Complexity Framework for Forbidden Subgraphs II: Edge Subdivision and the "H"-Graphs

Vadim Lozin
University of Warwick, Coventry, UK

Barnaby Martin
Durham University, UK

Sukanya Pandey
Utrecht University, The Netherlands

Daniël Paulusma
Durham University, UK

Mark Siggers
Kyungpook National University, Daegu, Republic of Korea

Siani Smith
University of Bristol, UK
Heilbronn Institute for Mathematical Research, Germany

Erik Jan van Leeuwen
Utrecht University, The Netherlands

Abstract

For a fixed set $\mathcal{H}$ of graphs, a graph G is $\mathcal{H}$-subgraph-free if G does not contain any $H \in \mathcal{H}$ as a (not necessarily induced) subgraph. A recent framework gives a complete classification on $\mathcal{H}$-subgraph-free graphs (for finite sets $\mathcal{H}$) for problems that are solvable in polynomial time on graph classes of bounded treewidth, NP-complete on subcubic graphs, and whose NP-hardness is preserved under edge subdivision. While a lot of problems satisfy these conditions, there are also many problems that do not satisfy all three conditions and for which the complexity in $\mathcal{H}$-subgraph-free graphs is unknown. We study problems for which only the first two conditions of the framework hold (they are solvable in polynomial time on classes of bounded treewidth and NP-complete on subcubic graphs, but NP-hardness is not preserved under edge subdivision). In particular, we make inroads into the classification of the complexity of four such problems: Hamilton Cycle, k-Induced Disjoint Paths, C_5-Colouring and Star 3-Colouring. Although we do not complete the classifications, we show that the boundary between polynomial time and NP-complete differs among our problems and also from problems that do satisfy all three conditions of the framework, in particular when we forbid certain subdivisions of the "H"-graph (the graph that looks like the letter "H"). Hence, we exhibit a rich complexity landscape among problems for $\mathcal{H}$-subgraph-free graph classes.

2012 ACM Subject Classification Mathematics of computing → Graph theory; Theory of computation → Graph algorithms analysis; Theory of computation → Problems, reductions and completeness

Keywords and phrases forbidden subgraph, complexity dichotomy, edge subdivision, treewidth

Digital Object Identifier 10.4230/LIPIcs.ISAAC.2024.47

Related Version *Full Version*: `https://doi.org/10.48550/arXiv.2211.14214`

Acknowledgements We are grateful to Matthew Johnson, Jelle J. Oostveen and Hans Bodlaender for useful discussions.

35th International Symposium on Algorithms and Computation (ISAAC 2024).
Editors: Julián Mestre and Anthony Wirth; Article No. 47; pp. 47:1–47:18
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Introduction

Graph containment relations, such as the (topological) minor and induced subgraph relations, have been extensively studied both from a graph-structural and algorithmic point of view. In this paper, we focus on the *subgraph relation*. If a graph H can be obtained from a graph G by a sequence of vertex deletions and edge deletions, then G contains H as a *subgraph*; otherwise, G is H*-subgraph-free*. For a set of graphs $\mathcal{H}$, a graph G is $\mathcal{H}$*-subgraph-free* if G is H-subgraph-free for every $H \in \mathcal{H}$; if $\mathcal{H} = \{H_1, \ldots, H_p\}$, then we also write that G is $(H_1, \ldots, H_p)$-subgraph-free. Graph classes closed under deletion of edge and vertices are called *monotone* [2, 8], and every monotone graph class $\mathcal{G}$ can be characterized by a unique (and possibly infinite) set of forbidden induced subgraphs $\mathcal{H}_{\mathcal{G}}$. We determine the complexity of two connectivity problems Hamilton Cycle and k-Induced Disjoint Paths, and two colouring problems C_5-Colouring and Star 3-Colouring on $\mathcal{H}$-subgraph-free graphs for various families $\mathcal{H}$. We focus on families $\mathcal{H}$ consisting of certain *subdivided "H"-graphs* $\mathbb{H}_i$, where $\mathbb{H}_1$ looks like the letter "H" (see Fig. 1 for the definition of the graphs $\mathbb{H}_i$). At first sight, these problems appear to have not much in common. Moreover, the graphs $\mathbb{H}_i$ might also seem arbitrary. However, these problems turn out to be well suited for a combined study, as they fit in a more general framework, in which the graphs $\mathbb{H}_i$ play a crucial role.

Context

If a graph problem is computationally hard, it is natural to restrict the input to some special graph class. Ideally we would like to know exactly which properties P such a graph class $\mathcal{G}$ must have such that any hard graph problem that satisfies some conditions C becomes easy on $\mathcal{G}$. The distinction between "*easy*" and "*hard*" means, in this paper, P versus NP-complete, but could also mean P versus Π_{2k}^{P}-complete [13], or almost-linear versus at-least-quadratic [20]. We first discuss some natural conditions C.

A graph is *subcubic* if every vertex has degree at most 3, or equivalently if is $K_{1,4}$-subgraph-free, where $K_{1,4}$ denotes the 5-vertex star. For $p \geq 1$, the *p-subdivision* of an edge $e = uv$ of a graph G replaces e by a path of $p + 1$ edges with endpoints u and v. The *p-subdivision* of a graph G is the graph obtained from G after p-subdividing each edge; see also Fig. 1. For a graph class $\mathcal{G}$ and an integer p, we let $\mathcal{G}^p$ be the class consisting of the p-subdivisions of the graphs in $\mathcal{G}$. A graph problem Π is hard *under edge subdivision of subcubic graphs* if for every $j \geq 1$ there is an $\ell \geq j$ such that: if Π is hard for the class $\mathcal{G}$ of subcubic graphs, then Π is hard for $\mathcal{G}^\ell$. We can now say that a graph problem Π has property:

- C1 if Π is easy for every graph class of bounded tree-width;
- C2 if Π is hard for subcubic graphs (or equivalently, $K_{1,4}$-subgraph-free graphs);
- C3 if Π is hard under edge subdivision of subcubic graphs;
- C4 if Π is hard for planar graphs;
- C5 if Π is hard for planar subcubic graphs.

We say that Π is a C123-problem if it satisfies C1, C2 and C3, while for example Π is a C1$\not{3}$-problem if it satisfies C1 but not C3, and so on.

Classical results of Robertson and Seymour [29] yield the following two *meta-classifications*. For all sets $\mathcal{H}$, a C14-problem Π is easy on $\mathcal{H}$-minor-free graphs if $\mathcal{H}$ contains a planar graph, or else it is hard. For all sets $\mathcal{H}$, a C15-problem Π is easy on $\mathcal{H}$-topological-minor-free graphs if $\mathcal{H}$ contains a planar subcubic graph, or else it is hard. No meta-classification for the induced subgraph relation exists (apart from a limited one [20] that is a direct consequence of the treewidth dichotomy [26]). However, for the subgraph relation, known results on Independent Set [3], Dominating Set [3], Long Path [3], Max-Cut [22] and List

Figure 1 [7] Left: A graph in $\mathcal{S}$: the graph $S_{3,3,3}+P_2+P_3+P_4$, where $S_{3,3,3}$ is the 2-subdivision of the claw $K_{1,3}$. Right: the "H"-graph $\mathbb{H}_1$ and the graph $\mathbb{H}_3$; for $i \geq 2$, the graph $\mathbb{H}_i$ ($i \geq 2$) is obtained from $\mathbb{H}_1$ by $(i-1)$-subdividing the edge that joins the middle vertices of the two P_3s.

Colouring [17] for monotone graph classes that are *finitely defined* (so, where the associated set of forbidden subgraphs $\mathcal{H}$ is finite) were recently unified and extended in [20]. This led to a new meta-classification, where the set $\mathcal{S}$ consists of all graphs, in which every connected component is either a path or a subcubic tree with exactly one vertex of degree 3 (see Fig. 1).

▶ **Theorem 1** ([20]). *For any* finite *set of graphs $\mathcal{H}$, a* C123*-problem Π is easy on $\mathcal{H}$-subgraph-free graphs if $\mathcal{H}$ contains a graph from $\mathcal{S}$, or else it is hard.*

The easy part of Theorem 1 holds because a class of $\mathcal{H}$-subgraph-free graphs satisfies C1 if and only if $\mathcal{H}$ contains a graph from $\mathcal{S}$ [28]. The hard part follows from combining C2 and C3, as discussed below. In [20], 20 C123-problems were identified on top of the five above.

