COMBINATORICA
Bolyai Society – Springer-Verlag

Combinatorica 23pp.
DOI: 10.1007/s00493-017-3553-8

THREE-COLORING AND LIST THREE-COLORING OF GRAPHS WITHOUT INDUCED PATHS ON SEVEN VERTICES

FLAVIA BONOMO*, MARIA CHUDNOVSKY†, PETER MACELI, OLIVER SCHAUDT, MAYA STEIN‡, MINGXIAN ZHONG

Received December 23, 2015
Revised August 26, 2016

In this paper we present a polynomial time algorithm that determines if an input graph containing no induced seven-vertex path is 3-colorable. This affirmatively answers a question posed by Randerath, Schiermeyer and Tewes in 2002. Our algorithm also solves the list-coloring version of the 3-coloring problem, where every vertex is assigned a list of colors that is a subset of $\{1,2,3\}$, and gives an explicit coloring if one exists.

1. Introduction

A *k-coloring* of a graph $G=(V,E)$ is a function $f\colon V\to\{1,\dots,k\}$ such that $f(v)\neq f(w)$ whenever $vw\in E$. The *vertex coloring problem*, whose input is a graph G and a natural number k, consists of deciding whether G is k-colorable or not. This well-known problem is one of Karp's 21 NP-complete problems [16] (unless $k=2$; then the problem is solvable in linear time). Stockmeyer [24] proved that the problem remains NP-complete even if $k\geq 3$ is fixed, and Maffray and Preissmann proved that it remains NP-complete for triangle-free graphs [19].

List variations of the vertex coloring problem can be found in the literature. For a survey on that kind of related problems, see [25]. In the *list-coloring problem*, every vertex v comes equipped with a list of permitted

Mathematics Subject Classification (2000): 05C15, 05C37, 05C85

* Partially supported by MathAmSud Project 13MATH-07, UBACyT Grant 20020130100808BA, CONICET PIP 112-201201-00450CO and ANPCyT PICT 2012-1324.

† Partially supported by NSF grants IIS-1117631, DMS-1001091 and DMS-1265803.

‡ Partially supported by MathAmSud Project 13MATH-07, Fondecyt grant 1140766, and Millenium Nucleus Information and Coordination in Networks.

colors $L(v)$, and we require the coloring to respect these lists, i.e., $f(v) \in L(v)$ for every v in V. For a positive integer k, the *k-list-coloring problem* is a particular case in which $|L(v)| \leq k$ for each v in V, but the union of the lists can be an arbitrary set. If the size of the list assigned to each vertex is at most two (i.e., 2-list-coloring), the instance can be solved in $O(|V|+|E|)$ time [6,7,26], by reducing the problem to a 2-SAT instance, which Aspvall, Plass and Tarjan [1] showed can be solved in linear time (in the number of variables and clauses). The *list k-coloring problem* is a particular case of k-list-coloring, in which the lists associated to each vertex are a subset of $\{1,\dots,k\}$. Since list k-coloring generalizes k-coloring, it is NP-complete as well.

Because of the notorious hardness of k-coloring, efforts were made to understand the problem on restricted graph classes. Some of the most prominent such classes are the classes of H-free graphs, i.e., graphs containing no induced subgraph isomorphic to H, for some fixed graph H. Kamiński and Lozin [15] and independently Král, Kratochvíl, Tuza, and Woeginger [17] proved that for any fixed $k, g \geq 3$, the k-coloring problem is NP-complete for the class of graphs containing no cycle of length less than g. As a consequence, if the graph H contains a cycle, then k-coloring is NP-complete for $k \geq 3$ for the class of H-free graphs.

The *claw* is the complete bipartite graph $K_{1,3}$. A theorem of Holyer [12] together with an extension due to Leven and Galil [18] imply that if a graph H contains a claw, then for every fixed $k \geq 3$, the k-coloring problem is NP-complete for the class of H-free graphs.

Combined, these two results only leave open the complexity of the k-coloring problem for the class of H-free graphs where H is a fixed acyclic claw-free graph, i.e., a disjoint union of paths. There is a nice recent survey by Hell and Huang on the complexity of coloring graphs without paths and cycles of certain lengths [10] and another nice survey by Golovach et al. [8]. We denote a path and a cycle on t vertices by P_t and C_t, respectively.

The strongest known results related to our work are due to Huang [13], who proved that 4-coloring is NP-complete for P_7-free graphs, and that 5-coloring is NP-complete for P_6-free graphs. On the positive side, Hoàng, Kamiński, Lozin, Sawada, and Shu [11] have shown that k-coloring can be solved in polynomial time on P_5-free graphs for any fixed k. Huang [13] conjectures that 4-coloring is polynomial-time solvable for P_6-free graphs. This conjecture, if true, thus settles the last remaining open case of the complexity of k-coloring P_t-free graphs for any fixed $k \geq 4$. On the other hand, for $k=3$ it is not known whether there exists a t such that 3-coloring is NP-complete for P_t-free graphs. Randerath and Schiermeyer [21] gave a

k \ t	4	5	6	7	8	...
3	$O(m)$ [5]	$O(n^\alpha)$ [20]	$O(mn^\alpha)$ [21]	**P**	?	...
4	$O(m)$ [5]	P [11]	?	NPC [13]	NPC	...
5	$O(m)$ [5]	P [11]	NPC [13]	NPC	NPC	...
6	$O(m)$ [5]	P [11]	NPC	NPC	NPC	...
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\ddots$

Table 1. Table of known complexities of the k-coloring problem in P_t-free graphs. Here, n is the number of vertices in the input graph, m the number of edges, and α is the matrix multiplication exponent known to satisfy $2 \leq a < 2.376$ [4]. The boldfaced complexity is the topic of this paper, while '?' stands for open problems.

polynomial time algorithm for 3-coloring P_6-free graphs. Later, Golovach et al. [9] showed that the list 3-coloring problem can be solved efficiently for P_6-free graphs. Some of these results are summarized in Table 1.

We show that the 3-coloring problem for P_7-free graphs is polynomial, answering positively a question first posed in 2002 by Randerath et al. [21,22]. Our algorithm even works for the list 3-coloring problem. This is not trivial: there are cases where k-coloring and list k-coloring have different complexities (unless P = NP). For instance, in the class of $\{P_6, C_5\}$-free graphs, 4-coloring can be solved in polynomial time [3] while list 4-coloring is NP-complete [14]. Our main theorem reads as follows.

Theorem 1. *One can decide whether a given P_7-free graph G has a list 3-coloring, and find such a coloring (if it exists) in polynomial time. The running time of the proposed algorithm is $O(|V(G)|^{21}(|V(G)|+|E(G)|))$.*

The algorithm given by Theorem 1 is based on the following ideas. First we apply some preprocessing techniques and compute a small 2-dominating set (i.e., a set such that every vertex has distance at most two to some vertex of the set). Then we use a controlled enumeration based on a structural analysis of the considered graphs, in order to reduce the problem to a polynomial number of instances of list 3-coloring in which the size of the list of each vertex is at most two. These instances, in turn, can be solved via 2-SAT.

2. Notation and preliminaries

We start by establishing some notation and preliminary results. A *stable set* in a graph G is a subset of pairwise non-adjacent vertices of G. Let X and Y be two sets of vertices of G. We say that X is *complete to Y* if every vertex in X is adjacent to every vertex in Y, and that X is *anticomplete to* Y if no vertex of X is adjacent to a vertex of Y.

If in a graph coloring context each of the vertices v in G is assigned a list $L(v) \subseteq \{1,2,3\}$ of possible colors, we call $L = \{L(v) \colon v \in V(G)\}$ a *palette* of G. A palette L' is a *subpalette* of L if $L'(v) \subseteq L(v)$ for each $v \in V(G)$. Given a graph G and a palette L, we say that a 3-coloring c of G is a *coloring of* (G,L) if $c(v) \in L(v)$ for all $v \in V(G)$. We also say that c is a coloring of G *for the palette* L. We say that (G,L) is *colorable* if there exists a coloring of (G,L). We denote by $(G,\mathcal{L})$ a graph G and a collection $\mathcal{L}$ of palettes of G. We say $(G,\mathcal{L})$ is *colorable* if (G,L) is colorable for some $L \in \mathcal{L}$. Further, c is a *coloring* of $(G,\mathcal{L})$ if c is a coloring of (G,L) for some $L \in \mathcal{L}$.

An *update* of the list of a vertex v *from* w means we delete an entry from the list of v that appears as the unique entry of the list of a neighbor w of v. Clearly, such an update does not change the colorability of the graph. If a palette L' is obtained from a palette L by updating repeatedly until for every vertex v, if v has a neighbor u with $L'(u) = \{i\}$, then $i \notin L'(v)$, we say we obtained L' from L by *updating*. For a fixed $w \in V(G)$ if a palette L' is obtained from a palette L by repeatedly updating vertices v from vertices w' that are connected to w by a path all whose vertices have current lists of size one, and continuing to do so until no candidates for updating are left, then we say we obtained palette L' from palette L by *updating from* w. Finally, if in either of these two procedures we update all vertices v except those from a fixed set T, we say we obtained L' by *updating except on* T.

