THÈSE

Pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ GRENOBLE ALPES

École doctorale : MSTII - Mathématiques, Sciences et technologies de l'information, Informatique
Spécialité : Mathématiques et Informatique
Unité de recherche : Laboratoire des Sciences pour la Conception, l'Optimisation et la Production de Grenoble

Propriétés structurelles et géométriques des graphes fortement symétriques

Structural and geometrical properties of highly symmetric graphs

Présentée par :

Ugo GIOCANTI

Direction de thèse :

Louis ESPERET — Directeur de thèse
DIRECTEUR DE RECHERCHE, CNRS DELEGATION ALPES

Stephan THOMASSÉ — Co-directeur de thèse
PROFESSEUR DES UNIVERSITES, ENS de Lyon

Rapporteurs :

JOHANNES CARMESIN
FULL PROFESSOR, TECHNISCHE UNIVERSITÄT BERGAKADEMIE
AGELOS GEORGAKOPOULOS
FULL PROFESSOR, UNIVERSITY OF WARWICK

Thèse soutenue publiquement le **9 juillet 2024**, devant le jury composé de :

EMMANUEL JEANDEL, PROFESSEUR DES UNIVERSITES, UNIVERSITE DE LORRAINE	Président
LOUIS ESPERET, DIRECTEUR DE RECHERCHE, CNRS DELEGATION ALPES	Directeur de thèse
STEPHAN THOMASSÉ, PROFESSEUR DES UNIVERSITES, ENS DE LYON	Co-directeur de thèse
JOHANNES CARMESIN, FULL PROFESSOR, TECHNISCHE UNIVERSITÄT BERGAKADEMIE	Rapporteur
AGELOS GEORGAKOPOULOS, FULL PROFESSOR, UNIVERSITY OF WARWICK	Rapporteur
KOLJA KNAUER, MAITRE DE CONFERENCES HDR, UNIVERSITAT DE BARCELONA	Examinateur
FRANCIS LAZARUS, DIRECTEUR DE RECHERCHE, CNRS DELEGATION ALPES	Examinateur
FRANÇOIS DAHMANI, PROFESSEUR DES UNIVERSITES, UNIVERSITE GRENOBLE ALPES	Examinateur

Acknowledgements

First of all, I would like to thank my two PhD advisors Louis Esperet and Stéphan Thomassé. Louis, it was a real pleasure working with you during this three years in Grenoble, both from a scientific and a human point of view. Thank you very much for your availability, your kindness and your great advices, I really feel that I learnt a lot (both scientifically and about myself) by your side. Stéphan, we mostly worked together before the thesis, and during the first year, and I thank you for having shared your motivation and your incredible scientific intuition with me. I also thank you for the discussions we had when I was a student of the ENS de Lyon, which significantly encouraged me doing theoretical computer science.

I am grateful to Johannes Carmesin and Agelos Georgakopoulos who accepted to be reviewers of this thesis, and for their serious, encouraging and very instructive reports, which significantly improved the quality of the manuscript. I also thank François Dahmani, Emmanuel Jeandel, Kolja Knauer and Francis Lazarus for their presence in the jury, and for their valuable comments and questions.

I would like to thank all of my coauthors, it has been a pleasure to work with each of you. In particular, there are a few colleagues that I would like to mention, with who I had the pleasure to work during the last few years, and whose interactions with have been particularly enjoyable: Édouard Bonnet, Jérémie Chalopin, Victor Chepoi, Solène Esnay, Clément Legrand-Duchesne, Étienne Moutot, Jean-Sébastien Sereni and Matěj Stehlík. I forgot to mention many people, and thus I would like to thank more generally everyone I had the pleasure to scientifically interact with during and before my thesis, a full list would be to long to fit here.

These three years in Grenoble have been a great moment, both inside and outside G-SCOP lab. I would like to thank every member from the "Optimisation Combinatoire" team, and more generally from G-SCOP, we shared great moments and I will miss all of you. In particular, I thank Félix Klingelhoeffer and Pierre Hoppenot with who I had the pleasure to share my office, and also everyone who accepted to follow me and play in the G-SCOP Union Football team during three years! There are too much things and people I will miss: the fencing club of *la Rapière*, *la coinche du Sub'* and more generally *le Sub'*, the breaks in *la cafét'*, or in *le Sun*, the incredible proximity with the mountains, the *FCG* games and so many other things...

