].$$

The edges of the hypergraph H are defined using the auxiliary coloring χ as follows

$$H := \{abc \in [n]^{(3)} : a < b < c \text{ and } (\chi(ab), \chi(ac), \chi(bc)) \in P\}.$$

One can observe that by definition, P paints H and therefore paints any subgraph of H. Hence, since P is $\mathcal{F}$-deficient, H is $\mathcal{F}$-free. Moreover, the probability that a triple $abc \in [n]^{(3)}$ is an edge in H is given by

$$\mathbb{P}(abc \in H) = \sum_{(i,j,k)\in P} x_i x_j x_k.$$

It can be shown, by a standard application of concentration inequalities, that for each η, with high probability the hypergraph H is (d, η)-dense for $d = \sum_{(i,j,k)\in P} x_i x_j x_k$ when n is taken sufficiently large.

The construction above naturally motivates the next definition. Denote the standard $(r-1)$-simplex by

$$\mathbb{S}_r := \{(x_1, \ldots, x_r) \in [0,1]^r : x_1 + \ldots + x_r = 1\}.$$

Definition 2.2. *A* weighting of a palette P *is a vector* $\mathbf{x} = (x_i)_{i\in C(P)} \in \mathbb{S}_{c(P)}$. *Given a palette* P *with a weighting* $\mathbf{x}$, *set*

$$\lambda_P(\mathbf{x}) := \sum_{(i,j,k)\in P} x_i x_j x_k .$$

We define the palette Lagrangian Λ_P *of* P *as*

$$\Lambda_P := \max_{\mathbf{x}\in\mathbb{S}_{c(P)}} \lambda_P(\mathbf{x}).$$

As defined in Section 1, the set of values obtained as the Lagrangian of a palette is denoted by $\Lambda_{\mathrm{pal}} := \{\Lambda_P : P \text{ is a palette}\}$. A consequence of the construction shown above is the following folklore result.

Fact 2.3. *Let* P *be a palette and let* $\mathcal{F}$ *be a family of* 3*-graphs such that* P *is* $\mathcal{F}$*-deficient. Then* $\pi_{\therefore}(\mathcal{F}) \geqslant \Lambda_P$.

A folklore conjecture in the area, stated formally in [31], is that every lower bound for a uniform Turán density should be obtained by a palette construction. In a recent breakthrough, Lamaison [24] proved an approximate version of the conjecture showing that for every family $\mathcal{F}$, the Turán density $\pi_{\therefore}(\mathcal{F})$ can be approximated by a sequence of palette Lagrangians. In this paper, we prove in some sense the converse of the conjecture: Every palette Lagrangian is the uniform Turán density of some finite family $\mathcal{F}$.

The proof of our main result proceeds by transferring the original problem to a Turán-type problem for palettes. For this it is crucial that the property of being $\mathcal{F}$-deficient is invariant under homomorphisms.

Fact 2.4. *Let P and Q be two palettes such that there exists a homomorphism $\psi : Q \to P$, and let $\mathcal{F}$ be a family of 3-graphs. If P is $\mathcal{F}$-deficient, then Q is $\mathcal{F}$-deficient.*

Proof. Suppose, to the contrary, that Q paints some $F \in \mathcal{F}$. Then there exists an ordering $\prec$ of $V(F)$ and a coloring $\chi : V(F)^{(2)} \to C(Q)$ satisfying (2.1). Hence, by definition, the ordering $\prec$ and the coloring $\psi \circ \chi : V(F)^{(2)} \to C(P)$ witness that P paints F, contradicting the assumption that P is $\mathcal{F}$-deficient. □

The first consequence of Fact 2.4 is that the property of being $\mathcal{F}$-deficient is closed under taking subpalettes. In particular, this allows us to define the palette Turán density of families of 3-graphs. Given a family $\mathcal{F}$ of 3-graphs, we define the *palette extremal number of $\mathcal{F}$* by

$$\mathrm{ex}_{\mathrm{pal}}(n, \mathcal{F}) := \max\{e(P) : P \text{ is an } \mathcal{F}\text{-deficient palette with } c(P) = n\},$$

i.e., the maximum number of patterns an $\mathcal{F}$-deficient palette with n colors can have. Similarly as for hypergraphs, one can show that the quantity $\mathrm{ex}_{\mathrm{pal}}(n, \mathcal{F})/n^3$ converges to a limit.

Proposition 2.5. *The limit $\lim\limits_{n\to\infty} \frac{ex_{pal}(n,\mathcal{F})}{n^3}$ exists.*

Proof. Recall that a palette P is non-degenerate if every pattern of P contains 3 distinct colors. For a family $\mathcal{F}$ of 3-graphs and an integer $n \geqslant 3$, we define the parameter $g(n, \mathcal{F})$ by

$$g(n, \mathcal{F}) := \max\{e(P) : P \text{ is a non-degenerate } \mathcal{F}\text{-deficient palette with } c(P) = n\}.$$

We claim that the quantity $\frac{g(n,\mathcal{F})}{n(n-1)(n-2)}$ is non-increasing. Indeed, let P be a non-degenerate $\mathcal{F}$-deficient palette on $n + 1$ colors that has $g(n + 1, \mathcal{F})$ patterns. Take a random subset $U \subseteq C(P)$ of size n and let $P[U]$ be the induced subpalette on this set of colors. Then, by Fact 2.4, we have that $P[U]$ is $\mathcal{F}$-deficient and consequently $e(P[U]) \leqslant g(n, \mathcal{F})$. Hence,

$$g(n, \mathcal{F}) \geqslant \mathbb{E}(e(P[U])) = \frac{n-2}{n+1} g(n + 1, \mathcal{F}),$$

which implies that $\frac{g(n+1,\mathcal{F})}{(n+1)n(n-1)} \leqslant \frac{g(n,\mathcal{F})}{n(n-1)(n-2)}$. Since every non-negative non-increasing sequence has a limit, we obtain that $\lim\limits_{n\to\infty} \frac{g(n,\mathcal{F})}{n(n-1)(n-2)}$ exists. The proposition now follows by the simple observation that $g(n, \mathcal{F}) \leqslant \mathrm{ex}_{\mathrm{pal}}(n, \mathcal{F}) \leqslant g(n, \mathcal{F}) + 3n^2$. □

We define the limit obtained in Proposition 2.5 as the *palette Turán density of a family $\mathcal{F}$*,

$$\pi_{\mathrm{pal}}(\mathcal{F}) := \lim_{n\to\infty} \frac{\mathrm{ex}_{\mathrm{pal}}(n, \mathcal{F})}{n^3}.$$

A consequence of the work in [24, 31] is that $\pi_{\mathrm{pal}}(\mathcal{F}) = \pi_{\therefore}(\mathcal{F})$ for finite $\mathcal{F}$ (see Section 7 for more details). Therefore, to show Theorem 1.1, it suffices to show that every palette Lagrangian is attained as a palette Turán density of a finite family.

By taking the uniform weighting $\mathbf{x} = (x_i)_{i\in C(P)}$ defined by $x_i = 1/c(P)$, we obtain that

$$d(P) \leqslant \Lambda_P\,. \tag{2.2}$$

Moreover, if $\mathbf{y} = (y_i)_{i\in C(P)}$ is an optimal weighting for the palette P, then one can approximate Λ_P by taking a sequence of blow-ups $\{S^{(\ell)}\}_{\ell\in\mathbb{N}}$ of P with partition structure $C(S^{(\ell)}) = \bigcup_{i\in C(P)} V_i^{(\ell)}$ such that $|V_i^{(\ell)}|/c(S^{(\ell)}) \to y_i$. Conversely, every blow-up of P yields a weighting for P. Hence, it follows that

$$\Lambda_P = \lim_{n\to\infty} \max\{d(S) :\ S \text{ is a blow-up of } P \text{ with } c(S) = n\}. \tag{2.3}$$

The equality in (2.3) and the fact that the property of being $\mathcal{F}$-deficient is blow-up invariant (Fact 2.4) hint at a way to obtain a palette Lagrangian as a palette Turán density: finding a finite family $\mathcal{F}$ such that the extremal constructions for $\mathrm{ex}_{\mathrm{pal}}(n, \mathcal{F})$ are exactly the family of blow-ups of P. Unfortunately, such a family does not always exist (see Section 4), but by adding an extra family of extremal constructions with the same palette Lagrangian, one can achieve such a goal. Given a palette P, we define the *reverse palette* $\mathrm{rev}(P)$ of P as

$$\mathrm{rev}(P) = \{(c, b, a) :\ (a, b, c) \in P\}.$$

That is, $\mathrm{rev}(P)$ is the palette obtained by reversing the order of the patterns of P. We say that a palette P is *reduced* if for every proper subpalette $Q \subsetneq P$, we have $\Lambda_Q < \Lambda_P$. We follow the graph convention and let $\mathrm{EX}_{\mathrm{pal}}(n, \mathcal{H}) = \{Q :\ e(Q) = \mathrm{ex}_{\mathrm{pal}}(n, \mathcal{H}) \text{ and } c(Q) = n\}$ be the set of extremal palettes. The following is the main technical result of this paper.

Theorem 2.6. *Let P be a reduced palette. There exists a finite family $\mathcal{H}$ such that P is $\mathcal{H}$-deficient and for all $n \in \mathbb{N}$*

$$EX_{pal}(n, \mathcal{H}) \subseteq \{Q :\ Q \text{ is a blow-up of } P \text{ or a blow-up of } \mathrm{rev}(P) \text{ and } c(Q) = n\}\,. \tag{2.4}$$

In particular, it immediately follows that

$$ex_{pal}(n, \mathcal{H}) = \max\{e(Q) :\ Q \text{ is a blow-up of } P \text{ with } c(Q) = n\}\,.$$

We remark that P is not necessarily isomorphic to $\mathrm{rev}(P)$. As a simple example, consider $P = \{(1, 2, 3), (1, 3, 2)\}$ and $\mathrm{rev}(P) = \{(2, 3, 1), (3, 2, 1)\}$. One can verify that in this case, $P \ncong \mathrm{rev}(P)$.

Organization. The paper is organized as follows. The proof of Theorem 2.6 relies on three main ingredients. The first is a palette variant of the removal lemma introduced in [2], so in Section 3 we use a regularity lemma for palettes to prove counting and removal lemmata for palettes painting graphs. The second is a structural Ramsey result in Section 4, dedicated to the problem of distinguishing palettes based on the graphs they can paint. The third component is a stability argument based on the work of [30] (Sections 5 and 6). For a brief outline of the proof of Theorem 2.6, the reader may refer to the introduction of Section 6. Finally, in Section 7, we present a proof of Theorem 1.1.

§3 Regularity lemma for palettes

For graphs, the following infinite removal lemma was shown in [3] (and a hypergraph analogue in [5]).

Lemma 3.1. *Given a (possibly infinite) family of graphs $\mathcal{F}$ and $\alpha > 0$, there are $M, n_0 \in \mathbb{N}$ and $\beta > 0$ so that the following holds for every graph G on $n \geqslant n_0$ vertices. If, for every $F \in \mathcal{F}$ with $v(F) \leqslant M$, G contains fewer than $\beta n^{v(F)}$ copies of F, then G can be made $\mathcal{F}$-free by removing at most αn^2 edges.*

The aim of this section is prove a version of this lemma for palettes. Instead of counting the number of copies of some $F \in \mathcal{F}$, we need to count the number of ways that a palette P paints F. This is made precise in the following definition.

Definition 3.2. *Let F be a 3-graph and P a palette. The number of ways that P paints F is defined as the number of maps $\varphi \colon \partial F \to C(P)$ for which there exists a total ordering $<$ of $V(F)$ such that $(<, \varphi)$ is a painting of F.*

We are now prepared to state the aforementioned removal lemma.

Lemma 3.3 (Palette Removal Lemma)**.** *Given a (possibly infinite) family of 3-graphs $\mathcal{F}$ and $\alpha > 0$, there are $M = M_{3.3}, N = N_{3.3} \in \mathbb{N}$ and $\beta = \beta_{3.3} > 0$ such that the following holds for every palette P on $n \geqslant N$ colors. If, for every $s \in [\binom{M}{2}]$, P paints the 3-graphs $F \in \mathcal{F}$ with $|\partial F| = s$ and $v(F) \leqslant M$ in less than βn^s ways, then there is an $\mathcal{F}$-deficient palette $Q \subseteq P$ with $|P \smallsetminus Q| \leqslant \alpha n^3$.*

The proof of Lemma 3.3 is given at the conclusion of this Section. Similar to the proof of Lemma 3.1 for graphs, it will require a regularity theory for palettes. In some ways it is helpful to consider a palette merely as essentially an oriented 3-graph, since the number of degenerate patterns is $O(n^2)$. Let us define what it means for sets of colors in a palette to be ε-regular. Let P be a palette and suppose that $W_1, W_2, W_3 \subseteq C(P)$ are non-empty. We set $E(W_1, W_2, W_3) = (W_1 \times W_2 \times W_3) \cap P$ and $e(W_1, W_2, W_3) = |E(W_1, W_2, W_3)|$[1]. Then the density of P induced on (W_1, W_2, W_3) is given by

$$d(W_1, W_2, W_3) := \frac{e(W_1, W_2, W_3)}{|W_1||W_2||W_3|}\,.$$

Definition 3.4. *We call (V_1, V_2, V_3) ε-regular if, for all $W_1 \subseteq V_1, W_2 \subseteq V_2, W_3 \subseteq V_3$ with $|W_i| \geqslant \varepsilon|V_i|$, $i \in [3]$, we have*

$$|d(W_1, W_2, W_3) - d(V_1, V_2, V_3)| \leqslant \varepsilon\,.$$

As discussed above, the most important feature of the above definition is that it is sensitive to order, and the regularity of (V_1, V_2, V_3) has no bearing on the regularity of (V_2, V_1, V_3). However, we can derive many properties of these regular color sets by applying corresponding results in the unoriented hypergraph setting (see Lemma 3.10 below), so we give this definition as well. Let $H = (V, E)$ be a 3-graph and suppose that $X_1, X_2, X_3 \subseteq V$ are non-empty. We set

$$E(X_1, X_2, X_3) = \{xyz \in E : \; x \in X_1, y \in X_2, z \in X_3\}$$

and $e(X_1, X_2, X_3) = |E(X_1, X_2, X_3)|$. Then the density of H induced on X_1, X_2, X_3 is given by

$$d(X_1, X_2, X_3) := \frac{e(X_1, X_2, X_3)}{|X_1||X_2||X_3|}\,.$$

Definition 3.5. *Suppose $H = (V, E)$ is a 3-graph and that $X_1, X_2, X_3 \subseteq V(H)$ are non-empty. We call X_1, X_2, X_3 ε-regular if, for all $Y_1 \subseteq X_1, Y_2 \subseteq X_2, Y_3 \subseteq X_3$ with $|Y_i| \geqslant \varepsilon|X_i|$, $i \in [3]$, we have*

$$|d(Y_1, Y_2, Y_3) - d(X_1, X_2, X_3)| \leqslant \varepsilon\,.$$

As usual we will be interested in partitioning $C(P)$ into a large (but bounded) number of parts so that most (V_i, V_j, V_k) are regular, so we also need the standard notions of equipartitions and refinements.

Definition 3.6. *Given a set C, an equipartition $\mathcal{A}$ of C is a partition $C = \dot\bigcup_{i\in[t]} V_i$, so that $||V_i| - |V_{i'}|| \leqslant 1$ for all $i, i' \in [t]$. A refinement of $\mathcal{A}$ is an equipartition $\mathcal{B} = \dot\bigcup_{i\in[t]} \dot\bigcup_{j\in[\ell]} V_{i,j}$ with $V_{i,j} \subseteq V_i$ for all $i \in [t]$ and $j \in [\ell]$. We also identify $\mathcal{A}$ with the family of partition classes, i.e., $\mathcal{A} = \{V_i : \; i \in [t]\}$.*

Our palette analogue of the well-known Szemerédi regularity lemma is the following.

Theorem 3.7. *For all $\varepsilon > 0$ and $m \in \mathbb{N}$ there exist $M = M_{3.7}, N = N_{3.7} \in \mathbb{N}$ so that given any palette Q with $c(Q) \geqslant N$ there is an equipartition $\mathcal{A} = \{V_i : i \in [t]\}$ of $C(Q)$ so that*

[1]For ease of notation we suppress the dependency on P when it is clear from the context.

(1) $m \leqslant t \leqslant M$ *and*

(2) the ordered triple (V_i, V_j, V_k) *is* ε*-regular for all but* εt^3 *of* $(i,j,k) \in [t]^3$.

Moreover, given some equipartition $\mathcal{A}_0$ *of* $C(Q)$ *with at most* m *parts, there is an* $\mathcal{A}$ *as above which refines* $\mathcal{A}_0$.

In our application we need a strengthening of this; in particular, the ε-regularity obtained should be allowed to depend on the number of parts t in the partition, as follows.

Corollary 3.8. *For all non-increasing maps* $\mathcal{E} : \mathbb{N} \to (0,1]$ *and* $m \in \mathbb{N}$ *there are* $M = M_{3.8}$, $N = N_{3.8}$ *and* $\delta = \delta_{3.8} > 0$ *so that given any palette* Q *with* $c(Q) = n \geqslant N$ *there is are an equipartition* $\mathcal{A} = \{V_i : i \in [t]\}$ *of* $C(Q)$ *and an equipartition* $\mathcal{A}' = \{U_i : i \in [t]\}$ *of some subset of* $C(Q)$ *so that:*

(i) $m \leqslant t \leqslant M$,

(ii) $U_i \subseteq V_i$ *with* $|U_i| \geqslant \delta n$,

(iii) the ordered triple (U_i, U_j, U_k) *is* $\mathcal{E}(t)$*-regular for all* $(i,j,k) \in [t]^3$, *and*

(iv) we have $|d(U_i,U_j,U_k) - d(V_i,V_j,V_k)| < \mathcal{E}(0)$ *for all but* $\mathcal{E}(0)t^3$ *of* $(i,j,k) \in [t]^3$.

The proof of Theorem 3.7 follows along the standard technique of iterated refinement used for graph and weak hypergraph regularity, and the method to obtain Corollary 3.8 from Theorem 3.7 was developed in [2] for graph-testing problems. For the interested reader we include the proofs of Theorem 3.7 and Corollary 3.8 in Appendix A.

Before we prove Lemma 3.3, we show a counting lemma for the number of ways in which a palette paints a 3-graph. We use the following result, which counts the number of copies of a linear hypergraph inside a regularly partitioned hypergraph. We state only the 3-uniform case.

Lemma 3.9 ([22, Lemma 10])**.** *Given* $\gamma, d_0 > 0$ *and* $\ell \in \mathbb{N}$ *there are* $\varepsilon = \varepsilon_{3.9}(\gamma, d_0, \ell)$ *and* $N = N_{3.9} = (\gamma, d_0, \ell)$ *such that the following holds. Let* F *be a linear* 3*-graph with* $V(F) = [\ell]$ *and* H *an* ℓ*-partite* 3*-graph on parts* $V_1, \dots, V_\ell$ *with* $|V_i| \geqslant N$ *for each* $i \in [\ell]$. *Suppose that* $\{V_i\}_{i \in f}$ *is* ε*-regular with density* $d_f \geqslant d_0$ *for every* $f \in E(F)$. *Then the number of copies of* F *in* H *that map each vertex* $i \in [\ell]$ *to* V_i *is at least* $(1-\gamma)d_0^{e(F)} \prod_{i \in [\ell]} |V_i|$.

