Induced subgraphs of induced subgraphs of large chromatic number

António Girão*† Freddie Illingworth*† Emil Powierski*
Michael Savery*‡ Alex Scott*† Youri Tamitegama* Jane Tan*

8 March 2022

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. Moreover, we show an analogous statement where clique number is replaced by odd girth.

1 Introduction

It is a fundamental problem of graph theory to understand what structures must appear in graphs of large chromatic number. A straightforward reason for a graph to have large chromatic number is the presence of a large clique. However, this is not necessary as there are examples of triangle-free graphs with large chromatic number (see, for instance, Tutte [5] and Zykov [19] among many classical constructions).

If a graph has large chromatic number and small clique number, it is reasonable to ask whether it contains induced subgraphs with large chromatic number which are 'simple' in some way. A natural conjecture, attributed to Esperet (see [17]), asked whether all graphs with large chromatic number contain either a large clique or a triangle-free induced subgraph with large chromatic number (the non-induced version of this was proven by a beautiful argument of Rödl [12]). This conjecture was recently disproved by Carbonero, Hompe, Moore and Spirkl [4], who found a surprising new twist on a construction of Kierstead and Trotter [9], and proved the following.

*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

Theorem 1 (Carbonero, Hompe, Moore and Spirkl [4]). *There are graphs of arbitrarily large chromatic number that contain neither a K_4 nor an induced triangle-free subgraph of chromatic number greater than four.*

Let us write $\omega(H)$ for the *clique number* (the maximum number of vertices in a complete subgraph) of a graph H and $\chi(H)$ for the chromatic number of H. We say that a graph G is *F-free* if it does not contain an induced copy of F.

Briański, Davies and Walczak [3] extended the result of Carbonero, Hompe, Moore and Spirkl to cliques of prime order, in an ingenious paper proving the following.

Theorem 2 (Briański, Davies and Walczak [3]). *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 $\chi(H) \leqslant \omega(H)^{\omega(H)^2}$.*

As a consequence, they showed that there are classes of graphs that are χ-bounded but not polynomially χ-bounded, disproving another conjecture of Esperet [7] (see [16] and [17] for definitions and related discussion).

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. It is natural to ask whether anything is true for other graphs F. In other words, is there a graph G with small clique number and large chromatic number such that every induced subgraph of G with large chromatic number contains an induced copy of F? Perhaps surprisingly, our main theorem answers this in the affirmative for every nontrivial graph F. Indeed, G can be taken to have the same clique number as 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.*

It is interesting to consider our result in the context of Ramsey classes. A class $\mathcal{C}$ of graphs is a *Ramsey class* if for every $F \in \mathcal{C}$ and integer k there is some $G \in \mathcal{C}$ such that there is no partition of the vertices of G into k sets each inducing F-free subgraphs. Folkman [8] proved that the class $\mathcal{C}_r$ of K_r-free graphs is Ramsey for every $r \geqslant 3$. Theorem 3 proves a vastly stronger property: it shows that for every $F \in \mathcal{C}_r$ there are graphs in $\mathcal{C}_r$ with arbitrarily large chromatic number in which *no* induced subgraph of large chromatic number is F-free. We discuss this further in Section 6.

Theorem 3 gives a graph G with the same clique number as F. In the case when F is triangle-free, it is natural to ask whether we can take our graphs G to have the same girth as F. We conjecture that this is the case.

Conjecture 4. *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.*

It would already be interesting to prove this in the special case where F is the 5-cycle. We prove a weaker version of the conjecture, where girth is replaced by *odd girth*, the length of the shortest odd cycle.

Theorem 5. *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.*

By taking F to be a disjoint union of a finite collection of graphs, we get the following immediate corollary.

Corollary 6. *Let $\omega \geqslant 2$, and suppose that $\mathcal{F}$ is a finite collection of graphs all with clique number at most ω. There exists a constant $c_{\mathcal{F}}$ and graphs G of arbitrarily large chromatic number and clique number at most ω such that every induced subgraph of G with chromatic number at least $c_{\mathcal{F}}$ contains every member of $\mathcal{F}$ as an induced subgraph.*

One may wonder whether this can be extended to infinite collections of graphs (where, of course, we now allow G to be an infinite graph). In Section 5 we prove this is not possible.

The proofs of Theorems 3 and 5 follow a common framework, building on the arguments of Carbonero, Hompe, Moore and Spirkl [4], and Briański, Davies and Walczak [3]. Broadly speaking, all of these constructions first define an oriented graph with certain properties (such as large chromatic number) to act as a base, and then add edges between any two vertices that are joined by a directed path with length in some carefully chosen set. Then G is taken as the underlying undirected graph of this digraph. The two moving parts here are therefore the choice of base graph, and choice of allowable length set.

In the proof of Theorem 3, it is convenient to begin with the oriented Zykov graph (which was also used in [3] and [4]). However, for Theorem 5, we introduce an oriented version of graphs of Nešetřil and Rödl [11]. A key ingredient of the proof, which may be of independent interest, is to show that every cycle in these graphs has a large number of direction changes. Another new ingredient in both proofs is the use of B_h-sets to control the appearance of the graph F.

The rest of the paper is organised as follows. In Section 2, we present the two families of base graphs that we will use, together with their key properties. In Section 3.1 we discuss B_3-sets, which are used to define the allowable lengths, and then, in Section 3.2, we present the full construction establishing explicit bounds for the constant c_F. The proof of Theorem 5 is handled in Section 4, with a parallel structure: we discuss B_h-sets in Section 4.1 and then give the construction in Section 4.2. In Section 5 we discuss extensions to infinite graphs, and finally in Section 6 we consider our results in the context of Ramsey classes and indicate possible avenues for further investigation.

2 Orientations of classical constructions

In this section, we define two families of directed graphs, alongside their underlying undirected graphs, which provide the starting point of our constructions. They are based on classical constructions of graphs with large chromatic number and restricted clique number or girth. The first, denoted (Z_n), are the well known Zykov graphs [19] (see [20] for an English translation). We will make use of orientations $(\overrightarrow{Z_n})$ of the Zykov graphs which were defined by Kierstead and Trotter [9], and used in the constructions of Carbonero, Hompe, Moore and Spirkl [4] and Briański, Davies and Walczak [3]. A second family of graphs, which we denote (Y_n), comes from a construction due to Nešetřil and Rödl [11]. We orient these to obtain a sequence of digraphs $(\overrightarrow{Y_n})$ that we use in the proof of Theorem 5.

