M2 Internship:
On the chromatic number of graphs of bounded twin-width

Romain Bourneuf

Département d'Informatique
ENS de Lyon

Supervised by:
Stéphan Thomassé
Professor at ENS de Lyon

LIP, Lyon
From February 6th 2023 to July 7th 2023

1 Introduction

There are many ways to measure the complexity of a class of discrete structures. From an algorithmic standpoint, a class can be considered simple if computations are easier on this class. From a structural perspective, a class where all elements can be constructed with few natural operations starting from a small set of basic elements can be considered simple. From a counting viewpoint, a class can be considered simple if it only contains a small number of elements of each fixed size, for some natural definition of size. Complexity parameters that capture several such aspects are often of prime importance in discrete mathematics and computer science. Examples of such parameters include tree-width, clique-width and VC-dimension.

In [6], Bonnet, Kim, Thomassé and Watrigant introduced twin-width as a new parameter. Graphs of bounded twin-width exhibit desirable properties in terms of computational complexity (first-order model checking can be done in linear time), model theory (they are closed under first-order interpretations), enumerative combinatorics (they form small classes [4]), and structural decomposition (they generalize classes of bounded clique-width, and proper minor-closed classes).

When examining a class of graphs, another way to assess its complexity is to consider the chromatic number of the graphs in the class, and especially how it compares to their clique number. Indeed, while for every graph G we have $\omega(G) \leq \chi(G)$, early constructions by Blanche Descartes [11], Zykov [19] and Mycielski [16] show that there are triangle-free graphs with arbitrarily large chromatic number.

A hereditary class of graphs $\mathcal{C}$ (i.e. closed under induced subgraphs) is *χ-bounded* if there exists a function f such that $\chi(G) \leq f(\omega(G))$ for every $G \in \mathcal{C}$. When f can be chosen as a polynomial, the class $\mathcal{C}$ is *polynomially χ-bounded*. With the exception of VC-dimension, classes of graphs with bounded complexity parameter are usually χ-bounded. This is the case for bounded tree-width, as it implies bounded degeneracy, but also the case for rank-width, as shown in [12]. Building on this result, Bonamy and Pilipczuk [1] showed that graphs with bounded clique-width are polynomially χ-bounded. It was shown in [5] that graphs with bounded twin-width are χ-bounded. In the same paper, a polynomial bound was posed as an open problem, which would extend the polynomial χ-boundedness of graphs with bounded clique-width.

A natural step when trying to achieve polynomial bounds is to first look for quasi-polynomial ones, which was done by Pilipczuk and Sokołowski in [17].

1.1 Overview of my internship

The initial objective of my internship with Stéphan Thomassé was to extend the result of [17] to prove the polynomial χ-boundedness of graphs of bounded twin-width. To do so, we thoroughly studied [17] and [9]. Building on these two papers, we devised two new operations, the *delayed decomposition* and the *right module partition*. Using them, we proved that graphs of bounded twin-width are polynomially χ-bounded [7]. We then focused on a question from [9], asking whether the closure of a χ-bounded class under substitution and gluing along small subsets of vertices is χ-bounded. To prove that this was indeed the case, we considered a simplified version of this problem, but we still did not manage to prove anything. With the help of Nicolas Trotignon, we thought we had found a proof, but there was a subtle mistake in that proof. We then asked Édouard Bonnet, Julien Duron and Colin Geniet to join us to try to find a counterexample, which we eventually did. This led to a new construction of a hereditary class of triangle-free graphs with unbounded chromatic number, the *twincut graphs* [3]. These graphs have a strinkingly low structural complexity, as their twin-width is almost as small as possible for triangle-free graphs of high chromatic number (up to an additive constant of 1). Furthermore, each of these graphs has a pair of non-adjacent twins or an edgeless cutset of size at most 2. With Édouard Bonnet, Colin Geniet and Stéphan Thomassé, we then started working on generalizing delayed decompositions to other discrete structures than graphs, such as permutations, which are known to behave well with respect to twin-width. With that, we showed that every strict class of permutations is contained in a bounded power of the class of separable permutations.

During my internship, I also had the opportunity to travel and present our results. I went to the 1st Workshop on Twin-Width in Aussois, where I presented our result on the polynomial χ-boundedness of bounded twin-width graphs. Then, I attended the Fifth ANR Digraph meeting in Sète. Finally, I travelled to Bergen for the FPT Fest in the honour of Mike Fellows, where I also presented our result on the polynomial χ-boundedness of bounded twin-width graphs.

1.2 Organization of the report

In Section 2, I introduce some notations, as well as the concepts of twin-width and χ-boundedness, which will be central in this report. In Section 3, I present several operations on graph classes that preserve χ-boundedness, including the delayed decomposition and the right module partition, which we defined during my internship. In Section 4, I present the main part of the proof that graphs of bounded twin-width are polynomially χ-bounded. In Section 5, I define the twincut graphs and show some of their properties. Finally, in Section 6, I give a brief overview of the factorization result for permutations of bounded twin-width.

2 Preliminaries

2.1 Notations and conventions

If n is an integer, we denote by $[n]$ the set $\{1, 2, \ldots, n\}$.
In this report, all graphs are undirected and simple (no multiple edges, no self loop). When talking about a *class* of graphs, we mean a *hereditary class* of graphs, i.e. closed under induced subgraphs. If G is a graph, we denote by $V(G)$ its vertex set and $E(G)$ its edge set. If $S \subseteq V(G)$, we denote by $N(S)$ the set of neighbours of S, deprived of S, and by $N[S]$ the set $N(S) \cup S$. If v is a vertex, we write $N(v)$ for $N(\{v\})$ and $N[v]$ for $N[\{v\}]$. The *degree* of a vertex v is its number of neighbours, i.e. the size of $N(v)$. Two distinct vertices u, v such that $N(u) = N(v)$ are called *false twins*, and *true twins* if $N[u] = N[v]$. Note that false twins are not adjacent while true twins are adjacent. Two vertices are called *twins* if they are true twins or false twins. If X is a subset of vertices of G, we denote by $G[X]$ the graph induced by G on X. If u is a vertex of G, we denote by $G - u$ the graph $G[V(G) \setminus \{u\}]$. We denote by $\omega(G)$ the size of the largest clique in G. A *proper k-coloring* of G is a function $c : V(G) \to [k]$ such that for every edge uv, we have $c(u) \neq c(v)$. The *chromatic number* of G, denoted $\chi(G)$, is the smallest k such that G has a proper k-coloring. We say that $X \subseteq V(G)$ is a *module* of G if for every $y \in V(G) \setminus X$ we have either all edges between y and X, or no edge between y and X.

2.2 Twin-width

To define the notion of twin-width, we first need to introduce the notion of trigraph. A *trigraph* is a triple $G = (V, E, R)$ where V is a set of vertices, and E, R are disjoint sets of edges on V. We will refer to E as the set of *black edges* and to R as the set of *red edges*. Informally, the presence of a red edge between two vertices means that we are unsure of the presence of this edge in the graph. Every graph $G = (V, E)$ can be interpreted as the trigraph $G = (V, E, \emptyset)$. A *d-trigraph* is a trigraph where the red degree of each vertex is at most d. We now define the notion of contraction. If $G = (V, E, R)$ is a trigraph and u, v are distinct vertices (not necessarily adjacent), the *contraction* of u, v is the trigraph we obtain by merging u and v into a single vertex w, and updating the edges in the following way. If both u, v are adjacent to x with a black edge, then wx is a black edge. If there was no edge ux and no edge vx, then there is no edge wx. In all other cases, there is a red edge wx. The rest of edges (not incident to u or v) remains unchanged. Note that u and v are removed from the graph.

A *d-sequence* of a graph G is a sequence a d-trigraphs $G_n, G_{n-1}, \ldots, G_1$ where $G_n = G$, G_1 is the trigraph on one vertex, and each G_{i-1} is obtained from G_i by performing a contraction of two vertices. Note that for every $i \in [n]$, G_i has exactly i vertices. The *twin-width* of G, denoted by $tww(G)$, is the minimum integer d such that G admits a d-sequence. For instance, the graphs of twin-width 0 are exactly the cographs.

If v is a vertex of G_i, we denote by $v(G)$ the set of vertices of G that were eventually contracted into v. Then, there is a black edge uv in G_i if and only if the bipartite graph between $u(G)$ and $v(G)$ is a complete bipartite graph. Similarly, there is no edge uv in G_i if and only if there is no edge between $u(G)$ and $v(G)$.

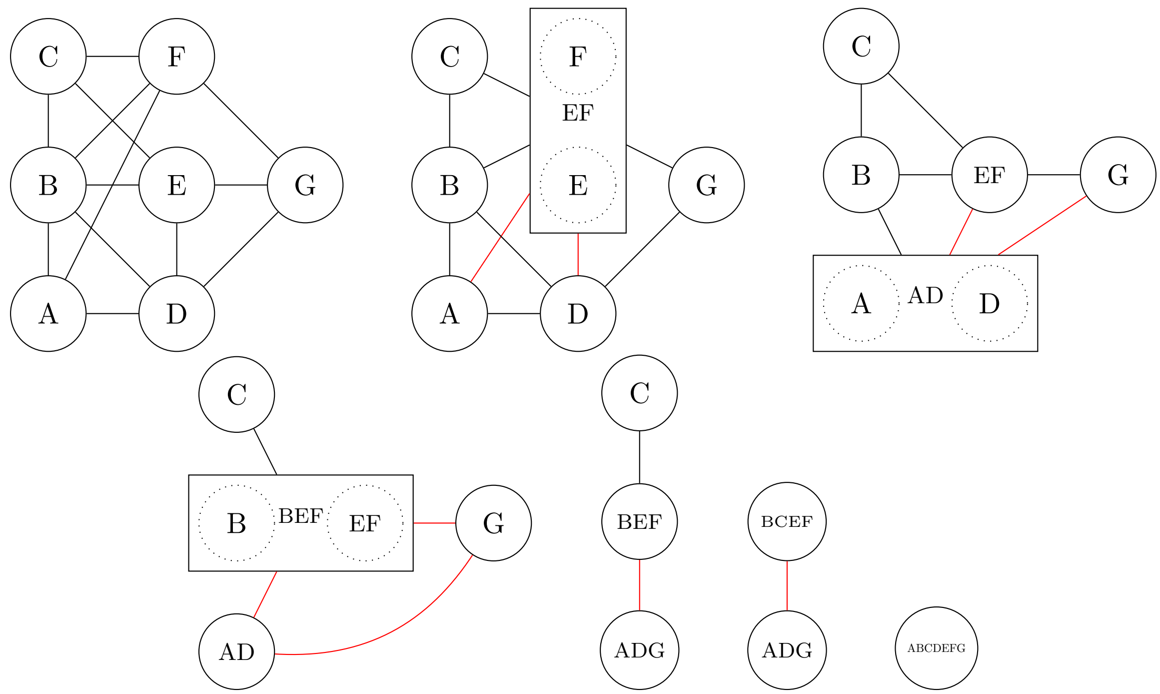

Figure 1: A 2-sequence witnessing that the original graph has twin-width at most 2.