Let us illustrate these notions with a quick example. Consider C_6 with lists $\{1\}$, $\{2,3\}$, $\{2\}$, $\{1,2\}$, $\{2,3\}$, $\{1,2\}$ (in this order). Then updating from v_1 gives lists $\{1\}$, $\{2,3\}$, $\{2\}$, $\{1,2\}$, $\{3\}$, $\{2\}$, while updating from v_1 except on $\{v_6\}$ leaves us with the initial lists. Note that updating can be carried out in $O(|V(G)|+|E(G)|)$ time.

By reducing to an instance of 2-SAT, which can be solved in linear time in the number of variables and clauses [1], several authors [6,7,26] independently proved the following.

Lemma 2. *If a palette L of a graph G is such that $|L(v)| \leq 2$ for all $v \in V(G)$, then a coloring of (G,L), or a determination that none exists, can be obtained in $O(|V(G)|+|E(G)|)$ time.*

Let G be a graph. A subset S of $V(G)$ is called *monochromatic* with respect to a given coloring c of G if $c(u) = c(v)$ for all $u,v \in S$. Let L be a palette of G, and Z a set of subsets of $V(G)$. We say that (G,L,Z) is *colorable* if there is a coloring c of (G,L) such that S is monochromatic with respect to c for all $S \in Z$.

A triple (G',L',Z') is a *restriction* of (G,L,Z) if

(a) G' is an induced subgraph of G,
(b) the palette L' is a subpalette of L restricted to the set $V(G')$, and
(c) Z' is a set of subsets of $V(G')$ such that if $S \in Z$ then $S \cap V(G') \subseteq S'$ for some $S' \in Z'$.

Let $\mathcal{R}$ be a set of restrictions of (G,L,Z). We say that $\mathcal{R}$ is *colorable* if at least one element of $\mathcal{R}$ is colorable. If $\mathcal{L}$ is a set of palettes of G, we write $(G,\mathcal{L},Z)$ to mean the set of restrictions (G,L',Z) where $L' \in \mathcal{L}$.

Note that if two sets S and S' are monochromatic and have a non-empty intersection, then $S \cup S'$ is monochromatic, too. Thus, for each triple (G,L,Z) there is an equivalent triple (G,L,Z') such that Z' contains only mutually disjoint sets. During our algorithm, we compute the set family Z such that the sets are mutually disjoint. Under this assumption, the proof of Lemma 2 can be easily modified to obtain the following generalization [23].

Lemma 3. *If a palette L of a graph G is such that $|L(v)| \leq 2$ for all $v \in V(G)$, and Z is a set of mutually disjoint subsets of $V(G)$, then a coloring of (G,L,Z), or a determination that none exists, can be obtained in $O(|V(G)| + |E(G)|)$ time.*

Proof. By traversing once each set in Z, create a vector r that maps each vertex v with a representative $r(v)$ on its set (the same representative for all the vertices in one set). Define $r(v) = v$ if v does not belong to a set in Z. Traversing the vector r once, iteratively for each $v \in V(G)$, update $L(r(v)) = L(r(v)) \cap L(v)$. If at some point $L(r(v)) = \emptyset$, return that no coloring exists. These steps can be performed in $O(|V(G)|)$ time.

If none of the lists $L(r(v))$ is empty, compute the 2-SAT formula that expresses the coloring problem of (G,L,Z), similarly as for (G,L) in Lemma 2. Namely, define for each vertex $v \in V(G)$ and each color $j \in L(r(v))$ the variable $x_{r(v)j}$ to model that vertex v gets color j. Notice that if v and w are in the same set of Z, then $r(v) = r(w)$, thus the sets of Z will be monochromatic in every coloring derived from a solution of the formula.

If $L(r(v)) = \{j\}$, add $(x_{r(v)j})$ as a clause, and if $L(r(v)) = \{j,k\}$, add $(x_{r(v)j} \vee x_{r(v)k})$ as a clause. This ensures every vertex gets a color. Finally, for each edge $vw \in E(G)$ and each color $j \in L(r(v)) \cap L(r(w))$, add the clause $(\neg x_{r(v)j} \vee \neg x_{r(w)j})$. This ensures two adjacent vertices get different colors. Notice also that two adjacent vertices in the same set of Z will produce an unfeasible formula, as desired.

The formula can be constructed in $O(|V(G)| + |E(G)|)$ time and has $O(|V(G)| + |E(G)|)$ variables and clauses. Since the algorithm that solves 2-SAT is linear in the number of variables and clauses [1], we are done. ∎

We write $N(S)$ for the set of vertices of $V(G)\setminus S$ with a neighbor in S. For disjoint sets of vertices S,T of $V(G)$, let $N_T(S)=N(S)\cap T$. If $S=\{s\}$ we just write $N_T(s)$. For a vertex set S, let $\overline{S}=S\cup N(S)$. If $\overline{S}=V(G)$, we say that S is *dominating* G, or is a *dominating set*. Moreover, if S is dominating and the subgraph induced by S is connected, then we call S a *connected dominating set*. If $\overline{S}$ dominates G, we call S *2-dominating*.

For a graph G with a palette L, call a (nonempty) 2-dominating set $S\subseteq V(G)$ which induces a connected subgraph a *seed* of (G,L), if $|L(v)|=1$ for each $v\in S$ and $|L(v)|=2$ for each $v\in N(S)$. Note that we do not require the palette L to be updated.

Observe that for any seed S, and for any two non-adjacent vertices $v,w\in N(S)$ the following holds.

(1) There is an induced v–w path of at least 3 vertices whose inner vertices all lie in S.

The next result is essential to our proof.

Theorem 4 (Camby and Schaudt [2]). *For all $t\geq 3$, any connected P_t-free graph has a connected dominating set whose induced subgraph is either P_{t-2}-free, or isomorphic to P_{t-2}.*

We use the following easy corollary of Theorem 4 in order to prove the existence of a small seed in P_7-free graphs that may be 3-colorable.

Corollary 5. *Every connected P_7-free graph G has either a connected 2-dominating set of size at most 3 or a complete subgraph of 4 vertices. The set or the subgraph can be found in $O(|V(G)|^3|E(G)|)$ time.*

Proof. We prove the first statement by applying Theorem 4 to the graph in question, say G. Let S be the connected dominating set of G whose induced subgraph, say H, is either a P_5 or P_5-free. If H is a P_5, the three non-leaf vertices of H form a connected 2-dominating set of G, as desired. Otherwise, another application of Theorem 4 shows that H has a connected dominating set S' whose induced subgraph is either a P_3 or P_3-free. If $|S'|\leq 3$, S' is a connected 2-dominating set of G of at most three vertices. Otherwise, as a connected P_3-free graph is complete, $|S'|\geq 4$ implies that G contains a complete subgraph on 4 vertices.

Now we turn to the second statement. It suffices to run through all triples T of vertices ($O(|V(G)|^3)$ triples), and check if there is a common neighbor v of T such that $T\cup\{v\}$ induces a complete subgraph ($O(|E(G)|)$ possible vertices v). If not, we check whether T induces a connected subgraph and all vertices of the graph are within distance 2 from T. We can test the second property by using two steps of a breadth-first-search (that has time complexity $O(|E(G)|)$). ∎

This corollary will help us to reduce in the next section the original instance to a polynomial number of simpler instances. In each of these, the vertices having lists of size 1 or 2 satisfy some structural properties and the vertices having lists of size 3 form a stable set. We will in turn solve these special instances in Section 3.1 by reducing them to a polynomial number of instances to which we can apply Lemma 3.

3. Proof of Theorem 1

Let G be a graph and v be a vertex of G. Observe first that if $G[N(v)]$ is not bipartite, then G is not 3-colorable. Observe also that if $G[N(v)]$ is a connected bipartite graph with bipartition U, W, then in every 3-coloring of G each of the sets U and W is monochromatic.

Let (G, L) be a list 3-coloring instance, such that for every $v \in V(G)$, $G[N(v)]$ is bipartite. We now describe a procedure that we call the *neighborhood reduction.*

If there is a vertex v with $|L(v)| = 3$ such that $G[N(v)]$ is connected, proceed as follows. Let U, W be a bipartition of $G[N(v)]$. We construct the graph G' we obtain from G by deleting v and replacing the neighborhood of v with an edge uw, where $N_{G'}(u) \cap V(G) = N_G(U) \cap V(G')$, and $N_{G'}(w) \cap V(G) = N_G(W) \cap V(G')$. In the case that W is empty, say, we can assume $U = \{u\}$, and we just define $G' = G - \{v\}$. The list of u is the intersection of all lists of vertices from U, and similar for w and W. Clearly, G admits a coloring for L if and only if G' admits a coloring for the new palette.

We iterate the above procedure until $G[N(v)]$ is disconnected for each vertex v with $|L(v)| = 3$. The term neighborhood reduction refers to the whole process until it stops.

The following claim says that this reduction preserves the property of being P_t-free, for $t \geq 3$.

Claim 6. *If G is a P_t-free graph ($t \geq 3$), then the graph obtained from the neighborhood reduction is P_t-free.*

Proof. It suffices to consider one reduction step. Let us say we contracted the neighborhood of the vertex v in G, and obtained the graph G'. It remains to show that G' is still P_t-free.