The graphs (X_n), where X is Y or Z, and their orientations $(\overrightarrow{X_n})$ are chosen to satisfy the following properties:

(a) $\chi(X_n) = n$,

(b) $\overrightarrow{X_n}$ is acyclic, that is, it contains no directed cycle, and

(c) for $u, v \in V(\overrightarrow{X_n})$, there is at most one directed path from u to v in $\overrightarrow{X_n}$.

There is one further condition that we need the Y_n to satisfy so we can rule out short odd cycles. We define a *change of direction* in a digraph to be a pair of edges that share a vertex but do not form a directed path. For a given positive integer g we ask that $\overrightarrow{Y_n}$ satisfies:

(d) for every cycle C in Y_n, the corresponding oriented cycle $\overrightarrow{C}$ in $\overrightarrow{Y_n}$ has at least g changes of direction.

Strictly speaking, $\overrightarrow{Y_n}$ and condition (d) depend on the choice of positive integer g, but we suppress this dependence in the notation. Note that if $g \geqslant 3$, then property (d) implies both (b) and (c).

By property (c), it is natural to view $V(\overrightarrow{X_n}) = V(X_n)$ 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 $\overrightarrow{X}_n$. For vertices $u < v$ we can then 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 $u < v < w$.

As the properties listed above are the only properties that we will require, one could in fact use any sequence of directed graphs $(\overrightarrow{X_n})$ and corresponding sequence (X_n) of underlying undirected graphs that satisfy the first three conditions to prove Theorem 3, or all four conditions to prove Theorem 5. In particular, the oriented Nešetřil-Rödl graphs in Section 2.2 suffice for both theorems. However, for clarity, we also briefly introduce the significantly simpler Zykov graphs below.

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.

2.1 Zykov graphs

The graphs $(\overrightarrow{Z_n})$ are defined inductively, with $\overrightarrow{Z_1}$ being the one-vertex graph. Given $\overrightarrow{Z_n}$, we construct $\overrightarrow{Z_{n+1}}$ by taking n disjoint copies of $\overrightarrow{Z_n}$, and for each tuple $(v_1, \ldots, v_n)$ with

v_i in the ith copy of $\overrightarrow{Z_n}$, we add a new vertex v together with edges $\overrightarrow{v_iv}$ for each i. Let Z_n be the underlying undirected graph associated to $\overrightarrow{Z_n}$.

It is straightforward to check using elementary inductive arguments that the graphs Z_n and $\overrightarrow{Z_n}$ have properties (a), (b) and (c).

2.2 The Nešetřil-Rödl construction

For $s \geqslant 2$ a hypergraph is *s-uniform* if every edge contains exactly s vertices. For our purposes, a *cycle of length* $\ell \geqslant 2$ in a hypergraph is a sequence of distinct vertices $v_0, \dots, v_{\ell-1}$ and a sequence of distinct edges $e_0, \dots, e_{\ell-1}$ such that $v_i \in e_i \cap e_{i+1}$ for all i, where the indices are taken modulo ℓ. Note that in particular we consider two distinct edges with at least two vertices in their intersection to give rise to a cycle of length 2. The *girth* of a hypergraph is the length of its shortest cycle. The *chromatic number* of a hypergraph is the minimum number of colours required to colour its vertices so that no edge is monochromatic.

The construction of Nešetřil and Rödl uses uniform hypergraphs with large girth and large chromatic number. Such hypergraphs were shown to exist by Erdős and Hajnal [6] via a probabilistic argument.

Theorem 7 (Erdős and Hajnal [6]). *For all $g, k, s \geqslant 2$, there is an s-uniform hypergraph which has chromatic number at least k and girth at least g.*

The digraph $\overrightarrow{Y}$ and graph Y constructed in the proof of the following lemma with parameters g and n will be the digraph $\overrightarrow{Y_n}$ and graph Y_n used in the proof of Theorem 5, where we recall that in the notation we suppress dependence on g, and that properties (b) and (c) follow from property (d) when $g \geqslant 3$. The construction in the following lemma is standard, as is the argument to determine the chromatic number of the graph Y_n, but we give both for completeness. The additional fact that the graphs Y_n can be oriented to produce directed graphs $\overrightarrow{Y_n}$ satisfying property (d) seems to be a new observation and thus the following lemma may be of independent interest.

Lemma 8. *For every $g \geqslant 3$ and $n \geqslant 2$, there is a graph Y with an orientation $\overrightarrow{Y}$ satisfying properties* (a) *and* (d)*, that is $\chi(Y) = n$ and for every cycle C in Y, the corresponding oriented cycle $\overrightarrow{C}$ in $\overrightarrow{Y}$ contains at least g changes of direction.*

Proof. To begin, let B_0 be an n-partite graph with partition $A_0^1 \cup \cdots \cup A_0^n$ such that

- $|A_0^i| = n - 1$ for each i,
- the edges of B_0 form a matching, and
- for any two distinct parts A_0^i and A_0^j, there is exactly one edge between them.

To obtain $\overrightarrow{B_0}$, we direct edges from A_0^i to A_0^j for all $i < j$. It is clear that $\overrightarrow{B_0}$ has property (d), as B_0 is acyclic.

Proceeding inductively, for $1 \leqslant i \leqslant n$, given a digraph $\overrightarrow{B_{i-1}}$ and underlying n-partite graph B_{i-1} with partition $A_{i-1}^1 \cup \cdots \cup A_{i-1}^n$, we construct $\overrightarrow{B_i}$ such that B_i is an n-partite graph with partition $A_i^1 \cup \cdots \cup A_i^n$. Let $\mathcal{H}_i$ be an $|A_{i-1}^i|$-uniform hypergraph with girth at least g and chromatic number at least n, which exists by Theorem 7. We build $\overrightarrow{B_i}$ by taking $|E(\mathcal{H}_i)|$ copies of $\overrightarrow{B_{i-1}}$ and identifying vertices in the ith class of each copy with the vertices in some edge of $\mathcal{H}_i$. More specifically, for every edge e of $\mathcal{H}_i$, we include a copy of $\overrightarrow{B_{i-1}}$, denoted by $\overrightarrow{B_{i-1}}(e)$, and arbitrarily identify the part of this copy corresponding to A_{i-1}^i with e. Thus, for distinct edges e_1 and e_2 of $\mathcal{H}_i$, $\overrightarrow{B_{i-1}}(e_1)$ and $\overrightarrow{B_{i-1}}(e_2)$ are vertex disjoint if e_1 and e_2 are disjoint, and share exactly one vertex if e_1 and e_2 intersect. The resulting digraph is $\overrightarrow{B_i}$.