2.3 χ-boundedness

A class of graphs $\mathcal{C}$ is *χ-bounded* if there exists a function f such that $\chi(G) \leq f(\omega(G))$ for every $G \in \mathcal{C}$. When f can be chosen as a polynomial, the class $\mathcal{C}$ is *polynomially χ-bounded*. Polynomial χ-boundedness is one of the tamest behaviours a class can have, and it generally indicates that the class has a lot of structure. There are classes which are χ-bounded but not polynomially χ-bounded, see Briański, Davies and Walczak [8]. The field is developing at a very fast pace: for a recent survey, see Scott and Seymour [18].

As an example, the class of graphs corresponding to $f(x) = x$ is the celebrated class of perfect graphs. This class can alternatively be defined by forbidden induced subgraphs (odd holes and antiholes) as well as it enjoys some structural results based on some elementary decompositions. This leads to three main directions of research: Which forbidden induced subgraphs give polynomial χ-boundedness? Which structural parameters give polynomial χ-boundedness? Which operations preserve χ-boundedness? These three questions are intimately connected. A canonical example of polynomially χ-bounded class, cographs, are altogether P_4-free graphs (no induced path of length 4), have clique-width at most 2 and are the closure under substitutions of the graphs of size at most 2.

From the forbidden induced subgraph perspective, the Perfect Graph Theorem, proved by Chudnovsky, Robertson, Seymour and Thomas [10] is definitely the masterpiece of the field. From the point of view of operations preserving (polynomial) χ-boundedness, the landscape is less developed. In this direction, Chudnovsky, Penev, Scott, and Trotignon [9] showed that the substitution closure of a (polynomially) χ-bounded class is also (polynomially) χ-bounded. They also showed that if a class is (polynomially) χ-bounded, its closure under the operation of gluing on a small set of vertices is also (polynomially) χ-bounded.

3 Extensions of graph classes

Our approach to prove that graphs of bounded twin-width are polynomially χ-bounded is the following. We want to show that each graph of twin-width at most d can be constructed using some operations, starting from some class of basic graphs. If these operations preserve polynomial χ-boundedness and if the basic class is polynomially χ-bounded, we immediately get the desired result. This structural approach is very common in the study of graph classes. In this section, we will explore several operations that preserve χ-boundedness.

3.1 Substitution closure

When working on χ-boundedness, one of the most natural classes is the class of perfect graphs. A graph G is *perfect* if for every induced subgraph H of G, we have $\chi(H) = \omega(H)$. One can wonder which operations preserve perfectness. In [14], Lovász proved the so-called Weak Perfect Graph Theorem, stating that the complement of a perfect graph is also a perfect graph. In the same paper, he also proved that perfect graphs are closed under substitution.

Given two vertex-disjoint graphs G and H, and a vertex u of G, by *substituting* H for u in G, we mean deleting u and joining every vertex of H to all neighbours of u in G. Formally, G' is obtained by *substituting* H for u in G if the following conditions hold:

- $V(G') = (V(G) \setminus \{u\}) \cup V(H)$,
- $G'[V(G) \setminus \{u\}] = G - u$,
- $G'[V(H)] = H$,
- For $v \in V(G) \setminus \{u\}$ and $w \in V(H)$, $vw \in E(G')$ if and only if $uv \in E(G)$.

Given a class of graphs $\mathcal{C}$, the *substitution closure* of $\mathcal{C}$, denoted by $\mathcal{C}_s$, is the class of all graphs we can obtain as follows: we start from a graph of $\mathcal{C}$, we substitute a vertex of this graph by a graph of $\mathcal{C}$, and we repeat this process any number of times. For instance, if $\mathcal{C}$ is the class of complete graphs and edgeless graphs, then $\mathcal{C}_s$ is the class of cographs. Since perfect graphs are closed under substitution, if $\mathcal{C}$ is the class of perfect graphs, then $\mathcal{C} = \mathcal{C}_s$. This indicates that the substitution closure behaves well with respect to χ-boundedness. More generally, Chudnovsky, Penev, Scott and Trotignon [9] proved the following result.

Theorem 3.1 ([9]) *If $\mathcal{C}$ is a class of graphs which is polynomially χ-bounded with function $\chi \leq \omega^k$, then $\mathcal{C}_s$ is polynomially χ-bounded with function $\chi \leq \omega^{3k+11}$.*

To say it with words, this means that the operation of substitution preserves polynomial χ-boundedness.

3.2 Tree decompositions

We now describe another (equivalent) way of defining the substitution closure of a class of graphs, based on tree decompositions. This definition is more natural to work with and easier to generalize.

To distinguish the elements of our (rooted) tree decompositions from the vertices of our graphs, we will speak about *nodes* of trees. Also we will make use of *parent* (first node of the path to the root), *ancestors* (nodes of the path to the root), *children*, *siblings* (nodes with same parent), *grandchildren* (children of a child), and *cousins* (non sibling grandchildren).

Given a class of graphs $\mathcal{C}$, a *$\mathcal{C}$-tree-decomposition* is a pair (T, g) in which T is a rooted tree and g is a function associating to every internal node x of T a graph $g(x) \in \mathcal{C}$ whose vertices are the children of x. The *realization* $R(T, g)$ is the graph such that:

- its vertex set is the set of leaves L of T,
- two vertices $x, y \in L$ are joined by an edge if, given that z is their closest ancestor in T and x', y' are the respective children of z which are the ancestors of x, y, the edge $x'y'$ belongs to $g(z)$.

Given a class $\mathcal{C}$, its substitution closure $\mathcal{C}_s$ is the class of all $R(T, g)$ where (T, g) is a $\mathcal{C}$-tree-decomposition. Let us give some intuition on why this is the case. Let (T, g) be a $\mathcal{C}$-tree-decomposition, with r the root of T, and $G = R(T, g)$. To build G using substitutions, we start from the graph $g(r) \in \mathcal{C}$. Then, for every node x that is a child of r, we substitute x by $g(x) \in \mathcal{C}$. We then keep doing this substitution process, sweeping all levels of the tree T. After having done so, we obtain a graph whose vertex set is indeed the set of leaves L of T, and it is straightforward to check that the adjacency relation in the graph we obtain corresponds to the adjacency relation in $R(T, g)$.

This was the "constructive" point of view on the substitution closure. There is a dual point of view, which can be described as a "decomposition" point of view. We now sketch it. Let (T, g) be a $\mathcal{C}$-tree-decomposition with root r. Given any node x_i that is a child of r in T, let $X_i \subseteq L$ be the set of leaves with ancestor x_i. Then, X_i is a module in G. Hence, to decompose our graph, we can partition its vertex set into modules, and ask that the quotient graph, which we obtain by keeping one vertex in each module, is in $\mathcal{C}$. This will be the case since this graph is precisely $g(r)$. We further require that the graph induced on each module is in $\mathcal{C}_s$.

In general, the aim of graph decompositions is to describe graphs that are somehow simple. One way to do so is to say that a graph is simple if we can partition its vertex set into parts, such that the "quotient graph" between the parts – for some adequate definition of quotient graph – is already known to be simple, and such that the graph induced on each part is also already known to be simple. This is exactly what happens for the substitution closure, and we will see it appear again several times.

3.3 Delayed tree decompositions

Tree decompositions are only useful when the graphs we consider have modules. However, graphs of bounded twin-width need not have modules. For this reason, we introduced a new decomposition technique, the *delayed tree decomposition*, which is more suited to the study of graphs of bounded twin-width.

Recall that a *grandchild* of a node x in a tree is a child of one of its children. Given a class of graphs $\mathcal{C}$, a *$\mathcal{C}$-delayed-tree-decomposition* is a pair (T_d, g) in which:

- T_d is a rooted tree in which every leaf has a parent with only one child,
- g is a function associating to every node $x \in T_d$ that has grandchildren a graph $g(x) \in \mathcal{C}$ such that the vertices of $g(x)$ are the grandchildren of x (hence the delayed).

The *delayed realization* $R_d(T_d, g)$ is the graph such that:

- its vertex set is the set of leaves L_d of T_d,
- two vertices $x, y \in L_d$ are joined by an edge if, given that z is their closest ancestor in T and x', y' are the respective grandchildren of z which are the ancestors of x, y, the edge $x'y'$ belongs to $g(z)$.

Given a class $\mathcal{C}$, we denote by $\mathcal{C}_d$ the class of all $R_d(T_d, g)$ where (T_d, g) is a $\mathcal{C}$-delayed-tree-decomposition. We call $\mathcal{C}_d$ the *delayed extension* of $\mathcal{C}$. It is not strictly speaking a closure since applying it twice can produce more graphs than applying it once.