To see this, suppose Q is an induced P_t in G'. Since G is P_t-free, it follows that $V(Q) \cap \{u, w\}$ is non-empty. Note that if Q contains both u and w, then

u,w are consecutive on Q. So (in any case) we can write Q as $Q_1-Q_2-Q_3$, where $V(Q_2)\subseteq\{u,w\}$ and Q_1,Q_3 avoid $\{u,w\}$. We can assume that Q_1,Q_3 are not empty, as otherwise it is easy to substitute Q_2 with one or two vertices in $U\cup W$, and thus find an induced P_t in G, a contradiction.

Observe that Q_1,Q_3 each have exactly one vertex q_1,q_3 in $N(V(Q_2))$. If $|V(Q_2)|=1$, we may assume both these vertices lie in $N(U)$, and we can substitute $Q_2=u$ with either a common neighbor of q_1,q_3, or with a path u_1-v-u_2 with $u_1\in U\cap N(q_1)$ and $u_2\in U\cap N(q_3)$. This gives an induced P_t in G, a contradiction.

So assume $|V(Q_2)|=2$, and without loss of generality $Q_2=u-w$, q_1 is adjacent to u and q_3 to w. Then q_1 is anticomplete to W and has a neighbor u_1 in U, and q_3 is anticomplete to U and has a neighbor w_1 in W. We can thus replace Q_2 with the path u_1-w_1 if they are adjacent, or with the path u_1-v-w_1 if they are not. This gives an induced P_t in G, yielding the final contradiction. ∎

Let G^* be a connected P_7-free graph with a palette L^*. We preprocess first the instance by applying the neighborhood reduction according to the input palette L^*, but, in order to simplify the following presentation and discussion of our algorithm, after that preprocessing, we do not take the input palette L^* into account. Instead, we consider the palette L^1 with $L^1(v)=\{1,2,3\}$ for each vertex v. We intersect the current lists with L^* at the very end of the first phase of the algorithm only.

Here is an overview over the steps taken in the algorithm.

(a) Assert that for every vertex v of G^*, $G^*[N(v)]$ is bipartite. Otherwise, we can report that G^* is not 3-colorable.
(b) Reduce the instance so that the neighborhood of every vertex v with $|L^*(v)|=3$ is disconnected. Let G be the graph obtained. By Claim 6, G is P_7-free.
(c) Apply Corollary 5 to G and obtain a connected 2-dominating set S_1 of size at most 3. (Notice that as we have asserted that every vertex has a bipartite neighborhood, G cannot contain a complete subgraph of size 4).
(d) For each feasible coloring of S_1 do the following to (G,L^1).
    (1) Update the lists of all remaining vertices to get a palette L^2 and a larger seed S_2. The set S_2 is the largest connected superset of S_1 containing only vertices with lists of size 1.
    (2) By guessing a partial coloring of the graph, obtain an equivalent set of palettes $\mathcal{L}_3$.
    (3) After another iteration, obtain a refined equivalent set of palettes $\mathcal{L}_4$.

(4) For each palette $L \in \mathcal{L}_4$, intersect L with the input palette L^* and obtain a palette L'.
(5) Update, and apply Lemma 11 to check for colorability.

We now describe the individual steps in more detail. The first step as well as the neighborhood reduction can be performed in $O(|V(G^*)|(|V(G^*)|+|E(G^*)|))$ time. The complexity associated to Corollary 5 is $O(|V(G)|^3|E(G)|)$ time. As we report that the graph is not 3-colorable otherwise, we may assume that $G[N(v)]$ is bipartite for every vertex v of G, and that we have obtained a 2-dominating connected set S_1 of G of size at most 3. For technical reasons, if S_1 is a singleton, we add one of its neighbors to S_1. Thus, we can assume that $|S_1| \geq 2$. We will go through all possible 3-colorings of S_1, and check for each whether it extends to a coloring of G which respects the palette L^*. This is clearly enough for deciding whether (G, L^*) is colorable.

So from now on, assume the coloring on S_1 is fixed and that for every other vertex v of G we have $L^1(v) = \{1,2,3\}$. We update the lists of all vertices in G. Note that updating can be done in $O(|V(G)|+|E(G)|)$ time, because each edge vw needs to be checked at most once (either updating v from w or updating w from v). After updating to palette L^2, consider the largest connected set S_2 of vertices with lists of size 1 that contains S_1. We claim that S_2 is a seed for (G, L^2). Indeed, since $\overline{S_1}$ dominates G, so does $\overline{S_2}$. Also, all vertices in $N(S_2)$ must have lists of size 2, since they are adjacent, but do not belong to S_2. So S_2 is a seed.

In the case that two adjacent vertices of S_2 have the same entry on their list, we abort the algorithm for that sub-instance and report that the current 3-coloring of S_1 does not lead to a valid 3-coloring of G.

Claim 7. *For every vertex v in $N(S_2)$ there is an induced path on at least 3 vertices contained in $S_2 \cup \{v\}$ having v as an endpoint.*

Proof. This holds since S_2 is connected, $|S_2| \geq |S_1| \geq 2$, and v is not adjacent to two vertices of S_2 that have different entries on their lists (because $|L^2(v)| = 2$ after updating). ∎

Now, in two steps $j = 3, 4$, we will refine the set of subpalettes of L^1 we are looking at, starting with $\mathcal{L}^2 = \{L^2\}$. At each step we replace the set $\mathcal{L}^{j-1}$ of palettes from the previous step with a set $\mathcal{L}^j$. More precisely, each element L of $\mathcal{L}^j$ is a subpalette of some element $Pred(L)$ of $\mathcal{L}^{j-1}$. We will argue below why it is sufficient to check colorability for the new set of palettes.

For each of the palettes L in $\mathcal{L}^j$, we will define a seed S_L and a set $T_L \subseteq N(S_L)$. We start with $S_{L^2} = S_2$ and T_{L^2} being the set of vertices

$x \in N(S_{L^2})$ for which there does not exist an induced path $x-y-z$ with $|L^2(y)|=3$ and $z \notin \overline{S_{L^2}}$. We will ensure for each palette L that $S_L \supseteq S_{Pred(L)}$ and $T_L \supseteq T_{Pred(L)}$. Furthermore, the seeds S_L and the sets T_L will have the following properties:

(A) for all $x \in N(S_L) \setminus T_L$, there is an induced path $x-y-z$ with $|L(y)|=3$ and $z \notin \overline{S_L}$, and for no $x \in T_L$ is there such a path; and
(B) for each vertex $v \in V(G) \setminus \overline{S_L}$ either $|L(v)|=1$ or $|L(v)|=3$.

Let us now get into the details of the procedure. Successively, for $j = 3,4$, we consider for each $L \in \mathcal{L}^{j-1}$ a set of subpalettes of L obtained by partitioning the possible colorings of induced paths $x-y-z$ with $x \in N(S_L) \setminus T_L$, $|L(y)|=3$ and $z \notin \overline{S_L}$ into a polynomial number of cases. The set $\mathcal{L}^j$ will be the union of all the sets of subpalettes corresponding to lists L in $\mathcal{L}^{j-1}$. The idea is to make the seed grow, and after these two steps, obtain a set of palettes we can deal with, and such that the graph admits a coloring for the original palette if and only if it admits a coloring for one of the palettes in the set.

For each $i \in \{1,2,3\}$, let $\mathcal{P}_i$ be the set of paths $x-y-z$ with $x \in N(S_L) \setminus T_L$, $|L(y)|=3$ and $z \notin \overline{S_L}$, and such that $i \notin L(x)$. We will order the paths of $\mathcal{P}_i$ non-increasingly by $|N(x) \setminus (N(y) \cup N(z) \cup \overline{S_L})|$, i.e., the number of vertices w (if any) such that $w-x-y-z$ is an induced path and $w \notin \overline{S_L}$.

We can compute and sort the paths of $\mathcal{P}_i$ in $O(|V(G)|^4)$ time. Moreover, this order of the paths induces an order on the set Y_i of vertices y that are midpoints of paths $x-y-z$ in $\mathcal{P}_i$. The vertices in Y_i are ordered by their first appearance as midpoints of the ordered paths in $\mathcal{P}_i$. Let $n_i = |Y_i|$, and $Y_i = \{y_{i,1}, \ldots, y_{i,n_i}\}$.

For each $i \in \{1,2,3\}$, we consider the following cases.

(a) All vertices in Y_i are colored i.
(b) There is a k, $1 \le k \le n_i$, such that the first $k-1$ vertices of Y_i are colored i, and the first path $x-y_{i,k}-z$ in $\mathcal{P}_i$ is colored such that the color of $y_{i,k}$ is different from i, the color of every vertex in $W = N(x) \setminus (N(y_{i,k}) \cup N(z) \cup \overline{S_L})$ is i, and the color of z is i if W is empty.
(c) There is a k, $1 \le k \le n_i$, such that the first $k-1$ vertices of Y_i are colored i, and the first path $x-y_{i,k}-z$ in $\mathcal{P}_i$ is colored such that the color of $y_{i,k}$ is different from i, the color of z is different from i if $W = N(x) \setminus (N(y_{i,k}) \cup N(z) \cup \overline{S_L})$ is empty, and if W is nonempty, there is a vertex w of W that gets a color different from i.

