

Exact Ramsey numbers of odd cycles via nonlinear optimisation

Jenssen, Matthew; Skokan, Jozef

DOI:

[10.1016/j.aim.2020.107444](https://doi.org/10.1016/j.aim.2020.107444)

License:

Creative Commons: Attribution-NonCommercial-NoDerivs (CC BY-NC-ND)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Jenssen, M & Skokan, J 2021, 'Exact Ramsey numbers of odd cycles via nonlinear optimisation', *Advances in Mathematics*, vol. 376, 107444. <https://doi.org/10.1016/j.aim.2020.107444>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

EXACT RAMSEY NUMBERS OF ODD CYCLES VIA NONLINEAR OPTIMISATION

MATTHEW JENSSEN[†], JOZEF SKOKAN[‡]

ABSTRACT. For a graph G , the k -colour Ramsey number $R_k(G)$ is the least integer N such that every k -colouring of the edges of the complete graph K_N contains a monochromatic copy of G . Let C_n denote the cycle on n vertices. We show that for fixed $k \geq 2$ and n odd and sufficiently large,

$$R_k(C_n) = 2^{k-1}(n-1) + 1.$$

This resolves a conjecture of Bondy and Erdős for large n . The proof is analytic in nature, the first step of which is to use the regularity method to relate this problem in Ramsey theory to one in nonlinear optimisation. This allows us to prove a stability-type generalisation of the above and establish a correspondence between extremal k -colourings for this problem and perfect matchings in the k -dimensional hypercube Q_k .

1. INTRODUCTION

One of the most well-known and extensively researched problems in combinatorics is that of determining the *Ramsey numbers* of graphs, defined as follows. Given graphs G_1, G_2, \dots, G_k , the *Ramsey number* $R(G_1, \dots, G_k)$ is the least integer N such that any colouring of the edges of the complete graph K_N on N vertices with k colours contains a monochromatic copy of G_i in the i -th colour for some i , $1 \leq i \leq k$. In the case where G_1, \dots, G_k are all isomorphic to the graph G , we call $R(G_1, \dots, G_k)$ the *k -colour Ramsey number* of G and denote it by $R_k(G)$. Broadly speaking the philosophy underpinning Ramsey theory is that large, potentially highly disordered, structures must contain ordered substructures.

Ramsey theory owes its name to the seminal paper of Frank Ramsey [Ram30] where it is shown that Ramsey numbers are finite. Some of the earliest and most famous examples of Ramsey numbers are those involving cliques. The systematic study of such Ramsey numbers began with a paper of Erdős and Szekeres [ES35] who considered the problem of determining $R_2(K_k)$, where K_k denotes the clique on k vertices. Erdős and Szekeres [ES35] and Erdős [Erd47] showed that $2^{k/2} \leq R_2(K_k) \leq 4^k$. The problem of improving these bounds has gained significant notoriety. Despite considerable effort over the past 80 years, the bases in the exponent in both of these bounds has not been improved.

This inertia has motivated the study of Ramsey numbers of graphs with a ‘simpler’ structure, where the problem may be more tractable. In this spirit, there has been a large body of research dedicated to the study of Ramsey numbers of graphs that are *sparse* in some sense

Date: October 5, 2020.

[†]School of Mathematics, University of Birmingham, Edgbaston, Birmingham B15 2TT, U.K.
E-mail: m.jenssen@bham.ac.uk.

[‡]Department of Mathematics, London School of Economics, Houghton Street, London WC2A 2AE, U.K. and Department of Mathematics, University of Illinois at Urbana-Champaign, 1409 W. Green Street, Urbana, IL 61801, USA *E-mail:* j.skokan@lse.ac.uk.

(e.g. they have bounded maximum degree). The cycle on n vertices C_n is a particularly simple example and was one of the earliest subjects in the study of Ramsey numbers of sparse graphs. The behaviour of the Ramsey number $R(C_{n_1}, C_{n_2})$ has been studied by several authors, including Bondy and Erdős [BE73], Faudree and Schelp [FS74], and Rosta [Ros73], and is now fully determined. For example it is known that

$$R_2(C_n) = \begin{cases} 2n - 1, & \text{if } n \geq 5 \text{ is odd,} \\ \frac{3n}{2} - 1, & \text{if } n \geq 6 \text{ is even.} \end{cases}$$

Results such as this one that exactly determine $R(G_1, G_2)$ for a pair of graphs G_1, G_2 are by now fairly plentiful. See Radziszowski [Rad94] for an excellent survey of such results.

However, in the case where more than two colours are involved such results are still extremely rare. The only non-trivial class of graphs for which the k -colour Ramsey number is exactly determined for arbitrary k is that of matchings (a result due to Cockayne and Lorimer [CL75]). In this paper we address the following conjecture attributed to Bondy and Erdős [BE73].

Conjecture 1.1. *If $k \geq 2$ and $n > 3$ is odd then*

$$R_k(C_n) = 2^{k-1}(n-1) + 1.$$

Note that the conjecture deals specifically with the case where n is odd. Odd and even cycles behave rather differently in this context due to the fact that an even cycle is bipartite whereas an odd cycle is not. We mention that Erdős and Graham [EG75] proved the bounds

$$(1.1) \quad 2^{k-1}(n-1) + 1 \leq R_k(C_n) \leq (k+2)!n,$$

for all $k \geq 2$ and all odd $n > 3$. In this paper we prove the following.

Theorem 1.2. *For any fixed $k \geq 2$ and odd n sufficiently large,*

$$R_k(C_n) = 2^{k-1}(n-1) + 1.$$

We therefore resolve Conjecture 1.1 for large n . We will in fact prove a stability-type strengthening of this result (see Theorem 3.2 below). Recently Day and Johnson [DJ17] showed that in the opposite regime, where we fix an odd n and let k be sufficiently large, one in fact has $R_k(C_n) > (n-1)(2+\varepsilon)^{k-1}$ for some $\varepsilon = \varepsilon(n) > 0$, and so Conjecture 1.1 is false when n is small with respect to k . The qualification that n is sufficiently large in Theorem 1.2 is therefore necessary, however due to the use of compactness arguments in the proof, we obtain no effective bound on how large n must be with respect to k .

In view of Theorem 1.2, let us call a k -colouring of the complete graph on $2^{k-1}(n-1)$ vertices which does not contain a monochromatic copy of C_n an *extremal k -colouring*. The lower bound in (1.1) was established by observing that one can naturally construct extremal k -colourings by induction. Indeed if there exists a k -colouring of the edges of the complete graph K_m with no monochromatic C_n , then by joining two such copies of K_m by edges of colour $k+1$, one obtains a $(k+1)$ -colouring of K_{2m} with no monochromatic C_n (here we use that C_n is non-bipartite). The base construction, for $k=1$, is simply a monochromatic clique of size $n-1$. It was believed that all extremal k -colourings come from such a doubling argument. We show that this is not the case, providing a classification of extremal k -colourings which exposes a surprising correspondence between extremal k -colourings and perfect matchings in the k -dimensional hypercube Q_k .

The first breakthrough towards Conjecture 1.1 was made by Łuczak [Luc99] who used the regularity method to show that the $k = 3$ case holds asymptotically i.e. that for n odd,

$$R_3(C_n) = 4n + o(n) \text{ as } n \rightarrow \infty.$$

Łuczak’s method of applying regularity in this setting has proven extremely fruitful (see e.g. [FL07a, GRSS07, KSS05, Luc99, ŁSS12]) and has since become a standard tool. We will come to describe the method in more detail as it provides the starting point for the present paper.

Building on Łuczak’s ideas, Kohayakawa, Simonovits and Skokan [KSS05] paired the regularity method with stability arguments to resolve Conjecture 1.1 for $k = 3$ and n large. The case where $k \geq 4$ remained open. Progress was made by Łuczak, Simonovits and Skokan [ŁSS12] who showed that for $k \geq 4$ and odd n ,

$$R_k(C_n) \leq k2^k n + o(n) \text{ as } n \rightarrow \infty.$$

We conclude this section with a broad overview of the proof method of Theorem 1.2. Let $\mathcal{G}_{n,k}$ denote the (finite) class of all cliques that admit a k -colouring with no monochromatic copy of C_n . Determining $R_k(C_n)$ is then equivalent to determining the maximum N such that $K_N \in \mathcal{G}_{n,k}$. Using the regularity method, we relate this problem to finding the maximum ℓ_1 -norm of an element in a certain compact subset \mathcal{S} of \mathbb{R}^{3k} . This allows us to import analytic and topological tools in support of our proof.

The relation is such that maximal elements of \mathcal{S} correspond to extremal k -colourings for Theorem 1.2. Therefore, by classifying the extremal points of \mathcal{S} we can classify the extremal k -colourings and prove a stability type strengthening of Theorem 1.2 (Theorem 3.2 below). Indeed we show that each perfect matching of the hypercube Q_k gives rise to a class of extremal k -colourings. Moreover, any extremal k -colouring must be ‘close’ to one such construction. We defer precise statements to Section 3.

The number of essentially different classes of extremal k -colourings is equal to the number of equivalence classes of perfect matchings in Q_k with respect to its automorphism group and this number is doubly exponential in k . Such a plethora of extremal constructions is usually forbidding when trying to establish stability type results in combinatorics. We believe that the fact we can overcome this obstacle is largely down to our analytic perspective.

2. NOTATION AND TERMINOLOGY

Let us collect some notation that we use throughout the paper. For $k \in \mathbb{N}$, we let $[k]$ denote the set $\{1, \dots, k\}$. For a set S , we let $\binom{S}{2}$ denote the set of all unordered pairs of distinct elements of S .

All graphs considered here will be finite. For a graph $G = (V, E)$, we let $v(G) = |V|$ and $e(G) = |E|$. For $v \in V$, we let $N_G(v)$ denote the neighbourhood of v in G and let $\delta(G)$ denote the minimum degree of G . For $X \subseteq V$ we denote by $G[X]$ the subgraph of G induced by the vertices of X . For disjoint subsets $A, B \subseteq V$, we denote by $G[A, B]$ the bipartite graph with vertex set $A \cup B$ and edge set $\{\{a, b\} \in E : a \in A, b \in B\}$, and we let $e_G(A, B)$ denote the size of this set. In the case where $A = \{v\}$, we write $G[v, B]$ instead of $G[\{v\}, B]$. All subscripts in the above notation may be suppressed if they are clear from the context. At times we may slightly abuse notation by writing $v \in G$ and $\{x, y\} \in G$ in lieu of $v \in V(G)$ and $\{x, y\} \in E(G)$ respectively.

Let $W = w_0 w_1 \dots w_\ell$ be a walk in G (that is a sequence of vertices w_0, w_1, \dots, w_ℓ such that $\{w_i, w_{i+1}\}$ is an edge of G for all $0 \leq i < \ell$). If all of the w_i are distinct then we call W a *path of length ℓ* . We may also refer to W as a $w_0 w_\ell$ -path to distinguish its endpoints. If all the w_i are distinct except $w_0 = w_\ell$ then we call W a *cycle of length ℓ* . We will also concatenate walks in the natural way. For example if $U = u_0 \dots u_m$ is a walk in G such that $u_m = w_0$, we let UW denote the walk $u_0 \dots u_m w_1 \dots w_\ell$. If x is a vertex such that $\{w_\ell, x\}$ is an edge of G then we let Wx denote the walk $w_0 \dots w_\ell x$.

A *k -coloured graph* is a graph $G = (V, E)$ equipped with some function $\varphi : E \rightarrow [k]$. Furthermore, for $i \in [k]$, we let G_i denote the subgraph $(V, \varphi^{-1}\{i\})$ of G . We call G_i the *i th colour class* of G .

A *digraph* $D = (V, A)$ consists of a set of vertices V and a set $A \subseteq V^2$ i.e. a set of ordered pairs from V which we call *directed edges*. For $v \in V$ we let $d^+(v)$ denote the size of the set $\{u : (v, u) \in A\}$ and call $d^+(v)$ the *outdegree* of v . Similarly we define the *indegree* of v as $d^-(v) = |\{u : (u, v) \in A\}|$.

For $x \in \mathbb{R}^d$ we let $\|x\|$ denote the ℓ_1 -norm of x i.e. $\|x\| = \sum_{i=1}^d |x_i|$. Furthermore, given $\varepsilon > 0$, we let $B_\varepsilon(x) := \{z \in \mathbb{R}^d : \|z - x\| < \varepsilon\}$, the open ball of radius ε centred at x . We let $\text{supp}(x)$ denote the support of x .

In the statements of theorems and lemmas it will be useful to use the notation $\alpha \ll \beta$ to mean that there is an increasing function $\alpha(x)$ so that the statement is valid for $0 < \alpha < \alpha(\beta)$. When we need to refer to this function at a later stage, we include the number of the lemma (or theorem) the function appears in as a subscript. For example, $\delta_{8.3}(x)$ denotes the implied function $\delta(x)$ from Lemma 8.3. Throughout the paper we omit the use of floor and ceiling symbols where they are not crucial.

3. A GRAPH DECOMPOSITION, EXTREMAL COLOURINGS AND STABILITY

In this section we describe the extremal colourings and give precise statements of the stability results referred to in the introduction. We also introduce some key concepts and results that will be used throughout the paper and give a more detailed overview of our proof methods. We begin by introducing a way of decomposing an arbitrary k -coloured graph. This decomposition will play a central role for us and is similar to a decomposition introduced in [LSS12].

3.1. A Graph Decomposition. Let G be a k -coloured graph. For each $i \in [k]$, we write $G_i = G'_i \cup G''_i$, where G'_i is the union of the bipartite components of G_i and G''_i is the union of the non-bipartite components of G_i . For each $i \in [k]$, write $V(G'_i) = V_0^i \cup V_1^i$ where V_0^i and V_1^i are the vertex classes of a bipartition of G'_i and set $V_*^i = V(G''_i)$. For $\tau \in \{0, 1, *\}^k$, let $V_\tau = \bigcap_{j=1}^k V_{\tau_j}^j$ and note that

$$V(G) = \bigcup_{\tau \in \{0, 1, *\}^k} V_\tau,$$

a disjoint union. We call $(V_\tau : \tau \in \{0, 1, *\}^k)$ a *profile partition* of G and we call the corresponding vector $(|V_\tau| : \tau \in \{0, 1, *\}^k)$ a *profile* of G . We will often denote a profile of G by $x(G)$. Note that G may admit multiple profile partitions since we made an arbitrary choice in choosing the bipartition $V(G'_i) = V_0^i \cup V_1^i$ for each $i \in [k]$.

3.2. Extremal Colourings and the Hypercube. For $k \in \mathbb{N}$, we let Q_k denote the k -dimensional hypercube i.e. the graph on vertex set $\{0, 1\}^k$ and edge set consisting of pairs differing in exactly one coordinate. It will be useful to think of an element $\tau \in \{0, 1, *\}^k$ as a subcube of the k -dimensional hypercube Q_k via the correspondence

$$\tau \leftrightarrow Q(\tau) := \{c \in \{0, 1\}^k : c_j = \tau_j \text{ if } \tau_j \in \{0, 1\}\}.$$

In other words we think of a coordinate j where $\tau_j = *$ as a ‘missing bit’ and let $Q(\tau)$ be the set of all possible ways of filling in these bits. For example, if $k = 3$ and $\tau = (0, *, *)$, then

$$Q(\tau) = \{(0, 0, 0), (0, 0, 1), (0, 1, 0), (0, 1, 1)\}.$$

We define the *weight* of τ to be the size of the set $\{i \in [k] : \tau_i = *\}$ (i.e. the number of missing bits) and denote it by $\omega(\tau)$. Note that $|Q(\tau)| = 2^{\omega(\tau)}$. In the language of the hypercube, $\omega(\tau)$ is the dimension of the subcube $Q(\tau)$. In particular if $\omega(\tau) = 1$, then we think of $Q(\tau)$ as an *edge* of Q_k .

We can now describe a class of extremal k -colourings in terms of perfect matchings in Q_k . Let \mathcal{M} be a perfect matching of Q_k . We identify each edge $e \in \mathcal{M}$ with the element $\tau \in \{0, 1, *\}^k$ for which $e = Q(\tau)$. Let $G = K_N$ where $N = 2^{k-1}(n-1)$ and let $V(G) = \bigcup_{\tau \in \mathcal{M}} V_\tau$ be a partition of $V(G)$ where $|V_\tau| = n-1$ for all $\tau \in \mathcal{M}$. For each $\tau \in \mathcal{M}$, colour all edges in $G[V_\tau]$ with the colour i , where i is the coordinate for which $\tau_i = *$. For $\tau, \sigma \in \mathcal{M}$, arbitrarily colour the edges between V_τ and V_σ with any colour j for which $\{\sigma_j, \tau_j\} = \{0, 1\}$ (i.e. edges τ, σ lie in opposite subcubes of Q_k of codimension 1 separated by the j th coordinate). It follows that each colour class of such a colouring is the disjoint union of cliques of size $n-1$ and a bipartite graph and therefore contains no monochromatic copy of C_n . We call such a colouring a *hypercube k -colouring with clique size $n-1$* . We emphasise that such a colouring assigns a unique colour to each edge of a clique on $2^{k-1}(n-1)$ vertices and so is indeed an extremal k -colouring (see Figure 1 for a depiction of two examples in the case $k = 3$).

Suppose now that we inductively construct a perfect matching on Q_k by taking two perfect matchings on a disjoint pair of subcubes of codimension 1. The class of hypercube k -colourings associated to such a matching includes the inductive colourings of Erdős and Graham [EG75] described in the introduction (and many more). However for $k \geq 4$, not all perfect matchings of Q_k decompose as the union of two matchings on a pair of codimension 1 subcubes, and so we obtain extremal k -colourings with a genuinely new structure. In particular, a novel feature that appears for $k \geq 4$ colours is that there exist extremal k -colourings that contain monochromatic cliques of size $n-1$ in all k possible colours.

3.3. Stability. In this subsection we state a theorem to the effect that the hypercube k -colourings considered in the previous subsection are the only extremal k -colourings for our problem. Moreover we assert that almost extremal colourings are in some sense ‘close’ to a hypercube colouring. Let us make this more precise.

Definition 3.1. Let G and H be k -coloured graphs with $V(H) \subseteq V(G)$. Let $\varepsilon > 0$, then we say that G is ε -close to H if $|G_i \Delta H_i| \leq \varepsilon v(G)^2$ for all $i \in [k]$.

Informally we may say that G and H as above are close in *edit distance*. We may now state the main result of this paper.