By creating a hypergraph which captures the structure of P, we can obtain a similar statement estimating the number of ways in which P paints a given 3-graph F.

Lemma 3.10. *Given* $\gamma, d_0 > 0$ *and* $s \in \mathbb{N}$ *there are* $\varepsilon = \varepsilon_{3.10}(\gamma, d_0, s)$ *and* $N = N_{3.10} = (\gamma, d_0, s)$ *such that the following holds. Let* F *be a* 3*-graph with* $|\partial F| = s$ *and* P *be a palette whose colours are partitioned into* $C(P) = V_1 \cup \ldots \cup V_s$ *where* $|V_i| \geqslant N$ *for every* $i \in [s]$. *Suppose there is an ordering* $<$ *of the vertices of* F *and a map* $\varphi \colon \partial F \to [s]$ *such that for* *each $uvw \in E(F)$ with $u < v < w$, the triple $(V_{\varphi(uv)}, V_{\varphi(uw)}, V_{\varphi(vw)})$ is ε-regular in P with density at least d_0. Then F is painted by P in at least $(1-\gamma)d_0^{e(F)} \prod_{uv \in \partial F} |V_{\varphi(uv)}|$ different ways.*

Proof. We write uvw for $\{u, v, w\}$ whenever $u < v < w$ for $u, v, w \in V(F)$. Let F^∂ be a 3-graph with vertex-set ∂F and all edges of the form $\{uv, uw, vw\}$ where $uvw \in E(F)$. Note that F^∂ is linear and $e(F) = e(F^\partial)$. For each $uv \in \partial F$ let V_i^{uv} be a copy of V_i, and label the copy of $a \in V_i$ as $a^{uv} \in V_i^{uv}$. Let H be a 3-graph with vertices $V(H) = \dot{\bigcup}_{\substack{uv \in \partial F \\ i \in [k]}} V_i^{uv}$ and edges

$$E(H) = \{a^{uv} b^{uw} c^{vw} \in V(H)^{(3)} \colon uv, uw, vw \in \partial F, \text{ and } (a, b, c) \in P\}\,.$$

In words, H is obtained by taking $s = |\partial F|$ copies of $C(P)$, indexed by $uv \in \partial F$, and then adding a 3-edge xyz only when x, y, and z are part of distinct copies and, when viewed as an ordered triple in $C(P)$, they give a pattern in P. Now if (V_i, V_j, V_k) is ε-regular in the palette P (Definition 3.4), then the vertex sets $(V_i^{uv}, V_j^{uw}, V_k^{vw})$ are ε-regular in H for each $uv, uw, vw \in \partial F$ (Definition 3.5). Consider in particular the s sets given by $V_{\varphi(uv)}^{uv}$ for $uv \in \partial F$. Since F^∂ is linear, Lemma 3.9 yields at least $(1-\gamma)d_0^{e(F^\partial)} \prod_{uv \in \partial F} |V_{\varphi(uv)}^{uv}|$ copies (or embeddings) of F^∂ in H that map each $uv \in V(F^\partial)$ to some vertex in $V_{\varphi(uv)}^{uv}$. Write Ψ for the collection of such embeddings. Let $\chi : V(H) \to C(P)$ be the projection map which sends a^{uv} to $\chi(a^{uv}) = a$. Note that for every $\psi \in \Psi$, the tuple $(<, \chi \circ \psi)$ is a painting of F using P and that any two distinct embeddings $\psi, \psi' \in \Psi$ get projected to distinct paintings, whence the number of such paintings is at least $(1-\gamma)d_0^{e(F)} \prod_{uv \in \partial F} |V_{\varphi(uv)}|$. $\square$

We are now ready to prove the removal lemma using the counting lemma.

Proof of Lemma 3.3. Let $\mathcal{P}_t(\mathcal{F})$ be the set of all palettes with color set $[t]$ which paint at least one $F \in \mathcal{F}$.

We define the map $\mathfrak{v}_{\mathcal{F}} \colon \mathbb{N} \to \mathbb{N}_{\geqslant 0}$ by

$$\mathfrak{v}_{\mathcal{F}}(t) = \max_{R \in \mathcal{P}_t(\mathcal{F})} \min\{|V(F)| \colon F \in \mathcal{F} \text{ and } R \text{ paints } F\}$$

(and $\mathfrak{v}_{\mathcal{F}}(t) = 0$ if $\mathcal{P}_t(\mathcal{F}) = \varnothing$). Since $\mathcal{P}_t(\mathcal{F})$ is finite for every t, the maximum exists. The idea is the following. Given a palette P that paints every small $F \in \mathcal{F}$ in few ways, we apply the palette regularity lemma. From this we obtain a "reduced" palette R with a constant number of colors t and a "cleaned" palette Q similar to P. If Q would still paint some $F \in \mathcal{F}$, then R would paint F and hence - using the definition of $\mathfrak{v}_{\mathcal{F}}$ - some $F' \in \mathcal{F}$ with few vertices. Then the counting lemma entails that P must paint F' as well. Let us formalize this argument.

Given $\alpha > 0$, define the function $\mathcal{E}\colon \mathbb{N}_0 \to [0,1]$ by

$$\mathcal{E}(t) = \begin{cases} 2\alpha/9 & \text{if } t = 0 \text{ and} \\ \varepsilon_{3.10}(1/2, 2\alpha/9, \binom{\mathfrak{v}_{\mathcal{F}}(t)}{2}) & \text{if } t \geqslant 1\,. \end{cases}$$

Note that we may assume that the functions $N_{3.10}(\gamma, d_0, s)$ and $\varepsilon_{3.10}(\gamma, d_0, s)$ are non-decreasing and non-increasing in s, respectively. Let $M_{3.8}, N_{3.8}, \delta = \delta_{3.8}$ be the constants given by an application of Corollary 3.8 with $\mathcal{E}$ and

$$m > 9/\alpha\,. \tag{3.1}$$

Finally, take

$$M = \max_{r \in [M_{3.8}]} \mathfrak{v}_{\mathcal{F}}(r)\,, \quad N = \max\Big\{N_{3.8}, \frac{1}{\delta} N_{3.10}\big(1/2, 2\alpha/9, \binom{M}{2}\big)\Big\}\,, \quad \beta \leqslant \frac{\left(\frac{2\alpha}{9}\right)^{\binom{M}{3}} \delta^{\binom{M}{2}}}{2}\,.$$

Now we argue that these choices of M, N, and β have the desired property. Given a palette P with $n \geqslant N$ colors, apply (the conclusion of) Corollary 3.8 to obtain equipartitions $\mathcal{A} = \{V_i : i \in [t]\}$ and $\mathcal{A}' = \{U_i : i \in [t]\}$ satisfying *(i)*-*(iv)*. Note that in particular, $|U_i| \geqslant \delta N \geqslant N_{3.10}\big(1/2, 2\alpha/9, \binom{M}{2}\big)$. We produce a 'reduced' palette R with $C(R) = [t]$ by including the pattern (i, j, k) in R if

(1) i, j, k are pairwise distinct,
(2) $|d(U_i, U_j, U_k) - d(V_i, V_j, V_k)| \leqslant \mathcal{E}(0) = 2\alpha/9$, and
(3) $d(U_i, U_j, U_k) > \mathcal{E}(0) = 2\alpha/9$.

Let $Q \subseteq P$ be the 'cleaned' version of P, where we delete all edges $(a, b, c) \in P$ with $a \in V_i, b \in V_j, c \in V_k$ when $(i, j, k) \notin R$. In this way, it is easy to see that Q is contained in a blow-up of R.

As in many other applications of the regularity lemma, it is not hard to check that

$$|P \smallsetminus Q| \leqslant \alpha n^3\,. \tag{3.2}$$

Indeed, there are at most $3t^2$ triplets of indices not satisfying *(1)*. Thus, due to (3.1), at most $3t^2(n/t)^3 \leqslant 3n^3/t \leqslant \alpha n^3/3$ patterns are deleted in this way. By Corollary 3.8 Part *(iv)* there are at most $\mathcal{E}(0)t^3 = 2\alpha t^3/9$ triplets (i, j, k) such that $|d(U_i, U_j, U_k) - d(V_i, V_j, V_k)| > \mathcal{E}(0)$, meaning that at most $n^3/t^3 \cdot 2\alpha t^3/9 \leqslant 2\alpha n^3/9$ patterns need to be deleted to ensure *(2)*. Finally, using that for the remaining triplets of indices *(2)* holds, in *(3)* we delete at most $4\alpha n^3/9$ edges.

Suppose for a contradiction that Q paints a hypergraph $F \in \mathcal{F}$. Since Q is contained in a blow-up of R, it follows that R paints F as well. Therefore, keeping in mind that $c(R) = t \leqslant M_{3.8}$, there is some hypergraph $F' \in \mathcal{F}$ painted by R with the additional property that $v(F') \leqslant \mathfrak{v}_{\mathcal{F}}(t) \leqslant M$. Let $s = |\partial F'| \leqslant \binom{M}{2}$ and let $\varphi : \partial F' \to C(R) = [t]$ be the coloring given by definition of painting (we leave the vertex ordering to be implicit). Due to *(1)*-*(3)* and *(iii)* in Corollary 3.8, the map φ satisfies the conditions of Lemma 3.10, when restricted to the subpalette $U_1 \cup \ldots \cup U_s$, which implies that Q paints F' in more than βn^s ways, a contradiction. □

§4 A Ramsey result

Given two distinct 3-graphs G and H such that neither is a subgraph of the other, there always exists a 3-graph F with the property that G contains a copy of F, but H is F-free. Indeed, one can take the 3-graph F to be G itself. It is somewhat natural to ask if the same is true for palettes. That is, for what pairs of palettes P and Q, does there exist a 3-graph F such that P paints F, but Q does not paint F? The goal of this section is to answer this question. We remark that similar considerations were mentioned in [24].

Given a palette P on n colors, recall that the *reverse palette* $\mathrm{rev}(P)$ of P as

$$\mathrm{rev}(P) = \{(c, b, a) : (a, b, c) \in P\},$$

that is, $\mathrm{rev}(P)$ is the palette obtained by reversing the order of the patterns of P. Note that a palette P paints a 3-graph F if and only if $\mathrm{rev}(P)$ paints F. Indeed, this can be seen by taking the ordering of the vertices in which P paints F and reversing it. A consequence of this observation is that no graph can distinguish a palette P from $\mathrm{rev}(P)$. It turns out that up to taking blow-ups, $\mathrm{rev}(P)$ is the only palette for which there is no 3-graph that distinguishes it from P. We remind the reader that a palette Q is contained in a blow-up of P if there exists a homomorphism $\psi : Q \to P$. The next lemma is the main result in this section.

Lemma 4.1. *Let P and Q be palettes such that Q is not contained in a blow-up of P nor in a blow-up of* $\mathrm{rev}(P)$*. Then there exists a* 3*-graph F such that P is F-deficient and Q paints F.*

The proof of Lemma 4.1 is completely Ramsey-theoretical and relies on a result of Nešetřil and Rödl [28] about Ramsey classes for ordered Steiner systems. We start with some preparations. An ordered k-graph $(H, <)$ is a pair where H is a k-graph and $<$ is a total ordering of $V(H)$. Given two ordered hypergraphs $(F, <)$ and $(H, <)$, we say that $(F, <)$ is a subgraph of $(H, <)$ if there exists an injective order-preserving map $\psi : V(F) \to V(H)$ that is a homomorphism, i.e., a map such that $\psi(x) < \psi(y)$ for $x < y$ and such that $\psi(f) \in E(H)$ for every edge $f \in E(F)$. Let $\binom{(H,<)}{(F,<)}$ denote the family of copies of $(F, <)$ in $(H, <)$. The next theorem shows that the class of ordered linear k-graphs is edge-Ramsey (see also [6, Lemma 2.12] and [33, Corollary 3.12]).

Theorem 4.2 ([28]). *Let $(G,<)$ be an ordered linear k-graph with $k \geqslant 2$, and let $r \geqslant 1$ be an integer. Then there exists an ordered linear k-graph $(H,<)$ and a family $\mathcal{G} \subseteq \binom{(H,<)}{(G,<)}$ of copies of $(G,<)$ in $(H,<)$ satisfying the following statements:*

(i) *For any r-coloring of the edges of H, there exists a monochromatic copy $(G',<) \in \mathcal{G}$.*
(ii) *For any two distinct copies $(G',<),(G'',<) \in \mathcal{G}$, it holds that either $|V(G') \cap V(G'')| \leqslant 1$ or $V(G') \cap V(G'') = e$ for an edge $e \in G_3$ for some $(G_3,<) \in \mathcal{G}$.*

We remark that for $k = 2$, a linear k-graph is just a graph, and condition (ii) translates to the fact that two distinct copies $(G',<)$ and $(G'',<) \in \mathcal{G}$ intersect either in an emptys set, a single vertex or in exactly one edge. Theorem 4.2 can be used to prove the following Ramsey result about systems of graphs. We note that similar statements were obtained previously in [1, 29].

Proposition 4.3. *Let $n, t \geqslant 1$ be integers, and let $(G,<)$ be an ordered graph where $G = \bigcup_{i=1}^{n} G_i$ is the union of n pairwise edge-disjoint ordered graphs with vertex set $V(G)$. Then there exists an ordered graph $(H,<)$ where $H = \bigcup_{i=1}^{n} H_i$ is the union of n pairwise edge-disjoint ordered graphs with vertex set $V(H)$ such that any t-coloring of H yields a set $X \subseteq V(H)$ with the following properties:*

(i) *For $1 \leqslant i \leqslant n$, we have $(H_i[X],<) \cong (G_i,<)$.*
(ii) *For $1 \leqslant i \leqslant n$, the graph $H_i[X]$ is monochromatic.*

Proof. For the sake of brevity, throughout the proof we will omit the total ordering $<$ from the notation and denote an ordered graph $(H,<)$ by H. We inductively construct ordered graphs $A^0,\dots,A^n$ such that for each $0 \leqslant j \leqslant n$, the ordered graph $A^j = \bigcup_{i=1}^{n} A_i^j$ is the union of n pairwise edge-disjoint ordered graphs on $V(A^j)$ as follows: Let $A^0 = G$ and $A_i^0 = G_i$ for $1 \leqslant i \leqslant n$. Suppose now that for $1 \leqslant j \leqslant n$ we have already defined the ordered graph A^{j-1} and want to define A^j. Apply Theorem 4.2 to the ordered graph A_j^{j-1} and t colors to obtain the ordered graph A_j^j and a system of copies $\mathcal{A}_j \subseteq \binom{A_j^j}{A_j^{j-1}}$ satisfying properties (i) and (ii) of the statement. In particular, for any two copies $B, B' \in \mathcal{A}_j$ of A_j^{j-1} we have that

$$|V(B) \cap V(B')| = 1 \quad \text{or} \quad V(B) \cap V(B') \subseteq e \tag{4.1}$$

for some edge e of some copy of A_j^{j-1} in $\mathcal{A}_j$. For $i \neq j$, let A_i^j be the ordered graph on $V(A_j^j)$ (with the same total ordering $<$ on $V(A_j^j)$) obtained by adding a copy of the ordered graph A_i^{j-1} to each set of vertices $V(B)$ with $B \in \mathcal{A}_j$, see Figure 4.1. By (4.1), each pair of copies $B, B' \in \mathcal{A}_j$ either intersects in a single vertex or in an edge of some copy of A_j^{j-1}. Hence, keeping in mind that the graphs A_i^{j-1} are pairwise edge-disjoint, all A_i^j

are pairwise edge-disjoint. We set $A^j = \bigcup_{i=1}^j A_i^j$. The key point of the construction is that the ordered graphs $A^0, \ldots, A^n$ satisfy the following claim.

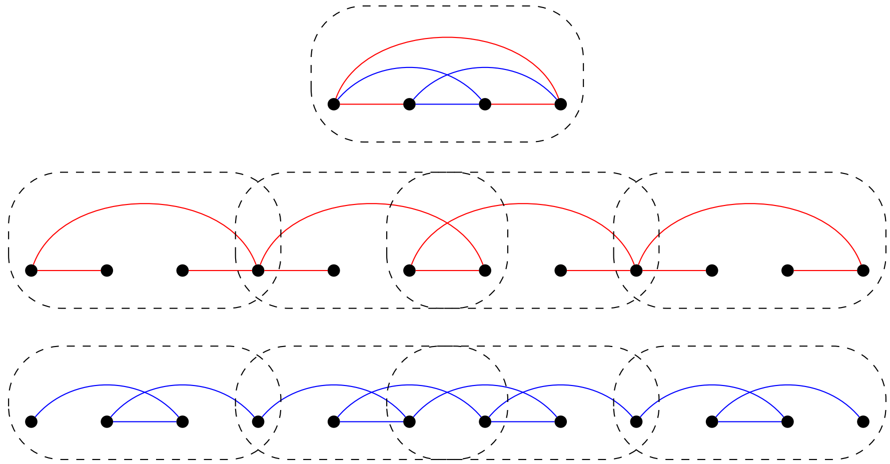

FIGURE 4.1. An example with $n = 2$. The ordered graph A_1^0 is in red and A_2^0 is in blue. On the second line we have the Ramsey graph A_1^1 with several copies of A_1^0 all intersecting in either an edge or a single vertex. Finally, in the last line we have the graph A_2^1 on the same set of vertices.

Claim 4.4. *Let $0 \leqslant j \leqslant n$. Then every t-coloring of A^n contains a copy of A^j such that the ordered graphs A_k^j are monochromatic for $j+1 \leqslant k \leqslant n$.*

Proof. We prove the statement by reverse induction on j. If $j = n$, then the statement is vacuously true. Now assume that for a t-coloring of A^n we obtain a copy $\tilde{A}^{j+1} = \bigcup_{i=1}^n \tilde{A}_i^{j+1}$ of A^{j+1} such that the ordered graphs $\tilde{A}_k^{j+1}$ are monochromatic for $j+2 \leqslant k \leqslant n$. Consider the restriction of the t-coloring to the ordered graph $\tilde{A}_{j+1}^{j+1}$. By construction and Theorem 4.2, there exists a monochromatic copy $\tilde{A}_{j+1}^j$ of A_{j+1}^j. For $i \neq j$, let $\tilde{A}_i^j = \tilde{A}_i^{j+1}[V(\tilde{A}_{j+1}^j)]$. It is easy to see that $\tilde{A}^j = \bigcup_{i=1}^n \tilde{A}_i^j$ is a copy of A^j. Moreover, since $\tilde{A}_k^j \subseteq \tilde{A}_k^{j+1}$ for $j+2 \leqslant k \leqslant n$, we have that $\tilde{A}_k^j$ is monochromatic for $j+1 \leqslant k \leqslant n$. This concludes the proof of the claim. □

Let $H := A^n$. Then by Claim 4.4, every t-coloring of H contains a copy of A^0 such that every A_i^0 is monochromatic. Since $A^0 = G$, properties (i) and (ii) follow. □

The second auxiliary result establishes the existence of an ordered linear k-graph with the property that, regardless of how one orders its vertices, there will always be an edge which according to the new ordering is arranged in either a strictly increasing or decreasing order (with respect to the original order). We remark that the problem becomes somewhat simpler if we drop the condition that the k-graph is linear. Indeed, in this case, one can construct a graph by simply taking the complete k-graph on $(k-1)^2+1$ vertices and applying the Erdős–Szekeres theorem [13].