In order to do that, we consider all choices of functions $f\colon \{1,2,3\} \to \{a,b,c\}$. For each of these choices, we generate a set $\mathcal{L}_f$ of subpalettes of L, and $\mathcal{L}^j$ will be the union of all sets $\mathcal{L}_f$. For fixed f the first step to obtain $\mathcal{L}_f$ consists of defining $\mathcal{L}_{i,f}$ for $i=1,2,3$ in the following way.

If $\mathcal{P}_i$ is empty, then set $\mathcal{L}_{i,f}=\{L\}$. Otherwise, the set is as follows.

If $f(i)=a$, set $\hat{L}(y)=\{i\}$ for every $y\in Y_i$ and $\hat{L}(v)=L(v)$ for every $v\in V(G)\setminus Y_i$. Set $\mathcal{L}_{i,f}=\{\hat{L}\}$.

If $f(i)\neq a$, for each $k\in\{1,\dots,n_i\}$, let x and z be such that $x-y_{i,k}-z$ is the first path in $\mathcal{P}_i$ having $y_{i,k}$ as midpoint, and let $W=N(x)\setminus(N(y_{i,k})\cup N(z)\cup\overline{S_L})$.

If $f(i)=b$, consider all subpalettes $\hat{L}$ of L which only differ from L on $W\cup\{y_{i,1},\dots,y_{i,k},z\}$, and satisfy $\hat{L}(y_{i,k})=\{i'\}$ for some $i'\neq i$, $\hat{L}(v)=\{i\}$ for all $v\in W\cup\{y_{i,1},\dots,y_{i,k-1}\}$, $|\hat{L}(z)|=1$, and $\hat{L}(z)=\{i\}$ if W is empty. Update these palettes $\hat{L}$ from $y_{i,k}$ except on T_L and let $\mathcal{L}_{i,f}$ be the set of all palettes found in this way, for every choice of k. Note that, in each palette, the updated list of x has size 1, and that the number of palettes generated this way is $O(|V(G)|)$.

If $f(i)=c$, if W is nonempty, for each $w\in W$ consider all subpalettes $\hat{L}$ of L which only differ from L on $\{y_{i,1},\dots,y_{i,k},z,w\}$, and satisfy $\hat{L}(v)=\{i\}$ for all $v\in\{y_{i,1},\dots,y_{i,k-1}\}$, $|\hat{L}(y_{i,k})|=|\hat{L}(z)|=|\hat{L}(w)|=1$, $\hat{L}(y_{i,k})\neq\{i\}$, and $\hat{L}(w)\neq\{i\}$. If W is empty, consider all subpalettes $\hat{L}$ of L which only differ from L on $\{y_{i,1},\dots,y_{i,k},z\}$, and satisfy $\hat{L}(v)=\{i\}$ for $v\in\{y_{i,1},\dots,y_{i,k-1}\}$, $|\hat{L}(y_{i,k})|=|\hat{L}(z)|=1$, $\hat{L}(y_{i,k})\neq\{i\}$, and $\hat{L}(z)\neq\{i\}$. Update these palettes $\hat{L}$ from $y_{i,k}$ except on T_L and let $\mathcal{L}_{i,f}$ be the set of all palettes found in this way, for every choice of k and of w (if such a w exists). Note that again, in each palette, the updated list of x has size 1, and that the number of palettes generated this way is $O(|V(G)|^2)$.

Finally, for each triple $(L_1,L_2,L_3)\in\mathcal{L}_{1,f}\times\mathcal{L}_{2,f}\times\mathcal{L}_{3,f}$ consider the palette $\tilde{L}$ obtained from intersecting the lists of L_1,L_2,L_3, taking intersections at each vertex. Update the palette $\tilde{L}$ from any vertex in S_L, except on T_L. Let $\mathcal{L}_f$ be the set of all palettes $\tilde{L}$ thus generated.

Observe that $|\mathcal{L}_f|=O(|V(G)|^6)$, since $|\mathcal{L}_{i,f}|=O(|V(G)|^2)$ for $i=1,2,3$.

For each $L'\in\mathcal{L}_f$, let $S_{L'}$ be a maximal connected set of vertices with list size 1 that contains S_L. Then $S_{L'}$ is a seed.

Note that for each $L'\in\mathcal{L}_f$, all vertices v in T_L satisfy $|L'(v)|=2$, since they were never updated. Let $T_{L'}$ be the union of T_L with all vertices $x\in N(S_{L'})$ which are not the starting point of an induced path $x-y-z$ with $|L'(y)|=3$ and $z\notin\overline{S_{L'}}$.

Clearly, $T_{L'}\subseteq N(S_{L'})$. Property (A) holds because of the way we defined $T_{L'}$, and because there are no new paths of the type described in (A) that start at vertices in T_L, as seeds grow and lists shrink. Property (B) holds because S_L was a seed satisfying Properties (A) and (B), and when defining palettes in $\mathcal{L}_f$ by the cases (a), (b), and (c), we have reduced the size of some vertex lists from 3 to 1, never to 2; then we only updated *from vertices* in S_L except on T_L, thus every vertex that got a list of size 1 by updating is connected to S_L by a path all whose vertices have current lists of size one, and is now in $S_{L'}$ and, consequently, every vertex that got a list of size 2 by updating is in $N(S_{L'})$.

Claim 8. *There is a coloring of G for the palette L^2 if and only if G has a coloring for at least one of the palettes in $\mathcal{L}^4$.*

Proof. Indeed, observe that when obtaining $\mathcal{L}^j$ from $\mathcal{L}^{j-1}$, we consider for each $L \in \mathcal{L}^{j-1}$ and for each $i \in \{1,2,3\}$ the possibility that all induced 3-vertex-paths that start in $N(S_L)$ and then leave $\overline{S_L}$ have their second vertex colored i (when $f(i)=a$). We also consider the possibility that there is such a path whose second vertex is colored with a different color (when $f(i)=b$ or $f(i)=c$). In that case, we consider separately the possible colorings of a fourth vertex w, if such a w exists. ∎

Note that $|\mathcal{L}^{j+1}| = O(|\mathcal{L}^j| \cdot |V(G)|^6)$ for each $j=2,3$. Since $|\mathcal{L}^2|=1$, the number of palettes in $\mathcal{L}^4$ is $O(|V(G)|^{12})$.

Next we prove that during the above described process, the union of our seed with the set T_L actually grows.

Claim 9. *For each $L \in \mathcal{L}^j$, we have $N(S_{Pred(L)}) \subset S_L \cup T_L$.*

Proof. Let $L' = Pred(L)$ and let f be the function used to produce L from L'. In order to see Claim 9, suppose there is a vertex $x' \in N(S_{L'}) \setminus (S_L \cup T_L)$. As $x' \notin S_L$ and $S_L \supseteq S_{L'}$, we know that $x' \in N(S_L)$. Furthermore, since $x' \notin T_L$, there is an induced path $x' - y' - z'$ with $|L(y')|=3$ and $z' \notin \overline{S_L}$. In particular, $z' \notin \overline{S_{L'}}$ and since lists only shrink, $|L'(y')|=3$. So $f(i) \neq a$, where i is such that $i \notin L(x') = L'(x')$. Thus $f(i) \in \{b,c\}$, and so there is an induced path $x - y - z$ with $x \in N(S_{L'})$, $y, z \notin N(S_{L'})$, $L(x) \neq \{i\}$, $L(y) \neq \{i\}$, and $|L(x)| = |L(y)| = |L(z)| = 1$, and thus $x, y, z \in S_L$. Since $y', z' \notin \overline{S_L}$, it follows that there are no edges between $\{y', z'\}$ and $\{x, y, z\}$. Also, since $x' \in N(S_L)$, there are no edges from x' to vertices $v \in \{x, y, z\}$ with $L(v) \subsetneq L(x')$. In other words, the only possible edge between $\{x, y, z\}$ and $\{x', y', z'\}$ is $x'z$, and if this edge is present, we have that $L(z) = \{i\}$. On the other hand, by (1), there is a path Q of at least 3 vertices connecting x and x' whose interior lies in $S_{L'}$ (in particular, the interior of Q is anticomplete to $\{y', z', y, z\}$). So, since G is P_7-free, the edge $x'z$ has to be present and thus we have $L(z) = \{i\}$.

Now, assume there is an extension of x–y–z to an induced path w–x–y–z with $w \notin \overline{S_{L'}}$. Then, as the sequence w–x–y–z–x'–y'–z' is not an induced P_7, there is an edge from w to one of x', y', z'. Observe if $|L(w)| = 1$, then $w \in S_L$ and neither wy' nor wz' is an edge. Hence either $|L(w)| \geq 2$, or $L(w) = \{i\}$, and in the latter case the only edge from w to $\{x', y', z'\}$ is wx'. As this happens for all possible choices of w, we see that $f(i) \neq c$, and thus $f(i) = b$. This means that for all possible w, w is adjacent to x'. But now, observe that

$$N(x) \setminus (N(y) \cup N(z) \cup \overline{S_{L'}}) \subsetneq N(x') \setminus (N(y') \cup N(z') \cup \overline{S_{L'}}),$$

since z is in the right hand side set, but not in the left hand side set. This is a contradiction to the choice of the path $x-y-z$ for the definition of L from L' and f.

