

Induced subgraphs of induced subgraphs of large chromatic number

António Girão^{*†} Freddie Illingworth^{*†} Emil Powierski^{*}
Michael Savery^{*‡} Alex Scott^{*†} Yuri Tamitegama^{*} Jane Tan^{*}

4 August 2023

Abstract

We prove that, for every graph F with at least one edge, there is a constant c_F such that there are graphs of arbitrarily large chromatic number and the same clique number as F in which every F -free induced subgraph has chromatic number at most c_F . This generalises recent theorems of Briański, Davies and Walczak, and Carbonero, Hompe, Moore and Spirkl. Our results imply that for every $r \geq 3$ the class of K_r -free graphs has a very strong vertex Ramsey-type property, giving a vast generalisation of a result of Folkman from 1970. We also prove related results for tournaments, hypergraphs and infinite families of graphs, and show an analogous statement for graphs where clique number is replaced by odd girth.

1 Introduction

Given a graph with large chromatic number, what can we say about its local structure? For example, are there induced subgraphs that appear in every graph with sufficiently large chromatic number? Perhaps the first question of this type one might ask is whether every graph with large chromatic number contains a large complete subgraph. However, it has been known since the 1940s that this is not the case: Tutte [8] and Zykov [41] gave constructions of triangle-free graphs with arbitrarily large chromatic number, and Erdős [10] showed that there exist graphs with both arbitrarily large girth and chromatic number. But this raises another question: if a graph with large chromatic number does not contain a large complete subgraph, then what can we say about its induced subgraphs?

It is helpful to use the language of χ -bounded classes: a hereditary class \mathcal{G} is χ -*bounded* if there is some function f such that $\chi(G) \leq f(\omega(G))$ for all $G \in \mathcal{G}$, where $\chi(G)$ denotes the *chromatic number* of G and $\omega(G)$ denotes the *clique number* (the maximum number of vertices in a complete subgraph of G). The study of χ -boundedness was strongly

^{*}Mathematical Institute, University of Oxford, Oxford OX2 6GG, UK

[†]Research supported by EPSRC grant EP/V007327/1

[‡]Heilbronn Institute for Mathematical Research, Bristol, UK

Email: {girao,illingworth,powierski,savery,scott,tamitegama,jane.tan}@maths.ox.ac.uk

influenced by an important paper of Gyárfás [14], that put forward a number of influential conjectures. There has been a burst of recent activity in this area, and many of the conjectures have recently been resolved (see [33, 36]).

An important strategy for showing that a hereditary class \mathcal{G} is χ -bounded has been that of ‘cleaning up’ its structure: given $G \in \mathcal{G}$ with very large chromatic number, we can attempt to find an induced subgraph that still has fairly large chromatic number, but is structurally simpler in some way. If this succeeds for all $G \in \mathcal{G}$, then to determine whether the class \mathcal{G} is χ -bounded we need only consider these structurally simpler graphs. A very strong version of this approach was proposed independently by Fred Galvin and Vojtěch Rödl (see [21]) and by Louis Esperet (see [36]), who made the striking conjecture that every graph with large chromatic number contains either a large clique or an induced triangle-free subgraph with large chromatic number. If this were true, it would have radical consequences for the study of χ -boundedness, as it would imply that whether or not a class is χ -bounded is determined by the triangle-free graphs in the class. However, Galvin and Rödl’s conjecture was recently disproved by Carbonero, Hompe, Moore and Spirkl [7], who proved the following via a surprising new twist on a construction of Kierstead and Trotter [17].

Theorem 1 (Carbonero, Hompe, Moore and Spirkl [7]). *There are graphs of arbitrarily large chromatic number that contain neither a K_4 nor a triangle-free induced subgraph of chromatic number greater than four.*

Shortly afterwards, Briański, Davies and Walczak [5] extended the result of Carbonero, Hompe, Moore and Spirkl to cliques of prime order in an ingenious paper proving the following.¹

Theorem 2 (Briański, Davies and Walczak [5]). *For every prime p , there are K_{p+1} -free graphs G of arbitrarily large chromatic number such that every K_p -free induced subgraph H of G satisfies $\chi(H) \leq \omega(H)^{\omega(H)^2}$.*

Theorems 1 and 2 show that there are K_{p+1} -free graphs of large chromatic number in which every induced subgraph of large chromatic number contains a copy of K_p . Thus, the ‘cleaning up’ strategy cannot be deployed in the straightforward way that Esperet’s conjecture suggests: we cannot simply drop to an induced subgraph with large chromatic number but smaller clique number. Nevertheless, we might still hope that we can get rid of some other graphs.

In this paper, we show that Theorems 1 and 2 are just the tip of the iceberg. Let us say that a graph G is *F -free* if it does not contain an induced copy of F . We will show that results such as Theorems 1 and 2 hold not just for cliques, but in fact for every nontrivial graph F .

Theorem 3. *For every graph F with at least one edge, there is a constant c_F and graphs G of arbitrarily large chromatic number and the same clique number as F such that every F -free induced subgraph of G is c_F -colourable.*

¹Building on this result, they further showed that there are classes of graphs that are χ -bounded but not polynomially χ -bounded, disproving another conjecture of Esperet [11].

It follows that there is no general ‘cleaning up’ strategy to get rid of induced subgraphs: it is necessary instead to use particular properties of the class under consideration, or to focus on other structural features.²

Theorem 3 is also interesting from the perspective of Ramsey theory. A class \mathcal{C} of graphs is a *vertex Ramsey class* if, for every $F \in \mathcal{C}$ and positive integer k , there is some $G \in \mathcal{C}$ such that every vertex colouring of G in k colours contains a monochromatic induced copy of F . Vertex Ramsey classes were first studied in the 1970s by Folkman [12] and Nešetřil and Rödl [24] as part of the development of a rich theory (see [22]). The general question is to determine which hereditary classes \mathcal{C} of graphs are vertex Ramsey. There has been substantial progress, particularly in the case when \mathcal{C} is determined by a single excluded induced subgraph, that is, when \mathcal{C} is the class of F -free graphs for some F (see [12, 24, 28, 30–32]), or when \mathcal{C} is defined by excluding a clique and a tree [16, 18]. Ramsey classes of infinite graphs have also been studied in [15, 19, 20].

One of the earliest results on vertex Ramsey classes is due to Folkman [12], who proved in 1970 that the class \mathcal{C}_r of K_r -free graphs is vertex Ramsey for every $r \geq 3$: in other words, for every K_r -free graph F and every positive integer k , there is a K_r -free graph H such that in every k -colouring of $V(H)$ some colour class contains an induced copy of F . Theorem 3 proves a vastly stronger property: it shows that for every K_r -free graph F there are K_r -free graphs with arbitrarily large chromatic number in which *every* induced subgraph of large chromatic number contains an induced copy of F . This difference resembles the difference between the theorems of Van der Waerden [40] and Szemerédi [39] on monochromatic arithmetic progressions in the integers.

Interestingly, there is no edge version of Theorem 3: Rödl [27] proved that, for all k , every graph with sufficiently large chromatic number contains a triangle-free subgraph with chromatic number k . In particular, while the class of triangle-free graphs is both vertex and edge Ramsey [25] and satisfies the strengthened vertex Ramsey property, it does not satisfy the analogous strengthened edge Ramsey property.

If F is triangle-free, then Theorem 3 guarantees a triangle-free G . It is natural to ask whether one can go further and guarantee that G has the same girth as F . We conjecture that this is the case, but have been unable to prove it. However, we can prove the following weaker statement, where girth is replaced by *odd girth* (the length of the shortest odd cycle). This implies that the class \mathcal{C}'_r of $\{C_3, C_5, \dots, C_{2r+1}\}$ -free graphs has the strengthened vertex Ramsey property mentioned above.

Theorem 4. *For every nonbipartite graph F , there is a constant c'_F and graphs G of arbitrarily large chromatic number and the same odd girth as F such that every F -free induced subgraph of G is c'_F -colourable.*

Theorems 3 and 4 both easily extend to finite families of graphs. Indeed, if \mathcal{F} is a finite family of graphs (at least one of which contains an edge) and $G_{\mathcal{F}}$ is the disjoint union of the members of \mathcal{F} , then applying Theorem 3 to $G_{\mathcal{F}}$ shows that there is some constant $c_{\mathcal{F}}$ and graphs G of arbitrarily large chromatic number and the same clique number as

²See, for example, the discussion of diameter and Conjecture 1.10 in [35].

the largest clique number of a graph in \mathcal{F} such that every induced subgraph of G with chromatic number greater than $c_{\mathcal{F}}$ contains every member of \mathcal{F} as an induced subgraph. Does this phenomenon occur for infinite families? We will say that a family of graphs \mathcal{F} is *durable* if there exists an infinite graph H with infinite chromatic number such that every infinite-chromatic induced subgraph of H contains every member of \mathcal{F} as an induced subgraph. All finite families of graphs are durable but the situation for infinite families is more complicated. We obtain an effective test for deciding whether a family of unbounded chromatic number is durable which, for example, gives the following result.

Theorem 5. *For every integer g , the family of graphs with girth at least g is not durable.*

In the bounded chromatic number case we obtain a link to χ -boundedness, which in particular implies the following.

Theorem 6. *The family of paths is durable.*

We also extend and generalise Theorem 3 to the settings of tournaments and hypergraphs (see Sections 5 and 6 for definitions not given below). For tournaments we prove the following.

Theorem 7. *For every tournament T , there is a constant C_T and tournaments S of arbitrarily large chromatic number such that every T -free subtournament of S has chromatic number at most C_T .*

For hypergraphs, there are two notions for chromatic number: the usual one where no edge is monochromatic and the strong one where no two vertices in an edge have the same colour. Pleasingly, our generalisation of Theorem 3 uses the usual notion of chromatic number for the constructed hypergraph, but the strong notion for its induced subhypergraphs. The statement also has a very general analogue of a clique. A hypergraph is said to be *strongly t -colourable* if its vertices can be t -coloured such that no edge contains two vertices of the same colour. A hypergraph *covers* a pair of vertices if it has an edge containing both of them. Theorem 8 applies to general (not necessarily uniform) hypergraphs, but we assume throughout this paper that all edges have size at least two.

Theorem 8. *For every hypergraph \mathcal{F} with at least one edge, there is a constant $c_{\mathcal{F}}$ and hypergraphs \mathcal{G} of arbitrarily large chromatic number such that every \mathcal{F} -free induced subhypergraph of \mathcal{G} is strongly $c_{\mathcal{F}}$ -colourable. Moreover, we can take \mathcal{G} such that if \mathcal{G} covers every pair from some $X \subseteq V(\mathcal{G})$, then $\mathcal{G}[X]$ is an induced subhypergraph of \mathcal{F} .*

It follows from the moreover part of the theorem that the edge sizes which appear in \mathcal{G} are the same as those which appear in \mathcal{F} , and that for all k the maximum size of a k -uniform clique in \mathcal{G} is the same as that in \mathcal{F} . In particular, this theorem contains Theorem 3.