Proposition 4.5. *For every integer $k \geqslant 2$, there exists an ordered linear k-graph $(H,<)$ with the following property. For any total ordering $\prec$ of $V(H)$, there exists an edge $\{x_1,\dots,x_k\}$ with $x_1 < \dots < x_k$ such that either*

(a) the edge $\{x_1,\dots,x_k\}$ is increasing in $\prec$, i.e., $x_1 \prec \dots \prec x_k$; or
(b) the edge $\{x_1,\dots,x_k\}$ is decreasing in $\prec$, i.e., $x_1 \succ \dots \succ x_k$.

Proof. Let S_k be the group of all permutations on k elements, and let $\mathrm{id}, \mathrm{rev} \in S_k$ be the permutations given by $\mathrm{id}(i) = i$ and $\mathrm{rev}(i) = k+1-i$. In other words, id and rev are the permutations that arrange the elements in increasing and decreasing order, respectively. Suppose that $\sigma \in S_k \setminus \{\mathrm{id}, \mathrm{rev}\}$. Then there exist integers $a_\sigma, b_\sigma, c_\sigma, d_\sigma \in [k]$, not necessarily distinct, satisfying $a_\sigma < b_\sigma$, $c_\sigma < d_\sigma$, and

$$\sigma(a_\sigma) < \sigma(b_\sigma) \quad \text{and} \quad \sigma(c_\sigma) > \sigma(d_\sigma). \tag{4.2}$$

For every permutation $\sigma \in S_k \setminus \{\mathrm{id}, \mathrm{rev}\}$, we construct an ordered linear k-graph $(G_\sigma,<)$ as follows. Let G_σ be a k-graph on $3k-3$ vertices consisting of three edges e_1, e_2, and e_3, with $|e_i \cap e_j| = 1$ for all $i \neq j$. Let $<$ be an ordering of $V(G_\sigma)$ with $e_1 = \{x_1,\dots,x_k\}$, $e_2 = \{y_1,\dots,y_k\}$, and $e_3 = \{z_1,\dots,z_k\}$, where the vertices are labeled in increasing order in $<$, such that

$$x_{a_\sigma} = z_{c_\sigma}, \quad x_{b_\sigma} = y_{a_\sigma}, \quad \text{and} \quad y_{b_\sigma} = z_{d_\sigma}. \tag{4.3}$$

Observe that such an ordering is always possible (e.g., see Figure 4.2). Let $(G,<)$ be the ordered k-graph obtained by taking the vertex-disjoint union of $(G_\sigma,<)$ for all permutations $\sigma \in S_k \setminus \{\mathrm{id}, \mathrm{rev}\}$. Our ordered linear k-graph $(H,<)$ is the k-graph obtained by applying Theorem 4.2 to $(G,<)$ with $t = k!$ colors.

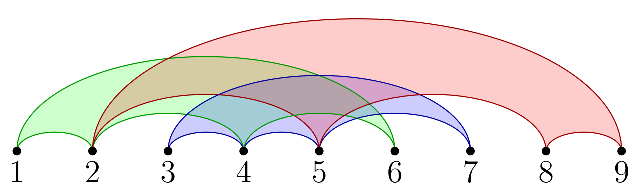

FIGURE 4.2. An example of G_σ for the permutation $\sigma \in S_4$ given by $\sigma(1) = 3$, $\sigma(2) = 1$, $\sigma(3) = 4$ and $\sigma(4) = 2$ and $a_\sigma = 2$, $b_\sigma = 3$, $c_\sigma = 1$ and $d_\sigma = 2$. The edges are given by $e_1 = \{1,2,4,6\}$ (green), $e_2 = \{3,4,5,7\}$ (blue) and $e_3 = \{2,5,8,9\}$ (red).

The importance of the k-graphs G_σ is illustrated in the next claim. Given an ordered edge $e = \{x_1, \ldots, x_k\}$ with $x_1 < \ldots < x_k$, a permutation $\sigma \in S_k$, and a total ordering $\prec$, we say that the edge e is *σ-compatible* with respect to the total ordering $\prec$ if

$$x_i \prec x_j \quad \text{if and only if} \quad \sigma(i) < \sigma(j). \tag{4.4}$$

In particular, the edge e is id-compatible with respect to $<$.

Claim 4.6. *Let $\sigma \in S_k \setminus \{\mathrm{id}, \mathrm{rev}\}$ and let $\prec$ be a total ordering of $V(G_\sigma)$. Then not all the edges of $(G_\sigma, <)$ are σ-compatible with respect to $\prec$.*

Proof. Suppose, for the sake of contradiction, that $\prec$ is a total ordering of $V(G_\sigma)$ such that all the edges are σ-compatible with respect to $\prec$. Let $e_1 = \{x_1, \ldots, x_k\}$, $e_2 = \{y_1, \ldots, y_k\}$, and $e_3 = \{z_1, \ldots, z_k\}$ be the edges of G_σ, where the vertices are labeled in increasing order with respect to $<$ and satisfy (4.3). Since all the edges are σ-compatible with respect to $\prec$, we have by (4.2), (4.3), and (4.4),

$$x_{a_\sigma} \prec x_{b_\sigma} = y_{a_\sigma} \prec y_{b_\sigma} = z_{d_\sigma} \prec z_{c_\sigma} = x_{a_\sigma},$$

which is a contradiction. This concludes the proof of the claim. □

We are now ready to prove that $(H, <)$ satisfies the statement. Let $\prec$ be a total ordering of $V(H)$, and let $S_k = \{\sigma_1, \ldots, \sigma_{k!}\}$ be a labeling of the $k!$ permutations. We define an auxiliary coloring $\chi : H \to [k!]$ of the edges of H as follows. For every edge $e \in H$, let $\chi(e) = i$ if the edge e is σ_i-compatible with respect to $\prec$. Since for every edge e, there exists a unique permutation that is compatible with respect to $\prec$, the auxiliary coloring χ is well defined.

By the construction of $(H, <)$ and Theorem 4.2, there exists a monochromatic copy of $(G, <)$ with respect to χ. In particular, this implies that there exists $\tau \in S_k$ such that every edge of G is τ-compatible with respect to $\prec$. Since G is the disjoint union of G_σ for $\sigma \in S_k \setminus \{\mathrm{id}, \mathrm{rev}\}$, we obtain by Claim 4.6 that $\tau \in \{\mathrm{id}, \mathrm{rev}\}$. If $\tau = \mathrm{id}$, then every edge of G satisfies (a). Otherwise, if $\tau = \mathrm{rev}$, then every edge of G satisfies (b). This concludes the proof of the proposition. □

We are now ready to prove Lemma 4.1.

Proof of Lemma 4.1. Let $t := c(P)$ and $n := c(Q)$ be the number of colors of P and Q. Let $m := e(Q)$ and enumerate the patterns of Q by $Q = \{q_1, \ldots, q_m\}$. We construct an ordered graph $(G, <)$ on the vertex set $[3m]$ with the natural order $<$ by taking $G = \bigcup_{j=1}^{m} T_j$ as the vertex-disjoint union of triangles T_j, with vertex set

$$V(T_j) = \{3j-2, 3j-1, 3j\},$$

for $1 \leqslant j \leqslant m$. We partition the edges of G into n edge-disjoint ordered graphs $G = \bigcup_{i=1}^{n} G_i$ as follows. For each $1 \leqslant j \leqslant m$, let $q_j = (a_j, b_j, c_j) \in Q$ be the j-th pattern of Q. We define the subgraphs G_i on the vertex set $[3m]$ by setting

$$E(G_i) = \bigcup_{j \in [m] : a_j = i} \{\{3j-2, 3j-1\}\} \cup \bigcup_{j \in [m] : b_j = i} \{\{3j-2, 3j\}\} \cup \bigcup_{j \in [m] : c_j = i} \{\{3j-1, 3j\}\} . \tag{4.5}$$

Informally speaking, each triangle T_j corresponds to the j-th pattern of Q and the graph G_i consists of all those pairs across all patterns which have the color i. It is easy to check that the G_i's are pairwise edge-disjoint and that $G = \bigcup_{i=1}^{n} G_i$ (e.g., see Figure 4.3).

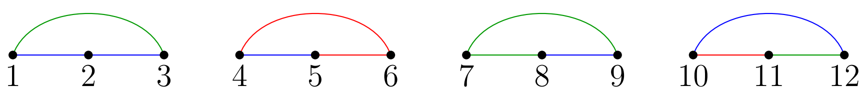

FIGURE 4.3. An example of G for the palette $Q = \{q_1, q_2, q_3, q_4\}$ given by $q_1 = (\text{blue}, \text{green}, \text{blue})$, $q_2 = (\text{blue}, \text{red}, \text{red})$, $q_3 = (\text{green}, \text{green}, \text{blue})$ and $q_4 = (\text{red}, \text{blue}, \text{green})$. The graph G consists of $m = 4$ triangles and it can be partitioned into $G_{\text{blue}} \cup G_{\text{green}} \cup G_{\text{red}}$ as shown in the picture.

We construct our desired 3-graph $F^{(3)}$ by applying Propositions 4.3 and 4.5. Let $(H, <)$ with $H = \bigcup_{i=1}^{n} H_i$ be the ordered graph obtained by applying Proposition 4.3 to the ordered graph $(G, <)$ with $G = \bigcup_{i=1}^{n} G_i$ and t-colors. Set $k := v(H)$ to be the number of vertices of H and let $(\mathcal{H}, <)$ be the linear k-graph obtained by Proposition 4.5. We construct the ordered graph $(A, <)$ with vertex set $V(A) = V(\mathcal{H})$ by replacing each edge $e \in \mathcal{H}$ with a copy $(H^e, <)$ of $(H, <)$. Since the k-graph $\mathcal{H}$ is linear, every two copies H^e and $H^{e'}$ in A intersect in at most one vertex. This in particular implies that $A = \bigcup_{i=1}^{n} A_i$ is the union of n edge-disjoint ordered graphs, where $A_i = \bigcup_{e \in \mathcal{H}} H_i^e$. Finally, the 3-graph $F := F^{(3)}$ is the hypergraph with vertex set $V(F) = V(A)$ and the edge set as follows. Let $\varphi : A \to [n]$ be the map defined by setting $\varphi(e) = i$ if and only if $e \in A_i$. With this notation in mind, the edge set of F is given by

$$F = \Big\{ \{x, y, z\} \in A^{(3)} : \ x < y < z, \ \{x, y\}, \ \{x, z\}, \ \{y, z\} \in A \\ \text{and } (\varphi(x, y), \varphi(x, z), \varphi(y, z)) \in Q \Big\} . \tag{4.6}$$

In other words, F is the 3-graph where the edges correspond to those triangles in A whose color pattern is given by the palette Q. Note that the construction given by (4.6) immediately gives us that Q paints F. Indeed, just take the natural order $<$ of $V(F)$ and consider the coloring $\chi : V(F)^{(2)} \to [n]$ given by

$$\chi(x,y) = \begin{cases} \varphi(x,y), & \text{if } \{x,y\} \in A, \\ 1, & \text{otherwise.} \end{cases}$$

We claim that P does not paint F. Suppose to the contrary that it does. Then there exists a total ordering $\lessdot$ of $V(F)$ and a coloring $\chi_P : V(F)^{(2)} \to [t]$ such that

$$(\chi_P(x,y), \chi_P(x,z), \chi_P(x,z)) \in P$$

for every $\{x,y,z\} \in F$ with $x \lessdot y \lessdot z$. Consider the restriction of χ_P to the edges of $A \subseteq V(F)^{(2)}$. By construction of A and Propositions 4.3 and 4.5, there exists a copy of $(G,<)$ with $G = G_1 \cup \ldots \cup G_n$ and vertex set $X = \{x_1, \ldots, x_{3m}\}$ with $x_1 < \ldots < x_{3m}$ such that

(i) For $1 \leqslant i \leqslant n$, the graph G_i is monochromatic with respect to χ_P
(ii) Either $x_1 \lessdot \ldots \lessdot x_{3m}$ or $x_1 \gtrdot \ldots \gtrdot x_{3m}$.

Note that by (4.6) and the definition of G, the induced graph $F[X]$ is just a matching of size m with edges $\{x_{3j-2}, x_{3j-1}, x_{3j}\}$ for $1 \leqslant j \leqslant m$ (see Figure 4.3). Let $\psi : C(Q) \to C(P)$ be the map defined by

$$\psi(i) = \chi_P(G_i),$$

i.e., $\psi(i)$ is the color of the monochromatic graph G_i. Claiming that this map gives a homomorphism from Q to P or a homomorphism from Q to $\mathrm{rev}(P)$, we split the proof into two cases.

<u>Case 1:</u> $x_1 \lessdot \ldots \lessdot x_{3m}$.

For $1 \leqslant j \leqslant m$, let $q_j = (a_j, b_j, c_j) \in Q$ be the j-th pattern of Q. By (4.5) we have that $\varphi(x_{3j-2}, x_{3j-1}) = a_j$, $\varphi(x_{3j-2}, x_{3j}) = b_j$ and $\varphi(x_{3j-1}, x_{3j}) = c_j$. Since χ_P is a witness that P paints F and we further have $x_{3j-2} \lessdot x_{3j-1} \lessdot x_{3j}$ and $\{x_{3j-2}, x_{3j-1}, x_{3j}\} \in F$, we infer that

$$(\psi(a_j), \psi(b_j), \psi(c_j)) = (\chi_P(x_{3j-2}, x_{3j-1}), \chi_P(x_{3j-2}, x_{3j}), \chi_P(x_{3j-1}, x_{3j})) \in P,$$

for $1 \leqslant j \leqslant m$. This in particular implies that ψ is a homomorphism from Q to P, which contradicts the assumption of the lemma.

<u>Case 2:</u> $x_1 \gtrdot \ldots \gtrdot x_{3m}$.

Similarly as in Case 1, since P paints F and $x_{3j} \lessdot x_{3j-1} \lessdot x_{3j-2}$, we have that

$$(\psi(c_j), \psi(b_j), \psi(a_j)) = (\chi_P(x_{3j-1}, x_{3j}), \chi_P(x_{3j-2}, x_{3j}), \chi_P(x_{3j-2}, x_{3j-1})) \in P,$$

for $1 \leqslant j \leqslant m$. This implies that $\psi(q_j) = (\psi(a_j), \psi(b_j), \psi(c_j)) \in \mathrm{rev}(P)$ and hence ψ is a homomorphism from Q to $\mathrm{rev}(P)$, which is again a contradiction. □

Remark 4.7. We observe that the same proof of Lemma 4.1 can be used to prove the statement for a finite family of palettes $\{P_1, \dots, P_k\}$. That is, given a palette Q that is not contained in a blow-up of P_i and $\mathrm{rev}(P_i)$ for $1 \leqslant i \leqslant k$, then there exists a 3-graph F such that Q paints F and none of the P_i's paints F.

§5 Properties of blow-ups of palettes

In this section, we discuss properties of palettes contained in a blow-up of a given palette P which will be necessary for the stability argument in Section 6. We begin by examining the interplay between P and $\mathrm{rev}(P)$. Let $C(P) = C(\mathrm{rev}(P)) = [t]$. The first observation is that since $(a, b, c) \in P$ if and only if $(c, b, a) \in \mathrm{rev}(P)$, it follows that P and $\mathrm{rev}(P)$ have the same Lagrange polynomial and, consequently, the same Lagrangian. Clearly, this still holds when inducing to any subset of colors.

Fact 5.1. $\Lambda_{P[U]} = \Lambda_{\mathrm{rev}(P)[U]}$ *for every* $U \subseteq [t]$.

It immediately follows that P is reduced if and only if $\mathrm{rev}(P)$ is reduced. Another simple observation is that if S_P and $S_{\mathrm{rev}(P)}$ are blow-ups of P and $\mathrm{rev}(P)$ on the same set of colors C and with the same partition structure $C = \bigcup_{i=1}^{t} V_i$, then $e(S_P) = e(S_{\mathrm{rev}(P)})$. In particular, this implies that the maximum blow-up of P on n colors has the same number of patterns as the maximum blow-up of $\mathrm{rev}(P)$ on n colors (something we have already used in the "in particular" part of Theorem 2.6).

The following observation shows that the Lagrangian of palettes is monotone with respect to homomorphisms.

Observation 5.2. *If Q and P are palettes and there is a homomorphism $\psi : Q \to P$, then*

$$\Lambda_Q \leqslant \Lambda_P \,.$$

Proof. Let $\mathbf{x} \in \mathbb{S}_{c(Q)}$ be a weighting of Q with $\lambda_Q(\mathbf{x}) = \Lambda_Q$. Then, for $d \in C(P)$, define

$$y_d = \sum_{a \in \psi^{-1}(d)} x_a .$$

Let $\mathbf{y} = (y_d)_{d \in C(P)}$ and note that $\mathbf{y} \in \mathbb{S}_{c(P)}$ because every $a \in C(Q)$ is in the preimage of exactly one $d \in C(P)$. Since ψ is a homomorphism and $\mathbf{x} \in \mathbb{S}_{c(Q)}$, it follows that

$$\begin{aligned}\Lambda_P \geqslant \lambda_P(\mathbf{y}) &= \sum_{(d,e,f)\in P} y_d y_e y_f \\ &= \sum_{(d,e,f)\in P} \left(\sum_{a\in\psi^{-1}(d)} x_a\right)\left(\sum_{b\in\psi^{-1}(e)} x_b\right)\left(\sum_{c\in\psi^{-1}(f)} x_c\right) \\ &\geqslant \sum_{(a,b,c)\in Q} x_a x_b x_c = \Lambda_Q.\end{aligned}$$

This concludes the proof. □

Recall that a palette P is *reduced* if for every proper subpalette $Q \subsetneq P$ we have $\Lambda_Q < \Lambda_P$. We conclude our discussion on the interplay between P and $\mathrm{rev}(P)$ by showing that if P is reduced and $P \not\cong \mathrm{rev}(P)$, then P is not contained in a blow-up of $\mathrm{rev}(P)$ and $\mathrm{rev}(P)$ is not contained in a blow-up of P.