For each edge e of $\mathcal{H}_i$, and for each $1 \leqslant j \leqslant n$, let $A_{i-1}^j(e)$ denote the jth part, A_{i-1}^j, of $\overrightarrow{B_{i-1}}(e)$. We view B_i as an n-partite graph with parts $A_i^1, \ldots, A_i^n$, where A_i^j is the union of $A_{i-1}^j(e)$ over all edges e of $\mathcal{H}_i$. Thus $A_i^i = V(\mathcal{H}_i)$ and for $j \neq i$ the union is disjoint, so part A_i^j has size $|E(\mathcal{H}_i)| \cdot |A_{i-1}^j|$. Note that for all $0 \leqslant i \leqslant n$ and $1 \leqslant j < j' \leqslant n$, every edge in $\overrightarrow{B_i}$ between A_i^j and $A_i^{j'}$ is directed towards $A_i^{j'}$. Hence, for every path in B_i between two distinct vertices of A_i^j, the corresponding oriented path in $\overrightarrow{B_i}$ contains a change of direction.

Let $\overrightarrow{Y} = \overrightarrow{B_n}$ and let Y be the underlying undirected graph. We first show that $\chi(Y) = n$. We have seen that Y is n-partite so it remains to show that there is no proper $(n-1)$-colouring of Y. Suppose for a contradiction that c is such a colouring. Since the chromatic number of $\mathcal{H}_i$ is at least n for each $1 \leqslant i \leqslant n$, in any colouring of a copy of $\mathcal{H}_i$ induced by c there is a monochromatic edge of $\mathcal{H}_i$. In $\mathcal{H}_n$ this monochromatic edge e_n is $A_{n-1}^n(e_n) \subseteq V(B_{n-1}(e_n))$. The copy of $\mathcal{H}_{n-1}$ on vertex set $A_{n-1}^{n-1}(e_n)$ has a monochromatic edge e_{n-1} which is the vertex set $A_{n-2}^{n-1}(e_{n-1}) \subseteq V(B_{n-2}(e_{n-1}))$. Continuing in this way, we obtain a chain of nested subgraphs

$$B_{n-1}(e_n) \supseteq B_{n-2}(e_{n-1}) \supseteq \cdots \supseteq B_0(e_1).$$

By construction, each part $A_0^1(e_1), \ldots, A_0^n(e_1)$ of $B_0(e_1)$ is monochromatic. By the pigeonhole principle, two parts must share a colour and thus c is not a proper vertex colouring of Y, which gives the required contradiction.

It remains to show by induction on i that for all $0 \leqslant i \leqslant n$, every cycle in $\overrightarrow{B_i}$ contains at least g changes of direction. The case $i = 0$ is noted above, so let $1 \leqslant i \leqslant n$ and let C be a cycle in B_i. If $V(C) \subseteq V(B_{i-1}(e))$ for some $e \in E(\mathcal{H}_i)$, then we are done by induction. Suppose therefore that C contains vertices from more than one copy $B_{i-1}(e)$ of B_{i-1}. Pick an edge of C and let $B_{i-1}(e_1)$ be the unique copy of B_{i-1} containing it. Now working along the edges of C in (either) cyclic order, let $B_{i-1}(e_j)$ be the unique copy of B_{i-1} containing the jth edge of the cycle. Consider the sequence $e_1, \ldots, e_{|C|}$ of not necessarily distinct edges of $\mathcal{H}_i$ produced. By assumption this sequence contains at least two distinct elements. It is straightforward to see that some subsequence of these edges forms a cycle in $\mathcal{H}_i$ with an appropriate sequence of vertices from C. By the properties of $\mathcal{H}_i$, it follows that there are at least g distinct edges in the sequence $e_1, \ldots, e_{|C|}$.

Consider a maximal path in C contained in $B_{i-1}(e_j)$ for some e_j. The path's (distinct) endpoints are in $A_{i-1}^i(e_j)$, so as noted above the corresponding oriented path in $\overrightarrow{B_{i-1}}(e_j)$ contains a change of direction. It follows that since C contains at least g such maximal paths, the corresponding oriented cycle in $\overrightarrow{B_i}$ contains at least g changes of direction, as required. □

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 gives

$$\begin{aligned} c_F &\leqslant (96f^4)^{|E(F)|+1}, \\ c_F &\leqslant \big((6+o(1))f^4\big)^{|E(F)|+1}. \end{aligned} \tag{1}$$

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 [3, 4, 9] start with the digraphs $(\overrightarrow{Z_n})$, and add new edges $\overrightarrow{uv}$ (or sometimes $\overrightarrow{vu}$ in [4]) between vertices $u < v$ whenever p does not divide $d(u, v)$ for some prime p. These papers consider the cases of arbitrary primes p, $p = 3$ and $p = 2$ respectively. Our construction is similar in that, for some prime p, we add edges to $\overrightarrow{Z_n}$ 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 $\overrightarrow{G}$ be the digraph so obtained, and G its underlying undirected graph. As Z_n is a subgraph of G, we know that $\chi(G) \geqslant n$. The extra edges will be added so that the following two properties hold:

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

This second property is particularly useful in light of the following folklore result, which is a standard tool for bounding chromatic number. We say that a digraph is *acyclic* if it contains no directed cycles.

Proposition 9. *Let $\overrightarrow{L}$ be an acyclic digraph with no directed path of length ℓ. If L is the underlying undirected graph, then $\chi(L) \leqslant \ell$.*

To see this, simply colour each vertex v in L by the length of the longest directed path in $\overrightarrow{L}$ starting at v.

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 of length i modulo p. By the second bullet point, H_i cannot contain a long directed path. Thus, Proposition 9 says that H_i has chromatic number bounded in terms of F. Taking a product colouring 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 $S = \{a_1 < a_2 < \cdots < a_k\}$ is a *B_3-set* if the sums

$$a_{i_1} + a_{i_2} + a_{i_3}, \qquad 1 \leqslant i_1 \leqslant i_2 \leqslant i_3 \leqslant k$$

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 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 [2], as a generalisation of Sidon sets [18], which are precisely the B_2-sets. In particular, every B_3-set is also a B_2-set. Let S be a B_3-set and let

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

be the *difference set of S*. The sets S and D satisfy an important property which motivated 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.