Our Focus

Many graph problems are not C123. See [6] and [7] for partial complexity classifications of the C$\not{1}$23-problems Subgraph Isomorphism and Steiner Forest, respectively, for H-subgraph-free graphs and [21] for partial complexity classifications of the C1$\not{2}$3-problems (Independent) Feedback Vertex Set, Connected Vertex Cover, Colouring (see also [18]) and Matching Cut for H-subgraph-free graphs (note that if a problem does not satisfy C2, then C3 is implied). Here, we consider the question: *Can we classify the complexity of C12$\not{3}$-problems on monotone graph classes?*

Why the Graphs $\mathbb{H}_i$

All C1-problems are easy on $\mathcal{H}$-subgraph-free graphs if $\mathcal{H}$ has a graph from $\mathcal{S}$ [28]. The infinite set $\mathcal{M} = \{C_3, C_4, \ldots, K_{1,4}, \mathbb{H}_1, \mathbb{H}_2, \ldots\}$ of minimal graphs not in $\mathcal{S}$ is a maximal antichain in the poset of connected graphs under the subgraph relation. Conditions C2 and C3 ensure that for every finite set $\mathcal{M}'$, C123-problems are hard on $\mathcal{M}'$-subgraph-free graphs if $\mathcal{M}' \subseteq \mathcal{M}$. If C3 is not satisfied, this is no longer guaranteed. Hence, a natural starting point to answer our research question is to determine for which finite subsets $\mathcal{M}' \subseteq \mathcal{M}$, C12-problems are still easy on $\mathcal{M}'$-subgraph-free graphs. So consider a C12-problem Π that is not C3. Let $\mathcal{M}'$ be a finite subset of $\mathcal{M}$. If $\mathcal{M}' = \{K_{1,4}\}$, then Π is hard for $\mathcal{M}'$-subgraph-free graphs due to C2. Hence, $\mathcal{M}'$ must contain at least one C_s or $\mathbb{H}_i$. The *girth* of a graph (that is not a forest) is the length of a shortest cycle in it. We say that Π has property:

- C2' if for all $g \geq 3$, Π is hard for subcubic graphs of girth at least g.

A graph is subcubic and of girth $g \geq 4$ if and only if it is $(K_{1,4}, C_3, \ldots, C_{g-1})$-subgraph-free. So if Π is not only C12, but even a C12'-problem, then Π is hard on $\mathcal{M}'$-subgraph-free graphs unless $\mathcal{M}'$ contains some $\mathbb{H}_i$. This makes studying the graphs $\mathbb{H}_i$ even more pressing.

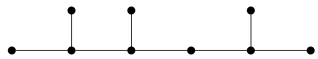

Figure 2 The tree T.

Our Testbed Problems

We take, as mentioned, the following four testbed problems:

(i) Hamilton Cycle, which is to decide if a graph G has a *Hamiltonian cycle*, i.e., a cycle through all vertices of G. This problem satisfies C1 [4], and it is NP-complete even for bipartite subcubic graphs of girth g, for every $g \geq 3$ [2]. Hence, it is even a C12'-problem.

(ii) k-Induced Disjoint Paths, which is to decide, given a graph G and pairwise disjoint vertex pairs $(s_1, t_1), (s_2, t_2), \ldots (s_k, t_k)$ for some fixed $k \geq 2$, if G has k *mutually induced* s_i-t_i-paths P^i, i.e., $P^1, \ldots, P^k$ are pairwise vertex-disjoint and there are no edges between vertices from different P^i and P^j. For every $k \geq 2$, this problem satisfies C1 due to Courcelle's Theorem [10] and also satisfies C2 [24]. Hence, it is a C12-problem for all $k \geq 2$.

(iii) C_5-Colouring, which is to decide if a graph G has a *homomorphism* (C_5-*colouring*) to the 5-cycle C_5 , i.e., a mapping $f : V(G) \to V(C_5)$ such that for every $uv \in E(G)$, it holds that $f(u)f(v) \in E(C_5)$. The problem satisfies C1 [12] and C2 [15]. Hence, it is a C12-problem.

(iv) Star 3-Colouring, which is to decide if a graph G has a *star* 3-*colouring*, i.e., a mapping $f : V(G) \to \{1, 2, 3\}$ such that for every i, the set U_i of vertices of G mapped to i is independent (so, f is a 3-*colouring*), and moreover, $U_1 \cup U_2$, $U_1 \cup U_3$, $U_2 \cup U_3$ all induce a disjoint union of stars. The problem satisfies C1 due to Courcelle's Theorem [10], and it is NP-complete for bipartite planar subcubic graphs of girth at least g, for every $g \geq 3$ [31]. Hence, it is even a C12'-problem.

We do not know if k-Induced Disjoint Paths and C_5-Colouring are C12', even though C_5-Colouring is NP-complete for graphs of maximum degree $6 \cdot 5^{13}$ and girth at least g, for all $g \geq 3$ (see Section 2.1).

All four problems violate C3. For $p \geq 3$, C_5-Colouring and Star 3-Colouring become true (all yes-instances) under p-subdivision, while Hamilton Cycle becomes false (all no-instances, unless we started with a cycle), and k-Induced Disjoint Paths reduces to the polynomial-time solvable problem k-Disjoint Paths [30, 33], which only requires the paths in a solution to be pairwise vertex-disjoint. See Section 3. We also note the following. First, when k is part of the input, Disjoint Paths and Induced Disjoint Paths are C123-problems [20]. Second, instead of C_5-Colouring we could have considered C_{2i+1}-Colouring, which is a C12-problem for all $i \geq 2$ [12, 15]. Third, Star-k-Colouring does not satisfy C2 for large k, as all subcubic graphs are star 10-colourable (as shown in Section 3).

Our Results

We show that the complexity of our four problems differ from each other and also from C123-problems, when we forbid certain graphs $\mathbb{H}_i$. We first show that C1-problems, and thus C12-problems, are easy on $(\mathbb{H}_\ell, \mathbb{H}_{\ell+1}, \ldots)$-subgraph-free graphs for every $\ell \geq 1$ and on $(\mathbb{H}_i, \mathbb{H}_{2i}, \mathbb{H}_{3i}, \ldots)$-subgraph-free graphs for every $i \geq 1$ (so, in particular if we forbid all even $\mathbb{H}_i$), as all these graph classes have bounded treewidth, as we show in Section 4. In contrast, any hard problem for bipartite graphs in which one partition class has maximum degree 2 is hard on $(\mathbb{H}_1, \mathbb{H}_3, \ldots)$-subgraph-free graphs (so, if we forbid all odd $\mathbb{H}_i$): every path between vertices of degree at least 3 has even length. The NP-hardness reduction in [1] shows that Star 3-Colouring is such a problem (see also Section 2.2).

The above results immediately give us Theorem 5. For the other three problems, we prove additional results. In Section 5 we show that Hamilton Cycle is polynomial-time solvable for $\mathbb{H}_\ell$-subgraph-free graphs if $\ell = 3$ by doing this for the superclass of T-subgraph-free graphs (T is the tree shown in Figure 2). For $\ell \in \{1, 2\}$ this was proven in [25]. On a side note, there exist trees T^* for which Hamilton Cycle is NP-complete over T^*-subgraph-free graphs. We refer to [23, 25] for examples of such trees T^*, which are not subdivided "H"-graphs $\mathbb{H}_i$.

In Section 6 we prove that for all $k \geq 2$, k-Induced Disjoint Paths is polynomial-time solvable for $\mathbb{H}_\ell$-subgraph-free graphs for $\ell \in \{1, 2\}$, but NP-complete for subcubic $(\mathbb{H}_4, \ldots, \mathbb{H}_\ell)$-subgraph-free graphs for all $\ell \geq 4$. For the first result, we first apply the algorithm for k-Disjoint Paths [30]. If this yields a solution that is not mutually induced, we apply a reduction rule and repeat the process on a smaller instance. For the second result, we carefully adapt the proof of [24] that shows that the problem of deciding if a subcubic graph contains an induced cycle between two given degree 2-vertices is NP-complete.

In Section 7 we determine all C_5-*critical* $\mathbb{H}_3$-subgraph-free graphs, which are not C_5-colourable unlike every proper subgraph of them. We show that this leads to a polynomial-time algorithm for $\mathbb{H}_3$-subgraph-free graphs that is even *certifying*. In contrast, the problem is NP-complete for the "complementary" class of $(\mathbb{H}_1, \mathbb{H}_2, \mathbb{H}_4, \mathbb{H}_5, \mathbb{H}_7, \mathbb{H}_8, \ldots)$-subgraph-free graphs (see Section 2.3).