Our constructions for Theorems 3, 4 and 8 all use a common framework: we start with a base structure of large chromatic number and add edges according to ‘distances’ in the base structure. Using generalised Sidon sets, we are able to choose these ‘distances’ to encode F or \mathcal{F} and give precise control over where its copies appear.

The rest of this paper is structured as follows. In Section 2 we describe the base graphs

used in the proofs of Theorems 3 and 4, then we complete the proofs of these theorems in Sections 3 and 4 respectively. Section 5 is concerned with tournaments and contains the proof of Theorem 7, then in Section 6 we turn our attention to hypergraphs and prove Theorem 8. Our results on the durability of infinite families of graphs are covered in Section 7, before we close with a discussion of some open problems in Section 8.

Note added. Since the original version of this paper was made available, Briański, Davies and Walczak have independently proved a strengthening of their Theorem 2, improving the bound on $\chi(H)$ and removing the restriction that p is prime. This can be found in the updated version of their arXiv paper [6].

2 Orientations of classical constructions

Our constructions for Theorems 3 and 4 start with orientations of graphs with high chromatic number and restricted clique number or girth. Traditionally the underlying graphs in such orientations have been the well known Zykov graphs [41]. Orientations of these were first considered by Kierstead and Trotter [17], and then used in the constructions of both Carbonero, Hompe, Moore and Spirkl [7] and Briański, Davies and Walczak [5]. For the proof of Theorem 3, we require precisely the following properties of the underlying graphs X_n and their orientations \vec{X}_n :

- (a) $\chi(X_n) = n$,
- (b) \vec{X}_n is acyclic, that is, it contains no directed cycle, and
- (c) for all $u, v \in V(\vec{X}_n)$, there is at most one directed path from u to v in \vec{X}_n .

The Zykov graphs and their orientations given by Kierstead and Trotter satisfy these conditions and so suffice for the proof of Theorem 3. In fact any X_n with chromatic number n and girth greater than $2n - 2$ has some orientation \vec{X}_n satisfying these three properties (see Lemma 10).

For the proof of Theorem 4 we require one further property in order to rule out short odd cycles (in particular, we cannot take Zykov graphs as the base graphs in the proof of Theorem 4). Define a *change of direction* in a digraph to be a pair of edges incident to a common vertex that are both directed towards or both directed away from that vertex. For a given positive integer g we ask that \vec{X}_n satisfies:

- (d) for every cycle in X_n , the corresponding oriented cycle in \vec{X}_n has at least g changes of direction.

Strictly speaking, \vec{X}_n and condition (d) depend on the choice of positive integer g , but we suppress this dependence in the notation. Note that if $g \geq 3$, then property (d) implies both (b) and (c). Lemma 10 below states that any X_n with chromatic number n and girth greater than $(g - 1)(n - 1)$ where $g \geq 3$ has an orientation \vec{X}_n satisfying all four properties.

By property (c), it is natural to view $V(\vec{X}_n) = V(X_n)$ as a poset with a strict partial

order based on *reachability*: for distinct vertices u and v , we write $u < v$ if there is a directed path from u to v in $\overrightarrow{X_n}$. Note that there is a correspondence between chains in this poset and sets of vertices that lie on a single directed path in $\overrightarrow{X_n}$. For vertices $u < v$ we can define the *distance* between u and v , written $d(u, v)$, to be the length of the unique directed path from u to v . The uniqueness of directed paths ensures that $d(u, w) = d(u, v) + d(v, w)$ for vertices $u < v < w$.

At this stage, the reader may like to proceed directly to Sections 3 and 4 for the proofs of our main theorems, keeping in mind that there exist (di)graphs with the required properties. We continue in this section with the promised constructions.

The following proposition, which is the acyclic case of the Gallai-Roy-Hasse-Vivater theorem, provides a strong relationship between orientations and chromatic number.

Proposition 9. *A graph L is k -colourable if and only if it has an acyclic orientation \overrightarrow{L} which contains no directed paths of length k .*

Proof. First suppose that L is k -colourable and let $c: V(L) \rightarrow \{1, \dots, k\}$ be a proper k -colouring. For each edge $e = uv$ of L , if $c(u) < c(v)$, then orient e from u to v and if $c(v) < c(u)$, then orient e from v to u . For any directed path $u_1 u_2 \dots u_t$ in the resulting orientation \overrightarrow{L} , we have $1 \leq c(u_1) < c(u_2) < \dots < c(u_t) \leq k$ and so the orientation is acyclic and the path has length at most $k - 1$.

On the other hand, if the acyclic orientation \overrightarrow{L} has no directed paths of length k , then we may simply colour each vertex v of L by the length of the longest directed path in \overrightarrow{L} ending at v . \square

Erdős [10] showed that for any integers $g, n \geq 2$ there are graphs of chromatic number n and girth at least g . This together with the following lemma (pointed out to us by Jarik Nešetřil) shows that there are indeed digraphs with properties (a) to (d).

Lemma 10. *Let $g \geq 3$ and $n \geq 2$ be integers. If a graph X_n has chromatic number n and girth greater than $(g - 1)(n - 1)$, then X_n has an orientation $\overrightarrow{X_n}$ satisfying properties (a), (b), (c) and (d).*

Proof. The graph X_n has chromatic number n so, by Proposition 9, has an acyclic orientation $\overrightarrow{X_n}$ which contains no directed paths of length n . In particular, $\overrightarrow{X_n}$ automatically satisfies properties (a) and (b). Consider a cycle C in X_n . The length of C is greater than $(g - 1)(n - 1)$ and every directed path in $\overrightarrow{X_n}$ has length at most $n - 1$, so C has at least g changes of direction in $\overrightarrow{X_n}$. Hence property (d) is satisfied. Finally, as $g \geq 3$, this implies that $\overrightarrow{X_n}$ satisfies property (c) also. \square

3 Clique number

In this section, we will give the proof of Theorem 3. We have not tried hard to optimise the bounds; however, writing $f = |V(F)|$, the proof gives³

$$\begin{aligned} c_F &\leq (32f^4)^{|E(F)|}, \\ c_F &\leq ((2 + o(1))f^4)^{|E(F)|}. \end{aligned} \tag{1}$$

We begin by sketching the proof strategy, and then discuss B_3 -sets in Section 3.1, before giving the details of the proof in Section 3.2.

The constructions in [5, 7, 17] start with oriented (Zykov) graphs (\vec{X}_n) satisfying properties (a) through (d), and add new edges $\vec{u}\vec{v}$ (or sometimes $\vec{v}\vec{u}$ in [7]) between vertices $u < v$ whenever p does not divide $d(u, v)$ for some prime p . These papers use arbitrary primes p , $p = 3$ and $p = 2$ respectively. Our construction is similar in that, for some prime p , we add edges to \vec{X}_n based on the residues modulo p of the distances between endpoints. However, the set of residues for which we add edges is now more sophisticated. Let \vec{G} be the digraph so obtained, and let G be its underlying undirected graph. As X_n is a subgraph of G , we know that $\chi(G) \geq n$. The extra edges will be added so that the following properties hold:

- $\omega(G) = \omega(F)$,
- \vec{G} is acyclic, and
- if there is a long (depending only on F) directed path in \vec{G} all of whose edges correspond to the same distance modulo p , then some vertices of this path induce a copy of F in G .

The second and third properties are particularly useful in light of Proposition 9.

Now taking G as above, suppose that H is an induced subgraph of G which does not contain F as an induced subgraph. Let H_i be the subgraph of H which consists of those edges corresponding to distance i modulo p . By the third bullet point, H_i cannot contain a long directed path. Thus, Proposition 9 and the second bullet point imply that H_i has chromatic number bounded in terms of F . Taking a product colouring (over the possible i) will then show that H itself must have chromatic number at most some c_F .

3.1 B_3 -sets

To guarantee that G has the same clique number as F our construction will utilise B_3 -sets. A set of integers $S = \{a_1 < a_2 < \dots < a_k\}$ is a B_3 -set if the sums

$$a_{i_1} + a_{i_2} + a_{i_3}, \quad 1 \leq i_1 \leq i_2 \leq i_3 \leq k$$

are all different. A simple example of a B_3 -set is the set of powers of four — it would suffice to use an initial segment of this for our arguments, although the resulting bound

³Using results from the updated version of Briański, Davies and Walczak's paper [6] (see the note at the end of the introduction), this can be improved to $c_F = O(f^9)$ by noting that the graph G we construct in Section 3.2 can be taken to be a subgraph of their $G_{n,p}$.

on c_F would be much worse. B_3 -sets (and more generally B_h -sets, which we discuss in Section 4.1) were introduced by Bose and Chowla [4] as a generalisation of Sidon sets [37], which are precisely the B_2 -sets (i.e. sets of integers where all pairwise differences are distinct). In particular, every B_3 -set is also a B_2 -set. Let S be a B_3 -set and let

$$D = \{s - s' : s' < s \text{ and } s, s' \in S\}$$

be the *difference set of S* . The sets S and D satisfy the following two claims which motivate our interest in B_3 -sets.

Claim 1 (Triangle fact). *Suppose that D contains not necessarily distinct d_1, d_2 such that $d_1 + d_2 \in D$. Then there are $b_1 < b_2 < b_3$ all in S such that $\{b_2 - b_1, b_3 - b_2\} = \{d_1, d_2\}$.*

Proof. Suppose that $d_1, d_2, d_1 + d_2$ are all in D . Then there are $x_1, x_2, x_3, y_1, y_2, y_3$ all in S with $d_1 = y_1 - x_1$, $d_2 = y_2 - x_2$ and $d_1 + d_2 = y_3 - x_3$. Hence,

$$\begin{aligned} (y_1 - x_1) + (y_2 - x_2) &= y_3 - x_3, \\ \Rightarrow y_1 + y_2 + x_3 &= x_1 + x_2 + y_3. \end{aligned}$$

But S is a B_3 -set, so y_1, y_2, x_3 must be x_1, x_2, y_3 in some order. The elements of D are positive integers so $x_i \neq y_i$ for all i . Hence y_1 is either x_2 or y_3 .

- If $y_1 = x_2$, then $b_1 = x_1$, $b_2 = y_1 = x_2$, and $b_3 = y_2$ satisfy $(b_2 - b_1, b_3 - b_2) = (d_1, d_2)$.
- If $y_1 = y_3$, then $y_2 = x_1$ (as $y_2 \neq x_2$) and so $b_1 = x_2$, $b_2 = x_1 = y_2$, and $b_3 = y_1$ satisfy $(b_2 - b_1, b_3 - b_2) = (d_2, d_1)$. \square

The preceding claim can be leveraged to cover more distances.