Theorem 3.2. *Let $k \geq 2$, let $\frac{1}{n} \ll \eta \ll \varepsilon \ll 1$, where n is odd, and let $N > (2^{k-1} - \eta)n$. Then if $G = K_N$ is k -coloured with no monochromatic copy of C_n , then $N \leq 2^{k-1}(n-1)$ and there exists a hypercube k -colouring H such that G is ε -close to H .*

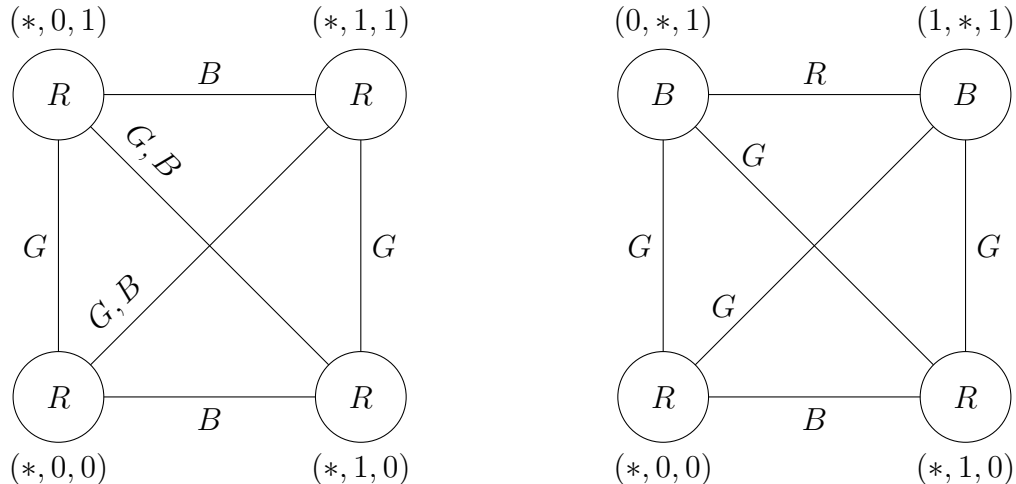


FIGURE 1. Two examples of an extremal k -colouring with $k = 3$. The two constructions correspond to two distinct perfect matchings in Q_3 , one representative from each of the two isomorphism classes of perfect matching in Q_3 . We use R, B, G to denote the colours (in place of 1, 2, 3 respectively). Each circle represents a clique of size $n - 1$ and is labelled by an edge of Q_3 . The letter at the centre of each clique represents the colour of the clique. A line joining two cliques is labelled by the colours that can be (arbitrarily) assigned to the edges joining those two cliques.

Note that Theorem 1.2 follows as an immediate corollary. The $k = 3$ case of Theorem 3.2 was proved in [KSS05] where the two classes of colourings the authors consider can be viewed as the colourings that arise from the two isomorphism classes of perfect matchings in Q_3 (see Figure 1). An interesting feature of Theorem 3.2 is that it deals with a wide variety of extremal colourings. Indeed if \mathcal{M}_1 and \mathcal{M}_2 are perfect matchings of Q_k that lie in distinct equivalence classes under the action of the automorphism group of Q_k , then it's not difficult to show that there are hypercube k -colourings associated to \mathcal{M}_1 that are not isomorphic to any hypercube k -colouring associated to \mathcal{M}_2 . It is also interesting to note that even though we can prove a stability statement around hypercube k -colourings, the structure of these colourings is not well understood. This is simply due to the fact that the structure of perfect matchings in the hypercube is not well understood. Indeed, even enumerating the perfect matchings (or their equivalence classes) in Q_k is a well-studied and difficult problem. Let $f(k)$ be the number of equivalence classes of perfect matchings in Q_k . It is clear that $f(3) = 2$ and so we obtain two essentially different extremal 3-colourings as in [KSS05]. Graham and Harary [GH88] showed that $f(4) = 8$ and recently Östergård and Pettersson [ÖP13] determined (with a large amount of computer time) $f(5)$, $f(6)$ and $f(7)$. The function $f(k)$ grows rather rapidly; it is amusing to note that already $f(7) = 607158046495120886820621$ and so we have this many essentially different classes of extremal 7-colourings. It was shown in [CGP97] that the number of perfect matchings in Q_k is $[(1 + o(1))k/e]^{2^{k-1}}$ (although this result in fact follows from a theorem in [PL86, p.312]). Since the automorphism group of Q_k has size $k!2^k$ it follows that $f(k) = [(1 + o(1))k/e]^{2^{k-1}}$ also.

3.4. Proof of Theorem 3.2: An Overview. The regularity method, discussed briefly in the introduction, plays a central role in our proof. We include an informal discussion of the method here, deferring details until later. We start with a definition.

Definition 3.3. Let F be a connected graph whose largest matching covers m vertices, then we call F a *connected matching of order m* . We distinguish a particular matching of largest size M_F in F and refer to an edge of M_F as a *matching edge* of F . If in addition F is non-bipartite, we call F an *odd connected matching of order m* .

The idea behind the regularity method is as follows. Suppose that G is a k -coloured complete graph on N vertices. Let G_1, \dots, G_k be its colour classes. We apply the multicolour version of the Regularity Lemma [Sze78] and obtain a regular partition of the vertex set $V(G)$ into $t + 1$ classes $V(G) = V_0 \cup \dots \cup V_t$. We construct an auxiliary graph R with vertex set $1, \dots, t$ and the edge set formed by pairs $\{i, j\}$ for which (V_i, V_j) is regular with respect to G_1, \dots, G_k . We colour each edge $\{i, j\}$ in R by the majority colour in the pair (V_i, V_j) . The crucial point is that if R contains a monochromatic odd connected matching of order greater than m , then G contains a monochromatic cycle C_ℓ where ℓ can take essentially any odd value smaller than mN/t . It follows that if G contains no monochromatic copy of C_n , then R cannot contain a monochromatic odd connected matching of order larger than nt/N . The advantage of this perspective is that forbidding a large connected matching is far more restrictive than forbidding a cycle of a given length. Indeed a cycle is itself an example of a connected matching, and so if a graph contains no connected matching of order greater than m then it contains no cycle of length greater than m . The following theorem of Erdős and Gallai [EG59] shows that this is a very strict condition.

Theorem 3.4. *Let $m \geq 3$. If G is a graph which contains no cycle of length greater than m , then $e(G) \leq m(v(G) - 1)/2$.*

The price one pays is that R is not a complete graph, however it can be chosen to be as dense as one likes. We are now able to state a theorem that is a major stepping stone toward the proof of Theorem 3.2.

Theorem 3.5. *Let $k \geq 2$ and let $\frac{1}{n} \ll \delta \ll \varepsilon \ll 1$. If G is a k -coloured graph with $v(G) = 2^{k-1}n$ and $e(G) \geq (1 - \delta) \binom{v(G)}{2}$ containing no monochromatic odd connected matching of order $\geq (1 + \delta)n$, then for any choice of profile $x(G)$ of G , there exists a hypercube k -colouring H with profile $x(H)$ satisfying*

$$\|x(G) - x(H)\| \leq \varepsilon n.$$

The proof of Theorem 3.5 occupies the majority of this paper. In the final section we show how Theorem 3.2 follows from Theorem 3.5 via combinatorial stability arguments and the regularity method.

The outline of the proof of Theorem 3.5 is as follows. Let G be as in the statement of Theorem 3.5 and let $x(G)$ denote a profile of G . Our starting point is to translate the combinatorial constraint of containing no large monochromatic odd connected matching into an analytic constraint on $x(G)$ of the form

$$(3.1) \quad F(x(G)) \leq 0,$$

where F is a quadratic form which we derive in the next section. We then view (3.1) as a constraint in an optimisation problem where we wish to maximise the objective function $\|x(G)\|$.

Recalling that $\|x(G)\| = v(G)$, we get a corresponding upper bound on the order of G . It turns out that optimal solutions to this optimisation problem correspond to the profiles of hypercube colourings. Solving the optimisation problem is the subject of Sections 5 and 6. In Section 7 we use compactness arguments to show that almost optimal solutions must be close in ℓ_1 -norm to the profile of a hypercube colouring. We then translate this analytic stability into the, more combinatorial, stability statement of Theorem 3.5.

Note that Theorem 3.5 will be applied to a reduced graph like the one described above. The focus of the final section is to show that if the profile of this reduced graph is close in ℓ_1 -norm to the profile of a hypercube colouring, then the original graph is close in edit distance to a hypercube colouring. In particular, we will show that the original graph can be covered by 2^{k-1} monochromatic cliques of equal size and a small ‘leftover set’ (Lemma 9.9). We then establish an ‘absorbing lemma’ (Lemma 9.12) that shows that for each vertex v of the leftover set, there is a monochromatic clique C , in colour i say, such that all of the edges between v and C have colour i . Informally, we think of the clique C ‘absorbing’ the vertex v . It follows that if the original graph has $> 2^{k-1}(n-1)$ vertices, one of the cliques and its absorbed vertices has size $\geq n$ and is easily seen to contain a monochromatic copy of C_n . Thus we obtain the exact value of the Ramsey number $R_k(C_n)$.

4. DERIVING THE ANALYTIC CONSTRAINTS

Given a k -coloured graph G , we will show how to translate the combinatorial constraint of containing no large monochromatic odd connected matching into an analytic constraint on the profile of G .

From here on, throughout the paper, we let $k \geq 2$ be a fixed integer. Let G be a k -coloured graph. First we distinguish between two types of edges of G . If $e \in E(G)$ is coloured with the colour j and lies in a bipartite component of G_j then we call e a *bipartite* edge. We call e *non-bipartite* otherwise. Let $(V_\tau : \tau \in \{0, 1, *\}^k)$ be a profile partition of G . We make two simple observations regarding the profile partition of a k -coloured graph.

Observation 4.1. If $e \in E(G)$ is a bipartite edge of colour j then it must have endpoints in parts V_τ, V_σ for some $\tau, \sigma \in \{0, 1, *\}^k$ such that $\tau_j = 0$ and $\sigma_j = 1$.

Observation 4.2. If $e \in E(G)$ is a non-bipartite edge of colour j then it must have endpoints in parts V_τ, V_σ for some (not necessarily distinct) $\tau, \sigma \in \{0, 1, *\}^k$ such that $\tau_j = \sigma_j = *$.

This motivates the following definitions.

Definition 4.3. We say that $\sigma, \tau \in \{0, 1, *\}^k$ are *distinguishable* if $\{\sigma_j, \tau_j\} = \{0, 1\}$ for some $j \in [k]$. We say that σ and τ are *indistinguishable* otherwise.

Definition 4.4. If $\sigma, \tau \in \{0, 1, *\}^k$ are such that either (i) σ, τ are distinguishable or (ii) $\sigma_j = \tau_j = *$ for some $j \in [k]$, then we say that σ and τ are *compatible*. We say that σ, τ are *incompatible* otherwise.

Viewing elements of $\{0, 1, *\}^k$ as subcubes of Q_k , we may reinterpret these definitions as follows.

Lemma 4.5. *Let $\sigma, \tau \in \{0, 1, *\}^k$. Then σ, τ are distinguishable if and only if $Q(\tau) \cap Q(\sigma) = \emptyset$. Furthermore, σ, τ are incompatible if and only if $|Q(\tau) \cap Q(\sigma)| = 1$.*

Proof. By the definition of the sets $Q(\tau), Q(\sigma)$ we have

$$Q(\tau) \cap Q(\sigma) = \{c \in \{0, 1\}^k : c_j = \tau_j \text{ if } \tau_j \in \{0, 1\} \text{ and } c_j = \sigma_j \text{ if } \sigma_j \in \{0, 1\}\}.$$

This is empty if and only if there exists a $j \in [k]$ such that $\sigma_j, \tau_j \in \{0, 1\}$ and $\sigma_j \neq \tau_j$ i.e. if and only if σ, τ are distinguishable. Let $T = \{i \in [k] : \sigma_i = \tau_i = *\}$. If σ, τ are indistinguishable then we see that $|Q(\tau) \cap Q(\sigma)| = 2^{|T|}$. Therefore $|Q(\tau) \cap Q(\sigma)| = 1$ if and only if σ, τ are indistinguishable and $T = \emptyset$ i.e. σ, τ are incompatible. \square

From now on, we let

$$D = \{\{\sigma, \tau\} : \sigma, \tau \in \{0, 1, *\}^k \text{ are distinguishable}\}.$$

It will also be convenient to make the following definition.

Definition 4.6. Let $\alpha > 0$ and let G be a graph such that $e(G) \geq \alpha \binom{v(G)}{2}$. Then we say that G is α -dense.

This next proposition provides the link between our combinatorial problem and a problem in nonlinear optimisation.

Proposition 4.7. Let $C > 1$, $0 < \delta < 1$ and let $n > 1/\delta$. Suppose that G is a $(1 - \delta)$ -dense, k -coloured graph with $v(G) = Cn$, containing no monochromatic odd connected matching of order $\geq (1 + \delta)n$. Let v be a profile of G and let $x = v/n$. Then the following hold:

(1)

$$\left(\sum_{\tau \in \{0, 1, *\}^k} x_\tau \right)^2 - 2 \sum_{\{\sigma, \tau\} \in D} x_\sigma x_\tau - \sum_{\tau \in \{0, 1, *\}^k} \omega(\tau) x_\tau \leq 2\delta k C^2.$$

(2) $x_\tau \leq 1 + 4\delta C^2$ whenever $\omega(\tau) = 1$.

(3) $x_\tau x_\sigma \leq 2\delta C^2$ whenever σ and τ are incompatible.

Proof. Let us first remind ourselves of the graph decomposition discussed in Subsection 3.1. For each $i \in [k]$, we write $G_i = G'_i \cup G''_i$, where G'_i is the union of the bipartite components of G_i and G''_i is the union of the non-bipartite components of G_i . For each $i \in [k]$, write $V(G'_i) = V_0^i \cup V_1^i$ where V_0^i and V_1^i are the vertex classes of a bipartition of G'_i and set $V_*^i = V(G''_i)$. For $\tau \in \{0, 1, *\}^k$, set $V_\tau = \bigcap_{j=1}^k V_{\tau_j}^j$. Let $v = (|V_\tau| : \tau \in \{0, 1, *\}^k)$ be the profile corresponding to this partition. Let $N = v(G)$ and note that

$$(4.1) \quad N = \sum_{\tau \in \{0, 1, *\}^k} v_\tau.$$

It follows from Observation 4.1 that the number of bipartite edges in G is at most $\sum_{\{\sigma, \tau\} \in D} v_\sigma v_\tau$. Letting e_* denote the number of non-bipartite edges in G we therefore have that

$$(4.2) \quad e_* \geq e(G) - \sum_{\{\sigma, \tau\} \in D} v_\sigma v_\tau.$$

Since $N \geq 1/\delta$, we have

$$(4.3) \quad e(G) \geq (1 - \delta) \binom{N}{2} \geq (1 - 2\delta) \frac{N^2}{2}.$$

Combining (4.1), (4.2) and (4.3) gives

$$(4.4) \quad e_* \geq \frac{1}{2} \left(\sum_{\tau \in \{0,1,*\}^k} v_\tau \right)^2 - \sum_{\{\sigma,\tau\} \in D} v_\sigma v_\tau - \delta N^2.$$

We now find a corresponding upper bound for e_* . Recall that for $\tau \in \{0,1,*\}^k$, the weight $\omega(\tau)$ of τ is defined to be the size of the set $\{i \in [k] : \tau_i = *\}$.

By assumption, for each $i \in [k]$, every connected component of G''_i has no matching on $(1 + \delta)n$ vertices and so in particular G''_i has no cycle of length greater than $(1 + \delta)n$. Theorem 3.4 therefore implies that

$$(4.5) \quad e(G''_i) \leq (1 + \delta) \frac{n}{2} |V_*^i|.$$

Observe that

$$(4.6) \quad |V_*^i| = \sum_{\{\tau \in \{0,1,*\}^k : \tau_i = *\}} v_\tau.$$

Since each non-bipartite edge of G belongs to $E(G''_i)$ for some i , (4.5) and (4.6) provide the upper bound

$$(4.7) \quad e_* = \sum_{i=1}^k e(G''_i) \leq (1 + \delta) \frac{n}{2} \sum_{i=1}^k |V_*^i| = (1 + \delta) \frac{n}{2} \sum_{\tau \in \{0,1,*\}^k} \omega(\tau) v_\tau.$$

Since $\omega(\tau) \leq k$ for all $\tau \in \{0,1,*\}^k$ by definition, (4.1) and (4.7) imply the bound

$$(4.8) \quad e_* \leq \frac{1}{2} \delta n k N + \frac{n}{2} \sum_{\tau \in \{0,1,*\}^k} \omega(\tau) v_\tau.$$

Recall that $x = v/n$. Comparing the bounds (4.4) and (4.8) and scaling the resulting inequality by $2/n^2$ yields

$$\left(\sum_{\tau \in \{0,1,*\}^k} x_\tau \right)^2 - 2 \sum_{\{\sigma,\tau\} \in D} x_\sigma x_\tau - 2\delta C^2 \leq \sum_{\tau \in \{0,1,*\}^k} \omega(\tau) x_\tau + \delta k C.$$

This establishes (1). Notice that if $\tau \in \{0,1,*\}^k$ is such that $\omega(\tau) = 1$, then $G[V_\tau]$ is monochromatic with all edges non-bipartite by Observations 4.1 and 4.2. $G[V_\tau]$ therefore contains no cycle of length greater than $(1 + \delta)n$ and so by Theorem 3.4 and the fact that G has at most $\delta \binom{N}{2}$ edges missing

$$\binom{v_\tau}{2} - \delta \binom{N}{2} \leq e(G[V_\tau]) \leq (1 + \delta) \frac{n}{2} v_\tau.$$

It follows that

$$x_\tau^2 \leq (1 + 2\delta) x_\tau + \delta C^2,$$

from which (2) follows. Finally, let us note that by Observations 4.1 and 4.2, if σ, τ are incompatible, then there can be no edges lying between V_σ and V_τ . Since G has at most $\delta \binom{N}{2}$ edges missing we must then have

$$v_\tau v_\sigma \leq 2\delta N^2$$

(note that this inequality also accounts for the case where $\sigma = \tau$) and so (3) follows. \square

Given a graph G its profile lies in the space $\mathbb{R}^{\{0,1,*\}^k}$ which we will denote by \mathbb{R}^* . In view of Proposition 4.7 we define the function $F : \mathbb{R}^* \rightarrow \mathbb{R}$ by

$$F(x) = \left(\sum_{\tau \in \{0,1,*\}^k} x_\tau \right)^2 - 2 \sum_{\{\sigma,\tau\} \in D} x_\sigma x_\tau - \sum_{\tau \in \{0,1,*\}^k} \omega(\tau) x_\tau.$$

Let us also define the following subsets of \mathbb{R}^* .

X(γ): For $\gamma \geq 0$, let $X(\gamma)$ denote the set of elements $x \in \mathbb{R}^*$ satisfying:

(X1) $F(x) \leq \gamma$

(X2) $x_\tau \leq 1 + \gamma$ whenever $w(\tau) = 1$.

(X3) $x_\tau x_\sigma \leq \gamma$ whenever σ and τ are incompatible.

(X4) $x_\tau \geq 0$ for all τ .

Now let G be as in the statement of Theorem 3.5 and let v be a profile of G . By the above proposition we have $v/n \in X(\delta k 2^{2k})$ whereas we also have $\|v\| = 2^{k-1}n$. We will show that for δ small, this means that v/n is an element of almost maximal norm in $X(\delta k 2^{2k})$. We will also show that elements of large norm in $X(\delta k 2^{2k})$ have a very specific structure (in fact they resemble the profile of a hypercube colouring) and so this imposes a lot of structure on v . For now we focus our attention on the set $X(0)$ which we denote simply by X . Later on, we use compactness arguments to relate properties of X and $X(\gamma)$ for γ small.

