Finding Large H-Colorable Subgraphs in Hereditary Graph Classes

Maria Chudnovsky
Princeton University, NJ, USA
mchudnov@math.princeton.edu

Jason King
Princeton University, NJ, USA
jtking@princeton.edu

Michał Pilipczuk
Institute of Informatics, University of Warsaw, Poland
michal.pilipczuk@mimuw.edu.pl

Paweł Rzążewski 🄿
Warsaw University of Technology, Faculty of Mathematics and Information Science, Poland
University of Warsaw, Institute of Informatics, Poland
p.rzazewski@mini.pw.edu.pl

Sophie Spirkl
Princeton University, NJ, USA
sspirkl@math.princeton.edu

─── **Abstract** ───

We study the MAX PARTIAL H-COLORING problem: given a graph G, find the largest induced subgraph of G that admits a homomorphism into H, where H is a fixed pattern graph without loops. Note that when H is a complete graph on k vertices, the problem reduces to finding the largest induced k-colorable subgraph, which for $k = 2$ is equivalent (by complementation) to ODD CYCLE TRANSVERSAL.

We prove that for every fixed pattern graph H without loops, MAX PARTIAL H-COLORING can be solved:

- in $\{P_5, F\}$-free graphs in polynomial time, whenever F is a threshold graph;
- in $\{P_5, \text{bull}\}$-free graphs in polynomial time;
- in P_5-free graphs in time $n^{\mathcal{O}(\omega(G))}$;
- in $\{P_6, 1\text{-subdivided claw}\}$-free graphs in time $n^{\mathcal{O}(\omega(G)^3)}$.

Here, n is the number of vertices of the input graph G and $\omega(G)$ is the maximum size of a clique in G. Furthermore, by combining the mentioned algorithms for P_5-free and for $\{P_6, 1\text{-subdivided claw}\}$-free graphs with a simple branching procedure, we obtain subexponential-time algorithms for MAX PARTIAL H-COLORING in these classes of graphs.

Finally, we show that even a restricted variant of MAX PARTIAL H-COLORING is NP-hard in the considered subclasses of P_5-free graphs, if we allow loops on H.

1 Introduction

Many computational graph problems that are (NP-)hard in general become tractable in restricted classes of input graphs. In this work we are interested in *hereditary* graph classes, or equivalently classes defined by forbidding induced subgraphs. For a set of graphs $\mathcal{F}$, we say that a graph G is *$\mathcal{F}$-free* if G does not contain any induced subgraph isomorphic to a graph from $\mathcal{F}$. By forbidding different sets $\mathcal{F}$ we obtain graph classes with various structural properties, which can be used in the algorithmic context. This highlights an interesting interplay between structural graph theory and algorithm design.

Perhaps the best known example of this paradigm is the case of the MAXIMUM INDEPENDENT SET problem: given a graph G, find the largest set of pairwise non-adjacent vertices in G. It is known that the problem is NP-hard on F-free graphs unless F is a forest whose every component is a path or a subdivided claw [2]; here, a *claw* is a star with 3 leaves. However, the remaining cases, when F is a *subdivided claw forest*, remain largely unexplored despite significant effort. Polynomial-time algorithms have been given for P_5-free graphs [24], P_6-free graphs [20], claw-free graphs [26, 29], and fork-free graphs [3, 25]. While the complexity status in all the other cases remains open, it has been observed that relaxing the goal of polynomial-time solvability leads to positive results in a larger generality. For instance, for every $t \in \mathbb{N}$, MAXIMUM INDEPENDENT SET can be solved in time $2^{\mathcal{O}(\sqrt{tn \log n})}$ in P_t-free graphs [4]. The existence of such a *subexponential-time algorithm* for F-free graphs is excluded under the Exponential Time Hypothesis whenever F is not a subdivided claw forest (see e.g. the discussion in [27]), which shows a qualitative difference between the negative and the potentially positive cases. Also, Chudnovsky et al. [10] recently gave a quasi-polynomial-time approximation scheme (QPTAS) for MAXIMUM INDEPENDENT SET in F-free graphs, for every fixed subdivided claw forest F.

The abovementioned positive results use a variety of structural techniques related to the considered graph classes, for instance: the concept of *Gyárfás path* that gives useful separators in P_t-free graphs [4, 6, 10], the dynamic programming approach based on potential maximal cliques [24, 20], or structural properties of claw-free and fork-free graphs that relate them to line graphs [25, 26, 29]. Some of these techniques can be used to give algorithms for related problems, which can be expressed as looking for the largest (in terms of the number of vertices) induced subgraph satisfying a fixed property. For MAXIMUM INDEPENDENT SET this property is being edgeless, but for instance the property of being acyclic corresponds to the MAXIMUM INDUCED FOREST problem, which by complementation is equivalent to FEEDBACK VERTEX SET. Work in this direction so far focused on properties that imply bounded treewidth [1, 17] or, more generally, that imply sparsity [27].

A different class of problems that admits an interesting complexity landscape on hereditary graphs classes are coloring problems. For fixed $k \in \mathbb{N}$, the *k-COLORING* problem asks whether the input graph admits a proper coloring with k colors. For every $k \geqslant 3$, the problem is NP-hard on F-free graphs unless F is a forest of paths (a *linear forest*) [18]. The classification of the remaining cases is more advanced than in the case of MAXIMUM INDEPENDENT SET, but not yet complete. On one hand, Hoàng et al. [22] showed that for every fixed k, k-COLORING

is polynomial-time solvable on P_5-free graphs. On the other hand, the problem becomes NP-hard already on P_6-free graphs for all $k \geqslant 5$ [23]. The cases $k = 3$ and $k = 4$ turn out to be very interesting. 4-Coloring is polynomial-time solvable on P_6-free graphs [14] and NP-hard in P_7-free graphs [23]. While there is a polynomial-time algorithm for 3-Coloring in P_7-free graphs [5], the complexity status in P_t-free graphs for $t \geqslant 8$ remains open. However, relaxing the goal again leads to positive results in a wider generality: for every $t \in \mathbb{N}$, there is a subexponential-time algorithm with running time $2^{\mathcal{O}(\sqrt{tn \log n})}$ for 3-Coloring in P_t-free graphs [19], and there is also a polynomial-time algorithm that given a 3-colorable P_t-free graph outputs its proper coloring with $\mathcal{O}(t)$ colors [12].

We are interested in using the toolbox developed for coloring problems in P_t-free graphs to the setting of finding maximum induced subgraphs with certain properties. Specifically, consider the following Maximum Induced k-Colorable Subgraph problem: given a graph G, find the largest induced subgraph of G that admits a proper coloring with k colors. While this problem clearly generalizes k-Coloring, for $k = 1$ it boils down to Maximum Independent Set. For $k = 2$ it can be expressed as Maximum Induced Bipartite Subgraph, which by complementation is equivalent to the well-studied Odd Cycle Transversal problem: find the smallest subset of vertices that intersects all odd cycles in a given graph. While polynomial-time solvability of Odd Cycle Transversal on P_4-free graphs (also known as *cographs*) follows from the fact that these graphs have bounded cliquewidth (see [15]), it is known that the problem is NP-hard in P_6-free graphs [16]. The complexity status of Odd Cycle Transversal in P_5-free graphs remains open [9, Problem 4.4]: resolving this question was the original motivation of our work.

Our contribution. Following the work of Groenland et al. [19], we work with a very general form of coloring problems, defined through homomorphisms. For graphs G and H, a *homomorphism* from G to H, or an H-*coloring* of G, is a function $\phi \colon V(G) \to V(H)$ such that for every edge uv in G, we have $\phi(u)\phi(v) \in E(H)$. We study the Max Partial H-Coloring problem defined as follows: given a graph G, find the largest induced subgraph of G that admits an H-coloring. Note that if H is the complete graph on k vertices, then an H-coloring is simply a proper coloring with k colors, hence this formulation generalizes the Maximum Induced k-Colorable Subgraph problem. Unless stated explicitly, we will always assume that the pattern graph H does not have loops, hence an H-coloring is a proper coloring with $|V(H)|$ colors.
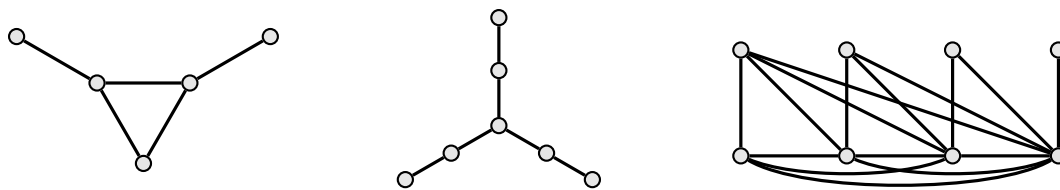
Figure 1 A bull, a 1-subdivided claw, and an example threshold graph.