The above results yields the following state-of-the-art summaries:

▶ **Theorem 2.** Hamilton Cycle *is polynomial-time solvable for $(\mathbb{H}_\ell, \mathbb{H}_{\ell+1}, \ldots)$-subgraph-free graphs ($\ell \geq 1$), for $(\mathbb{H}_i, \mathbb{H}_{2i}, \mathbb{H}_{3i}, \ldots)$-subgraph-free graphs ($i \geq 1$) and for $\mathbb{H}_\ell$-subgraph-free graphs ($\ell \in \{1, 2, 3\}$).*

▶ **Theorem 3.** *For all $k \geq 2$, k-*Induced Disjoint Paths *is polynomial-time solvable for $\mathbb{H}_\ell$-subgraph-free graphs ($\ell \in \{1, 2\}$), for $(\mathbb{H}_\ell, \mathbb{H}_{\ell+1}, \ldots)$-subgraph-free graphs ($\ell \geq 1$) and for $(\mathbb{H}_i, \mathbb{H}_{2i}, \mathbb{H}_{3i}, \ldots)$-subgraph-free graphs ($i \geq 1$), but* NP*-complete for subcubic $(\mathbb{H}_4, \ldots, \mathbb{H}_\ell)$-subgraph-free graphs ($\ell \geq 4$).*

▶ **Theorem 4.** C_5-Colouring *is polynomial-time solvable for $\mathbb{H}_3$-subgraph-free graphs, for $(\mathbb{H}_\ell, \mathbb{H}_{\ell+1}, \ldots)$-subgraph-free graphs ($\ell \geq 1$) and for $(\mathbb{H}_i, \mathbb{H}_{2i}, \mathbb{H}_{3i}, \ldots)$-subgraph-free graphs ($i \geq 1$), but* NP*-complete for $(\mathbb{H}_1, \mathbb{H}_2, \mathbb{H}_4, \mathbb{H}_5, \mathbb{H}_7, \mathbb{H}_8, \ldots)$-subgraph-free graphs.*

▶ **Theorem 5.** Star 3-Colouring *is polynomial-time solvable for $(\mathbb{H}_\ell, \mathbb{H}_{\ell+1}, \ldots)$-subgraph-free graphs ($\ell \geq 1$) and $(\mathbb{H}_i, \mathbb{H}_{2i}, \mathbb{H}_{3i}, \ldots)$-subgraph-free graphs ($i \geq 1$), but* NP*-complete for $(\mathbb{H}_1, \mathbb{H}_3, \mathbb{H}_5 \ldots)$-subgraph-free graphs.*

We note that the complexity classifications above indeed differ except perhaps for Hamilton Cycle and k-Induced Disjoint Paths. Hence, Theorems 2–5 give clear evidence of a rich landscape for C12-problems on $\mathcal{H}$-subgraph-free graphs. In Section 8 we discuss open problems resulting from our study.

2 Some Basic Results

In this section, we provide further details for some statements made in Section 1.

2.1 C_5-Colouring for Bounded Degree and Large Girth

The k-Colouring problem is to decide if a graph G has a *k-colouring*, which is a mapping $c : V(G) \to \{1, \ldots, k\}$ such that $c(u) \neq c(v)$ for any two adjacent vertices u and v of G. We need a result of Emden-Weinert, Hougardy and Kreuter:

▶ **Theorem 6** ([14]). *For all $k \geq 3$ and all $g \geq 3$, k-Colouring is NP-complete for graphs with girth at least g and with maximum degree at most $6k^{13}$.*

We now repeat the proof of Chudnovsky et al. [9], which comes down to replacing each edge of an input graph G of 5-Colouring, which we may assume has girth at least g and maximum degree at most $6 \cdot 5^{13}$ due to Theorem 6, by a path of length 3. This yields a new graph G' of girth at least g, such that G and G' have the same maximum degree. Hence, we derive the following result.

▶ **Proposition 7.** *For every $g \geq 3$, C_5-Colouring is NP-complete for graphs with girth at least g and with maximum degree at most $6 \cdot 5^{13}$.*

2.2 The Standard NP-hardness Reduction to Star-3-Colouring

For reference, we explain the gadget from Albertson et al. [1] that yields the following result.

▶ **Theorem 8** ([1]). *Star 3-Colouring is NP-complete for planar bipartite graphs in which one partition class has size 2.*

Proof. Reduce from 3-Colourability which is known to be NP-complete even for planar graphs [11]. Let G be a planar graph. Replace each edge e by three new vertices a_e, b_e, c_e that are made adjacent only to the two end-vertices of e in G. Let G' be the resulting graph. Then every vertex of $V(G') \setminus V(G)$ has degree 2 in G. Moreover, G' is planar and bipartite with partition classes $V(G') \setminus V(G)$ and $V(G)$. It remains to observe that G has a 3-colouring if and only if G' has a star 3-colouring. ◀

2.3 The Standard NP-hardness Reduction to C_5-Colouring

We make the following observation.

▶ **Proposition 9.** *C_5-Colouring is NP-complete for $(\mathbb{H}_1, \mathbb{H}_2, \mathbb{H}_4, \mathbb{H}_5, \ldots)$-subgraph-free graphs.*

Proof. It is well known [19] and easy to see that there is a reduction from K_5-Colouring, which is to decide if a graph has a K_5-colouring, that is, a homomorphism from G to the complete graph K_5 on five vertices. This problem is well known to be NP-complete [19]. Let G be a graph, and let G' be the 2-subdivision of G. We note that G' is $(\mathbb{H}_1, \mathbb{H}_2, \mathbb{H}_4, \mathbb{H}_5, \ldots)$-subgraph-free (but may contain many instances of $\mathbb{H}_\ell$ where $\ell = 0 \mod 3$). Moreover, G has a K_5-colouring if and only if G' has a C_5-colouring. ◀

3 The Four Testbed Problems Do Not Satisfy C3

In this section we show that none of our four problems satisfy C3. We use the following notation in this section: for a graph G and an integer $p \geq 1$, let G_p be the p-subdivision of G (which we recall is the graph obtained from G after subdividing each edge of G exactly p times).

▶ **Proposition 10.** Hamilton Cycle *does not satisfy C3.*

Proof. We observe that for every graphs G and every $p \geq 1$, G_p is a no-instance of Hamilton Cycle unless G was a cycle. ◀

▶ **Proposition 11.** k-Induced Disjoint Paths *does not satisfy C3.*

Proof. Under any kind of subdivision, k-Induced Disjoint Paths reduces to k-Disjoint Paths over the same graph, which is in P for all $k \geq 2$, as shown in [33] for $k = 2$ and in [30] for every $k \geq 3$. ◀

▶ **Proposition 12.** C_5-Colouring *does not satisfy C3.*

Proof. We first prove that for all $p \geq 4$, and for all $x, y \in V(C_5)$, there is a walk of length p in C_5 from x to y. First let $p = 4$. To walk a distance of zero: walk two forward then two back. To walk at distance one (without loss of generality) forward: walk four backward. To walk at distance two (without loss of generality) forward: walk one back, one forward, and two forward. Now let $p = 5$. To walk a distance of zero: walk five forward. To walk at distance one (without loss of generality) forward: walk two forward, two back and one forward. To walk at distance two (without loss of generality) forward: walk one back, one forward, and three back. Finally, let $p \geq 6$. Keep moving one forward then one back until one of the two previous cases applies.

Now let G be a graph. We give each vertex in G a label from $\{1, \ldots, 5\}$. From the above it follows that for every $p \geq 3$, we can extend c to a homomorphism from G_p to C_5; in other words, G_p is a yes-instance of C_5-Colouring. ◀

▶ **Proposition 13.** Star 3-Colouring *does not satisfy C3.*

Proof. Let G be a graph. We show that for all $p \geq 3$, G_p is a yes-instance of Star 3-Colouring. We do this by giving each vertex in G a label from $\{1, 2, 3\}$. The resulting labelling c might not be a 3-colouring, but this is not important: we will show that we can extend c to a star 3-colouring of G_p as follows.

Consider an edge e in G and let P be the corresponding path (of $p + 1$ edges) in G_p. It suffices to give two star 3-colourings of this path, so that the first three vertices are distinct colours and the last three vertices are distinct colours: one in which the first and last vertices are the same colour and one in which they are a different colour. Let us identify a 3-colouring of P by a sequence of length $p + 1$ over $\{1, 2, 3\}$. If $p + 1$ is a multiple of three, then use $(123)^{\frac{p+1}{3}}$ for the different colour and $(123)^{\frac{p+1}{3}-1}231$ for the same colour. If $p + 1$ is 1 mod 3, then use $(123)^{\frac{p}{3}-1}2132$ for the different colour and $(123)^{\frac{p}{3}}1$ for the same colour. If $p + 1$ is 2 mod 3, then use $(123)^{\frac{p-1}{3}}12$ for the different colour and $(123)^{\frac{p-1}{3}}21$ for the same colour. ◀

We finish this section with another small observation. Namely, we cannot generalise our result for Star 3-Colouring to Star k-Colouring for any $k \geq 3$, as for large k the problem no longer satisfies C2. In fact, we prove even a stronger statement. A k-colouring of a graph G is said to be *injective* if for every vertex $u \in V(G)$, every neighbour of u is assigned a different colour, or in other words, the union of any two colour-classes induce a disjoint union of isolated vertices and edges. So, any injective k-colouring is a star k-colouring (but the reverse does not necessarily hold, for instance the P_3 is star 2-colourable but has no injective 2-colouring).