Clearly, I would never have been interested about mathematics and computer science without great teachers during high school (and even before), and I would like to credit three of them who particularly encouraged me: Alexis Barnier, Jean-Luc Joly and Thierry Audibert. Eventually, I would like to thank my family and especially my brother and my parents, who always supported me even though they were completely outsiders from the academic world, and who always encouraged all my choices.

Introduction

Many statements in combinatorics or group theory have the following form: "if X is a combinatorial object satisfying a property (P), then X has the following structure: ...". Among them, one can think about Robertson-Seymour Structure Theorem in graph theory, which explicitly describes the structure of graphs excluding a fixed minor, or about Maschke's planarity theorem, which establishes the complete list of the finite planar Cayley graphs. In this thesis, I expose a number of results of this type whenever one considers objects X which are highly symmetric. More precisely, I focus in Chapter 1 on locally finite graphs that are *quasi-transitive*, i.e., that admit, up to applying an automorphism, only finitely many types of vertices (see Chapter 1 for a formal definition). The class of quasi-transitive graphs is not only interesting because it contains the rich class of *Cayley graphs*, but also because it allows to avoid the rigidity present in algebraic structures, and thus one can relax and adapt many reasonings. Quasi-transitive graphs behave particularly well with *canonical tree-decompositions*, a useful tool allowing to decompose in a unique way a given graph into simpler pieces in a tree-like fashion. In particular, I present in Chapter 1 proofs of the following decomposition results:

- Every planar locally finite 3-connected quasi-transitive graph admits a canonical tree-decomposition whose edge-separations correspond to cycle-separations, and where every part admits a vertex-accumulation-free planar embedding.

- Every locally finite quasi-transitive graph excluding the countable clique as a minor admits a canonical tree-decomposition whose parts are finite or planar.

I will also provide a number of graph theoretic applications of the second result, and relate it to recent results lying at the intersection of graph theory and coarse geometry. Chapter 2 is devoted to the study of some specific classes of finitely generated groups. I introduce concepts and questions from the field of symbolic dynamics of groups, and try to establish some connections with the notions and results studied in Chapter 1. In particular, I will provide some applications of these results to study questions related to the domino problem, and the aperiodicity of tilings in groups.

Even though many notions in this manuscript are connected to concepts that originally come from algebra or geometric group theory, I would like to point out that most (if not all) of the proofs I present rely on combinatorial arguments, and this manuscript is first intended to people that are not experts in group theory. In particular, Chapter 1 only deals with graph theoretic concepts, and should be understandable to the reader whose knowledge in algebra does not go further than the definitions of a group and of a group action. Even though many questions and results from Chapter 1 are deeply connected to concepts coming from group theory, I did my best to postpone every definition related to group theoretic objects to Chapter 2, including the definitions of Cayley graphs and of group presentations.

Most of the results and proofs presented in this manuscript come from the papers [EGLD23, EGLD23, EG24b, AEGH24, EGM24, Gio24] (some of them are still in preparation), on which I worked during my PhD. However, their content is not always presented the exact same way here, and there is not necessarily a one-to-one correspondence between an article and a section. More precise references to these papers and their content are given at the beginning of each section.

Contents

Chapter 1

Structure of locally finite quasi-transitive graphs

Notation: In the whole chapter, we will use capital latin letters to denote graphs, and capital greek letters to denote groups. In particular, the letter G will always be used here to denote a graph.

For every $n \in \mathbb{N}$, we will use the notation $[n]$ to denote the set of integers $\{1, \ldots, n\}$.

1 Introduction

In an early work from 1896, Maschke [Mas96] established the complete list of all finite planar Cayley graphs. A group admitting a planar Cayley graph is called a *planar group*, and Maschke showed in the same paper that the planar finite groups are exactly the countable groups of isometries of the 2-dimensional sphere $\mathbb{S}^2$. Based on the works of Wilkie [Wil66] and MacBeath [Mac67], Zieschang, Volgt and Coldewey [ZVC80] established the complete list of infinite countable groups admitting a planar simplified Cayley complex (see Section 13 for a definition), which are exactly the countable groups for which there exists a planar Cayley graph with a *vertex-accumulation-free* planar embedding. We also refer to [MS83, Section III.5] for a complementary work on such groups. The study of infinite planar Cayley graphs has attracted a lot of attention, and many other structural properties of such graphs have been studied. Among other results, planar groups have been proved to be finitely presented [Dro06], and thus by a result of Dunwoody [Dun85] they are also accessible. Many specific techniques have been developed to study their general structure [Dro06, Dun09, GH15, GH23].