Fix a pattern graph H without loops. We prove that Max Partial H-Coloring can be solved:

(R1) in $\{P_5, F\}$-free graphs in polynomial time, whenever F is a threshold graph;

(R2) in $\{P_5, \text{bull}\}$-free graphs in polynomial time;

(R3) in P_5-free graphs in time $n^{\mathcal{O}(\omega(G))}$; and

(R4) in $\{P_6, 1\text{-subdivided claw}\}$-free graphs in time $n^{\mathcal{O}(\omega(G)^3)}$.

Here, n is the number of vertices of the input graph G and $\omega(G)$ is the size of the maximum clique in G. Also, recall that a graph G is a *threshold graph* if $V(G)$ can be partitioned into an independent set A and a clique B such that for each $a, a' \in A$, we have either $N(a) \supseteq N(a')$ or $N(a) \subseteq N(a')$. There is also a characterization via forbidden induced subgraphs: threshold graphs are exactly $\{2P_2, C_4, P_4\}$-free graphs, where $2P_2$ is an induced matching of size 2. Figure 1 depicts a bull, a 1-subdivided claw, and an example threshold graph.

Further, we observe that by employing a simple branching strategy, an $n^{\mathcal{O}(\omega(G)^\alpha)}$-time algorithm for MAX PARTIAL H-COLORING in $\mathcal{F}$-free graphs can be used to give also a subexponential-time algorithm in this setting, with running time $n^{\mathcal{O}(n^{\alpha/(\alpha+1)})}$. Thus, results (R3) and (R4) imply that for every fixed irreflexive H, the MAX PARTIAL H-COLORING problem can be solved in time $n^{\mathcal{O}(\sqrt{n})}$ in P_5-free graphs and in time $n^{\mathcal{O}(n^{3/4})}$ in $\{P_6, 1\text{-subdivided claw}\}$-free graphs. This in particular applies to the ODD CYCLE TRANSVERSAL problem. We note here that Dabrowski et al. [16] proved that ODD CYCLE TRANSVERSAL in $\{P_6, K_4\}$-free graphs is NP-hard and does not admit a subexponential-time algorithm under the Exponential Time Hypothesis. Thus, it is unlikely that any of our algorithmic results – the $n^{\mathcal{O}(\omega(G))}$-time algorithm and the $n^{\mathcal{O}(\sqrt{n})}$-time algorithm – can be extended from P_5-free graphs to P_6-free graphs.

All our algorithms work in a weighted setting, where instead of just maximizing the size of the domain of an H-coloring, we maximize its total *revenue*, where for each pair $(u, v) \in V(G) \times V(H)$ we have a prescribed revenue yielded by sending u to v. This setting allows encoding a broader range of coloring problems. For instance, list variants can be expressed by giving negative revenues for forbidden assignments (see e.g. [21, 28]). Also, our algorithms work in a slightly larger generality than stated above, see Section 5 for precise statements.

Finally, we investigate the possibility of extending our algorithmic results to pattern graphs with possible loops. We show an example of a graph H with loops, for which MAX PARTIAL H-COLORING is NP-hard and admits no subexponential-time algorithm under the ETH even in very restricted subclasses of P_5-free graphs, including $\{P_5, \text{bull}\}$-free graphs. This shows that whether the pattern graph is allowed to have loops has a major impact on the complexity of the problem.

Full version. In this extended abstract we focus on proving results (R3) and (R4). Results (R1) and (R2), as well as of the abovementioned lower bound, are proved in the full version of the paper, which is available on arXiv [8]. Also, the main branching step is given here in a simplified form that is sufficient for results (R3) and (R4), but not for results (R1) and (R2).

Our techniques. The key element of our approach is a branching procedure that, given an instance (G, rev) of MAX PARTIAL H-COLORING, where rev is the revenue function, produces a relatively small set of instances Π such that solving (G, rev) reduces to solving all the instances in Π. Moreover, every instance $(G', \mathsf{rev}') \in \Pi$ is simpler in the following sense: either it is an instance of MAX PARTIAL H'-COLORING for H' being a proper induced subgraph of H (hence it can be solved by induction on $|V(H)|$), or for any connected graph F on at least two vertices, G' is F-free provided we assume that G is $F^{\bullet\!-\!\circ}$-free. Here, $F^{\bullet\!-\!\circ}$ is the graph obtained from F by adding a universal vertex y and a degree-1 vertex x adjacent only to y. In particular we have $\omega(G') < \omega(G)$, so applying the branching procedure exhaustively in a recursion scheme yields a recursion tree of depth bounded by $\omega(G)$. Now, for results (R3) and (R4) we respectively have $|\Pi| \leqslant n^{\mathcal{O}(1)}$ and $|\Pi| \leqslant n^{\mathcal{O}(\omega(G)^2)}$, giving bounds of $n^{\mathcal{O}(\omega(G))}$ and $n^{\mathcal{O}(\omega(G)^3)}$ on the total size of the recursion tree and on the overall time complexity.

For result (R1) we apply the branching procedure not exhaustively, but a constant number of times: if the original graph G is $\{P_5, F\}$-free for some threshold graph F, it suffices to apply the branching procedure $\mathcal{O}(|V(F)|)$ times to reduce the original instances to a set of edgeless instances, which can be solved trivially. As $\mathcal{O}(|V(F)|) = \mathcal{O}(1)$, this gives recursion tree of polynomial size, and hence a polynomial-time complexity due to always having $|\Pi| \leqslant n^{\mathcal{O}(1)}$ in this setting. For result (R2), we show that two applications of the branching procedure reduce the input instance to a polynomial number of instances that are P_4-free, which can be solved in polynomial time due to P_4-free graphs (also known as *cographs*) having cliquewidth at most 2. However, these applications are interleaved with a reduction to the case of *prime graphs* – graphs with no non-trivial modules – which we achieve using dynamic programming on the modular decomposition of the input graph. This is in order to apply some results on the structure of prime bull-free graphs [11, 13], so that P_4-freeness is achieved at the end.

Let us briefly discuss the key branching procedure. The first step is finding a useful dominating structure that we call a *monitor*: a subset of vertices M of a connected graph G is a monitor if for every connected component C of $G - M$, there is a vertex in M that is complete to C. We prove that in a connected P_6-free graph there is always a monitor that is the closed neighborhood of a set of at most three vertices. After finding such a monitor $N[X]$ for $|X| \leqslant 3$, we perform a structural analysis of the graph centered around the set X. This analysis shows that there exists a subset of $\mathcal{O}(|V(H)|)$ vertices such that after guessing this subset and the H-coloring on it, the instance can be partitioned into several separate subinstances, each of which has a strictly smaller clique number. This structural analysis, and in particular the way the separation of subinstances is achieved, is inspired by the polynomial-time algorithm of Hoàng et al. [22] for k-COLORING in P_5-free graphs.

Other related work.   We remark that very recently and independently of us, Brettell et al. [7] proved that for every fixed $s, t \in \mathbb{N}$, the class of $\{K_t, sK_1 + P_5\}$-free graphs has bounded mim-width. Here, *mim-width* is a graph parameter that is less restrictive than cliquewidth, but the important aspect is that a wide range of vertex-partitioning problems, including the MAX PARTIAL H-COLORING problem considered in this work, can be solved in polynomial time on every class of graphs where the mim-width is universally bounded and a corresponding decomposition can be computed efficiently. The result of Brettell et al. thus shows that in P_5-free graphs, the mim-width is bounded by a function of the clique number. This gives an $n^{f(\omega(G))}$-time algorithm for MAX PARTIAL H-COLORING in P_5-free graphs (for fixed H), for some function f. However, the proof presented in [7] gives only an exponential upper bound on the function f, which in particular does not imply the existence of a subexponential-time algorithm. To compare, our reasoning leads to an $n^{\mathcal{O}(\omega(G))}$-time algorithm and a subexponential-time algorithm with complexity $n^{\mathcal{O}(\sqrt{n})}$.