▶ **Proposition 14.** *For $k \geq 10$, all subcubic graphs have an injective 10-colouring.*

Proof. It suffices to prove the statement for $k = 10$. We do this by induction. For the base case, a graph with one vertex is star 10-colourable. Now take a vertex v in a graph G and assume $G \setminus \{v\}$ has an injective 10-colouring. As G is subcubic, v has at most three neighbours, each of which have at most two more neighbours each. Thus there are at most nine vertices whose colour we wish to avoid. As we have ten colours in total, this means that we can safely colour v. ◀

4 Bounded Treewidth Results

A graph G contains H as a *minor* if G can be modified to H by a sequence of vertex deletions, edge deletions and edge contractions; if not, then G is *H-minor-free.*

▶ **Proposition 15.** *For every $\ell \geq 1$, the class of $(\mathbb{H}_\ell, \mathbb{H}_{\ell+1}, \ldots)$-subgraph-free graphs has bounded treewidth.*

Proof. For $\ell \geq 1$, a $(\mathbb{H}_\ell, \mathbb{H}_{\ell+1}, \ldots)$-subgraph-free graph is $\mathbb{H}_\ell$-minor-free. For every forest F, all F-minor-free graphs have pathwidth, and thus treewidth, at most $|V(F)| - 2$ [5]. ◀

▶ **Proposition 16.** *For every $n \geq 1$, the class of $(\mathbb{H}_n, \mathbb{H}_{2n}, \mathbb{H}_{3n}, \ldots)$-subgraph-free graphs has bounded treewidth.*

Proof. If a class of graphs has unbounded treewidth, then every grid appears as a minor in some graph [29]. Let us explain the argument for $n = 2$ first. We consider that the 3×3-grid appears as a minor in some graph G in our class and let f be the minor map from G to the 3×3-grid. Consider the three vertices in the 3×3-grid that form the central row as u, v, w (in succession). Choose $u' \in f^{-1}(u), v' \in f^{-1}(v), w' \in f^{-1}(w)$ so that u', v', w' have degree greater than 2, noting that such vertices must exist. If the distance in G between u' and v' is even, of length $2i$, then there is an $\mathbb{H}_{2i}$ subgraph in G with central path from u' to v'. If the distance in G between v' and w' is even, of length $2i$, then there is an $\mathbb{H}_{2i}$ subgraph in G with central path from v' to w'. Else, there is a path of even length $4i$ from u' to w' and then there is an $\mathbb{H}_{4i}$ subgraph in G with central path from u' to w'.

For $(\mathbb{H}_n, \mathbb{H}_{2n}, \ldots)$-subgraph-free graphs, we consider the Abelian group $(\mathbb{Z}/n\mathbb{Z})$. The Davenport constant of an Abelian group G is the minimum d so that any sequence of elements of G contains a non-empty consecutive subsequence of zero-sum (that adds to the identity element 0). It is known that for $(\mathbb{Z}/n\mathbb{Z})$ the Davenport constant is n (see page 24 in [16]). Take an $(n+1) \times (n+1)$-grid and consider some row not at the top or bottom of the grid with vertices $w_1, \ldots, w_{n+1}$ in succession. Consider some $w'_1 \in f^{-1}(w_1), \ldots, w'_{n+1} \in f^{-1}(w_{n+1})$ where f is the minor map as before, and the distances x_i between w'_{i+1} and w'_i. Using the Davenport constant, there is a subsequence $x_j, \ldots, x_{j'}$ $(j' > j)$ such that $x_j + \ldots + x_{j'} = 0 \bmod n$. Now choose $w'_j, \ldots, w'_{j'+1}$ as the central path in some $\mathbb{H}_{in}$. ◀

5 Hamilton Cycle

In this section we show Theorem 17. Due to the page limit we have omitted the proofs of some of the claims in the proof of Theorem 17.

▶ **Theorem 17.** Hamilton Cycle *is polynomial-time solvable for T-subgraph-free graphs.*

Proof. Let G be a T-subgraph-free graph. We call vertices of degree 2 in G *white* and vertices of degree at least 3 *black*. The *black graph* is a subgraph of G induced by black vertices and a *black component* is a connected component in the black graph.

We first describe some helpful rules to solve the problem and a set of reductions simplifying the input graph, i.e. reductions transforming G into a graph G' that has fewer edges and/or vertices and that has a Hamiltonian cycle if and only if G has. We emphasize that by deleting an edge or a vertex from an H-subgraph-free graph, we obtain an H-subgraph-free graph again.

We start with some obvious rules:

(R1) if the graph has vertices of degree 0 or 1, then stop: G has no Hamiltonian cycle.

(R2) if the graph contains a vertex adjacent to more than two white vertices, then stop: G has no Hamiltonian cycle.

(R3) if the graph is disconnected, then stop: G has no Hamiltonian cycle.

(R4) if the graph contains a vertex v adjacent to exactly two white vertices, then delete the edges connecting v to all other its neighbours (if there are any).

Now we introduce a reduction applicable to a graph G containing an induced subgraph shown on the left in Figure 3, in which vertices a, b, c have degree 3 in G. The reduction depends on the degree of x. If the degree of x is also 3, the reduction consists in deleting the edges ab and xc. Otherwise, it transforms the graph as shown in Figure 3. We refer to this reduction as the *diamond reduction* and denote it by (R5).

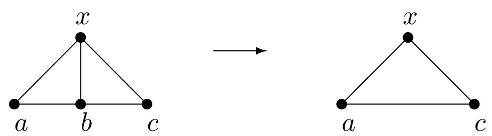

Figure 3 The diamond reduction: it is applicable to a graph G containing an induced subgraph shown on the left, in which vertices a, b, c have degree 3 in G. If the degree of x is also 3, the reduction consists in deleting the edges ab and xc. Otherwise, the reduction consists in deleting vertex b and introducing the edge ac.

We omit the proof of the next two claims.

▷ Claim 18. Let G' be a graph obtained from G by the diamond reduction. Then G has a Hamiltonian cycle if and only if G' has a Hamiltonian cycle. Moreover, if G is T-subgraph-free, then so is G'.

In Figure 4, we illustrate the *butterfly reduction* denoted by (R6).

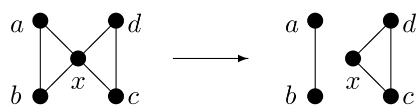

Figure 4 The butterfly reduction: it is applicable to a graph G with an induced subgraph shown on the left, in which vertices a, b, c have degree 3 in G, and moreover, a and b have white neighbours.

▷ Claim 19. Let G' be a graph obtained from G by the butterfly reduction. Then G has a Hamiltonian cycle if and only if G' has a Hamiltonian cycle.

In our algorithm we implement the above rules and reductions whenever they are applicable. We now develop more reductions allowing us to bound the number of vertices in black components. We assume that none of the above rules and reductions is applicable to G.

We omit the proof of the next claim.

▷ Claim 20. Let x be a vertex of degree at least 13. If the neighbourhood of x does not contain two adjacent vertices of degree 3, then G has no Hamiltonian cycle. Otherwise, G has a Hamiltonian cycle if and only if $G - x$ has.

Application of Claim 20 to vertices of large degree either shows that G has no Hamiltonian cycle or reduces the input graph to a graph of maximum degree 12. We will refer to this reduction as the *large degree reduction* and will denote it by (R7).

We omit the proof of the next claim.

▷ Claim 21. The black graph has no induced paths of length 8.

Since graphs of diameter D and maximum degree Δ have fewer than $\frac{\Delta}{\Delta-2}(\Delta-1)^D$ vertices, we conclude that after eliminating vertices of large degree, every black component has fewer than $\frac{12}{10}11^7$ vertices.

To develop more rules and reductions, assume G has a Hamiltonian cycle C. We can further assume that not all vertices of the graph are black, since otherwise the graph contains fewer than $\frac{12}{10}11^7$ vertices, in which case we can solve the problem by brute-force. A sequence of consecutive vertices of C surrounded by white vertices will be called a *black interval*. Observe that each black interval consists of at least two vertices (according to (R4)).

Let K be a black component of G. We will call the vertices of K that have white neighbours the *contact vertices*. Note that K may consists of one or more intervals. Each interval gives rise to exactly two contact vertices. Hence, the number of contact vertices in K is even.