One of the basic properties of Cayley graphs is that they are (vertex-)transitive. In this first chapter, we will focus on the class of quasi-transitive graphs which is slightly more general than the class of transitive graphs. A graph is *quasi-transitive* if the action of its automorphism group on its vertex set induces only finitely many orbits. In particular, note that every finite graph is quasi-transitive, hence this notion becomes interesting only when studying infinite graphs. Intuitively, quasi-transitive graphs have to be thought as "graphs having many symmetries". They appear in a natural way in the celebrated Švarc–Milnor lemma [Š55, Mil68], which in the special case of locally finite graphs implies that if a finitely generated group Γ induces a quasi-transitive group action on a graph G such that every vertex has a finite stabilizer in Γ, then Γ is finitely generated and every locally finite Cayley graph of Γ is quasi-isometric to G.

A central concept in structural graph theory is the notion of *tree-decompositions*, which appears in particular in the statement of the Graph Minor Structure theorem of Robertson Seymour [RS03]. The graph minor structure theorem basically states that for every finite graph H, there exists a constant $g \geqslant 1$ such that every graph excluding H as a minor admits a tree-decomposition of bounded adhesion whose parts are graphs that "almost embed" in the closed orientable surface of genus g. More generally, tree-decompositions are a convenient tool to decompose a graph in smaller parts glued one to the other in a tree-like way. A natural question is the following: if the graph G we study has non-trivial symmetries, can we make these symmetries apparent in the tree-decomposition? For example, do graphs avoiding a fixed minor have a tree-decomposition as above, but with the additional constraint that the decomposition is *canonical*, i.e., invariant under the action of the automorphism group of G?

In Sections 2 and 3, we will introduce basic concepts related to quasi-transitive graphs and canonical tree-decompositions and present some general results about them. We will survey in Section 4 many known characterisations of quasi-transitive graphs of bounded treewidth. Then we will explain in Section 5 how the results from [Ham15, Ham18b] imply a structure theorem for planar quasi-transitive graphs, which can be seen as an analogue of the structure theorem from [Dro06] for general planar groups. In Section 6, we will introduce tangles, a notion introduced in the tenth paper of the Graph Minor series [RS91], and playing a central role in the proof of the Graph Minor Structure theorem. We will present there the main ideas from [Gro16] which uses tangles as a building block to prove a general structure theorem for decomposing finite 3-connected graphs into parts of higher connectivity. In particular, we will explain how these ideas can be adapted when working on locally finite quasi-transitive graphs. We will crucially rely on the results from this part to prove in Section 7 the main result from this first chapter, that is a structure theorem for locally finite quasi-transitive graphs that exclude any countable graph (and thus the countable infinite clique K_∞) as a minor. This result can be thought as a tailored version of the Graph Minor Structure Theorem of Robertson and Seymour [RS03] in the special case of locally finite quasi-transitive graphs, and its proof is mainly based on a combination of results of Thomassen [Tho92] and Grohe [Gro16]. We will present applications of this structure theorem in Section 8 and will then discuss in Section 9 different properties that have been recently studied and that generalize minor-exclusion in quasi-transitive graphs.

The proofs presented in Section 5 are unpublished results for which I am the sole author. The results of Sections 6, 7, 8 together with their proofs come from the paper [EGLD23], co-authored with Louis Esperet and Clément Legrand-Duchesne, while the content of Section 9 mainly comes from the paper [EG24a], a joint work with Louis Esperet. Section 10 contains a discussion about some of the questions we asked at the end of [EGLD23], and the results presented or stated are joint work with Tara Abrishami, Louis Esperet and Matthias Hamann [AEGH24].

2 Quasi-transitive graphs and quasi-isometries

2.1 Graphs

The graphs we will consider in this manuscript will always be *simple*, i.e., loopless without multi-edges (except for Cayley graphs, which will be be defined in Chapter 2), and unless specified they will always be connected. In this thesis, a simple graph G consists in a pair $(V(G), E(G))$ where $V(G)$ denotes the *vertex set* of G and $E(G) \subseteq \binom{V(G)}{2}$ denotes its *edge-set*. We will often use the notation xy instead of $\{x, y\}$ to denote the edges of a graph. For every vertex $v \in V(G)$, we will use the notation $\deg_G(v)$ (or simply $\deg(v)$ if the context is clear) to denote its *degree*. Unless specified, we will only work on *locally finite* graphs, that is graphs in which every vertex has a finite degree. In particular, when they are connected, such graphs have countably many vertices and edges.