Observe that by definition of $R_d(T_d, g)$, we never look at the adjacency between nodes that are siblings. Hence, such edges in the $g(x)$ can help make $g(x)$ simpler, but they are meaningless with respect to $R_d(T_d, g)$.

Delayed tree decompositions are in spirit very similar to tree decompositions, the main difference being on the definition of the "quotient graphs", the $g(x)$, which are now defined on the grandchildren of a node and not on its children, thus offering more flexibility to decompose module-free graphs. This similarity is reflected in the following lemma.

Lemma 3.2 *Every graph in $\mathcal{C}_d$ is the edge union of two graphs in $\mathcal{C}_s$.*

Proof. Let $R_d(T_d, g)$ be a delayed realization of a $\mathcal{C}$-delayed-tree-decomposition. We define the *parity* of a node of T_d as its distance to the root modulo 2. For instance, the root is even. We partition the set of edges xy of $R_d(T_d, g)$ into *odd edges* and *even edges* depending of the parity of the closest ancestor of x, y in T_d. Hence $R_d(T_d, g)$ is the edge union of two spanning subgraphs G_o (odd edges) and G_e (even edges). The key-fact is that G_o and G_e are both a realization of a $\mathcal{C}$-tree-decomposition. To see this, we form two subsets of nodes T_o (odd nodes) and T_e (even nodes) of T_d. Note that since every leaf l of T_d has a parent l' with only one child, either l or l' belongs to T_e (or to T_o). Hence, to every leaf l (and thus to every vertex of $R_d(T_d, g)$) one can identify a leaf of T_o and a leaf of T_e. Now, g restricts naturally as a function g_e of the internal nodes of T_e in which $g_e(x)$ is a graph on the children of x in T_e. The same holding for g_o, we finally obtain that $R_d(T_d, g) = R(T_e, g_e) \cup R(T_o, g_o)$. ∎

Corollary 3.3 *If $\mathcal{C}$ is polynomially χ-bounded with function ω^k, then $\mathcal{C}_d$ is polynomially χ-bounded with function ω^{6k+22}.*

Proof. Let $G \in \mathcal{C}_d$. By Lemma 3.2, G is the edge union of G_o and G_e with $G_o, G_e \in \mathcal{C}_s$. Then, $\chi(G) \leq \chi(G_o) \cdot \chi(G_e)$ (take the product of the two colorings). By Theorem 3.1, we have $\chi(G_o) \leq \omega(G_o)^{3k+11} \leq \omega(G)^{3k+11}$ and similarly for $\chi(G_e)$. Thus, $\chi(G) \leq \omega(G)^{6k+22}$. ∎

3.4 Obtaining a delayed tree decomposition

In this section, we prove that every graph G has a delayed tree decomposition and we describe a canonical way to compute one when given an order on the vertices. Graphs of bounded twin-width naturally come with such an order, which gives a natural way of obtaining a delayed tree decomposition.

Let G be a graph on vertex set $V(G) = \{v_1, \ldots, v_n\}$. The tree T_d is defined via a sequence of refining partitions $P_0, P_1, \ldots, P_k$ of $V(G)$, starting with the root vertex $P_0 = \{V(G)\}$ and ending on $P_k = \{\{v_1\}, \ldots, \{v_n\}\}$, the partition into singletons corresponding to the vertices. We now describe how to construct the refinement P_{i+1} of $P_i = \{B_1, \ldots, B_m\}$, that is how a part B_j is further refined. In this process, all parts will consist of consecutive vertices in the ordering $v_1, \ldots, v_n$.

For this, we define a partition $P(I)$ of an arbitrary interval of consecutive vertices $I = \{v_i, \ldots, v_j\}$.

- If $i = j$, it is not refined.
- If I is a module in G, we divide it into two parts $\{v_i, \ldots, v_{\lfloor (i+j)/2 \rfloor}\}$ and $\{v_{\lfloor (i+j)/2 \rfloor + 1}, \ldots, v_j\}$.
- If I is not a module in G, we partition it into maximal subsets S of consecutive vertices such that S is a module in $G[(V \setminus I) \cup S]$. We call these parts *local modules*. In other words $v_s, v_{s+1} \in I$ are in the same local module of I if and only if they have the same neighbours in $V \setminus I$.

To form the refinement P_{i+1} of $P_i = \{B_1, \ldots, B_m\}$, we just refine every B_j into $P(B_j)$. Observe that since $P_0 = \{V(G)\}$ is a module, we have $P_1 = \{\{v_1, \ldots, v_{\lceil n/2 \rceil}\}, \{v_{\lceil n/2 \rceil + 1}, \ldots, v_n\}\}$. We stop the process when $P_k = P_{k-1}$, which happens when all parts are singletons. We keep the two identical partitions into singletons P_{k-1}, P_k instead of simply stopping at step $k-1$ so that each leaf has a parent with only one child. The partition of a module into two (near) equal intervals is arbitrary, we could for instance partition I into $\{v_i\}$ and $\{v_{i+1}, \ldots, v_j\}$.

We next consider the tree T_d corresponding to this decomposition process, where the nodes at depth i correspond to the parts of P_i, and the children of a node $x \in P_i$ are the parts $y \in P_{i+1}$ such that $y \subseteq x$ (we usually identify the nodes of T_d to subsets of $V(G)$). The leaves of T_d are the vertices v_i of G. Moreover, the parent of a leaf v_i is v_i since $P_{k-1} = P_k$. We now describe how to structure T_d as a *delayed tree decomposition* in order to encode the graph G. Consider any node x of T_d that has two grandchildren y, z which are not siblings (recall that edges between siblings are meaningless). Let y' be the parent of y and z' the parent of z. By construction of y and z, we have that y is a module in $G[(x \setminus y') \cup y]$ and that z is a module in $G[(x \setminus z') \cup z]$. This implies that y and z are modules in $G[y \cup z]$. In other words, cousins are modules with respect to each other. Therefore, either we have all edges between y and z, or no edge between them.

From this observation, we define a function g associating to every node $x \in T_d$ a graph $g(x)$ whose vertex set is the set of grandchildren of x and such that yz is an edge of $g(x)$ if y, z are cousins and there exists an edge between y and z in G (and thus y is fully joined to z).

It is then straightforward to check that we indeed have $G = R_d(T_d, g)$.

Given a class $\mathcal{C}$, let $\mathcal{C}_q$ be the class of all graphs that appear as a quotient graph in a delayed decomposition of some element of $\mathcal{C}$. By Corollary 3.3, to prove that $\mathcal{C}$ is polynomially χ-bounded, it suffices to prove that $\mathcal{C}_q$ is polynomially χ-bounded. This is what we will do to prove that graphs of bounded twin-width are polynomially χ-bounded. The operation we introduce in the next section will help us analyze the quotient graphs in that case.

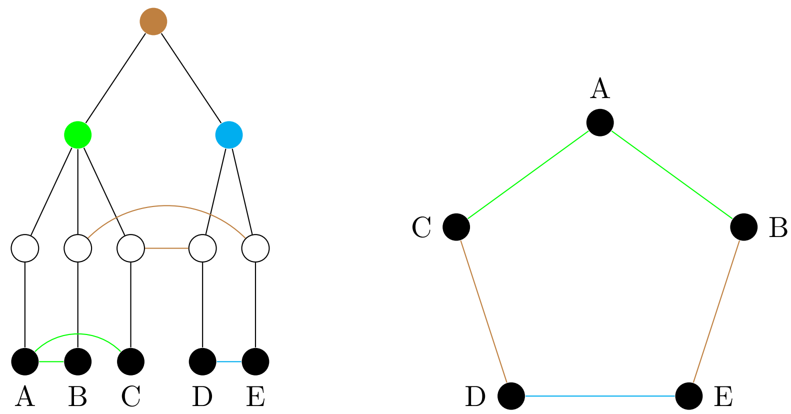

Figure 2: The canonical delayed decomposition tree of C_5 obtained from the order A, B, C, D, E on the vertices, and its corresponding realization. The edges of every $g(x)$ are drawn in the same color as x. Note that all $g(x)$ are cographs.

3.5 Right extension

In this section, we define an extension of a class of graphs $\mathcal{C}$ which preserves χ-boundedness, and even polynomial χ-boundedness when the twin-width of $\mathcal{C}$ is bounded.

Given a graph G, a *right module partition* (RMP) is a partition $V_1, \ldots, V_k$ of the vertices of G such that

1. Each V_i is a stable set,
2. For every $i < j$, V_i is a module with respect to V_j (i.e. V_i is a module in $G[V_i \cup V_j]$).

Note that every graph G has a trivial RMP where each V_i consists of a single vertex. Therefore, there should be some limitations to the definition of RMP, for instance by imposing restrictions on the quotient graph between the parts. A first attempt to define the quotient is to consider a class of graphs $\mathcal{C}$, and insist that every induced subgraph intersecting every V_j on at most one vertex (called a *transversal*) belongs to $\mathcal{C}$. Unfortunately, even RMP with forests transversals are not χ-bounded. To see this, consider $S_{n,2}$, the n-th shift graph, whose vertex set is $\{(i,j), 1 \leq i < j \leq n\}$ and such that there is an edge between (i,j) and (i',j') if and only if $j = i'$. Observe that the graphs $S_{n,2}$ have unbounded chromatic number and are triangle-free. However, the partition $(V_2, \ldots, V_n)$ where $V_j = \{(i,j), 1 \leq i < j\}$ for $2 \leq j \leq n$ is an RMP (with V_1 empty) such that the only neighbours of $(i,j) \in V_j$ in parts V_k where $k < j$ are in V_i. Thus if (i,j) belongs to a transversal, its degree is at most one with respect to the vertices in V_k with $k < j$. Hence, all transversals are forests, while the graphs $S_{n,2}$ are not χ-bounded.