Our next goal is to classify elements of maximal ℓ_1 -norm in X . To describe these elements we need a definition.

Definition 4.8. Call a set $\mathcal{A} \subseteq \{0, 1, *\}^k$ *distinguishable* if every pair of distinct elements of \mathcal{A} are distinguishable and also $\omega(\tau) \geq 1$ for all $\tau \in \mathcal{A}$.

The requirement that elements have weight at least 1 is for notational convenience later in the paper. Viewing elements of $\{0, 1, *\}^k$ as subcubes of Q_k , a distinguishable set is simply a collection of disjoint subcubes of Q_k (of dimension at least 1). If this collection covers the whole cube we give it a special name.

Definition 4.9. Call a distinguishable set $\mathcal{A} \subseteq \{0, 1, *\}^k$ a *decomposition* if $\bigcup_{\tau \in \mathcal{A}} Q(\tau) = \{0, 1, *\}^k$.

Let us quickly record a simple result concerning distinguishable sets which will become useful later.

Lemma 4.10. *Let $\mathcal{A} \subset \{0, 1, *\}^k$ be a distinguishable set. Then*

$$\sum_{\tau \in \mathcal{A}} 2^{\omega(\tau)} \leq 2^k,$$

with equality if and only if \mathcal{A} is a decomposition.

Proof. This is simply the observation that a distinguishable set \mathcal{A} is a collection of disjoint subcubes of Q_k and so the sum of their sizes $\sum_{\tau \in \mathcal{A}} 2^{\omega(\tau)}$ is bounded by the size of Q_k . Moreover we have equality if and only if these subcubes cover all of Q_k i.e. \mathcal{A} is a decomposition. \square

We define the following subset of \mathbb{R}^* .

O: Let O denote the set of elements $x \in \mathbb{R}^*$ satisfying:

(O1) $\text{supp}(x)$ is a decomposition where $\omega(\tau) = 1$ or 2 for all $\tau \in \text{supp}(x)$.

(O2) For all $\tau \in \text{supp}(x)$, if $\omega(\tau) = 1$ then $x_\tau = 1$ and if $\omega(\tau) = 2$ then $x_\tau = 2$.

It is easy to check that $O \subseteq X$. The next proposition asserts that O is the set of elements of maximal ℓ_1 -norm in X .

Proposition 4.11. *If $x \in X$, then $\|x\| \leq 2^{k-1}$ with equality if and only if $x \in O$.*

We note that the ‘if’ statement in the above proposition is immediate. Indeed if $x \in O$ then $\text{supp}(x)$ is a decomposition so that $\sum_{\tau \in \text{supp}(x)} 2^{\omega(\tau)} = 2^k$ by Lemma 4.10. Moreover $2^{\omega(\tau)} = 2x_\tau$ for all $\tau \in \text{supp}(x)$ by (O1) and (O2).

Definition 4.12. If $x \in X$ is such that $\|x\| = \sup_{z \in X} \|z\|$ then we say that x is an *optimal point* of X .

We note that since X is compact, optimal points of X exist. The proof of Proposition 4.11 is split over the next three sections.

5. COMPRESSIONS AND A SPHERICAL CONSTRAINT

In this section we make the first steps towards a proof of Proposition 4.11. Broadly speaking we apply the combinatorial technique of ‘shifting’ or ‘compression’ to transform the complicated non-linear constraint in the definition of $X = X(0)$ into a spherical constraint which is much more amenable to analysis. In Section 6 we apply optimisation tools to this transformed problem. We begin with a simple lemma concerning elements of $\{0, 1, *\}^k$.

Lemma 5.1. *If $\sigma, \tau \subseteq \{0, 1, *\}^k$ are indistinguishable and compatible and $\omega(\tau) = 1$, then $Q(\tau) \subseteq Q(\sigma)$. In particular if $\omega(\sigma) = 1$ also, then $\sigma = \tau$.*

Proof. Since σ, τ are indistinguishable and compatible we have $|Q(\tau) \cap Q(\sigma)| \geq 2$ by Lemma 4.5. However, $|Q(\tau)| = 2$ and so it follows that $Q(\tau) \subseteq Q(\sigma)$. If $\omega(\sigma) = 1$ also, then clearly $Q(\sigma) = Q(\tau)$ i.e. $\sigma = \tau$. \square

Definition 5.2. Let $x \in \mathbb{R}^*$. If all pairs of (not necessarily distinct) elements of $\text{supp}(x)$ are compatible, then we say that x has *compatible support*.

Let us note that condition (X3) (with $\gamma = 0$) in the definition of the set X is simply the condition that elements of X have compatible support. In particular, if $x \in X$ and $\tau \in \{0, 1, *\}^k$ has weight 0, then $x_\tau = 0$ since τ is not compatible with itself.

The following lemma establishes an important property of optimal points. For $\tau \in \{0, 1, *\}^k$ we let $e_\tau \in \mathbb{R}^*$ denote the standard unit vector whose entries are all 0 except the entry labelled τ which is 1.

Lemma 5.3. *If $x \in X$ is an optimal point, then $F(x) = 0$.*

Proof. Suppose for contradiction that $F(x) < 0$. Assume first that there exists $\tau \in \text{supp}(x)$ with $\omega(\tau) \geq 2$. By the continuity of F we may choose $\alpha > 0$ small enough so that $F(x + \alpha e_\tau) < 0$. Let $x' = x + \alpha e_\tau$. Since $\text{supp}(x') = \text{supp}(x)$ and x_τ remains unchanged for τ with $\omega(\tau) = 1$, it is clear that $x' \in X$. However, $\|x'\| = \|x\| + \alpha > \|x\|$ contradicting the fact that x is optimal.

We may assume then that $\text{supp}(x)$ consists only of elements of weight 1 and therefore is a distinguishable set by Lemma 5.1 and the fact that x has compatible support. It follows from the definition of F that

$$F(x) = \sum_{\tau \in \text{supp}(x)} (x_\tau^2 - x_\tau) < 0,$$

and so $x_\tau < 1$ for some $\tau \in \text{supp}(x)$. As before there exists some $\alpha > 0$ sufficiently small so that $F(x + \alpha e_\tau) < 0$. Let $x' = x + \alpha e_\tau$. If we pick α small enough so that $x'_\tau = x_\tau + \alpha \leq 1$ also, then again we have $x' \in X$ with $\|x'\| > \|x\|$ contradicting the optimality of x . \square

We now describe the transformations alluded to at the beginning of this section. They will be of great use in simplifying our analysis of optimal points of X .

Definition 5.4. Let $x \in \mathbb{R}^*$ and let $\pi, \rho \in \{0, 1, *\}^k$ be distinct. We define the (π, ρ) -compression of x , denoted $x(\pi, \rho)$, as follows:

- If $\omega(\rho) \geq 2$, or if $\omega(\rho) \leq 1$ and $x_\pi + x_\rho < 1$, then let $x(\pi, \rho)$ be the vector x' with coordinates: $x'_\pi = 0$, $x'_\rho = x_\pi + x_\rho$ and $x'_\tau = x_\tau$ for all $\tau \in \{0, 1, *\}^k \setminus \{\pi, \rho\}$.
- If $\omega(\rho) \leq 1$ and $x_\pi + x_\rho \geq 1$ then let $x(\pi, \rho)$ be the vector x' with coordinates: $x'_\pi = x_\pi + x_\rho - 1$, $x'_\rho = 1$ and $x'_\tau = x_\tau$ for all $\tau \in \{0, 1, *\}^k \setminus \{\pi, \rho\}$.

If $x(\pi, \rho) = x$ then we say that x is (π, ρ) -compressed.

Let $x \in X$ be an optimal point, we will be interested in instances where $x(\pi, \rho)$ is also an optimal point of X . We observe that if $x \in \mathbb{R}^*$ and $\pi, \rho \in \{0, 1, *\}^k$ are distinct then $\|x(\pi, \rho)\| = \|x\|$. However, if $x \in X$ then it does not follow in general that $x(\pi, \rho) \in X$.

We only consider (π, ρ) -compressions in the case where π and ρ are indistinguishable since it is easier to track the effect of the compression on the value of the function F in this case (see the proof of the next lemma). It will therefore be useful to associate to each point $x \in X$, the digraph $D(x) = (V(x), E(x))$ where $V(x) = \text{supp}(x)$ and

$$E(x) = \{(\pi, \rho) : \pi, \rho \in \text{supp}(x) \text{ are distinct, indistinguishable and } x(\pi, \rho) \in X\}.$$

In particular if $x \in X$ is (π, ρ) -compressed, where π and ρ are distinct and indistinguishable, then $(\pi, \rho) \in E(x)$. We draw attention to the fact that edges of $D(x)$ only occur between indistinguishable pairs. Conversely, the following lemma shows that, when $x \in X$ is optimal, at least one edge occurs between any indistinguishable pair in $D(x)$.

Lemma 5.5. *Let $x \in X$ be optimal and suppose that $\pi, \rho \in V(x)$, are indistinguishable and distinct. Then one of the following holds:*

- (i) x is (π, ρ) -compressed, $x_\rho = 1$, $\omega(\rho) = 1$, $\omega(\pi) \geq 2$ and $(\rho, \pi) \notin E(x)$,

- (ii) x is (ρ, π) -compressed, $x_\pi = 1$, $\omega(\pi) = 1$, $\omega(\rho) \geq 2$ and $(\pi, \rho) \notin E(x)$,
- (iii) (ρ, π) and (π, ρ) both lie in $E(x)$.

Proof. Recall that

$$F(x) = \left(\sum_{\tau \in \{0,1,*\}^k} x_\tau \right)^2 - 2 \sum_{\{\sigma,\tau\} \in D} x_\sigma x_\tau - \sum_{\tau \in \{0,1,*\}^k} \omega(\tau) x_\tau,$$

where D is the set of unordered distinguishable pairs from $\{0, 1, *\}^k$. Since π, ρ are indistinguishable, the term $x_\rho x_\pi$ does not appear in the sum $\sum_{\{\sigma,\tau\} \in D} x_\sigma x_\tau$. We may therefore express $F(x)$ in the form

$$(5.1) \quad F(x) = \left(\sum_{\tau \in \{0,1,*\}^k} x_\tau \right)^2 - Ax_\rho - Bx_\pi - C,$$

where A, B and C do not depend on x_ρ or x_π and $A, B \geq 0$. Suppose that $A \geq B$ and let $x' = x(\pi, \rho)$. Let us show that $(\pi, \rho) \in E(x)$ i.e. $x' \in X$. By (5.1) we have $F(x') \leq F(x)$ and so x' satisfies (X1) in the definition of X . By the definition of (π, ρ) -compression it is clear that x' also satisfies (X2) and (X4). Since $\rho \in \text{supp}(x)$, we also have $\text{supp}(x') \subseteq \text{supp}(x)$. Since x has compatible support the same is true for x' i.e. x' satisfies (X3) and so $x' \in X$. Note that since compressions preserve the ℓ_1 -norm, x' is also an optimal point of X .

In the case $A = B$, an identical argument shows that $(\rho, \pi) \in E(x)$ also, and so (iii) holds.

Suppose then that $A > B$. In this case, looking again at (5.1), we see that if x is not (π, ρ) -compressed then we in fact have $F(x') < F(x) = 0$, contradicting Lemma 5.3. We conclude that x is (π, ρ) -compressed. Suppose $\omega(\rho) \geq 2$, then by the definition of (π, ρ) -compression we have $x_\pi = x'_\pi = 0$ contradicting the fact that $\pi \in \text{supp}(x)$, and so $\omega(\rho) = 1$. Since π and ρ are compatible, indistinguishable and distinct, it follows from Lemma 5.1 that $\omega(\pi) \geq 2$. Let $x'' = x(\rho, \pi)$. It follows that $x''_\rho = 0$ and $x''_\pi = x_\rho + x_\pi$ and so by (5.1), $F(x'') > F(x) = 0$. We conclude that $x'' \notin X$ i.e. $(\rho, \pi) \notin E(x)$. The fact that $x_\rho = 1$ follows from the fact that $\omega(\rho) = 1$, x is (π, ρ) -compressed and $\pi \in \text{supp}(x)$. Thus (i) holds, and similarly if $A < B$ then (ii) holds. \square

We obtain the following immediate corollary.

Corollary 5.6. *Let $x \in X$ be an optimal point and suppose that I is an independent set in $D(x)$. Then I is a distinguishable set.*

Definition 5.7. We call an optimal point $x \in X$ *compressed* if it is (π, ρ) -compressed for all $(\pi, \rho) \in E(x)$.

We now show that compressed optimal points of X exist. In fact we show that given any optimal point of $x \in X$ we may obtain a compressed optimal point by applying a finite number of compressions to x . The simpler structure of compressed optimal points will make it easier to bound their ℓ_1 -norm which is the goal of Proposition 4.11.

Lemma 5.8. *Compressed optimal points of X exist.*

Proof. Let x be an arbitrary optimal point of X and define a sequence x_0, x_1, x_2, \dots of elements of X recursively as follows: Set $x_0 = x$. Having chosen x_0, \dots, x_t , if x_t is compressed

then stop the sequence at x_t . If not, then there exists $(\pi, \rho) \in E(x_t)$ such that x_t is not (π, ρ) -compressed. By Lemma 5.5, we must therefore have that $x(\pi, \rho)$ and $x(\rho, \pi)$ are both optimal points of X . Note that by the definition of $D(x_t)$, ρ and π are indistinguishable and $\{\pi, \rho\} \subseteq \text{supp}(x_t)$. Since x_τ has compatible support, it follows from Lemma 5.1 that either $\omega(\pi) \geq 2$ or $\omega(\rho) \geq 2$. If $\omega(\rho) \geq 2$ then set $x_{t+1} = x(\pi, \rho)$, if not (so that $\omega(\pi) \geq 2$) set $x_{t+1} = x(\rho, \pi)$. In either case x_{t+1} is an optimal point of X satisfying $|V(x_{t+1})| = |V(x_t)| - 1$. Since $0 \leq |V(x)| \leq 3^k$ for all $x \in X$ it follows that the sequence must terminate in at most 3^k steps. \square

Having discovered compressed optimal points, we now explore some of their properties. First we need a definition.

Definition 5.9. A *star* is a digraph with vertex set $\{\rho, \pi_1, \dots, \pi_m\}$ (for some $m \geq 0$) and edge set $\{(\rho, \pi_1), \dots, (\rho, \pi_m)\}$. We refer to ρ as the *root* of the star and we call π_1, \dots, π_m *leaves*. Note that we have included the possibility of a star with no leaves.

Lemma 5.10. *Let $x \in X$ be a compressed optimal point, then $D(x)$ is a disjoint union of stars. Moreover if ρ is a root of positive outdegree then $\omega(\rho) \geq 2$ and if π is a leaf then $\omega(\pi) = 1$ and $x_\pi = 1$.*

Proof. It suffices to prove the following:

- (1) If $(\rho, \pi) \in E(x)$ then $\omega(\rho) \geq 2$, $\omega(\pi) = 1$, $x_\pi = 1$ and $(\pi, \rho) \notin E(x)$.
- (2) If $(\rho_1, \pi), (\rho_2, \pi) \in E(x)$ then $\rho_1 = \rho_2$.

Suppose that $(\rho, \pi) \in E(x)$, in particular ρ and π are indistinguishable. If $(\pi, \rho) \in E(x)$ also, then since x is compressed we have by definition that $x(\rho, \pi) = x = x(\pi, \rho)$. However, from the definition of compression we see that the only way we can have $x(\rho, \pi) = x(\pi, \rho)$ is if $\omega(\pi) = \omega(\rho) = 1$. But then by Lemma 5.1, $\pi = \rho$, a contradiction. We conclude that $(\pi, \rho) \notin E(x)$ and so (1) follows from Lemma 5.5.

Suppose now that $(\rho_1, \pi), (\rho_2, \pi) \in E(x)$. By (1) we know that $\omega(\pi) = 1$ and $\omega(\rho_i) \geq 2$ for $i = 1, 2$. By Lemma 5.1 it follows that $Q(\pi) \subseteq Q(\rho_1) \cap Q(\rho_2)$ and so ρ_1, ρ_2 are indistinguishable by Lemma 4.5. If $\rho_1 \neq \rho_2$ then by Lemma 5.5 we have that either $(\rho_1, \rho_2) \in E(x)$ or $(\rho_2, \rho_1) \in E(x)$, but this contradicts (1). \square

Recall that $V(x) = \text{supp}(x)$. Given a compressed optimal point $x \in X$ let

$$L(x) = \{\tau \in V(x) : d^-(\tau) > 0\}$$

and

$$R(x) = V(x) \setminus L(x).$$

By Lemma 5.10, $L(x)$ and $R(x)$ are the set of leaves and the set of roots of $D(x)$ respectively.

Lemma 5.11. *Let $x \in X$ be a compressed optimal point. Then*

$$F(x) = \sum_{\tau \in R(x)} (x_\tau^2 + (2d^+(\tau) - \omega(\tau))x_\tau).$$

Proof. Lemmas 5.5 and 5.10 together show that for any indistinguishable pair $\pi, \rho \in V(x)$, where $\pi \neq \rho$, exactly one of (π, ρ) and (ρ, π) is in $E(x)$ and so

$$F(x) = \sum_{\tau \in V(x)} x_\tau^2 + 2 \sum_{(\sigma, \tau) \in E(x)} x_\sigma x_\tau - \sum_{\tau \in V(x)} \omega(\tau) x_\tau.$$

By Lemma 5.10 we may write

$$\sum_{(\sigma, \tau) \in E(x)} x_\sigma x_\tau = \sum_{\rho \in R(x)} x_\rho \left(\sum_{\pi: (\rho, \pi) \in E(x)} x_\pi \right) = \sum_{\rho \in R(x)} d^+(\rho) x_\rho.$$

Moreover by Lemma 5.10 we have $\sum_{\tau \in L(x)} (x_\tau^2 - \omega(\tau)x_\tau) = 0$. The result follows. \square

The key feature here is that for a compressed optimal point x , the constraint equation $F(x) = 0$ is spherical. This allows us to more easily apply standard optimisation techniques and this will be the concern of the next section. For now it will be useful for us to establish some degree conditions on the vertices of $D(x)$ for a compressed optimal point $x \in X$.

Lemma 5.12. *Let $x \in X$ be a compressed optimal point, then $d^+(\sigma) \leq 2^{\omega(\sigma)-1}$ for all $\sigma \in V(x)$.*

Proof. Suppose that $\rho \in V(x)$ is such that $d^+(\rho) > 0$. By Lemma 5.10, ρ is the root of a star in $D(x)$. Let L be the set of leaves of this star (so in particular $|L| = d^+(\rho)$ and $\omega(\pi) = 1$ for all $\pi \in L$). Note that L is an independent set in $D(x)$ and therefore it is a distinguishable set by Corollary 5.6. In particular, the sets $Q(\pi)$ where $\pi \in L$ are pairwise disjoint by Lemma 4.5. Note further that for each $\pi \in L$, ρ and π are indistinguishable and compatible and hence $Q(\pi) \subseteq Q(\rho)$ by Lemma 5.1. It follows that

$$2d^+(\rho) = \sum_{\pi \in L} |Q(\pi)| \leq |Q(\rho)| = 2^{\omega(\rho)}.$$

\square

We can now bootstrap, using the previous two lemmas to establish a much stronger degree condition. The idea behind the proof of the following lemma is readily explained however it is notationally laborious. The idea is that if $x \in X$ is a compressed optimal point and a star in $D(x)$ with root ρ has $> \omega(\rho)$ leaves, then one can contradict the optimality of x by replacing this star with a collection of stars whose roots have smaller weight. First let us generalise an earlier notation.