We equip a graph G with its shortest-path metric d_G. For any set of vertices $X \subseteq V(G)$, the *neighborhood* of X in G is denoted by

$$N_G(X) := \{u \in V(G) \setminus X : \exists v \in X, uv \in E(G)\}.$$

When the graph G is clear from the context we will drop the subscript and write $N(X)$ instead of $N_G(X)$. For each $v \in V(G)$, we set $N_G(v) := N_G(\{v\})$. For every $k \in \mathbb{N}$, $X \subseteq V(G)$, we also let $B_{G,k}(X) := \{u \in V(G) : \exists v \in X, d_G(u, v) \leqslant k\}$ or simply $B_k(X)$ when the context is clear be the *ball of radius* k around v. Again, for each $v \in V(G)$ and $k \in \mathbb{N}$, we set $B_k(v) := B_k(\{v\})$. For every graph G and every subset of vertices $X \subseteq V(G)$, we denote by $G[X]$ the *subgraph of G induced by X*, which is the graph with vertex set X whose edge set consists of all the pairs uv such that $uv \in E(G)$. We let $G - X := G[V(G) \setminus X]$. For every $E' \subseteq E(G)$, we also let $G - E' := (V(G), E(G) \setminus E')$. The *girth* (length of a smallest cycle, if it exists) of a graph G is denoted by $\mathrm{girth}(G) \in \mathbb{N} \cup \{\infty\}$, and its *diameter* (the supremum of the distances between pairs of vertices of G) is denoted by $\mathrm{diam}(G) \in \mathbb{N} \cup \{\infty\}$.

A *proper vertex-coloring* of G is a mapping $c : V(G) \to \mathcal{C}$, where $\mathcal{C}$ denotes a (possibly infinite) set of colors such that for every $uv \in E(G)$, $c(u) \neq c(v)$. A *proper edge-coloring* of G is a mapping $c : E(G) \to \mathcal{C}$ such for every two distinct edges $e, e' \in E(G)$ with a common endpoint, $c(e) \neq c(e')$. The *line graph* $L(G)$ of G is the graph with vertex set $E(G)$ and where we add an edge $ee' \in E(L(G))$ exactly when e and e' are distinct edges with a common endpoint. We denote with $\chi(G) \in \mathbb{N} \cup \{\infty\}$ the *chromatic number* of G, that is the infimum over the number of colors required in any proper vertex-coloring of G. Similarly we denote with $\chi'(G) \in \mathbb{N} \cup \{\infty\}$ the *chromatic index* of G, that is the infimum over the number of colors required in any proper edge-coloring of G. As we will mainly consider in this manuscript graphs with bounded degree, their chromatic number and chromatic index will always be finite.

For each $n \in \mathbb{N}$, we let K_n denote the complete graph with vertex set $[n]$. We also let K_∞ be the *countable clique* (sometimes also denoted by $K_{\aleph_0}$ in the literature), that is the infinite complete graph with vertex set $\mathbb{N}$.

2.2 Minors and models

Given two graphs G, H, we say that H is a *minor* of G if it can be obtained from G after removing some vertices and edges, and contracting edges. A *model* of H in G is a family $(V_v)_{v\in V(H)}$ of pairwise-disjoint vertex subsets of G such that each V_v induces a nonempty connected subgraph of G, and for each $uv \in E(H)$, there exists $u' \in V_u, v' \in V_v$ such that $u'v' \in E(G)$. Note that H is a minor of G if and only if there is a model of H in G. When $V(H) \subseteq V(G)$, a model $(V_v)_{v\in V(H)}$ of H in G is said to be *faithful* if for each $v \in V(H), v \in V_v$. H is a *faithful minor* of G if it admits a faithful model in G. For example, if H is a subgraph of G, the set $(\{v\})_{v\in V(H)}$ constitutes a faithful model of H in G.

2.3 Connectedness

For every $k \geqslant 0$, a graph G is *k-connected* if it has at least $k+1$ vertices and for every subset $S \subseteq V(G)$ of at most $k-1$ vertices, the graph $G - S$ is connected. We say that a graph is *quasi-4-connected* if it is 3-connected and for every set $S \subseteq V(G)$ of size 3 such that $G - S$ is not connected, $G - S$ has exactly two connected components and one of them consists of a single vertex.