For this reason, we introduce a stronger notion of RMP, meant to preserve χ-boundedness. If $\mathcal{P} = (V_1, \ldots, V_k)$ is an RMP of a graph G, for every $1 \leq j_1 < j_2 < \ldots < j_\ell \leq k$ and every $W_{j_1} \subseteq V_{j_1}, \ldots, W_{j_\ell} \subseteq V_{j_\ell}$, all non-empty, we denote by $G/\{W_{j_1}, \ldots, W_{j_\ell}\}$ the graph on vertex set $[\ell]$ such that there is an edge ii' if and only if there is an edge between W_{j_i} and $W_{j_{i'}}$. We call such a graph a *transversal minor* of $(G, \mathcal{P})$. Given a class $\mathcal{C}$, an RMP such that all transversal minors are in $\mathcal{C}$ is called a *$\mathcal{C}$-RMP*. The class of graphs G admitting a $\mathcal{C}$-RMP is denoted by $RM(\mathcal{C})$ and is called the *right extension* of $\mathcal{C}$.

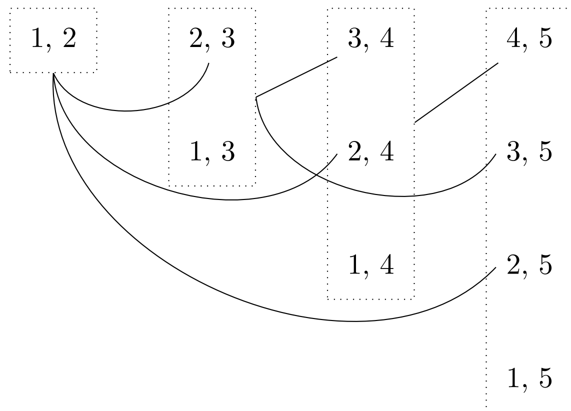

Figure 3: The RMP for $S_{5,2}$. Here, an edge means that we have all edges from the stable set on the left to the vertex on the right. Observe that every transversal is a forest, however we can form every graph on 4 vertices as a transversal minor.

The following result proves that right extensions behave well with respect to the chromatic number. It is is not used in the proof that bounded twin-width graphs are polynomially χ-bounded, hence we postpone its proof to Appendix A.

Proposition 3.4 *If $\mathcal{C}$ is a χ-bounded class of graphs, then $RM(\mathcal{C})$ is also χ-bounded.*

It is worth noting that the proof of this proposition does not provide a polynomial bound if the class $\mathcal{C}$ is polynomially χ-bounded. We could not prove (or disprove) that polynomial χ-boundedness is preserved by RMP. However, this is the case when the twin-width of $\mathcal{C}$ is bounded, as shown in the following section.

4 Polynomial χ-boundedness of bounded twin-width

With all these operations in hand, we move to the proof that classes of graphs of bounded twin-width are polynomially χ-bounded. The proof we give heavily relies on some definitions, tools and techniques introduced by Pilipczuk and Sokołowski in [17], where they prove the quasi-polynomial χ-boundedness of graphs of bounded twin-width.

4.1 Twin-width and matrices

In this section, we will adopt the matrix point of view of twin-width, where every graph G is represented via its symmetric adjacency matrix $(a_{u,v})$ where u, v are over couples of vertices. The entry $a_{u,v}$ is 1 if uv is an edge, 0 if uv is not an edge, and $*$ if $u = v$. The addition of the $*$ symbol slightly simplifies some technicalities, but is not necessary for the argument.

A 01$*$-matrix is *horizontal* if all its rows are constant. It is *vertical* if all its columns are constant. It is *constant* if it is both horizontal and vertical. It is *mixed* if it is neither horizontal nor vertical, or if it has at least 2 rows and 2 columns and contains a $*$ entry. A *corner* in a matrix M is a mixed 2×2 contiguous submatrix of M.

Lemma 4.1 ([6]) *A matrix is mixed if and only if it contains a corner.*

A *row partition* of a matrix M is a partition of the rows of M in which each part of the partition consists of consecutive rows. We define *column partitions* in a similar way. A *division* $\mathcal{D}$ of M is a pair $(\mathcal{R},\mathcal{C})$ where $\mathcal{R}$ is a row partition of M and $\mathcal{C}$ is a column partition of M. If $\mathcal{R}$ and $\mathcal{C}$ both have the same number of parts, say k, we say that $(\mathcal{R},\mathcal{C})$ is a *k-division* of M. In this case, we index the row blocks and the column blocks of $\mathcal{D}$ with integers from $[k]$ in the natural order of the blocks. If $\mathcal{D}$ is a k-division of M, for $i,j \in [k]$, we denote by $\mathcal{D}[i,j]$ the intersection of the i-th row block with the j-th column block, which we call a *zone* of $\mathcal{D}$. If $R_i \in \mathcal{R}$ and $C_j \in \mathcal{C}$, we also adopt the notation $[R_i, C_j]$ for the zone $\mathcal{D}[i,j]$. It is a contiguous submatrix of M.

We say that a zone of a matrix is *mixed* if it is mixed as a submatrix. If M is symmetric, we say that a division $(\mathcal{R},\mathcal{C})$ is *symmetric* if $\mathcal{R}$ and $\mathcal{C}$ partition rows and columns in the same way (i.e. $\mathcal{R}$ is the transpose of $\mathcal{C}$).

Let $\mathcal{D}$ be a d-division of a matrix M. We say that $\mathcal{D}$ is a *d-mixed minor* if each zone of $\mathcal{D}$ is mixed. If M does not have any d-mixed minor, we say that M is *d-mixed free*. The twin-width parameter and mixed-minor freeness are functionally equivalent, see [6], but here we only need the following bound.

Theorem 4.2 ([6]) *Let G be a graph and t be a positive integer. If the twin-width of G is at most t, then there is a vertex ordering of G for which the adjacency matrix of G is $(2t+2)$-mixed free.*

We also recall the following result, which is a direct consequence of the Marcus-Tardos theorem, see [15]:

Theorem 4.3 *For every positive integer d, there is a constant mt_d such that for every d-mixed free matrix M and every k-division of M, the number of mixed zones is at most $mt_d \cdot k$*

Here is the key-definition of [17]. A d-division $\mathcal{D}$ of a matrix M is a *d-almost mixed minor* if for every $i \neq j \in [d]$, the zone $\mathcal{D}[i,j]$ is mixed. If M does not have any d-almost mixed minor, we say that M is *d-almost mixed free*. By extension, a graph is *d-almost mixed free* if we can order its vertices in such a way that its adjacency matrix is d-almost mixed free.

Observe that every d-almost mixed free matrix is also d-mixed free. Conversely, every d-mixed free matrix is also $2d$-almost mixed free. Indeed, if M has a $2d$-almost mixed minor, then merging the first $d+1$ row blocks, and the last $d+1$ column blocks gives a d-mixed minor of M. Note that every submatrix of a d-(almost) mixed free matrix is also d-(almost) mixed free, hence every subgraph of a d-almost mixed free graph is also d-almost mixed free.

4.2 Twin-width and right module partitions

Let G be a graph with an RMP $\mathcal{P} = (V_1, \dots, V_k)$. We say that $(G,\mathcal{P})$ is *d-almost mixed free*, if for every coarsening $\mathcal{P}'$ of $\mathcal{P}$ into d parts $(V'_1, \dots, V'_d)$, where each V'_i consists of consecutive parts of $\mathcal{P}$, some zone $[V'_i, V'_j]$, where $i \neq j$, is not mixed in the adjacency matrix of G. Note that we only speak here of restrictions on symmetric divisions of G, which encompass much larger classes than bounded twin-width.

Lemma 4.4 *If $(G,\mathcal{P})$ is d-almost mixed free, then $\omega(G/\mathcal{P}) \leq \omega(G)^d$.*

Proof. We write $\omega := \omega(G)$ and denote $\mathcal{P} = (V_1, \dots, V_k)$. Let ϕ such that $\omega(G/\mathcal{P}) \leq \phi(\omega, d)$. We have $\phi(\cdot, 1) = 0$ (empty graph) and $\phi(1, \cdot) = 1$ (edgeless graph). We assume $\omega \geq 2$ and $d \geq 2$ and show that $\phi(\omega, d) = \phi(\omega - 1, d) + \phi(\omega, d - 1) + 1$ upper bounds $\omega(G/\mathcal{P})$. We can restrict ourselves to a maximal clique of $G/\mathcal{P}$, so we can assume that there is an edge between V_i and V_j whenever $i < j \in [k]$.

Let us consider the smallest ℓ such that $\omega(G[V_1 \cup \cdots \cup V_\ell]) = \omega$. Denote by Y the set $V_1 \cup \cdots \cup V_\ell$, and consider any V_i where $i \geq \ell + 1$. Note that Y is not a module with respect to V_i, as some vertex in V_i would dominate Y, hence forming a clique of size $\omega + 1$ in G. Conversely, if V_i is a module with respect to Y, since $\mathcal{P}$ is an RMP, and there exists an edge between all pairs of parts, V_i would dominate Y, with the same contradiction.

Consider the graph $G' = G[V_{\ell+1} \cup \cdots \cup V_k]$ and its RMP $\mathcal{P}' = (V_{\ell+1}, \dots, V_k)$. We claim that $(G', \mathcal{P}')$ is $d-1$-almost mixed free, otherwise any $d-1$-almost mixed minor coarsening $(V'_1, \dots, V'_{d-1})$ of $\mathcal{P}'$ could be extended to the d-almost mixed minor $(Y, V'_1, \dots, V'_{d-1})$ of $(G, \mathcal{P})$.