We conclude that there is no extension of $x-y-z$ to an induced path $w-x-y-z$. But then, the fact that $L(z) = \{i\}$ implies that again, $f(i) \neq c$, and thus, $f(i) = b$. The existence of the edge $x'z$ gives a contradiction to the choice of the path x–y–z for the definition of L from L' and f. This proves Claim 9. ∎

Next, we prove that two steps of performing the above procedure suffice to take care of all paths on three vertices that start in the boundary of the current seed, and then leave the seed.

Claim 10. *For each $L \in \mathcal{L}^4$, we have $N(S_L) \subset T_L$.*

Proof. Suppose there are $L \in \mathcal{L}^4$ and $x \in N(S_L)$ such that $x \notin T_L$. Then by (A) there is a path $x-y-z$ with $|L(y)| = 3$ and $z \notin \overline{S_L}$. Clearly y and z are anticomplete to S_L. Let $L' = Pred(L)$ and $L'' = Pred(L')$. Choose an induced path P from x to $N(S_{L''})$ with all vertices but x in S_L, say it ends in $x'' \in N(S_{L''})$. By Claim 9, $N(S_{L''}) \subseteq S_{L'} \cup T_{L'}$. On the other hand, as $T_{L'} \cap S_L = \emptyset$, $x'' \in S_{L'}$. In particular, $x \neq x''$.

Let x_1 be the neighbor of x in P. Since $x \notin S_L \cup T_L$, by Claim 9, $x_1 \in S_L \setminus S_{L'}$. As the subpath of P from x_1 to x'' goes from $S_L \setminus S_{L'}$ to $S_{L'}$, it contains a vertex x' in $N(S_{L'})$. The vertex x' may be x_1, but x' is different from x'' because $x'' \in S_{L'}$. As x' is in the subpath from x_1 to x'', $x' \neq x$. Summing up, x, x' and x'' are three distinct vertices, and so P together with the path $x-y-z$ and the path provided by Claim 7 for x'' gives a path on at least 7 vertices, a contradiction. ∎

By Claim 8, (G, L^2) is colorable if and only if (G, L) is colorable for some $L \in \mathcal{L}^4$. For each $L \in \mathcal{L}^4$ our aim is to check whether there is a coloring of (G, L). This we will do, after some more discussion, with the help of

Lemma 11 below. So from now on, let $L \in \mathcal{L}^4$ be fixed. Let X be the set of all vertices in $V(G) \setminus \overline{S_L}$ with lists of size 1, and set $Y = V(G) \setminus (\overline{S_L} \cup X)$. By construction, $|L(y)| = 3$ for each $y \in Y$.

By Claim 10, no vertex of $N(S_L)$ is the starting point of an induced path $x - y - z$ with $y \in Y$ and $z \in X \cup Y$. In other words, for each $y \in Y$, all edges between $N(y) \cap \overline{S_L}$ and $N(y) \setminus \overline{S_L}$ are present.

Now we intersect L with the given input palette L^*, and then update. Let L' be the resulting palette. We may assume that $|L'(v)| \geq 1$ for all $v \in V(G)$, otherwise we may safely report that (G, L') is not colorable, and thus L does not lead to a feasible coloring of (G, L^*). Let Y' be the set of vertices y of Y such that $|L'(y)| = 3$. We noticed that for each $y \in Y$, all the edges between $N(y) \cap \overline{S_L}$ and $N(y) \setminus \overline{S_L}$ are present. Since Y' is a subset of the vertices v such that $|L^*(v)| = 3$ and we have applied the neighborhood reduction at the beginning of the algorithm and the graph did not change, for $y \in Y'$ one of the sets $N(y) \cap \overline{S_L}$ or $N(y) \setminus \overline{S_L}$ must be empty. Since $\overline{S_L} \supseteq \overline{S_2}$ is a dominating set, we conclude that $N(y) \setminus \overline{S_L} = \emptyset$, and thus

$$N(y) \subseteq \overline{S_L} \quad \text{for each} \quad y \in Y'. \tag{2}$$

Consider the set S' of all vertices that are connected to S_L by a (possibly trivial) path containing only vertices with lists L' of size 1. Note that S' is a seed. In particular, $S_L \subseteq S'$ and by (2), we have $N(y) \subseteq \overline{S'}$ for every $y \in Y'$. That is, Y' is a stable set anticomplete to $V(G) \setminus (\overline{S'} \cup Y')$.

We are now in a situation where the following lemma applies, solving the remaining problem.

Lemma 11. *Let G be a connected P_7-free graph with a palette L. Let S be a seed of G such that if $v \in S$ and $w \in N(S)$ are adjacent, then they do not share list entries. Assume that the set X of vertices having lists of size 3 is stable and anticomplete to $V(G) \setminus (\overline{S} \cup X)$. Assume also that no vertex in X has a connected neighborhood. Then we can decide whether G has a coloring for L in $O(|V(G)|^9(|V(G)| + |E(G)|))$ time.*

The next subsection is devoted to the proof of Lemma 11. Since we have $|\mathcal{L}^4| = O(|V(G)|^{12})$ many lists to consider, and need to apply Lemma 11 to each of these, the total running time of the whole algorithm amounts to $O(|V(G)|^{21}(|V(G)| + |E(G)|))$.

3.1. Proof of Lemma 11

Let G, L, S and X be as in the statement of Lemma 11.

In this proof we make extensive use of the concept of monochromatic set constraints as defined in Section 2. Note that (G,L) is colorable if and only if the triple $(G,L,\emptyset)$ is colorable. Our aim is to define a set $\mathcal{R}$ of restrictions of $(G,L,\emptyset)$ with the property that in any element of $\mathcal{R}$ there are no vertices with list of size 3, and $(G,L,\emptyset)$ is colorable if and only if $\mathcal{R}$ is colorable. Moreover, $\mathcal{R}$ has polynomial size and is computable in polynomial time.

If $X=\emptyset$, we simply let $\mathcal{R}=\{(G,L,\emptyset)\}$. Otherwise, for all $i=1,2,3$, let D_i be the set of vertices $v\in N(S)$ with $L(v)=\{1,2,3\}\setminus\{i\}$, and for $x\in X$, let $N_i(x)=N(x)\cap D_i$, for $i=1,2,3$. Observe that, under the hypothesis of Lemma 11, for every $d\in D_i$ and for every $s\in S\cap N(d)$, we have $L(s)=\{i\}$. By the same hypothesis, no vertex of X has neighbors in S.

If $N(x)\subseteq D_i$ for some $x\in X$ and some $i\in\{1,2,3\}$, then setting $L(x)=\{i\}$ does not change the colorability of $(G,L,\emptyset)$, so we may assume that for every $x\in X$ at least two of the sets $N_1(x),N_2(x),N_3(x)$ are non-empty. Let X_1 be the set of vertices $x\in X$ for which $N_2(x)$ is not complete to $N_3(x)$; for every $x\in X_1$ fix two vertices $n_2(x)\in N_2(x)$ and $n_3(x)\in N_3(x)$ such that $n_2(x)$ is non-adjacent to $n_3(x)$. Define similarly X_2 and $n_1(x),n_3(x)$ for every $x\in X_2$, and X_3 and $n_1(x),n_2(x)$ for every $x\in X_3$. Since no vertex of X has a connected neighborhood and X is a stable set and anticomplete to $V(G)\setminus(\overline{S}\cup X)$, it follows that $X=X_1\cup X_2\cup X_3$.

Before we state the coloring algorithm, we need some auxiliary statements. For a path P with ends u,v let $P^*=V(P)\setminus\{u,v\}$ denote the interior vertices of P.

Claim 12. *Let $i,j\in\{1,2,3\}$, $i\neq j$, and let $u_i,v_i\in D_i$ and $u_j,v_j\in D_j$, such that $\{u_i,v_i,u_j,v_j\}$ is a stable set. Then there exists an induced path P with ends $a,b\in\{u_i,v_i,u_j,v_j\}$ such that*
(a) $\{a,b\}\neq\{u_i,u_j\}$ and $\{a,b\}\neq\{v_i,v_j\}$,
(b) P^* is contained in S and, in particular, $|L(v)|=1$ for every $v\in P^*$, and
(c) P^* is anticomplete to $\{u_i,v_i,u_j,v_j\}\setminus\{a,b\}$.