- 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 $x_1 = y_2$, then $b_1 = x_2$, $b_2 = x_1 = y_2$ and $b_3 = y_1$ satisfy $(b_2 - b_1, b_3 - b_2) = (d_2, d_1)$.

Hence, we may assume that $y_1 \neq x_2$ and $x_1 \neq y_2$. If $y_1 = x_1$, then $d_1 = 0$, which is impossible as $0 \notin D$. Thus $y_1 = y_3$ and we can similarly deduce that $x_1 = x_3$. But then $d_1 + d_2 = y_3 - x_3 = d_1$ and so $d_2 = 0$, which is a contradiction. □

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 $d_1, \ldots, d_\ell$ such that every sum of the form*

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

with $1 \leqslant i_1 \leqslant i_2 \leqslant \ell$ is in D. Then there are $b_1 < \cdots < b_{\ell+1}$ 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 $\ell \geqslant 3$. By the induction hypothesis, there are $b_1 < \cdots < b_\ell$ 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 gives that there is $b \in S$ such that either $b_{\ell-1} - b = d_\ell$ or $b - b_\ell = d_\ell$. In particular, d_ℓ is associated with either $b_{\ell-1}$ or b_ℓ.

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 d_ℓ is associated with either b_1 or b_ℓ. As $\ell \geqslant 3$, the integers $b_1, b_{\ell-1}$, and b_ℓ are all distinct, so d_ℓ is associated with b_ℓ and there is $b \in S$ with $b - b_\ell = d_\ell$ as desired. □

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

Theorem 10 (Bose, Chowla)**.** *For any prime power q, there is a B_3-subset of $[q^3 - 1]$ of size q.*

Given a natural number f, using the well known fact that gap between consecutive primes $q < q'$ is $o(q)$ (see, for instance, [1] for an up-to-date bound), Theorem 10 implies that $[N]$ contains a B_3-set of size at least f for $N = (1 + o(1))f^3$. Bertrand's postulate gives a non-asymptotic bound: the smallest prime q greater than f is less than $2f$ and so Theorem 10 gives a B_3-subset of $[q^3] \subseteq [8f^3]$ of size at least f.

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 $\big[\lfloor p/6 \rfloor\big]$ contains a B_3-set of size f. By the discussion at the end of Section 3.1, it suffices to have $\lfloor p/6 \rfloor \geqslant \min\big\{(1+o(1))f^3, 8f^3\big\}$. Using the result on prime gaps and Bertrand's postulate we may choose such a p with $p \leqslant \min\big\{(6+o(1))f^3, 96f^3\big\}$.

Pick a B_3-set $S' \subseteq \big[\lfloor p/6 \rfloor\big]$ of size f and set $S = 3S' = \{3s' \colon s' \in S'\} \subseteq [(p-1)/2]$. Let the elements of S be $s_1, \ldots, s_f$ in ascending order and let $D = \{s_j - s_i \colon i < j\}$ be the difference set of S. Note that S is itself a B_3-set and so S and D satisfy Claims 1 and 2.

Taking a copy F^* of F with vertex set S, we write

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

to denote the set of distances corresponding to edges of F. Note that $E \subseteq D$. Let $\overline{D} = D + p\mathbb{Z}$ and $\overline{E} = E + p\mathbb{Z}$.

Recall the (di)graphs $\overrightarrow{Z_n}$ and Z_n from Section 2.1, 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 $\overrightarrow{Z_n}$. We now define the required graph G by adding edges to $\overrightarrow{Z_n}$ as follows.

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

This produces a digraph $\overrightarrow{G}$ which is acyclic, since the presence of the edge $\overrightarrow{uv}$ implies that $u < v$ in $\overrightarrow{Z_n}$. Let G be the underlying undirected graph of $\overrightarrow{G}$. It is useful to view each edge of $\overrightarrow{G}$ as having one of two types: those inherited from $\overrightarrow{Z_n}$ (corresponding to distance 1), and those which correspond to distance i for some $i \in \overline{E}$.

Before proving that G has the properties stated in Theorem 3, we record some facts about D and E. By definition, every element in S is divisible by 3. Hence, no two elements in D are consecutive, and D does not contain 0, 1 or 2. Since further $S \subseteq [(p-1)/2]$, we have $E \subseteq D \subseteq [(p-1)/2]$. Therefore, $\overline{D}$ and $\overline{E}$ do not contain 0, 1 or 2, nor any two consecutive numbers. In particular, $\overline{E} \cup \{1\}$ does not contain two consecutive numbers.

We now verify that G satisfies the conditions of Theorem 3. Since G contains Z_n as a subgraph, it has chromatic number at least n. Next we show that G has no clique on $\omega(F)+1$ vertices. Our main tool will be the following, which follows almost immediately from the clique fact.

Claim 3. *Let ℓ be a natural number and let $d_1, \ldots, d_\ell$ be not necessarily distinct elements of D. Suppose that every sum of the form in (2) is in $\overline{D}$. Then there are $b_1 < \cdots < b_{\ell+1}$ 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. Let d and d' be not necessarily distinct elements of D such that $d + d' \in \overline{D}$. Since $D \subseteq [(p-1)/2]$ we have $0 \leqslant d+d' \leqslant p-1$. But then $d+d' \in D$, so $0 \leqslant d+d' \leqslant (p-1)/2$.

Letting ℓ and $d_1, \ldots, d_\ell$ be as in the statement of the claim, we can repeatedly apply the last inequalities to show that the sums

$$\sum_{i_1 \leqslant j \leqslant i_2} d_j, \qquad 1 \leqslant i_1 \leqslant i_2 \leqslant \ell$$

are all between 0 and $p-1$, and thus are all in D. The claim now follows from the clique fact. □

Suppose for a contradiction that G contains a clique on $\omega(F)+1$ vertices. Fix such a clique and label the vertices as $v_1, \ldots, v_{\omega(F)+1}$, where $v_1 < \cdots < v_{\omega(F)+1}$ under the ordering of the vertices of $\overrightarrow{Z_n}$. Then $d(v_i, v_j) \in \overline{E} \cup \{1\}$ for all $1 \leqslant i < j \leqslant \omega(F)+1$. Let $d_i \equiv d(v_i, v_{i+1}) \bmod p$ be chosen so that $0 \leqslant d_i \leqslant p-1$. Then $d_i \in E \cup \{1\}$. The presence of the clique in G implies that every sum of the form in (2) is in $\overline{E} \cup \{1\}$. We will show that in fact none of these sums are equal to 1.