Thus $k = \omega(G/\mathcal{P})$ satisfies by induction that $k \leq \ell + \phi(\omega, d-1)$. And since the first $\ell - 1$ parts do not contain a clique of size ω, we have $k \leq \phi(\omega - 1, d) + 1 + \phi(\omega, d-1) = \phi(\omega, d)$. Setting $\psi(\cdot, \cdot) = \phi(\cdot, \cdot) - 1$, we have that $\psi(\omega, d) = \psi(\omega - 1, d) + \psi(\omega, d - 1)$. Moreover, we both have $\psi(\omega, 1) = -1$ and $\psi(1, d) = 0$, so $\psi(\omega, d) \leq \binom{\omega+d-2}{d-1} \leq \omega^{d-1}$. ∎

Proposition 4.5 *Let $\mathcal{C}$ be a class of graphs with polynomial χ-bounding function $f(x) = x^c$. If $\mathcal{P}$ is a $\mathcal{C}$-RMP of G such that $(G, \mathcal{P})$ is d-almost mixed free, then $\chi(G) \leq \omega^{cd}$.*

Proof. Since $\omega(G/\mathcal{P}) \leq \omega(G)^d$, we have $\chi(G) \leq \chi(G/\mathcal{P}) \leq \omega(G/\mathcal{P})^c \leq \omega(G)^{cd}$. ∎

4.3 Proof of polynomial χ-boundedness of bounded twin-width

To prove that a class $\mathcal{C}$ is polynomially χ-bounded, a strategy is to show that every graph G of $\mathcal{C}$ has a vertex-partition or an edge-partition into a bounded number of graphs, each of them belonging to some known polynomially χ-bounded class $\mathcal{C}'$. We will use this argument several times here.

Theorem 4.6 *The class of d-almost mixed free graphs is polynomially χ-bounded.*

Proof. The proof is by induction on d. For $d = 2$, if G has a 2-almost mixed free adjacency matrix, G is a cograph (see [17]), hence G is perfect so the property holds. Now, let $d \geq 3$ and consider a graph G, which is d-almost mixed free with respect to the vertex ordering $v_1, \dots, v_n$. We first partition the set of vertices $V' = \{v_2, \dots, v_{n-1}\}$ into four subsets $V'_{00}, V'_{01}, V'_{10}, V'_{11}$ according to their neighbourhood in $\{v_1, v_n\}$. For instance, $V'_{01} = (V' \setminus N(v_1)) \cap N(v_n)$. It suffices to show that all four graphs $G[V'_{ij}]$ belong to some polynomially χ-bounded class. Hence, without loss of generality, we can assume that V' is a module in G. We now consider the delayed decomposition tree (T_d, g) of $G' := G[V']$ and consider the class $\mathcal{C}$ containing all $g(x)$ and their induced subgraphs. By Corollary 3.3, we just have to show that $\mathcal{C}$ is polynomially χ-bounded. The graphs $g(x)$ are obtained by starting from an interval $I = \{v_s, \dots, v_t\}$ of vertices of G', partitioning it into local modules $L_1, \dots, L_k$, and then partitioning each local module into local submodules. Let H be the graph we obtain from $G[I]$ by removing all edges inside all local modules L_i. Note that H can also be obtained by substituting the vertices of

$g(x)$ by stable sets. This does not change χ and ω. Therefore, to show that $\mathcal{C}$ is polynomially χ-bounded, it suffices to prove the following.

Claim 4.7 *Let $I = \{v_s, \dots, v_t\}$ be any interval of vertices of G'. Consider its partition into local modules $L_1, \dots, L_k$ and denote by H the graph on vertex set I obtained from $G[I]$ by deleting the edges inside the local modules. Then, H is a bounded (vertex and edge) union of graphs from hereditary polynomially χ-bounded classes.*

Note that the claim holds when I is a non trivial module since it is cut into two parts. Indeed, in this case H is bipartite and thus χ-bounded by 2. Assume now that we have local modules $L_1, \dots, L_k$, and consider all pairs $i < j$ such that $G[L_i, L_j]$ is mixed (call *mixed pairs*). The L_i's form a k-division of the adjacency matrix of $G[I]$, which is d-mixed free. Thus, by Theorem 4.3, the graph R on vertex set $[k]$ whose edges correspond to mixed pairs has at most $\frac{mt_d}{2} \cdot k$ edges. In particular, R can be vertex-colored into $(mt_d + 1)$ classes. In other words, one can partition the set of local modules into $(mt_d + 1)$ subsets, in which local modules are not pairwise mixed. We denote by L' such a subset of local modules. To prove our claim, we just have to show that $H' := H[L']$ belongs to a polynomially χ-bounded class of graphs.

Observe that for every $i < j$ and $L_i, L_j \in L'$, we have that L_i is a module in $H[L_i \cup L_j]$ (denoted $L_i \to L_j$) or L_j is a module in $H[L_i \cup L_j]$, that is $L_j \to L_i$. Note that if we both have $L_j \to L_i$ and $L_i \to L_j$, we have all edges or no edge between L_i and L_j. We now define two subgraphs $H'_{\rightarrow}$ and $H'_{\leftarrow}$ of H': in $H'_{\rightarrow}$ we only keep the edges of H' between pairs $L_i \to L_j$ where $i < j$, and in $H'_{\leftarrow}$ we only keep the edges of H' between pairs $L_i \leftarrow L_j$ where $i < j$. Note that $H' = H'_{\rightarrow} \cup H'_{\leftarrow}$, and thus we just have to show that (for instance) $H'_{\rightarrow}$ belongs to a polynomially χ-bounded class of graphs.

The graph $H'_{\rightarrow}$ with the partition L' is a right module partition. Note that the same holds for $H'_{\leftarrow}$ if we reverse the order of the local modules. We further partition $H'_{\rightarrow}$: let us say that a local module L_i is *left* if $i > 1$ and there is a vertex v_j among $v_1, v_2, \dots, v_{s-1}$ (i.e. to the left of I) which distinguishes L_{i-1} from L_i. Precisely, v_j is not joined in the same way to the last vertex of L_{i-1} and to the first of L_i. If L_i (with $i > 1$) is not left, then it is *right* (and indeed some vertex v_j to the right of I distinguishes L_{i-1} from L_i). We neglect L_1 in this definition (it only adds 1 to the chromatic number of H'). We now partition L' into L'_{ri} containing all right local modules L_i of L', and L'_{le} containing the left local modules. Again, by vertex partition, we just have to show that the RMP $H'_{\rightarrow,ri}$ which is the induced restriction of $H'_{\rightarrow}$ to L'_{ri} is polynomially χ-bounded. To apply Proposition 4.5, we first show that the transversal minors of $(H'_{\rightarrow,ri}, L'_{ri})$ are $d-1$-almost mixed free, which by induction implies that they are polynomially χ-bounded. Then we argue that $(H'_{\rightarrow,ri}, L'_{ri})$ has no large almost mixed minor. It suffices here to show that it is $2d$-almost mixed free.

Claim 4.8 ([17]) *Every transversal minor of $(H'_{\rightarrow,ri}, L'_{ri})$ is $d-1$-almost mixed free.*

Claim 4.9 *The pair $(H'_{\rightarrow,ri}, L'_{ri})$ is $2d$-almost mixed free.*

We postpone the proofs of these two claims to Appendix B and Appendix C

∎

By Theorem 4.2 and since every t-mixed-free matrix is $2t$-almost mixed-free, we finally obtain the following.

Theorem 4.10 *For every $d \in \mathbb{N}$, the class of graphs of twin-width at most d is polynomially χ-bounded.*

5 Twincut graphs

5.1 Motivation

The following was shown in [5].

Theorem 5.1 ([5]) *If G has twin-width t, then $\chi(G) \leq (t+2)^{\omega(G)-1}$.*

In the previous section, we showed that this bound is far from being tight for large values of ω, as the optimal bound is polynomial in ω and not exponential. However, this polynomial bound does not help us understand how the chromatic number of graphs of small twin-width and small clique number behaves. In particular, the following is an immediate corollary of Theorem 5.1.

Theorem 5.2 *Every triangle-free graph of twin-width t has chromatic number at most $t+2$.*

At the 2022 Barbados workshop, Bonnet [2] questioned the sharpness of the bound. We introduce a new explicit sequence of triangle-free graphs G_k, which we call *twincut graphs*, satisfying $\chi(G_k) = k$ and $tww(G_k) \leq k-1$. Using twincut graphs, we obtain here a near optimal answer (and suspect that Theorem 5.2 could be improved to $t+1$ when $t \geqslant 1$). Twincut graphs also give a new construction of triangle-free graphs with unbounded chromatic number and strikingly low structural complexity.

5.2 Construction

We now present an inductive construction of the *twincut graphs*.

A *structured tree* is a pair (T, g) where T is a rooted tree and g is a function defined on the internal nodes v of T (i.e. non leaves) such that $g(v)$ is a graph whose vertices are the children of v in T. A *branch* in T is a path from the root to one of the leaves of T. The *realization* $R(T, g)$ of (T, g) is the graph defined on vertex set $V(T) \cup B$, where B is the set of branches of T. The edges of $R(T, g)$ first consist of all uv where u, v are children of z and uv is an edge of $g(z)$. At this point, the graph $R(T, g)$ is simply the disjoint union of all $g(z)$ and some isolated vertices (such as B and the root). Next, we connect each *branch vertex* $b \in B$ to all the vertices of T in the branch b. Observe that the edges of T are not edges of $R(T, g)$. Note that this realization has nothing to do with the realization of the tree decompositions we mentioned previously, even though the notations are similar.

Note that when T has only one (root) node, it is also a leaf. In particular g is empty (T has no internal node) and therefore $R(T, g)$ is obtained from T by adding a single vertex which is adjacent to the root. Hence, $R(T, g)$ is K_2. We now move to the construction itself.

First, G_1 is defined as the graph on one vertex. Assuming that $G_1, \dots, G_{k-1}$ have been built, the graph G_k is defined as the realization of the following structured tree (T_k, g_k): the tree T_k has $k-1$ levels (the root being at level 1), and for each node v at level $i < k-1$, we give $|V(G_{i+1})|$ children to v and set $g_k(v) = G_{i+1}$. For instance T_2 consists only of its root, and its realization G_2 is K_2 as explained above. Then, T_3 has a root r with two children c, c' which are linked in $g_3(r) = G_2$. The realization adds a vertex x connected to r, c, and a vertex y connected to r, c', thus creating a 5-cycle $rxcc'y$, hence $G_3 = C_5$. The graph G_4 has $1 + 2 + 10 + 10 = 23$ vertices, see Figure 4.