Definition 5.13. Let $\sigma \in \{0, 1, *\}^k$ and let $W = \{i \in [k] : \sigma_i = *\}$. Then for $S \subseteq W$ define

$$Q(\sigma; S) = \{\tau \in \{0, 1, *\}^k : \tau_i \in \{0, 1\} \text{ for } i \in W \setminus S \text{ and } \tau_i = \sigma_i \text{ otherwise}\}.$$

Note that elements of $Q(\sigma; S)$ are pairwise distinguishable and that $Q(\sigma; \emptyset)$ is simply the set $Q(\sigma)$. We may think of $Q(\sigma; S)$ as a decomposition of $Q(\sigma)$ into ‘parallel’ subcubes of dimension $|S|$.

Lemma 5.14. *Let $x \in X$ be a compressed optimal point, then $d^+(\sigma) \leq \omega(\sigma)$ for all $\sigma \in V(x)$.*

Proof. By Lemma 5.10 we may write $V(x) = S_1 \cup \dots \cup S_q$, a disjoint union where each S_i is the vertex set of a star in $D(x)$. Suppose that there exists $\sigma \in V(x)$ such that $d^+(\sigma) > \omega(\sigma)$. Without loss of generality assume σ is the root of S_1 . By Lemma 5.11 we then have that $\omega(\sigma) < d^+(\sigma) \leq 2^{\omega(\sigma)-1}$ and so $\omega(\sigma) \geq 3$. Without loss of generality assume that $\sigma_1 = \sigma_2 = *$. By Lemma 5.10 we have $x_\tau = 1$ for all $\tau \in L(x)$ and so

$$(5.2) \quad \|x\| = |L(x)| + \sum_{\tau \in R(x)} x_\tau.$$

We proceed by modifying x , being careful to stay within the set X . Take $\pi \in Q(\sigma; \{1, 2\})$ and note that $\omega(\pi) = 2$. Consider now the element $x' \in \mathbb{R}^*$ defined as follows. Let $x'_\pi = x_\sigma$, $x'_\tau = 1$ for all $\tau \in Q(\sigma; \{1\})$, $x'_\tau = x_\tau$ for all $\tau \in S_2 \cup \dots \cup S_q$ and $x'_\tau = 0$ otherwise. We now check that $x' \in X$. Clearly $x'_\tau \leq 1$ whenever $\omega(\tau) = 1$. Let $S'_1 = \{\pi\} \cup Q(\sigma; \{1\})$ and note that $\text{supp}(x') = S'_1 \cup S_2 \cup \dots \cup S_q$. Now, if $\tau \in S'_1$ we have $Q(\tau) \subseteq Q(\sigma)$ and since $\sigma \in S_1$, we know that σ , and hence also τ , is distinguishable from each element of $S_2 \cup \dots \cup S_k$. Since $\tau_1 = *$ for each $\tau \in S'_1$ we see that S'_1 contains no incompatible pairs. It follows that x' has compatible support (i.e. condition (X3) holds for x'). Finally, by a calculation similar to that in the proof of Lemma 5.11 and the fact that ρ and σ are distinguishable for all but exactly two elements $\rho \in Q(\sigma; \{1\})$, we have

$$(5.3) \quad \begin{aligned} F(x') &= x_\sigma^2 + 2x_\sigma + \sum_{\tau \in R(x) \setminus \{\sigma\}} (x_\tau^2 + (2d^+(\tau) - \omega(\tau))x_\tau) \\ &= F(x) - (2d^+(\sigma) - \omega(\sigma) - 2)x_\sigma. \end{aligned}$$

Recalling that $d^+(\sigma) > \omega(\sigma)$ we have $2d^+(\sigma) - \omega(\sigma) > \omega(\sigma) \geq 3$. Since $\sigma \in \text{supp}(x)$, we also have $x_\sigma > 0$ and so (5.3) implies that $F(x') < F(x) = 0$. Thus we do indeed have $x' \in X$. Note that

$$(5.4) \quad \|x'\| = |L(x)| - d^+(\sigma) + |Q(\sigma; \{1\})| + \sum_{\tau \in R(x)} x_\tau,$$

and observe that $d^+(\sigma) \leq 2^{\omega(\sigma)-1} = |Q(\sigma; \{1\})|$ by Lemma 5.11. It now follows from (5.2) and (5.4) that $\|x'\| \geq \|x\|$, and so x' is an optimal point of X . However we have shown that $F(x') < 0$ contradicting Lemma 5.3. \square

Gathering all the information we have obtained on compressed optimal points, we show that a proof of the following proposition is almost enough to deduce Proposition 4.11. Let us remind ourselves that in the definition of a distinguishable set (Definition 4.8), we require all elements of the set to have weight at least 1.

Proposition 5.15. *Let $\mathcal{A} \subseteq \{0, 1, *\}^k$ be a distinguishable set and let $\Delta : \mathcal{A} \rightarrow \mathbb{Z}_{\geq 0}$ be such that $\Delta(\tau) \leq \omega(\tau)$ for all $\tau \in \mathcal{A}$, and $\Delta(\tau) = 0$ whenever $\omega(\tau) = 1$. Suppose that $x \in \mathbb{R}^*$ is a vector with $\text{supp}(x) = \mathcal{A}$ satisfying*

(1)

$$\sum_{\tau \in \mathcal{A}} (x_\tau^2 + (2\Delta(\tau) - \omega(\tau))x_\tau) = 0,$$

(2) $x_\tau \leq 1$ whenever $\omega(\tau) = 1$.

Then

$$\sum_{\tau \in \mathcal{A}} x_\tau \leq 2^{k-1} - \sum_{\tau \in \mathcal{A}} \Delta(\tau).$$

Furthermore we have equality only if $x \in O$ and $\Delta \equiv 0$.

A proof of Proposition 5.15 will be the focus of the next section, for now we note that it has the following corollary

Corollary 5.16. *O is the set of compressed optimal points of X . In particular, if $x \in X$ then $\|x\| \leq 2^{k-1}$.*

Proof. Let $x \in X$ be a compressed optimal point. Note that $R(x)$ is an independent set in the digraph $D(x)$ and hence by Corollary 5.6, $R(x)$ is a distinguishable set. Let $\Delta : R(x) \rightarrow \mathbb{Z}_{\geq 0}$ be the map given by $\tau \mapsto d^+(\tau)$. Let x' be the element of \mathbb{R}^* supported on $R(x)$ such that $x'_\tau = x_\tau$ for all $\tau \in R(x)$. Note that by Lemmas 5.3, 5.11, 5.14 and by the definition of the set X , we have that Δ and x' satisfy the conditions in the statement of Proposition 5.15. Assuming Proposition 5.15, it therefore follows that

$$(5.5) \quad \sum_{\tau \in R(x)} x_\tau \leq 2^{k-1} - \sum_{\tau \in R(x)} d^+(\tau),$$

with equality only if $x' \in O$ and $d^+(\tau) = 0$ for all $\tau \in R(x)$. The latter condition implies that $x' = x$ and so we have equality in (5.5) only if $x \in O$. By Lemma 5.10,

$$\|x\| = \sum_{\tau \in R(x)} x_\tau + \sum_{\tau \in R(x)} d^+(\tau),$$

and so it follows that $\|x\| \leq 2^{k-1}$ with equality only if $x \in O$. The result follows by noting that for all $z \in O$, $\|z\| = 2^{k-1}$ and z is compressed (since for every distinct pair $\pi, \rho \in \text{supp}(x)$, π and ρ are distinguishable). □

It is now clear that Proposition 4.11 would follow if we could also prove the following.

Proposition 5.17. *If $x \in X$ is an optimal point, then x is compressed.*

We prove Proposition 5.17 in Section 7.

6. CONSTRAINED OPTIMISATION AND A PROOF OF PROPOSITION 5.15

In this section we prove Proposition 5.15 thus finalising the main stepping stone toward a proof of Proposition 4.11. We use standard tools from the theory of convex optimisation to exploit the spherical constraint found in the previous section. This will lead us to consider the possible distributions of weights in distinguishable sets which we optimise over in a separate argument. The main tools that we borrow are the Karush-Kuhn-Tucker (KKT) conditions along with Slater's constraint qualification. Below is a statement of the result we use, phrased to match our needs (see [BV04, p. 244] for a detailed account).

Theorem 6.1 (KKT + Slater's Condition). *Let $f, g_1, \dots, g_r : \mathbb{R}^m \rightarrow \mathbb{R}$ be convex, differentiable functions and let*

$$S = \{x \in \mathbb{R}^m : g_i(x) \leq 0 \text{ for } i = 1, \dots, r\}.$$

Suppose that there exists an $x_0 \in \mathbb{R}^m$ such that $g_i(x_0) < 0$ for $i = 1, \dots, r$. Then if $x^ \in S$ is such that*

$$f(x^*) = \sup_{x \in S} f(x),$$

then there exist $\lambda_1, \dots, \lambda_r \in \mathbb{R}$ such that

- (i) $\nabla f(x^*) = \sum_{i=1}^r \lambda_i \nabla g_i(x^*)$,
- (ii) $\lambda_i \geq 0$, $i = 1, \dots, r$,
- (iii) $\lambda_i g_i(x^*) = 0$, $i = 1, \dots, r$.

In view of the statement of Proposition 5.15 it is natural to apply Theorem 6.1 to establish the following.

Lemma 6.2. Let $\alpha_1, \dots, \alpha_m$ be integers where $\alpha_i = 1$ for $i = 1, \dots, \ell$, ($0 \leq \ell \leq m$) and let x_1, \dots, x_m be real numbers satisfying

- (i) $\sum_{i=1}^m (x_i^2 - \alpha_i x_i) \leq 0$,
- (ii) $x_i \leq 1$ for $i = 1, \dots, \ell$.

If $\sum_{i=1}^m \alpha_i^2 > m$, then

$$\sum_{i=1}^m x_i \leq \ell + \frac{1}{2} \left(\sum_{i=\ell+1}^m \alpha_i + \sqrt{(m-\ell) \sum_{i=\ell+1}^m \alpha_i^2} \right),$$

with equality only if $x_i = 1$ for $i \leq \ell$ and $x_i = \frac{1}{2} \left(\alpha_i + \sqrt{\frac{1}{(m-\ell)} \sum_{i=\ell+1}^m \alpha_i^2} \right)$ for $i > \ell$.

If instead $\sum_{i=1}^m \alpha_i^2 \leq m$, then

$$\sum_{i=1}^m x_i \leq \frac{1}{2} \left(\sum_{i=1}^m \alpha_i + \sqrt{m \sum_{i=1}^m \alpha_i^2} \right),$$

with equality only if $x_i = \frac{1}{2} \left(\alpha_i + \sqrt{\frac{1}{m} \sum_{i=1}^m \alpha_i^2} \right)$ for all i .

Proof. Note first that if $\alpha_i = 0$ for all i (so in particular $\ell = 0$) then inequality (i) implies that $x_i = 0$ for all i in which case there's nothing to prove. Suppose then that this is not the case and define functions $f, g_1, \dots, g_{\ell+1} : \mathbb{R}^m \rightarrow \mathbb{R}$ as follows. Let $f(x) = \sum_{i=1}^m x_i$, $g_i(x) = x_i - 1$ for $i = 1, \dots, \ell$ and $g_{\ell+1}(x) = \sum_{i=1}^m (x_i^2 - \alpha_i x_i)$. Note that the functions just defined are all convex and differentiable. Let $x_0 = (\alpha_1/2, \dots, \alpha_m/2)$, the centre of the spherical region described by (i) and observe that $g_i(x_0) < 0$ for $i = 1, \dots, \ell + 1$. Let $S = \{x \in \mathbb{R}^m : g_i(x) \leq 0 \text{ for } i = 1, \dots, \ell + 1\}$. S is compact and f is continuous hence we may pick $x^* \in S$ such that

$$f(x^*) = \sup_{x \in S} f(x).$$

Let $x^* = (x_1, \dots, x_m)$. By Theorem 6.1, there exist real numbers $\lambda_1, \dots, \lambda_\ell$ and Λ such that the following hold (for notational convenience we define $\lambda_j = 0$ for $j > \ell$):

- (1) $\Lambda(2x_i - \alpha_i) + \lambda_i = 1$ for all $i \in [m]$,
- (2) $\Lambda \geq 0$ and $\lambda_i \geq 0$ for all $i \in [m]$,
- (3) $\Lambda(\sum_{i=1}^m (x_i^2 - \alpha_i x_i)) = 0$ and $\lambda_i(x_i - 1) = 0$ for all $i \in [m]$.

We consider three cases depending on the value of Λ . First let us suppose that $\Lambda = 0$. In this case, by (1) we must have $\lambda_i = 1$ for all $i \in [m]$. Recalling that $\lambda_j = 0$ for $j > \ell$ by definition, we must also have $\ell = m$ and so $\alpha_i = 1$ for all $i \in [m]$. Moreover, it follows from (3) that $x_i = 1$ for all $i \in [m]$ and so we're done.

By (2), we may now assume that $\Lambda > 0$ and so we may rewrite (1) as

$$(6.1) \quad x_i = \frac{1}{2} \left(\frac{1 - \lambda_i}{\Lambda} + \alpha_i \right) \text{ for all } i \in [m].$$

Moreover, $\sum_{i=1}^m (x_i^2 - \alpha_i x_i) = 0$ by (3) which by (6.1) gives

$$(6.2) \quad \frac{1}{\Lambda^2} \sum_{i=1}^m (1 - \lambda_i)^2 = \sum_{i=1}^m \alpha_i^2.$$

Now, note that for $i \leq \ell$ we have $\alpha_i = 1$ and $x_i \leq 1$ and so by (6.1) we have

$$(6.3) \quad 1 - \Lambda \leq \lambda_i \text{ for } i \leq \ell.$$

If $\Lambda < 1$ then by (6.3) we have $\lambda_i > 0$ for $i \leq \ell$ and so by (3), $x_i = 1$ for $i \leq \ell$ and so in fact by (6.1)

$$1 - \lambda_i = \Lambda \text{ for } i \leq \ell.$$

Recalling that $\alpha_i = 1$ for $i \leq \ell$ and $\lambda_i = 0$ for $i > \ell$ by definition, (6.2) then gives

$$(6.4) \quad \frac{1}{\Lambda} = \sqrt{\frac{1}{m - \ell} \sum_{i=\ell+1}^m \alpha_i^2}.$$

From (6.1) it now follows that

$$x_i = \frac{1}{2} \left(\alpha_i + \sqrt{\frac{1}{m - \ell} \sum_{i=\ell+1}^m \alpha_i^2} \right) \text{ for } i > \ell,$$

and so

$$\sum_{i=1}^m x_i = \ell + \frac{1}{2} \left(\sum_{i=\ell+1}^m \alpha_i + \sqrt{(m - \ell) \sum_{i=\ell+1}^m \alpha_i^2} \right).$$

Recalling that $\Lambda < 1$, it follows from (6.4) that $\sum_{i=\ell+1}^m \alpha_i^2 > m$.

It remains to consider the case where $\Lambda \geq 1$. Recall that if $\lambda_i > 0$ for some i then $x_i = 1$ by (3). Moreover we must have $i \in [j]$ by definition and so $\Lambda = 1 - \lambda_i$ by (6.1). However this contradicts the assumption that $\Lambda \geq 1$ and so we conclude that $\lambda_i = 0$ for all $i \in [m]$. It follows from (6.2) that

$$(6.5) \quad \frac{1}{\Lambda} = \sqrt{\frac{1}{m} \sum_{i=1}^m \alpha_i^2},$$

so that by (6.1),

$$x_i = \frac{1}{2} \left(\alpha_i + \sqrt{\frac{1}{m} \sum_{i=1}^m \alpha_i^2} \right) \text{ for all } i \in [m].$$

The result follows, noting that by (6.5) we have $\sum_{i=1}^m \alpha_i^2 \leq m$ in this case. \square

We are almost ready to prove Proposition 5.15, but first we need the following inequality.

Lemma 6.3. *If $\alpha_1, \dots, \alpha_m \geq 2$ are integers, then*

$$\sum_{i=1}^m \alpha_i + \sqrt{m \sum_{i=1}^m \alpha_i^2} \leq \sum_{i=1}^m 2^{\alpha_i},$$

and equality holds if and only if $\alpha_i = 2$ for all i .