We remark that the techniques used by Brettell et al. [7] also rely on revisiting the approach of Hoàng et al. [22], and they similarly observe that this approach can be used to apply induction based on the clique number of the graph.

2   Preliminaries

Graphs.   For a graph G, the vertex and edge sets of G are denoted by $V(G)$ and $E(G)$, respectively. The *open neighborhood* of a vertex u is the set $N_G(u) \coloneqq \{v \colon uv \in E(G)\}$, while the *closed neighborhood* is $N_G[u] \coloneqq N_G(u) \cup \{u\}$. This notation is extended to sets of vertices: for $X \subseteq V(G)$, we set $N_G[X] \coloneqq \bigcup_{u \in X} N_G[u]$ and $N_G(X) \coloneqq N_G[X] - X$. We may omit the subscript if the graph G is clear from the context. By C_t, P_t, and K_t we respectively denote the cycle, the path, and the complete graph on t vertices.

The *clique number* $\omega(G)$ is the size of the largest clique in a graph G. A clique K in G is *maximal* if no proper superset of K is a clique.

For $s, t \in \mathbb{N}$, the *Ramsey number* of s and t is the smallest integer k such that every graph on k vertices contains either a clique of size s or an independent set of size t. It is well-known that the Ramsey number of s and t is bounded from above by $\binom{s+t-2}{s-1}$, hence we will denote $\mathsf{Ramsey}(s, t) := \binom{s+t-2}{s-1}$.

For a graph G and $A \subseteq V(G)$, by $G[A]$ we denote the subgraph of G induced by A. We write $G - A := G[V(G) - A]$. We say that F is an *induced subgraph* of G if there is $A \subseteq V(G)$ such that $G[A]$ is isomorphic to F; this containment is *proper* if in addition $A \neq V(G)$. For a family of graphs $\mathcal{F}$, a graph G is $\mathcal{F}$-*free* if G does not contain any induced subgraph from $\mathcal{F}$. If $\mathcal{F} = \{H\}$, then we may speak about H-*free* graphs as well.

If G is a graph and $A \subseteq V(G)$ is a subset of vertices, then a vertex $u \notin A$ is *complete* to A if u is adjacent to all the vertices of A, and u is *anti-complete* to A if u has no neighbors in A. We will use the following simple claim several times.

▶ **Lemma 1.** *Suppose G is a graph, A is a subset of its vertices such that $G[A]$ is connected, and $u \notin A$ is a vertex that is neither complete nor anti-complete to A in G. Then there are vertices $a, b \in A$ such that $u - a - b$ is an induced P_3 in G.*

Proof. Since u is neither complete nor anticomplete to A, both the sets $A \cap N(u)$ and $A - N(u)$ are non-empty. As A is connected, there exist $a \in A \cap N(u)$ and $b \in A - N(u)$ such that a and b are adjacent. Now $u - a - b$ is the desired induced P_3. ◀

For a graph F, by $F^\bullet$ we denote the graph obtained from F by adding a *universal vertex*: a vertex adjacent to all the other vertices. Similarly, by $F^{\bullet\circ}$ we denote the graph obtained from F by adding first an isolated vertex, say x, and then a universal vertex, say y. Note that thus y is adjacent to all the other vertices of $F^{\bullet\circ}$, while x is adjacent only to y.

H-colorings.   For graphs H and G, a function $\phi \colon V(G) \to V(H)$ is a *homomorphism* from G to H if for every $uv \in E(G)$, we also have $\phi(u)\phi(v) \in E(H)$. Note that a homomorphism from G to the complete graph K_t is nothing else than a proper coloring of G with t colors. Therefore, a homomorphism from G to H will be also called an H-*coloring* of G, and we will refer to vertices of H as colors. Note that we will always assume that H is a simple graph without loops, so no two adjacent vertices of G can be mapped by a homomorphism to the same vertex of H. To stress this, we will call such H an *irreflexive pattern graph*.

A *partial homomorphism* from G to H, or a *partial H-coloring* of G, is a partial function $\phi \colon V(G) \rightharpoonup V(H)$ that is a homomorphism from $G[\mathsf{dom}\,\phi]$ to H, where $\mathsf{dom}\,\phi$ denotes the domain of ϕ.

Suppose that with graphs G and H we associate a *revenue function* $\mathsf{rev} \colon V(G) \times V(H) \to \mathbb{R}$. Then the *revenue* of a partial H-coloring ϕ is defined as

$$\mathsf{rev}(\phi) := \sum_{u \in \mathsf{dom}\,\phi} \mathsf{rev}(u, \phi(u)).$$

In other words, for $u \in V(G)$ and $v \in V(H)$, $\mathsf{rev}(u, v)$ denotes the revenue yielded by assigning $\phi(u) := v$.

We now define the main problem studied in this work. In the following, we consider the graph H fixed.

MAX PARTIAL H-COLORING

Input: Graph G and a revenue function $\mathsf{rev} \colon V(G) \times V(H) \to \mathbb{R}$
Output: A partial H-coloring ϕ of G that maximizes $\mathsf{rev}(\phi)$

An *instance* of the MAX PARTIAL H-COLORING problem is a pair (G, rev) as above. A *solution* to an instance (G, rev) is a partial H-coloring of G, and it is *optimum* if it maximizes $\mathsf{rev}(\phi)$ among solutions. By $\mathrm{OPT}(G, \mathsf{rev})$ we denote the maximum possible revenue of a solution to the instance (G, rev).

Let us note one aspect that will be used later on. Observe that in revenue functions we allow negative revenues for some assignments. However, if we are interested in maximizing the total revenue, there is no point in using such assignments: if $u \in \mathsf{dom}\, \phi$ and $\mathsf{rev}(u, \phi(u)) < 0$, then just removing u from the domain of ϕ increases the revenue. Thus, optimal solutions never use assignments with negative revenues. Note that this feature can be used to model list versions of partial coloring problems, where each vertex $v \in V(G)$ is assigned a list of colors $L(v) \subseteq V(H)$ and can only be mapped to a vertex from $L(v)$.

3 Monitors in P_6-free graphs

In this section we prove an auxiliary result about finding useful separators in P_6-free graphs. The desired property is expressed in the following definition.

▶ **Definition 2.** *Let G be a connected graph. A subset of vertices $M \subseteq V(G)$ is a* monitor *in G if for every connected component C of $G - M$, there exists a vertex $w \in M$ that is complete to C.*

Let us note the following property of monitors.

▶ **Lemma 3.** *If M is a monitor in a connected graph G, then every maximal clique in G intersects M. In particular, $\omega(G - M) < \omega(G)$.*

Proof. If K is a clique in $G - M$, then K has to be entirely contained in some connected component C of $G - M$. Since M is a monitor, there exists $w \in M$ that is complete to C. Then $K \cup \{w\}$ is also a clique in G, hence K cannot be a maximal clique in G. ◀

We now prove that in P_6-free graphs we can always find easily describable monitors.

▶ **Lemma 4.** *Let G be a connected P_6-free graph. Then for every $u \in V(G)$ there exists a subset of vertices X such that $u \in X$, $|X| \leqslant 3$, $G[X]$ is a path whose one endpoint is u, and $N_G[X]$ is a monitor in G.*

Lemma 4 follows immediately from the following statement applied for $t = 6$.

▶ **Lemma 5.** *Let $t \in \{4, 5, 6\}$, G be a connected P_6-free graph, and $u \in V(G)$ be a vertex such that in G there is no induced P_t with u being one of the endpoints. Then there exists a subset X of vertices such that $u \in X$, $|X| \leqslant t - 3$, $G[X]$ is a path whose one endpoint is u, and $N_G[X]$ is a monitor in G.*

Proof. We proceed by induction on t. The base case for $t = 4$ will be proved directly within the analysis.