Indeed, suppose that $d_i = 1$ for some i. If $i \neq \omega(F)$, then d_{i+1} and $d_{i+1}+d_i$ are consecutive numbers in $\overline{E} \cup \{1\}$, while if $i = \omega(F)$, then d_{i-1} and $d_{i-1} + d_i$ are consecutive. Both of these are impossible. Thus $d_i \neq 1$ for all i. Sums of more than one d_i are greater than 1. Hence all sums of the form in (2) are in $\overline{E}$, and hence in $\overline{D}$, while the d_i are all in E and thus in D. It now follows from Claim 3 that there are $b_1 < \cdots < b_{\omega(F)+1}$ all in S such that either $b_{i+1} - b_i = d_i$ for all i, or $b_{i+1} - b_i = d_{\omega(F)+1-i}$ for all i. All the pairwise distances between elements of $B = \{b_1, \ldots, b_{\omega(F)+1}\}$ are of the form in (2) and so are in $\overline{E}$. But since $B \subseteq S \subseteq [(p-1)/2]$, they must actually all be in E. By the definition of E, these vertices induce a clique in F^* of size greater than $\omega(F)$, which is a contradiction. Therefore $\omega(G) \leqslant \omega(F)$.

We now turn to the last required property of G, that all of its induced subgraphs of chromatic number greater than 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 $i \in E \cup \{1\}$. Label the first vertex of this path by v_1. Recall that the elements of S are $s_1, \dots, s_f$ in ascending order. Let v_2 be the first vertex on this path to which the path in $\vec{Z_n}$ from v_1 has length congruent to $s_2 - s_1$ modulo 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 $d(v_2, v_3) \equiv s_3 - s_2 \bmod p$, 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 to ensure 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 $d(v_i, v_j) \equiv s_j - s_i \bmod p$. Then, by the construction of G, 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 $i \in E \cup \{1\}$, 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 $\chi_i \colon V(H_i) \to [pf]$ for each i. Then the product colouring $\chi \colon V(H) \to [(pf)]^{|E|+1}$ given by $\chi(v) = (\chi_i(v) \colon i \in E \cup \{1\})$ is a proper colouring of H. Hence, we may take c_F to be $(pf)^{|E|+1}$. Finally, we note that $|E| = |E(F)|$ and recall that $p \leqslant \min\big\{(6 + o(1))f^3, 96f^3\big\}$ to obtain

$$c_F \leqslant \min\Big\{\big((6+o(1))f^4\big)^{|E(F)|+1}, (96f^4)^{|E(F)|+1}\Big\}.$$

4 Odd girth

The proof of Theorem 5 shares some ingredients with that of Theorem 3. We will pay particular attention to the differences.

4.1 B_h-sets

Given a natural number $h \geqslant 4$, we can define the notion of B_h-sets analogously to B_3-sets. That is, a set of integers $S = \{a_1 < a_2 < \cdots < a_k\}$ is a *B_h-set* if the sums

$$a_{i_1} + a_{i_2} + \cdots + a_{i_h}, \qquad 1 \leqslant i_1 \leqslant i_2 \leqslant \cdots \leqslant i_h \leqslant k$$

are all different. The powers of h form such a set (see [2] for a more efficient construction). Note that a B_h-set is also a $B_{h'}$-set for any $h' < h$, and in particular is a Sidon set.

Let S be a B_h-set and let D be its difference set as before. The importance of B_h-sets for our proof of Theorem 5 is the following claim. Its statement is convenient for induction purposes and, in fact, we will only make use of the 'moreover' part. By a circuit, we mean a closed walk in which vertices may be repeated but edges may not.

Claim 4 (Cycle fact)**.** *Let $2 \leqslant \ell \leqslant h$ be integers. Suppose that D contains not necessarily distinct $d_1, \ldots, d_\ell$ such that for some $1 \leqslant s \leqslant \ell - 1$ we have*

$$d_1 + \cdots + d_s = d_{s+1} + \cdots + d_\ell. \tag{3}$$

Let M be the multigraph on vertex set S in which we add an edge e_i between the pair of vertices at distance d_i for each $1 \leqslant i \leqslant \ell$. 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 $\ell \geqslant 3$. Let $x_1, \ldots, x_\ell, y_1, \ldots, y_\ell \in S$ be such that $d_i = y_i - x_i$ for all $1 \leqslant i \leqslant \ell$. From equation (3) we obtain

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

and so, as S is a B_h-set and $\ell \leqslant h$, $y_1, \ldots, y_s, x_{s+1}, \ldots, x_\ell$ must be $x_1, \ldots, x_s, y_{s+1}, \ldots, y_\ell$ in some order. In particular, $y_1 = x_i$ for some $1 \leqslant i \leqslant s$ or $y_1 = y_j$ for some $s+1 \leqslant j \leqslant \ell$. In the first case we may assume by relabelling that $y_1 = x_2$ ($y_1 \neq x_1$ as $0 \notin D$). Then $d \coloneqq d_1 + d_2 = y_2 - x_1$ which is in D as $y_2 > x_2 = y_1 > x_1$. Moreover, $d + d_3 + \cdots + d_s = d_{s+1} + \cdots + d_\ell$. Thus, by induction, the edges of the multigraph obtained from M by deleting e_1 and e_2 and adding an edge e between y_2 and x_1 can be decomposed into a collection of circuits. Adding back the edges e_1 and e_2, which form a path from y_2 to x_1, 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 $d_2 + \cdots + d_s = d_{s+2} + \cdots + d_\ell$. 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 $d \coloneqq d_1 - d_{s+1} = x_{s+1} - x_1 \in D$ and $d + d_2 + \cdots + d_s = d_{s+2} + \cdots + d_\ell$. By induction, there is a decomposition of the edges of the multigraph obtained from M by deleting e_{s+1} and e_1 and adding a new edge e incident with x_{s+1} and x_1. Adding back the edges e_{s+1} and e_1, which form a path between x_{s+1} and x_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. □

4.2 Proof of Theorem 5

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 $\left[\lfloor p/h^2 \rfloor\right]$ contains a B_h-set of size f. Let S' be such a set and let $S = hS' = \{hs' \colon s' \in S'\} \subseteq \left[\lfloor p/h \rfloor\right]$. Let the elements of S be $s_1, \ldots, s_f$ in ascending order and let $D = \{s_j - s_i \colon i < j\}$ be the difference set of S. Note that S is itself a B_h-set and so S and D satisfy the cycle fact.