Claim 2 (Clique fact). *Let ℓ be a natural number and suppose that D contains not necessarily distinct d_1, \dots, d_ℓ such that every sum of the form*

$$\sum_{i_1 \leq j \leq i_2} d_j \tag{2}$$

with $1 \leq i_1 \leq i_2 \leq \ell$ is in D . Then there are $b_1 < \dots < b_{\ell+1}$ all in S such that either $b_{i+1} - b_i = d_i$ for all i , or $b_{i+1} - b_i = d_{\ell+1-i}$ for all i .

Proof. We proceed by induction on ℓ . The case $\ell = 1$ is trivial and the case $\ell = 2$ is the triangle fact, so suppose $\ell \geq 3$. By the induction hypothesis, there are $b_1 < \dots < b_\ell$ such that either $b_{i+1} - b_i = d_i$ for all $i < \ell$, or $b_{i+1} - b_i = d_{\ell-i}$ for all $i < \ell$. We may assume that $b_{i+1} - b_i = d_i$ for all $i < \ell$, as the other case follows by a symmetric argument.

Let us say that a difference $d \in D$ is associated with $b \in S$ if there is some $c \in S$ such that $|b - c| = d$. As S is a Sidon set, $d_{\ell-1}$ is only associated with $b_{\ell-1}$ and b_ℓ . Thus, applying the triangle fact to $d_{\ell-1}$ and d_ℓ we find that there is $b \in S$ such that either $b_{\ell-1} - b = d_\ell$ or $b - b_\ell = d_\ell$. In the latter case we are done, so assume that $b_{\ell-1} - b = d_\ell$.

Similarly, $d := \sum_{j=1}^{\ell-1} d_j$ is only associated with b_1 and b_ℓ , so by applying the triangle fact to d and d_ℓ we find that there is some $b' \in S$ such that either $b_1 - b' = d_\ell$ or $b' - b_\ell = d_\ell$. In the latter case we are done so we assume that $b_1 - b' = d_\ell$. But then $b' < b_1 < b_{\ell-1}$ are all associated with d_ℓ while S is a Sidon set. This is a contradiction, as required. \square

The existence of large B_3 -sets is guaranteed by the following theorem of Bose and Chowla [4, Thm. 2]. We use the standard shorthand $[n] = \{1, \dots, n\}$.

Theorem 11 (Bose and Chowla [4]). *Let q be one more than a prime power. There is a B_3 -subset of $[q^3]$ of size q that contains both 1 and 2.*

Given a natural number f , using the well known fact that the gap between consecutive primes $q < q'$ is $o(q)$ (see, for instance, [2] for an up-to-date bound), Theorem 11 implies that for $N = (1 + o(1))f^3$, there is a B_3 -subset of $[N]$ of size at least f that contains both 1 and 2. Bertrand's postulate gives a non-asymptotic bound: the smallest prime $p = q - 1$ greater than $f - 1$ is less than $2f - 1$ and so Theorem 11 gives a B_3 -subset of $[q^3] \subseteq [8f^3]$ of size at least f that contains both 1 and 2.

3.2 Proof of Theorem 3

Fix a graph F with at least one edge and let $f = |V(F)|$. Let p be a prime number chosen to be large enough that $\lfloor (p-1)/2 \rfloor$ has a B_3 -subset S of size f that contains 1 and 2. By the discussion at the end of Section 3.1, it suffices to have $(p-1)/2 \geq \min\{(1+o(1))f^3, 8f^3\}$. Using the result on prime gaps and Bertrand's postulate we may choose such a p with $p \leq \min\{(2+o(1))f^3, 32f^3\}$.

Let the elements of S be $s_1 = 1, s_2 = 2, s_3, \dots, s_f$ in ascending order and let the difference set of S be $D = \{s_j - s_i : i < j\}$. As F has at least one edge we may take a copy F^* of F with vertex set S which includes the edge s_1s_2 . We write

$$E = \{s_j - s_i : i < j \text{ and } s_i s_j \in E(F^*)\}$$

for the set of distances corresponding to edges of F^* . Since $s_1 = 1$ and $s_2 = 2$, we have $1 \in E$. Note that $E \subseteq D$. Let $\overline{D} = D + p\mathbb{Z}$ and $\overline{E} = E + p\mathbb{Z}$.

For a fixed large integer n , recall the (di)graphs \overrightarrow{X}_n and X_n from Section 2, and that for vertices $u < v$ (in the reachability ordering) of these graphs we write $d(u, v)$ for the length of the unique directed path from u to v in \overrightarrow{X}_n . We now define a digraph \overrightarrow{G} with vertex set $V(\overrightarrow{X}_n)$ as follows.

The edge \overrightarrow{uv} is present if $u < v$ and $d(u, v) \in \overline{E}$.

The digraph \overrightarrow{G} is acyclic, since the presence of the edge \overrightarrow{uv} implies that $u < v$ in the reachability order. Let G be the underlying undirected graph of \overrightarrow{G} .

We now verify that G satisfies the conditions of Theorem 3. Since $1 \in E$, the graph X_n is a subgraph of G and so $\chi(G) \geq n$. Next we show that $\omega(G) \leq \omega(F)$. To this end, suppose that vertices $v_1, \dots, v_{\ell+1}$ form a clique in G . Since any two of these vertices are $<$ -comparable, we may assume that $v_1 < \dots < v_{\ell+1}$. For $i \in [\ell]$ let $d_i \in \{0, \dots, p-1\}$ be the residue modulo p of $d(v_i, v_{i+1})$, and note that

$$d(v_{i_1}, v_{i_2+1}) \equiv \sum_{i_1 \leq j \leq i_2} d_j \pmod{p}$$

for all $1 \leq i_1 \leq i_2 \leq \ell$. Since $v_1, \dots, v_{\ell+1}$ form a clique in G , all these distances $d(v_{i_1}, v_{i_2+1})$, and hence all sums of the form in (2), are in $\overline{E} \subseteq \overline{D}$. We claim that they are actually in D . It is clear that $d_j \in D$ for all j , and since $D \subseteq [(p-1)/2]$ we have that the two-term sums of the form in (2) are in D as well. Iterating this shows that all sums of the form in (2) are in D .

It now follows from the clique fact that there are $b_1 < \dots < b_{\ell+1}$ all in S such that either $b_{i+1} - b_i = d_i$ for all i , or $b_{i+1} - b_i = d_{\ell+1-i}$ for all i . All pairwise distances between elements of $B = \{b_1, \dots, b_{\ell+1}\}$ are of the form in (2) and so are in \overline{E} . But since $B \subseteq S$, they must all actually be in E . By the definition of E and the fact that S is a Sidon set, the vertices $b_1, \dots, b_{\ell+1}$ induce a clique in F^* and so $\ell + 1 \leq \omega(F)$, as required.

We now turn to the last required property of G , that all of its induced subgraphs of chromatic number greater than some c_F contain an induced copy of F . Start by colouring each edge uv of G (and \vec{G}), where $u < v$, with the colour $d(u, v) \bmod p$. Let \vec{P} be a directed path in \vec{G} of length pf which is monochromatic in colour i for some $i \in E$. Label the first vertex of this path by v_1 . Recall that the elements of S are $s_1 < \dots < s_f$. Let v_2 be the first vertex on path \vec{P} such that $d(v_1, v_2) \equiv s_2 - s_1 \pmod{p}$. Since p is prime and $i \neq 0$, there are at most $p-1$ edges of \vec{P} between v_1 and v_2 on this path. Next, define v_3 to be the first vertex of \vec{P} such that $d(v_2, v_3) \equiv s_3 - s_2 \pmod{p}$, and continue in this manner to define v_4, \dots, v_f . Note that for each i there are at most $p-1$ edges between v_i and v_{i+1} on path \vec{P} , so the path is long enough that these vertices can be found.

The way in which we have chosen the vertices v_i means that, for each pair $i < j$, we have $d(v_i, v_j) \equiv s_j - s_i \pmod{p}$. Then, by the construction of G and the fact that S is a Sidon set, the edge $v_i v_j$ is present in the graph exactly when there is an edge between s_i and s_j in F^* . It follows that G contains an induced copy of F with vertex set $\{v_1, \dots, v_f\}$.

Now let H be an induced subgraph of G which does not contain an induced copy of F , and let \vec{H} be the subdigraph of \vec{G} induced by the same set of vertices. For each $i \in E$, let H_i be the subgraph of H consisting of the i -coloured edges and define \vec{H}_i similarly. By the above, \vec{H}_i contains no directed paths of length pf and hence H_i is pf -colourable by Proposition 9. Take such a proper colouring $\chi_i: V(H_i) \rightarrow [pf]$ for each i . Then the product colouring $\chi: V(H) \rightarrow [pf]^{|E|}$ given by $\chi(v) = (\chi_i(v): i \in E)$ is a proper colouring of H . Hence, we may take c_F to be $(pf)^{|E|}$. Finally, we note that $|E| = |E(F)|$ and recall that $p \leq \min\{(2 + o(1))f^3, 32f^3\}$ to obtain the bounds given by (1).

4 Odd girth

The proof of Theorem 4 shares some ingredients with that of Theorem 3. In this exposition, we will pay particular attention to the differences.

4.1 B_h -sets

Given a natural number $h \geq 2$, we can define the notion of B_h -sets analogously to B_3 -sets. That is, a set of integers $S = \{a_1 < a_2 < \dots < a_k\}$ is a *B_h -set* if the sums

$$a_{i_1} + a_{i_2} + \dots + a_{i_h}, \quad 1 \leq i_1 \leq i_2 \leq \dots \leq i_h \leq k$$

are all different. Note that a B_h -set is also a $B_{h'}$ -set for any $h' < h$, and in particular is a Sidon set. The powers of h form a B_h -set but Bose and Chowla [4, Thm. 2] provide a more efficient construction.

Theorem 12 (Bose and Chowla [4]). *Let h be a positive integer and q be one more than a prime power. There is a B_h -subset of $[q^h]$ of size q that contains both 1 and 2.*

The importance of B_h -sets for our proof of Theorem 4 is the following claim. We will only make use of the ‘moreover’ part but we need the main statement to apply induction. Let S be a B_h -set and let D be its difference set as before. By a *circuit*, we mean a closed walk in which vertices may be repeated but edges may not.

Claim 3 (Cycle fact). *Let $2 \leq \ell \leq h$ be integers. Suppose that D contains not necessarily distinct d_1, \dots, d_ℓ such that for some $1 \leq s \leq \ell - 1$ we have*

$$d_1 + \dots + d_s = d_{s+1} + \dots + d_\ell. \quad (3)$$

Let M be the multigraph on vertex set S built by adding, for each $1 \leq i \leq \ell$, an edge e_i between the pair of vertices whose difference is d_i . Then the edge set of M can be decomposed into a collection of circuits. Moreover, if ℓ is odd, then the graph obtained from M by deleting duplicate edges contains an odd cycle.