In the following, by *slabs* we mean connected components of the graph $G - N_G[u]$. We shall say that a vertex $w \in N_G(u)$ is *mixed* on a slab C if w is neither complete nor anti-complete to C. A slab C is *simple* if there exists a vertex $w \in N_G(u)$ that is complete to C, and *difficult* otherwise.

Note that since G is connected, for every difficult slab D there exists some vertex $w \in N_G(u)$ that is mixed on D. Then, by Lemma 1, we can find vertices $a, b \in D$ such that $u - w - a - b$ is an induced P_4 in G. If $t = 4$ then no such induced P_4 can exists, so we infer that in this case there are no difficult slabs. Then $N_G[u]$ is a monitor, so we may set $X := \{u\}$. This proves the claim for $t = 4$; from now on we assume that $t \geqslant 5$.

Let us choose a vertex $v \in N_G(u)$ that maximizes the number of difficult slabs on which v is mixed. Suppose there is a difficult slab D' such that v is anti-complete to D'. As we argued, there exists a vertex $v' \in N_G(u)$ such that v' is mixed on D'; clearly $v' \neq v$. By the choice of v, there exists a difficult slab D such that v is mixed on D and v' is anti-complete to D. By applying Lemma 1 twice, we find vertices $a, b \in D$ and $a', b' \in D'$ such that $v - a - b$ and $v' - a' - b'$ are induced P_3s in G. Now, if v and v' were adjacent, then $b - a - v - v' - a' - b'$ would be an induced P_6 in G, a contradiction. Otherwise $b - a - v - u - v' - a' - b'$ is an induced P_7 in G, again a contradiction (see Figure 2).
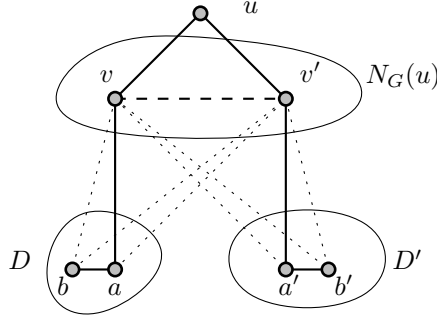
■ **Figure 2** The graph G in the proof of Lemma 5 when v anti-complete to some difficult slab D'. Dotted lines show non-edges. The edge vv' might be present.

We conclude that v is mixed on every difficult slab. Let

$$A := \{v\} \cup \bigcup_{D:\ \text{difficult slab}} V(D).$$

Then $G[A]$ is connected and P_6-free. Moreover, in $G[A]$ there is no P_{t-1} with one endpoint being v, because otherwise we would be able to extend such an induced P_{t-1} using u, and thus obtain an induced P_t in G with one endpoint being u. Consequently, by induction we find a subset $Y \subseteq A$ such that $|Y| \leqslant (t-1) - 3 = t - 4$, $G[Y]$ is a path with one of the endpoints being v, and $N_{G[A]}[Y]$ is a monitor in $G[A]$. Let $X := Y \cup \{u\}$. Then $|X| \leqslant t - 3$ and $G[X]$ is a path with u being one of the endpoints.

We verify that $N_G[X]$ is a monitor in G. Consider any connected component C of $G - N_G[X]$. As $N_G[X] \supseteq N_G[u]$, C is contained in some slab D. If D is simple, then by definition there exists a vertex $w \in N_G[u] \subseteq N_G[X]$ that is complete to D, hence also complete to C. Otherwise D is difficult, hence C is a connected component of $G[A] - N_{G[A]}[Y]$. Since $N_{G[A]}[Y]$ is a monitor in $G[A]$, there exists a vertex $w \in N_{G[A]}[Y] \subseteq N_G[X]$ that is complete to C. This completes the proof.　　◀

We remark that no statement analogous to Lemma 4 can hold for P_7-free graphs, even if from X we only require that $N_G[X]$ intersects all the maximum-size cliques in G (which is implied by the property of being a monitor, see Lemma 3). Consider the following example. Let G be a graph obtained from the union of $n + 1$ complete graphs $K^{(0)}, \ldots, K^{(n)}$, each on n vertices, by making one vertex from each of the graphs $K^{(1)}, \ldots, K^{(n)}$ adjacent to a different vertex of $K^{(0)}$. Then G is P_7-free, but the minimum size of a set $X \subseteq V(G)$ such that $N_G[X]$ intersects all maximum-size cliques in G is n.

4   Branching

We now present the core branching step that is used by all our algorithms. This part is inspired by the approach of Hoàng et al. [22]. We will rely on the following two graph families; see Figure 3. For $t \in \mathbb{N}$, the graph S_t is obtained from the star $K_{1,t}$ by subdividing every edge once. Then $L_1 := P_3$ and for $t \geqslant 2$ the graph L_t is obtained from S_t by making all the leaves of S_t pairwise adjacent.

■ **Figure 3** Graphs S_4 and L_4.

▶ **Lemma 6.** *Let H be a fixed irreflexive pattern graph. Suppose we are given integers s, t and an instance (G, rev) of* MAX PARTIAL H-COLORING *such that G is connected and $\{P_6, L_s, S_t\}$-free. Denoting $n := |V(G)|$, one can in time $n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$ construct a subgraph G' of G with $V(G') = V(G)$ and a set Π consisting of at most $n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$ revenue functions with domain $V(G) \times V(H)$ such that the following conditions hold:*

(C1) *The graph G' is $\{P_6, L_s, S_t\}$-free. Moreover, if G is $F^\bullet$-free for some connected graph F on at least two vertices, then G' is F-free.*

(C2) *We have $\mathrm{OPT}(G, \mathsf{rev}) = \max_{\mathsf{rev}' \in \Pi} \mathrm{OPT}(G', \mathsf{rev}')$. Moreover, for any $\mathsf{rev}' \in \Pi$ for which the maximum is reached, every optimum solution ϕ to (G', rev') is also an optimum solution to (G, rev) with $\mathsf{rev}(\phi) = \mathsf{rev}'(\phi)$.*

We remark that the statement above is a simplified variant of the lemma, and it is sufficient for proving results (R3) and (R4), but not for results (R1) and (R2). In the full variant, presented in the full version of the paper, solving the instance (G, rev) is reduced to solving a list Π of *pairs* of instances. Each pair $((G_1, \mathsf{rev}_1), (G_2, \mathsf{rev}_2)) \in \Pi$ satisfies the following: (G_1, rev_1) is an instance of MAX PARTIAL H'-COLORING for some proper induced subgraph H' of G; and if G_2 contains some induced connected graph F on at least two vertices, then G contains not only an induced $F^\bullet$, but even an induced $F^{\bullet-\circ}$. This gives a stronger reduction of structure upon application of Lemma 6, which is vitally used in the proofs of results (R1) and (R2).

The remainder of this section is devoted to the proof of Lemma 6. We fix the irreflexive pattern graph H and consider an input instance (G, rev). We find it more didactic to first perform an analysis of (G, rev), and only provide the algorithm at the end. Thus, the correctness will be clear from the previous observations.

Since G is connected, by Lemma 4 there exists $X \subseteq V(G)$ such that $|X| \leqslant 3$ and $N[X]$ is a monitor in G. Note that such a set X can be found in polynomial time by checking all subsets of $V(G)$ of size at most 3. In case $|X| < 3$, we may add arbitrary vertices to X so that $|X| = 3$, note that the property of being a monitor still holds. Let us arbitrarily enumerate the vertices of X as $\{x_1, x_2, x_3\}$.

We partition $V(G) - X$ into A_1, A_2, A_3, A_4 as follows (see Figure 4):

$$A_1 := N(x_1) - X, \ A_2 := N(x_2) - (X \cup A_1), \ A_3 := N(x_3) - (X \cup A_1 \cup A_2), \ A_4 := V(G) - N[X].$$

Note that $\{A_1, A_2, A_3\}$ is a partition of $N(X)$. For $i \in \{1, 2, 3\}$, denote $A_{>i} := \bigcup_{j=i+1}^{4} A_j$ and observe that x_i is complete to A_i and anti-complete to $A_{>i}$. Moreover, we have the following.