Proposition 5.3. *Let P be reduced. If there is a homomorphism $\psi : P \to \mathrm{rev}(P)$, then $P \cong \mathrm{rev}(P)$.*

Proof. First, we check that ψ must be surjective. Indeed, if $\mathrm{Im}(\psi) \subsetneq C(\mathrm{rev}(P))$, then

$$\Lambda_P \leqslant \Lambda_{\mathrm{rev}(P)[\mathrm{Im}(\psi)]} = \Lambda_{P[\mathrm{Im}(\psi)]} < \Lambda_P,$$

where the first inequality follows from Observation 5.2, the equality follows from Fact 5.1, and the final inequality follows from P being reduced. This is a contradiction, and therefore ψ is surjective. Since P and $\mathrm{rev}(P)$ have the same number of colors and patterns, the surjective homomorphism ψ must be an isomorphism, concluding the proof. □

The next two results deal with properties of reduced palettes P. Roughly speaking, the first one states that any palette Q which is contained in a blow-up of P and has density $d(Q) = e(Q)/c(Q)^3$ very close to Λ_P must have a positive proportion of colors in each class of the partition structure.

Proposition 5.4. *Given a reduced palette P with t colors, there are $\beta = \beta_{5.4} > 0$ and $\varepsilon = \varepsilon_{5.4} > 0$ such that the following holds. Suppose that Q is a palette that is contained in a blow-up P' of P with partition structure $C(P') = \bigcup_{i=1}^t V_i$ and $c(P') = c(Q) := n$. If in addition we have $d(Q) \geqslant \Lambda_P - \varepsilon$, then $|V_i| \geqslant \beta n$ for every $i \in [t]$.*

Proof. Suppose for the sake of contradiction, that the statement does not hold. Then, for each integer $m \in \mathbb{N}$, there exist a blow-up $P'^{(m)}$ of P with partition structure $C^{(m)} :=$

$C(P'^{(m)}) = \bigcup_{i=1}^{t} V_i^{(m)}$ and a palette $Q^{(m)} \subseteq P'^{(m)}$ such that $d(Q^{(m)}) \geqslant \Lambda_P - \frac{1}{m}$ and

$$|V_{i^{(m)}}^{(m)}| < \frac{1}{m} c(Q^{(m)}) \tag{5.1}$$

for some sequence of indices $i^{(m)} \in [t]$. By applying the pigeonhole principle, we obtain a subsequence (which we reindex using m again) such that $i^{(m)}$ is constant, say $i^{(m)} = t$. Then, for each $m \in \mathbb{N}$, we have

$$\Lambda_{P[[t-1]]} \geqslant \Lambda_{Q^{(m)}[C^{(m)} \smallsetminus V_t^{(m)}]} \geqslant d(Q^{(m)}[C^{(m)} \smallsetminus V_t^{(m)}]) \geqslant d(Q^{(m)}) - \frac{1}{m} \geqslant \Lambda_P - \frac{2}{m}, \tag{5.2}$$

where the first inequality follows from Observation 5.2 applied to $P[[t-1]]$ and $Q^{(m)}[C^{(m)} \smallsetminus V_t^{(m)}]$, the second inequality comes from (2.2), the third from (5.1), and the fourth from the choice of $Q^{(m)}$. Since P is reduced, we must have $\Lambda_{P[[t-1]]} < \Lambda_P$, which contradicts (5.2) for m large enough. □

Given a palette P and two (distinct) colors $a, b \in C(P)$, we say that b *dominates* a if, for every pattern $p \in P$ containing a, any substitution of the color a with the color b results in a pattern $p' \in P$. As an example, suppose that $1, 2 \in C(P)$ and 2 dominates 1. Then $(1, 1, x) \in P$ implies that $(1, 2, x)$, $(2, 1, x)$, and $(2, 2, x)$ are all in P. Although it is straightforward to verify the following lemma, we include a proof for the convenience of the reader.

Lemma 5.5. *For a reduced palette P with $C(P) = [t]$ there are no $a, b \in [t]$ such that b dominates a.*

Proof. First note that we may assume that there is no $c \in [t]$ with $(c, c, c) \in P$. Otherwise $d(P[\{c\}]) = 1$, whence P being induced would imply $C(P) = \{c\}$, and we would be done.

Now suppose, for the sake of contradiction, that there are $a, b \in [t]$ such that b dominates a. For $\mathbf{z} \in \mathbb{S}_t$ we can write

$$\lambda_P(\mathbf{z}) = z_a f_a(\mathbf{z}') + z_b f_b(\mathbf{z}') + z_a^2 f_{a,a}(\mathbf{z}') + z_b^2 f_{b,b}(\mathbf{z}') + z_a z_b f_{a,b}(\mathbf{z}') + g(\mathbf{z}')$$

for some polynomials f_a, f_b, $f_{a,a}$, $f_{b,b}$, $f_{a,b}$, and g in $\mathbf{z}' = (z_c)_{c \in [t] \smallsetminus \{a,b\}}$. Let $\mathbf{x}$ be an optimal weighting of P witnessing $\lambda_P(\mathbf{x}) = \Lambda_P$. The hypothesis that b dominates a implies that $f_b \geqslant f_a$, as well as $f_{b,b} \geqslant f_{a,a}$, and $2f_{b,b} \geqslant f_{a,b}$, where the 2 appears since (a, b, x) and (b, a, x) are both 'covered' by (b, b, x). We claim that the weighting $\mathbf{y} \in \mathbb{S}_t$ given by $y_a = 0$, $y_b = x_a + x_b$, and $y_k = x_k$ for $k \in [t] \smallsetminus \{a, b\}$ satisfies

$$\lambda_P(\mathbf{y}) \geqslant \lambda_P(\mathbf{x}) = \Lambda_P. \tag{5.3}$$

Indeed,

$$\begin{aligned}\lambda_P(\mathbf{y}) &= (x_a + x_b) f_b(\mathbf{y}') + (x_a + x_b)^2 f_{b,b}(\mathbf{y}') + g(\mathbf{y}') \\ &\geqslant x_a f_b(\mathbf{x}') + x_b f_b(\mathbf{x}') + x_b^2 f_{b,b}(\mathbf{x}') + x_a^2 f_{b,b}(\mathbf{x}') + 2x_a x_b f_{b,b}(\mathbf{x}') + g(\mathbf{x}') \\ &\geqslant \lambda_P(\mathbf{x}),\end{aligned}$$

where $\mathbf{y}' = (y_c)_{c\in[t]\smallsetminus\{a,b\}} = (x_c)_{c\in[t]\smallsetminus\{a,b\}} = \mathbf{x}'$. On the other hand, note that $y_a = 0$ and therefore $\lambda_P(\mathbf{y}) \leqslant \Lambda_{P[[t]\smallsetminus\{a\}]} < \Lambda_P$, where we use the fact that P is reduced in the last inequality. This contradicts (5.3), which concludes the proof. $\square$

We finish the section by introducing a concept used in [30] that will be important in the stability argument. Let P be a palette with $C(P) = [t]$. A palette R contained in a blow-up of P is *rigid (with respect to P)* if there exists a partition of its colors $C(R) = \bigcup_{i=1}^t U_i$ satisfying the following: If R is contained in a blow-up S of P with partition structure $C(S) = \bigcup_{i=1}^t V_i$, then there exists an automorphism $h : [t] \to [t]$ of P such that $U_i \subseteq V_{h(i)}$ for $i \in [t]$. In other words, a palette R is rigid if there is essentially a unique way to embed it in a blow-up of P. The next result shows that if P is reduced, then rigid palettes always exist for sufficiently many colors. Recall that a palette is non-degenerate if every pattern has exactly 3 colors.

Lemma 5.6. *Let P be a reduced palette with color set $[t]$. Then there exists an integer $M := M_{5.6}(P)$ and a rigid palette $R \subseteq [M]^3$ with partition structure $[M] = \bigcup_{i=1}^t U_i$ such that*

(i) R is non-degenerate.
(ii) For $i \in [t]$, we have $|U_i| \geqslant 3t$.
(iii) Any blow-up R' of R on $M+1$ colors is rigid.

Moreover, if $P \ncong \mathrm{rev}(P)$, then the palette R is not contained in a blow-up of $\mathrm{rev}(P)$.

Proof. Since P is reduced, there exists a real number $\delta > 0$ such that $\Lambda_Q < \Lambda_P - \delta$ for every proper subset $Q \subsetneq P$. Let ε, β be the constants given by Proposition 5.4. We will choose a sufficiently large M satisfying the following conditions:

(a) $M \gg \frac{1}{\varepsilon}, \frac{1}{\delta}, t, \frac{1}{\beta}$.
(b) There exists a blow-up $\tilde{R}$ of P on M colors with partition structure $C(\tilde{R}) = \bigcup_{i=1}^t U_i$ such that $d(\tilde{R}) > \Lambda_P - \min\{\delta/2, \varepsilon/2\}$.

Note that condition (b) can always be satisfied because of (2.3). Let R be the palette obtained by removing every degenerate edge from $\tilde{R}$. It is not difficult to see from conditions

(a) and (b) that

$$d(R) \geqslant \frac{e(\tilde{R}) - 3M^2}{M^3} \geqslant d(\tilde{R}) - 3/M \geqslant \Lambda_P - \min\{\delta, \varepsilon\}. \tag{5.4}$$

We claim that the palette R is a rigid palette satisfying properties (i), (ii), and (iii).

We first check properties (i) and (ii). Property (i) follows immediately from the construction since we deleted all degenerate edges from $\tilde{R}$. To see that property (ii) holds, note that by (5.4), we have $d(R) \geqslant \Lambda_P - \varepsilon$. Hence, Proposition 5.4 gives us that $|U_i| \geqslant \beta M \geqslant 3t$, where the last inequality holds due to our choice of M (condition (a)).

We now proceed to prove that R is rigid. Clearly, R is contained in a blow-up of P (namely $\tilde{R}$). Let S be a blow-up of P with partition structure $C(S) = \bigcup_{j=1}^{t} V_j$ and let $\psi : C(R) \to C(S)$ be an embedding (i.e., an injective homomorphism) of R into S. We define a mapping $h : [t] \to [t]$ by letting $h(i)$ be an arbitrary index in $[t]$ such that

$$|\psi(U_i) \cap V_{h(i)}| \geqslant 3$$

for $i \in [t]$. Such a choice of h always exists because $|U_i| \geqslant 3t$ for $i \in [t]$. Let $Y_i \subseteq U_i$ be the preimage of $\psi(U_i) \cap V_{h(i)}$ under ψ, i.e., the subset of U_i with $\psi(Y_i) \subseteq V_{h(i)}$. Our goal is to prove that $h : [t] \to [t]$ is an automorphism of P and that $\psi(U_i) \subseteq V_{h(i)}$ for every $i \in [t]$ (that is, $Y_i = U_i$). We will do that in several claims.

Claim 5.7. *The map h is a homomorphism from P to P.*

Proof. Note that if $(i_1, i_2, i_3) \in P$, then there is a pattern $(a, b, c) \in Y_{i_1} \times Y_{i_2} \times Y_{i_3}$ in R. Indeed, by definition, R contains all the non-degenerate patterns in $U_{i_1} \times U_{i_2} \times U_{i_3}$. And since $|Y_i| = |\psi(U_i) \cap V_{h(i)}| \geqslant 3$ for all $i \in [t]$, even if $i_1 = i_2 = i_3$, there is some non-degenerate pattern in $Y_{i_1} \times Y_{i_2} \times Y_{i_3}$. However, since ψ is a homomorphism, the pattern $(\psi(a), \psi(b), \psi(c)) \in V_{h(i_1)} \times V_{h(i_2)} \times V_{h(i_3)}$ is a pattern of S. This implies that $(h(i_1), h(i_2), h(i_3)) \in P$ and consequently, $h : P \to P$ is a homomorphism. □

Claim 5.8. *The homomorphism $h : [t] \to [t]$ of P is bijective.*

Proof. We construct an auxiliary blow-up S^h of P with partition structure $C(S^h) = \bigcup_{j=1}^{t} V_j^h$ defined by

$$V_j^h = \bigcup_{i \in h^{-1}(j)} \psi(U_i).$$

That is, the sets V_j^h are the union of the images of U_i such that $h(i) = j$. It is clear that $C(S^h) \subseteq C(S)$. Moreover, note that some of the sets V_i^h might be empty.

We claim that the map $\psi : C(R) \to C(S^h)$ is an embedding of R into S^h such that $\psi(U_i) \subseteq V_{h(i)}^h$. The latter part holds by definition. To see that it is an embedding, let $(a, b, c) \in U_{i_1} \times$

$U_{i_2} \times U_{i_3}$ be a pattern of R. Then $(i_1, i_2, i_3) \in P$ and hence, since h is a homomorphism, we have that $(h(i_1), h(i_2), h(i_3)) \in P$. This implies that $(\psi(a), \psi(b), \psi(c)) \in V^h_{h(i_1)} \times V^h_{h(i_2)} \times V^h_{h(i_3)}$ is a pattern of S^h and $\psi : R \to S^h$ is an embedding.

Suppose, for the sake of contradiction, that h is not bijective. Then there exists an index $j \in [t]$ such that $V^h_j = \varnothing$. This implies that there exists a homomorphism $\varphi : R \to P[[t] \setminus \{j\}]$. Hence, by (2.2), Observation 5.2, and the definition of δ, we have

$$d(R) \leqslant \Lambda_R \leqslant \Lambda_{P[[t]\setminus\{j\}]} < \Lambda_P - \delta,$$

which contradicts (5.4). □

The last two claims show that $h : [t] \to [t]$ is an automorphism of P. Suppose, without loss of generality, that $h(i) = i$ (and henceforth we index both U_i and V_i by $i \in [t]$).

Claim 5.9. *For all $i \in [t]$ we have $\psi(U_i) \subseteq V_i$.*

Proof. Suppose to the contrary that there exist indices i, j and a color $x \in U_i$ such that $\psi(x) \in V_j$. We claim that j dominates i in P. Let $p \in P$ be a pattern containing the color i. Suppose, without loss of generality, that p is of the form $p = (i, k, \ell)$, where $k, \ell \in [t]$ (k and ℓ might be equal to i or j). Then we can choose distinct $b \in Y_k$ and $c \in Y_\ell$ (which are also both distinct from x) and so the triple $(x, b, c) \in U_i \times Y_k \times Y_\ell$ is a pattern of R. This implies that $(\psi(x), \psi(b), \psi(c)) \in V_j \times V_k \times V_\ell$ is a pattern of S. Hence, by construction, we have $(j, k, \ell) \in P$. That is, by substituting the color i with the color j, we still have a pattern of P. However, since P is reduced, by Lemma 5.5 there are no pairs $\{i, j\}$ with j dominating i, which yields a contradiction. □

Claims 5.7–5.9 show that R is a rigid configuration. Next we prove that R satisfies property (iii). First observe that the above argument verifying the rigidity of R used only properties (i) and (ii) of Lemma 5.6 and (5.4). Thus, it is sufficient to check these for any palette R' on $M + 1$ colors obtained by taking a blow-up of R. Viewing R' as a blow-up of P, it has partition structure $C(R') = \bigcup_{i=1}^t V'_i$, where $|V_i| \leqslant |V'_i| \leqslant |V_i| + 1$ for all $i \in [t]$ and so R' inherits properties (i) and (ii) from R. Moreover, by our choice of M in condition (a), we have

$$d(R') \geqslant \frac{e(\tilde{R}) - 3M^2}{(M+1)^3} \geqslant (d(\tilde{R}) - 3/M)\frac{M^3}{(M+1)^3} \geqslant \Lambda_P - \min\{\delta, \varepsilon\}.$$

Thus, by the same proof, the palette R' is rigid.

Finally, to prove the moreover part, suppose that $P \ncong \mathrm{rev}(P)$ and that there exists a homomorphism $\varphi : R \to \mathrm{rev}(P)$. Let $C(R) = \bigcup_{i=1}^t W_i$ be the partition structure of R as a subpalette of a blow-up of $\mathrm{rev}(P)$, i.e., $W_i = \varphi^{-1}(i)$. We construct a map $\xi : [t] \to [t]$ by letting $\xi(i)$ be an index such that

$$|U_i \cap W_{\xi(i)}| \geqslant 3.$$

Such an index always exists since $|U_i| \geqslant 3t$. For $i \in [t]$ let $Z_i = U_i \cap W_{\xi(i)}$. We claim that ξ is a homomorphism from P to $\mathrm{rev}(P)$. Let $(i_1, i_2, i_3) \in P$ be a pattern. Since R is obtained by deleting just the non-degenerate edges of a blow-up of P and $|Z_{i_1}|, |Z_{i_2}|, |Z_{i_3}| \geqslant 3$, there exists a pattern $(a, b, c) \in Z_{i_1} \times Z_{i_2} \times Z_{i_3}$. This implies that $(a, b, c) \in W_{\xi(i_1)} \times W_{\xi(i_2)} \times W_{\xi(i_3)}$ and consequently $(\xi(i_1), \xi(i_2), \xi(i_3)) \in \mathrm{rev}(P)$. Hence, $\xi : P \to \mathrm{rev}(P)$ is a homomorphism. A contradiction now follows from Proposition 5.3 and the fact that P is reduced. This concludes the proof of Lemma 5.6. □

§6 Stability argument

The main goal of this section is to provide a proof of Theorem 2.6. We remark that our approach is very similar to the approach in [30], but adapted to our needs. We begin with an outline of the proof. For a reduced palette $P \subseteq [t]^3$ with t colors, define the family $\mathcal{F}(P)$ of unpaintable 3-graphs as follows:

$$\mathcal{F}(P) := \{F : F \text{ is a 3-graph and } P \text{ is } F\text{-deficient}\}\,. \tag{6.1}$$

The first observation is that if Q is an $\mathcal{F}(P)$-deficient palette, then Q is contained in a blow-up of P or in a blow-up of $\mathrm{rev}(P)$ (see Lemma 4.1). Unfortunately, the family $\mathcal{F}(P)$ might be infinite in size and hence cannot be used directly as a witness for Theorem 2.6. To circumvent this issue, we choose an appropriate integer M and truncate the family $\mathcal{F}(P)$ to the subfamily $\mathcal{F}(P)_M$ of 3-graphs with at most M vertices, i.e.,

$$\mathcal{F}(P)_M := \{F \in \mathcal{F}(P) : v(F) \leqslant M\}\,. \tag{6.2}$$

Let Q be an $\mathcal{F}(P)_M$-deficient palette that maximizes the number of patterns among all $\mathcal{F}(P)_M$-deficient palettes; in other words, suppose that $Q \in \mathrm{EX}_{\mathrm{pal}}(n, \mathcal{F}(P)_M)$. An application of the removal lemma for palettes (see Lemma 3.3) implies that Q is very close to a blow-up of P or a blow-up of $\mathrm{rev}(P)$ (see Corollary 6.1 below). We then complete the proof using a stability argument, showing that any extremal palette sufficiently close to a blow-up of P must actually be a blow-up of P (see Lemma 6.2). Now let us proceed with the details.

We remind the reader that given two palettes P and Q on n colors, the edit distance $|P \triangle Q| = |P \smallsetminus Q| + |Q \smallsetminus P|$ is the minimum number of patterns that must be deleted or added to transform P into Q. Moreover, we say that P is α-close to Q if $|P \triangle Q| \leqslant \alpha n^3$.