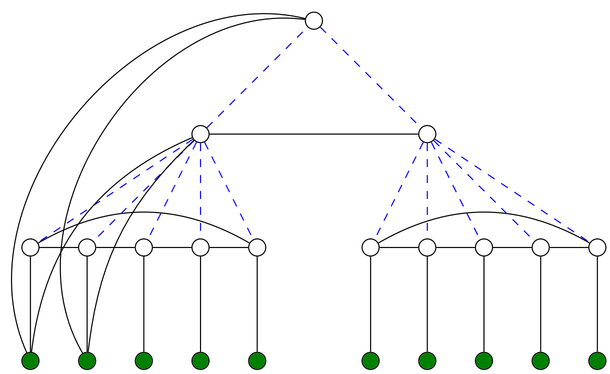

Figure 4: The 4-chromatic triangle-free graph G_4. The tree T_4 is represented with dashed blue edges (which are *not* actual edges of G_4). Every green vertex is adjacent to all vertices in a branch of T_4. We explicitly represented these edges for the two leftmost green vertices.

5.3 Properties of twincut graphs

Proposition 5.3 *For every integer $k \geq 1$, G_k is triangle-free.*

Proof. This can be seen by induction on k. G_1 is triangle-free since it has a single vertex. G_{k+1} is obtained from the disjoint union of copies of $G_1, G_2, \ldots, G_k$, which by the induction hypothesis is triangle-free, by adding vertices adjacent to an independent set. Indeed each new vertex b in G_{k+1} is adjacent to at most one vertex in each copy of the graphs $G_1, G_2, \ldots, G_k$, hence cannot create a triangle. Thus G_{k+1} is itself triangle-free. ∎

Twincut graphs have unbounded chromatic number, with a similar argument to the one used for Zykov graphs, and the additional twist of finding a rainbow independent set along a branch of the structured tree.

Proposition 5.4 *For every integer $k \geq 1$, we have $\chi(G_k) = k$.*

Proof. The proof is again by induction on k. The case $k = 1$ holds since G_1 is a 1-vertex graph. Now, let $k \geq 1$ and suppose $\chi(G_\ell) = \ell$ for $\ell \leq k$. Fix c a proper coloring of G_{k+1}. In the underlying structured tree T_{k+1}, we will pick a branch which uses k distinct colors. Assume by induction that $v_1, \ldots, v_\ell$ is a path in T_{k+1} starting from the root v_1 such that the colors $c(v_i)$ are all distinct. By construction of G_{k+1}, the children of v_ℓ induce a copy of $G_{\ell+1}$, which is $(\ell + 1)$-chromatic. Thus, there is a child $v_{\ell+1}$ whose color is distinct from $c(v_1), \ldots, c(v_\ell)$, with which we extend the path. Once this process reaches a leaf of T_{k+1}, we obtain a branch b whose vertices use k distinct colors, hence the vertex b, which is connected exactly to this branch, needs one additional color. Thus, c uses at least $k + 1$ colors, so $\chi(G_{k+1}) \geq k + 1$.

Conversely, if we color in G_{k+1} all branch vertices by $k + 1$ and remove them from G_{k+1}, we are left with the disjoint union of all graphs $g(v)$, i.e., copies of $G_1, \ldots, G_k$ which are k-colorable by induction. This yields a $(k + 1)$-coloring of G_{k+1}. ∎

Proposition 5.5 *For every k, the twin-width of G_k is at most $k-1$.*

Proof. Observe that G_1 and G_2 both have twin-width 0, and $G_3 = C_5$ has twin-width 2. Let $k \geq 4$ and assume that the property holds for every $j < k$. Let v be a vertex at level $k-1$ in T_{k+1}. Then, the vertices of $g(v)$ are leaves of T_{k+1}. Furthermore, the leaves of T_{k+1} are in bijection with the branches of T_{k+1} so if ℓ is the leaf corresponding to branch b, we can denote x_b by x_ℓ without any ambiguity. We have that $g(v) = G_k$ has twin-width at most $k-1$ by induction hypothesis. Consider a contraction sequence of G_k achieving $tww(G_k) \leq k-1$. We mimic this contraction sequence inside $g(v)$, but before doing any contraction between ℓ and ℓ', we first contract x_ℓ and $x_{\ell'}$. By doing this, at any point, the maximum red degree of a x_ℓ is at most 2 (two x_ℓ that are merged have all their ancestors in common except one), and the maximum red degree of a vertex of $g(v)$ is its red degree in the original contraction sequence plus one, i.e. at most k. We do this for every vertex v at level $k-1$.

Let T'_{k+1} be the tree we get by starting from T_{k+1} and removing all nodes at level k. The graph we obtain after having contracted all the $g(v)$ corresponds to the realization of this tree, except that for every x_b now has a pending red edge where the pending vertex corresponds to a contracted $g(v)$. Next, consider a vertex w at level $k-2$ in T'_{k+1}. Then, $g(w) = G_{k-1}$ has twin-width at most $k-2$ by induction hypothesis. We once again mimic an optimal sequence inside $g(w)$, except that we first contract the vertices pending from the x_b, then the x_b and only then we do the contraction in $g(w)$. When doing this, the maximum red degree of a pending vertex is 2, the maximum red degree of a x_b is 3, and the maximum red degree of a vertex of $g(w)$ is its red degree in the original contraction sequence plus one, i.e. at most $k-1$. We finally contract the vertex corresponding to $g(w)$ to the vertex pending from the leaf, which creates no red edge, and decreases the red degree of the node corresponding to the contraction of all branch vertices by 1. We do this for all w at level $k-2$. By doing so, we once again remove a level to our tree, while keeping a single pending red edge for every x_b. Iterating this, we can contract G_{k+1} to a single vertex, while keeping the red degree at most k at each step. ∎

In particular, by Theorem 5.2, this means that the twincut graphs are triangle-free graphs with unbounded chromatic number whose structural complexity is close to being optimal. This is also reflected in the value of other width parameters, such as the tree-width or the rank-width.

5.4 Operations that preserve χ-boundedness

We define (yet) another operation on graphs, the *gluing* operation. If G_1, G_2 are graphs with inclusion-wise incomparable vertex sets, let $C = V(G_1) \cap V(G_2)$. Assume that $G_1[C] = G_2[C]$. Let G the the graph such that $V(G) = V(G_1) \cup V(G_2)$, with adjacency as follows.

- $G[V(G_1)] = G_1$,
- $G[V(G_2)] = G_2$,
- There is no edge between $V(G_1) \setminus C$ and $V(G_2) \setminus C$.

We then say that G is obtained by *gluing* G_1 and G_2 *along* C. In this case, we also say that G is obtained by gluing G_1 and G_2 on $|C|$ vertices. If $\mathcal{C}$ is a class of graphs and k is an integer, the *k-gluing closure* of $\mathcal{C}$, denoted by $\mathcal{C}_g^k$ is the class of graphs we can obtain by iteratively gluing graphs of $\mathcal{C}$ on at most k vertices. In [9], the authors show that for every k, if $\mathcal{C}$ is (polynomially) χ-bounded, then $\mathcal{C}_g^k$ is also (polynomially) χ-bounded. In the same paper, they prove that the closure of a χ-bounded class under substitution is χ-bounded, and that substitutions further preserve polynomial χ-boundedness. Trying to merge these two operations, they posed the following problem, also mentioned in [18]:

Problem 5.6 *Is the closure of a χ-bounded class under substitution and gluing along a bounded number of vertices also χ-bounded?*

This problem can be simplified to a seemingly trivial case, where the base class consists only of the complete graphs on one and two vertices, where we can glue on at most 2 vertices, under the extra constraint that these two vertices are not adjacent, and where we can only substitute a vertex by two non-adjacent vertices. Call $\mathcal{T}$ the resulting class. It is immediate that $\mathcal{T}$ is the class of graphs that verify the following property: all their induced subgraphs with at least 3 vertices have false twins (two non-adjacent vertices with the same neighborhood), or an edgeless vertex cutset of size at most two. Since $\mathcal{T}$ is hereditary and since the triangle is not in $\mathcal{T}$, all graphs in $\mathcal{T}$ are triangle-free. Hence, to give a negative answer to Problem 5.6, it suffices to prove that graphs of $\mathcal{T}$ can have arbitrarily high chromatic number.

Somewhat surprisingly, the twincut graphs are in $\mathcal{T}$, which gives a negative answer to Problem 5.6, as well as it highlights even more the low structural complexity of the twincut graphs.

Proposition 5.7 *The graphs G_k are in $\mathcal{T}$.*

We more generally show that $\mathcal{T}$ is closed under the realization of structured trees, which immediately implies Proposition 5.7.

Lemma 5.8 *Let (T, g) be a structured tree such that every $g(v)$ is in $\mathcal{T}$. Then $R(T, g) \in \mathcal{T}$.*

Proof. For a node v of T, let $T(v)$ be the subtree rooted at v, i.e., the subtree consisting of all descendants of v. Equipped with the restriction of g, $T(v)$ is a structured tree. We prove by induction on T, starting from the leaves, that for all nodes v, the realization $R(T(v), g)$ is in $\mathcal{T}$. For the sake of brevity, let us denote this realization of a subtree by $R_v = R(T(v), g)$.