$\triangleright$ **Claim 7.** For every connected graph F and $i \in \{1, 2, 3, 4\}$, if $G[A_i]$ contains an induced F, then G contains an induced $F^\bullet$.

Proof. Suppose $B \subseteq A_i$ induces F in G. If $i \in \{1, 2, 3\}$ then $B \cup \{x_i\}$ induces $F^\bullet$ in G, hence assume that $i = 4$. Since F is connected, B is entirely contained in one connected component C of $G[A_4]$. As $N[X]$ is a monitor in G, there exists a vertex $w \in N[X]$ that is complete to C. Now $B \cup \{w\}$ induces $F^\bullet$ in G. $\triangleleft$
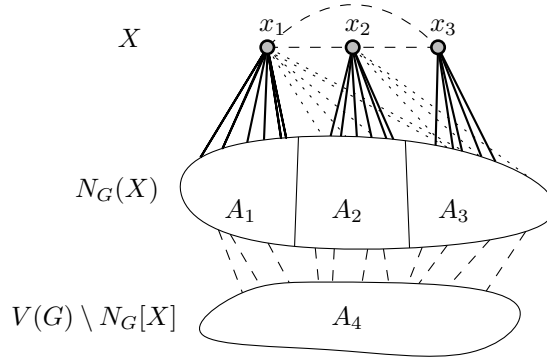
Figure 4 The partition on $V(G)$ in the proof of Lemma 6. Solid and dotted lines respectively indicate that a vertex is complete or anti-complete to a set. Dashed edges might, but do not have to exist.

The next claim contains the core combinatorial observation of the proof.

$\triangleright$ **Claim 8.** Let ϕ be a solution to the instance (G, rev). Then for every $i \in \{1, 2, 3\}$ and $v \in V(H)$, there exists a set $S \subseteq A_i$ such that:
- $|S| < \mathsf{Ramsey}(s, t)$ and $S \subseteq A_i \cap \phi^{-1}(v)$; and
- every vertex $u \in A_{>i}$ that has a neighbor in $A_i \cap \phi^{-1}(v)$, also has a neighbor in S.

Proof. Let S be the smallest set contained in $A_i \cap \phi^{-1}(v)$ and satisfying the second condition, it exists, as this condition is satisfied by $A_i \cap \phi^{-1}(v)$. Note that since H is irreflexive, it follows that $\phi^{-1}(v)$ is an independent set in G, hence S is independent as well.

Suppose for contradiction that $|S| \geqslant \mathsf{Ramsey}(s, t)$. By minimality, for every $u \in S$ there exists $u' \in A_{>i}$ such that u is the only neighbor of u' in S. Let $S' := \{u' : u \in S\}$. Since $|S'| = |S| \geqslant \mathsf{Ramsey}(s, t)$, in $G[S']$ we can either find a clique K' of size s or an independent set I' of size t; denote $K := \{u : u' \in K'\}$ and $I := \{u : u' \in I'\}$. In the former case, we find that $\{x_i\} \cup K \cup K'$ induces the graph L_s in G, a contradiction. Similarly, in the latter case we have that $\{x_i\} \cup I \cup I'$ induces S_t in G, again a contradiction. This completes the proof of the claim. $\triangleleft$

Claim 8 suggests the following notion. A *guess* is a function $R: V(H) \to 2^{N[X]}$ satisfying that:
- for each $v \in V(H)$, $R(v)$ is a subset of $N[X]$ such that $|R(v) \cap A_i| < \mathsf{Ramsey}(s, t)$ for all $i \in \{1, 2, 3\}$; and
- sets $R(v)$ are pairwise disjoint for different $v \in V(H)$.

Let $\mathcal{R}$ be the family of all possible guesses; then we easily have the following.

$\triangleright$ **Claim 9.**   We have that $|\mathcal{R}| \leqslant n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$ and $\mathcal{R}$ can be enumerated in time $n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$.

Proof. For each $v \in V(H)$, the number of choices for $R(v)$ in a guess R is bounded by $2^3 \cdot n^{3 \cdot \mathsf{Ramsey}(s,t)}$: the first factor corresponds to the choice of $R(v) \cap X$, while the second factor bounds the number of choices of $R(v) \cap A_i$ for $i \in \{1, 2, 3\}$. Since the guess R is determined by choosing $R(v)$ for each $v \in V(H)$ and $|V(H)|$ is considered a constant, the number of different guesses is bounded by $\left(2^3 \cdot n^{3 \cdot \mathsf{Ramsey}(s,t)}\right)^{|V(H)|} = n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$. Clearly, they can be also enumerated in time $n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$. $\triangleleft$

Now, we say that a guess R is *compatible* with a solution ϕ to (G, rev) if the following conditions hold for every $v \in V(H)$:

(C1) $R(v) \subseteq \phi^{-1}(v)$;

(C2) $R(v) \cap X = \phi^{-1}(v) \cap X$; and

(C3) for all $i \in \{1, 2, 3\}$ and $u \in A_{>i}$, if u has a neighbor in $\phi^{-1}(v) \cap A_i$, then u also has a neighbor in $R(v) \cap A_i$.

The following statement follows immediately from Claim 8.

$\triangleright$ **Claim 10.**   For every solution ϕ to the instance (G, rev), there exists a guess $R \in \mathcal{R}$ that is compatible with ϕ.

Consider a guess $R \in \mathcal{R}$. We define a set $B^R \subseteq V(G) \times V(H)$ of *disallowed pairs* for R as follows. We include a pair $(u, v) \in V(G) \times V(H)$ in B^R if any of the following conditions holds:

(D1) $u \in X$ and $u \notin R(v)$;

(D2) $u \in R(v')$ for some $v' \in V(H)$ that is different from v;

(D3) u has a neighbor in G that belongs to $R(v')$ for some $v' \in V(H)$ such that $vv' \notin E(H)$; or

(D4) $u \in A_i - R(v)$ for some $i \in \{1, 2, 3\}$ and there exists $u' \in A_{>i}$ such that $uu' \in E(G)$ and $N_G(u') \cap A_i \cap R(v) = \emptyset$.

Intuitively, B^R contains assignments that contradict the supposition that R is compatible with a considered solution.

Based on B^R, we define a new revenue function $\mathsf{rev}^R \colon V(G) \times V(H) \to \mathbb{R}$ as follows:

$$\mathsf{rev}^R(u, v) = \begin{cases} -1 & \text{if } (u, v) \in B^R; \\ \mathsf{rev}(u, v) & \text{otherwise.} \end{cases}$$

The intuition is that disallowing a pair (u, v) is modelled by assigning a negative revenue to the corresponding assignment. This forbids optimum solutions from using this assignment.

We define a subgraph G' of G as follows: $V(G') := V(G)$ and $E(G')$ comprises all edges of G whose both endpoints belong to the same set A_i, for some $i \in \{1, 2, 3, 4\}$. Thus, in G' the vertices of X are isolated, and there are no edges between any A_i and A_j for $i \neq j$, nor between any A_i and X. For every guess $R \in \mathcal{R}$, we may consider a new instance (G', rev^R) of MAX PARTIAL H-COLORING. In the following two claims we establish the relationship between solutions to the instance (G, rev) and solutions to instances (G', rev^R) for $R \in \mathcal{R}$. The proofs essentially boil down to a verification that all the previous definitions work as expected. In particular, the key point is that the modification of revenues applied when constructing rev^R implies automatic satisfaction of all the constraints associated with edges that were present in G, but got removed in G'.

▷ **Claim 11.** For every guess $R \in \mathcal{R}$, every optimum solution ϕ to the instance (G', rev^R) is also a solution to the instance (G, rev), and moreover $\mathsf{rev}^R(\phi) = \mathsf{rev}(\phi)$.