Proof. We prove that the edge set of M can be decomposed into a collection of circuits by induction on ℓ . If $\ell = 2$, then we have $d_1 = d_2$ and M is just a pair of edges between the same two vertices.

Suppose that $\ell \geq 3$. Let $x_1, \dots, x_\ell, y_1, \dots, y_\ell \in S$ be such that $d_i = y_i - x_i$ for all $1 \leq i \leq \ell$. From equation (3) we obtain

$$y_1 + \dots + y_s + x_{s+1} + \dots + x_\ell = x_1 + \dots + x_s + y_{s+1} + \dots + y_\ell$$

and so, as S is a B_h -set and $\ell \leq h$, $y_1, \dots, y_s, x_{s+1}, \dots, x_\ell$ must be $x_1, \dots, x_s, y_{s+1}, \dots, y_\ell$ in some order. In particular, $y_1 = x_i$ for some $1 \leq i \leq s$ or $y_1 = y_j$ for some $s+1 \leq j \leq \ell$. In the first case we may assume by relabelling that $y_1 = x_2$ ($y_1 \neq x_1$ as $0 \notin D$). Then $d := d_1 + d_2 = y_2 - x_1$ which is in D as $y_2 > x_2 = y_1 > x_1$. Moreover, $d + d_3 + \dots + d_s = d_{s+1} + \dots + d_\ell$. Thus, by induction, the edges of the multigraph obtained from M by deleting e_1 and e_2 and adding an edge e between x_1 and y_2 can be decomposed into a collection of circuits. Adding back the edges e_1 and e_2 (which form a path between x_1 and y_2) in place of e , we obtain a suitable decomposition of the edges of M .

In the second case we may assume by relabelling that $y_1 = y_{s+1}$. If $x_1 = x_{s+1}$, then $d_1 = d_{s+1}$ and hence $d_2 + \dots + d_s = d_{s+2} + \dots + d_\ell$. By induction, the edges of M

without e_1 and e_{s+1} can be decomposed into a collection of circuits. Adding the circuit consisting of e_1 and e_{s+1} , we obtain a decomposition of all edges of M . If $x_1 \neq x_{s+1}$, then without loss of generality $x_{s+1} > x_1$ and we have $d := d_1 - d_{s+1} = x_{s+1} - x_1 \in D$ and $d + d_2 + \dots + d_s = d_{s+2} + \dots + d_\ell$. By induction, there is a decomposition of the edges of the multigraph obtained from M by deleting e_1 and e_{s+1} and adding a new edge e incident to x_1 and x_{s+1} . Adding back the edges e_1 and e_{s+1} (which form a path between x_1 and x_{s+1}) in place of e , we obtain a suitable decomposition of the edges of M .

For the ‘moreover’ part of the claim, note that if M has an odd number of edges, then the collection of circuits we obtain must contain a circuit with an odd number of edges. Any such circuit contains a cycle of odd length, which completes the proof of the claim. \square

4.2 Proof of Theorem 4

Fix a graph F with at least one odd cycle, let $f = |V(F)|$ and let h be the odd girth of F . Let p be a prime large enough that $\lceil [p/h] \rceil$ has a B_h -subset S of size f that contains 1 and 2. Let the elements of S be $s_1 = 1, s_2 = 2, \dots, s_f$ in ascending order and let the difference set of S be $D = \{s_j - s_i : i < j\}$. As F has at least one edge we may take a copy F^* of F with vertex set S which includes the edge $s_1 s_2$. We write

$$E = \{s_j - s_i : i < j \text{ and } s_i s_j \in E(F^*)\}$$

for the set of distances corresponding to edges of F^* . Since $s_1 = 1$ and $s_2 = 2$, we have $1 \in E$. Note that $E \subseteq D$. Let $\overline{D} = D + p\mathbb{Z}$ and $\overline{E} = E + p\mathbb{Z}$.

Let \overrightarrow{X}_n be a digraph given by applying Lemma 10 with $g = h$. Recall that (a) the underlying graph X_n has chromatic number n , (b) the orientation \overrightarrow{X}_n is acyclic, (c) for any pair of vertices u, v , there is at most one directed path from u to v in \overrightarrow{X}_n , and (d) for every cycle in X_n , the corresponding oriented cycle in \overrightarrow{X}_n has at least h changes of direction. For vertices $u < v$ (in the reachability ordering), write $d(u, v)$ for the length of the unique directed path from u to v in \overrightarrow{X}_n . We now define a digraph \overrightarrow{G} with vertex set $V(\overrightarrow{X}_n)$ as follows.

The edge \overrightarrow{uv} is present if $u < v$ and $d(u, v) \in \overline{E}$.

As before, the digraph \overrightarrow{G} produced is acyclic. We now show that G , its underlying undirected graph, satisfies the properties stated in Theorem 4. Firstly since $1 \in E$, the graph X_n is a subgraph of G so $\chi(G) \geq n$. Next we check that every induced subgraph of G with sufficiently large chromatic number contains an induced copy of F . The argument is almost identical to the corresponding argument at the end of Section 3.2. We colour each edge uv of G (and \overrightarrow{G}), where $u < v$, with the colour $d(u, v) \bmod p$. As before, any monochromatic directed path in \overrightarrow{G} of length pf gives rise to an induced copy of F . Let H be an induced subgraph of G which does not contain an induced copy of F . Each subgraph of H consisting of i -coloured edges is pf -colourable by Proposition 9 and so, by a product colouring, H is itself c'_F -colourable for $c'_F = (pf)^{|E|}$.

It remains to check that the odd girth of G is at least h . Suppose for a contradiction that G contains an odd cycle $v_0 v_1 \dots v_{\ell-1}$ of length $\ell < h$. Since \overrightarrow{G} contains no directed

cycles, without loss of generality the path $v_{\ell-1}v_0v_1$ in G is not a directed path in \vec{G} . We may further assume that since ℓ is odd, the path $v_0v_1v_2$ in G is a directed path in \vec{G} from v_0 to v_2 . For each edge v_iv_{i+1} in this cycle (where here and throughout we take subscript addition in the v_i to be modulo ℓ), there is a directed path \vec{P}_i from v_i to v_{i+1} or from v_{i+1} to v_i in \vec{X}_n depending on the direction of the edge between v_i and v_{i+1} in \vec{G} . For each i let P_i be the undirected path underlying \vec{P}_i and concatenate these paths P_i in the natural way to obtain a walk W in X_n which begins at v_0 , visits $v_1, \dots, v_{\ell-1}$ in turn, and then finally returns to v_0 .

For each traversal of an edge in the walk, when the direction of traversal is the same as the direction of the edge in \vec{X}_n , we consider the traversal of this edge in the walk to be ‘coloured black’. When the two directions are different we consider the traversal of the edge to be ‘coloured red’. For each of the paths P_i , the walk’s traversals corresponding to P_i are all the same colour, and we have assumed that the colours for P_0 and P_1 are both black, so the walk changes colour at most $\ell - 2$ times.

Consider the subgraph L of G consisting of all vertices and edges contained in the walk W . We examine two cases based on whether or not L contains a cycle.

Case 1: Suppose that L is acyclic. Then L is a tree. Let $e = uv$ be an edge of L . The graph $L - e$ (i.e. the graph with vertex set $V(L)$ and edge set $E(L) \setminus \{e\}$) has exactly two components, one of which contains u and the other of which contains v . It is clear that the traversals of the edge e in the walk W alternate between traversing from u to v and traversing from v to u . Moreover, since the walk starts and ends at v_0 , there are an equal number of traversals of e of in each direction. In other words, the total number of traversals of e coloured black is equal to the total number coloured red. Since this is true for all edges e , the total number of black traversals of edges in the walk is equal to the total number of red traversals.

Let $S_1 = \{0 \leq i \leq \ell - 1 : \overrightarrow{v_iv_{i+1}} \in E(\vec{G})\}$ and $S_2 = \{0 \leq i \leq \ell - 1 : \overrightarrow{v_{i+1}v_i} \in E(\vec{G})\}$ so that S_1 and S_2 form a partition of $\{0, 1, \dots, \ell - 1\}$. For $0 \leq i \leq \ell - 1$, let d_i be the length of the path P_i . Then the total number of black traversals is equal to the sum of the d_i for $i \in S_1$ and the total number of red traversals is the sum of the d_i for $i \in S_2$. Hence by the above

$$\sum_{i \in S_1} d_i = \sum_{i \in S_2} d_i.$$

Since the edges v_iv_{i+1} are all present in G , we have $d_i \in \overline{E}$ for all i . Let $\overline{d}_i \equiv d_i \pmod{p}$ be chosen so that $0 \leq \overline{d}_i \leq p - 1$. Then $\overline{d}_i \in E$. The sums $\sum_{i \in S_1} \overline{d}_i$ and $\sum_{i \in S_2} \overline{d}_i$ are the same modulo p and each have fewer than h terms, all of which are at most p/h . It follows that, in fact,

$$\sum_{i \in S_1} \overline{d}_i = \sum_{i \in S_2} \overline{d}_i.$$

Noting that $E \subseteq D$, we can apply the cycle fact to show that the graph on vertex set S with edges between pairs of vertices at distances $\overline{d}_0, \dots, \overline{d}_{\ell-1}$ contains an odd cycle of length at most ℓ . Since $\overline{d}_0, \dots, \overline{d}_{\ell-1} \in E$, by the definition of E and the fact that S is

a Sidon set, F^* contains a cycle of this length, which is a contradiction. This completes the analysis of the case where L does not contain a cycle.

Case 2: Suppose that L contains a cycle. Let Γ be the first cycle produced by the walk, and label its vertices as c_0, c_1, \dots, c_{r-1} in cyclic order around Γ , where c_0 is the vertex amongst these that W arrives at first. Let $C = \{c_0, \dots, c_{r-1}\}$. We will show that Γ contains fewer than h changes of direction, which contradicts property (d) of \vec{X}_n . In what follows, we are only concerned with the portion of walk W from its first visit to c_0 to the point when cycle Γ is formed. Let W' be this segment of W . In particular, c_0 is the first vertex of W' and the final edge traversal in W' completes cycle Γ .