Taking a copy F^* of F with vertex set S, we write

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

to denote the set of distances corresponding to edges of F. Note that $E \subseteq D$. By definition, every element of S is divisible by h. Hence, all elements of D and E are divisible by h. Let $\overline{D} = D + p\mathbb{Z}$ and $\overline{E} = E + p\mathbb{Z}$.

We now define the required graph G as follows. Start with the digraph $\overrightarrow{Y_n}$ defined in Section 2.2 where every cycle has at least h changes of direction. We add the following edges.

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

As before, the digraph $\overrightarrow{G}$ produced is acyclic. We now show that G satisfies the properties stated in Theorem 5. Firstly, the graph Y_n is a subgraph of G so $\chi(G) \geqslant n$. 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 $c'_F = (pf)^{|E|+1}$.

It remains to check that the odd girth of G is at least h. Suppose for a contradiction that G contains an odd cycle $v_0 v_1 \ldots v_{\ell-1}$ of length $\ell < h$. Since $\overrightarrow{G}$ contains no directed cycles, without loss of generality the path $v_{\ell-1} v_0 v_1$ in G is not a directed path in $\overrightarrow{G}$. We may further assume that since ℓ is odd, the path $v_0 v_1 v_2$ in G is a directed path in $\overrightarrow{G}$ from v_0 to v_2. For each edge $v_i v_{i+1}$ in this cycle (where here and throughout we take subscript addition in the v_i's to be modulo ℓ), there is a directed path $\overrightarrow{P_i}$ from v_i to v_{i+1} or from v_{i+1} to v_i in $\overrightarrow{Y_n}$ depending on the direction of the edge between v_i and v_{i+1} in $\overrightarrow{G}$. For each i let P_i be the undirected path underlying $\overrightarrow{P_i}$ and concatenate these paths P_i in the natural way to obtain a walk W in Y_n which begins at v_0, visits $v_1, \ldots, v_{\ell-1}$ 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 $\overrightarrow{Y_n}$, 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 $S_1 = \left\{0 \leqslant i \leqslant \ell - 1 \colon \overrightarrow{v_i v_{i+1}} \in E(\overrightarrow{G})\right\}$ and $S_2 = \left\{0 \leqslant i \leqslant \ell - 1 \colon \overrightarrow{v_{i+1} v_i} \in E(\overrightarrow{G})\right\}$ so that S_1 and S_2 form a partition of $\{0, 1, \ldots, \ell - 1\}$. For $0 \leqslant i \leqslant \ell - 1$, 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 $v_i v_{i+1}$ are all present in G, we have $d_i \in \overline{E} \cup \{1\}$ for all i. Let $\overline{d_i} \equiv d_i \bmod p$ be chosen so that $0 \leqslant \overline{d_i} \leqslant p - 1$ for all i. As $E \subseteq \left[\lfloor p/h \rfloor\right]$, we have $\overline{d_i} \in E \cup \{1\}$ for all i. Now 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 have fewer than h terms, all of which are at most p/h. It follows that, in fact,

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

We claim that 1 occurs the same number of times in these two sums. If this is not the case, then since there are fewer than h terms in total and every element of E is divisible by h, the two sums must have different values modulo h, which is a contradiction. Now since ℓ is odd and $0 \notin \overline{E}$, in total there are at least two terms which are not 1, at least one on either side of the equation. We now subtract all the 1's from both sides and relabel to get some $d'_1, d'_2, \ldots, d'_{\ell'} \in E$ such that

$$\sum_{i=1}^{t} d'_i = \sum_{i=t+1}^{\ell'} d'_i,$$

where $\ell' \leqslant \ell$ is odd, and $1 \leqslant t \leqslant \ell' - 1$.

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 $d'_1, \ldots, d'_{\ell'}$ contains an odd cycle of length at most ℓ'. Since $d'_1, \ldots, d'_{\ell'} \in E$, by the definition of E the copy F^* of F on vertex set S 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 $c_0, c_1, \ldots, c_{r-1}$ in cyclic order around Γ, where c_0 is the vertex amongst these that W arrives at first. Let $C = \{c_0, \ldots, c_{r-1}\}$. We will show that Γ contains fewer than h changes of direction, which contradicts the properties of $\overrightarrow{Y_n}$. 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 $0 \leqslant i \leqslant r-1$ 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 $\{c_ic_{i+1}, c_{i+1}c_{i+2}, \ldots, c_{j-1}c_j\}$ and $\{c_ic_{i-1}, c_{i-1}c_{i-2}, \ldots, c_{j+1}c_j\}$. 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 $0 \leqslant i \leqslant r-1$, 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_ic_{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_ic_{i+1}. Similarly to above, at the first time after the formation of P' at which either c_ic_{i+1} has been traversed or all of $c_ic_{i-1}, \ldots, c_{i+2}c_{i+1}$ 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_ic_{i-1}. These facts combine to give the following behaviour of W': it starts at c_0, carries out a (possibly empty) walk from c_0 to itself, avoiding all vertices of $C \setminus \{c_0\}$, then traverses an edge of Γ to c_1 or c_{r-1} and proceeds in this manner. More formally, there are vertices $\gamma_0, \ldots, \gamma_{s-1}$ 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 sequence of vertices $\gamma_0, \gamma_1, \ldots, \gamma_{s-1}$ forms 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_ic_{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 $\overrightarrow{Y_n}$ with fewer than h changes of direction, which gives the required contradiction. Therefore, the odd girth of G is at least h.

5 Extensions to infinite graphs

By taking a countable union of graphs, Theorem 3 has the following corollary.

Corollary 11. *For every graph F with at least one edge, there is a constant c_F and a graph G of infinite chromatic number such that $\omega(G) = \omega(F)$ and every F-free induced subgraph of G has chromatic number at most c_F.*

As in Corollary 6, it is straightforward to replace F with a finite family of graphs; but what about an infinite family? For example, we might ask whether there exists a (countably infinite) graph G with infinite chromatic number such that every induced subgraph of G with infinite chromatic number contains an induced copy of every (finite) graph. It follows from Lemma 12 that such a G cannot exist. Recall that the *union* of a family of graphs $\mathcal{F}$ is the (possibly infinite) graph consisting of pairwise vertex-disjoint copies of the graphs in $\mathcal{F}$ with no edges between copies.