Proof. Let $\alpha = (\alpha_1, \dots, \alpha_m)$. We induct on the value of $S_\alpha := \sum_{i=1}^m (2^{\alpha_i-2} - 1)$. If $S_\alpha = 0$, then $\alpha_i = 2$ for all i , so that

$$\sum_{i=1}^m \alpha_i + \sqrt{m \sum_{i=1}^m \alpha_i^2} = 4m = \sum_{i=1}^m 2^{\alpha_i}.$$

Suppose then that $S_\alpha > 0$ so that $\alpha_j \geq 3$ for some $j \in [m]$. Without loss of generality assume that $j = 1$. Define a new sequence of integers $\alpha' = (\alpha'_1, \dots, \alpha'_{m+1})$, as follows: Let $\alpha'_1 = \alpha'_2 = \alpha_1 - 1$ and $\alpha'_i = \alpha_{i-1}$ for $i = 3, 4, \dots, m+1$. Note that $\alpha'_i \geq 2$ for all i and $S_{\alpha'} = S_\alpha - 1$ and so by the inductive hypothesis

$$(6.6) \quad \sum_{i=1}^{m+1} \alpha'_i + \sqrt{(m+1) \sum_{i=1}^{m+1} \alpha_i'^2} \leq \sum_{i=1}^{m+1} 2^{\alpha'_i},$$

Note that

$$(6.7) \quad \sum_{i=1}^{m+1} 2^{\alpha'_i} = \sum_{i=1}^m 2^{\alpha_i} \quad \text{and} \quad \sum_{i=1}^{m+1} \alpha'_i - \sum_{i=1}^m \alpha_i = \alpha_1 - 2 > 0,$$

and also

$$(6.8) \quad (m+1) \sum_{i=1}^{m+1} \alpha_i'^2 - m \sum_{i=1}^m \alpha_i^2 = \sum_{i=1}^m \alpha_i^2 + (m+1)(\alpha_1^2 - 4\alpha_1 + 2) \geq \sum_{i=1}^m \alpha_i^2 - (m+1),$$

where in the last inequality we used the fact that $\alpha^2 - 4\alpha + 2 \geq -1$ for $\alpha \geq 3$. Note that since $\alpha_i \geq 2$ for all i , we certainly have that $\sum_{i=1}^m \alpha_i^2 > m+1$. It follows then from (6.8) that

$$(6.9) \quad (m+1) \sum_{i=1}^{m+1} \alpha_i'^2 > m \sum_{i=1}^m \alpha_i^2.$$

Combining (6.6), (6.7), and (6.9) we have

$$\sum_{i=1}^m \alpha_i + \sqrt{m \sum_{i=1}^m \alpha_i^2} < \sum_{i=1}^{m+1} \alpha'_i + \sqrt{(m+1) \sum_{i=1}^{m+1} \alpha_i'^2} \leq \sum_{i=1}^{m+1} 2^{\alpha'_i} = \sum_{i=1}^m 2^{\alpha_i}$$

as required. Note the strict inequality, and so we only have equality in the case where $\alpha_i = 2$ for all i . \square

Proof of Proposition 5.15. Consider first the case where $\sum_{\tau \in \mathcal{A}} (\omega(\tau) - 2\Delta(\tau))^2 > |\mathcal{A}|$. Suppose that ℓ elements of \mathcal{A} have weight 1 and let $\mathcal{A}' = \{\tau \in \mathcal{A} : \omega(\tau) \geq 2\}$. Note that by Lemma 4.10 we have $\sum_{\tau \in \mathcal{A}} 2^{\omega(\tau)} \leq 2^k$ and hence

$$(6.10) \quad \sum_{\tau \in \mathcal{A}'} 2^{\omega(\tau)} \leq 2^k - 2\ell.$$

Applying Lemma 6.2, recalling that $0 \leq \Delta(\tau) \leq \omega(\tau)$ for all $\tau \in \mathcal{A}$, that $\Delta(\tau) = 0$ whenever $\omega(\tau) = 1$ and using (6.10) and Lemma 6.3 we have

$$(6.11) \quad \sum_{\tau \in \mathcal{A}} x_\tau \leq \ell + \frac{1}{2} \left(\sum_{\tau \in \mathcal{A}'} (\omega(\tau) - 2\Delta(\tau)) + \sqrt{|\mathcal{A}'| \sum_{\tau \in \mathcal{A}'} (\omega(\tau) - 2\Delta(\tau))^2} \right)$$

$$(6.12) \quad \leq \ell + \frac{1}{2} \left(\sum_{\tau \in \mathcal{A}'} \omega(\tau) + \sqrt{|\mathcal{A}'| \sum_{\tau \in \mathcal{A}'} \omega(\tau)^2} \right) - \sum_{\tau \in \mathcal{A}} \Delta(\tau)$$

$$(6.13) \quad \leq \ell + \sum_{\tau \in \mathcal{A}'} 2^{\omega(\tau)-1} - \sum_{\tau \in \mathcal{A}} \Delta(\tau)$$

$$(6.14) \quad \leq 2^{k-1} - \sum_{\tau \in \mathcal{A}} \Delta(\tau).$$

We analyse the conditions for equality to hold. For equality to hold in (6.12) it must be the case that for all $\tau \in \mathcal{A}$, either $\Delta(\tau) = 0$ or $\Delta(\tau) = \omega(\tau)$. By Lemma 6.3, for equality to hold in (6.13) it must be the case that $\omega(\tau) = 2$ for all $\tau \in \mathcal{A}'$. It now follows from Lemma 6.2 that for equality to also hold in (6.11), we must have $x_\tau = 1$ whenever $\omega(\tau) = 1$, $x_\tau = 2$ for all $\tau \in \mathcal{A}'$ such that $\Delta(\tau) = 0$ and $x_\tau = 0$ for all $\tau \in \mathcal{A}'$ such that $\Delta(\tau) = \omega(\tau)$. However, since each x_τ is non-zero by assumption we conclude that $\Delta(\tau) = 0$ for all $\tau \in \mathcal{A}$ i.e. $\Delta \equiv 0$. Finally, for equality to hold in (6.14) we must have equality in (6.10) and so \mathcal{A} is a decomposition by Lemma 4.10. It follows that $x \in O$.

It remains to consider the case where $\sum_{\tau \in \mathcal{A}} (\omega(\tau) - 2\Delta(\tau))^2 \leq |\mathcal{A}|$. By Lemma 6.2 and Lemma 4.10 we then have

$$(6.15) \quad \begin{aligned} \sum_{\tau \in \mathcal{A}} x_\tau &\leq \frac{1}{2} \left(|\mathcal{A}| + \sum_{\tau \in \mathcal{A}} \omega(\tau) \right) - \sum_{\tau \in \mathcal{A}} \Delta(\tau) \\ &\leq \frac{1}{2} \left(|\mathcal{A}| + \sum_{\tau \in \mathcal{A}} 2^{\omega(\tau)-1} \right) - \sum_{\tau \in \mathcal{A}} \Delta(\tau) \\ &\leq 2^{k-1} - \sum_{\tau \in \mathcal{A}} \Delta(\tau). \end{aligned}$$

For equality to hold in (6.15), we must have that $|\mathcal{A}| = 2^{k-1}$ and so \mathcal{A} is a decomposition consisting only of elements of weight 1. It follows that $\Delta(\tau) = 0$ and $x_\tau \leq 1$ for all $\tau \in \mathcal{A}$. If equality holds throughout the above, we then have that $x_\tau = 1$ for all $\tau \in \mathcal{A}$ and so $x \in O$. \square

7. ANALYTIC AND COMBINATORIAL STABILITY

In this section we prove Proposition 5.17 thus concluding our proof of Proposition 4.11. Note that Proposition 4.11 classifies the optimal points of X . We use compactness arguments to prove a result to the effect that ‘almost optimal’ points of X must be close (in ℓ_1 -norm) to a genuine optimal point of X . Furthermore, compactness allows us to derive similar properties for $X(\gamma)$ when γ is small. We then investigate what implications this has in our original combinatorial setting and complete the proof of Theorem 3.5.

Proof of Proposition 5.17. Let H be the matrix with rows and columns indexed by $\{0, 1, *\}^k$ where

$$H_{\sigma\tau} = \begin{cases} 1 & \text{if } \sigma, \tau \text{ are indistinguishable,} \\ 0 & \text{if } \sigma, \tau \text{ are distinguishable.} \end{cases}$$

Note that in particular, all diagonal entries of H are equal to 1. Let $w = (\omega(\tau) : \tau \in \{0, 1, *\}^k) \in \mathbb{R}^*$, then for $x \in \mathbb{R}^*$ we may write

$$F(x) = x^T Hx - w^T x.$$

Suppose now that $x \in X$ is an optimal point. By the proof of Lemma 5.8, there is a finite sequence $x = x_0, x_1, \dots, x_m$ of distinct optimal points of X where x_m is compressed, and for $i = 0, \dots, m-1$, $x_{i+1} = x_i(\pi_i, \rho_i)$ for some indistinguishable pair $\pi_i, \rho_i \in \text{supp}(x_i)$. Moreover we know that $\omega(\rho_i) \geq 2$ and that $\pi_i \notin \text{supp}(x_{i+1})$ for all i .

Suppose that x is not compressed so that $m \geq 1$. Let $y = x_{m-1}$, $z = x_m$ and let $\pi = \pi_{m-1}$, $\rho = \rho_{m-1}$. Since $z = y(\pi, \rho)$, it follows from the definition of compression that $z = y + \alpha(e_\rho - e_\pi)$ for some $\alpha > 0$. Let $p = e_\pi - e_\rho$. It follows, by the Taylor expansion of F , that

$$(7.1) \quad F(y) = F(z + \alpha p) = F(z) + \alpha p^T \nabla F(z) + \alpha^2 p^T H p.$$

Recall that $F(y) = F(z) = 0$ by Lemma 5.3. Furthermore by direct calculation we also have $p^T H p = 0$. It follows from (7.1) that $p^T \nabla F(z) = 0$ i.e.

$$(7.2) \quad \frac{\partial F}{\partial x_\pi}(z) = \frac{\partial F}{\partial x_\rho}(z).$$

Let I_ρ be the set of elements of $\{0, 1, *\}^k$ that are indistinguishable from ρ excluding ρ itself. Define I_π similarly. From the definition of F we have

$$(7.3) \quad \frac{\partial F}{\partial x_\rho}(z) = 2z_\rho + 2 \sum_{\tau \in I_\rho} z_\tau - \omega(\rho).$$

Since z is a compressed optimal point we have $z \in O$ by Corollary 5.16. Since $\omega(\rho) \geq 2$ and $\rho \in \text{supp}(z)$ we conclude that in fact $\omega(\rho) = 2$ and so $z_\rho = 2$. Moreover since $\text{supp}(z)$ is a distinguishable set we conclude that $z_\tau = 0$ for all $\tau \in I_\rho$. It follows from (7.3) that $\frac{\partial F}{\partial x_\rho}(z) = 2$ and hence from (7.2) that

$$(7.4) \quad \frac{\partial F}{\partial x_\pi}(z) = 2z_\pi + 2 \sum_{\tau \in I_\pi} z_\tau - \omega(\pi) = 2.$$

Since $z \in O$ we know that for all $\tau \in \text{supp}(z)$, $\omega(\tau) = 1$ or 2 and $z_\tau = \omega(\tau)$. Let w_1, w_2 be the number of elements of $I_\pi \cap \text{supp}(z)$ with weights 1, 2 respectively. Since $\pi \notin \text{supp}(z)$, we can then infer from (7.4) that

$$(7.5) \quad 2w_1 + 4w_2 - \omega(\pi) = 2.$$

We also know that $\text{supp}(z)$ is a decomposition and so

$$(7.6) \quad Q(\pi) = \bigcup_{\tau \in \text{supp}(z)} (Q(\tau) \cap Q(\pi)) = \bigcup_{\tau \in I_\pi \cap \text{supp}(z)} (Q(\tau) \cap Q(\pi)) \subseteq \bigcup_{\tau \in I_\pi \cap \text{supp}(z)} Q(\tau).$$

The second equality comes from the fact that $Q(\tau) \cap Q(\pi) = \emptyset$ whenever τ and π are distinguishable. Comparing the cardinality of the sets in (7.6) yields

$$(7.7) \quad 2^{\omega(\pi)} \leq \sum_{\tau \in I_\pi \cap \text{supp}(z)} 2^{\omega(\tau)} = 2w_1 + 4w_2.$$

Note also that $\rho \in I_\pi \cap \text{supp}(z)$ and $\omega(\rho) = 2$ so that $w_2 \geq 1$. Using (7.5), this last observation implies that $\omega(\pi) \geq 2$ whereas combining (7.5) and (7.7) we have

$$(7.8) \quad 2^{\omega(\pi)} - \omega(\pi) \leq 2w_1 + 4w_2 - \omega(\pi) = 2$$

We deduce that $\omega(\pi) = 2$ and so we have equality throughout (7.8), in particular we have equality in (7.7) and so also in (7.6). Note that $|Q(\pi)| = |Q(\rho)|$ since $\omega(\pi) = \omega(\rho) = 2$. Recalling that $\rho \in I_\pi \cap \text{supp}(z)$ equality in (7.6) would therefore imply that $Q(\pi) = Q(\rho)$ i.e. $\pi = \rho$. This is a contradiction and so x must be compressed. \square

Proposition 4.11 has the following corollary that says an almost optimal point of X must be close in norm to an actual optimal point of X .

Proposition 7.1. *Let $\eta \ll \varepsilon$. If $x \in X$ satisfies $\|x\| > 2^{k-1} - \eta$, then there exists an $x^* \in O$ such that $\|x - x^*\| < \varepsilon$.*

Proof. Consider the set

$$\tilde{X} := X \setminus \bigcup_{x^* \in O} B_\varepsilon(x^*).$$

\tilde{X} is compact and so $\sup_{z \in \tilde{X}} \|z\| = \|\tilde{x}\|$ for some $\tilde{x} \in \tilde{X}$. By the definition of \tilde{X} , $\tilde{x} \notin O$ and so by Proposition 4.11, $\|\tilde{x}\| = 2^{k-1} - \eta$ for some $\eta > 0$. It follows that if $x \in X$ satisfies $\|x\| > 2^{k-1} - \eta$ then $x \notin \tilde{X}$ and so $x \in B_\varepsilon(x^*)$ for some $x^* \in O$. \square

The following lemma allows us to relate properties of X and $X(\gamma)$ for γ small.

Lemma 7.2. *Let $\gamma \ll \eta$. If $x \in X(\gamma)$, then there exists $x_0 \in X$ for which $\|x - x_0\| < \eta$.*

Proof. Let $(\gamma_i)_{i \in \mathbb{N}}$ be a strictly decreasing sequence tending to 0, and let $X_i = X(\gamma_i)$ for $i \in \mathbb{N}$. Then X_1, X_2, \dots is a decreasing sequence of compact sets i.e. $X_{i+1} \subseteq X_i$ for $i \in \mathbb{N}$. Consider the set

$$U = \bigcup_{z \in X} B_\eta(z),$$

an open set containing X . Note that $(X_i \setminus U)_{i \in \mathbb{N}}$ is also a decreasing sequence of compact sets and also

$$\bigcap_{i=1}^{\infty} (X_i \setminus U) = \left(\bigcap_{i=1}^{\infty} X_i \right) \setminus U = X \setminus U = \emptyset.$$

By Cantor's Intersection Theorem (see [Rud76, Theorem 2.36, p.38]) it follows that $X_m \setminus U = \emptyset$ for some $m \in \mathbb{N}$. In other words, if $x \in X(\gamma_m)$ then $x \in U$ so that $x \in B_\eta(x_0)$ for some $x_0 \in X$. The result follows by taking $\gamma \leq \gamma_m$. \square

Corollary 7.3. *Let $\gamma \ll \varepsilon$. If $x \in X(\gamma)$ satisfies $\|x\| = 2^{k-1}$, then there exists an $x^* \in O$ such that $\|x - x^*\| < \varepsilon$.*

Proof. Given $\varepsilon > 0$, let $\eta = \min\{\eta_{7.1}(\varepsilon/2), \varepsilon/2\}$ and suppose that $\gamma \leq \gamma_{7.2}(\eta)$. Suppose that $x \in X(\gamma)$ satisfies $\|x\| = 2^{k-1}$. By Lemma 7.2 there exists an $x_0 \in X$ such that $\|x_0 - x\| < \eta$ and so $\|x_0\| > \|x\| - \eta = 2^{k-1} - \eta$. It follows from Proposition 7.1 that there exists an $x^* \in O$ such that $\|x_0 - x^*\| < \varepsilon/2$ and so

$$\|x - x^*\| \leq \|x - x_0\| + \|x_0 - x^*\| < \eta + \varepsilon/2 \leq \varepsilon. \quad \square$$

Let

$$O^* = \{x \in O : \omega(\tau) = 1 \text{ for all } \tau \in \text{supp}(x)\}.$$

In words, O^* is the set of all elements $x \in \mathbb{R}^*$ such that x is supported on a perfect matching of Q_k and all non-zero entries of x are equal to 1. We can also view O^* as the set of profiles of hypercube colourings normalised by clique size. Our aim is to use the stability-type statement of Corollary 7.3 to prove Theorem 3.5 in the following form.

Theorem 7.4. *Let $\frac{1}{t} \ll \delta \ll \varepsilon \ll 1$. If G is a $(1 - \delta)$ -dense, k -coloured graph with $v(G) = 2^{k-1}t$, containing no monochromatic odd connected matching of order $\geq (1 + \delta)t$, then for any choice of profile $x(G)$ of G , there exists some $x^* \in O^*$ such that*

$$\|x(G)/t - x^*\| < \varepsilon.$$

First we need the following two colour Ramsey result which is a direct consequence of the more general Theorem 1.8 in [BLS⁺12].

Lemma 7.5. *Let $\frac{1}{t} \ll \delta \ll \varepsilon \ll 1$. If H is a $(1 - \delta)$ -dense, 2-coloured graph with $v(H) \geq (\frac{3}{2} + \varepsilon)t$, then H contains a monochromatic cycle of length $\geq (1 + \delta)t$. \square*

Recall that by the definition of a connected matching, Definition 3.3, a cycle of length m is a connected matching of order $\geq m - 1$.

We remark that in the statements of Theorem 7.4 and Lemma 7.5, we do not assume that t is an integer.

Proof of Theorem 7.4. Given $0 < \varepsilon < 1/4$, let $\gamma = \gamma_{7.3}(\varepsilon)$ and $\delta' = \delta_{7.5}(\varepsilon)$. Suppose that $\delta < \min\{\gamma k^{-1}2^{-2k}, \delta'2^{-2k}\}$ and that $t \geq \max\{t_{7.5}(\delta'), \delta^{-1}\}$. Let G be a k -coloured graph as in the statement of Theorem 7.4. Let $x(G)$ be any choice of profile for G and let the corresponding profile partition be $(V_\tau : \tau \in \{0, 1, *\}^k)$. Note that $\|x(G)/t\| = 2^{k-1}$ and by Proposition 4.7, we have that $x(G)/t \in X(\delta k 2^{2k}) \subseteq X(\gamma)$. By Corollary 7.3 there exists an element $x^* \in O$ such that

$$(7.9) \quad \|x(G)/t - x^*\| < \varepsilon.$$

Suppose that $x^* \in O \setminus O^*$, then $x_\tau = 2$ for some $\tau \in \{0, 1, *\}^k$ such that $\omega(\tau) = 2$. It follows from (7.9) that $x(G)_\tau = |V_\tau| > (2 - \varepsilon)n \geq (3/2 + \varepsilon)n$. Let $H = G[V_\tau]$. By the definition of V_τ , H is a 2-coloured graph. Moreover since G has at most $\delta \binom{v(G)}{2} \leq \delta' \binom{v(H)}{2}$ edges missing, the same is true for H . It follows by Lemma 7.5 that H contains a monochromatic cycle of length $\geq (1 + \delta')t$ and so in particular H contains a monochromatic connected matching of order $> (1 + \delta)t$. However, by the definition of $V_\tau = V(H)$, any monochromatic component of H is contained in a non-bipartite monochromatic component of G . Thus G contains a monochromatic odd connected matching of order $> (1 + \delta)t$ contrary to assumption. We conclude that $x^* \in O^*$. \square

8. THE REGULARITY METHOD

In this section we discuss the tools and results we need from the regularity method. Our starting point is Szemerédi's Regularity Lemma [Sze78] which we discuss briefly now.

Let G be a graph and let A, B be disjoint subsets of $V(G)$. We call

$$d_G(A, B) := \frac{e_G(A, B)}{|A||B|}$$

the *density* of the pair (A, B) . For $\delta > 0$, we say that the pair (A, B) is δ -*regular with respect to* G if, for every $A' \subseteq A$ and $B' \subseteq B$ satisfying $|A'| \geq \delta|A|$ and $|B'| \geq \delta|B|$, we have

$$|d_G(A', B') - d_G(A, B)| < \delta.$$

If, for $d > 0$, we also have that $|N_G(a) \cap B| \geq d|B|$ for all $a \in A$ and $|N_G(b) \cap A| \geq d|A|$ for all $b \in B$, then we say that (A, B) is (δ, d) -*super-regular with respect to* G . If the graph G is clear from the context we may omit it from the above notation. The following is a version of Szemerédi's Regularity Lemma that appears as Theorem 1.18 in [KS96]. Recall that given a k -coloured graph G and $i \in [k]$, we let G_i denote the i th colour class of G .