If v is a leaf, then R_v is simply an edge, which is in $\mathcal{T}$. Let now v be an internal node with children $u_1, \dots, u_\ell$, and assume that each R_{u_i} is in $\mathcal{T}$. Recall also that $g(v)$ is assumed to be in $\mathcal{T}$. We construct R_v as follows. First, in each R_{u_i}, create a copy u'_i of u_i by substituting u_i with a stable set of size two, and call R'_{u_i} the resulting graph. Next, take $g(v)$ and add to it an isolated vertex standing for v. We then glue each R'_{u_i} successively with this graph, by identifying u'_i with v, and identifying the occurrences of u_i in R'_{u_i} and in $g(v)$. This corresponds to gluing along a stable set of size two. Hence, we constructed R_v starting from $g(v), R_{u_1}, \dots, R_{u_\ell}$, by substituting with and gluing on stable sets of size at most two, thereby proving that $R_v \in \mathcal{T}$. ∎

6 Factoring permutations into separable ones

The scope of twin-width extends beyond graphs. In particular, one can define the twin-width of a permutation. A seminal result by Guillemot and Marx [13] is that strict ("hereditary") classes of permutation have bounded twin-width – this paper actually preceded and inspired the definition of twin-width. Using a variant of delayed decompositions, we show that every strict class of permutations is contained in a bounded power of the class of separable permutations. This structural result illustrates the versatility of both twin-width and delayed decompositions. I will explain the result we obtain, but I will not give any proof in this report.

6.1 Patterns

If $n \leq m$ are two integers, we say that a permutation $\pi \in S_n$ is a *pattern* of $\sigma \in S_m$ if there is a strictly increasing function f from $[n]$ to $[m]$ such that $\pi(i) < \pi(j)$ if and only if $\sigma(f(i)) < \sigma(f(j))$ for all $i < j$ in $[n]$. Another way to characterize patterns is to associate to every permutation $\sigma \in S_n$ its $n \times n$ matrix $A(\sigma) = (a_{ij})$ in which a_{ij} is equal to 1 whenever $j = \sigma(i)$ and equal to 0 otherwise. Then, π is a pattern of σ if and only if $A(\pi)$ is a submatrix of $A(\sigma)$. Informally, patterns in permutations correspond to induced subgraphs in graphs. For instance, 12345 is a pattern of σ if σ contains an increasing subsequence of length five. A crucial achievement in permutation patterns is the Guillemot-Marx algorithm [13] which decides if a permutation π is a pattern of σ in time $f(\pi).|\sigma|$, where $|\sigma|$ is the length of σ.

Patterns offer a very easy way to analyze the complexity of a permutation: roughly speaking, a permutation is simple if it does not contain a given small pattern. It is more convenient here to speak of *classes of permutations* which are assumed closed with respect to patterns. In [13], Guillemot and Marx introduced a measure of complexity of a permutation, which they simply call width. Twin-width is a generalization of this notion to other structures, including graphs. Although their presentation is quite different, one may check that the *width* of a permutation σ in the work of Guillemot and Marx is precisely the twin-width of σ. With this in mind, a major result of Guillemot and Marx— the duality between patterns and twin-width—can be restated as follows.

Theorem 6.1 ([13, Theorem 4.1]) *For any permutation τ, there is $c_\tau \in \mathbb{N}$ such that if σ avoids τ as pattern, then σ has twin-width at most c_τ.*

The existence of a gap between the class of all permutations and any smaller class is illustrated by the Marcus-Tardos theorem [15] (answering the Stanley-Wilf conjecture): every strict class of permutations has at most exponential growth, whereas the class of all permutations reaches factorial growth.

Strict classes of permutations are therefore very simple, both from the algorithmic and counting point of view. The next natural question is then to construct them from a basic class using some operations. In the case of permutations, one of the simplest operations is the product (or composition). Furthermore, *separable permutations*, which are the permutations avoiding the patterns 2413 and 3142 constitute a very restricted class (as they are equivalently the closure under substitutions of the permutations 12 and 21). In that sense, they can be seen as the permutation counterpart of cographs. We can now state our result.

Theorem 6.2 *For every permutation τ there exists k_τ such that every permutation avoiding τ as a pattern is a product of at most k_τ separable permutations.*

Note that this result can be seen as a constructive version of the Marcus-Tardos theorem [15] (which is used in the proof). To prove this result, thanks to Theorem 6.1, it suffices to prove that any permutation of bounded twin-width can be written as a product of some bounded number of separable permutations. To do so, one of the main tools we use is a generalization of delayed decompositions to permutations, so that we only have to analyze the "quotient permutations". I will not give more details about the proof of Theorem 6.2 in this report. However, we are currently finalizing the write up of a paper proving Theorem 6.2.

Acknowledgements

I am deeply grateful to Stéphan Thomassé for dedicating me so much time, for always being available for discussions and reflections, and for the invaluable knowledge and personal growth he has imparted. Additionally, I extend my thanks to the entire MC2 team for creating a delightful environment during my internship, and for the wonderful moments we shared together.

References

[1] Marthe Bonamy and Michał Pilipczuk. Graphs of bounded cliquewidth are polynomially χ-bounded. *Advances in Combinatorics*, (8):21pp, 2020.

[2] Édouard Bonnet. Problem 3. *Open problems for the second 2022 Barbados workshop,* `https://web.math.princeton.edu/~tunghn/2022openproblems.pdf`, 2022.

[3] Édouard Bonnet, Romain Bourneuf, Julien Duron, Colin Geniet, Stéphan Thomassé, and Nicolas Trotignon. A tamed family of triangle-free graphs with unbounded chromatic number, 2023.

[4] Edouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width II: small classes. In *ACM-SIAM Symposium on Discrete Algorithms (SODA21)*, Alexandria, United States, January 2021. 37 pages, 9 figures.

[5] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: Max Independent Set, Min Dominating Set, and Coloring. In Nikhil Bansal, Emanuela Merelli, and James Worrell, editors, *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, volume 198 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 35:1–35:20, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[6] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: Tractable FO model checking. *J. ACM*, 69(1), nov 2021.

[7] Romain Bourneuf and Stéphan Thomassé. Bounded twin-width graphs are polynomially χ-bounded, 2023.

[8] Marcin Briański, James Davies, and Bartosz Walczak. Separating polynomial χ-boundedness from χ-boundedness, 2022.

[9] Maria Chudnovsky, Irena Penev, Alex Scott, and Nicolas Trotignon. Substitution and χ-boundedness. *Journal of Combinatorial Theory, Series B*, 103(5):567–586, 2013.

[10] Maria Chudnovsky, Neil Robertson, Paul Seymour, and Robin Thomas. The strong perfect graph theorem. *Annals of Mathematics*, 164:51–229, 2006.

[11] Blanche Descartes. Solution to advanced problem no. 4526. *Amer. Math. Monthly*, 61, 1954.

[12] Zdeněk Dvořák and Daniel Král'. Classes of graphs with small rank decompositions are χ-bounded. *European Journal of Combinatorics*, 33(4):679–683, 2012.

[13] Sylvain Guillemot and Dániel Marx. Finding small patterns in permutations in linear time. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, pages 82–101, 2014.

[14] L. Lovász. Normal hypergraphs and the perfect graph conjecture. *Discrete Mathematics*, 2(3):253–267, 1972.

[15] Adam Marcus and Gábor Tardos. Excluded permutation matrices and the Stanley-Wilf conjecture. *Journal of Combinatorial Theory - Series A*, 107(1):153–160, July 2004.

[16] Jan Mycielski. Sur le coloriage des graphes. *Colloq. Math.*, 3:161–162, 1955.

[17] Michał Pilipczuk and Marek Sokołowski. Graphs of bounded twin-width are quasi-polynomially χ-bounded, 2022.

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

[19] Alexander A. Zykov. On some properties of linear complexes. *Amer. Math. Soc. Transl.*, 79:163–188, 1952.

A Proof of Proposition 3.4

We will express our result in the language of ordered graphs (graphs with a total order on vertices). An RMP for an ordered graph must respect the order, that is the parts of the partition must consist of consecutive vertices.

The proof of Proposition 3.4 is heavily based on the following lemma about the clique number. Given an ordered graph G and an RMP $\mathcal{P} = (V_1, \dots, V_k)$, we denote by $G/\mathcal{P}$ the ordered graph obtained by contracting all parts of $\mathcal{P}$, i.e. $G/\{V_1, \dots, V_k\}$. This is a particular transversal minor of G, with the property that $\chi(G) \leq \chi(G/\mathcal{P})$. A class $\mathcal{C}$ of ordered graphs is *h-free* if it does not contain some ordered graph on h vertices.

Lemma A.1 *There exists a function ϕ such that for every h-free class $\mathcal{C}$ and every ordered graph G with a $\mathcal{C}$-right module partition $\mathcal{P}$, we have $\omega(G/\mathcal{P}) \leq \phi(\omega(G), h)$*

Proof. The proof is by induction on $\omega := \omega(G)$ and h. If $h = 1$ or $\omega = 1$ then G is edgeless, so we can set $\phi(x, 1) = \phi(1, y) = 1$. Now, $\omega \geq 2$ and $h \geq 2$, and we assume that we proved the existence of $\phi(\omega - 1, h)$ and $\phi(\omega, h - 1)$. We denote $\mathcal{P} = (V_1, \dots, V_k)$ and consider an ordered graph H on vertices $v_1, \dots, v_h$ which is not in $\mathcal{C}$. Observe that we can restrict ourselves to the case where $G/\mathcal{P}$ is a clique, as we can only consider a maximal clique of $G/\mathcal{P}$.

Thus, for every $i < j \in [k]$, there is an edge between V_i and V_j. Let D_k be a subset of V_k of minimal size such that there is an edge between V_i and D_k for every $i < k$. By minimality of D_k, for every $d \in D_k$ there exists $i_d \in [k-1]$ such that there is an edge between V_{i_d} and d (and thus d dominates V_{i_d}) but no edge between V_{i_d} and $D_k \setminus \{d\}$. We consider two cases depending on the size of D_k.