Proof. Recall that ϕ is a solution to (G, rev) if and only if ϕ is a partial H-coloring of G. Hence, we need to prove that for every $uu' \in E(G)$ with $u, u' \in \mathsf{dom}\,\phi$, we have $\phi(u)\phi(u') \in E(H)$. Denote $v := \phi(u)$ and $v' := \phi(u')$ and suppose for contradiction that $vv' \notin E(H)$. Since ϕ is an optimum solution to (G', rev^R), we have $\mathsf{rev}^R(u, v) \geqslant 0$, which implies that $(u, v) \notin B^R$. Similarly $(u', v') \notin B^R$. We now consider cases depending on the alignment of u and u' in G.

If $u, u' \in A_i$ for some $i \in \{1, 2, 3, 4\}$ then $uu' \in E(G')$, so the supposition $vv' \notin E(H)$ would contradict the assumption that ϕ is a solution to (G', rev^R).

Suppose $u \in A_i$ and $u' \in A_j$ for $i, j \in \{1, 2, 3, 4\}$, $i \neq j$; by symmetry, assume $i < j$. As $vv' \notin E(H)$, we infer that u' does not have any neighbors in $R(v)$ in G, for otherwise we would have $(u', v') \in B^R$ by (D3). As $uu' \in E(G)$, $u \in A_i$, and $u' \in A_{>i}$, this implies that $(u, v) \in B^R$ by (D4), a contradiction.

Finally, suppose that $\{u, u'\} \cap X \neq \emptyset$, say $u \in X$. Since $(u, v) \notin B^R$, by (D1) we infer that $u \in R(v)$. Then, by (D3), $vv' \notin E(H)$ and $uu' \in E(G)$ together imply that $(u', v') \in B^R$, a contradiction.

This completes the proof that ϕ is a solution to (G, rev). To see that $\mathsf{rev}^R(\phi) = \mathsf{rev}(\phi)$ note that ϕ, being an optimum solution to (G', rev^R), does not use any assignments with negative revenues in rev^R, while $\mathsf{rev}(u, v) = \mathsf{rev}^R(u, v)$ for all (u, v) satisfying $\mathsf{rev}^R(u, v) \geqslant 0$. ◁

▷ **Claim 12.** If ϕ is a solution to (G, rev) that is compatible with a guess $R \in \mathcal{R}$, then ϕ is also a solution to (G', rev^R) and $\mathsf{rev}^R(\phi) = \mathsf{rev}(\phi)$.

Proof. As ϕ is a solution to (G, rev), it is a partial H-coloring of G. Since G' is a subgraph of G with $V(G') = V(G)$, ϕ is also a partial H-coloring of G'. Hence ϕ is a solution to (G', rev^R).

To prove that $\mathsf{rev}^R(\phi) = \mathsf{rev}(\phi)$ it suffices to show that $(u, \phi(u)) \notin B^R$ for every $u \in \mathsf{dom}\,\phi$, since functions rev^R and rev differ only on the pairs from B^R. Suppose otherwise, and consider cases depending on the reason for including $(u, \phi(u))$ in B^R. Denote $v := \phi(u)$.

First, suppose $u \in X$ and $u \notin R(v)$. By (C2) we have $u \notin R(v) \cap X = \phi^{-1}(v) \cap X \ni u$, a contradiction.

Second, suppose $u \in R(v')$ for some $v' \neq v$. By (C1) we have $v = \phi(u) = v'$, again a contradiction.

Third, suppose that u has a neighbor u' in G such that $u' \in R(v')$ for some $v' \in V(H)$ satisfying $vv' \notin E(H)$. By (C1), we have $u' \in \mathsf{dom}\,\phi$ and $\phi(u') = v'$. But then $\phi(u)\phi(u') = vv' \notin E(H)$ even though $uu' \in E(G)$, a contradiction with the assumption that ϕ is a partial H-coloring of G.

Fourth, suppose that $u \in A_i - R(v)$ for some $i \in \{1, 2, 3\}$ and there exists $u' \in A_{>i}$ such that $uu' \in E(G)$ and $N_G(u') \cap R(v) \cap A_i = \emptyset$. Observe that since $u \in A_i \cap \phi^{-1}(v)$ and $uu' \in E(G)$, by (C3) u' has a neighbor in $R(v) \cap A_i$ in the graph G. This contradicts the supposition that $N_G(u') \cap R(v) \cap A_i = \emptyset$.

As in all the cases we have obtained a contradiction, this concludes the proof of the claim. ◁

Let now $\Pi := \{\mathsf{rev}^R \colon R \in \mathcal{R}\}$. Then, condition (C2) can be easily derived from Claim 11 and Claim 12, while condition (C1) is implied by Claim 7. Note here that G' is $\{P_6, L_s, S_t\}$-free, because it is a disjoint union of induced subgraphs of G. Finally, from Claim 9 we infer that $|\Pi| = |\mathcal{R}| \leqslant n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$ and Π can be constructed in time $n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$, because given $R \in \mathcal{R}$ it is straightforward to construct rev^R in polynomial time. Hence Π satisfies all the requested properties, and this completes the proof of Lemma 6.

5   Corollaries for subclasses of P_6-free graphs

In this section we prove results (R3) and (R4) promised in Section 1. The idea is to apply
Lemma 6 exhaustively, until the considered instance becomes trivial. The main point is that
with each application the clique number of the graph drops, hence we naturally obtain an
upper bound of the form $n^{f(\omega(G))}$ for the total size of the recursion tree, hence also on the
running time.

The following statement captures the idea of exhaustive applying Lemma 6 in a recursive
scheme. For convenience, we formulate the statement so that s and t are given on input.

▶ **Theorem 13.** *Let H be a fixed irreflexive pattern graph. There exists an algorithm
that given $s, t \in \mathbb{N}$ and an instance (G, rev) of* MAX PARTIAL H-COLORING *where G is
$\{P_6, L_s, S_t\}$-free, solves this instance in time $n^{\mathcal{O}(\mathsf{Ramsey}(s,t) \cdot \omega(G))}$.*

Proof. If G is not connected, then for every connected component C of G we apply the
algorithm recursively to $(C, \mathsf{rev}|_{V(C)})$. If ϕ_C is the obtained optimum solution to this instance,
we may output $\phi := \bigcup_C \phi_C$. It is clear that ϕ constructed in this way is an optimum solution
to (G, rev).

Assume then that G is connected. If G consists of only one vertex, say u, then we may
simply output $\phi := \{(u, v)\}$ where v maximizes $\mathsf{rev}(u, v)$, or $\phi := \emptyset$ if $\mathsf{rev}(\cdot)$ has no positive
value in its range. Hence, assume that G has at least two vertices, in particular $\omega(G) \geqslant 2$.
We now apply Lemma 6 to G. Thus, in time $n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$ we obtain a subgraph G' of G
with $V(G) = V(G')$ and a suitable set of revenue functions Π satisfying $|\Pi| \leqslant n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$.
Recall here that G' is $\{P_6, L_s, S_t\}$-free. Moreover, if we set $F = K_{\omega(G)}$ then G is $F^\bullet$-free, so
Lemma 6 implies that G' is F-free. This means that $\omega(G') < \omega(G)$.

Next, for every $\mathsf{rev}' \in \Pi$ we recursively solve the instance (G', rev'). Lemma 6 implies
that if among the obtained optimum solutions to instances (G', rev') we pick the one with
the largest revenue, then this solution is also an optimum solution to (G, rev).

We are left with analyzing the running time. Recall that every time we recurse into
subproblems constructed using Lemma 6, the clique number of the currently considered
graph drops by at least one. Since recursing on a disconnected graph yields connected graphs
in subproblems, we conclude that the total depth of the recursion tree is bounded by $2 \cdot \omega(G)$.
In every recursion step we branch into $n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$ subproblems, hence the total number of
nodes in the recursion tree is bounded by $\left(n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}\right)^{2 \cdot \omega(G)} = n^{\mathcal{O}(\mathsf{Ramsey}(s,t) \cdot \omega(G))}$. The
internal computation in each subproblem take time $n^{\mathcal{O}(\mathsf{Ramsey}(s,t))}$, hence the total running
time is indeed $n^{\mathcal{O}(\mathsf{Ramsey}(s,t) \cdot \omega(G))}$.   ◀

Note that since both L_3 and S_2 contain P_5 as an induced subgraph, every P_5-free graph
is $\{P_6, L_3, S_2\}$-free. Hence, from Theorem 13 we may immediately conclude the following
statement, where the setting of P_5-free graphs is covered by the case $s = 3$ and $t = 2$.