Suppose that as we travel along W' , there is an occasion on which we arrive at c_i for some $0 \leq i \leq r-1$ and the next vertex in C that we visit is c_j for some $j \notin \{i-1, i, i+1\}$, where here and throughout we take addition in the subscripts of the c_i to be modulo r . Note that W' does not terminate when it reaches c_j since the final edge it traverses must be an edge of Γ . This portion of the walk contains a path P from c_i to c_j which avoids every vertex in $C \setminus \{c_i, c_j\}$. Consider the sets of edges $\{c_i c_{i+1}, c_{i+1} c_{i+2}, \dots, c_{j-1} c_j\}$ and $\{c_i c_{i-1}, c_{i-1} c_{i-2}, \dots, c_{j+1} c_j\}$. Each of these sets of edges form a cycle with path P , so at the first time after the formation of P at which all the edges in one set have been traversed, the graph walked so far contains a cycle. Since the union of these sets is the edge set of Γ , this occurs strictly before the formation of Γ , which is a contradiction. Thus, after W' visits c_i , the next vertex of C that it visits is one of c_{i-1} , c_i and c_{i+1} .

Now suppose that after visiting c_i for some $0 \leq i \leq r-1$, the next vertex of C that walk W' visits is c_{i+1} , and suppose further that it arrives at c_{i+1} via an edge other than $c_i c_{i+1}$. Clearly it also does not arrive at c_{i+1} via edge $c_{i+2} c_{i+1}$, so it arrives via an edge not in Γ and hence W' does not terminate when it reaches c_{i+1} . The portion of the walk between c_i and c_{i+1} contains a path P' from c_i to c_{i+1} which avoids every vertex in $C \setminus \{c_i, c_{i+1}\}$ and avoids the edge $c_i c_{i+1}$. Similarly to above, at the first time after the formation of P' at which either $c_i c_{i+1}$ has been traversed or all of $c_i c_{i-1}, \dots, c_{i+2} c_{i+1}$ have been traversed, the graph walked by W' contains a cycle. This occurs strictly before the formation of Γ , which is a contradiction.

Similarly, if the next vertex of C that W' visits after c_i is c_{i-1} , then it arrives at c_{i-1} via the edge $c_i c_{i-1}$. These facts combined imply that there are vertices $\gamma_0, \dots, \gamma_{s-1}$ of C , with $\gamma_0 = c_0$ and γ_{j+1} a neighbour of γ_j in Γ , such that W' has the following form. It is a (possibly empty) walk from γ_0 to itself avoiding $C \setminus \{\gamma_0\}$, followed by a traversal of the edge $\gamma_0 \gamma_1$, followed by a (possibly empty) walk from γ_1 to itself avoiding $C \setminus \{\gamma_1\}$, followed by a traversal of the edge $\gamma_1 \gamma_2$, and so on, concluding with a traversal of the edge $\gamma_{s-2} \gamma_{s-1}$, which completes cycle Γ . Note that the vertices $\gamma_0, \gamma_1, \dots, \gamma_{s-1}$ form a walk W'' on Γ traversing all the edges of Γ . Let the traversals of each edge have the same colour in W'' as the corresponding traversals in W' . Note that W'' changes colour at most as many times as W' .

We say that a vertex $c_i \in C$ is *crossed* by the walk W'' if c_{i-1}, c_i, c_{i+1} are consecutive vertices of W'' in either order. If there are two vertices of C not crossed by W'' , then W'' does not traverse all edges of Γ , so there is at most one vertex not crossed by W'' .

Now, if the edges $c_{i-1}c_i$ and $c_i c_{i+1}$ are a change of direction, then when W'' crosses c_i the traversals of these two edges are of different colours. It follows that every time W'' crosses c_i , it changes colour between the traversals of the two edges. Hence, since W'' changes colour at most $h - 2$ times and all but at most one of the vertices in C are crossed, there are at most $h - 1$ changes of direction in Γ . We have thus obtained a cycle in $\overrightarrow{X_n}$ with fewer than h changes of direction, which gives the required contradiction. Therefore, the odd girth of G is at least h .

5 Tournaments

We now prove Theorem 7, which is restated below. Our strategy is akin to that used by Alon, Pach and Solymosi in [1]. Recall that a *tournament* is a complete graph in which each edge has an orientation, and that a tournament is *transitive* if, for all distinct vertices u, v, w , the presence of the edges \overrightarrow{uv} and \overrightarrow{vw} in the tournament implies the presence of the edge \overrightarrow{uw} . The *chromatic number* $\chi(T)$ of a tournament T is the smallest possible number of parts in a partition of its vertex set in which each part induces a transitive tournament. We will prove the following corollary of Theorem 3.

Theorem 7. *For every tournament T , there is a constant C_T and tournaments S of arbitrarily large chromatic number such that every T -free subtournament of S has chromatic number at most C_T .*

Note that there are certain tournaments which are contained in every tournament of sufficiently large chromatic number; all such tournaments were described explicitly in [3]. Before proving the corollary, we state two results we will need. The first is the following standard corollary of Dilworth's theorem [9].

Lemma 13. *Let (a_1, a_2, \dots, a_k) be a sequence of distinct elements and let $<$ be a strict total order on them. If this sequence has no decreasing (with respect to $<$) subsequence of length greater than m , then it can be partitioned into at most m increasing subsequences.*

We will also use the following result originally proved by Rödl and Winkler in [29]. Recall that an *ordering* of a graph or digraph G is a strict total ordering of $V(G)$, and an *ordered (di)graph* is a pair $(G, <)$ where G is a (di)graph and $<$ is an ordering of G .

Theorem 14 (Rödl and Winkler [29]). *For any ordered graph $(B, <)$, there exists a graph F such that, for every ordering $<'$ of F , $(B, <)$ is an ordered induced subgraph of $(F, <')$.*

We are now ready to prove Theorem 7.

Proof of Theorem 7. Fix an ordered tournament $(T, <)$ with vertices $x_1 < \dots < x_{|T|}$. Let $(B, <)$ be the ordered back-edge graph of $(T, <)$, that is, the ordered graph on the same vertex set as T and with the same ordering as T , where $x_i x_j \in E(B)$ for $i < j$ exactly when $\overrightarrow{x_j x_i} \in E(T)$. Let F be the graph obtained by applying Theorem 14 to $(B, <)$. Now apply Theorem 3 to F . We obtain a constant $C_T = c_F$ and a graph L of (arbitrarily large) chromatic number n with $\omega(L) = \omega(F) = \omega$ such that every induced

subgraph of L with chromatic number greater than C_T contains an induced copy of F . Fix an arbitrary ordering $<'$ of L and define an ordered tournament $(S, <')$ with the same vertex set as L and the same ordering as L by orienting \vec{xy} for $x <' y$ if $xy \notin E(L)$ and \vec{yx} otherwise (so that $(L, <')$ is the ordered back-edge graph of $(S, <')$).

We now show that $\chi(S) \geq n/\omega$, and hence that $\chi(S)$ is arbitrarily large. Consider a transitive subtournament A in S and write its vertices as a sequence $(a_1, a_2, \dots, a_{|A|})$ such that $\vec{a_i a_j} \in E(S)$ for each $i < j$. Observe that this sequence has no decreasing (with respect to $<'$) subsequence of length $\omega + 1$ as this would give rise to an $(\omega + 1)$ -clique in L . Hence, by Lemma 13, we may partition the sequence into at most ω increasing subsequences. Each of these corresponds to an independent set in L and hence, $\omega \cdot \chi(S) \geq n$ as required.

It remains to show that every subtournament R of S with $\chi(R) > C_T$ contains a copy of T . If R is such a subtournament, then clearly the induced subgraph R' of L with the same vertex set as R has $\chi(R') > C_T$. Hence, by our assumptions on L , R' contains an induced copy of F . Consider the ordering of F given by the restriction of $<'$ to this copy. By the construction of F , under this ordering it contains an induced ordered copy of $(B, <)$ which guarantees the existence of a copy of T in R as required. This completes the proof of the theorem. \square

6 Hypergraphs

In this section we prove Theorem 8, which is restated below. Recall that in this paper all hypergraph edges have size at least two. For a hypergraph \mathcal{G} and a set $X \subseteq V(\mathcal{G})$ of vertices of \mathcal{G} , the *subhypergraph of \mathcal{G} induced on X* , denoted $\mathcal{G}[X]$, is the hypergraph with vertex set X and edge set $\{e: e \subseteq X, e \in E(\mathcal{G})\}$. For hypergraphs \mathcal{F} and \mathcal{G} , we will say that \mathcal{G} is *\mathcal{F} -free* if it does not contain \mathcal{F} as an induced subhypergraph. Recall from the introduction that a hypergraph is said to be *strongly t -colourable* if its vertices can be t -coloured such that no edge contains two vertices of the same colour, and that a hypergraph *covers* a pair of vertices if it has an edge containing both of them.

Theorem 8. *For every hypergraph \mathcal{F} with at least one edge, there is a constant $c_{\mathcal{F}}$ and hypergraphs \mathcal{G} of arbitrarily large chromatic number such that every \mathcal{F} -free induced subhypergraph of \mathcal{G} is strongly $c_{\mathcal{F}}$ -colourable. Moreover, we can take \mathcal{G} such that if \mathcal{G} covers every pair from some $X \subseteq V(\mathcal{G})$, then $\mathcal{G}[X]$ is an induced subhypergraph of \mathcal{F} .*

The proof of Theorem 8 resembles those of Theorems 3 and 4 in that we start with an oriented base (hyper)graph with large chromatic number and certain ‘reachability’ properties, to which we add edges based on ‘distances’ between vertices modulo some prime p . As before, we will ascertain which edges to add to the base hypergraph by placing a copy of \mathcal{F} on a B_3 -set, an approach reminiscent of that used in the earlier proofs. Once we have constructed \mathcal{G} , the arguments we use to show it has the desired properties also bear strong similarities to those used above.

We now introduce an analogue of reachability and distances between vertices in the

hypergraph setting. Given a hypergraph \mathcal{G} and a total order \prec on its vertex set, we construct the \prec -digraph \vec{G} on vertex set $V(\mathcal{G})$ by adding to $E(\vec{G})$ the directed edges $\overrightarrow{v_1v_2}, \overrightarrow{v_2v_3}, \dots, \overrightarrow{v_{a-1}v_a}$ for each edge $e = \{v_1 \prec v_2 \prec \dots \prec v_a\}$ of \mathcal{G} . The \prec -graph G is the underlying undirected graph of \vec{G} . If \mathcal{G} has girth at least 3, then no two distinct edges intersect in more than one vertex, and it follows that every pair of distinct vertices has at most one directed edge between them in \vec{G} . Also, in this setting every edge in \vec{G} comes from a unique edge of \mathcal{G} .

The base hypergraph for our construction should have large chromatic number as well as an ordering \prec on its vertex set such that its \prec -digraph is suitable for defining reachability and distance. The lemma below states that such hypergraphs exist.

Lemma 15. *For all integers $k \geq 2$, $g \geq 3$ and $n \geq 2$, there exists a k -uniform hypergraph \mathcal{Y}_n with chromatic number n and girth at least g , and an order \prec on its vertex set such that, for every cycle in its \prec -graph Y_n , the corresponding oriented cycle in its \prec -digraph \vec{Y}_n has at least g changes of direction.*

Proof. Nešetřil and Rödl [26] proved that for all integers $k \geq 2$, $g \geq 3$ and $n \geq 2$ there is a k -uniform hypergraph with chromatic number n and girth at least g which is strongly a -colourable⁴ for $a = (k - 1)n + 1$.