Lemma 12. *If an infinite graph G has unbounded clique number and no infinite clique, then it contains the union of all finite cliques as an induced subgraph.*

Proof. Let G have unbounded clique number but no infinite clique. Let $(a_i)_{i\geqslant 0}$ 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_0} in G on vertex set V_0. By identifying a copy of $K_{a_0+a_1}$ in G, we can find a copy of K_{a_1} on vertex set V_1 disjoint from V_0. Continuing in this manner, we obtain a sequence $(V_i)_{i\geqslant 0}$ of pairwise disjoint sets of vertices of G such that V_i induces a copy of K_{a_i} for each i.

Fix a vertex $v_0 \in V_0$, then for each $i \geqslant 1$ remove at most half of the vertices from V_i so that v_0 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_0, 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 $|V_i'| \geqslant |V_i|/2^{a_0}$ and have the property that for each $i \geqslant 1$ and $v \in V_0$, v is either adjacent to every vertex in V_i' or not adjacent to any of them.

For each $i \geqslant 1$ we can then define a vector $b_i \in \{0,1\}^{a_0}$ by letting the jth entry be 1 if the jth vertex of V_0 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_0 corresponding to 1 entries, for some $t \geqslant a_0/2$ we obtain a copy of K_t in G which is adjacent to every vertex in $\bigcup_{i \geqslant 1} V_i'$. Otherwise we obtain a copy of K_t with no neighbours in this union for some $t \geqslant a_0/2$.

Repeat this process for V_1', then what remains of V_2', and so on. If the a_i are chosen appropriately, then it follows that there exists an induced subgraph of G consisting of a copy of K_n on vertex set W_n for each $n \in \mathbb{N}$, where the sets W_n are pairwise disjoint, such that for each n the edges between W_n and $\cup_{i>n} W_i$ are either all present or all not present. If these edges are present for infinitely many values of n, then G contains an infinite clique, which contradicts our assumptions. Hence there is an induced subgraph of G consisting of a union of arbitrarily large cliques, which in particular contains the union of all finite cliques as an induced subgraph. □

Despite this result, we might hope for a weaker statement to be true. For example, does there exist, for a natural number r, an infinite-chromatic graph G such that every infinite-chromatic induced subgraph of G contains an induced copy of every K_r-free graph? Theorem 13 answers this in the negative for all r and goes further by showing that, in fact, any such G has infinite-chromatic induced subgraphs which are F-free for graphs F of arbitrarily large girth. In what follows, for $g \geqslant 3$ we will denote by $\mathcal{G}_g$ the class of all (finite) graphs with girth at least g.

Theorem 13. *For all $g \geqslant 3$, every infinite graph with infinite chromatic number has an induced subgraph with infinite chromatic number which is F-free for some graph F with girth at least g.*

Proof. Let $g \geqslant 3$ and let G be an infinite-chromatic graph. We may assume that G contains an induced copy of every graph in $\mathcal{G}_g$, and we may further assume that G does not contain an infinite clique, as this would constitute an infinite-chromatic subgraph which does not contain, for example, P_3 as an induced subgraph. By Lemma 12, if G has unbounded clique number, then again it has an induced subgraph with no infinite-chromatic and P_3-free subgraph.

Hence we now assume that $\omega(G)$ is finite. Suppose for a contradiction that every infinite-chromatic induced subgraph of G contains every graph in $\mathcal{G}_g$ as an induced subgraph. 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. This contains every graph in $\mathcal{G}_g$ as an induced subgraph by assumption, but has clique number strictly less than that of G. Repeating this process finitely many times, we obtain an infinite-chromatic induced subgraph of G containing an induced copy of every graph in $\mathcal{G}_g$ in which the neighbourhood of every vertex is finitely colourable.

Let F_2 be a graph with chromatic number 2 and girth greater than g. By assumption, G contains an induced copy of F_2, say on vertex set V_2. Let c_2 be the (finite) chromatic number of $G[V_2 \cup \Gamma(V_2)]$, then let F_3 be a graph with chromatic number at least $c_2 + 3$ and girth greater than g. Fix an induced copy of F_3 in G and note that the part which remains after removing any vertices in V_2 or its neighbourhood has chromatic number at least 3 and girth greater than g. Let V_3 be the set of vertices of the copy of F_3 which remain.

Continue in this manner: at each step let c_i be the chromatic number of the graph induced on the union of $V_2, \dots, V_i$ and their neighbourhoods, then let F_{i+1} be a graph with chromatic number at least $c_i + i + 1$ and girth greater than g. Fix an induced copy of F_{i+1} in G and note that the graph induced on V_{i+1} – the set of vertices of the copy which remain after those in $V_2, \dots, V_i$ or their neighbourhoods are removed – has chromatic number at least $i + 1$ and girth greater than g. We obtain pairwise disjoint subsets $(V_i)_{i \geqslant 2}$ of $V(G)$ such that the subgraph of G induced on V_i has girth greater than g and chromatic number at least i, and such that there are no edges between V_i and V_j for $i \neq j$. This completes the proof of the theorem. □

By extending the argument used to prove Lemma 12 we can obtain even stronger results than Theorem 13, such as the following.

Theorem 14. *Let $g \geqslant 3$, and let G be an infinite graph which contains every graph in $\mathcal{G}_g$ as an induced subgraph. Then either G contains an infinite clique, or it contains the disjoint union of all graphs $\mathcal{G}_g$ as an induced subgraph.*

Note that the *independence number* of a hypergraph is the size of the largest set of vertices which does not contain an edge. An alternative formulation of Theorem 7 states that for all $g, s \geqslant 2$ and $\varepsilon > 0$, there exists an s-uniform hypergraph $\mathcal{H} = (V, \mathcal{E})$ which has girth at least g and independence number at most $\varepsilon|V|$.

Proof. Enumerate $\mathcal{G}_g$ as $(A'_i)_{i \geqslant 0}$, then for each $i \geqslant 0$ let A_i be the union of $A'_0, \dots, A'_i$. We now build a family of graphs $(H_i)_{i \geqslant 0}$ and positive constants $(\varepsilon_i)_{i \geqslant 0}$ as follows. For $i \geqslant 0$, given H_j and ε_j for all $j < i$, let ε_i be small in terms of $(|V(H_j)|)_{j<i}$, and let $\mathcal{H}_i = (V_i, \mathcal{E}_i)$ be an $|A_i|$-uniform hypergraph with girth at least g and independence number at most $\varepsilon_i|V_i|$. Form H_i on the vertex set V_i by arbitrarily placing a copy of A_i in each edge of $\mathcal{H}_i$. It is clear that the graphs H_i all have girth at least g.