The following is a corollary of Lemmata 3.3 and 4.1.

Corollary 6.1. *Given a palette P and $\alpha > 0$, there exist $M = M_{6.1}, N = N_{6.1} \in \mathbb{N}$ such that the following holds for every palette Q on $n \geqslant N$ colors. If Q is $\mathcal{F}(P)_M$-deficient, then it is α-close to being contained in a blow-up of P or a blow-up of* $\mathrm{rev}(P)$.

Proof. Given $\alpha > 0$, apply Lemma 3.3 to the family $\mathcal{F}(P)$ to obtain $M, N \in \mathbb{N}$ and $\beta > 0$. The conclusion of Lemma 3.3 entails that for every $\mathcal{F}(P)_M$-deficient palette Q with $c(Q) = n \geqslant N$, there exists an α-close palette Q' that does not paint $\mathcal{F}(P)$. Suppose for the sake of contradiction that Q' is not contained in a blow-up of P or $\mathrm{rev}(P)$. Then Lemma 4.1 yields a 3-graph F painted by Q' but not by P. In particular, $F \in \mathcal{F}(P)$, a contradiction. □

The next lemma is the key technical result of this section. We prove it using a stability argument similar to the one in [30].

Lemma 6.2. *Let P be a reduced palette with $C(P) = [t]$. There are $N := N_{6.2}(P) \in \mathbb{N}$, $M := M_{6.2}(P) \in \mathbb{N}$ and $\alpha := \alpha_{6.2}(P) > 0$ such that for all integers $n \geqslant N$ and $m \geqslant M$, the following holds. If $Q \in EX_{pal}(n, \mathcal{F}(P)_m)$ and Q is $\frac{\alpha}{2}$-close to a blow-up of P, then Q is a blow-up of P.*

Proof. Let $M_1 = M_{5.6}$ be the integer obtained from Lemma 5.6, and $\varepsilon = \varepsilon_{5.4}, \beta = \beta_{5.4} > 0$ be the constants given by Proposition 5.4. Set α and γ as real numbers such that

$$\alpha = \min\{(\beta/4)^{2M_1^2}, \varepsilon/3\} \quad \text{and} \quad \gamma = (\beta/4)^{2M_1}. \tag{6.3}$$

Consider the family of palettes $\mathcal{R}$ given by

$$\mathcal{R} := \{R : c(R) \leqslant M_1 + 3 \text{ and}$$
$$R \text{ is neither contained in a blow-up of } P \text{ nor in a blow-up of } \mathrm{rev}(P)\}.$$

For every $R \in \mathcal{R}$, by Lemma 4.1, there exists a 3-graph F_R such that R paints F_R and P does not paint F_R (i.e., $F_R \in \mathcal{F}(P)$). Let $M_2 := \max_{R \in \mathcal{R}}\{v(F_R)\}$. Such an M_2 always exists since $\mathcal{R}$ is a finite family. We will show that the statement of the lemma holds for M_2 as M and $N \in \mathbb{N}$ large enough. Fix integers $n \geqslant N$ and $m \geqslant M_2$. An immediate consequence of our choice is the following:

$$\text{If } Q \text{ does not paint } \mathcal{F}_m(P), \text{ then } Q \text{ is } \mathcal{R}\text{-free.} \tag{6.4}$$

Given $n \in \mathbb{N}$, let S be a blow-up of P with partition structure $[n] = \bigcup_{i=1}^t V_i$, i.e.,

$$S = \{(x, y, z) \in V_i \times V_j \times V_k : (i, j, k) \in P\}.$$

Given any such S we can define the set of *missing patterns* A and the set of *bad patterns* B of Q (with respect to S) by

$$A := S \setminus Q \quad \text{and} \quad B := Q \setminus S, \tag{6.5}$$

i.e., the patterns from S missing in Q and the patterns in Q that are not in our target blow-up S. Let $Q \in \mathrm{EX}_{\mathrm{pal}}(n, \mathcal{F}_{M_2}(P))$ be a maximum palette that does not paint $\mathcal{F}_{M_2}(P)$ and is $\frac{\alpha}{2}$-close to a blow-up $S \subseteq [n]^3$ of P. Since $|Q \triangle S| \leqslant \frac{\alpha}{2}n^3$, we have $|A| + |B| \leqslant \frac{\alpha}{2}n^3$. Moreover, because S is a blow-up of P, it does not paint $\mathcal{F}_M(P)$. Therefore, by the maximality of Q, we have that $|A| \leqslant |B| \leqslant \frac{\alpha}{2}n^3$.

It will be useful later to compare P not to S but instead to some blow-up S' of P with n colors that minimizes $|B|$ among all blow-ups of P with n colors. Fortunately, we can still check that such S' is not too far from Q. Indeed, if there exists a blow-up S' with partition structure $[n] = \bigcup_{i=1}^{t} V_i'$ such that the missing patterns $B' = Q \smallsetminus S'$ of Q with respect to S' satisfy $|B'| < |B|$, then by the maximality of Q, we have $|Q \triangle S'| = |A'| + |B'| < 2|B| \leqslant \alpha n^3$. Hence, at the marginal cost of a slightly larger edit distance, we can replace S by S' and assume that the partition $[n] = \bigcup_{i=1}^{t} V_i$ minimizes the number of bad patterns. To ease the notation, we write S and A and B instead of S', A', and B'. In particular, we now have $|Q \triangle S| = |A| + |B| \leqslant \alpha n^3$ and

$$|A| \leqslant |B| \leqslant \alpha n^3 \,. \tag{6.6}$$

Finally, note that by the maximality of Q, there exists some N so that we have $d(Q) \geqslant \Lambda_P - \alpha$ whenever $c(Q) = n \geqslant N$ (since such density can be achieved by a maximal blow-up of P on n colors). This implies by (6.3) that $d(S) \geqslant \Lambda_P - 3\alpha \geqslant \Lambda_P - \varepsilon$. Hence, by Proposition 5.4, we have

$$|V_i| \geqslant \beta n \tag{6.7}$$

for $i \in [t]$.

We claim that $|A| = |B| = 0$. Suppose, to the contrary, that $|B| > 0$. Given a color $x \in [n]$, let $\deg_B(x)$ be the number of patterns in B containing x (note that we do not count multiplicities, i.e., the pattern (x, x, y) counts only once). We define the *maximum degree* $\Delta(B)$ as the maximum of $\deg_B(x)$ over $x \in [n]$. We will first show that the palette of bad patterns B does not have a large maximum degree.

Claim 6.3. $\Delta(B) \leqslant \gamma n^2$.

Proof. Suppose that $\Delta(B) > \gamma n^2$ and let $x \in [n]$ such that $\deg_B(x) > \gamma n^2$. Suppose, without loss of generality, that $x \in V_1$. For every $k \in [t]$, we define a set $B_x^{(k)}$ as follows. Let $B^{(1)} := B$. For $k \neq 1$, consider the partition $[n] = \bigcup_{i=1}^{t} V_i^{(k)}$ with $V_i^{(k)} = V_i$ for $i \notin \{1, k\}$, and $V_1^{(k)} = V_1 \smallsetminus \{x\}$ as well as $V_k^{(k)} = V_k \cup \{x\}$. That is, $[n] = \bigcup_{i=1}^{t} V_i^{(k)}$ is the partition obtained by moving the color x from V_1 to V_k. Let $S^{(k)}$ be the blow-up of P with this partition structure, and let $B^{(k)} = Q \smallsetminus S^{(k)}$ be the new set of bad patterns of Q

with respect to $S^{(k)}$. Now we define

$$B_x^{(k)} := \{q \in B^{(k)} : x \in q\}$$

for $k \in [t]$. Recall that we chose S to minimize the number $|B| = |Q \setminus S|$, so we have that $|B| \leqslant |B^{(k)}|$ for $k \in [t]$. Since $|B^{(k)}| - |B| = |B_x^{(k)}| - |B_x^{(1)}|$, we obtain that

$$|B_x^{(k)}| \geqslant |B_x^{(1)}| = \deg_B(x) > \gamma n^2 . \tag{6.8}$$

for $k \in [t]$.

Let $R_\star$ be the non-degenerate rigid palette on M_1 colors obtained by Lemma 5.6 with partition $[M_1] = \bigcup_{i=1}^t U_i$. For each t-tuple $\mathbf{X} = (p_1, \dots, p_t) \in \prod_{k=1}^t B_x^{(k)}$ of patterns, we define a palette $R_\star^{\mathbf{X}}$ on $M_1 + 1$ colors as follows. Let $X = \{p_1, \dots, p_t\}$ be the palette containing the t elements of $\mathbf{X}$. Setting $W_i = V_i \cap (C(X) \setminus \{x\})$, we have the partitioning $C(X) = \left(\bigcup_{i=1}^t W_i\right) \cup \{x\}$. Consider an injective map $\iota : C(X) \to [M_1 + 1]$ such that $\iota(W_i) \subseteq U_i$ and $\iota(x) = M_1 + 1$. Such an embedding is always possible since $|U_i| \geqslant 3t$ for $i \in [t]$ (Property (ii) of Lemma 5.6) and $|W_i| \leqslant 3t$. Let $\iota(X)$ be the palette with patterns given by $\iota(p)$ for every $p \in X$. The palette $R_\star^{\mathbf{X}}$ is defined as the union $R_\star^{\mathbf{X}} = R_\star \cup \iota(X)$.

Subclaim 6.4. *$R_\star^{\mathbf{X}} \in \mathcal{R}$.*

Proof. Clearly, $c(R_\star^{\mathbf{X}}) \leqslant M_1 + 3$. We claim that $R_\star^{\mathbf{X}}$ is neither contained in a blow-up of P nor in a blow-up of $\mathrm{rev}(P)$. First, note that it suffices to check the claim only for P. Indeed, by Lemma 5.6, if $P \ncong \mathrm{rev}(P)$, then the palette $R_\star$ is not contained in a blow-up of $\mathrm{rev}(P)$. Since $R_\star \subseteq R_\star^{\mathbf{X}}$, this in particular implies that $R_\star^{\mathbf{X}}$ is not contained in a blow-up of $\mathrm{rev}(P)$.

Now suppose for the sake of contradiction that $R_\star^{\mathbf{X}}$ is contained in a blow-up S' of P with partition structure $\bigcup_{i=1}^t V_i'$. In particular, $R_\star \subseteq S'$, and by the definition of rigidity, there exists some automorphism $h : [t] \to [t]$ of P with $U_i \subseteq V'_{h(i)}$ for all $i \in [t]$. We may assume without loss of generality that $h(i) = i$, i.e., that $U_i \subseteq V_i'$ for $i \in [t]$. Suppose that the last color $M_1 + 1$ of $R_\star^{\mathbf{X}}$ is contained in V_k' for some index $k \in [t]$. Let p_k be the k-th pattern of X. From the fact that $\iota(x) = M_1 + 1 \in V_k'$ and $U_i \subseteq V_i'$, we obtain that the patterns p_k and $\iota(p_k)$ respect the same underlying structure in the blow-ups $S^{(k)}$ and S' of P, respectively. That is, $p_k \in V_a^{(k)} \times V_b^{(k)} \times V_c^{(k)}$ if and only if $\iota(p_k) \in V_a' \times V_b' \times V_c'$. Since we assumed that $R_\star^{\mathbf{X}} \subseteq S'$, we have that in particular $\iota(p_k) \in S'$. This implies that $p_k \in S^{(k)}$, which is a contradiction since $p_k \in B_x^{(k)} = Q \setminus S^{(k)}$. Therefore, $R_\star^{\mathbf{X}}$ is not contained in a blow-up of P. $\square$

Our goal now is to lower bound the cardinality of the set $A = S \setminus Q$ to obtain a contradiction. We say that an (injective) embedding $f : R_\star^{\mathbf{X}} \to Q \cup S$ is *good* if $f(M_1+1) = x$

and $f(\iota(p_k)) = p_k$ for $k \in [t]$. That is, if f embeds $\iota(X)$ into X. Fix a pair $(f, \mathbf{X})$ where f is a good embedding. By Subclaim 6.4 and (6.4), we obtain that $f(R_\star^{\mathbf{X}}) \nsubseteq Q$. Therefore, $f(R_\star^{\mathbf{X}}) \cap A \neq \varnothing$, i.e., there exists a pattern in $f(R_\star^{\mathbf{X}})$ that is a missing pattern. Let $\xi(f, \mathbf{X})$ be such a pattern. Since $X \subseteq Q$, $R_\star^{\mathbf{X}} = R_\star \cup \iota(X)$, and f embeds $\iota(X)$ into X we must have that $\xi(f, \mathbf{X}) \in f(R_\star)$ and so in particular $x \notin \xi(f, \mathbf{X})$. By property (i) of Lemma 5.6, this implies that $\xi(f, \mathbf{X})$ is a non-degenerate pattern (i.e., it has three distinct colors). For this reason, we will only estimate the number of non-degenerate missing patterns not containing x.

For $i \in [t]$, let $c_i(\mathbf{X})$ be the number of distinct colors of U_i present in the patterns of $\iota(X)$, and let $c(\mathbf{X}) = \sum_{i=1}^t c_i(\mathbf{X})$. Note that $c(\mathbf{X}) \leqslant 2t$ since each of the t patterns in X contains x. For a fixed $\mathbf{X} \in \prod_{k=1}^t B_x^{(k)}$, the number of good embeddings f is at least the number of ways to select, for each $i \in [t]$, $(|U_i| - c_i(\mathbf{X}))$ vertices from $V_i \smallsetminus \{X \cap V_i\}$. Therefore, for N large enough, the number of distinct pairs $(f, \mathbf{X})$ can be lower bounded by

$$\begin{aligned}\sum_{\mathbf{X} \in \prod_{k=1}^t B_x^{(k)}} \prod_{i=1}^t \left(\frac{|V_i|}{2}\right)^{|U_i| - c_i(\mathbf{X})} &\geqslant \sum_{\mathbf{X} \in \prod_{k=1}^t B_x^{(k)}} \left(\frac{\beta n}{2}\right)^{M_1 - c(\mathbf{X})} \\ &\geqslant \left(\frac{\beta n}{2}\right)^{M_1 - 2t} \prod_{k=1}^t |B_x^{(k)}| \geqslant \gamma^t n^{M_1} \left(\frac{\beta}{2}\right)^{M_1 - 2t}, \end{aligned} \tag{6.9}$$

where we use (6.7) and (6.8). Moreover, for a fixed non-degenerate missing pattern p not containing x, the number of pairs $(f, \mathbf{X})$ such that $p \in f(R_\star^{\mathbf{X}})$ is at most $n^{M_1 - 3}$ (since the pattern p and the vertex x are fixed). Therefore, the number of missing patterns can be estimated by

$$|A| \geqslant \gamma^t n^3 \left(\frac{\beta}{2}\right)^{M_1 - 2t} > \alpha n^3$$

by our choice of α and γ in (6.3). However, this contradicts (6.6), which concludes the proof. □

By applying a similar argument as in the last claim, one can obtain the following.

Claim 6.5. *For every bad pattern $q_B \in B$, there exist at least $(\beta/4)^{M_1} n^2$ non-degenerate missing patterns $q_A \in A$ such that $|C(q_A) \cap C(q_B)| = 1$.*

Proof. Fix $q_B = (a, b, c) \in B$. Let $\chi : C(q_B) \to [t]$ be the t-coloring such that $q_B \in V_{\chi(a)} \times V_{\chi(b)} \times V_{\chi(c)}$. Let $R_\star$ be the non-degenerate rigid palette on M_1 colors obtained by Lemma 5.6 with partition $[M_1] = \bigcup_{i=1}^t U_i$. We define the palette $R_\star^{q_B} \subseteq [M_1 + c(q_B)]^3$ as follows. Let $C(q_B) = \{x_1, \ldots, x_{c(q_B)}\}$ (note that $c(q_B)$ can be either 2 or 3, and that $\{a, b, c\} = \{x_1, \ldots, x_{c(q_B)}\}$ as sets). For $1 \leqslant k \leqslant c(q_B)$, let R_k be a palette with color set $[M_1] \cup \{M_1 + k\}$ obtained from $R_\star$ by blowing up a color in the set $U_{\chi(x_k)}$

to form the new vertex $M_1 + k$. Let $p_B \in \{M_1+1, \ldots, M_1 + c(q_B)\}^3$ be the pattern obtained by sending each color x_k in the pattern q_B to the color $M_1 + k$. Then we define $R_\star^{q_B} := \left(\bigcup_{k=1}^{c(q_B)} R_k\right) \cup \{p_B\}$ and remark that in this construction the only pattern in $R_\star^{q_B}$ which meets $\{M+1, \ldots, M + c(q_B)\}$ in more than one color is p_B.

Subclaim 6.6. $R_\star^{q_B} \in \mathcal{R}$.

Proof. As in the proof of Subclaim 6.4, it suffices to check that $R_\star^{q_B}$ is not contained in a blow-up of P. Suppose, for the sake of contradiction, that $R_\star^{q_B}$ is contained in a blow-up S' of P with partition structure $\bigcup_{i=1}^t V_i'$. This in particular implies that $R_\star \subseteq S'$ and $R_k \subseteq S'$ for each $1 \leqslant k \leqslant c(q_B)$. By the rigidity of $R_\star$, we may assume without loss of generality that $U_i \subseteq V_i'$ for $i \in [t]$. Moreover, by property (iii) of Lemma 5.6, we also obtain that $M_1 + k \in V'_{\chi(x_k)}$. Thus, and since $q_B \in V_{\chi(a)} \times V_{\chi(b)} \times V_{\chi(c)}$, we obtain that $p_B \in V'_{\chi(a)} \times V'_{\chi(b)} \times V'_{\chi(c)}$. This implies that $(\chi(a), \chi(b), \chi(c)) \in P$, which contradicts the fact that $q_B \in B$ is a bad pattern. $\square$

Our goal now is to lower bound the cardinality of the missing patterns $A = S \smallsetminus Q$. We define the subsets $A_0, A_1 \subseteq A$ as

$$A_0 := \{q_A \in A : c(q_A) = 3,\ C(q_A) \cap C(q_B) = \varnothing\}\,,$$
$$A_1 := \{q_A \in A : c(q_A) = 3,\ |C(q_A) \cap C(q_B)| = 1\}\,.$$

That is, A_0 consists of the non-degenerate missing patterns that are disjoint from q_B, and A_1 consists of those that intersect q_B in exactly one color. We say that an embedding $f : R_\star^{q_B} \to Q \cup S$ is *good* if f sends p_B to q_B. Let Ψ be the set of all good embeddings f. Fix $f \in \Psi$. By Subclaim 6.6 and (6.4), we obtain that $f(R_\star^{q_B}) \nsubseteq Q$. Therefore, $f(R_\star^{q_B})$ must contain a missing pattern $\xi(f) \in A$. Since $q_B \in B$, we obtain that $\xi(f) \in f(R_\star^{q_B} \smallsetminus \{p_B\})$. By recalling that the only pattern in $R_\star^{q_B}$ which meets $C(p_B)$ in more than one color is p_B, this implies that $\xi(f) \in A_0 \cup A_1$. Using (i) and (iii) of Lemma 5.6, the number of good embeddings f can be lower bounded by

$$|\Psi| \geqslant \prod_{i=1}^{t} \left(\frac{|V_i|}{2}\right)^{|U_i|} \geqslant \left(\frac{\beta n}{2}\right)^{M_1},$$

where we use (6.7). We can split the argument into two cases:

Case 1: For at least $|\Psi|/2$ embeddings f, the pattern $\xi(f)$ is in A_0.