Theorem 8.1 (Multicolour Regularity Lemma). *For all $\delta > 0$ and $k, \ell \in \mathbb{N}$ there exists $L = L(\delta, k, \ell)$ and $M = M(\delta, k, \ell)$ such that the following holds. For all k -coloured graphs G on at least M vertices, $V(G)$ may be partitioned into sets V_0, V_1, \dots, V_t such that*

- $\ell \leq t \leq L$;
- $|V_0| < \delta v(G)$ and $|V_1| = |V_2| = \dots = |V_t|$;
- apart from at most $\delta \binom{t}{2}$ exceptional pairs, the pairs (V_i, V_j) , $1 \leq i < j \leq t$, are δ -regular with respect to G_s for $s = 1, \dots, k$.

We now state some technical lemmas related to Łuczak's method of connected matchings. First we need a definition. (It might be useful at this point to recall the definition of a connected matching of order m , Definition 3.3.)

Definition 8.2. Let $\delta, d \in [0, 1]$ and $q, m \geq 1$ be integers.

- Let F be a graph on vertex set $[q]$ and let U_1, \dots, U_q be disjoint sets of size m . We call a graph H on vertex set $\bigcup_{i \in [q]} U_i$ a (δ, m) -*regular blow-up* of F if whenever $\{i, j\} \in E(F)$, we have that (U_i, U_j) is a δ -regular pair.
- If in addition to the above, $d(U_i, U_j) \geq d$ for each edge $\{i, j\}$ of F then we say that H has *minimum density* d .
- Suppose that F is a connected matching and H is a (δ, m) -regular blow-up of F with minimum density d . If for each matching edge $\{i, j\}$ of F , the pair (U_i, U_j) is in fact (δ, d) -super-regular in H , then we say that H is a (δ, d, m) -*super-regular blow-up* of F .

Versions of the following two lemmas abound in the literature (e.g. [KSS05], [Łuc99]), but here we give statements tailored to our needs. However since they are not new, we defer their proofs to the Appendix.

Lemma 8.3. *Let $q \geq 4$ and suppose that $\frac{1}{m} \ll \delta \ll d$. Let F be a connected matching of order q such that every vertex of F is incident to a matching edge and let H be a (δ, d, m) -super-regular blow-up of F . Then the following holds:*

If $i, j \in V(F)$ and there is an ij -path of length r in F , then for every pair of vertices $u \in U_i, w \in U_j$, there exists a uw -path of length ℓ in H for each $3q \leq \ell \leq (1 - 6\delta)qm$ such that $\ell \equiv r \pmod{2}$.

Lemma 8.4. Let $q \geq 4$ and let $\frac{1}{m} \ll \delta \ll d$. Let F be an odd connected matching of order q and suppose that H is a (δ, m) -regular blow-up of F with minimum density d . Then H contains a cycle of length ℓ for each odd $3q \leq \ell \leq (1 - 6\delta)qm$.

We borrow the following fact.

Fact 8.5. ([GRSS07, Lemma 9]). Let H be a $(1 - \delta)$ -dense graph on t vertices. Then H has a subgraph H' such that $v(H') \geq (1 - \sqrt{\delta})t$ and $\delta(H') \geq (1 - 2\sqrt{\delta})t$.

We will also need the following two standard facts whose proofs we omit.

Fact 8.6. Let $0 < \delta \leq 1/2$ and let (A, B) be a δ -regular pair with density d . Suppose that $A' \subseteq A, B' \subseteq B$ such that $|A'| \geq (1 - \delta)|A|, |B'| \geq (1 - \delta)|B|$. Then (A', B') is 2δ -regular with density $d' > d - \delta$. Moreover, if (A, B) is in fact (δ, β) -super-regular for some $\beta > 0$, then (A', B') is $(2\delta, \beta - \delta)$ -super-regular. \square

Fact 8.7. Let $0 < \delta \leq 1/2$ and let (A, B) be a δ -regular pair with density d . Then there exist $A' \subseteq A, B' \subseteq B$ such that $|A'| = (1 - \delta)|A|, |B'| = (1 - \delta)|B|$ and (A', B') is $(2\delta, d - 2\delta)$ -super-regular. \square

9. PROOF OF MAIN THEOREM

In this final section we prove our main result Theorem 3.2, and therefore also Theorem 1.2. The idea is to invoke Theorem 7.4 to show that the profile of a certain reduced graph is close in ℓ_1 -norm to the profile of a hypercube colouring (we refer the reader back to Section 3.2 for the definition of a hypercube colouring). We then translate this information to show that the original graph is close in edit distance to a hypercube colouring.

The stability-type methods of this section require some care due to the plethora of extremal constructions. Luckily hypercube colourings share enough common features for these methods to be viable and surprisingly we require no case analysis.

First let us state a result which is a corollary of a classical theorem of Bondy [Bon71].

Theorem 9.1. Let G be a graph on at least three vertices with minimum degree $> v(G)/2$, then G is pancyclic i.e. G contains cycles of all lengths $3 \leq \ell \leq v(G)$.

Proof of Theorem 3.2. Let $0 < \varepsilon < 2^{-4k}$, let

$$(9.1) \quad \eta < \delta \leq \min \left\{ \frac{1}{9} \delta_{7.4}^2(\varepsilon), \delta_{8.3} \left(\frac{1}{k} \right), \delta_{8.4} \left(\frac{1}{k} \right) \right\},$$

let $t_0 \geq \max\{t_{7.4}(\delta), \delta^{-1/2}\}$ and let $L = L_{8.1}(\delta, k, 2^k t_0)$. Let n be odd with

$$(9.2) \quad n \geq \max\{Lm_{8.3}(\delta), Lm_{8.4}(\delta), M_{8.1}(\delta, k, 2^k t_0)\}.$$

Finally let G be a k -coloured copy of K_N where $N > (2^{k-1} - \eta)n$ and assume that

$$(\dagger) \quad G \text{ contains no monochromatic copy of } C_n.$$

Applying Theorem 8.1 to G we obtain a partition of $V(G)$ into sets V_0, V_1, \dots, V_{t_1} such that

$$(i) \quad 2^k t_0 \leq t_1 \leq L;$$

- (ii) $|V_0| < \delta N$ and $|V_1| = |V_2| = \dots = |V_{t_1}|$;
- (iii) apart from at most $\delta \binom{t_1}{2}$ exceptional pairs, the pairs (V_i, V_j) , $1 \leq i < j \leq t_1$, are δ -regular with respect to G_s for $s = 1, \dots, k$.

It follows that for $i \in [t_1]$,

$$(9.3) \quad m := |V_i| \geq \frac{(1 - \delta)N}{t_1}.$$

We construct a *reduced graph* R_1 with vertex set $[t_1]$ and edge set formed by pairs $\{u, w\}$ for which (V_u, V_w) is δ -regular with respect to G_i for $i = 1, \dots, k$. It follows from (iii) of the above that R_1 is $(1 - \delta)$ -dense. Fact 8.5 allows us to find a subgraph $R \subseteq R_1$ satisfying $v(R) \geq (1 - \sqrt{\delta})t_1$ and $\delta(R) \geq (1 - 2\sqrt{\delta})t_1$. Let $t = v(R)$ and assume without loss of generality that $V(R) = [t]$. We k -colour R by colouring an edge $\{u, w\}$ with the least colour i for which

$$(9.4) \quad d_{G_i}(V_u, V_w) \geq \frac{1}{k}.$$

Let $t' = t/2^{k-1}$ and note that by (9.3) and the definition of t ,

$$(9.5) \quad mt' \geq (1 - 2\sqrt{\delta})n.$$

Suppose that R contains a monochromatic odd connected matching F of order $q \geq (1 + 3\sqrt{\delta})t'$. Then G contains a monochromatic (δ, m) -regular blow-up of F with minimum density d for some $d \geq 1/k$ by (9.4). Note that since $t_1 \leq L$ we have $m \geq m_{8.4}(\delta)$ by (9.3) and (9.2). It follows from Lemma 8.4 that G contains a monochromatic copy of C_n since n is odd and

$$3q \leq 3L \leq n \leq (1 - 6\delta)(1 + 3\sqrt{\delta})mt' \leq (1 - 6\delta)qm,$$

contradicting (\dagger) . We conclude that R contains no such odd connected matching. Let $(W_\tau : \tau \in \{0, 1, *\}^k)$ be a profile partition of R (as defined in Section 3.1) and let $x(R) = (|W_\tau| : \tau \in \{0, 1, *\}^k)$ be the corresponding profile. It follows by Theorem 7.4 that there exists $x^* \in O^*$ such that

$$(9.6) \quad \|x(R)/t' - x^*\| < \varepsilon.$$

This tells us a lot about the structure of R , indeed it is ‘close to’ a hypercube colouring. In particular, we know that R can be covered by 2^{k-1} monochromatic cliques of approximately equal size and a small leftover set. Our next goal is to establish some more detailed properties of R and to show that these properties transfer to G . In particular we will show that, like R , G can be covered by 2^{k-1} monochromatic cliques of approximately equal size and a small leftover set (Lemma 9.9). We then prove an ‘absorbing lemma’ (Lemma 9.12) allowing us to deal with the leftover set in G .

By the definition of O^* we have that

$$\text{supp}(x^*) = \mathcal{M} \subseteq \{0, 1, *\}^k,$$

for some perfect matching \mathcal{M} of the hypercube Q_k and $x_\tau^* = 1$ for all $\tau \in \mathcal{M}$. Let

$$W = R \setminus \bigcup_{\tau \in \mathcal{M}} W_\tau.$$

We will treat W as a ‘leftover set’ of vertices of R and study only the structure of $R \setminus W$. Note that by (9.6) we have

$$(9.7) \quad (1 - \varepsilon)t' < |W_\tau| < (1 + \varepsilon)t' \text{ for all } \tau \in \mathcal{M},$$

and so by removing at most $2\varepsilon t'$ vertices from each part W_τ , where $\tau \in \mathcal{M}$, and absorbing these removed vertices into W , we may assume that these parts W_τ all have the same size $> (1 - \varepsilon)t'$. Note that even after this absorption we have

$$|W| = t - \sum_{\tau \in \mathcal{M}} |W_\tau| < t - 2^{k-1}(1 - \varepsilon)t' = \varepsilon t.$$

We make a couple of observations regarding the colouring of R with respect to these vertex classes. For $j \in [k]$ we let

$$(9.8) \quad I_j = \{\tau \in \mathcal{M} : \tau_j = *\}.$$

Lemma 9.2. *Let $\tau \in I_j$. Then $R[W_\tau]$ is monochromatic in the colour j and has minimum degree at least $(1 - 2^{k+1}\sqrt{\delta})|W_\tau|$.*

Proof. By the definition of the profile partition, for each colour $i \neq j$, each pair $v, w \in W_\tau$ must lie in the same vertex class in an induced bipartite subgraph of R_i . It follows that if $\{v, w\} \in E(R)$ then it cannot receive the colour i and hence must receive colour j . Since $\delta(R) \geq (1 - 2\sqrt{\delta})t$ we have

$$\delta(R[W_\tau]) \geq |W_\tau| - 1 - 2\sqrt{\delta}t \geq (1 - 2^{k+1}\sqrt{\delta})|W_\tau|$$

where for the last inequality we used (9.7). □

Definition 9.3. For $\sigma, \tau \in \{0, 1, *\}^k$, we let $\Delta(\sigma, \tau) = \{i \in [k] : \{\sigma_i, \tau_i\} = \{0, 1\}\}$.

Lemma 9.4. *Let $\sigma, \tau \in \mathcal{M}$ be distinct, then*

- (i) *Each edge of $R[W_\sigma, W_\tau]$ receives a colour from the set $\Delta(\sigma, \tau)$;*
- (ii) *$R[W_\sigma, W_\tau]$ has minimum degree $\geq (1 - 2^{k+1}\sqrt{\delta})|W_\sigma|$, in particular $R[W_\sigma, W_\tau]$ is connected and contains a perfect matching.*

Proof. Let $\sigma \in I_j, \tau \in I_\ell$. Suppose that $j \neq \ell$ so that $\{\sigma_i, \tau_i\} \neq \{*\}$ for each $i \in [k]$. By the definition of the profile partition, for each colour $i \notin \Delta(\sigma, \tau)$, each pair $v \in W_\sigma, w \in W_\tau$ must lie in either the same vertex class in an induced bipartite subgraph of R_i or they lie in different connected components of R_i . It follows that if $\{v, w\} \in E(R)$ then it must receive a colour from $\Delta(\sigma, \tau)$. Similarly, if $j = \ell$ then each edge of $R[W_\sigma, W_\tau]$ must receive a colour from $\Delta(\sigma, \tau) \cup \{j\}$. However, by Lemma 9.2, $R[W_\tau]$ and $R[W_\sigma]$ are both monochromatic in the colour j and both have minimum degree at least $(1 - 2^{k+1}\sqrt{\delta})|W_\sigma| > |W_\sigma|/2$ (recall that $|W_\sigma| = |W_\tau|$). It follows, by Theorem 9.1 for example, that $R[W_\tau]$ and $R[W_\sigma]$ are both Hamiltonian and non-bipartite. Using (9.7), we deduce that if an edge of $R[W_\sigma, W_\tau]$ receives the colour j , then R contains a monochromatic odd connected matching in the colour j of order at least

$$|W_\sigma| + |W_\tau| - 2 \geq 2(1 - \varepsilon)t' - 2 \geq (1 + 3\sqrt{\delta})t',$$

which we showed previously was not the case. Part (i) of the lemma follows. If $v \in W_\tau$ then, since $\delta(R) \geq (1 - 2\sqrt{\delta})t$, we have

$$|N(v) \cap W_\sigma| \geq |W_\sigma| - 2\sqrt{\delta}t \geq (1 - 2^{k+1}\sqrt{\delta})|W_\sigma|.$$

Similarly if $w \in W_\sigma$, then $|N(w) \cap W_\tau| \geq (1 - 2^{k+1}\sqrt{\delta})|W_\sigma|$. Since $1 - 2^{k+1}\sqrt{\delta} > 1/2$, it follows that $R[W_\sigma, W_\tau]$ is connected and (e.g. by Hall's theorem) contains a perfect matching. \square

Let Γ denote the k -coloured multigraph on vertex set \mathcal{M} where we have an edge between σ and τ in each colour j for which $R[W_\sigma, W_\tau]$ contains an edge of colour j . Note that since $\delta(R) > (1 - 2\sqrt{\delta})t$ and $|W_\tau| > (1 - \varepsilon)t' > 2\sqrt{\delta}t$ for all $\tau \in \mathcal{M}$, $R[W_\sigma, W_\tau]$ contains an edge for every $\sigma, \tau \in \mathcal{M}$. Let Γ^* denote the subgraph of Γ where we keep only those edges that occur as the unique edge between a given pair of vertices in Γ .

Lemma 9.5. *For each $j \in [k]$, the vertices of Γ_j^* can be covered by a matching $T_j \subseteq \Gamma_j^*$ and the set I_j . Moreover I_j is a set of isolated vertices in Γ_j .*

Proof. Fix $j \in [k]$. If $\sigma \in I_j$ then σ is an isolated vertex in Γ_j by Lemma 9.4(i). If $\sigma \notin I_j$ then we may assume without loss of generality that $\sigma_j = 0$. Let σ' be the element of $\{0, 1, *\}^k$ such that $\sigma'_j = 1$ and $\sigma'_i = \sigma_i$ for all $i \neq j$. Let H be the graph on \mathcal{M} with edge set $\{\{\rho, \pi\} : \Delta(\rho, \pi) = \{j\}\}$. By Lemma 9.4(i) we have $H \subseteq \Gamma_j^*$. The neighbours of σ in H are precisely those elements of \mathcal{M} that are indistinguishable from σ' i.e. those elements of \mathcal{M} (viewed as edges of Q_k) that intersect $Q(\sigma')$. If $\sigma' \in \mathcal{M}$, then since \mathcal{M} is a perfect matching, we must have that σ' is the unique neighbour of σ in H and similarly σ is the unique neighbour of σ' (i.e. they form an isolated edge in H). If $\sigma' \notin \mathcal{M}$, then precisely two elements of \mathcal{M} intersect $Q(\sigma')$ (i.e. the degree of σ in H is two). It follows that H is the disjoint union of cycles (where we consider an edge a cycle) and the independent set I_j . Since H is bipartite with bipartition $\{\tau \in \mathcal{M} : \tau_j = 0 \text{ or } *\} \cup \{\tau \in \mathcal{M} : \tau_j = 1\}$, the cycles in H are all even. The result follows. \square

Let $j \in [k]$, then for each $\{\sigma, \tau\} \in T_j$ (T_j as in the statement of Lemma 9.5), we may fix a monochromatic perfect matching $M_{\sigma\tau}^j$ in the colour j in $R[W_\sigma, W_\tau]$ by Lemmas 9.4(ii) and 9.5. Let

$$\mathcal{T}_j = \bigcup_{\{\sigma, \tau\} \in T_j} M_{\sigma\tau}^j.$$

and note that \mathcal{T}_j is a matching in R , monochromatic in the colour j , which covers the vertex set $\bigcup_{\tau \notin I_j} W_\tau$. The following corollary hints at an important common feature of all hypercube colourings.

Corollary 9.6. *Given $j \in [k]$ and $\rho \in \mathcal{M} \setminus I_j$, there exists $\pi \in \mathcal{M}$ such that $R[W_\rho, W_\pi]$ contains a monochromatic connected perfect matching in the colour j whose matching edges are edges of \mathcal{T}_j .*

Proof. Since $\rho \notin I_j$, by Lemma 9.5 there must exist $\pi \in \mathcal{M}$ such that $\{\rho, \pi\}$ is an edge of $T_j \subseteq \Gamma_j^*$. By the definition of Γ^* , we have that $R[W_\rho, W_\pi]$ is monochromatic in the colour j . The result follows from the definition of \mathcal{T}_j and Lemma 9.4(ii). \square

It will be useful to prune the sets V_i for $i \in R$ in such a way that if $\{x, y\}$ is an edge of the matching \mathcal{T}_j then $G_j[V_x, V_y]$ is super-regular.

Lemma 9.7. *For each $i \in R$ there exists $V'_i \subseteq V_i$ such that*

- (i) $|V'_i| = (1 - 2^k\delta)m$ for all $i \in R$,
- (ii) $G_j[V'_x, V'_y]$ is $2^{k+1}\delta$ -regular with density $\geq \frac{1}{k+1}$ for all $j \in [k], \{x, y\} \in R_j$,

(iii) $G_j[V'_x, V'_y]$ is $(2^{k+1}\delta, \frac{1}{k+1})$ -super-regular for all $j \in [k], \{x, y\} \in \mathcal{T}_j$.

Proof. For $i \in R$ we define a sequence of subsets V_i^0, \dots, V_i^k of V_i recursively. Let $V_i^0 := V_i$ for all $i \in R$. Suppose that for all $i \in R$ we have found $V_i^\ell \subseteq V_i$ with the following properties.

- (a) $|V_i^\ell| \geq (1 - 2^\ell\delta)m$ for all $i \in R$,
- (b) $G_j[V_x^\ell, V_y^\ell]$ is $2^\ell\delta$ -regular with density $\geq 1/k - 2^\ell\delta$ for all $j \in [k], \{x, y\} \in R_j$,
- (c) $G_j[V_x^\ell, V_y^\ell]$ is $(2^\ell\delta, 1/k - 2^{\ell+1}\delta)$ -super-regular for all $j \in [\ell], \{x, y\} \in \mathcal{T}_j$.