By identifying an induced copy of H_0 in G, then an induced copy of the union of $|V(H_0)| + 1$ copies of H_1, then an induced copy of the union of $|V(H_0)| + |V(H_1)| + 1$ copies of H_2, and so on, we can obtain a sequence $(V_i)_{i \geqslant 0}$ of pairwise disjoint sets of vertices of G such that V_i induces a copy of H_i for each i. It now follows from an argument very similar to that used in Lemma 12, using the fact that any set of at least $\varepsilon_i|V_i|$ vertices of H_i contains an induced copy of A_i, that either G contains an infinite clique or G contains as an induced subgraph the union of the graphs in some subsequence of $(A_i)_{i \geqslant 0}$. This completes the proof of the theorem. □

6 Ramsey classes and open problems

As noted in the introduction, it is interesting to compare our results to the theory of Ramsey classes. Recall that a class $\mathcal{C}$ of graphs is a *Ramsey class* if for every $F \in \mathcal{C}$ and integer k there is some G such that there is no partition of the vertices of G into k sets each inducing F-free subgraphs. For graphs F and G define $\binom{G}{F}$ to be the hypergraph on vertex set $V(G)$ whose edges correspond to induced copies of F in G. Note that an independent set in $\binom{G}{F}$ is exactly the vertex set of some F-free induced subgraph of G. Hence, an equivalent formulation of a class $\mathcal{C}$ being Ramsey is that for every $F \in \mathcal{C}$ there are $G \in \mathcal{C}$ for which $\binom{G}{F}$ has arbitrarily large chromatic number.

The general question for Ramsey classes is to determine which hereditary classes $\mathcal{C}$ of graphs have this Ramsey property (see [15]). For example, there has been substantial progress in the case when $\mathcal{C}$ is determined by a single excluded subgraph, that is when $\mathcal{C}$ is the class of F-free graphs for some F (see [8, 10, 13, 14]).

Provided F has at least one edge, any proper vertex colouring of G will be a proper vertex colouring of $\binom{G}{F}$ and so $\chi\big(\binom{G}{F}\big) \leqslant \chi(G)$. Of course, the $G \in \mathcal{C}$ that witness the fact that $\chi\big(\binom{G}{F}\big)$ can be arbitrarily large may have much larger chromatic number than $\binom{G}{F}$. An interesting corollary of Theorem 3 is that for the class of K_r-free graphs one can in fact ensure G's chromatic number is at most a multiplicative factor (depending only on F) larger than the chromatic number of $\binom{G}{F}$. Indeed, consider the G given by Theorem 3: for any independent set I in $\binom{G}{F}$ we have $\chi(G[I]) \leqslant c_F$ and so $\chi(G) \leqslant c_F \cdot \chi\big(\binom{G}{F}\big)$.

Theorem 3 shows that the class of K_r-free graphs has a far stronger property than just being a Ramsey class, and it is natural to ask which other families have this property.

Question 15. 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?

Any class $\mathcal{C}$ having the property in Question 15 must be a Ramsey class. Theorem 5 states that for each integer $g \geqslant 2$ the class of graphs with odd girth at least g has this property while Conjecture 4 asserts that the class of graphs with girth at least g does too. As cycles are 2-connected, it follows from a theorem of Nešetřil and Rödl [10] that these two classes are both Ramsey classes. A good starting point in Question 15 would be to resolve the case when $\mathcal{C}$ is determined by a single excluded graph:

Question 16. For which graphs H does the class of H-free graphs have the property that for every H-free graph F there is a constant c_F and H-free graphs G of arbitrarily large chromatic number such that every F-free induced subgraph of G is c_F-colourable?

We also recall our conjecture from the introduction concerning girth.

Conjecture 4. *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.*

The graphs $\overrightarrow{Z_n}$ and Z_n described in Section 2 are unsuitable for tackling this problem, just as in the odd girth case. The graphs $\overrightarrow{Y_n}$ and Y_n are more promising, but unfortunately 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, as stated in Section 1, a resolution of this conjecture even in the special case where F is the 5-cycle would be of interest, as such an argument would likely overcome many of the difficulties involved in proving the full conjecture.

References

[1] 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.

[2] R. C. Bose and S. Chowla. Theorems in the additive theory of numbers. *Commentarii Mathematici Helvetici*, **37**:141–147, 1962.

[3] M. Briański, J. Davies and B. Walczak. Separating polynomial χ-boundedness from χ-boundedness. *arXiv:2201.08814* preprint, 2022.

[4] 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. *arXiv:2201.08204* preprint, 2022.

[5] B. Descartes. k-chromatic graphs without triangles. *American Mathematical Monthly*, **61**(5):352–353, 1954.

[6] P. Erdős and A. Hajnal. On chromatic number of graphs and set-systems. *Acta Mathematica Academiae Scientiarum Hungarica*, **17**:61–99, 1966.

[7] L. Esperet. Graph colorings, flows and perfect matchings. Habilitation thesis, Université Grenoble Alpes, 2017.

[8] J. Folkman. Graphs with monochromatic complete subgraphs in every edge coloring. *SIAM Journal on Applied Mathematics*, **18**(1):19–24, 1970.

[9] H. A. Kierstead and W. T. Trotter. Colorful induced subgraphs. *Discrete Mathematics*, **101**:165–169, 1992.

[10] J. Nešetřil and V. Rödl. Partitions of vertices. *Commentationes Mathematicae Universitatis Carolinae*, **17**(1):85–95, 1976.

[11] 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.

[12] V. Rödl. On the chromatic number of subgraphs of a given graph. *Proceedings of the American Mathematical Society*, **64**:370–371, 1977.

[13] V. Rödl, N. Sauer and X. Zhu. Ramsey families which exclude a graph. *Combinatorica*, **15**(4):589–596, 1995.

[14] 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.

[15] N. Sauer. On the Ramsey property of families of graphs. *Transactions of the Amertican Mathematical Society*, **374**(3):785–833, 1995.

[16] A. Scott. Graphs of large chromatic number. In: *Proceedings of the ICM 2022*, to appear.

[17] A. Scott and P. Seymour. A survey of χ-boundedness. *Journal of Graph Theory*, **95**(3):473–504, 2020.

[18] S. Sidon. Ein satz über trigonometrische polynome und seine anwendung in der theorie der fourier-reihen. *Mathematische Annalen*, **106**:536–539, 1932.

[19] A. Zykov. On some properties of linear complexes (in Russian). *Matematicheskiĭ Sbornik*, **24**:163–188, 1949.

[20] A. Zykov. On some properties of linear complexes (English translation). *American Mathematical Society*, **79**, 1952.