Fix integers $k \geq 2$, $g \geq 3$, and $n \geq 2$, and let $a = (k - 1)n + 1$. By the result of Nešetřil and Rödl, there is a k -uniform hypergraph \mathcal{Y}_n with chromatic number n and girth greater than $(g - 1)(a - 1)$ which is strongly a -colourable. Fix a strong a -colouring $c: V(\mathcal{Y}_n) \rightarrow \{1, \dots, a\}$ of \mathcal{Y}_n . Let \prec be any ordering of $V(\mathcal{Y}_n)$ in which every vertex of colour i precedes every vertex of colour j whenever $i < j$. Let Y_n and \vec{Y}_n respectively be the \prec -graph and \prec -digraph of \mathcal{Y}_n .

Consider any edge \overrightarrow{uv} in \vec{Y}_n . There is some edge of \mathcal{Y}_n containing both u and v and so $c(u) \neq c(v)$. By the choice of \prec , we must have $c(u) < c(v)$. Letting $u_1u_2 \dots u_t$ be any directed path in \vec{Y}_n , by the preceding discussion we have $1 \leq c(u_1) < \dots < c(u_t) \leq a$. Hence, the length of any directed path in \vec{Y}_n is at most $a - 1$.

Consider any cycle C in Y_n : each edge of this cycle comes from an edge of a cycle in \mathcal{Y}_n and so C has length greater than $(g - 1)(a - 1)$. Every directed path in \vec{Y}_n has length at most $a - 1$, so C has at least g changes of direction in \vec{Y}_n , as required. \square

We are now ready to prove Theorem 8.

Proof of Theorem 8. Fix a hypergraph \mathcal{F} with at least one edge. Let f be the number of vertices of \mathcal{F} and let $n \geq 2$. Let $m \geq 2$ be the minimum edge size of \mathcal{F} . Let \mathcal{Y}_n be an m -uniform hypergraph of chromatic number at least n and girth at least four given by Lemma 15, with an ordering \prec on its vertex set such that in the associated digraph \vec{Y}_n every cycle has at least three changes of direction. In particular, in \vec{Y}_n there are no directed cycles and there is at most one directed path between any pair of vertices. If

⁴Although the strongly a -colourable condition does not appear in their theorem statement, it is emphasised in their proof.

there is a directed path from u to v in \vec{Y}_n , then we write $u < v$ and define the distance between u and v , written $d(u, v)$, to be the length of this path.

We will define a hypergraph \mathcal{G} satisfying the conditions of Theorem 8 by adding edges to \mathcal{Y}_n as follows. Start by picking a prime p such that $\lfloor p/(4m) \rfloor$ contains a B_3 -set S' of size f , then define $S = 2m \cdot S' = \{2ms' : s' \in S'\} \subseteq \lfloor (p-1)/2 \rfloor$. Let the elements of S be s_1, \dots, s_f in ascending order and let $D = \{s_j - s_i : i < j\}$ be the difference set of S . Note that all elements of D are positive multiples of $2m$ and are at most $(p-1)/2$.

Taking a copy \mathcal{F}^* of \mathcal{F} with vertex set S , we define

$$E = \{(s_{i_2} - s_{i_1}, s_{i_3} - s_{i_2}, \dots, s_{i_a} - s_{i_{a-1}}) : i_1 < i_2 < \dots < i_a \text{ and } \{s_{i_1}, s_{i_2}, \dots, s_{i_a}\} \in E(\mathcal{F}^*)\}$$

to be the analogue of the sets of allowable distances from the proofs of Theorems 3 and 4. Note that every entry of every tuple in E is in D . Let $\overline{D} = D + p\mathbb{Z}$ and $\overline{E} = \{(d_1 + k_1p, \dots, d_{a-1} + k_{a-1}p) : k_1, \dots, k_{a-1} \in \mathbb{Z} \text{ and } (d_1, \dots, d_{a-1}) \in E\}$. Construct hypergraph \mathcal{G} by *adding* edges to \mathcal{Y}_n as follows.

$$\begin{aligned} \text{The edge } \{v_1, v_2, \dots, v_a\} \text{ is added if } v_1 < v_2 < \dots < v_a \text{ and} \\ (d(v_1, v_2), d(v_2, v_3), \dots, d(v_{a-1}, v_a)) \in \overline{E}. \end{aligned}$$

There are two types of edges in \mathcal{G} : those which appear in \mathcal{Y}_n and the newly added ones. By construction, if two vertices are in a common edge, then they are $<$ -comparable. The edges which appear in \mathcal{Y}_n are of the form $\{v_1, \dots, v_m\}$ where $d(v_i, v_{i+1}) = 1$ for all i . Hence, if vertices $u < v$ are both in some edge of \mathcal{Y}_n , then $d(u, v) \in [m-1]$. On the other hand, if $u < v$ are both in an edge that has been newly added, then $d(u, v) \in \overline{D}$ and so $d(u, v) \geq 2m$. In particular, if vertices $u < v$ are both in some edge of \mathcal{G} , then $d(u, v)$ is in $\overline{D} \cup [m-1]$ and so is not divisible by p .

We will now show that this hypergraph satisfies the conditions of Theorem 8. Firstly, since it contains \mathcal{Y}_n , it certainly has chromatic number at least n . Next, we will show that every \mathcal{F} -free induced subhypergraph of \mathcal{G} has bounded strong chromatic number. We begin by defining L to be the graph with vertex set $V(\mathcal{G})$ where uv is an edge if u and v are in some common edge of \mathcal{G} . For each edge uv of L , u and v are $<$ -comparable. Orient the edge from u to v if $u < v$, and from v to u otherwise. Denote the resulting digraph by \vec{L} .

All edges \vec{uv} of \vec{L} have endpoints which satisfy $d(u, v) \in \overline{D} \cup [m-1]$, and we colour the edges of \vec{L} (and L) by the residue modulo p of this distance. We claim that the vertex set of a monochromatic directed path of length pf in \vec{L} has a subset which induces a copy of \mathcal{F} in \mathcal{G} . Indeed, by repeatedly jumping at most $p-1$ steps along such a path we can find a sequence of vertices $u_1 < u_2 < \dots < u_f$ in the path such that $d(u_i, u_{i+1}) \equiv s_{i+1} - s_i \pmod{p}$ for all i . Note that, for $i < j$, $d(u_i, u_j) \equiv s_j - s_i \pmod{p}$ and so $d(u_i, u_j) \in \overline{D}$. Hence, $d(u_i, u_j)$ is never 1 and so $\{u_1, \dots, u_f\}$ does not contain an edge of \mathcal{Y}_n . Thus, by the construction of \mathcal{G} , the vertices u_1, \dots, u_f induce a copy of \mathcal{F} , as claimed.

Now let \mathcal{H} be an \mathcal{F} -free induced subhypergraph of \mathcal{G} . By the above, for each $i \in D \cup [m-1]$ there is no i -coloured directed path of length pf in $\vec{L}[V(\mathcal{H})]$, so the i -coloured subgraph

of $L[V(\mathcal{H})]$ is (pf) -colourable by Proposition 9. A product colouring demonstrates that $L[V(\mathcal{H})]$ is $(pf)^{p-1}$ -colourable. This colouring is a strong colouring of \mathcal{H} . Hence we may take $c_{\mathcal{F}} = (pf)^{p-1}$, which is a constant that only depends on \mathcal{F} .

It remains to show the moreover part of the statement. Let $X = \{v_1, \dots, v_t\}$ be a set of vertices every pair of which is covered by \mathcal{G} . If $t \leq 1$, then the result is clear so assume $t \geq 2$. It follows from the fact that \mathcal{G} covers every pair from X that each pair is $<$ -comparable; we may assume that $v_1 < \dots < v_t$. Define $d_i = d(v_i, v_{i+1})$ for all i and note that since $d(v_i, v_j) \in \overline{D} \cup [m-1]$ for all $i < j$, all sums of the form

$$\sum_{i_1 \leq j \leq i_2} d_j \tag{4}$$

with $1 \leq i_1 \leq i_2 \leq t-1$ are in $\overline{D} \cup [m-1]$.

Suppose that $d_i \leq m-1$ for some i . If $i < t-1$, then d_{i+1} and $d_i + d_{i+1}$ differ by at most $m-1$. This cannot occur if $d_{i+1} \in \overline{D}$, so we conclude that $d_{i+1} \leq m-1$. Similarly, if $i > 1$ then $d_{i-1} \leq m-1$. Repeating this argument we find that $d_1, \dots, d_{t-1} \in [m-1]$. Since the minimal positive element of \overline{D} is at least $2m$, it follows that $d_1 + d_2 \notin \overline{D}$, so $d_1 + d_2 \in [m-1]$. Applying this repeatedly we find that all sums of the form in (4) are in $[m-1]$, and hence no pair of vertices from X appear together in an edge of \mathcal{G} which is not in \mathcal{Y}_n . In particular, \mathcal{Y}_n covers all pairs from X .

Let e be an edge of \mathcal{Y}_n containing both v_1 and v_2 , and suppose there is some v_a not in e . Let e_1, e_2 be edges of \mathcal{Y}_n such that e_i contains v_a and v_i . If $e_1 = e_2$, then the edges e and e_1 form a 2-cycle in \mathcal{Y}_n , while if $e_1 \neq e_2$, then the edges e, e_1, e_2 form a 3-cycle in \mathcal{Y}_n . As \mathcal{Y}_n has girth at least four, neither of these cases is possible. We deduce that $X \subseteq e$, and in particular that $t \leq m$. If $t < m$, then the subhypergraph of \mathcal{G} induced on X has no edges and thus is clearly an induced subhypergraph of \mathcal{F} . If $t = m$, then the subhypergraph of \mathcal{G} induced on these vertices has a single edge of size m , which again is clearly an induced subhypergraph of \mathcal{F} .