In our next claim, whose proof we omit, we show that for T-subgraph-free graphs, the number of intervals is at most 2.

▷ Claim 22. Any black component consists of at most two intervals.

By Claim 22, if G has a Hamiltonian cycle, then every black component has two or four contact vertices.

(R8) If a black component K has exactly two contact vertices, check if K has a Hamiltonian path connecting the contact vertices. If such a path does not exist, then stop: the input graph has no Hamiltonian cycle. Otherwise, choose arbitrarily a Hamiltonian path connecting the contact vertices, include the edges of the path in the solution and delete all other edges from K.

Rule (R8) destroys black components with two contact vertices, i.e. after its implementation all vertices in such components become white.

Now we discuss the case where each black component has exactly four contact vertices. Let K be such a component with contact vertices v_1, v_2, v_3, v_4. If G has a Hamiltonian cycle, then the vertices of K can be partitioned into two parts each of which forms a path connecting a pair of contact vertices. We will call such a partition a *pairing* (of contact vertices) and will refer to a pairing as the set of edges in the two paths. Also, we will say that two pairings are of the same type, if they pair the contact vertices in the same way. Clearly, if all possible pairings in K have the same type, then it is irrelevant which one to choose, since non-contact vertices of K have no neighbours outside of K.

The above discussion justifies the following two rules.

(R9) If a black component K with four contact vertices does not admit any pairing, then stop: the input graph has no Hamiltonian cycle.

(R10) If in a black component K with four contact vertices all possible pairings have the same type, then choose arbitrarily any such pairing and delete all other edges from K. If this procedure disconnects the graph, then stop: the input graph has no Hamiltonian cycle.

Finally, we analyse the situation when each black component of G admits pairings of at least two different types.

▷ Claim 23. If each black component of the (connected) graph G admits pairings of at least two different types, then G has a Hamiltonian cycle.

Proof. Let K be a black component with contact vertices v_1, v_2, v_3, v_4 and let B and R be two pairings of different types, say B pairs v_1 with v_2 and v_3 with v_4, while R pairs v_1 with v_3 and v_2 with v_4. Assume that

(1) the deletion of all edges of K except for the edges of B disconnects the graph into two components C_{12} (containing vertices v_1 and v_2) and C_{34} (containing vertices v_3 and v_4), and

(2) the deletion of all edges of K except for the edges of R disconnects the graph into two components C_{13} (containing vertices v_1 and v_3) and C_{24} (containing vertices v_2 and v_4).

Note that (1) separates v_1 from v_3 and v_4, while (2) separates v_1 from v_4. Therefore, after the deletion of *all* edges of K vertex v_1 is separated from all other contact vertices. In other words, after the deletion of all edges of K, vertices v_1, v_2, v_3, v_4 belong to pairwise different connected components, say V_1, V_2, V_3, V_4, respectively.

We observe that in each connected component V_i vertex v_i has degree 1 (it is adjacent to a white vertex only). Any other vertex of odd degree in V_i (if there is any) is black, i.e. appears in some black component K'. In the graph $G[K']$ the number of odd vertices is even (by the Handshake lemma). Attaching to $G[K']$ four white neighbours changes the parity of exactly four vertices of K' and hence leaves the number of vertices of K' with odd degrees in the graph G even. Since all vertices of K' belong to only one of the components V_i, we conclude that in each component V_i the number of vertices of odd degree is odd. This is not possible by the Handshake lemma and hence either (1) or (2) is not valid, i.e. we can keep one of the pairings and delete all other edges of K without disconnecting G. This operation destroys K, i.e. makes all vertices of K white.

Applying the above arguments to all black components, one by one, we transform G into a connected graph in which all vertices are white, i.e. to a Hamiltonian cycle. ◁

We summarize the discussion in the following algorithm to solve the problem.

1. Apply rules and reductions (R1) – (R7) as long as they are applicable.
2. If the algorithm did not stop at Step 1 and the graph has fewer than $\frac{12}{10}11^7$ vertices, then solve the problem by brute-force. Otherwise, check the number of contact vertices in black components. If there is a black component with the number of contact vertices different from 2 or 4, then stop: G has no Hamiltonian cycle.
3. If the algorithm did not stop at Step 2, then apply (R8) to black components with two contact vertices, and (R9) and (R10) to black components with four contact vertices.
4. If the algorithm did not stop at Step 3, then find a Hamiltonian cycle according to Claim 23.

Reductions (R8), (R9), (R10) can be implemented in constant time, because the number of vertices in each black component is bounded by a constant. It is also obvious that all other rules, and hence all steps of the algorithm can be implemented in polynomial time. The correctness of the algorithm follows from the proofs of the claims. ◀

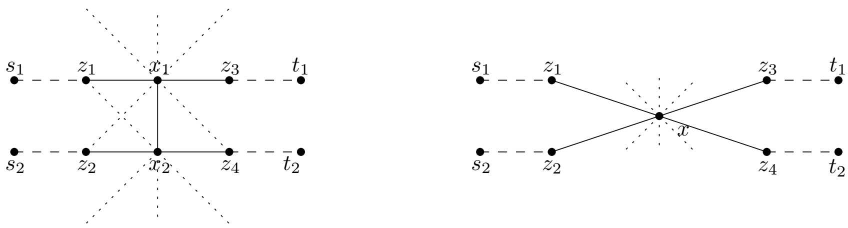

Figure 5 Rule 1. Possible connections in our subgraph (left). What we replace this subgraph with (right). Dotted lines are possible additional edges.

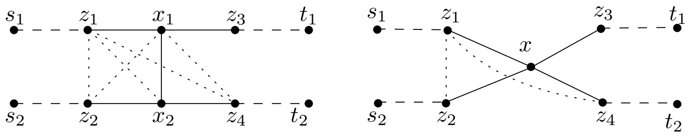

Figure 6 Rule 2. Possible connections in our subgraph (left). What we replace this subgraph with (right). Dotted lines are possible additional edges.

6 k-Induced Disjoint Paths

The case $H = \mathbb{H}_1$ follows from the observation that solutions of k-Induced Disjoint Paths with long paths are solutions of k-Disjoint Paths, which is polynomial-time solvable [30]. We omit the proof details. The case $H = \mathbb{H}_2$ is more involved.

▶ **Theorem 24.** *For all $k \geq 2$, k-*Induced Disjoint Paths *is polynomial-time solvable for $\mathbb{H}_1$-subgraph-free graphs.*

▶ **Theorem 25.** *For all $k \geq 2$, k-*Induced Disjoint Paths *is polynomial-time solvable for $\mathbb{H}_2$-subgraph-free graphs.*

Proof. First, branch on all $O(2^k n^{3k})$ options (so a polynomial number, as k is fixed) of solution paths that have at most three internal vertices. For each branch, we remove the guessed solution paths and the neighbours of the vertices on these paths. Let k still be the number of terminal pairs. We now only look for solution paths with at least four vertices. Branch on all $O(n^{4k})$ options of choosing the first two vertices a_z, b_z on the solution path from every terminal $z \in \{s_i, t_i\}$ for $i \in \{1, \dots, k\}$. In each branch, we remove all other neighbours of z, a_z from the graph, so every terminal z now has degree 1, while a_z has degree 2. We discard the branch if (†) $\{a_z, b_z\} \cap \{a_{z'}, b_{z'}\} \neq \emptyset$ for some terminals z, z' or one of a_z, b_z is the same or neighbours one of $a_{z'}, b_{z'}$ for some terminals z, z' not from the same terminal pair.

We now start a recursive procedure. We first preprocess the input. If b_{s_i} and b_{t_i} are adjacent for some $i \in \{1, \dots, k\}$, then we remove the solution path $s_i, a_{s_i}, b_{s_i}, b_{t_i}, a_{t_i}, t_i$ and their neighbours from the graph. If b_z and $b_{z'}$ are adjacent for some terminals z, z' that do not form a terminal pair, we discard the branch.

We run the polynomial-time algorithm for k-Disjoint Paths from [30] on the remaining terminal pairs. If this results in a no-answer, we discard the branch. Else, we found a solution $P_1, \dots, P_k$. We may assume that each path P_i is induced, or we may shortcut it. If $P_1, \dots, P_k$ is also a solution of k-Induced Disjoint Paths, then we return "yes". Otherwise, there is (say) an edge (x_1, x_2) between paths $x_1 \in P_1$ and $x_2 \in P_2$. We pick x_1 such that it is closest to t_1 on P_1 and under that condition we pick x_2 such that it is closest to t_2 on P_2.