For a pattern $q_A \in A_0$, the number of $f \in \Psi$ such that $q_A \in f(R_\star^{q_B})$ is at most n^{M_1-3}. Therefore,

$$|A| \geqslant |A_0| \geqslant \frac{|\Psi|}{2n^{M_1-3}} \geqslant \frac{n^3}{2}\left(\frac{\beta}{2}\right)^{M_1} > \alpha n^3,$$

by our choice of α in (6.3). However, this contradicts (6.6).

Case 2: For at least $|\Psi|/2$ embeddings f, the pattern $\xi(f)$ is in A_1.

For a pattern $q_A \in A_1$, the number of $f \in \Psi$ such that $q_A \in f(R_\star^{q_B})$ is at most n^{M_1-2} (since $|C(q_A) \cap C(q_B)| = 1$ and $c(q_A) = 3$). Therefore,

$$|A_1| \geqslant \frac{|\Psi|}{2n^{M_1-2}} \geqslant \frac{n^2}{2}\left(\frac{\beta}{2}\right)^{M_1} > \left(\frac{\beta}{4}\right)^{M_1} n^2 .$$

This concludes the proof of the claim. □

To finish the proof, we double count the number of pairs (q_A, q_B) where $q_B \in B$ is a bad pattern and $q_A \in A$ is a non-degenerate missing pattern such that $|C(q_A) \cap C(q_B)| = 1$. Fix $q_B \in B$. By Claim 6.5, there exist at least $(\beta/4)^{M_1} n^2$ patterns q_A. Therefore, there are at least $(\beta/4)^{M_1} n^2 |B|$ such pairs (q_A, q_B). On the other hand, for a fixed $q_A \in A$, there are at most $3\Delta(B)$ patterns q_B such that $|C(q_A) \cap C(q_B)| = 1$. Consequently, by Claim 6.3, the number of pairs (q_A, q_B) is at most $3\gamma n^2 |A|$. Combining the two bounds yields, by (6.3), that

$$|A| \geqslant \frac{(\beta/4)^{M_1}}{3\gamma} |B| > |B|,$$

which contradicts (6.6). Therefore $|A| = |B| = 0$ and consequently $Q = S$, concluding the proof of the lemma. □

We are now able to prove Theorem 2.6 as described in the outline at the beginning of the section.

Proof of Theorem 2.6. We start by setting up the constants. Let $M_{6.2}^P$, $N_{6.2}^P$, $\alpha_{6.2}^P$ and $M_{6.2}^{\mathrm{rev}(P)}$, $N_{6.2}^{\mathrm{rev}(P)}$, $\alpha_{6.2}^{\mathrm{rev}(P)}$ be the constants obtained by applying Lemma 6.2 to the palettes P and $\mathrm{rev}(P)$, respectively (Lemma 6.2 applies to $\mathrm{rev}(P)$ because, as remarked in Section 5, it is reduced if and only if P is). Let $M_{6.1}$ and $N_{6.1}$ be the constants obtained by applying Corollary 6.1 to the palette P with $\alpha = \frac{1}{2}\min\{\alpha_{6.2}^P, \alpha_{6.2}^{\mathrm{rev}(P)}\}$. Set $N_0 := \max\{N_{6.2}^P, N_{6.2}^{\mathrm{rev}(P)}, N_{6.1}\}$. Consider the family of palettes $\mathcal{R}$ given by

$$\begin{aligned}\mathcal{R} := \{R :\ & R \text{ is a palette with } c(R) \leqslant N_0 \\ & \text{and } R \text{ is not contained in a blow-up of } P \text{ nor in a blow-up of } \mathrm{rev}(P)\} .\end{aligned}$$

For every $R \in \mathcal{R}$, by Lemma 4.1, there exists a 3-graph F_R such that R paints F_R and P does not paint F_R. Let $M_4 := \max_{R \in \mathcal{R}}\{v(F_R)\}$ and note that this is well-defined since $\mathcal{R}$ is finite. Since $\{F_R\}_{R\in\mathcal{R}} \subseteq \mathcal{F}(P)$ and by our choice of M_4, every palette in $\mathcal{R}$ paints a 3-graph in $\mathcal{F}(P)_{M_4}$. In other words, we have that

$$\text{if } Q \text{ is } \mathcal{F}_{M_4}(P)\text{-deficient, then } Q \text{ is } \mathcal{R}\text{-free.} \tag{6.10}$$

Set $M := \max\{M^P_{6.2}, M^{\mathrm{rev}(P)}_{6.2}, M_{6.1}, M_4\}$. We claim that the family $\mathcal{H} := \mathcal{F}(P)_M$ satisfies the hypothesis of the theorem.

Let $n \in \mathbb{N}$ and let $Q \in \mathrm{EX}_{\mathrm{pal}}(n, \mathcal{H})$ be a palette with $c(Q) = n$ that maximizes the number of patterns among all $\mathcal{H}$-deficient palettes. Suppose first that $n \leqslant N_0$. By (6.10), the palette Q is $\mathcal{R}$-free, and then the definition of $\mathcal{R}$ and $c(Q) \leqslant N_0$ entail that Q is contained in a blow-up of P or in a blow-up of $\mathrm{rev}(P)$ and so Q is among the palettes which are considered in the right-hand side of (2.4). Hence, we may assume that $n \geqslant N_0$. Now, Corollary 6.1 implies that Q is α-close to being contained in a blow-up of P of $\mathrm{rev}(P)$. Therefore, by Lemma 6.2, Q is a blow-up of P or $\mathrm{rev}(P)$. $\square$

§7 Proof of main theorem

In this section we prove Theorem 1.1. As mentioned before, it essentially follows from Theorem 2.6 and the fact that $\pi_{\mathrm{pal}}(\mathcal{F}) = \pi_{\therefore}(\mathcal{F})$ holds for finite families. This equality in turn follows from the work in [31] (which is implicit in [34]) and [24]. To expand on this argument, we recall the notion of reduced hypergraphs from [34]. Essentially, reduced hypergraphs capture the setting that one arrives at after applying hypergraph regularity.

Definition 7.1. *A reduced 3-graph is a triple $\mathscr{A} = (I, \{\mathcal{P}^{ij}\}_{ij \in I^{(2)}}, \{\mathscr{A}^{ijk}\}_{ijk \in I^{(3)}})$ consisting of a finite index set I, a collection of pairwise disjoint sets of vertices $\{\mathcal{P}^{ij}\}_{ij \in I^{(2)}}$, and a collection of 3-partite 3-graphs $\{\mathscr{A}^{ijk}\}_{ijk \in I^{(ijk)}}$ such that for every $ijk \in I^{(3)}$ the vertex classes of $\mathscr{A}^{ijk}$ are $\mathcal{P}^{ij}$, $\mathcal{P}^{ik}$, and $\mathcal{P}^{jk}$. If $d \in [0,1]$ and $e(\mathscr{A}^{ijk}) \geqslant d|\mathcal{P}^{ij}||\mathcal{P}^{ik}||\mathcal{P}^{jk}|$ holds for all $ijk \in I^{(3)}$, we say that $\mathscr{A}$ is $(d, \therefore)$-dense.*

To ease the notation, we often simply write a reduced 3-graph as $\mathscr{A} = (I, \mathcal{P}^{ij}, \mathcal{A}^{ijk})$. By the hypergraph embedding lemma, in order to find a copy of a 3-graph F in the original host 3-graph H, it is sufficient to find a "reduced map" of F to a suitable reduced 3-graph. This is made formal with the following definition.

Definition 7.2. *A reduced map from a 3-graph F to a reduced 3-graph $\mathscr{A} = (I, \mathcal{P}^{ij}, \mathcal{A}^{ijk})$ is a pair (λ, φ) such that*

(i) $\lambda\colon V(F) \longrightarrow I$ and $\varphi\colon \partial F \longrightarrow \bigcup_{ij \in I^{(2)}} \mathcal{P}^{ij}$, where ∂F denotes the set of all pairs of vertices covered by an edge of F;
(ii) if $uv \in \partial F$, then $\lambda(u) \neq \lambda(v)$ and $\varphi(uv) \in \mathcal{P}^{\lambda(u)\lambda(v)}$;
(iii) if $uvw \in E(F)$, then $\varphi(uv)\varphi(uw)\varphi(vw) \in E(\mathscr{A}^{\lambda(u)\lambda(v)\lambda(w)})$.

If some such reduced map exists, we say that $\mathscr{A}$ contains a reduced image of F, and otherwise $\mathscr{A}$ is called F-free.

Given a family $\mathcal{F}$ of 3-graphs we say that a reduced 3-graph $\mathscr{A}$ is $\mathcal{F}$-free if it is F-free for all $F \in \mathcal{F}$. One can now define the Turán density of a family of 3-graphs with respect to reduced 3-graphs.

Definition 7.3. *If $\mathcal{F}$ is a family of 3-graphs, then*

$$\pi^{rd}_{\therefore}(\mathcal{F}) = \sup\big\{d \in [0,1] \colon \text{For every } m \in \mathbb{N} \text{ there is a } (d,\therefore)\text{-dense,}$$
$$\mathcal{F}\text{-free, reduced 3-graph with an index set of size } m\big\}\,.$$

The key behind almost all the progress on the uniform Turán problem in the past decade is that an argument based on the hypergraph regularity method yields the following result.

Theorem 7.4 (Theorem 3.3 in [31], implicit in [34])**.** *If $\mathcal{F}$ is a finite family of 3-graphs, then*

$$\pi^{rd}_{\therefore}(\mathcal{F}) = \pi_{\therefore}(\mathcal{F}).$$

For the next lemma, we need to set up the following notation. Let $\mathscr{A} = (U, \mathcal{P}^{ij}, \mathscr{A}^{ijk})$ be a reduced 3-graph and let $\mathcal{S}^{ij} \subseteq \mathcal{P}^{ij}$ be multisets, for all $ij \in U^{(2)}$. For all $ij \in U^{(2)}$, set $(\mathcal{S}^{ij})' = \{(x,r) : x \in \mathcal{S}^{ij}, r \in [\ell(x)]\}$, where $\ell(x)$ is the multiplicity of x in $\mathcal{S}^{ij}$. Next, for all $ijk \in U^{(3)}$, let $(\mathscr{A}^{ijk})'$ be the 3-partite 3-graph with vertex classes $(\mathcal{S}^{ij})'$, $(\mathcal{S}^{ik})'$, and $(\mathcal{S}^{jk})'$ and edge set

$$\{(x,a)(y,b)(z,c) : xyz \in \mathscr{A}^{ijk}, a \in [\ell(x)], b \in [\ell(y)], c \in [\ell(z)]\}\,.$$

The following was a crucial technical lemma used in [24], to obtain a palette from a reduced 3-graph with the appropriate dependence on ε and m.

Lemma 7.5. *For all $\varepsilon > 0$ there is some $s \in \mathbb{N}$ such that for all $m \in \mathbb{N}$ there is some $N \in \mathbb{N}$ with the following property. Every reduced 3-graph $\mathscr{A} = ([N], \mathcal{P}^{ij}, \mathscr{A}^{ijk})$ with density $d \in [0,1]$ contains an index subset $U \subseteq [N]$ with $|U| \geqslant m$ and multisets $\mathcal{S}^{ij} \subseteq \mathcal{P}^{ij}$, for all $ij \in U^{(2)}$, such that each $|\mathcal{S}^{ij}| = s$ and the reduced 3-graph $(U, (\mathcal{S}^{ij})', (\mathscr{A}^{ijk})')$ is $(d-\varepsilon, \therefore)$-dense.*

By applying Theorem 7.4, the following theorem suffices to complete the proof of Theorem 1.1.

Theorem 7.6. *Let P be a palette. Then there exists a finite family $\mathcal{H}$ of 3-graphs so that $\pi^{rd}_{\therefore}(\mathcal{H}) = \Lambda_P$.*

Proof. If P is not reduced, there exists a reduced palette $P' \subsetneq P$ with $\Lambda_{P'} = \Lambda_P$ that we could consider instead, so we may assume that P is reduced. Then Theorem 2.6 yields a finite family $\mathcal{H}$ such that P is $\mathcal{H}$-deficient and for all $n \in \mathbb{N}$

$$\mathrm{ex}_{\mathrm{pal}}(n, \mathcal{H}) = \max\{e(Q) : Q \text{ is a blow-up of } P \text{ and } c(Q) = n\}\,. \tag{7.1}$$

We will show that for every $\varepsilon > 0$ we have $\Lambda_P - \varepsilon \leqslant \pi^{\mathrm{rd}}_{\therefore}(\mathcal{H}) \leqslant \Lambda_P + 2\varepsilon$. So let $\varepsilon > 0$ be given. First we show that $\Lambda_P - \varepsilon \leqslant \pi^{\mathrm{rd}}_{\therefore}(\mathcal{H})$. By (2.3), there is some n_0 such that for every $n \in \mathbb{N}$ with $n \geqslant n_0$ there is a palette Q with n colors which is a blow-up of P, attains the maximum on the right-hand side in (7.1), and satisfies $\frac{e(Q)}{n^3} \geqslant \Lambda_P - \varepsilon$. Now it follows from (2.2), Fact 2.3, and Theorem 7.4 that

$$\Lambda_P - \varepsilon \leqslant \frac{e(Q)}{n^3} \leqslant \Lambda_Q \leqslant \pi_{\therefore}(\mathcal{H}) = \pi^{\mathrm{rd}}_{\therefore}(\mathcal{H})\,.$$

Next we show that $\pi^{\mathrm{rd}}_{\therefore}(\mathcal{H}) - 2\varepsilon \leqslant \Lambda_P$. First we observe that for all $s \in \mathbb{N}$ we have

$$\frac{\mathrm{ex}_{\mathrm{pal}}(s, \mathcal{H})}{s^3} \leqslant \Lambda_P\,. \tag{7.2}$$

Indeed, by (7.1) and (2.2), there is a blow-up of Q with $c(Q) = s$ such that $\mathrm{ex}_{\mathrm{pal}}(s, \mathcal{H})/s^3 = d(Q) \leqslant \Lambda_Q$. Since Q is a blow-up of P, Observation 5.2 entails that $\Lambda_Q \leqslant \Lambda_P$, and (7.2) follows.

Now let m be the maximum number of vertices of any $H \in \mathcal{H}$ and let $s \in \mathbb{N}$ be given by applying Lemma 7.5 with ε. Next, let R be the Ramsey number $R_3(m; 2^{s^3})$ (i.e., R is the smallest integer such that any coloring of the 3-edges of $K_R^{(3)}$ with 2^{s^3} colors contains a monochromatic $K_m^{(3)}$). Finally, let $N \in \mathbb{N}$ be as guaranteed by (the conclusion of) Lemma 7.5 applied to R here instead of m there. Let $\mathscr{A} = ([N], \mathcal{P}^{ij}, \mathscr{A}^{ijk})$ be a $(\pi^{\mathrm{rd}}_{\therefore}(\mathcal{H}) - \varepsilon, \therefore)$-dense $\mathcal{H}$-free reduced 3-graph. The conclusion of Lemma 7.5 provides an index subset $U \subseteq [N]$ with $|U| \geqslant R$ and multisets $\mathcal{S}^{ij} \subseteq \mathcal{P}^{ij}$ with $|\mathcal{S}^{ij}| = s$, for all $ij \in U^{(2)}$, such that the reduced 3-graph $(U, (\mathcal{S}^{ij})', (\mathscr{A}^{ijk})')$ is $(\pi^{\mathrm{rd}}_{\therefore}(\mathcal{H}) - 2\varepsilon, \therefore)$-dense. For all $ij \in U^{(2)}$ we identify $(\mathcal{S}^{ij})'$ (arbitrarily) with $[s]$. Then each $(\mathscr{A}^{ijk})'$ can be viewed as one of the s^3 possible subsets of $[s]^3$. This yields a 2^{s^3}-coloring of $U^{(3)}$, whence our choice of R provides an index set $U' \subseteq U$ with $|U'| = m$ and a subset $G \subseteq [s]^3$ so that for each $ijk \in (U')^{(3)}$ with $i < j < k$, $(\mathscr{A}^{ijk})'$ corresponds to G (under the fixed identifications of S^{ij}, S^{ik}, and S^{jk} with $[s]$). Now G can naturally be interpreted as a palette G' with $C(G') = [s]$ and $E(G') = G$. Since $e((\mathscr{A}^{ijk})') \geqslant (\pi^{\mathrm{rd}}_{\therefore}(\mathcal{H}) - 2\varepsilon)s^3$, it follows that $d(G') \geqslant \pi^{\mathrm{rd}}_{\therefore}(\mathcal{H}) - 2\varepsilon$. Further, it can easily be checked that if G' paints any $H \in \mathcal{H}$, this would entail a reduced image of H in $(U, (\mathcal{S}^{ij})', (\mathscr{A}^{ijk})')$ and thus in $\mathscr{A}$. Hence, $\mathscr{A}$ being $\mathcal{H}$-free, we know that G' is $\mathcal{H}$-deficient. Therefore, using (7.2), we get

$$\pi^{\mathrm{rd}}_{\therefore}(\mathcal{H}) - 2\varepsilon \leqslant \frac{e(G')}{s^3} \leqslant \frac{\mathrm{ex}_{\mathrm{pal}}(s, \mathcal{H})}{s^3} \leqslant \Lambda_P\,.$$

This concludes the proof of the theorem. □

§8 Concluding remarks

In this work, we obtain that the Lagrangian of any finite palette is attained as the uniform Turán density of a finite family of 3-graphs. A consequence of this result combined with [24] is that

$$\Lambda_{\mathrm{pal}} \subseteq \Pi_{\therefore,\mathrm{fin}} \subseteq \Pi_{\therefore,\infty} \subseteq \bar{\Lambda}_{\mathrm{pal}}\,. \tag{8.1}$$

It would be interesting to determine which of these inclusions are strict. In addition, if $\Pi_{\therefore}$ denotes the set of uniform Turán densities of single 3-graphs, is it true that $\Pi_{\therefore,} \subsetneq \Pi_{\therefore,\mathrm{fin}} \subsetneq \Pi_{\therefore,\infty}$? In [30], it was proved that the set of Turán densities of possibly infinite families of k-graphs, $\Pi_{\infty}^{(k)}$ is uncountable and closed for $k \geqslant 3$. One could ask whether similar statements hold for $\Pi_{\therefore,\infty}$. We remark that a direct application of the methods used in [30] does not seem to work here.