On the other hand, if $d_i > m-1$ for all i , then all the sums of the form in (4) are in \overline{D} . Writing \overline{d}_i for the residue modulo p of d_i , we have $\overline{d}_i \in D$ for all i and all sums of the form

$$\sum_{i_1 \leq j \leq i_2} \overline{d}_j$$

with $1 \leq i_1 \leq i_2 \leq t-1$ are in \overline{D} . We can deduce from the fact that $D \subseteq [(p-1)/2]$ that in fact all of these sums are in D . It now follows from the clique fact that there exist $b_1 < \dots < b_t$ in S such that $b_{i+1} - b_i = \overline{d}_i$ for all i . Since $\overline{d}_i \neq 1$ for all i , the subhypergraph of \mathcal{G} induced on X does not include any edge of \mathcal{Y}_n , so by the construction of \mathcal{G} and the fact that S is a Sidon set, the subhypergraph of \mathcal{G} induced on X is isomorphic to that of \mathcal{F}^* induced on $\{b_1, \dots, b_t\}$, and the claim follows. \square

7 Extensions to infinite families of graphs

Recall that the *disjoint union* of a family of graphs is the (possibly infinite) graph consisting of pairwise vertex-disjoint copies of the graphs in the family with no edges between copies. As noted in the introduction, for a finite family of graphs \mathcal{F} (at least one of which contains an edge), an easy consequence of Theorem 3 is that there is a constant $c_{\mathcal{F}}$ and graphs G of arbitrarily large chromatic number and the same clique number as the largest clique number of a graph in \mathcal{F} such that every induced subgraph of G with chromatic number greater than $c_{\mathcal{F}}$ contains every member of \mathcal{F} as an induced subgraph. Taking a disjoint union of the graphs G of arbitrarily large chromatic number gives the following corollary.

Corollary 16. *Let $\omega \geq 2$, and suppose that \mathcal{F} is a finite family of finite graphs all with clique number at most ω . There exists a constant $c_{\mathcal{F}}$ and an infinite graph G of infinite chromatic number and clique number at most ω such that every induced subgraph of G with chromatic number greater than $c_{\mathcal{F}}$ contains every member of \mathcal{F} as an induced subgraph.*

Does this phenomenon occur for infinite families of graphs? To this end we will say that a (possibly infinite) family of graphs \mathcal{F} is *durable* if there exists an infinite graph G with infinite chromatic number such that every infinite-chromatic induced subgraph of G contains every member of \mathcal{F} as an induced subgraph. If such a G exists we will call it a *witness* of \mathcal{F} 's durability. We note a few properties of witnesses and durability.

- A subfamily of a durable family is durable.
- Every infinite-chromatic induced subgraph of a witness is also a witness.
- Every witness contains every graph in \mathcal{F} as an induced subgraph.

Corollary 16 says that every finite family of graphs is durable. Moreover, if at least one graph in the finite family contains an edge, then there is a witness whose clique number is no bigger than the largest clique number of a member of the family.

We are interested in which infinite families of finite graphs are durable. There certainly are some durable infinite families, for example any family whose members are all disjoint unions of complete graphs is durable as witnessed by the disjoint union of K_1, K_2, K_3, \dots . This example suggests that for a durable family with unbounded chromatic number, the disjoint union of the members of the family might be a witness of the family's durability. If this were the case, then this disjoint union would be a minimal witness. The following result confirms this hypothesis.

Theorem 17. *Let \mathcal{F} be a countable family of finite graphs. Let $G_{\mathcal{F}}$ be the disjoint union of the members of \mathcal{F} .*

- (a) *If \mathcal{F} is durable, then every witness of \mathcal{F} 's durability contains $G_{\mathcal{F}}$ or K_{∞} as an induced subgraph.*
- (b) *If $\{\chi(F) : F \in \mathcal{F}\}$ is unbounded, then \mathcal{F} is durable if and only if $G_{\mathcal{F}}$ is a witness of \mathcal{F} 's durability.*

Theorem 17(b) shows that if a family of unbounded chromatic number is durable, then

it has a witness whose clique number is no bigger than the largest clique number of a member of the family. An immediate application of Theorem 17 is Theorem 5: the family, \mathcal{F}_g , of finite graphs of girth at least g is not durable. Indeed, this family has unbounded chromatic number and so if it were durable, then $G_{\mathcal{F}_g}$ would witness this. However, the disjoint union of graphs of girth at least $g + 1$ is an infinite-chromatic induced subgraph of $G_{\mathcal{F}_g}$ that does not contain the cycle of length g .

The following lemma will be needed in the proof of Theorem 17.

Lemma 18. *Every infinite graph either contains an infinite clique, has bounded clique number or contains the disjoint union of all finite cliques as an induced subgraph.*

Proof. Let G be an infinite graph with unbounded clique number. Let $(a_i)_{i \geq 1}$ be a sequence of natural numbers in which each term is taken to be large relative to all previous terms. Fix a copy of K_{a_1} in G on vertex set V_1 . By identifying a copy of $K_{a_1+a_2}$ in G , we can find a copy of K_{a_2} on vertex set V_2 disjoint from V_1 . Continuing in this manner, we obtain a sequence $(V_i)_{i \geq 1}$ of pairwise disjoint sets of vertices of G such that V_i induces a copy of K_{a_i} for each i .

Select a vertex $v \in V_1$, and for each $i \geq 2$ remove at most half of the vertices from V_i so that v is either adjacent to every remaining vertex or not adjacent to any of them. Repeat this process for each of the other vertices in V_1 , removing at most half of the vertices currently in each V_i at each step, so that at the end of the process the remaining sets of vertices $V'_i \subseteq V_i$ satisfy $|V'_i| \geq |V_i|/2^{a_1}$ and have the property that for all $i \geq 2$ and $v \in V_1$, v is either adjacent to every vertex in V'_i or not adjacent to any of them.

For each $i \geq 2$ we can now define a vector $b_i \in \{0, 1\}^{a_1}$ by letting the j th entry be 1 if the j th vertex of V_1 is adjacent to every vertex in V'_i and 0 otherwise. There are only finitely many such vectors, so there exists a sequence (i_k) such that all vectors b_{i_k} agree. By relabelling, we may therefore assume that all vectors b_i are equal to some b . If at least half the entries in b are 1's, then by considering the vertices in V_1 corresponding to 1 entries, for some $t \geq a_1/2$ we obtain a copy of K_t in G , every vertex of which is adjacent to every vertex in $\bigcup_{i \geq 2} V'_i$. Otherwise, for some $t \geq a_1/2$ we obtain a copy of K_t none of whose vertices have a neighbour in this union.

Repeat this process for V'_2 , then what remains of V'_3 , and so on. If the a_i are chosen appropriately, then it follows that there exists a sequence of pairwise disjoint sets of vertices of G , $(W_n)_{n \geq 1}$, such that W_n induces a copy of K_n in G and the possible edges between W_n and $\bigcup_{i > n} W_i$ are either all present in G or all not present in G . If these edges are present for infinitely many n , then G contains an infinite clique. If the edges are present for only finitely many n , then there is an induced subgraph of G consisting of a disjoint union of arbitrarily large cliques, which in particular contains the disjoint union of all finite cliques as an induced subgraph. \square

We are ready to prove Theorem 17. For a vertex v in a (possibly infinite) graph G , let $\Gamma(v) := \{u \in V(G) : uv \in E(G)\}$ be the *neighbourhood* of v in G . Similarly, for $V \subseteq V(G)$ we write $\Gamma(V) := \{u \in V(G) : uv \in E(G) \text{ for some } v \in V\}$.

Proof of Theorem 17. We first show that (a) implies (b) in the statement of the theorem. Suppose that \mathcal{F} is durable with $\{\chi(F) : F \in \mathcal{F}\}$ unbounded, and let G be a witness of \mathcal{F} 's durability. By part (a), G contains $G_{\mathcal{F}}$ or K_{∞} as an induced subgraph. In the former case, since $\{\chi(F) : F \in \mathcal{F}\}$ is unbounded, $G_{\mathcal{F}}$ has infinite chromatic number and is therefore a witness by one of the properties of witnesses noted above. In the latter case, again by the properties noted above, K_{∞} is a witness of \mathcal{F} 's durability and hence every member of \mathcal{F} is an induced subgraph of K_{∞} . Thus, \mathcal{F} is a family of cliques and it follows from the fact that $\{\chi(F) : F \in \mathcal{F}\}$ is unbounded that $G_{\mathcal{F}}$ is a witness.

It remains to prove the first part of the theorem. Suppose that \mathcal{F} is durable and let G be a witness. By Lemma 18, G either contains an infinite clique, has bounded clique number or contains the disjoint union of all finite cliques as an induced subgraph. In the first case we are done. In the third case, the disjoint union of all finite cliques is a witness of \mathcal{F} 's durability, so all members of \mathcal{F} are disjoint unions of cliques. In particular, the disjoint union of all finite cliques contains $G_{\mathcal{F}}$ as an induced subgraph and so we are done.

Hence, we are left with the case where G has bounded clique number. We may assume that the graph induced on the neighbourhood of each vertex of G is finitely colourable. Indeed, otherwise there exists $v \in V(G)$ such that $G[\Gamma(v)]$ is an infinite-chromatic induced subgraph of G . In this case $G[\Gamma(v)]$ is a witness of \mathcal{F} 's durability which has clique number strictly less than that of G . Repeating this process finitely many times, we obtain an infinite-chromatic induced subgraph of G that is a witness of \mathcal{F} 's durability and in which the neighbourhood of every vertex is finitely colourable.

Enumerate \mathcal{F} as $(F_i)_{i \geq 1}$. Since G is a witness, G contains an induced copy of F_1 , say on vertex set V_1 . By assumption, the chromatic number of the subgraph of G induced on $V_1 \cup \Gamma(V_1)$ is finite. It follows that $G - (V_1 \cup \Gamma(V_1))$ is an infinite-chromatic induced subgraph of G and therefore contains every graph in \mathcal{F} as an induced subgraph. We may now find an induced copy of F_2 in this graph, say on vertex set V_2 , and remove every vertex in V_2 or its neighbourhood from $G - (V_1 \cup \Gamma(V_1))$ to obtain a new infinite-chromatic induced subgraph of G . Continuing in this manner we obtain an induced subgraph of G consisting of the disjoint union of the members of \mathcal{F} , that is, G contains $G_{\mathcal{F}}$ as an induced subgraph. \square

Theorem 17 gives an effective test for whether a family of unbounded chromatic number is durable. Still, it is not clear whether there are durable families of unbounded chromatic number whose members are not all disjoint unions of cliques. Theorem 17 shows that the existence of such a family is equivalent to a positive answer to the following.

Question 19. *Does there exist a sequence of connected finite graphs $(F_i)_{i \geq 1}$, not all of which are cliques, and positive integers $(c_i)_{i \geq 1}$ such that $\chi(F_i) \rightarrow \infty$ and every F_i -free subgraph of F_j is c_i -colourable for all $i < j$?*

We now turn to infinite families with bounded chromatic number. The next result provides some durable families of forests. A graph F is χ -*bounding* if the class of F -free graphs is χ -bounded. It is folklore that any χ -bounding graph must be a forest and the celebrated Gyárfás–Sumner conjecture [13, 38] asserts the converse. We remark that a

beautiful proof of Gyárfás [14] shows that every path is χ -bounding and so the following theorem, combined with the first bullet point above, implies Theorem 6.