By Fact 8.7, for each edge $\{u, w\}$ in the matching $\mathcal{T}_{\ell+1}$ (so in particular $\{u, w\} \in R_{\ell+1}$) there exists $V_u^{\ell+1} \subseteq V_u^\ell$ and $V_w^{\ell+1} \subseteq V_w^\ell$ such that $|V_u^{\ell+1}| = (1 - 2^\ell\delta)|V_u^\ell|$, $|V_w^{\ell+1}| = (1 - 2^\ell\delta)|V_w^\ell|$ and $G_{\ell+1}[V_u^{\ell+1}, V_w^{\ell+1}]$ is $(2^{\ell+1}\delta, 1/k - 2^{\ell+2}\delta)$ -super-regular. If i is not incident to any edge of $\mathcal{T}_{\ell+1}$ then simply set $V_i^{\ell+1} = V_i^\ell$. Note that by (a), for all $i \in R$,

$$|V_i^{\ell+1}| \geq (1 - 2^\ell\delta)|V_i^\ell| \geq (1 - 2^\ell\delta)^2m \geq (1 - 2^{\ell+1}\delta)m.$$

For $j \in [k]$ and $\{x, y\} \in R_j$, by (b) and Fact 8.6 we have that $G_j[V_x^{\ell+1}, V_y^{\ell+1}]$ is $2^{\ell+1}\delta$ -regular with density $\geq 1/k - 2^{\ell+1}\delta$. Using (c) and Fact 8.6, it also follows that $G_j[V_x^{\ell+1}, V_y^{\ell+1}]$ is $(2^{\ell+1}\delta, 1/k - 2^{\ell+2}\delta)$ -super-regular for all $j \in [\ell], \{x, y\} \in \mathcal{T}_j$. We have shown that the sets $V_i^{\ell+1}, i \in R$, satisfy (a)-(c) (with ℓ replaced by $\ell + 1$). The result follows by letting V'_i be any subset of V_i^k of size $(1 - 2^k\delta)m$ for all $i \in R$ and appealing to Fact 8.6, noting that $1/(k+1) \leq 1/k - 2^{k+2}\delta$. □

Given $\sigma \in \mathcal{M}$, let

$$\widetilde{W}_\sigma = \bigcup_{i \in W_\sigma} V'_i \subseteq V(G),$$

and let

$$\widetilde{W} = V(G) \setminus \bigcup_{\tau \in \mathcal{M}} \widetilde{W}_\tau.$$

As with W , we think of \widetilde{W} as a small leftover set of vertices. Let $m' := (1 - 2^k\delta)m$ and note that by (9.5), $m't' \geq (1 - 3\sqrt{\delta})n$ and so by (9.7)

$$(9.9) \quad |\widetilde{W}_\tau| \geq (1 - \varepsilon)m't' \geq (1 - 2\varepsilon)n \text{ for all } \tau \in \mathcal{M}.$$

We also have

$$(9.10) \quad |\widetilde{W}| \leq N - 2^{k-1}(1 - \varepsilon)m't' = N - (1 - \varepsilon)(1 - 2^k\delta)mt \leq 2\varepsilon N.$$

Where for the last inequality we recalled (9.3).

We can now establish our first piece of structure on the graph G . We show that almost all of $V(G)$ can be covered by 2^{k-1} monochromatic cliques of equal size. First we make a quick definition.

Definition 9.8. If $\tau \in \{0, 1, *\}^k$ has weight 1, we let $c(\tau)$ denote the unique element of $i \in [k]$ for which $\tau_i = *$.

We note that for $\tau \in \mathcal{M}$ and $j \in [k]$, we have $c(\tau) = j$ if and only if $\tau \in I_j$ (as defined in (9.8)).

Lemma 9.9. For all $\sigma \in \mathcal{M}$, $G[\widetilde{W}_\sigma]$ is monochromatic in the colour $c(\sigma)$.

Proof. Suppose that $G[\widetilde{W}_\sigma]$ contains an edge $\{x, y\}$ of colour $j \neq c(\sigma)$ (so that $\sigma \notin I_j$). By Corollary 9.6 there exists $\tau \in \mathcal{M}$ such that $R_j[W_\sigma, W_\tau]$ contains a connected perfect matching, F , whose matching edges are edges of \mathcal{T}_j . Let $q := v(F)$, then by (9.7) we have $2(1 - \varepsilon)t' \leq q \leq 2(1 + \varepsilon)t'$. By Lemma 9.7 we see that $G_j[\widetilde{W}_\sigma, \widetilde{W}_\tau]$ contains a spanning $(2^{k+1}\delta, 1/(k+1), m')$ -super-regular blow-up of F (with the V_i' playing the role of the U_i in Definition 8.2). Suppose that $x \in V_a'$ and $y \in V_b'$, then a and b lie on the same side of the bipartition of the connected graph F and so F contains an ab -path of even length. Note that $m' \geq n_{8.3}(\delta)$, by (9.3) and (9.2). We may therefore apply Lemma 8.3 to deduce that $G_j[\widetilde{W}_\sigma, \widetilde{W}_\tau]$ contains a path of length $n - 1$ joining x and y since $n - 1$ is even and

$$(9.11) \quad (1 - 6 \cdot 2^k \delta)qm' \geq 2(1 - 6 \cdot 2^k \delta)(1 - \varepsilon)m't' \geq n - 1 \geq 3L \geq 3q,$$

where we used (9.1), (9.2) and (9.9). Together with the edge $\{x, y\}$ this creates a monochromatic copy of C_n in G , contrary to assumption (\dagger) . \square

Our aim now is to say something about the edges of G lying between \widetilde{W} and the rest of the graph (see Lemma 9.12 below). Recall that Γ is the k -coloured multigraph on vertex set \mathcal{M} where we have an edge between σ and τ in each colour j for which $R[W_\sigma, W_\tau]$ contains an edge of colour j . With this graph in mind, we make the following definition.

Definition 9.10. Let \mathcal{M}' be a perfect matching of Q_k and let $\varphi : \mathcal{M} \rightarrow \mathcal{M}'$ be a bijection such that $c(\varphi(\tau)) = c(\tau)$ for all $\tau \in \mathcal{M}$. Suppose that Ψ is a k -coloured multigraph on vertex set \mathcal{M} . We call φ an *admissible labelling* of Ψ if for all $\sigma, \tau \in \mathcal{M}$, the edges between σ, τ in Ψ only take colours from the set $\Delta(\varphi(\sigma), \varphi(\tau))$.

Note that by Lemma 9.4(i) the identity map $\iota : \mathcal{M} \rightarrow \mathcal{M}$ is an admissible labelling of Γ . The following lemma gives a useful way of generating new admissible labellings of Γ . For $\tau \in \{0, 1, *\}^k$, such that $\tau_j \in \{0, 1\}$, we let $\tau^j := (\tau_1, \dots, \tau_{j-1}, 1 - \tau_j, \tau_{j+1}, \dots, \tau_k)$ i.e. τ^j denotes τ with the j th coordinate flipped.

Lemma 9.11. *Let φ be an admissible labelling of Γ . Let $j \in [k]$ and let C be the vertex set of a component of Γ_j such that $\tau_j \neq *$ for all $\tau \in C$. Let φ' be the function on \mathcal{M} given by $\varphi'(\tau) = \varphi(\tau)^j$ for all $\tau \in C$, $\varphi'(\tau) = \varphi(\tau)$ otherwise. Then φ' is an admissible labelling of Γ .*

Proof. Let us first check that the image of φ' is a perfect matching of Q_k . First note that by the definition of φ' and the fact that φ is admissible, each element of $\varphi'(\mathcal{M})$ has weight 1 and $c(\varphi'(\tau)) = c(\varphi(\tau)) = c(\tau)$ for all $\tau \in \mathcal{M}$. To show that $\varphi'(\mathcal{M})$ is a perfect matching, it thus suffices to show that if $\sigma, \tau \in \mathcal{M}$ are distinct, then $\varphi'(\sigma), \varphi'(\tau)$ are disjoint (i.e. $\Delta(\varphi'(\sigma), \varphi'(\tau)) \neq \emptyset$). We do this by considering an edge between σ and τ in Γ and showing that if it has the colour i then $i \in \Delta(\varphi'(\sigma), \varphi'(\tau))$. Note that this in fact suffices to show that φ' is admissible.

Suppose then that there is an edge between σ, τ in Γ in the colour i . Since φ is admissible we have $i \in \Delta(\varphi(\sigma), \varphi(\tau))$. Suppose that $i \neq j$ then by the definition of φ' , $\varphi'(\tau)_i = \varphi(\tau)_i$ and $\varphi'(\sigma)_i = \varphi(\sigma)_i$ and so $i \in \Delta(\varphi'(\sigma), \varphi'(\tau))$ also. Suppose then that $i = j$, so that either $\sigma, \tau \in C$ or $\sigma, \tau \in \mathcal{M} \setminus C$. If $\sigma, \tau \in C$, then $\varphi'(\tau)_i = 1 - \varphi(\tau)_i$, $\varphi'(\sigma)_i = 1 - \varphi(\sigma)_i$ and if $\sigma, \tau \in \mathcal{M} \setminus C$, then $\varphi'(\tau)_i = \varphi(\tau)_i$, $\varphi'(\sigma)_i = \varphi(\sigma)_i$. In either case $i \in \Delta(\varphi'(\sigma), \varphi'(\tau))$. \square

We are now in a position to prove our ‘absorbing lemma’ which associates each vertex in \widetilde{W} to some class \widetilde{W}_σ in G .

Lemma 9.12. *Let $v \in \widetilde{W}$. Then there exists $\sigma \in \mathcal{M}$ such that $G[v, \widetilde{W}_\sigma]$ is monochromatic in the colour $c(\sigma)$.*

Proof. Suppose otherwise, then for each $\sigma \in \mathcal{M}$ there exists a $u \in \widetilde{W}_\sigma$ such that the edge $\{v, u\}$ receives a colour $j_\sigma \neq c(\sigma)$. We augment the multigraph Γ in the following way. We add the vertex v to Γ and for each $\sigma \in \mathcal{M}$ we add an edge between v and σ in the colour j_σ . Let us call this augmented multigraph Γ^+ .

Claim 9.13. Γ^+ contains a monochromatic odd cycle.

Proof of Claim. Suppose otherwise and choose an admissible labelling φ of Γ that minimises the function

$$S(\varphi) = \sum_{j \in [k]} |\{\tau \in \mathcal{M} : j_\tau = j \text{ and } \varphi(\tau)_i = 1\}|.$$

Suppose that $S(\varphi) > 0$, then there exists a colour $j \in [k]$ and an element $\sigma \in \mathcal{M}$ for which $j_\sigma = j$ and $\varphi(\sigma)_j = 1$. Let C denote the component of Γ_j containing the vertex σ and note that by the definition of admissibility C is bipartite with parts $\{\tau \in C : \varphi(\tau)_j = 0\}$ and $\{\tau \in C : \varphi(\tau)_j = 1\}$. Note that since C is connected in Γ_j this is the unique bipartition of C . Since Γ_j^+ is bipartite by assumption we must therefore have that $\varphi(\tau)_j = 1$ for all $\tau \in C$ such that $j_\tau = j$. Let φ' denote the function on \mathcal{M} given by $\varphi'(\tau) = \varphi(\tau)^j$ for all $\tau \in C$, $\varphi'(\tau) = \varphi(\tau)$ otherwise. By Lemma 9.11, φ' is an admissible labelling of Γ , however $S(\varphi') < S(\varphi)$ contradicting the minimality of φ . We conclude that $S(\varphi) = 0$ i.e.

$$(9.12) \quad \text{For all } i \in [k], \tau \in \mathcal{M}, \text{ if } j_\tau = i \text{ then } \varphi(\tau)_i = 0.$$

Since $\varphi(\mathcal{M})$ is a perfect matching of Q_k , there must exist $\rho \in \mathcal{M}$ such that the edge $\varphi(\rho)$ is incident to the vertex $(1, 1, \dots, 1)$ (formally $Q(\varphi(\rho))$ contains $(1, 1, \dots, 1)$). Without loss of generality suppose $\varphi(\rho) = (*, 1, \dots, 1)$. However, whatever value j_ρ takes, we contradict (9.12). This concludes the proof of the claim. \square

Suppose that Γ^+ contains a monochromatic odd cycle in the colour j . Since Γ_j is bipartite and Γ_j^+ is not, there must exist $\sigma, \tau \in \mathcal{M}$ such that σ, τ lie in opposite parts of the bipartition of a connected component in Γ_j and the edges $\{v, \sigma\}, \{v, \tau\}$ both have colour j in Γ^+ . By the definition of Γ^+ , there exist vertices $u \in \widetilde{W}_\sigma, w \in \widetilde{W}_\tau$ such that $\{v, u\}$ and $\{v, w\}$ both have colour j in G and $j \notin \{c(\sigma), c(\tau)\}$ i.e. $\sigma, \tau \notin I_j$. Suppose that $u \in V'_a$ and $w \in V'_b$ then by Lemmas 9.4(ii) and 9.5 and the definition of \mathcal{T}_j , a and b lie in opposite parts of a bipartite connected matching, F , in R_j whose matching edges span F and are edges of \mathcal{T}_j (in particular there is an ab -path of odd length in F). Moreover we may assume that F spans the vertex sets W_σ, W_τ in R_j and so by (9.7), $2(1 - \varepsilon)t' \leq v(F) \leq v(R) \leq L$. By Lemma 9.7, we have a $(2^{k+1}\delta, 1/(k+1), m')$ -super-regular blow-up of F in G_j . By Lemma 8.3 (using inequalities as in (9.11) and noting that $n - 2$ is odd) there exists a path of length $n - 2$ joining u and w in G_j . This together with the edges $\{v, u\}, \{v, w\}$ forms a monochromatic copy of C_n in G contrary to assumption (\dagger). This concludes the proof of Lemma 9.12. \square

Using Lemma 9.12 we may define a function $f : \widetilde{W} \rightarrow \mathcal{M}$ where $f(v)$ is an element of \mathcal{M} such that $G[v, \widetilde{W}_{f(v)}]$ is monochromatic in the colour $c(f(v))$. For each $\tau \in \mathcal{M}$, let

$U_\tau = \widetilde{W}_\tau \cup f^{-1}(\{\tau\})$. By (9.9), (9.10), (9.3) and Lemma 9.9 we have that

$$(9.13) \quad \delta(G_{c(\tau)}[U_\tau]) \geq (1 - 2^{k+1}\varepsilon)|U_\tau| \text{ for all } \tau \in \mathcal{M}.$$

Note that the sets U_τ , $\tau \in \mathcal{M}$, partition the vertex set of G and so if $N \geq 2^{k-1}(n-1) + 1$ then by the pigeonhole principle there exists $\sigma \in \mathcal{M}$ such that $|U_\sigma| \geq n$. However, by (9.13) and Theorem 9.1, it follows that U_σ contains a monochromatic copy of C_n in the colour $c(\sigma)$, contrary to assumption (\dagger) . We therefore have that $N \leq 2^{k-1}(n-1)$. Note that at this point we have done enough to prove Theorem 1.2.

It remains to show that G is close in edit distance to a hypercube colouring. Recall that $|\widetilde{W}| \leq 2\varepsilon N$ and so there are at most $2\varepsilon N^2$ edges of G incident to \widetilde{W} . We now aim to show that $G \setminus \widetilde{W}$ is close to a hypercube colouring. Recall that we have partitioned the vertex set of $G \setminus \widetilde{W}$ into the monochromatic, equally sized cliques $\{\widetilde{W}_\tau : \tau \in \mathcal{M}\}$. For $\sigma \in \mathcal{M}$, we showed that $|\widetilde{W}_\sigma| \geq (1 - 2\varepsilon)n$ and \widetilde{W}_σ is monochromatic in the colour $c(\sigma)$. First note that at most $2\varepsilon n$ vertices of $G \setminus \widetilde{W}_\sigma$ have more than $2\varepsilon n$ neighbours in \widetilde{W}_σ in the colour $c(\sigma)$ else we immediately find a monochromatic C_n in the colour $c(\sigma)$ in G . It follows that there are at most $2\varepsilon nN$ edges leaving the clique \widetilde{W}_σ in the colour $c(\sigma)$. Over all $\tau \in \mathcal{M}$, there are therefore at most $2^k \varepsilon nN < 3\varepsilon N^2$ edges in total leaving a clique \widetilde{W}_τ in the colour $c(\tau)$.

Let Φ now be the multigraph on vertex set \mathcal{M} where we have an edge between σ and τ in the colour j for each $j \notin \{c(\sigma), c(\tau)\}$ for which $G[\widetilde{W}_\sigma, \widetilde{W}_\tau]$ contains a matching of two edges in the colour j . First we observe that to complete the proof it suffices to show that there exists an admissible labelling φ of Φ (recall Definition 9.10). Indeed suppose that this is the case, then since φ is admissible, for each pair of distinct $\sigma, \tau \in \mathcal{M}$ and each $j \notin \Delta(\varphi(\sigma), \varphi(\tau)) \cup \{c(\sigma), c(\tau)\}$, we have that $G_j[\widetilde{W}_\sigma, \widetilde{W}_\tau]$ contains no matching of two edges and hence contains at most $|\widetilde{W}_\sigma| < n$ edges in total. It follows that there is a hypercube colouring H associated to the perfect matching $\varphi(\mathcal{M})$ of Q_k , where H has vertex set $V(G) \setminus \widetilde{W}$, such that for each $i \in [k]$,

$$|G_i \Delta H_i| \leq 2\varepsilon N^2 + |(G \setminus \widetilde{W})_i \Delta H_i| \leq 2\varepsilon N^2 + 3\varepsilon N^2 + n \binom{2^{k-1}}{2} \leq 6\varepsilon N^2.$$

The $2\varepsilon N^2$ term accounts for edges of G_i incident to \widetilde{W} , the $3\varepsilon N^2$ term accounts for edges of G_i leaving a clique \widetilde{W}_τ where $c(\tau) = i$, and the $n \binom{2^{k-1}}{2}$ term accounts for edges of G_i lying between pairs $\widetilde{W}_\tau, \widetilde{W}_\sigma$ for which $i \notin \Delta(\varphi(\sigma), \varphi(\tau)) \cup \{c(\sigma), c(\tau)\}$. We have thus shown that G is 6ε -close to H . It remains to show that we have the desired labelling of Φ .

Claim 9.14. Φ contains no monochromatic odd cycle.

Proof of Claim. Suppose otherwise and let $\sigma_1 \dots \sigma_\ell$ be an odd cycle in Φ in the colour j . This allows us to fix a matching of size two in graphs $G_j[\widetilde{W}_{\sigma_i}, \widetilde{W}_{\sigma_{i+1}}]$ for $i = 1, \dots, \ell$ (where $\sigma_{\ell+1} := \sigma_1$). Let S be the subset of vertices of G covered by these matchings and note that $|S| < 2^{k+1}$. We first aim to build a short even path in G_j with endpoints in \widetilde{W}_{σ_1} and $\widetilde{W}_{\sigma_\ell}$.