2.4 Rays and ends

A *ray* in a graph G is an infinite simple one-way path $P = (v_1, v_2, \ldots)$. A *subray* P' of P is a ray of the form $P' = (v_i, v_{i+1}, \ldots)$ for some $i \geqslant 1$. We say that a ray *lives* in a set $X \subseteq V(G)$ if one of its subrays is included in X. We define an equivalence relation $\sim$ over the set of rays $\mathcal{R}(G)$ by letting $P \sim P'$ if and only if for every finite set of vertices $S \subseteq V(G)$, there is a component of $G - S$ that contains infinitely many vertices from both P and P'. Note that this is equivalent to saying that for any finite set $S \subseteq V(G)$, P and P' are living in the same component of $G - S$. The *ends* of G are the elements of $\mathcal{R}(G)/\sim$, the equivalence classes of rays under $\sim$. For every $X \subseteq V(G)$, we say that an end ω *lives in* X if one of its rays lives in X.

We now introduce the notion of accessibility in graphs considered by Thomassen and Woess [TW93]. To distinguish it from the related notion in groups (see Chapter 2), we will call it vertex-accessibility in the remainder of the manuscript. When there is a finite set X of vertices of G, two distinct components C_1, C_2 of $G - X$, and two distinct ends ω_1, ω_2 of G such that for each $i = 1, 2$, ω_i lives in C_i, we say that X *separates* ω_1 and ω_2. A graph G is *vertex-accessible* if there is an integer k such that for any two distinct ends ω_1, ω_2 in G, there is a set of at most k vertices that separates ω_1 and ω_2. The *degree* of an end ω is the supremum number $k \in \mathbb{N} \cup \{\infty\}$ of pairwise-disjoint rays that belong to ω. By a result of Halin [Hal65], this supremum is a maximum i.e., if an end ω has infinite degree, then there exists an infinite countable family of pairwise-disjoint rays belonging to ω. An end is *thin* if it has finite degree, and *thick* otherwise. It is an easy exercise to check that for every end ω of finite degree k and every end $\omega' \neq \omega$, there is a set of size at most k that separates ω from ω'.

The interested reader is referred to Chapter 8 in [Die17] for more background and important results in infinite graph theory.

2.5 Closed walks

A *walk* in a graph G is a finite sequence of vertices $W = (v_1, \dots, v_k)$ where for each $i \in [k-1]$, $v_i v_{i+1} \in E(G)$. We call W a *closed walk* when $v_1 = v_k$ or $v_1 v_k \in E(G)$, and we call it a *cycle* if moreover its vertices are pairwise distinct. We let $\mathcal{W}(G)$ denote the set of closed walks of G. If W is a closed walk that contains a *spur*, i.e. if there is some i such that $v_{i-1} = v_{i+1}$, then we say that $W' = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_k)$ is obtained from W by *deleting the spur*. The inverse operation of *adding a spur* consists in adding a neighbor of v between v_i and v_{i+1} in the walk W, so that the obtained walk is $W' = (v_1, \dots, v_i, v, v_i, v_{i+1}, \dots, v_k)$. The *rotation* of W is the walk $(v_2, v_3 \dots, v_k, v_1)$, and the *reflection* of W is the walk $(v_k, v_{k-1} \dots, v_2, v_1)$.

If $W = (v_1, \dots, v_k)$ and $W' = (v'_1, \dots, v'_\ell)$ are two walks such that $v_k = v'_1$, then their *sum* is the walk $W \cdot W' := (v_1, \dots, v_k = v'_1, \dots, v'_\ell)$. We will say that a set of closed walks $\mathcal{W}$ *generates* another set of closed walks $\mathcal{W}'$ if every element of $\mathcal{W}'$ can be obtained from elements of $\mathcal{W}$ by adding and deleting spurs, and performing sums, reflections and rotations.

2.6 Group actions on graphs

Unless stated otherwise, in this manuscript we will always denote with $\cdot$ the binary operation of a group and denote groups with capital greek letters. For simplicity, we will denote a group $(\Gamma, \cdot)$ by Γ. We will denote the identity element of Γ with 1_Γ and for every $\gamma \in \Gamma$, we set $\gamma^0 := 1_\Gamma$, and $\gamma^i := \gamma \cdot \gamma^{i-1}$ for all $i \geqslant 1$. We also set $\gamma^i := (\gamma^{-1})^{-i}$ for all $i \leqslant 0$, where γ^{-1} denotes the inverse of γ in Γ.