Theorem 20. *The family of χ -bounding graphs is durable.*

Proof. Let G be a triangle-free graph with infinite chromatic number. Let H be an induced subgraph of G with infinite chromatic number. Fix a χ -bounding graph F and note that there is some constant c_F such that every triangle-free F -free graph is c_F -colourable. But H is triangle-free and has chromatic number greater than c_F , so it contains F as an induced subgraph. Hence, G witnesses the durability of the family of χ -bounding graphs. \square

It is possible that every family with bounded chromatic number is durable (although we suspect this is not the case). An answer to the following very natural question would resolve this.

Question 21. *Is the family of k -colourable graphs durable? That is, for each positive integer k , is there a graph G_k with infinite chromatic number such that every infinite-chromatic induced subgraph of G contains every finite k -colourable graph as an induced subgraph?*

It would be natural to ask in addition for the graph G_k to have clique number k to align with Corollary 16.

8 Conclusion

Having considered tournaments, hypergraphs and infinite families of graphs, we now focus on the setting of graphs. As discussed in the introduction, Theorem 3 shows that the class \mathcal{C}_r of K_r -free graphs has a far stronger property than just being a vertex Ramsey class, and it is natural to ask which other families have this property.

Question 22. *Which hereditary graph classes \mathcal{C} have the property that for every $F \in \mathcal{C}$ there is a constant c_F and graphs $G \in \mathcal{C}$ of arbitrarily large chromatic number such that every F -free induced subgraph of G is c_F -colourable?*

A good starting point for tackling Question 22 would be to resolve the case when \mathcal{C} is determined by a single excluded graph.

A case of particular interest is when \mathcal{C} is the class of graphs with girth at least g for some $g \geq 3$. Theorem 4 asserts that the class of graphs with *odd girth* at least g has the stated property. As mentioned in the introduction we conjecture that the same is true for girth.

Conjecture 23. *For every graph F with at least one cycle, there exists a constant b_F and graphs G of arbitrarily large chromatic number and the same girth as F such that every F -free induced subgraph of G is b_F -colourable.*

For this conjecture to hold, the class of graphs with girth at least g must certainly be vertex Ramsey. Happily, since cycles are 2-connected this follows from a theorem of Nešetřil and Rödl [24].

The graphs \overrightarrow{X}_n and X_n described in Section 2 can be chosen to not contain short cycles. However, if we continue to build our graphs by adding edges \overrightarrow{uv} for $u < v$ simply based on residues modulo p of $d(u, v)$, we will unavoidably introduce many 4-cycles. For this reason, a resolution of this conjecture even in the special case where F is the 5-cycle would be very interesting, as such an argument would likely overcome many of the difficulties involved in proving the full conjecture.

Another area for further research is determining the optimal constant c_F in the statement of Theorem 3.

Question 24. *For each graph F containing an edge, what is the smallest value of c_F for which Theorem 3 holds? Can we take c_F bounded by a function of $\chi(F)$?*

Clearly c_F must be at least $\chi(F) - 1$ for all F , and at least $\chi(F)$ if F is not vertex-critical. Nešetřil [23] has conjectured that the optimal c_F is at most $\chi(F)$ for all F .

For triangles, it follows from results in [7, 34] that the optimal c_{K_3} is either 3 or 4 (in particular, the bound $\chi(F) - 1$ for vertex-critical graphs is not always attained). For larger cliques, Briański, Davies and Walczak [6] showed that the optimal c_{K_m} is $O(m^3)$; as noted at the start of Section 3, the optimal c_F is $O(|V(F)|^9)$ in general.

Acknowledgements

The authors would like to thank Jarik Nešetřil for helpful discussions and comments, including pointing out a simpler way of obtaining the base structures in Section 2 and Lemma 15, and for suggesting Question 24.

References

- [1] N. Alon, J. Pach and J. Solymosi. Ramsey-type theorems with forbidden subgraphs. *Combinatorica*, **21**(2):155–170, 2001.
- [2] R. C. Baker, G. Harman and J. Pintz. The difference between consecutive primes, II. *Proceedings of the London Mathematical Society*, **83**(3):532–562, 2001.
- [3] E. Berger, K. Choromanski, M. Chudnovsky, J. Fox, M. Loebl, A. Scott, P. Seymour and S. Thomassé. Tournaments and colouring. *Journal of Combinatorial Theory, Series B*, **103**(1):1–20, 2013.
- [4] R. C. Bose and S. Chowla. Theorems in the additive theory of numbers. *Commentarii Mathematici Helvetici*, **37**:141–147, 1962.
- [5] M. Briański, J. Davies and B. Walczak. Separating polynomial χ -boundedness from χ -boundedness. *arXiv:2201.08814v1* preprint, 2022.
- [6] M. Briański, J. Davies and B. Walczak. Separating polynomial χ -boundedness from χ -boundedness. *arXiv:2201.08814v2* preprint, 2022.

- [7] A. Carbonero, P. Hompe, B. Moore and S. Spirkl. A counterexample to a conjecture about triangle-free induced subgraphs of graphs with large chromatic number. *Journal of Combinatorial Theory, Series B*, **158**:63–69, 2023.
- [8] B. Descartes. k -chromatic graphs without triangles. *American Mathematical Monthly*, **61**(5):352–353, 1954.
- [9] R. P. Dilworth. A decomposition theorem for partially ordered sets. *Annals of Mathematics*, **51**(1):161–166, 1950.
- [10] P. Erdős. Graph theory and probability. *Canadian Journal of Mathematics*, **11**:34–38, 1959.
- [11] L. Esperet. Graph colorings, flows and perfect matchings. Habilitation thesis, Université Grenoble Alpes, 2017.
- [12] J. Folkman. Graphs with monochromatic complete subgraphs in every edge coloring. *SIAM Journal on Applied Mathematics*, **18**(1):19–24, 1970.
- [13] A. Gyárfás. On Ramsey covering-numbers. In *Infinite and finite sets (Colloquium, Keszthely, 1973; dedicated to P. Erdős on his 60th birthday)*, Vol. II, Colloquia Mathematica Societatis János Bolyai, Vol. 10, pages 801–816. North-Holland, Amsterdam, 1975.
- [14] A. Gyárfás. Problems from the world surrounding perfect graphs. In *Proceedings of the International Conference on Combinatorial Analysis and its Applications (Pokrzywna, 1985)*, Zastosowania Matematyki, **19**:413–441, 1987.
- [15] A. Hajnal and P. Komjáth. Embedding graphs into colored graphs. *Transactions of the American Mathematical Society*, **307**(1):395–409, 1988.
- [16] H. A. Kierstead. Classes of graphs that are not vertex Ramsey. *SIAM Journal on Discrete Mathematics*, **10**(3):373–380, 1997.
- [17] H. A. Kierstead and W. T. Trotter. Colorful induced subgraphs. *Discrete Mathematics*, **101**:165–169, 1992.
- [18] H. A. Kierstead and Y. Zhu. Classes of graphs that exclude a tree and a clique and are not vertex Ramsey. *Combinatorica*, **16**(4):493–504, 1996.
- [19] P. Komjáth and V. Rödl. Coloring of universal graphs. *Graphs and Combinatorics*, **2**(1):55–60, 1986.
- [20] P. Komjáth and S. Shelah. A consistent edge partition theorem for infinite graphs. *Acta Mathematica Hungarica*, **61**(1-2):115–120, 1993.
- [21] J. Nešetřil. *Teorie grafů*. Praha: Státní Nakladatelství Technické Literatury, 1st edition, 1979.
- [22] J. Nešetřil. Ramsey theory. In *Handbook of Combinatorics*, volume 2, pages 1331–1403. MIT Press, 1996.
- [23] J. Nešetřil, 2023. Personal communication.
- [24] J. Nešetřil and V. Rödl. Partitions of vertices. *Commentationes Mathematicae Universitatis Carolinae*, **17**(1):85–95, 1976.
- [25] J. Nešetřil and V. Rödl. The Ramsey property for graphs with forbidden complete subgraphs. *Journal of Combinatorial Theory, Series B*, **20**(3):243–249, 1976.
- [26] J. Nešetřil and V. Rödl. A short proof of the existence of highly chromatic hypergraphs without short cycles. *Journal of Combinatorial Theory, Series B*, **27**:225–227, 1979.

- [27] V. Rödl. On the chromatic number of subgraphs of a given graph. *Proceedings of the American Mathematical Society*, **64**:370–371, 1977.
- [28] V. Rödl, N. Sauer and X. Zhu. Ramsey families which exclude a graph. *Combinatorica*, **15**(4):589–596, 1995.
- [29] V. Rödl and P. Winkler. A Ramsey-type theorem for orderings of a graph. *SIAM Journal on Discrete Mathematics*, **2**(3):402–406, 1989.
- [30] V. Rödl and N. Sauer. The Ramsey property for families of graphs which exclude a given graph. *Canadian Journal of Mathematics*, **44**(5):1050–1060, 1992.
- [31] N. Sauer. On the Ramsey property of families of graphs. *Transactions of the American Mathematical Society*, **374**(3):785–833, 1995.
- [32] N. W. Sauer and X. Zhu. Graphs which do not embed a given graph and the Ramsey property. In *Sets, graphs and numbers (Budapest, 1991)*, volume 60 of *Colloquia Mathematica Societatis János Bolyai*, pages 631–636. North-Holland, Amsterdam, 1992.
- [33] A. Scott. Graphs of large chromatic number. In: *Proceedings of the ICM 2022*, to appear.
- [34] A. Scott and P. Seymour. Induced subgraphs of graphs with large chromatic number. I. Odd holes. *Journal of Combinatorial Theory, Series B*, **121**:68–84, 2016.
- [35] A. Scott and P. Seymour. Induced subgraphs of graphs with large chromatic number. VI. Banana trees. *Journal of Combinatorial Theory, Series B*, **145**:487–510, 2020.
- [36] A. Scott and P. Seymour. A survey of χ -boundedness. *Journal of Graph Theory*, **95**(3):473–504, 2020.
- [37] S. Sidon. Ein Satz über trigonometrische Polynome und seine Anwendung in der Theorie der Fourier-Reihen. *Mathematische Annalen*, **106**:536–539, 1932.
- [38] D. P. Sumner. Subtrees of a graph and the chromatic number. In *The theory and applications of graphs (Kalamazoo, Michigan, 1980)*, pages 557–576. Wiley, New York, 1981.
- [39] E. Szemerédi. On sets of integers containing no k elements in arithmetic progression. *Acta Arithmetica*, **27**:199–245, 1975.
- [40] B. L. van der Waerden. Beweis einer Baudetschen Vermutung. *Nieuw Archief voor Wiskunde*, **15**:212–216, 1927.
- [41] A. Zykov. On some properties of linear complexes (in Russian). *Matematicheskii Sbornik*, **24**:163–188, 1949. For an English translation see *American Mathematical Society Translations*, **79**, 1952.