Let $x \in S \cap \widetilde{W}_{\sigma_1}$ and suppose that for some $2 \leq r < \ell$ there exists $y \in \widetilde{W}_{\sigma_r}$ such that G_j contains an xy -path P_r of length $r-1+2L(r-2)$ where $|P_r \cap S \cap \widetilde{W}_{\sigma_r}| = 1$ and $P_r \cap S \cap \widetilde{W}_{\sigma_s} = \emptyset$ for $r < s \leq \ell$ (note that this does indeed hold for $r = 2$). We may then pick $w \in \widetilde{W}_{\sigma_r} \cap S$ and $z \in \widetilde{W}_{\sigma_{r+1}} \cap S$ such that $\{w, z\}$ is an edge of $G_j[\widetilde{W}_{\sigma_r}, \widetilde{W}_{\sigma_{r+1}}]$ and $w \neq y$ (here we are using

that we have a matching of size two available to us by the definition of Φ). By the definition of Φ , σ_r is not in I_j and so by Corollary 9.6 there exists $\pi \in \mathcal{M}$ such that $R_j[W_{\sigma_r}, W_{\pi}]$ contains a connected perfect matching, F , whose matching edges are edges of \mathcal{T}_j . By Lemma 9.7 we see that $G_j[\widetilde{W}_{\sigma_r}, \widetilde{W}_{\pi}]$ contains a spanning $(2^{k+1}\delta, 1/(k+1), m')$ -super-regular blow-up of F where $2(1-\varepsilon)t' \leq v(F) \leq 2(1+\varepsilon)t'$ by (9.7). Moreover w and y lie in the same part in the bipartition of this blow-up. By calculations similar to those made previously, we may apply Lemma 8.3 to deduce that $G_j[\widetilde{W}_{\sigma_r}, \widetilde{W}_{\pi}]$ contains an yw -path Q of length $2L$. Moreover, since $|P_r \cup S| \leq 2^k L$ and using Fact 8.6 it is easy to ensure that Q only intersects $P_r \cup S$ at its endpoints. It follows that $P_{r+1} := P_r Q z$ is an xz -path of length $r + 2L(r - 1)$ where $|P_{r+1} \cap S \cap \widetilde{W}_{\sigma_{r+1}}| = 1$ and $P_{r+1} \cap S \cap \widetilde{W}_{\sigma_s} = \emptyset$ for $r + 1 < s \leq \ell$. It follows by recursion that there exists $u \in \widetilde{W}_{\sigma_\ell}$ and an xu -path P_ℓ of length $p := \ell - 1 + 2L(\ell - 2)$ and $|P_\ell \cap S \cap \widetilde{W}_{\sigma_\ell}| = 1$. Note that the length of P_ℓ is even.

Finally, let $\{v, t\} \subseteq S$ be an edge in $G_j[\widetilde{W}_{\sigma_\ell}, \widetilde{W}_{\sigma_1}]$ where $v \in \widetilde{W}_{\sigma_\ell}$ and $v \neq u$. If $x = t$ then applying Lemma 8.3 as above we find a uv -path Q_0 in the colour j of length $n - p - 1$ intersecting $P_\ell \cup S$ only at its endpoints. It follows that $P_\ell Q_0 x$ is a monochromatic copy of C_n contradicting (\dagger) . Similarly, if $x \neq t$, we find a uv -path Q_1 of length $2L$ and a tx -path Q_2 of length $n - p - 2L - 1$ both in the colour j so that $P_\ell Q_1 t Q_2$ is a monochromatic copy of C_n contradicting (\dagger) . \square

We now construct an admissible labelling of Φ recursively. Suppose that \mathcal{M}' is a perfect matching of Q_k and that $\psi : \mathcal{M} \rightarrow \mathcal{M}'$ is some bijection. Let $\sigma, \tau \in \mathcal{M}$ and suppose there is an edge f in Φ between σ and τ with colour j not in $\Delta(\psi(\sigma), \psi(\tau))$. We will call such an edge ‘bad’ (with respect to ψ).

Let $\{f_1, \dots, f_t\}$ be the set of edges of Φ that are bad with respect to the identity map $\iota : \mathcal{M} \rightarrow \mathcal{M}$ and note that ι is an admissible labelling of $\Phi \setminus \{f_1, \dots, f_t\}$. Suppose now that φ_i is an admissible labelling of $\Phi^i := \Phi \setminus \{f_1, \dots, f_t\}$ for some $1 \leq i \leq t$. Suppose that f_i is bad with respect to φ_i and that f_i has colour j and lies between $\sigma, \tau \in \mathcal{M}$. Note that $j \notin \{c(\sigma), c(\tau)\}$ by the definition of Φ . Moreover by the admissibility of φ_i we have $c(\sigma) = c(\varphi_i(\sigma))$ and $c(\tau) = c(\varphi_i(\tau))$. Since f_i is bad it follows that we must have $\varphi_i(\sigma)_j = \varphi_i(\tau)_j \in \{0, 1\}$. Let us show that σ, τ lie in separate components of Φ_j^i (the j th colour class of Φ^i). Suppose otherwise and take a path in Φ_j^i joining σ and τ . Since φ_i is admissible for Φ^i and $\varphi_i(\sigma)_j = \varphi_i(\tau)_j$ this path must have even length. It follows that f_i completes this path to a monochromatic odd cycle in Φ contradicting Claim 9.14. Let C then denote the component of Φ_j^i containing τ (so that $\sigma \notin C$). Let φ_{i-1} be the function on \mathcal{M} given by $\varphi_{i-1}(\rho) = \varphi_i(\rho)^j$ for all $\rho \in C$, $\varphi_{i-1}(\rho) = \varphi_i(\rho)$ otherwise. By Lemma 9.11, φ_{i-1} is an admissible labelling of Φ^i . Since

$$j \in \Delta(\varphi_i(\sigma), \varphi_i(\tau)^j) = \Delta(\varphi_{i-1}(\sigma), \varphi_{i-1}(\tau)),$$

we also have that φ_{i-1} is an admissible labelling of Φ^{i-1} . If f_i is not bad with respect to φ_i we simply let $\varphi_{i-1} = \varphi_i$. Running this recursion to the end we obtain an admissible labelling φ_0 of Φ as required. \square

10. CONCLUDING REMARKS

A simple adaptation of the proof method in this paper proves the following generalisation of Theorem 1.2.

Theorem 10.1. *For all $k \geq 3$ there exists N_k such that the following holds. If $N_k \leq n_1 \leq n_2 \dots \leq n_k$ are all odd then*

$$R(C_{n_1}, \dots, C_{n_k}) = 2^{k-1}(n_k - 1) + 1.$$

The off-diagonal case has been well-studied. Erdős *et al.* [EFRS76] determined the value of $R(C_n, C_{\ell_1}, C_{\ell_2})$ and $R(C_n, C_{\ell_1}, C_{\ell_2}, C_{\ell_3})$ for ℓ_i fixed and n sufficiently large. In a similar vein, as a corollary to a more general result in the study of Ramsey goodness, Allen, Brightwell and Skokan [ABS13] determined the value of $R(C_n, C_{\ell_1}, \dots, C_{\ell_k})$ for ℓ_i fixed and odd satisfying $\ell_i > 2^i$ for $1 \leq i \leq k$ and n sufficiently large. In [FL07b], Figaj and Łuczak asymptotically determine the Ramsey number of a triple of large cycles with any fixed combination of parities for the cycle lengths. In the case where not all of the cycles have the same parity, Ferguson [Fer15a, Fer15b, Fer15c] strengthened the asymptotic results of [FL07b] to exact results. It would be interesting to extend the methods of the present paper to such a mixed parity setting. More generally, we would like to investigate whether the analytic approach presented here has wider applications in Ramsey theory.

Acknowledgements. We would like to thank Julia Böttcher for helpful discussions and comments. We would also like to thank the anonymous referees for a very careful reading of this paper.

REFERENCES

- [ABS13] P. Allen, G. Brightwell, and J. Skokan. Ramsey-goodness—and otherwise. *Combinatorica*, 33(2):125–160, 2013.
- [BE73] J. A. Bondy and P. Erdős. Ramsey numbers for cycles in graphs. *J. Combinatorial Theory Ser. B*, 14:46–54, 1973.
- [BLS⁺12] F. Benevides, T. Łuczak, A. Scott, J. Skokan, and M. White. Monochromatic cycles in 2-coloured graphs. *Combinatorics, Probability and Computing*, 21(1-2):57–87, 2012.
- [Bon71] J. A. Bondy. Pancyclic graphs i. *J. Combinatorial Theory Ser. B*, 11(1):80–84, 1971.
- [BV04] S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, 2004.
- [CGP97] L.H. Clark, J. C. George, and T. D. Porter. On the number of l -factors in the n -cube. *Congressus Numerantium*, pages 67–70, 1997.
- [CL75] E. J. Cockayne and P. J. Lorimer. The Ramsey number for stripes. *J. Austral. Math. Soc.*, 19:252–256, 1975.
- [DJ17] A. N. Day and J. R. Johnson. Multicolour Ramsey Numbers of Odd Cycles. *J. Combinatorial Theory Ser. B*, 124:56–63, 2017.
- [EFRS76] P. Erdős, R. J. Faudree, C. C. Rousseau, and R. H. Schelp. Generalized ramsey theory for multiple colors. *J. Combinatorial Theory Ser. B*, 20(3):250–264, 1976.
- [EG59] P. Erdős and T. Gallai. On maximal paths and circuits of graphs. *Acta Math. Acad. Sci. Hungar.*, 10:337–356 (unbound insert), 1959.
- [EG75] P. Erdős and R. L. Graham. On partition theorems for finite graphs. In *Infinite and finite sets (Colloq., Keszthely, 1973; dedicated to P. Erdős on his 60th birthday)*, Vol. I, pages 515–527. Colloq. Math. Soc. János Bolyai, Vol. 10. North-Holland, Amsterdam, 1975.
- [Erd47] P. Erdős. Some remarks on the theory of graphs. *Bulletin of the American Mathematical Society*, 53(4):292–294, 1947.

- [ES35] P. Erdős and G. Szekeres. A combinatorial problem in geometry. *Compositio Mathematica*, 2:463–470, 1935.
- [Fer15a] D. Ferguson. The ramsey number of mixed-parity cycles i. *arXiv preprint arXiv:1508.07154*, 2015.
- [Fer15b] D. Ferguson. The ramsey number of mixed-parity cycles ii. *arXiv preprint arXiv:1508.07171*, 2015.
- [Fer15c] D. Ferguson. The ramsey number of mixed-parity cycles iii. *arXiv preprint arXiv:1508.07176*, 2015.
- [FL07a] A. Figaj and T. Łuczak. The Ramsey number for a triple of long even cycles. *J. Combin. Theory Ser. B*, 97(4):584–596, 2007.
- [FL07b] A. Figaj and T. Łuczak. The Ramsey number for a triple of large cycles. *arXiv preprint arXiv:0709.0048*, 2007.
- [FS74] R. J. Faudree and R. H. Schelp. All Ramsey numbers for cycles in graphs. *Discrete Math.*, 8:313–329, 1974.
- [GH88] N. Graham and F. Harary. The number of perfect matchings in a hypercube. *Applied Mathematics Letters*, 1(1):45–48, 1988.
- [GRSS07] A. Gyárfás, M. Ruszinkó, G. N. Sárközi, and E. Szemerédi. Three-color Ramsey numbers for paths. *Combinatorica*, 27(1):35–69, 2007.
- [KS96] J. Komlós and M. Simonovits. Szemerédi’s regularity lemma and its applications in graph theory. In *Combinatorics, Paul Erdős is eighty, Vol. 2 (Keszthely, 1993)*, volume 2 of *Bolyai Soc. Math. Stud.*, pages 295–352. János Bolyai Math. Soc., Budapest, 1996.
- [KSS05] Y. Kohayakawa, M. Simonovits, and J. Skokan. The 3-colored Ramsey number of odd cycles. In *Proceedings of GRACO2005*, volume 19 of *Electron. Notes Discrete Math.*, pages 397–402 (electronic), Amsterdam, 2005. Elsevier.
- [LSS12] T. Łuczak, M. Simonovits, and J. Skokan. On the multi-colored Ramsey numbers of cycles. *J. Graph Theory*, 69(2):169–175, 2012.
- [Luc99] T. Łuczak. $R(C_n, C_n, C_n) \leq (4 + o(1))n$. *J. Combin. Theory Ser. B*, 75(2):174–187, 1999.
- [ÖP13] P. R. J. Östergård and V. H. Pettersson. Enumerating perfect matchings in n-cubes. *Order*, pages 1–15, 2013.
- [PL86] M. D. Plummer and L. Lovász. *Matching theory*, volume 29. Elsevier, 1986.
- [Rad94] S. P. Radziszowski. Small Ramsey numbers. *Electron. J. Combin.*, 1:Dynamic Survey 1, 30 pp. (electronic), 1994.
- [Ram30] F. P. Ramsey. On a Problem of Formal Logic. *Proc. London Math. Soc.*, S2-30(1):264, 1930.
- [Ros73] V. Rosta. On a Ramsey-type problem of J. A. Bondy and P. Erdős. I, II. *J. Combinatorial Theory Ser. B*, 15:94–104; *ibid.* 15 (1973), 105–120, 1973.
- [Rud76] W. Rudin. *Principles of Mathematical Analysis*. International series in pure and applied mathematics. McGraw-Hill, 1976.
- [Sze78] E. Szemerédi. Regular partitions of graphs. In *Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976)*, volume 260 of *Colloq. Internat. CNRS*, pages 399–401. CNRS, Paris, 1978.

APPENDIX A. PROOF OF LEMMAS 8.3 AND 8.4

We use the following simple property of regular pairs which appears as Lemma 5 in [FL07b].

Lemma A.1. *Let $1/m \ll \delta \ll d$ and let $G = (V_1, V_2)$ be a (δ, d) -super-regular pair with $|V_1| = |V_2| = m$. Then for each pair $u \in V_1, w \in V_2$, G contains a uw -path of length ℓ for each odd $3 \leq \ell \leq 2(1 - 5\delta)m$.*

Lemma 8.3. *Let $q \geq 4$ and suppose that $\frac{1}{m} \ll \delta \ll d$. Let F be a connected matching of order q such that every vertex of F is incident to a matching edge and let H be a (δ, d, m) -super-regular blow-up of F . Then the following holds:*

If $i, j \in V(F)$ and there is an ij -path of length r in F , then for every pair of vertices $u \in U_i, w \in U_j$, there exists a uw -path of length ℓ in H for each $3q \leq \ell \leq (1 - 6\delta)qm$ such that $\ell \equiv r \pmod{2}$.

Proof. Take $i, j \in V(F)$ and let $u \in U_i, w \in U_j$. Let T be a spanning tree of F which includes every matching edge of F . Note that T contains a closed walk $W = y_0 \dots y_p$, where $y_1 = y_p = i$ and W covers each edge of T exactly twice, in particular $p = 2(q - 1)$ (note that $q = v(F)$). Using basic properties of regular pairs we can find a path $\widetilde{W} = w_0 \dots w_p$ in H where $u = w_0$ and $w_t \in U_{y_t}$ for all t . Let $P = x_0 \dots x_r$ be a path of length r in F where $x_0 = i, x_r = j$. Again, using basic properties of regular pairs we can find a path $\widetilde{P} = v_0 \dots v_r$ in H where $v_0 = w_p, v_r = w, v_t \in U_{x_t}$ for all t and \widetilde{P} intersects \widetilde{W} only in the vertex w_p . Letting $Q = \widetilde{W}\widetilde{P}$, it follows that Q is a uw -path in H of length $r + p = r + 2(q - 1) \equiv r \pmod{2}$. Suppose that $\{a, b\}$ is a matching edge of F so that (U_a, U_b) is (δ, d) -super-regular in H . Note that Q visits each set U_i in H at most 3 times and so there exist $U'_a \subseteq U_a \setminus Q, U'_b \subseteq U_b \setminus Q$ such that $|U'_a| = |U'_b| = m - 3$. Note that (U'_a, U'_b) is certainly $(2\delta, d/2)$ -super-regular by Fact 8.6. By construction, we may pick consecutive vertices w_t, w_{t+1} of \widetilde{W} (and hence Q) such that $w_t \in U_a, w_{t+1} \in U_b$. By super-regularity we may then pick vertices $u_a \in N(w_{t+1}) \cap U'_a, u_b \in N(w_t) \cap U'_b$ such that $\{u_a, u_b\}$ is an edge of H . Applying Lemma A.1 to (U'_a, U'_b) and vertices u_a, u_b , it follows that we can find a $q_t q_{t+1}$ -path in H which intersects Q only at its endpoints and we can choose this path to have any odd length $1 \leq \ell \leq 2(1 - 5\delta)(m - 3) + 2$. Note that letting such a path replace the edge $\{q_t, q_{t+1}\}$ in Q does not change the parity of the length of Q . Applying the same argument to each matching edge of F we see that H contains uw -paths of each length $r + 2(q - 1) \leq \ell \leq r + 2(q - 1) + \frac{q}{2} \cdot 2(1 - 6\delta)m$ for which $\ell \equiv r \pmod{2}$. The result follows. \square

Lemma 8.4. *Let $q \geq 4$ and let $\frac{1}{m} \ll \delta \ll d$. Let F be an odd connected matching of order q and suppose that H is a (δ, m) -regular blow-up of F with minimum density d . Then H contains a cycle of length ℓ for each odd $3q \leq \ell \leq (1 - 6\delta)qm$*

Proof. Since F is non-bipartite it contains an odd cycle C . Since the largest matching in F has $q/2$ edges it follows that $|C| \leq q + 1$. Let $T \subseteq F$ be a minimal tree that contains every matching edge of F . It is easy to show that T must have $< 2q$ vertices. Let W be a closed walk in T which traverses each edge of T precisely twice (so in particular W has even length). Since W and C must intersect, we can augment the walk W by C to obtain a closed walk $W' = x_1 \dots x_p x_1$ in F where p is odd and $p \leq 3q$ by the above. Note that by Facts 8.6 and 8.7, we can find $H' \subseteq H$ such that H' is a $(2\delta, d/2, (1 - \delta)m)$ -super-regular blowup of F . Let U_j denote the vertex class of H' corresponding to the vertex j in F for each $j \in V(F)$. Using basic properties of regular pairs, we can find an odd cycle $D = v_1 \dots v_p v_1$ in H' where $v_j \in U_{x_j}$ for all j .

Suppose that $\{a, b\}$ is a matching edge of F so that (U_a, U_b) is $(2\delta, d/2)$ -super-regular. By construction, we may pick consecutive vertices v_t, v_{t+1} of D such that $v_t \in U_a, v_{t+1} \in U_b$. Note that D visits each set U_i in H at most 3 times. We may therefore apply Lemma A.1 as we did in the proof of Lemma 8.3 to find a $v_t v_{t+1}$ -path Q in H' such that Q intersects D only at its endpoints and we can choose Q to have any odd length $1 \leq \ell \leq 2(1 - 5\delta)[(1 - \delta)m - 3] + 2$. Applying the same argument to each matching edge of F we see that H contains an odd cycle of each odd length $p \leq \ell \leq p + \frac{q}{2} \cdot 2(1 - 6\delta)m$. The result follows. \square