Proof. Note that each of u_i,u_j,v_i,v_j has a neighbor in S, and $G[S]$ is connected. Let P be an induced path with $P^*\subseteq S$ that connects u_i with v_i. If P is not as desired, at least one of u_j,v_j has a neighbor on P. Let p be the neighbor of u_j or v_j on P that is closest to v_i; by symmetry we may assume p is a neighbor of u_j. Note that p is not adjacent to $u_i,v_i\in D_i$, because p is already adjacent to $u_j\in D_j$. Hence, if $u_j-p-P-v_i$ is not as desired, then v_j must have a neighbor on $p-P-v_i$. Among all such neighbors, let p' be the one that is closest to p (possibly $p'=p$). As before, p' is not adjacent to any of $u_i,v_i\in D_i$, and thus, $u_j-p-P-p'-v_j$ is the desired path. ∎

Claim 13. *Let* $\{i,j,k\} = \{1,2,3\}$. *Let* $x,y \in X_i$, *let* $n_j \in N_j(x)$ *and* $n_k \in N_k(x)$ *such that* n_j *is non-adjacent to* n_k. *Then there is an edge between* $\{x,n_j,n_k\}$ *and* $\{y,n_j(y),n_k(y)\}$.

Proof. Assume there is no such edge. Then in particular, vertices $n_j,n_j(y),n_k,n_k(y)$ are distinct, and we can apply Claim 12 to obtain a path P with $P^* \subseteq S$ that connects two vertices from $\{n_j,n_j(y),n_k,n_k(y)\}$ in way that P^*, together with n_j-x-n_k and $n_j(y)-y-n_k(y)$, forms an induced path of length at least 7, a contradiction. ∎

Next we distinguish between several types of colorings of G, and show how to reduce the list sizes assuming that a coloring of a certain type exists. For this, let $\{i,j,k\}=\{1,2,3\}$. We call a coloring c of a restriction (G',L',Z') of (G,L,Z)

(A) a *type A coloring with respect to* i if there exists an induced path $n_j–x–n_k–z–m_j$ with $x,z \in X_i$, $n_j \in N_j(x)$, $m_j \in N_j(z)$, and $n_k \in N_k(x) \cap N_k(z)$ such that $c(n_j)=i$, $c(x)=j$ and $c(z)=k$ (this implies $c(n_k)=c(m_j)=i$), or the same with the roles of j and k reversed;

(B) a *type B coloring with respect to* i if it is not a type A coloring with respect to i, and there exists an induced path $x-n_k-z-m_j$ with $x,z \in X_i \cap V(G')$, $m_j \in N_j(z)$, $n_k \in N_k(x) \cap N_k(z)$ such that $c(x)=j$ and $c(z)=k$ (this implies $c(n_k)=c(m_j)=i$), or the same with the roles of j and k reversed;

(C) a *type C coloring with respect to* i if it is not a type A or type B coloring, and there exist $z \in X_i \cap V(G')$, $m_j \in N_j(z)$ and $n_k \in N_k(z)$ such that $c(m_j)=c(n_k)=i$.

We will show in Claim 14 how to refine the instances to test if a graph admits a type A coloring with respect to a color i; in Claim 15 how to refine the instances to test if a graph admits a type B coloring with respect to i under the assumption that it does not admit a type A coloring with respect to i; in Claim 16 how to refine the instances to test if a graph admits a type C coloring with respect to i under the assumption that it does not admit a type A or type B coloring with respect to i; finally, in Claim 17 we show how to refine the instances to test if a graph admits a coloring under the assumption that it does not admit a type A, or type B, or type C coloring with respect to i. After the claims, we describe how to combine them in order to obtain the desired list of restrictions of the original instance.

Claim 14. *Let* (G',L',Z') *be a restriction of* (G,L,Z). *There exists a set* $\mathcal{L}_i$ *of* $O(|V(G)|^3)$ *subpalettes of* L' *such that*
(a) $|L''(v)| \leq 2$ *for every* $L'' \in \mathcal{L}_i$ *and* $v \in X_i \cap V(G')$, *and*

(b) (G',L',Z') *admits a type A coloring with respect to* i *if and only if* $(G',\mathcal{L}_i,Z')$ *is colorable.*

Moreover, $\mathcal{L}_i$ *can be constructed in* $O(|V(G)|^4)$ *time.*

For every $x,z \in X_i \cap V(G')$ and $n_j \in N_j(x)$ for which there are $n_k \in N_k(x) \cap N_k(z)$ and $m_j \in N_j(z)$ such that $n_j - x - n_k - z - m_j$ is an induced path, we construct a palette $L'' = L_{x,z,n_j}$ depending on x,z,n_j; for the same case with triples $x,z \in X_i \cap V(G')$, $n_k \in N_k(x)$, and the roles of j and k reversed, we construct in an analogous way a palette $L'' = L'_{x,z,n_k}$ depending on x,z,n_k. The set $\mathcal{L}_i$ will be the set of all palettes L'' obtained in this way. So the number of palettes in $\mathcal{L}_i$ is $O(|V(G)|^3)$.

For x,z,n_j as above (we will assume the first case in the definition, the other case is analogous), we define L'' by setting $L''(x) = \{j\}$, $L''(z) = \{k\}$, $L''(n_j) = \{i\}$, and leaving $L''(v) = L'(v)$ for all $v \in V(G') \setminus \{x,z,n_j\}$. Update $N_j(z)$ from z, and $N_k(x)$ from x. Let n_k and m_j be such that $n_k \in N_k(x) \cap N_k(z)$, $m_j \in N_j(z)$, and $n_j – x – n_k – z – m_j$ is an induced path. Note that after updating, $L''(n_k) = L''(m_j) = \{i\}$. Now, for each vertex $v \in D_j \cup D_k$ that has a neighbor $v' \in \{x,z,n_j,n_k,m_j\}$, update v from each such neighbor v'. Next, for every vertex $y \in X_i \cap V(G')$, if $n_j(y)$ or $n_k(y)$ now has list size 1, then update y from both $n_j(y)$ and $n_k(y)$, and also update y from m_j, n_j and n_k in the case that y is adjacent to any of them. Call the obtained palette L'' (slightly abusing notation). By the way we updated, it only takes $O(|V(G)|)$ time to compute this palette. The total time for constructing all palettes for $\mathcal{L}_i$ thus amounts to $O(|V(G)|^4)$.

In order to see Claim 14 (a), we need to show that $|L''(y)| \leq 2$ for all $y \in X_i \cap V(G')$. For contradiction, suppose $|L''(y)| = 3$ for some $y \in X_i \cap V(G')$. By Claim 13, there must be edges between $\{x,n_j,n_k\}$ and $\{y,n_j(y),n_k(y)\}$, and also between $\{z,m_j,n_k\}$ and $\{y,n_j(y),n_k(y)\}$. By the way we updated L'', the only possibly edges between these sets are those connecting $n_j(y)$ with x, and $n_k(y)$ with z. Consequently, $n_j(y)x$ and $n_k(y)z$ are both edges, and so $m_j - z - n_k(y) - y - n_j(y) - x - n_j$ is a P_7, a contradiction.

For Claim 14 (b), first note that by construction, if $(G',\mathcal{L}_i,Z')$ is colorable then (G',L',Z') has a type A coloring with respect to i. On the other hand, if c is a type A coloring of (G',L',Z') with respect to i, then there is an induced path $n_j – x – n_k – z – m_j$ with $x,z \in X_i$, $n_j, m_j \in N_j(x)$, and $n_k \in N_k(x)$ such that $c(n_j) = c(m_j) = c(n_k) = i$, $c(x) = j$, and $c(z) = k$ (or the same with the roles of j and k reversed). Since updating does not change the set of possible colorings for a list, c is a coloring for the list $L'' = L_{x,z,n_j}$ (respectively, $L'' = L_{x,z,n_k}$). So $\mathcal{L}_i$ is as required for Claim 14 (b). ∎

Claim 15. *Let (G',L',Z') be a restriction of (G,L,Z) that does not admit a type A coloring. There exists a set $\mathcal{L}_i$ of $O(|V(G)|^2)$ subpalettes of L' such that*
(a) $|L''(v)| \leq 2$ for every $L'' \in \mathcal{L}_i$ and $v \in X_i \cap V(G')$, and
(b) (G',L',Z') admits a type B coloring with respect to i if and only if $(G',\mathcal{L}_i,Z')$ is colorable.
Moreover, $\mathcal{L}_i$ can be constructed in $O(|V(G)|^3)$ time.

Proof. For every $x,z \in X_i \cap V(G')$ for which there exist $n_k \in N_k(x) \cap N_k(z)$ and $m_j \in N_j(z)$ such that $x-n_k-z-m_j$ is an induced path, we construct a palette $L'' = L_{x,z}$, depending on x and z. For the case with the roles of j and k reversed, we construct analogously a palette $L'' = L'_{x,z}$. The set $\mathcal{L}_i$ will be the set of all palettes L'' obtained in this way. So the number of palettes in $\mathcal{L}_i$ is $O(|V(G)|^2)$.

Given a pair of vertices x, z in $X_i \cap V(G')$ satisfying the hypothesis, let n_k and m_j such that $n_k \in N_k(x) \cap N_k(z)$, $m_j \in N_j(z)$, and $x-n_k-z-m_j$ is an induced path. Let M be the set of all $n \in N_j(x)$ for which $n-x-n_k-z-m_j$ is an induced path.