We say that $d \in [0,1)$ is a *jump* in a set $X \subseteq [0,1]$ if there exists some $\varepsilon > 0$ such that $(d, d+\varepsilon) \cap X = \varnothing$. Erdős [14] showed that for every $k \geqslant 2$, 0 is a jump in $\Pi_{\infty}^{(k)}$. On the other hand, Frankl and Rödl [15] proved that for every $k \geqslant 3$ there is some $d \in [0,1)$ that is not a jump in $\Pi_{\infty}^{(k)}$, disproving the famous Erdős jumping conjecture. For the uniform Turán density, Reiher, Rödl and Schacht [34] showed that 0 is a jump for $\Pi_{\therefore,\infty}$. A consequence of our work is that every non-jump in $\Pi_{\infty}^{(3)}$ yields a non-jump in $\Pi_{\therefore,\mathrm{fin}}$ (see [21]).

Note that the palettes considered here, as well as in [21] and [24], have finitely many colors. One might ask what the situation looks like for a (countably) infinite palette, which is a pair $P = (C, E)$ consisting of an infinite set of colors C and an infinite set of patterns $E \subseteq C^3$.

To define the Lagrangian of an infinite palette, note that the Lagrangian of a (finite) hypergraph F is (up to scaling) simply the maximum edge density a blow-up of F can have. Following this spirit, a reasonable way to define the "Lagrangian" of an infinite palette P is

$$\begin{aligned}\Lambda_P = \sup\{d \in [0,1] :\ & \text{for every } \eta > 0 \text{ and } n \in \mathbb{N}, \text{ there is} \\ & \text{a } (d,\eta)\text{-dense 3-graph } H \text{ with } v(H) \geqslant n \text{ such that } P \text{ paints } H\}\,.\end{aligned}$$

From [21] it follows that for finite palettes this definition is equivalent to our previous definition. Now one can ask whether Theorem 1.1 still holds for every infinite palette P. Lamaison and Wu [25] announced that there exists a 3-graph F such that there is no finite palette P that is F-deficient and satisfies $\Lambda_P = \pi_{\therefore}(F)$. Note that this means that $\Pi_{\therefore,\mathrm{fin}} \subsetneq \Lambda_{\mathrm{pal}}$. Setting $\Lambda_{\mathrm{pal},\infty} = \{\Lambda_P : P \text{ is a finite or infinite palette}\}$, it would be curious if in fact any of the sets $\Lambda_{\mathrm{pal},\infty}$, $\Pi_{\therefore,\mathrm{fin}}$, and $\Pi_{\therefore,\infty}$ are equal.

References

[1] F. G. Abramson and L. A. Harrington, *Models without indiscernibles*, J. Symbolic Logic **43** (1978), no. 3, 572–600. MR503795 ↑4

[2] N. Alon, E. Fischer, M. Krivelevich, and M. Szegedy, *Efficient testing of large graphs*, Combinatorica **20** (2000), no. 4, 451–476. ↑2, 3, A, A

[3] N. Alon and A. Shapira, *Every monotone graph property is testable*, Proceedings of the thirty-seventh annual acm symposium on theory of computing, 2005, pp. 128–137. ↑3

[4] N. Alon and J. H Spencer, *The probabilistic method*, 4th ed., Wiley Series in Discrete Mathematics and Optimization, John Wiley & Sons, Nashville, TN, 2016 (en). ↑A

[5] C. Avart, V. Rödl, and M. Schacht, *Every monotone 3-graph property is testable*, SIAM Journal on Discrete Mathematics **21** (2007), no. 1, 73–92. ↑3

[6] V. Bhat, J. Nešetřil, C. Reiher, and V. Rödl, *A Ramsey class for Steiner systems*, J. Combin. Theory Ser. A **154** (2018), 323–349. MR3718069 ↑4

[7] W. G. Brown and M. Simonovits, *Digraph extremal problems, hypergraph extremal problems, and the densities of graph structures*, Discret. Math. **48** (1984), 147–162. ↑1

[8] P. Erdős and A. Hajnal, *On Ramsey like theorems. Problems and results*, Combinatorics (Proc. Conf. Combinatorial Math., Math. Inst., Oxford, 1972), 1972, pp. 123–140. MR337636 ↑2

[9] P. Erdős and M. Simonovits, *A limit theorem in graph theory*, Studia Sci. Math. Hungar. **1** (1966), 51–57. MR205876 ↑1

[10] P. Erdős and V. T. Sós, *On Ramsey-Turán type theorems for hypergraphs*, Combinatorica **2** (1982), no. 3, 289–295. MR698654 ↑1

[11] P. Erdős and A. H. Stone, *On the structure of linear graphs*, Bull. Amer. Math. Soc. **52** (1946), 1087–1091. MR18807 ↑1

[12] P. Erdős, *Problems and results on graphs and hypergraphs: similarities and differences* (1990), 12–28. MR1083590 ↑1

[13] P. Erdös and G. Szekeres, *A combinatorial problem in geometry*, Compositio Math. **2** (1935), 463–470. MR1556929 ↑4

[14] P. Erdös, *On extremal problems of graphs and generalized graphs*, Israel Journal of Mathematics **2** (1964), no. 3, 183–190. ↑8

[15] P. Frankl and V. Rödl, *Hypergraphs do not jump*, Combinatorica **4** (1984), 149–159. ↑1, 8

[16] F. Garbe, D. Iľkovič, D. Kráľ, F. Kučerák, and A. Lamaison, *Hypergraphs with uniform Turán density equal to 8/27* (2024), available at arXiv:2407.05829. ↑1

[17] F. Garbe, D. Král', and A. Lamaison, *Hypergraphs with minimum positive uniform Turán density*, Israel J. Math. **259** (2024), no. 2, 701–726. MR4732978 ↑1

[18] R. Glebov, D. Král', and J. Volec, *A problem of Erdős and Sós on 3-graphs*, Israel J. Math. **211** (2016), no. 1, 349–366. MR3474967 ↑1, 1

[19] G. Katona, T. Nemetz, and M. Simonovits, *On a problem of Turán in the theory of graphs*, Mat. Lapok **15** (1964), 228–238 (Hungarian, with English and Russian summaries). MR172263 ↑1

[20] P. Keevash, *Hypergraph Turán problems* (2011), 83–139. MR2866732 ↑1
[21] D. King, M. Sales, and B. Schülke, *Lagrangians are attained as uniform Turán densities* (2024). ↑1, 1, 8
[22] Y. Kohayakawa, B. Nagle, V. Rödl, and M. Schacht, *Weak hypergraph regularity and linear hypergraphs*, Journal of Combinatorial Theory, Series B **100** (2010), no. 2, 151–160. ↑3.9
[23] D. Král', F. Kučerák, A. Lamaison, and G. Tardos, *Uniform Turán density—palette classification*. In preparation. ↑1
[24] A. Lamaison, *Palettes determine uniform Turán density* (2024). ↑1, 2, 2, 4, 7, 7, 8, 8
[25] A. Lamaison and Z. Wu, *Relating the Turán density and the uniform Turán density of hypergraphs*. In preparation. ↑8
[26] ———, *The uniform Turán density of large stars* (2024), available at arXiv:2409.03699. ↑1
[27] T. S. Motzkin and E. G. Straus, *Maxima for graphs and a new proof of a theorem of Turán*, Canadian J. Math. **17** (1965), 533–540. MR0175813 ↑1
[28] J. Nešetřil and V. Rödl, *Strong Ramsey theorems for Steiner systems*, Trans. Amer. Math. Soc. **303** (1987), no. 1, 183–192. MR896015 ↑1, 4, 4.2
[29] ———, *The partite construction and Ramsey set systems*, 1989, pp. 327–334. Graph theory and combinatorics (Cambridge, 1988). MR1001405 ↑4
[30] O. Pikhurko, *On possible Turán densities*, Israel Journal of Mathematics **201** (201204). ↑1, 2, 5, 6, 6, 8
[31] Chr. Reiher, *Extremal problems in uniformly dense hypergraphs*, European J. Combin. **88** (2020), 103117, 22. MR4111729 ↑1, 2, 2, 2, 7, 7.4
[32] C. Reiher, V. Rödl, and M. Schacht, *Hypergraphs with vanishing Turán density in uniformly dense hypergraphs*, J. Lond. Math. Soc. (2) **97** (2018), no. 1, 77–97. MR3764068 ↑1
[33] C. Reiher and V. Rödl, *The girth ramsey theorem* (2023). ↑4
[34] Chr. Reiher, V. Rödl, and M. Schacht, *On a Turán problem in weakly quasirandom 3-uniform hypergraphs*, J. Eur. Math. Soc. (JEMS) **20** (2018), no. 5, 1139–1159. MR3790065 ↑1, 1, 7, 7.4, 8
[35] V. Rödl, *On universality of graphs with uniformly distributed edges*, Discrete Math. **59** (1986), no. 1-2, 125–134. MR837962 ↑2
[36] P. Turán, *Eine Extremalaufgabe aus der Graphentheorie*, Mat. Fiz. Lapok **48** (1941), 436–452 (Hungarian, with German summary). MR18405 ↑1

Appendix A. Weak Palette Regularity

Before we begin with details we provide a brief overview. Our proof of Theorem 3.7 will be as expected. We will define an energy function q which takes as argument a partition of $C(Q)$ and returns an element of $[0,1]$, and show both that refining a partition cannot decrease the value of q, and that if $\mathcal{A}$ fails to be ε-regular then there exists a refinement which substantially increases q. Once we have proven Theorem 3.7, Corollary 3.8 will be obtained through the following intermediary result.

Theorem A.1. *For all non-increasing $\mathcal{E}(r) : \mathbb{N} \to (0,1]$ and m there exists $M = M_{A.1}, N = N_{A.1}$ so that given any palette Q with $c(Q) \geqslant N$ there is an equipartition $\mathcal{A} = \{V_i : i \in [t]\}$ of $C(Q)$ with refinement $\mathcal{B} = \{V_{i,j} : i \in [t], j \in [\ell]\}$ so that:*

(1) *$m \leqslant t$ and $t\ell \leqslant M$,*
(2) *all but $\mathcal{E}(0)t^3$ of $(i_1, i_2, i_3) \in [t]^3$ have the ordered triple $(V_{i_1}, V_{i_2}, V_{i_3})$ $\mathcal{E}(0)$-regular,*
(3) *for all $(i_1, i_2, i_3) \in [t]^3$, all but $\mathcal{E}(t)\ell^3$ of $(j_1, j_2, j_3) \in [\ell]^3$ have the ordered triple $(V_{i_1,j_1}, V_{i_2,j_2}, V_{i_3,j_3})$ $\mathcal{E}(t)$-regular, and*
(4) *for all but $\mathcal{E}(0)t^3$ of $(i_1, i_2, i_3) \in [t]^3$, all but $\mathcal{E}(0)\ell^3$ of $(j_1, j_2, j_3) \in [\ell]^3$ have $|d(V_{i_1,j_1}, V_{i_2,j_2}, V_{i_3,j_3}) - d(V_{i_1}, V_{i_2}, V_{i_3})| < \mathcal{E}(0)$.*

The relationship between our Theorem 3.7, Corollary 3.8, and Theorem A.1 is entirely analogous to those found in [2] between Lemma 3.3 (the traditional Szemerédi graph regularity lemma), Lemma 4.1, and Corollary 4.2 there. Our proofs use the same strategies.

Let us begin now with the details of Theorem 3.7. We follow closely the probabilistic techniques presented in [4], working here with palettes in place of graphs. For ease of notation we will still refer to colors as vertices, and patterns as edges on these vertices. We are interested in partitions of $C(Q)$, and we let $n = c(Q)$ throughout. We will, in intermediate steps, allow our equipartitions to include a small set V_0 of exceptional vertices. As part of a partition of $C(Q)$ we consider V_0 as composed of singleton vertices, so that $\mathcal{B} = U_0 \cup U_1 \cup \ldots \cup U_{t'}$ refines $\mathcal{A} = V_0 \cup V_1 \cup \ldots \cup V_t$ (denoted $\mathcal{B} \prec \mathcal{A}$) so long as, for each $i \in [t]$, V_i is obtained exactly as the union of some U_j together with some vertices from U_0, and $V_0 \subseteq U_0$. Recalling Definition 3.4, we say that a partition $\mathcal{A} = V_0 \cup V_1 \cup \ldots \cup V_t$ is ε-regular if for all but εt^3 of $(i_1, i_2, i_3) \in [t]^3$, the ordered triple $(V_{i_1}, V_{i_2}, V_{i_3})$ is ε-regular. When we index over $V \in \mathcal{A}$ for a partition $\mathcal{A}$, we mean to take one term for each part of $\mathcal{A}$, denoted V. With these notions in mind we can now define the energy q.

Definition A.2. *Suppose that $V_1, V_2, V_3 \subseteq C(Q)$ with $c(Q) = n$. Then*

$$q(V_1, V_2, V_3) := \frac{|V_1||V_2||V_3|}{n^3} d^2(V_1, V_2, V_3)\,.$$

If $\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3$ are partitions of $C(Q)$, we set

$$q(\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3) := \sum_{V_1 \in \mathcal{A}_1, V_2 \in \mathcal{A}_2, V_3 \in \mathcal{A}_3} q(V_1, V_2, V_3)\,.$$

We will use $q(\mathcal{A})$ to refer to $q(\mathcal{A}, \mathcal{A}, \mathcal{A})$ along a single partition.

Since $q(V_1, V_2, V_3) \leqslant \frac{|V_1||V_2||V_3|}{n^3}$ it is immediate that $q(\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3) \in [0, 1]$. The following technical lemma captures two more important properties we will require of the function q.

Lemma A.3. *If $\mathcal{B}_1 \prec \mathcal{A}_1, \mathcal{B}_2 \prec \mathcal{A}_2$, and $\mathcal{B}_3 \prec \mathcal{A}_3$ then $q(\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3) \geqslant q(\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$. Furthermore, if $V_1, V_2, V_3 \subseteq C(Q)$ are so that (V_1, V_2, V_3) is not ε-regular, then there are partitions $V_1 = V_1^1 \cup V_1^2, V_2 = V_2^1 \cup V_2^2, V_3 = V_3^1 \cup V_3^2$ so that*

$$q(V_1^1 \cup V_1^2, V_2^1 \cup V_2^2, V_3^1 \cup V_3^2) \geqslant q(V_1, V_2, V_3) + \varepsilon^5 \frac{|V_1||V_2||V_3|}{n^3}\,.$$

Proof. For the first part, it suffices to check the case when $\mathcal{A}_1 = V_1, \mathcal{A}_2 = V_2, \mathcal{A}_3 = V_3$ consist of a single set each, since any $\mathcal{B}_i$ can be obtained by successive refinement in this way. In this case we define a random variable Z as follows. Select vertices $x_1 \in V_1, x_2 \in V_2, x_3 \in V_3$ uniformly at random and let $U_1 \in \mathcal{B}_1, U_2 \in \mathcal{B}_2, U_3 \in \mathcal{B}_3$ be the unique parts of their respective partitions so that $x_1 \in U_1, x_2 \in U_2, x_3 \in U_3$ before setting $Z = d(U_1, U_2, U_3)$. We can directly compute both the expectation

$$\begin{aligned}\mathbb{E}(Z) &= \sum_{U_1\in\mathcal{B}_1, U_2\in\mathcal{B}_2, U_3\in\mathcal{B}_3} \frac{|U_1||U_2||U_3|}{|V_1||V_2||V_3|} d(U_1, U_2, U_3) \\ &= \sum_{U_1\in\mathcal{B}_1, U_2\in\mathcal{B}_2, U_3\in\mathcal{B}_3} \frac{e(U_1, U_2, U_3)}{|V_1||V_2||V_3|} \\ &= d(V_1, V_2, V_3)\end{aligned}$$

and the second moment

$$\begin{aligned}\mathbb{E}(Z^2) &= \sum_{U_1\in\mathcal{B}_1, U_2\in\mathcal{B}_2, U_3\in\mathcal{B}_3} \frac{|U_1||U_2||U_3|}{|V_1||V_2||V_3|} d^2(U_1, U_2, U_3) \\ &= \frac{n^3}{|V_1||V_2||V_3|} \sum_{U_1\in\mathcal{B}_1, U_2\in\mathcal{B}_2, U_3\in\mathcal{B}_3} \frac{|U_1||U_2||U_3|}{n^3} d^2(U_1, U_2, U_3) \\ &= \frac{n^3}{|V_1||V_2||V_3|} q(\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3)\,.\end{aligned}$$

Combining these yields

$$0 \leqslant \mathrm{Var}(Z) = \mathbb{E}(Z^2) - \mathbb{E}(Z)^2 = \frac{n^3}{|V_1||V_2||V_3|} \Big(q(\mathcal{B}_1, \mathcal{B}_2, \mathcal{B}_3) - q(V_1, V_2, V_3) \Big)$$

as needed.

Now we proceed to the 'furthermore' part. Let V_i^1 be witness sets to the failure of ε-regularity and V_i^2 their complements, so that $|V_i^1| \geqslant \varepsilon |V_i|$ but $|d(V_1^1, V_2^1, V_3^1) - d(V_1, V_2, V_3)| \geqslant \varepsilon$. Let Z be the random variable defined above on the new partition obtained of V_1, V_2, V_3, where $\mathcal{B}_i = V_i^1 \cup V_i^2$. Then Chebyshev's inequality shows

$$\varepsilon^3 \leqslant \frac{|V_1^1||V_2^1||V_3^1|}{|V_1||V_2||V_3|} \leqslant \Pr(|Z - \mathbb{E}(Z)| \geqslant \varepsilon) \leqslant \frac{\mathrm{Var}(Z)}{\varepsilon^2}$$

and we are done, having computed $\mathrm{Var}(Z)$ above. □

The furthermore part of the previous lemma shows that irregular triples can be refined to increase the value of q by a small amount. Next we will show that if $\mathcal{A}$ has many irregular triples, by refining each we can increment q by a small constant depending only on ε, while still controlling the order of the partitions we create.