Let z_1, z_3 be the two neighbours of x_1 on P_1 and z_2, z_4 the two neighbours of x_2 on P_2. We let $S = \{z_1, x_1, z_3, z_2, x_2, z_4\}$. Observe that S contains no terminal by † and the preprocessing. If any $z \in \{z_1, z_2, z_3, z_4\}$ has two neighbours outside of S, then G has a $\mathbb{H}_2$ as a subgraph. Thus we may assume (‡) that each $z \in \{z_1, z_2, z_3, z_4\}$ has at most one neighbour not in S.

By the choice of (x_1, x_2) and as P_1 is induced, z_3 has no neighbours in S except x_1. Suppose the edge (z_1, z_2) exists and one of $\{x_1, x_2\}$ has a neighbour outside of S. Then there is an $\mathbb{H}_2$ with middle path x_1, z_1, z_2 since $s_2 \notin S$. Suppose the edge (z_1, z_4) exists and one of $\{x_1, x_2\}$ has a neighbour outside of S. Then there is an $\mathbb{H}_2$ with middle path x_1, z_1, z_4 since $s_2 \notin S$. Now either the edges (z_1, z_2) and (z_1, z_4) do not exist (see Figure 5), or at least one of them exists and x_1, x_2 have no neighbours outside S (see Figure 6). In the former case, we apply Rule 1, while in the latter case, we apply Rule 2; see Fig. 5 and 6 for their description.

Rule 1 is safe: Suppose that we have a solution to k-Induced Disjoint Paths in G. If this solution uses no vertices in S, then it is already a solution to k-Induced Disjoint Paths in G'. Thus, it must use some vertex in S. If the solution does not use x_1 nor x_2, then recall that by ‡, each of z_1, z_2, z_3, z_4 has at most one neighbour outside of S, and thus the solution must avoid thus S entirely, a contradiction. If the solution uses both x_1 and x_2, then it must use the edge (x_1, x_2). We can substitute the edge (x_1, x_2) in the solution to k-Induced Disjoint Paths in G with x to obtain a solution to k-Induced Disjoint Paths in G'. Hence, without loss of generality, suppose the solution uses x_1. We can substitute this for x to obtain a solution to k-Induced Disjoint Paths in G', unless some other solution path runs through a neighbour q of x_2. Note q cannot be a terminal due to our preprocessing. Hence it has two neighbours p and r on this other solution path, and these are outside of $\{z_1, x_1, z_3\}$ because this path must avoid x_1 and any of its neighbours. But now p, q, r, q, x_2, x_1, z_1, x_1, z_3 forms an $\mathbb{H}_2$ (with middle path q, x_2, x_1), a contradiction.

Suppose we have a solution to k-Induced Disjoint Paths in G'. If this solution does not involve x, then it maps to a solution of k-Induced Disjoint Paths in G. Suppose now it does involve x. Suppose mapping x to either of x_1 or x_2 does not produce a solution to k-Induced Disjoint Paths in G. Then mapping x to either the edge (x_1, x_2) (or the symmetric (x_2, x_1)) must produce a solution to k-Induced Disjoint Paths in G.

Rule 2 is safe: Suppose we have a solution to k-Induced Disjoint Paths in G. If it uses no vertices in S, then it is already a solution to k-Induced Disjoint Paths in G'. Thus, it must use some vertex in S. Suppose the edge (z_1, z_2) exists, and the solution uses (z_1, z_2). Then by ‡, the solution does not use any other vertex from S and we can keep this edge to obtain a solution for k-Induced Disjoint Paths in G'. Suppose the edge (z_1, z_4) exists and the solution uses (z_1, z_4). Then by ‡, the solution does not use any other vertex from S and we can keep this edge to obtain a solution for k-Induced Disjoint Paths in G'.

If the solution uses both x_1 and x_2, then it must use the edge (x_1, x_2), and we can substitute (x_1, x_2) in the solution to k-Induced Disjoint Paths in G with x to obtain a solution to k-Induced Disjoint Paths in G'. Suppose it uses neither x_1 nor x_2. Then by ‡ and the fact that S is used, the solution must use either the edge (z_1, z_4) or (z_1, z_2) and we are in a previous case. Hence, without loss of generality, suppose the solution uses x_1. We can substitute this for x to obtain a solution to k-Induced Disjoint Paths in G'. This is safe, as x_1, x_2 have no neighbours outside S.

Suppose we have a solution to k-Induced Disjoint Paths in G'. If this solution does not involve x then it maps to a solution of k-Induced Disjoint Paths in G. Suppose now it does involve x. Suppose mapping x to either of x_1 or x_2 does not produce a solution to k-Induced Disjoint Paths in G. Then mapping x to either the edge (x_1, x_2) (or the symmetric (x_2, x_1)) must produce a solution to k-Induced Disjoint Paths in G.

Next, we show that any graph G' obtained after applying Rule 1 or 2 is also $\mathbb{H}_2$-subgraph-free. Suppose G' has an $\mathbb{H}_2$. Then this $\mathbb{H}_2$ must contain x. If x is a leaf in $\mathbb{H}_2$, then G already had this $\mathbb{H}_2$ involving either x_1 or x_2. Suppose x is a degree-3 vertex in this $\mathbb{H}_2$. If the neighbours of x in the $\mathbb{H}_2$ were all neighbours of x_1 or all neighbours of x_2 in G, then G already had this $\mathbb{H}_2$, a contradiction. Let z_1' and z_2' be the leafs of the $\mathbb{H}_2$ adjacent to x in G'.

Suppose z_1' and z_2' are both adjacent to x_2 and both not to x_1. Then the middle vertex of the $\mathbb{H}_2$ is only adjacent to x_1. Ideally, we would replace x by x_1, z_1' by z_1 and z_2' by z_3. This does not work if (say) z_1 is part of the $\mathbb{H}_2$. However, z_1' and z_2' are both not z_1, as z_1 is adjacent to x_1, and we would contradict our assumption on the adjacency of z_1' and z_2'. We now consider three cases, depending on where z_1 is in the $\mathbb{H}_2$.

Suppose z_1 is a leaf of the $\mathbb{H}_2$. By ‡ and the inducedness of paths, its neighbouring degree-3 vertex cannot be one of z_2, z_3, z_4. Hence, this must be the unique neighbour p of z_1 outside S. The other neighbours q, r of p on the $\mathbb{H}_2$, where r is the middle vertex, are both not z_3, as P_1 is induced. Hence, q, p, z_1, p, r, x_1, x_2, x_1, z_3 form an $\mathbb{H}_2$, a contradiction.

Suppose that z_1 is the middle vertex of the $\mathbb{H}_2$. By ‡, the other degree-3 vertex of the $\mathbb{H}_2$ cannot be z_2 or z_4, so it must be the unique neighbour p of z_1 outside S. The other neighbours q, r of p on the $\mathbb{H}_2$, which are both leafs of the $\mathbb{H}_2$, are both not z_3 since P_1 is induced. Hence, G has a $\mathbb{H}_2$ formed by q, p, r, p, z_1, x_1, x_2, x_1, z_3, a contradiction.

Suppose that z_1 is a degree-3 vertex of the $\mathbb{H}_2$. Let p be the unique neighbour of z_1 outside S; it is unique by ‡. Then one of p, z_2, z_3 must be the middle vertex of the $\mathbb{H}_2$ and the other two the leafs neighbouring z_1. If the middle vertex is z_2, then z_1', x_2, z_2', x_2, z_2, z_1, p, z_1, z_4 is a $\mathbb{H}_2$ in G, a contradiction. The other cases are similar. This concludes the argument when z_1' and z_2' are both adjacent to x_2 and both not to x_1.

Suppose instead that, say z_1', is adjacent to x_1 and the other, z_2', is adjacent to x_2. Let x', x'', z_1'', z_2'' form the remaining vertices of the $\mathbb{H}_2$ where x, x', x'' and z_1'', x'', z_2'' are both paths of length 2 in this $\mathbb{H}_2$. Thus, z_1', x, z_2', x, x', x'' and z_1'', x'', z_2'' form the $\mathbb{H}_2$ in G'. Without loss of generality, suppose x' was adjacent to x_1 in G. Now it is clear that z_1', x_1, x_2, x_1, x', x'' and z_1'', x'', z_2'' formed an $\mathbb{H}_2$ in G.