Define L'' by setting $L''(x) = \{j\}$, $L''(z) = \{k\}$, $L''(n_k) = L''(m_j) = \{i\}$, and $L''(n) = \{k\}$ for all $n \in M$, and leaving $L''(v) = L'(v)$ for all $v \in V(G') \setminus (\{x,z,n_k,m_j\} \cup M)$. Now, for each vertex $v \in D_j \cup D_k$ that has a neighbor v' in $\{x,z,m_j,n_k\}$, update v from each such neighbor v'. Next, for every vertex $y \in X_i \cap V(G')$, if $n_j(y)$ or $n_k(y)$ now has list size 1, then update y from both $n_j(y)$ and $n_k(y)$, and also update y from m_j and n_k in the case that y is adjacent to either of them. Call the obtained palette L''. Note that by the way we updated, it takes $O(|V(G)|)$ time to compute this palette. The total time for constructing all palettes for $\mathcal{L}_i$ thus amounts to $O(|V(G)|^3)$.

In order to see Claim 15 (a), we need to show that $|L''(y)| \leq 2$ for all $y \in X_i \cap V(G')$. For contradiction, suppose $|L''(y)| = 3$ for some $y \in X_i \cap V(G')$. Then $n_j(y) \notin M \cup \{m_j\}$ and $n_k(y) \neq n_k$. By Claim 13, it follows that $n_k(y)$ is adjacent to z, and by the way we updated L'', the only other possible edge between $\{x,n_k,z,m_j\}$ and $\{y,n_j(y),n_k(y)\}$ would be $xn_j(y)$. However, since $n_j(y) \notin M$, we deduce that $n_j(y)$ is non-adjacent to x. Let s be a neighbor of $n_j(y)$ in S with $L(s) = \{j\}$. Then s is anticomplete to $\{n_k,x,y,z,n_k(y)\}$. So $x-n_k-z-n_k(y)-y-n_j(y)-s$ is a P_7, a contradiction.

For Claim 15 (b), note that by construction, if $(G',\mathcal{L}_i,Z')$ is colorable then (G',L',Z') has a type B coloring with respect to i. On the other hand, if c is a type B coloring of (G',L',Z') with respect to i, then there is an induced path $x–n_k–z–m_j$ with $x,z \in X_i$, $m_j \in N_j(x)$, and $n_k \in N_k(x) \cap N_k(z)$ such that $c(m_j) = c(n_k) = i$, $c(x) = j$, and $c(z) = k$ (or the same with the roles of j and k reversed). Since c is not a type A coloring, it follows that $c(v) = k$

for all v in M. Since updating does not change the set of possible colorings for a list, c is a coloring for $L''=L_{x,z}$. So $\mathcal{L}_i$ is as required for Claim 15 (b). ∎

Claim 16. *Let (G',L',Z') be a restriction of (G,L,Z) that does not admit a type A or type B coloring. There exists a set $\mathcal{L}_i$ of $O(|V(G)|^2)$ subpalettes of L' such that*
(a) $|L''(v)|\leq 2$ for every $L''\in\mathcal{L}_i$ and $v\in X_i\cap V(G')$, and
(b) (G',L',Z') admits a type C coloring with respect to i if and only if $(G',\mathcal{L}_i,Z')$ is colorable.
Moreover, $\mathcal{L}_i$ can be constructed in $O(|V(G)|^4)$ time.

Proof. For every $z\in X_i\cap V(G')$ having non-adjacent neighbors $m_j\in N_j(z)$ and $n_k\in N_k(z)$, we construct two families of palettes, one for each of the possible colors j,k of z in a type C coloring, z, m_j, n_k are as in the definition of a type C coloring. We only describe how to obtain the family of palettes L'' with $L''(z)=\{k\}$; the definition of the family of palettes L'' with $L''(z)=\{j\}$ is analogous, with the roles of j and k reversed.

Let N_z be the set of vertices n_k in $N_k(z)$ having a non-neighbor in $N_j(z)$. For each such vertex n_k, let $W=W_{z,n_k}$ be the set of all $w\in X_i\cap V(G')$ such that there exists an induced path $w-n_k-z-m_j$ with $m_j\in N_j(z)$. We will order the vertices of N_z non-increasingly by $|W|$. We can compute and sort the vertices of N_z in $O(|V(G)|^3)$ time.

For each $n_k\in N_z$, define $L''=L_{z,n_k}$ by setting $L''(z)=L''(w)=\{k\}$ for all $w\in W$, $L''(n_k)=\{i\}$, $L''(n'_k)=\{j\}$ for every $n'_k\in N_z$ having an index lower than the index of n_k in N_z, and leaving $L''(v)=L'(v)$ for all the remaining vertices. Update each vertex of $N_j(z)$ from z. Now, for each vertex v that has a neighbor in $\{z\}\cup N_k(z)\cup N_j(z)\cup W$, update v from each such neighbor v'. Next, for every vertex $y\in X_i\cap V(G')$, if $n_j(y)$ or $n_k(y)$ now has list size 1, then update y from both $n_j(y)$ and $n_k(y)$. Call the obtained palette L''. Note that by the way we updated, it takes $O(|V(G)|^2)$ time to compute this palette. The number of palettes L_{z,n_k} is $O(|V(G)|^2)$, and the same for the case with the roles of j and k reversed. Then $\mathcal{L}_i$, the set of all palettes obtained in this way, has cardinality $O(|V(G)|^2)$, and can be constructed in $O(|V(G)|^4)$ time. We may assume that $|L''(v)|\geq 1$ for all $v\in V(G')$, otherwise we detect that the palette L'' does not lead to a feasible solution to L'.

In order to see Claim 16 (a), we need to show that $|L''(y)|\leq 2$ for all $y\in X_i\cap V(G')$. For contradiction, suppose $|L''(y)|=3$ for some $y\in X_i\cap V(G')$. Let m_j be a non-neighbor of n_k in $N_j(z)$. Note that by the way we updated, $L''(m_j)=\{i\}$. Claim 13 guarantees an edge between $\{z,m_j,n_k\}$ and $\{y,n_j(y),n_k(y)\}$. By the way we updated L'', $n_j(y)\neq m_j$, $n_k(y)\neq n_k$, z is not adjacent to $n_j(y)$, and there is no edge between $\{m_j,n_k\}$ and

$\{y,n_j(y),n_k(y)\}$. So z is adjacent to $n_k(y)$. Since $n_k(y)$ is not adjacent to m_j, $n_k(y)$ belongs to N_z, and as it has two colors in its list L'', its index is greater than the index of n_k in N_z. As y is adjacent to $n_k(y)$ and not to $\{m_j,n_k\}$, $y \in W_{z,n_k(y)} \setminus W_{z,n_k}$. Since $|W_{z,n_k}| \geq |W_{z,n_k(y)}|$, there is a vertex $x \in W_{z,n_k} \setminus W_{z,n_k(y)}$. By definition, $L''(x) = \{k\}$, thus x is not adjacent to $\{y,n_j(y)\}$. Let s be a neighbor of $n_j(y)$ in S with $L(s) = \{j\}$. Then s is anticomplete to $\{n_k,x,y,z,n_k(y)\}$. So $x-n_k-z-n_k(y)-y-n_j(y)-s$ is a P_7, a contradiction.

For Claim 16 (b), note that by construction, if $(G',\mathcal{L}_i,Z')$ is colorable then (G',L',Z') has a type C coloring with respect to i. On the other hand, if c is a type C coloring of (G',L',Z') with respect to i, then there is a path n_k-z-m_j with $z \in X_i \cap V(G')$, $m_j \in N_j(z)$, $n_k \in N_k(z)$, and $c(m_j)=c(n_k)=i$. Assume $c(z) = k$ (the case $c(z) = j$ is analogous), and consider the path $n_k - z - m_j$ that minimizes the index of n_k in N_z. Since $c(m'_j)=i$ for every m'_j in $N_j(z)$, it follows that $c(n'_k)=j$ for every $n'_k \in N_z$ having a lower index than the index of n_k in N_z.

Since c is not a type A or B coloring, for every vertex $w \in W_{z,n_k}$ we have $c(w) = c(z) = k$. Since updating does not change the set of possible colorings for a list, c satisfies the palette $L''=L_{z,n_k}$. So $\mathcal{L}_i$ is as required for Claim 16 (b). ∎

Claim 17. *Let (G',L',Z') be a restriction of (G,L,Z). Assume that (G',L',Z') does not admit a type A, type B, or type C coloring with respect to i (i.e., no coloring with a vertex x of $X_i \cap V(G')$ having neighbors colored i both in $N_j(x)$ and $N_k(x)$). Let Y_i be the set of vertices $x \in X_i \cap V(G')$ such that $N_i(x)=\emptyset$, and let $Z_i = \bigcup_{y \in Y_i}\{N_j(y),N_k(y)\}$. Then (G',L',Z') is colorable if and only if $(G' \setminus Y_i, L', Z' \cup Z_i)$ is colorable, and any 3-coloring of $(G' \setminus Y_i, L', Z' \cup Z_i)$ can be extended to a 3-coloring of (G',L',Z') in $O(|V(G)|)$ time.*

Proof. It is enough to prove that for every coloring c of (G',L',Z') and every $x \in X_i \cap V(G')$ such that $N_i(x) = \emptyset$, the sets $N_j(x)$ and $N_k(x)$ are monochromatic with respect to c. Supposing this is false, we may assume that for some coloring c there are vertices $u,v \in N_j(x)$ with $c(u) = i$ and $c(v)=k$. Since there are no type A or type B colorings and c is not of type C, it follows that $c(w)=j$ for every $w \in N_k(x)$. But then x has neighbors of all three colors, contrary to the fact that c is a coloring. ∎