▶ **Corollary 14.** *For any fixed $s, t \in \mathbb{N}$ and irreflexive pattern graph H,* MAX PARTIAL
H-COLORING *can be solved in $\{P_6, L_s, S_t\}$-free graphs in time $n^{\mathcal{O}(\omega(G))}$. This in particular
applies to P_5-free graphs.*

Next, we observe that the statement of Theorem 13 can be also used for non-constant s
to obtain an algorithm for the case when the graph L_s is not excluded.

▶ **Corollary 15.** *For any fixed $t \in \mathbb{N}$ and irreflexive pattern graph H,* MAX PARTIAL H-
COLORING *can be solved in $\{P_6, S_t\}$-free graphs in time $n^{\mathcal{O}(\omega(G)^t)}$.*

Proof. Observe that since the graph L_s contains a clique of size s, every graph G is actually $L_{\omega(G)+1}$-free. Therefore, we may apply the algorithm of Theorem 13 for $s := \omega(G)+1$. Note here that $\omega(G)$ can be computed in time $n^{\omega(G)+\mathcal{O}(1)}$ by verifying whether G has cliques of size $1, 2, 3, \dots$ up to the point when the check yields a negative answer. Since for $s = \omega(G)+1$ and fixed t we have

$$\mathsf{Ramsey}(s,t) = \binom{s+t-2}{t-1} \leqslant \mathcal{O}(\omega(G)^{t-1}),$$

the obtained running time is $n^{\mathcal{O}(\mathsf{Ramsey}(s,t)\cdot\omega(G))} \leqslant n^{\mathcal{O}(\omega(G)^t)}$. ◄

Let us note that an algorithm with running time $n^{\mathcal{O}(\omega(G)^\alpha)}$, for some constant α, can be used within a simple branching strategy to obtain a subexponential-time algorithm.

▶ **Lemma 16.** *Let H be a fixed irreflexive graph and suppose MAX PARTIAL H-COLORING can be solved in time $n^{\mathcal{O}(\omega(G)^\alpha)}$ on $\mathcal{F}$-free graphs, for some family of graphs $\mathcal{F}$ and some constant $\alpha \geqslant 1$. Then MAX PARTIAL H-COLORING can be solved in time $n^{\mathcal{O}(n^{\alpha/(\alpha+1)})}$ on $\mathcal{F}$-free graphs.*

Proof. Let (G, rev) be the input instance, where G has n vertices. We define threshold $\tau := \left\lfloor n^{\frac{1}{\alpha+1}} \right\rfloor$.

The algorithm first checks whether G contains a clique on τ vertices. This can be done in time $n^{\tau+\mathcal{O}(1)} \leqslant n^{\mathcal{O}(n^{1/(\alpha+1)})}$ by verifying all subsets of τ vertices in G. If there is no such clique then $\omega(G) < \tau$, so we can solve the problem using the assumed algorithm in time $n^{\mathcal{O}(\omega(G)^\alpha)} \leqslant n^{\mathcal{O}(\tau^\alpha)} \leqslant n^{\mathcal{O}(n^{\alpha/(\alpha+1)})}$. Hence, suppose that we have found a clique K on τ vertices.

Observe that since H is irreflexive, in any partial H-coloring ϕ of G only at most $|V(H)|$ vertices of K can be colored, that is, belong to $\mathsf{dom}\,\phi$. We recurse into $\binom{\tau}{\leqslant|V(H)|} \leqslant n^{|V(H)|}$ subproblems: in each subproblem we fix a different subset $A \subseteq K$ with $|A| \leqslant |V(H)|$ and recurse on the graph $G_A := G - (K - A)$ with revenue function $\mathsf{rev}_A := \mathsf{rev}|_{V(G_A)}$. Note here that G_A is $\mathcal{F}$-free. From the above discussion it is clear that $\mathrm{OPT}(G, \mathsf{rev}) = \max_{A \subseteq K, |A| \leqslant |V(H)|} \mathrm{OPT}(G_A, \mathsf{rev}_A)$. Therefore, the algorithm may return the solution with the highest revenue among those obtained in recursive calls.

As for the running time, observe that in every recursive call, the algorithm either solves the problem in time $n^{\mathcal{O}(n^{\alpha/(\alpha+1)})}$, or recurses into $n^{|V(H)|} = n^{\mathcal{O}(1)}$ subcalls, where in each subcall the vertex count is decremented by at least $\left\lfloor n^{\frac{1}{\alpha+1}} \right\rfloor$. It follows that the depth of the recursion is bounded by $\mathcal{O}(n^{\alpha/(\alpha+1)})$, hence the total number of nodes in the recursion tree is at most $n^{\mathcal{O}(n^{\alpha/(\alpha+1)})}$. Since the time used for each node is bounded by $n^{\mathcal{O}(n^{\alpha/(\alpha+1)})}$, the total running time of $n^{\mathcal{O}(n^{\alpha/(\alpha+1)})}$ follows. ◄

By combining Corollary 14 and Corollary 15 with Lemma 16 we conclude the following.

▶ **Corollary 17.** *For any fixed $s, t \in \mathbb{N}$ and irreflexive pattern graph H, MAX PARTIAL H-COLORING can be solved in*
1. *$\{P_6, L_s, S_t\}$-free graphs in time $n^{\mathcal{O}(\sqrt{n})}$ (this in particular applies to P_5-free graphs),*
2. *$\{P_6, S_t\}$-free graphs in time $n^{\mathcal{O}(n^{t/(t+1)})}$.*

6   Open problems

The following question, which originally motivated our work, still remains unresolved.

▶ **Question 1.** *Is* Odd Cycle Transversal *polynomial-time solvable in P_5-free graphs?*

Note that our work stops short of giving a positive answer to this question: we give an algorithm with running time $n^{\mathcal{O}(\omega(G))}$, a subexponential-time algorithm, and polynomial time algorithms for the cases when either a threshold graphs or a bull is additionally forbidden. Therefore, we are hopeful that the answer to the question is indeed positive.

One aspect of our work that we find particularly interesting is the possibility of treating the clique number $\omega(G)$ as a progress measure for an algorithm, which enables bounding the recursion depth in terms of $\omega(G)$. This approach naturally leads to algorithms with running time of the form $n^{f(\omega(G))}$ for some function f, that is, polynomial-time for every fixed clique number. By Lemma 16, having a polynomial function f in the above gives a subexponential-time algorithm, at least in the setting of Max Partial H-Coloring for irreflexive H. However, looking for algorithms with time complexity $n^{f(\omega(G))}$ seems to be another relaxation of the goal of polynomial-time solvability, somewhat orthogonal to subexponential-time algorithms [4, 6, 19] or approximation schemes [10]. Note that our work and the recent work of Brettell et al. [7] actually show two different methods of obtaining such algorithms: using direct recursion, or via dynamic programming on branch decompositions of bounded mim-width. It would be interesting to investigate this direction in the context of Maximum Independent Set in P_t-free graphs. A concrete question would be:

▶ **Question 2.** *Is there a polynomial-time algorithm for* Maximum Independent Set *in $\{P_t, K_t\}$-free graphs, for every fixed t?*

In all our algorithms, we state the time complexity assuming that the pattern graph H is fixed. This means that the constants hidden in the $\mathcal{O}(\cdot)$ notation in the exponent may – and do – depend on the size of H. In the language of parameterized complexity, this means that we give XP algorithms for the parameterization by the size of H. It is natural to ask whether this state of art can be improved to the existence of FPT algorithms, that is, with running time $f(H) \cdot n^c$ for some computable function f and universal constant c, independent of H. This is not known even for the case of k-Coloring P_5-free graphs, so let us re-iterate the old question of Hoàng et al. [22].