Finally, suppose that x is the degree-2 vertex in $\mathbb{H}_2$. Let z_1', x', z_2', x', x, x'', z_1'', x'', z_2'' be the paths that form the $\mathbb{H}_2$ in G'. Suppose, without loss of generality, that x' was adjacent to x_1 in G. If x'' was also adjacent to x_1 in G, then z_1', x', z_2', x', x_1, x'', z_1'', x'', z_2'' are paths that form an $\mathbb{H}_2$ in G. Suppose now that x'' was adjacent to x_2 but not x_1 in G and we may also assume that x' is adjacent to x_1 but not x_2. Now z_1', x', z_2', x', x_1, x_2, z_2, x_2, z_4 are paths that form a $\mathbb{H}_2$ in G, unless $\{z_2, z_4\} \cap \{z_1', z_2'\} \neq \emptyset$. Without loss of generality, suppose $z_2 = z_1'$. Note that $z_2 \neq s_2$ (recall that S does not contain any terminal). Let p be the next vertex on the path from t_2 to s_2 after z_2. Then p, z_2, x_2, z_2, x', x_1, z_1, x_1, z_3 is an $\mathbb{H}_2$ in G (note that $\{z_1, z_3\} \cap \{x', z_2, p\} = \emptyset$), a contradiction.

Finally, note that x_1 and x_2 cannot be z or a_z for some terminal z, as these vertices have degree 1 and 2 respectively, while x_1, x_2 have degree at least 3. Moreover, $\{x_1, x_2\} \neq \{b_z, b_z'\}$ for some terminals z, z' by our preprocessing. Hence, Rules 1 and 2 preserve †.

We can recognize and apply Rules 1 and 2 in polynomial time. This decreases the size of the graph by one vertex and we recurse. Hence, our algorithm runs in polynomial time. ◀

For our next result we follow the proof from Section 2.4 in [24] by carefully p-subdividing *some* of the edges of that construction. We omit the proof details.

▶ **Theorem 26.** *For all $k \geq 2$, k-*Induced Disjoint Paths *is* NP*-complete for subcubic $(\mathbb{H}_4, \ldots, \mathbb{H}_\ell)$-subgraph-free graphs for all $\ell \geq 4$.*

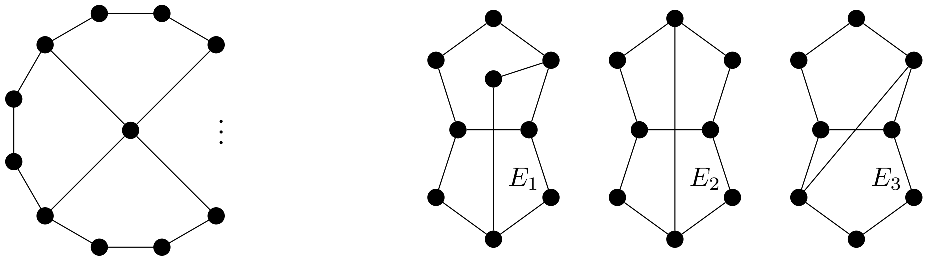

Figure 7 The C_5-flower F_n and the $\mathbb{H}_3$-subgraph-free C_5-critical graphs E_1, E_2 and E_3.

7 C₅-Colouring

In this section, we give our polynomial-time certifying algorithm for C_5-Colouring on $\mathbb{H}_3$-subgraph-free graphs. The *C_5-flower F_n* is the graph (see Figure 7) that we get from C_{3n} (for $n \geq 3$) by adding a new central vertex with an edge to every third vertex of C_{3n}. If n is odd, we call F_n an *odd C_5-flower*, and if it is even we call F_n an *even C_5-flower*. We refer to the graphs E_1, E_2 and E_3 shown in Figure 7 as *exceptional graphs*.

The following lemma (whose proof is a simple exercise) shows that all these graphs are *C_5-critical*, that is, they are not C_5-colourable but every proper subgraph of them is.

▶ **Lemma 27.** *The graph K_3, the odd flowers F_n for odd $n \geq 3$ and the exceptional graphs E_1, E_2 and E_3 are all $\mathbb{H}_3$-subgraph-free and C_5-critical.*

We can now show a structural result, which we use to prove our algorithmic result. We omit its proof.

▶ **Theorem 28.** *The only $\mathbb{H}_3$-subgraph-free C_5-critical graphs are K_3, odd flowers F_n ($n \geq 3$) and exceptional graphs E_1, E_2, E_3. Equivalently, the following statements all hold:*
1. *All $\mathbb{H}_3$-subgraph-free graphs of girth at least 6 are C_5-colourable.*
2. *The only $\mathbb{H}_3$-subgraph-free C_5-critical graphs of girth 5 are E_1, E_2 and odd C_5-flowers F_n.*
3. *The only $\mathbb{H}_3$-subgraph-free C_5-critical graph of girth 4 is E_3.*

▶ **Theorem 29.** *There exists a polynomial-time certifying algorithm for C_5-Colouring on $\mathbb{H}_3$-subgraph-free graphs.*

Proof. As every graph that does not map to C_5 must contain a C_5-critical subgraph, it suffices, due to Theorem 28, to detect the non-existence of the graphs K_3, E_1, E_2, E_3 and F_n (odd $n \geq 3$) in a $\mathbb{H}_3$-subgraph-free graph G. For the graphs K_3, E_1, E_2 and E_3 we can simply use brute force. To detect an odd C_5-flower F_n in polynomial time, we observe that for a fixed centre vertex, v_0 we can make an auxiliary graph on its neighbours putting an edge between two if there is a path on three edges between them in G. Now, G contains an odd C_5-flower with centre v_0 if and only if this auxiliary graph has an odd cycle. We can check this in polynomial time for each v_0, so can find an odd C_5-flower in G polynomial time. ◀

8 Conclusions

We took four classic problems, Hamilton Cycle, k-Induced Disjoint Paths, C_5-Colouring and Star 3-Colouring, that are "easy" on bounded treewidth, but for which we showed that their hardness on subcubic graphs is not preserved under edge subdivision. We gave polynomial and NP-completeness results for $\mathcal{H}$-subgraph-free graphs when $\mathcal{H}$ is some subset of $\{\mathbb{H}_1, \mathbb{H}_2, \ldots\}$, but we need to better understand the case $\mathcal{H} = \{\mathbb{H}_i\}$ ($i \geq 1$).

▶ **Open Problem 1.** *Is there a graph* $\mathbb{H}_\ell$ *such that* HAMILTON CYCLE *is* NP*-complete for* $\mathbb{H}_\ell$*-subgraph-free graphs?*

We note that the case $\mathbb{H}_3$ is the only missing case for obtaining a dichotomy for k-INDUCED DISJOINT PATHS on $\mathbb{H}_i$-subgraph-free graphs,

▶ **Open Problem 2.** *What is the complexity of* k-INDUCED DISJOINT PATHS *for* $\mathbb{H}_3$*-subgraph-free graphs?*

If C_5-COLOURING on $\mathbb{H}_i$-subgraph-free graphs is polynomial-time solvable when $i = 0 \bmod 3$, then we would get a dichotomy for C_5-COLOURING on $\mathbb{H}_i$-subgraph-free graphs based on $i \bmod 3$.

▶ **Open Problem 3.** *What is the complexity of* C_5-COLOURING *for* $\mathbb{H}_i$*-subgraph-free graphs, when* $i = 0 \bmod 3$*?*

If STAR 3-COLOURING on $\mathbb{H}_{2i}$-subgraph-free graphs is polynomial-time solvable for $i \geq 1$, then we would get a dichotomy for STAR 3-COLOURING on $\mathbb{H}_i$-subgraph-free graphs based on $i \bmod 2$.

▶ **Open Problem 4.** *What is the complexity of* STAR 3-COLOURING *for* $\mathbb{H}_{2i}$*-subgraph-free graphs for* $i \geq 1$*?*

Moreover, even though STAR k-COLOURING is not C2 for $k \geq 10$ (Proposition 14), this is not known for $4 \leq k \leq 9$. In particular, Shalu and Antony asked about the case $k = 4$ in [31], and we recall their open problem.

▶ **Open Problem 5.** *What is the complexity of* STAR 4-COLOURING *for subcubic graphs?*

We also still need to determine whether the C12-problems k-INDUCED DISJOINT PATHS and C_5-COLOURING are even C12' just like HAMILTON CYCLE and STAR 3-COLOURING. In order to know this, we must solve the following two problems.

▶ **Open Problem 6.** *What is the complexity of* k-INDUCED DISJOINT PATHS *for subcubic graphs of girth* g *for* $g \geq 3$*?*

▶ **Open Problem 7.** *What is the complexity of* C_5-COLOURING *for subcubic graphs of girth* g *for* $g \geq 3$*?*

We also do not know the complexity of k-INDUCED DISJOINT PATHS, for $k \geq 2$, on graphs of girth at least g with an additional degree bound, whereas the best degree bound for C_5-COLOURING is $6 \cdot 5^{13}$. Namely, for every $g \geq 3$, C_5-COLOURING is NP-complete for graphs with girth at least g and with maximum degree at most $6 \cdot 5^{13}$ (Theorem 6).