Let $Z = \emptyset$. Recall that our aim was to define a set $\mathcal{R}$ of restrictions of (G,L,Z) with the property that in any element of $\mathcal{R}$ there are no vertices with list of size 3, and such that (G,L,Z) is colorable if and only if $\mathcal{R}$ is colorable. We now construct $\mathcal{R}$ as follows. Apply Claims 14, 15, 16 and 17 with $i=1$ to (G,L,Z) to create sets $\mathcal{R}_2,\dots,\mathcal{R}_5$, each consisting of $O(|V(G)|^3)$ restrictions of (G,L,Z). For every $x\in X_1$ and every $(G',L',Z')\in\mathcal{R}_2\cup\mathcal{R}_3\cup\mathcal{R}_4$, we have that $|L'(x)|\le 2$. For $(G',L',Z')\in\mathcal{R}_5$, if $x\in X_1$ and $|L'(x)|=3$, then $N_j(x)\neq\emptyset$ for every $j\in\{1,2,3\}$. Repeat this with $i=2$ for every restriction in $\mathcal{R}_2\cup\mathcal{R}_3\cup\mathcal{R}_4\cup\mathcal{R}_5$, and then again with $i=3$ for every restriction obtained with $i=2$. This creates a set $\mathcal{R}'$ of $O(|V(G)|^9)$ restrictions. Finally, we construct $\mathcal{R}$ from $\mathcal{R}'$ by removing all restrictions that still contain lists which have size three for some vertex. Following Claims 14, 15, 16 and 17, the whole computation can be done in $O(|V(G)|^9\cdot|V(G)|)=O(|V(G)|^{10})$ time.

Let us say that $x\in X$ is *wide* if $N_1(x)\neq\emptyset, N_2(x)\neq\emptyset$ and $N_3(x)\neq\emptyset$. Due to the construction of $\mathcal{R}'$ it holds that if $|L'(x)|=3$ for some $(G',L',Z')\in\mathcal{R}'$, then x is wide.

It remains to show that (G,L,Z) is colorable if and only if $\mathcal{R}$ is colorable. By Claims 14, 15, 16 and 17, we know that if $\mathcal{R}$ is colorable then (G,L,Z) is colorable. Now assume that (G,L,Z) is colorable, and let c be a coloring of (G,L,Z). Consider any wide vertex x. Since the neighborhood of x can only have two distinct colors in total, there are two vertices $n_j\in N_j(x)$ and $n_k\in N_k(x)$ such that $c(n_j)=c(n_k)=i$, for some distinct $j,k\in\{1,2,3\}$. Then c is a type A, type B, or type C coloring with respect to i. Consequently, by Claims 14, 15, 16, there is a restriction $(G',L',Z')\in\mathcal{R}'$ that is colorable, and where for every wide vertex $y\in X\cap V(G')$ it holds that $|L'(y)|\le 2$. Therefore $(G',L',Z')\in\mathcal{R}$.

We now come to the running time analysis. Using Lemma 3, we can check in $O(|V(G)|+|E(G)|)$ time whether a given restriction (G',L',Z') of (G,L,Z) is colorable. Since we have $O(|V(G)|^9)$ many restrictions to consider, and these can be computed in $O(|V(G)|^{10})$ time, the total running time amounts to $O(|V(G)|^9(|V(G)|+|E(G)|))$. This completes the proof.

Acknowledgements. We are grateful to Paul Seymour, Juraj Stacho, and Daniël Paulusma for many fruitful discussions. We also thank Alex Scott for telling us about the problem that this paper solves. We thank four anonymous reviewers for their very careful reading and helpful comments.

References

[1] B. Aspvall, M. F. Plass and R. E. Tarjan: A linear-time algorithm for testing the truth of certain quantified boolean formulas, *Information Processing Letters* **8** (1979), 121–123.

[2] E. Camby and O. Schaudt: A new characterization of P_k-free graphs, *Algorithmica*, 2014. To appear.

[3] M. Chudnovsky, P. Maceli, J. Stacho and M. Zhong: 4-coloring P_6-free graphs with no induced 5-cycles, arXiv:1407.2487v1 [cs.DM], July 2014.

[4] D. Coppersmith and S. Winograd: Matrix multiplication via arithmetic progressions *J. Symbolic Computation* **9** (1990), 251–280.

[5] D. Corneil, Y. Perl and L. Stewart: Cographs: recognition, applications and algorithms, *Congressus Numerantium* **43** (1984), 249–258.

[6] K. Edwards. The complexity of colouring problems on dense graphs, *Theoretical Computer Science* **43** (1986), 337–343.

[7] P. Erdős, A. Rubin and H. Taylor: Choosability in graphs, *Congressus Numerantium* **26** (1979), 125–157.

[8] P. Golovach, M. Johnson, D. Paulusma and J. Song: A Survey on the Computational Complexity of Colouring Graphs with Forbidden Subgraphs, arXiv:1407.1482v6 [cs.CC], June 2015.

[9] P. Golovach, D. Paulusma and J. Song: Closing complexity gaps for coloring problems on H-free graphs, *Information and Computation* **237** (2014), 204–214.

[10] P. Hell and S. Huang: Complexity of coloring graphs without paths and cycles, arXiv:1310.0340v1 [cs.DM], September 2013.

[11] C. T. Hoàng, M. Kamiński, V. V. Lozin, J Sawada and X. Shu: Deciding k-colorability of P_5-free graphs in polynomial time, *Algorithmica* **57** (2010), 74–81.

[12] I. Holyer: The NP-completeness of edge-coloring, *SIAM Journal on Computing* **10** (1981), 718–720.

[13] S. Huang: Improved complexity results on k-coloring P_t-free graphs, in: *Proceedings of the International Symposium on Mathematical Foundations of Computer Science 2013*, volume 7551 of *Lecture Notes in Computer Science*, 551–558, 2013.

[14] S. Huang, M. Johnson and D. Paulusma: Narrowing the complexity gap for colouring (C_s, P_t)-free graphs, arXiv:1407.1480v1 [cs.CC], July 2014.

[15] M. Kamiński and V. V. Lozin: Coloring edges and vertices of graphs without short or long cycles, *Contributions to Discrete Mathematics* **2** (2007), 61–66.

[16] R. Karp: Reducibility among combinatorial problems, in: R. Miller and J. Thatcher, editors, *Complexity of Computer Computations*, 85–103. Plenum Press, New York, 1972.

[17] D. Král, J. Kratochvíl, Zs. Tuza and G. J. Woeginger: Complexity of coloring graphs without forbidden induced subgraphs, in: M. C. Golumbic, M. Stern, A. Levy, and G. Morgenstern, editors, *Proceedings of the International Workshop on Graph-Theoretic Concepts in Computer Science 2001*, volume 2204 of *Lecture Notes in Computer Science*, pages 254–262, 2001.

[18] D. Leven and Z. Galil: NP-completeness of finding the chromatic index of regular graphs, *Journal of Algorithms* **4** (1983), 35–44.

[19] F. Maffray and M. Preissmann: On the NP-completeness of the k-colorability problem for triangle-free graphs, *Discrete Mathematics* **162** (1996), 313–317.

[20] S. Mellin: Polynomielle farbungsalgorithmen für P_k-freie graphen, Diplomarbeit am Institut für Informatik, Universität zu Köln, 2002.

[21] B. Randerath and I. Schiermeyer: 3-Colorability $\in$ P for P_6-free graphs, *Discrete Applied Mathematics* **136** (2004), 299–313.

[22] B. Randerath, I. Schiermeyer and M. Tewes: Three-colorability and forbidden subgraphs. II: polynomial algorithms, *Discrete Mathematics* **251** (2002), 137–153.

[23] J. Stacho: Personal communication, 2014.

[24] L. Stockmeyer: Planar 3-colorability is polynomial complete. *ACM SIGACT News* **5** (1973), 19–25.

[25] Zs. Tuza: Graph colorings with local constraints – a survey, *Discussiones Mathematicae. Graph Theory* **17** (1997), 161–228.

[26] V. Vizing: Coloring the vertices of a graph in prescribed colors, *Metody Diskretnogo Analiza* **29** (1976), 3–10.

Flavia Bonomo

CONICET and Departamento
de Computación, FCEN
Universidad de Buenos Aires
Argentina
fbonomo@dc.uba.ar

Maria Chudnovsky

Princeton University
Princeton, NJ 08544, USA
mchudnov@math.princeton.edu

Peter Maceli

Department of Mathematics
and Statistics
Canisius College Buffalo
NY 14208, USA
peter.maceli@canisius.edu

Oliver Schaudt

Institut für Informatik
Universität zu Köln
Köln, Germany
schaudto@uni-koeln.de

Maya Stein

Departamento de Ingeniería Matemática
Universidad de Chile
Santiago, Chile
mstein@dim.uchile.cl

Mingxian Zhong

Columbia University
New York, NY 10027, USA
mz2325@columbia.edu