Lemma A.4. *Suppose $\mathcal{A}$ is an equipartition of $C(Q)$ into $V_0 \cup V_1 \ldots \cup V_t$ and $t \geqslant 6$. If there are εt^3 many $(i_1, i_2, i_3) \in [t]^3$ with the ordered triple $(V_{i_1}, V_{i_2}, V_{i_3})$ failing to be ε-regular and $|V_0| \leqslant \varepsilon n$,then there exists a refinement $\mathcal{B} = U_0 \cup U_1 \ldots \cup U_\ell$ of $\mathcal{A}$ with $q(\mathcal{B}) \geqslant q(\mathcal{A}) + \frac{\varepsilon^6}{16}$, $|U_0| \leqslant |V_0| + \frac{n}{2^t}$ and $\ell \leqslant 2t2^{3t^2}2^t$.*

Proof. Any time that V_i appears as part of an irregular triple (in first, second, or third position) with distinct indices, we will apply Lemma A.3 to partition V_i into V_i^1 and V_i^2. Formally, for every $i \in [t]$ and $(i_1, i_2, i_3) \in [t]^3$ with $i \in \{i_1, i_2, i_3\}$ and i_1, i_2, i_3 all distinct, we define a partition $\mathcal{V}^{(i)}_{(i_1,i_2,i_3)}$ of V_i. If $(V_{i_1}, V_{i_2}, V_{i_3})$ is ε-regular we let $\mathcal{V}^{(i)}_{(i_1,i_2,i_3)}$ be the trivial partition consisting of the single set V_i, and if $(V_{i_1}, V_{i_2}, V_{i_3})$ is not ε-regular we let $\mathcal{V}^{(i)}_{(i_1,i_2,i_3)}$ be the partition furnished by Lemma A.3. Each $\mathcal{V}^{(i)}_{(i_1,i_2,i_3)}$ consists of either 1 or 2 sets, so for fixed i the mutual refinement consists of at most 2^{3t^2} parts. Let $\tilde{\mathcal{B}}$ be the partition obtained by mutually refining each V_i in this manner, so that (since there at least εt^3 irregular triples and at most $\varepsilon 3t^2 \leqslant \frac{\varepsilon}{2} t^3$ of them have a repeated index)

$$q(\tilde{\mathcal{B}}) \geqslant q(\mathcal{A}) + \frac{\varepsilon}{2} t^3 \varepsilon^5 \frac{1}{8t^3}.$$

At this point $\tilde{\mathcal{B}}$ has incremented q as desired; all that remains is to balance the sizes of the member sets. $\tilde{\mathcal{B}}$ has at most $t2^{3t^2}$ parts, say $\tilde{U}_j$, together with the exceptional set V_0 remaining from $\mathcal{A}$. Let $\mathcal{B}$ be obtained from $\tilde{\mathcal{B}}$ by taking, from each $\tilde{U}_j$, a maximal disjoint family of sets of size $\frac{n}{t2^{3t^2}2^t}$, say $U_{j,k}$ for $0 \leqslant k \leqslant K_j$, and adding all of the leftover vertices to V_0 to form U_0. Formally,

$$U_0 = V_0 \cup \left(\bigcup_j \tilde{U}_j \setminus \left(\bigcup_{1 \leqslant k \leqslant K_j} U_{j,k} \right) \right).$$

The first part of Lemma A.3 gives $q(\mathcal{B}) \geqslant q(\tilde{\mathcal{B}}) \geqslant q(\mathcal{A}) + \frac{\varepsilon^6}{16}$. Overcounting the number of vertices discarded gives $|U_0| \leqslant |V_0| + \frac{n}{2^t}$, and flattening pairs (j,k) to a single index ℓ gives an equipartition $\mathcal{B}$ with at most $2t2^{3t^2}2^t$members, as needed. □

Finally, the following lemma will be used to redistribute the exceptional set V_0 amongst the other classes of the equipartition without destroying regularity.

Lemma A.5. *For all $\varepsilon > 0$ there exists $\gamma > 0$ so that the following holds. Suppose that (V_1, V_2, V_3) is ε-regular in $C(Q)$ and that $X_1, X_2, X_3 \subseteq C(Q)$ with $|X_i| \leqslant \gamma |V_i|$. Then $(V_1 \cup X_1, V_2 \cup X_2, V_3 \cup X_3)$ is 2ε-regular.*

Proof. We may assume that $\gamma < \varepsilon$; we will verify directly that $(V_1 \cup X_1, V_2 \cup X_2, V_3 \cup X_3)$ is 2ε-regular. Suppose that $W_i \subseteq (V_i \cup X_i)$ with $|W_i| \geqslant 2\varepsilon |V_i \cup X_i|$, and set $W_i^V := W_i \cap V_i$

with leftovers $W_i^X := W_i \setminus W_i^V$. Then $|W_i^V| \geqslant |W_i| - \gamma|V_i| \geqslant \varepsilon|V_i|$, so by the ε-regularity of (V_1, V_2, V_3) it follows that

$$|d(W_1^V, W_2^V, W_3^V) - d(V_1, V_2, V_3)| \leqslant \varepsilon\,.$$

Next observe the simple subset

$$e(V_1, V_2, V_3) \leqslant e(V_1 \cup X_1, V_2 \cup X_2, V_3 \cup X_3)$$

and union bounds

$$\begin{aligned} e(V_1 \cup X_1, V_2 \cup X_2, V_3 \cup X_3) \leqslant e(V_1, V_2, V_3) &+ |X_1||V_2 \cup X_2||V_3 \cup X_3| \\ &+ |V_1 \cup X_1||X_2||V_3 \cup X_3| \\ &+ |V_1 \cup X_1||V_2 \cup X_2||X_3|\,. \end{aligned}$$

Dividing through by $|V_1 \cup X_1||V_2 \cup X_2||V_3 \cup X_3|$ yields

$$\frac{1}{(1+\gamma)^3} d(V_1, V_2, V_3) \leqslant d(V_1 \cup X_1, V_2 \cup X_2, V_3 \cup X_3) \leqslant d(V_1, V_2, V_3) + 3\gamma$$

and repeating the same argument with W_i^V in place of V_i and W_i^X in place of X_i gives

$$\frac{1}{(1+\frac{\gamma}{\varepsilon})^3} d(W_1^V, W_2^V, W_3^V) \leqslant d(W_1, W_2, W_3) \leqslant d(W_1^V, W_2^V, W_3^V) + 3\frac{\gamma}{\varepsilon}\,.$$

Then for γ taken small enough as a function of ε, it follows that

$$|d(V_1, V_2, V_3) - d(V_1 \cup X_1, V_2 \cup X_2, V_3 \cup X_3)| \leqslant \varepsilon/2$$

and

$$|d(W_1, W_2, W_3) - d(W_1^V, W_2^V, W_3^V)| \leqslant \varepsilon/2\,.$$

Finally the triangle inequality shows

$$|d(W_1, W_2, W_3) - d(V_1 \cup X_1, V_2 \cup X_2, V_3 \cup X_3)| \leqslant 2\varepsilon$$

as needed. □

Proof of Theorem 3.7. We directly prove the 'more generally' part of the Theorem, which implies the first part by taking an arbitrary equipartition of m parts. Let γ be the result of Lemma A.5 applied to ε. By increasing the value of m we may assume that $\lceil \frac{16}{\varepsilon^6} \rceil \frac{1}{2^m} \leqslant \min(\varepsilon, \gamma/4)$ and $m \geqslant 6$ - this suffices for the general case. Inductively define a sequence $t_0 = 2m$ and $t_{i+1} = 2t_i 2^{3t_i^2} 2^{t_i}$. We will show that taking $M = N = t_{\lceil \frac{16}{\varepsilon^6} \rceil}$ suffices.

Indeed, let $\mathcal{A}^0$ be the initial equipartition with $s_0 = m$ parts and $V_0^0 = \varnothing$. Iterate the following process, beginning with $i = 0$. If $\mathcal{A}^i$ is not ε-regular, apply Lemma A.4 to obtain a refinement $\mathcal{A}^{i+1}$ with $q(\mathcal{A}^{i+1}) \geqslant q(\mathcal{A}^i) + \frac{\varepsilon^6}{16}$, so that $|V_0^{i+1}| \leqslant |V_0^i| + \frac{n}{2^{s_i}}$ and $\mathcal{A}^{i+1}$ has $s_{i+1} \leqslant t_{i+1}$ parts (since the inductive definition of t_i matches the output of Lemma A.4).

Since $q(\cdot) \in [0,1]$, it follows that in $\lceil \frac{16}{\varepsilon^6} \rceil$ steps we must find some $\mathcal{A}^j$ which is ε-regular, with $s \leqslant M$ parts, at which point our iterative process halts.

We can estimate the size of the exceptional set as

$$|V_0^j| \leqslant n \left(\left\lceil \frac{16}{\varepsilon^6} \right\rceil \frac{1}{2^m} \right) \leqslant \varepsilon n \,.$$

All that remains is to redistribute V_0^j amongst the other parts of the equipartition. By evenly distributing the vertices of V_0^j no part will receive more than $2\frac{|V_0^j|}{s} \leqslant \gamma$ many vertices with $2\frac{|V_0^j|/s}{|V_i|} \leqslant \gamma$ so that Lemma A.5 guarantees the 2ε-regularity of our equipartition. Applying the above proof to $\varepsilon' = \frac{\varepsilon}{2}$ would provide ε-regularity; we leave the proof as written for readability. □

Next we show, as done in [2], how to iterate Theorem 3.7 and then randomize to obtain Corollary 3.8. To do so we will require another property of q. If $\mathcal{B} \prec \mathcal{A}$ it follows from Lemma A.3 that $q(\mathcal{B}) \geqslant q(\mathcal{A})$; we will show that if $q(\mathcal{B})$ remains very close to $q(\mathcal{A})$, then most of the densities d found across parts of $\mathcal{B}$ are very close to those found in $\mathcal{A}$.

Lemma A.6. *Suppose equipartitions $\mathcal{A} = \{V_i : i \in [t]\}$ and refinement $\mathcal{B} = \{V_{i,j} : i \in [k], j \in [\ell]\}$ have $q(\mathcal{B}) - q(\mathcal{A}) \leqslant \varepsilon^4/64$. Then for all but εt^3 $(i_1, i_2, i_3) \in [t]^3$, all but $\varepsilon \ell^3$ $(j_1, j_2, j_3) \in [\ell]^3$ have $|d(V_{i_1}, V_{i_2}, V_{i_3}) - d(V_{i_1,j_1}, V_{i_2,j_2}, V_{i_3,j_3})| \leqslant \varepsilon$.*

Proof. First fix $(i_1, i_2, i_3) \in [t]^3$ and let Z be the random variable defined in Lemma A.3, over the refinement given by $\mathcal{B}$ of the sets $V_{i_1}, V_{i_2}, V_{i_3}$. If $L^{(i_1,i_2,i_3)} \subseteq [\ell]^3$ are those $(j_1, j_2, j_3) \in [\ell]^3$ with $|d(V_{i_1}, V_{i_2}, V_{i_3}) - d(V_{i_1,j_1}, V_{i_2,j_2}, V_{i_3,j_3})| \geqslant \varepsilon$, and $|L^{(i_1,i_2,i_3)}| \geqslant \varepsilon \ell^3$, then

$$\varepsilon \ell^3 \left(\frac{1}{2\ell} \right)^3 \leqslant \sum_{(j_1,j_2,j_3) \in L^{(i_1,i_2,i_3)}} \frac{|V_{i_1,j_1}||V_{i_2,j_3}||V_{i_3,j_3}|}{|V_{i_1}||V_{i_2}||V_{i_3}|} \leqslant \Pr(|Z - \mathbb{E}(Z)| \geqslant \varepsilon) \leqslant \frac{\mathrm{Var}(Z)}{\varepsilon^2}$$

and therefore, recalling we calculated $\mathrm{Var}(Z)$ in Lemma A.3,

$$q\Big(\dot\bigcup_{j_1 \in [\ell]} V_{i_1,j_1}, \dot\bigcup_{j_2 \in [\ell]} V_{i_2,j_2}, \dot\bigcup_{j_3 \in [\ell]} V_{i_3,j_3} \Big) - q(V_{i_1}, V_{i_2}, V_{i_3}) \geqslant \frac{\varepsilon^3}{8} \frac{|V_{i_1}||V_{i_2}||V_{i_3}|}{n^3} \,.$$

Next define $I \subseteq [t]^3$ as those $(i_1, i_2, i_3) \in t^3$ for which $|L^{(i_1,i_2,i_3)}| \geqslant \varepsilon \ell^3$ and suppose for contradiction that $|I| > \varepsilon t^3$. Then

$$\begin{aligned} q(\mathcal{B}) - q(\mathcal{A}) &\geqslant \sum_{(i_1,i_2,i_3) \in I} q\Big(\dot\bigcup_{j_1 \in [\ell]} V_{i_1,j_1}, \dot\bigcup_{j_2 \in [\ell]} V_{i_2,j_2}, \dot\bigcup_{j_3 \in [\ell]} V_{i_3,j_3} \Big) - q(V_{i_1}, V_{i_2}, V_{i_3}) \\ &> \varepsilon t^3 \frac{\varepsilon^3}{8} \frac{|V_{i_1}||V_{i_2}||V_{i_3}|}{n^3} \geqslant \frac{\varepsilon^4}{64} \end{aligned}$$

a contradiction. □

We can now iterate 3.7 to prove Theorem A.1.

Proof of Theorem A.1. Given any $m \in \mathbb{N}$ and $\varepsilon > 0$, let $M_{3.7}(m, \varepsilon)$ and $N_{3.7}(m, \varepsilon)$ denote the output of Theorem 3.7. Let $\mathcal{E}$ and m be as in the Theorem statement, and let $\varepsilon = \mathcal{E}(0)$. Set $M_0 = M_{3.7}(m, \varepsilon)$ and $N_0 = N_{3.7}(m, \varepsilon)$ before inductively defining

$$M_i = M_{3.7}\left(M_{i-1}, \frac{\mathcal{E}(M_{i-1})}{M_{i-1}^3}\right)$$
$$N_i = N_{3.7}\left(M_{i-1}, \frac{\mathcal{E}(M_{i-1})}{M_{i-1}^3}\right).$$

If we set $s = \left\lceil \frac{64}{\varepsilon^4} \right\rceil + 1$, then we claim that $M = M_s, N = N_s$ suffices.

First let $\mathcal{A}^0$ be an ε-regular equipartition of order $t_0 \in [m, M_0]$, provided by Theorem 3.7 and then perform the following iterative procedure, starting with $i = 0$ (where $t_{-1} = 0$ for convenience). Given $\mathcal{A}^i$ with order $t_i \in [t_{i-1}, M_i]$ we may apply Theorem 3.7 again to obtain a refinement $\mathcal{A}^{i+1}$ which is $\frac{\mathcal{E}(M_i)}{M_i^3}$-regular with order $t_{i+1} \in [t_i, M_{i+1}]$. In particular, at most $\frac{\mathcal{E}(M_{i-1})}{M_{i-1}^3} t_i^3$ of $(i_1, i_2, i_3) \in [t_i]^3$ have $(V_{i_1}, V_{i_2}, V_{i_3})$ fail to be $\mathcal{E}(M_{i-1})$-regular, since $\frac{\mathcal{E}(M_{i-1})}{M_{i-1}^3} \leqslant \mathcal{E}(M_{i-1})$. Let i be the first i so that $q(\mathcal{A}^i) - q(\mathcal{A}^{i-1}) \leqslant \frac{\varepsilon^4}{64}$, and set $\mathcal{A} = \mathcal{A}_{i-1}$ and $\mathcal{B} = \mathcal{A}_i$, with $t = t_{i-1}$ and $t\ell = t_i$. It remains to check items *(1)*-*(4)*.

Part *(1)* follows immediately by recalling that the t_i are increasing, so $m \leqslant t_0 \leqslant t \leqslant t\ell \leqslant M$. Part *(2)* also follows immediately using the monotonicity of $\mathcal{E}$, since all but $(\frac{\mathcal{E}(M_{i-2})}{M_{i-2}^3})[t_{i-1}]^3 \leqslant \mathcal{E}(0)[t_{i-1}]^3$ of $(i_1, i_2, i_3) \in [t_{i-1}]^3$ have $(V_{i_1} V_{i_2}, V_{i_3})$ as $\mathcal{E}(0)$-regular. For Part *(3)*, we have that the partition $\mathcal{B}$ has at most $\frac{\mathcal{E}(M_{i-1})}{M_{i-1}^3}(t\ell)^3$ pairs $((i_1, i_2, i_3), (j_1, j_2, j_3)) \in ([t]^3, [\ell]^3)$ with $(V_{i_1,j_1}, V_{i_2,j_2}, V_{i_3,j_3})$ failing to be $\mathcal{E}(M_{i-1})$-regular. Since $\frac{\mathcal{E}(M_{i-1})}{M_{i-1}^3}(t\ell)^3 \leqslant \mathcal{E}(M_{i-1})\ell^3$, there is certainly no $(i_1, i_2, i_3) \in [t]^3$ with more than $\mathcal{E}(t)\ell^3$ many $(j_1, j_2, j_3) \in [\ell]^3$ with $(V_{i_1,j_1}, V_{i_2,j_2}, V_{i_3,j_3})$ failing to be $\mathcal{E}(M_t)$-regular. Finally, Part *(4)* follows by direct application of Lemma A.6.

□

Finally we may sample from the refinement $\mathcal{B}$ of $\mathcal{A}$ to find the model vertex sets we require.

Proof of Corollary 3.8. We apply Theorem A.1 with $\mathcal{E}'(r) = \min\{\mathcal{E}(r), \frac{1}{4t^3}, \frac{\varepsilon}{8}\}$ and m to obtain $M_{A.1}$ and $N_{A.1}$ and claim that $M = M_{A.1}$, $N = N_{A.1}$, and $\delta = \frac{1}{2M}$ suffices. To that end, let Q be any palette with at least N colors, so that Theorem A.1 gives $\mathcal{B}$ refining $\mathcal{A}$ so that Parts *(1)*-*(4)* hold. For each $i \in [t]$ select $j_i \in [\ell]$, independently uniformly at random, and set $U_i = V_{i,j_i}$. We now show that parts *(i)*-*(iv)* are satisfied with positive probability.

Parts *(i)* and *(ii)* are immediate and always hold. For Part *(iii)*,

$$\mathbb{P}\left((U_{i_1},U_{i_2},U_{i_3}) \text{ is not } \varepsilon\text{-regular for some } (i_1,i_2,i_3)\in[t]^3\right) \leqslant t^3\mathcal{E}'(t) \leqslant \frac{1}{4}$$

by applying the union bound and Part *(3)* of Theorem A.1. Meanwhile,

$$\mathbb{E}(|(i_1,i_2,i_3)\in[t]^3 \text{ with } |d(U_{i_1},U_{i_2},U_{i_3})-d(V_{i_1},V_{i_2},V_{i_3})|\geqslant\varepsilon|) \leqslant \frac{\varepsilon}{8}t^3+\frac{\varepsilon}{8}t^3$$

by Part *(4)* of A.1, and therefore the probability that there are more than εt^3 $(i_1,i_2,i_3)\in[t]^3$ with $|d(U_{i_1},U_{i_2},U_{i_3})-d(V_{i_1},V_{i_2},V_{i_3})|\geqslant\varepsilon$ cannot exceed $\frac{1}{4}$. Then with probability at least $\frac{1}{2}$ both Part *(iii)* and *(iv)* are satisfied as well, so there exists such a choice of U_i and we are done. □

Mathematics Department, California Institute of Technology, Pasadena, USA
Email address: `dking@caltech.edu`

Fachbereich Mathematik, Universität Hamburg, Hamburg, Germany
Email address: `simon.piga@uni-hamburg.de`

Mathematics Department, University of California Irvine, Irvine, USA
Email address: `mtsales@uci.edu`

Extremal Combinatorics and Probability Group, Institute for Basic Science, Daejeon, South Korea
Email address: `schuelke@ibs.re.kr`