1. If $|D_k| > \phi(\omega, h-1)+1$, we show that we reach a contradiction. Select some $x \in D_k$ and consider the set $\mathcal{P}'$ of $|D_k| - 1$ parts V_{i_d} as defined above, except for the part V_{i_x} which is not selected in $\mathcal{P}'$. Since $G/\mathcal{P}'$ is a clique of size at least $\phi(\omega, h - 1) + 1$, it follows by induction that $(G[\cup \mathcal{P}'], \mathcal{P}')$ contains all ordered graphs of size $h - 1$ as transversal minors. In particular, the ordered graph $H' = H \setminus v_h$ is a transversal minor of $\mathcal{P}'$. If v_h is isolated in H, we reach a contradiction since $H' \cup x$ is isomorphic to H and is a transversal minor of $(G, \mathcal{P})$. Otherwise, observe that one can extend H' in all possible ways as a transversal minor of $\mathcal{P}$ by selecting some vertices in D_k. Indeed, every V_{i_d} corresponds to a vertex d in D_k which is only joined to V_{i_d}. We can then select vertices in D_k to extend H' to H, a contradiction.

2. If $|D_k| \leq \phi(\omega, h - 1) + 1$. For $d \in D_k$, let S_d be the set of neighbours of d in G. In particular, $\omega(G[S_d]) \leq \omega - 1$. Furthermore, $\mathcal{P}$ induces by restriction a $\mathcal{C}$-RMP $\mathcal{P}_d$ of $G[S_d]$. We then have $\omega(G[S_d]/\mathcal{P}_d) \leq \phi(\omega - 1, h)$. By taking the union over all $d \in D_k$, we deduce $\omega(G/\mathcal{P}) \leq \phi(\omega - 1, h) \cdot (\phi(\omega, h - 1) + 1) + 1$ (the additional +1 stands for the last class V_k which is not dominated).

Therefore, we can choose $\phi(\omega, h) = \phi(\omega - 1, h) \cdot (\phi(\omega, h - 1) + 1) + 1$. ∎

We are now ready to prove Proposition 3.4. If $\mathcal{C}$ has a χ-bounding function f, by the fact that the class of all graphs is not χ-bounded, there is a graph H of size h which is not in $\mathcal{C}$. Consider now any graph G in $RM(\mathcal{C})$ with clique number ω and $\mathcal{C}$-RMP $\mathcal{P}$. We have

$$\chi(G) \leq \chi(G/\mathcal{P}) \leq f(\omega(G/\mathcal{P})) \leq f(\phi(\omega(G), h)).$$

Therefore the function $f(\phi(\omega(G), h))$ is χ-bounding for $RM(\mathcal{C})$.

B Proof of Claim 4.8

As we will perform some operations on our graphs (such as deleting edges and contracting subsets of vertices), we show in the next results how mixed zones are affected. These are technical results with short and easy proofs. We do not give them here, but they can be found in [7]. Let M be a $01*$-matrix (not necessarily an adjacency matrix) with exactly two row blocks R, R' and two columns blocks C, C', each of size at least two.

Lemma B.1 *If all four zones of M are mixed, there is a corner intersecting all zones.*

The *contraction* $M' = M/\{R, R'; C, C'\}$ is the 2×2-matrix obtained by keeping a single value x for each of the 4 zones, with the following rule: x is the maximum value of the zone according to the order $0 < 1 < *$. Thus, we get $*$ as soon as there exists a $*$, and we get 0 only if the zone is full 0.

Lemma B.2 *If M' is mixed, then M is mixed.*

We keep the same notations as before, and assume moreover that none of the four zones of M is mixed (in particular M has no value $*$). The *horizontal-deletion* M_H of M is the matrix obtained from M by setting all values to 0 in each zone which is not vertical (or equivalently each zone which is horizontal and non-constant). We similarly define the *vertical-deletion* M_V.

Lemma B.3 *If M_H is mixed, then M is mixed.*

We now move to the proof of Claim 4.8

Proof. Assume for contradiction that we can find a sequence of local modules $L'_1, \ldots, L'_t$ in L'_{ri}, each of them containing a non empty subset of vertices $W_1, \ldots, W_t$, such that the graph $Q = H'_{\to,ri}/\{W_1, \ldots, W_t\}$ has a $d-1$-almost mixed minor. The vertices of Q are denoted $W = \{w_1, \ldots, w_t\}$, where W_i is contracted to w_i. Moreover, there exist two partitions of W into consecutive blocks of vertices $(R_1, \ldots, R_{d-1})$ and $(C_1, \ldots, C_{d-1})$ such that Q is mixed on the zone $[R_i, C_j]$ whenever $i \neq j$ (thus all R_i and C_j have size at least 2). We now show how to "lift" these partitions to G in order to get a contradiction.

Consider any partition $\mathcal{R}' = (R'_1, \ldots, R'_{d-1})$ of I (where parts consist of consecutive local modules) satisfying that $w_i \in R_j$ implies $L'_i \subseteq R'_j$. Similarly $\mathcal{C}' = (C'_1, \ldots, C'_{d-1})$ partitions I and $w_i \in C_j$ implies $L'_i \subseteq C'_j$. We now extend the partitions $\mathcal{R}', \mathcal{C}'$ of I to the whole vertex set V of G by first setting $R'_1 := R'_1 \cup \{v_1, \ldots, v_{s-1}\}$ and $C'_1 := C'_1 \cup \{v_1, \ldots, v_{s-1}\}$, and then adding a new part $\{v_{t+1}, \ldots, v_n\} = R'_d = C'_d$ to both $\mathcal{R}'$ and $\mathcal{C}'$. These new partitions are called $\mathcal{R}$ and $\mathcal{C}$. Observe that if we were working with $H'_{\to,le}$, we would have added the part $R'_0 = C'_0 = \{v_1, \ldots, v_{s-1}\}$ to both $\mathcal{R}'$ and $\mathcal{C}'$ and extended the parts R'_{d-1} and C'_{d-1} by adding $\{v_{t+1}, \ldots, v_n\}$.

We now show that $\mathcal{R}$ and $\mathcal{C}$ form a d-almost mixed minor for G, which will be our contradiction. We need to focus on two points: the added parts R'_d, C'_d should be mixed with respect to the others, and the original mixed zones $[R_i, C_j]$ of Q should yield mixed zones $[R'_i, C'_j]$ of G. We separate the two arguments:

- Consider first some zone $[R'_i, C'_d]$ where $i < d$ (similar argument for $[R'_d, C'_i]$). By the fact that R_i contains two vertices w_a, w_b, the part R'_i contains the right local modules L'_a, L'_b (where $a < b$). Focus now on the vertex v_j which is the first vertex of L'_b, and note that $v_{j-1} \in R'_i$ since $L'_a \subseteq R'_i$. Since L'_b is a right local module, there exists v_k, where $k > t$ such that v_k is differently joined to v_{j-1} and v_j. Recall that v_n is joined in the same way to v_{j-1} and v_j. Since $v_{j-1}, v_j \in R'_i$ and $v_k, v_n \in C'_d$, they witness the fact that $[R'_i, C'_d]$ is mixed.

- Now, consider any zone $[R'_i, C'_j]$ where $i, j < d$ and $i \neq j$. If the zone $[R_i, C_j]$ contains a $*$ (i.e. some w_a both belongs to R_i and C_j), then $[R'_i, C'_j]$ also contains a $*$. Otherwise, by Lemma 4.1, R_i contains two vertices w_a, w_b and C_j contains two vertices w_c, w_d such that $\{w_a, w_b\}, \{w_c, w_d\}$ is a corner. Moreover, since there is no $*$ value, we have $a < b < c < d$ or $c < d < a < b$. Without loss of generality, we assume $a < b < c < d$. By Lemma B.2, the restriction of the adjacency matrix of $H'_{\rightarrow,ri}$ on $[W_a \cup W_b, W_c \cup W_d]$ is mixed since its contraction is the corner $\{w_a, w_b\}, \{w_c, w_d\}$. So the submatrix $[L'_a \cup L'_b, L'_c \cup L'_d]$ is also mixed. By definition of $H'_{\rightarrow,ri}$, if L'_a (or L'_b) is not a module with respect to L'_c (or to L'_d), then the zone $[L'_a, L'_c]$ is set to 0. In other words, the adjacency matrix of $H'_{\rightarrow,ri}$ restricted to $[L'_a \cup L'_b, L'_c \cup L'_d]$, is the horizontal-deletion of the adjacency matrix of G. Thus the zone $[R'_i, C'_j]$ is mixed by Lemma B.3.

∎

C Proof of Claim 4.9

Proof. Assume for contradiction that we can find a coarsening $W'_1, \ldots, W'_{2d}$ of L'_{ri} which forms a $2d$-almost mixed minor of $H'_{\rightarrow,ri}$. We now set $W_i = W'_{2i-1} \cup W'_{2i}$ for all $i = 1, \ldots, d$. By Lemma 4.1 every (mixed) zone $[W_i, W_j]$ with $i \neq j$ of $H'_{\rightarrow,ri}$ contains a corner $\{w_a, w_b\}, \{w_c, w_d\}$. By Lemma B.1, we can assume that w_a, w_b, w_c, w_d belong respectively to $W'_{2i-1}, W'_{2i}, W'_{2j-1}, W'_{2j}$, hence to respective distinct local modules L'_a, L'_b, L'_c, L'_d. We assume without loss of generality that $i < j$, and thus $a < b < c < d$. By definition of $H'_{\rightarrow,ri}$, if L'_a (or L'_b) is not a module with respect to L'_c (or to L'_d), then the zone $[L'_a, L'_c]$ is set to 0. In other words, the adjacency matrix of $H'_{\rightarrow,ri}$ restricted to $[L'_a \cup L'_b, L'_c \cup L'_d]$, is the horizontal-deletion of the adjacency matrix of G. Hence the zone $[W_i, W_j]$ is mixed in G by Lemma B.3. Thus G restricted to $W_1, \ldots, W_d$ has a d-almost mixed minor, a contradiction.

∎