Finally, there exist other problems that are NP-complete for bipartite graphs in which one partition class has maximum degree 2 and thus on $(\mathbb{H}_1, \mathbb{H}_3, \ldots)$-subgraph-free graphs. One example of such a problem is MATCHING CUT [27]. Another example is ACYCLIC 3-COLOURING, for which we can show the same results as for STAR 3-COLOURING in Theorem 5 by using the same arguments. However, in contrast to STAR 3-COLOURING, we do not know if ACYCLIC 3-COLOURING satisfies C2 and we recall the following open problem from Shalu and Antony [32].

▶ **Open Problem 8.** *What is the complexity of* ACYCLIC 3-COLOURING *for subcubic graphs?*

References

1 Michael O. Albertson, Glenn G. Chappell, Henry A. Kierstead, André Kündgen, and Radhika Ramamurthi. Coloring with no 2-colored P_4's. *Electronic Journal of Combinatorics*, 11, 2004.

2 Vladimir E. Alekseev, Rodica Boliac, Dmitry V. Korobitsyn, and Vadim V. Lozin. NP-hard graph problems and boundary classes of graphs. *Theoretical Computer Science*, 389:219–236, 2007. doi:10.1016/J.TCS.2007.09.013.

3 Vladimir E. Alekseev and Dmitry V. Korobitsyn. Complexity of some problems on hereditary graph classes. *Diskretnaya Matematika*, 4:34–40, 1992.

4 Stefan Arnborg and Andrzej Proskurowski. Linear time algorithms for NP-hard problems restricted to partial k-trees. *Discrete Applied Mathematics*, 23:11–24, 1989. doi:10.1016/0166-218X(89)90031-0.

5 Daniel Bienstock, Neil Robertson, Paul D. Seymour, and Robin Thomas. Quickly excluding a forest. *Journal of Combinatoral Theory, Series B*, 52:274–283, 1991. doi:10.1016/0095-8956(91)90068-U.

6 Hans L. Bodlaender, Tesshu Hanaka, Yasuaki Kobayashi, Yusuke Kobayashi, Yoshio Okamoto, Yota Otachi, and Tom C. van der Zanden. Subgraph isomorphism on graph classes that exclude a substructure. *Algorithmica*, 82:3566–3587, 2020. doi:10.1007/S00453-020-00737-Z.

7 Hans L. Bodlaender, Matthew Johnson, Barnaby Martin, Jelle J. Oostveen, Sukanya Pandey, Daniël Paulusma, Siani Smith, and Erik Jan van Leeuwen. Complexity framework for forbidden subgraphs IV: The Steiner Forest problem. *Proc. IWOCA 2024, LNCS*, 14764:206–217, 2024. doi:10.1007/978-3-031-63021-7_16.

8 Rodica Boliac and Vadim V. Lozin. On the clique-width of graphs in hereditary classes. *Proc. ISAAC 2022, LNCS*, 2518:44–54, 2002. doi:10.1007/3-540-36136-7_5.

9 Maria Chudnovsky, Shenwei Huang, Pawe ł Rzążewsk, Sophie Spirkl, and Mingxian Zhong. Complexity of C_k-coloring in hereditary classes of graphs. *Proc. ESA 2019, LIPIcs*, 144:31:1–31:15, 2019. doi:10.4230/LIPICS.ESA.2019.31.

10 Bruno Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Information and Computation*, 85:12–75, 1990. doi:10.1016/0890-5401(90)90043-H.

11 David P. Dailey. Uniqueness of colorability and colorability of planar 4-regular graphs are NP-complete. *Discrete Mathematics*, 30:289–293, 1980. doi:10.1016/0012-365X(80)90236-8.

12 Rina Dechter and Judea Pearl. Tree clustering for constraint networks. *Artificial Intelligence*, 38:353–366, 1989. doi:10.1016/0004-3702(89)90037-4.

13 Tala Eagling-Vose, Barnaby Martin, Daniël Paulusma, and Siani Smith. Graph homomorphism, monotone classes and bounded pathwidth. *Proc. CiE 2024, LNCS*, 14773:233–251, 2024. doi:10.1007/978-3-031-64309-5_19.

14 Thomas Emden-Weinert, Stefan Hougardy, and Bernd Kreuter. Uniquely colourable graphs and the hardness of colouring graphs of large girth. *Combinatorics, Probability & Computing*, 7:375–386, 1998. URL: http://journals.cambridge.org/action/displayAbstract?aid=46667.

15 Anna Galluccio, Pavol Hell, and Jaroslav Nešetřil. The complexity of H-colouring of bounded degree graphs. *Discrete Mathematics*, 222:101–109, 2000. doi:10.1016/S0012-365X(00)00009-1.

16 Alfred Geroldinger and Imre Z. Ruzsa. *Combinatorial Number Theory and Additive Group Theory*. Birkhäuser, 2009.

17 Petr A. Golovach and Daniël Paulusma. List coloring in the absence of two subgraphs. *Discrete Applied Mathematics*, 166:123–130, 2014. doi:10.1016/J.DAM.2013.10.010.

18 Petr A. Golovach, Daniël Paulusma, and Bernard Ries. Coloring graphs characterized by a forbidden subgraph. *Discrete Applied Mathematics*, 180:101–110, 2015. doi:10.1016/J.DAM.2014.08.008.

19 Pavol Hell and Jaroslav Nesetril. On the complexity of H-coloring. *Journal of Combinatorial Theory, Series B*, 48:92–110, 1990. doi:10.1016/0095-8956(90)90132-J.

20 Matthew Johnson, Barnaby Martin, Jelle J. Oostveen, Sukanya Pandey, Daniël Paulusma, Siani Smith, and Erik Jan van Leeuwen. Complexity framework for forbidden subgraphs I: The framework. *CoRR*, 2211.12887, 2022.

21 Matthew Johnson, Barnaby Martin, Sukanya Pandey, Daniël Paulusma, Siani Smith, and Erik Jan van Leeuwen. Complexity framework for forbidden subgraphs III: When problems are tractable on subcubic graphs. *Proc. MFCS 2024, LIPIcs*, 272:57:1–57:15, 2023. doi:10.4230/LIPICS.MFCS.2023.57.

22 Marcin Kamiński. Max-Cut and containment relations in graphs. *Theoretical Computer Science*, 438:89–95, 2012. doi:10.1016/J.TCS.2012.02.036.

23 Nicholas Korpelainen, Vadim V. Lozin, Dmitriy S. Malyshev, and Alexander Tiskin. Boundary properties of graphs for algorithmic graph problems. *Theoretical Computer Science*, 412:3545–3554, 2011. doi:10.1016/J.TCS.2011.03.001.

24 Benjamin Lévêque, David Y. Lin, Frédéric Maffray, and Nicolas Trotignon. Detecting induced subgraphs. *Discrete Applied Mathematics*, 157:3540–3551, 2009. doi:10.1016/J.DAM.2009.02.015.

25 Vadim V. Lozin. The hamiltonian cycle problem and monotone classes. *Proceedings of IWOCA 2024, LNCS*, 14764:460–471, 2024. doi:10.1007/978-3-031-63021-7_35.

26 Vadim V. Lozin and Igor Razgon. Tree-width dichotomy. *European Journal of Combinatorics*, 103:103517, 2022. doi:10.1016/J.EJC.2022.103517.

27 Augustine M. Moshi. Matching cutsets in graphs. *Journal of Graph Theory*, 13:527–536, 1989. doi:10.1002/JGT.3190130502.

28 Neil Robertson and Paul D. Seymour. Graph minors. III. Planar tree-width. *Journal of Combinatorial Theory, Series B*, 36:49–64, 1984. doi:10.1016/0095-8956(84)90013-3.

29 Neil Robertson and Paul D. Seymour. Graph minors. V. Excluding a planar graph. *Journal of Combinatorial Theory, Series B*, 41:92–114, 1986. doi:10.1016/0095-8956(86)90030-4.

30 Neil Robertson and Paul D. Seymour. Graph minors. XIII. The Disjoint Paths problem. *Journal of Combinatorial Theory, Series B*, 63:65–110, 1995. doi:10.1006/JCTB.1995.1006.

31 M. A. Shalu and Cyriac Antony. Star colouring of bounded degree graphs and regular graphs. *Discrete Mathematics*, 345:112850, 2022. doi:10.1016/J.DISC.2022.112850.

32 M. A. Shalu and Cyriac Antony. Hardness transitions and uniqueness of acyclic colouring. *Discrete Applied Mathematics*, 345:77–98, 2024. doi:10.1016/J.DAM.2023.11.030.

33 Yossi Shiloach. A polynomial solution to the undirected Two Paths problem. *Journal of the ACM*, 27:445–456, 1980. doi:10.1145/322203.322207.