▶ **Question 3.** *Is there an* FPT *algorithm for k-*Coloring *in P_5-free graphs parameterized by k?*

While the above question seems hard, it is conceivable that FPT results could be derived in some more restricted settings considered in this work, for instance for $\{P_5, \text{bull}\}$-free graphs.

───── **References** ─────

1   Tara Abrishami, Maria Chudnovsky, Marcin Pilipczuk, Paweł Rzążewski, and Paul Seymour. Induced subgraphs of bounded treewidth and the container method. *CoRR*, abs/2003.05185, 2020. `arXiv:2003.05185`.

2   Vladimir E. Alekseev. The effect of local constraints on the complexity of determination of the graph independence number. *Combinatorial-algebraic methods in applied mathematics*, pages 3–13, 1982. (in Russian).

3   Vladimir E. Alekseev. Polynomial algorithm for finding the largest independent sets in graphs without forks. *Discret. Appl. Math.*, 135(1-3):3–16, 2004. `doi:10.1016/S0166-218X(02)00290-1`.

4   Gábor Bacsó, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Zsolt Tuza, and Erik Jan van Leeuwen. Subexponential-time algorithms for Maximum Independent Set in P_t-free and broom-free graphs. *Algorithmica*, 81(2):421–438, 2019. `doi:10.1007/s00453-018-0479-5`.

5   Flavia Bonomo, Maria Chudnovsky, Peter Maceli, Oliver Schaudt, Maya Stein, and Mingxian Zhong. Three-coloring and list three-coloring of graphs without induced paths on seven vertices. *Combinatorica*, 38(4):779–801, 2018. `doi:10.1007/s00493-017-3553-8`.

6   Christoph Brause. A subexponential-time algorithm for the Maximum Independent Set problem in P_t-free graphs. *Discret. Appl. Math.*, 231:113–118, 2017. `doi:10.1016/j.dam.2016.06.016`.

7   Nick Brettell, Jake Horsfield, and Daniël Paulusma. Colouring $sP_1 + P_5$-free graphs: a mim-width perspective. *CoRR*, abs/2004.05022, 2020. `arXiv:2004.05022`.

8   Maria Chudnovsky, Jason King, Michal Pilipczuk, Paweł Rzążewski, and Sophie Spirkl. Finding large H-colorable subgraphs in hereditary graph classes. *CoRR*, abs/2004.09425, 2020. `arXiv:2004.09425`.

9   Maria Chudnovsky, Daniël Paulusma, and Oliver Schaudt. Graph colouring: from structure to algorithms (dagstuhl seminar 19271). *Dagstuhl Reports*, 9(6):125–142, 2019. `doi:10.4230/DagRep.9.6.125`.

10 Maria Chudnovsky, Marcin Pilipczuk, Michał Pilipczuk, and Stéphan Thomassé. Quasi-polynomial time approximation schemes for the Maximum Weight Independent Set problem in H-free graphs. In *Proceedings of the 31^{st} ACM-SIAM Symposium on Discrete Algorithms, SODA 2020*, pages 2260–2278. SIAM, 2020. `doi:10.1137/1.9781611975994.139`.

11 Maria Chudnovsky and Shmuel Safra. The Erdős-Hajnal conjecture for bull-free graphs. *J. Comb. Theory, Ser. B*, 98(6):1301–1310, 2008. `doi:10.1016/j.jctb.2008.02.005`.

12 Maria Chudnovsky, Oliver Schaudt, Sophie Spirkl, Maya Stein, and Mingxian Zhong. Approximately coloring graphs without long induced paths. *Algorithmica*, 81(8):3186–3199, 2019. `doi:10.1007/s00453-019-00577-6`.

13 Maria Chudnovsky and Vaidy Sivaraman. Odd holes in bull-free graphs. *SIAM J. Discrete Math.*, 32(2):951–955, 2018. `doi:10.1137/17M1131301`.

14 Maria Chudnovsky, Sophie Spirkl, and Mingxian Zhong. Four-coloring P_6-free graphs. In *Proceedings of the 30^{th} Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019*, pages 1239–1256. SIAM, 2019. `doi:10.1137/1.9781611975482.76`.

15 Bruno Courcelle, Johann A. Makowsky, and Udi Rotics. Linear time solvable optimization problems on graphs of bounded clique-width. *Theory Comput. Syst.*, 33(2):125–150, 2000. `doi:10.1007/s002249910009`.

16 Konrad K. Dabrowski, Carl Feghali, Matthew Johnson, Giacomo Paesani, Daniël Paulusma, and Paweł Rzążewski. On cycle transversals and their connected variants in the absence of a small linear forest. *CoRR*, abs/1908.00491, 2019. Accepted to Algorithmica. `arXiv:1908.00491`.

17 Fedor V. Fomin, Ioan Todinca, and Yngve Villanger. Large induced subgraphs via triangulations and CMSO. *SIAM J. Comput.*, 44(1):54–87, 2015. `doi:10.1137/140964801`.

18 Petr A. Golovach, Daniël Paulusma, and Jian Song. Closing complexity gaps for coloring problems on H-free graphs. *Inf. Comput.*, 237:204–214, 2014. `doi:10.1016/j.ic.2014.02.004`.

19 Carla Groenland, Karolina Okrasa, Paweł Rzążewski, Alex D. Scott, Paul D. Seymour, and Sophie Spirkl. H-colouring P_t-free graphs in subexponential time. *Discret. Appl. Math.*, 267:184–189, 2019. `doi:10.1016/j.dam.2019.04.010`.

20 Andrzej Grzesik, Tereza Klimošová, Marcin Pilipczuk, and Michał Pilipczuk. Polynomial-time algorithm for Maximum Weight Independent Set on P_6-free graphs. In *Proceedings of the 30^{th} Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019*, pages 1257–1271. SIAM, 2019. `doi:10.1137/1.9781611975482.77`.

21 Gregory Z. Gutin, Pavol Hell, Arash Rafiey, and Anders Yeo. A dichotomy for minimum cost graph homomorphisms. *Eur. J. Comb.*, 29(4):900–911, 2008. `doi:10.1016/j.ejc.2007.11.012`.

22 Chính T Hoàng, Marcin Kamiński, Vadim Lozin, Joe Sawada, and Xiao Shu. Deciding k-colorability of P_5-free graphs in polynomial time. *Algorithmica*, 57(1):74–81, 2010.

23 Shenwei Huang. Improved complexity results on k-coloring P_t-free graphs. *Eur. J. Comb.*, 51:336–346, 2016. `doi:10.1016/j.ejc.2015.06.005`.

24 Daniel Lokshtanov, Martin Vatshelle, and Yngve Villanger. Independent set in P_5-free graphs in polynomial time. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*, pages 570–581. SIAM, 2014. `doi:10.1137/1.9781611973402.43`.

25 Vadim V. Lozin and Martin Milanič. A polynomial algorithm to find an independent set of maximum weight in a fork-free graph. *J. Discrete Algorithms*, 6(4):595–604, 2008. `doi:10.1016/j.jda.2008.04.001`.

26 George J. Minty. On maximal independent sets of vertices in claw-free graphs. *J. Comb. Theory, Ser. B*, 28(3):284–304, 1980. `doi:10.1016/0095-8956(80)90074-X`.

27 Jana Novotná, Karolina Okrasa, Michał Pilipczuk, Paweł Rzążewski, Erik Jan van Leeuwen, and Bartosz Walczak. Subexponential-time algorithms for finding large induced sparse subgraphs. In *Proceedings of the 14th International Symposium on Parameterized and Exact Computation, IPEC 2019*, volume 148 of *LIPIcs*, pages 23:1–23:11. Schloss Dagstuhl — Leibniz-Zentrum für Informatik, 2019. `doi:10.4230/LIPIcs.IPEC.2019.23`.

28 Karolina Okrasa and Paweł Rzążewski. Subexponential algorithms for variants of the homomorphism problem in string graphs. *J. Comput. Syst. Sci.*, 109:126–144, 2020. `doi:10.1016/j.jcss.2019.12.004`.

29 Najiba Sbihi. Algorithme de recherche d'un stable de cardinalité maximum dans un graphe sans étoile. *Discrete Mathematics*, 29(1):53–76, 1980. (in French).