An *automorphism* of a graph G is a graph isomorphism from G to itself (i.e., a bijection from $V(G)$ to $V(G)$ that maps edges to edges and non-edges to non-edges). The set of automorphisms of G has a natural group structure (as a subgroup of the symmetric group over $V(G)$); the group of automorphisms of G is denoted by $\mathrm{Aut}(G)$.

For a graph G and a group Γ, we will say that Γ acts by automorphisms on G (or simply that Γ acts on G when the context is clear) if every element of Γ induces an automorphism γ of G, such that the induced application $\Gamma \to \mathrm{Aut}(G)$ is a group morphism. We will usually use the left multiplicative notation $\gamma \cdot x$ instead of $\gamma(x)$ for $\gamma \in \Gamma$, $x \in V(G)$. Note that the opposite convention was adopted in [EGLD23], however it is more common in the literature to associate the left action of Γ on G with the action of a group Γ by translation on its Cayley graphs (see Section 12), thus we decided to stick to this convention here. For every $X \subseteq V(G), \Gamma' \subseteq \Gamma$ and $\gamma \in \Gamma$, we let $\gamma \cdot X := \gamma(X) = \{\gamma \cdot x : x \in X\}$ and $\Gamma' \cdot X := \bigcup_{\gamma \in \Gamma'} \gamma \cdot X$. We denote the set of orbits of $V(G)$ under the action of Γ by G/Γ (Γ naturally induces an equivalence relation on $V(G)$, relating elements in the same orbit of Γ). For every subset $X \subseteq V(G)$ we let $\mathrm{Stab}_\Gamma(X) := \{\gamma \in \Gamma : \gamma \cdot X = X\}$ denote the *stabilizer* of X, which is always a subgroup of Γ. For each $x \in X$, we let $\Gamma_x := \mathrm{Stab}_\Gamma(\{x\})$.

2.7 Quasi-transitive graphs

We call the action of Γ on G *vertex-transitive* (or simply *transitive*) when there is only one orbit in G/Γ, i.e., when for every two vertices $u, v \in V(G)$ there exists an element $\gamma \in \Gamma$ such that $\gamma \cdot u = v$. The action of Γ on G is called *quasi-transitive* if there is only a finite number of orbits in G/Γ. We say that G is *transitive* (respectively *quasi-transitive*) if it admits a transitive (respectively quasi-transitive) group action. Note that every quasi-transitive locally finite graph has bounded degree.

It was proved, first by Freudenthal [Fre44] for finitely generated groups and then in the more general graph-theoretic context [DJM93], that the number of ends of a quasi-transitive graph is either $0, 1, 2$ or ∞. A graph with a single end is said to be *one-ended*.

2.8 Quasi-isometries

Let (X, d_X) and (Y, d_Y) be two metric spaces. We say that X is *quasi-isometric* to Y if there is a map $f : X \to Y$ and constants $\varepsilon \geqslant 0$, $\lambda \geqslant 1$, and $C \geqslant 0$ such that (i) for any $y \in Y$ there is $x \in X$ such that $d_Y(y, f(x)) \leqslant C$, and (ii) for every $x_1, x_2 \in X$,

$$\frac{1}{\lambda} d_X(x_1, x_2) - \varepsilon \leqslant d_Y(f(x_1), f(x_2)) \leqslant \lambda d_X(x_1, x_2) + \varepsilon.$$

It is not difficult to check that the definition is symmetric, and we often simply say that X and Y are quasi-isometric. If condition (i) is omitted in the definition above, we say that f is a *quasi-isometric embedding of X in Y*. If $\varepsilon = 0$, then we call f a *bilipschitz* mapping. We will only consider in this manuscript quasi-isometries between graphs equipped with their shortest path metric. Since graphs are uniformly discrete metric spaces (in the sense that any two distinct elements lie at distance at least 1 apart), any *injective* quasi-isometric embedding of a graph G in a graph H is also a bilipschitz embedding of G in H. In particular, if G has bounded degree and $f : G \to H$ is a quasi-isometric embedding, then for every $y \in V(H)$, the set $f^{-1}(y) := \{x \in V(G) : f(x) = y\}$ must have uniformly bounded diameter, and thus as G has bounded degree, it also has uniformly bounded size. Hence, if G admits a quasi-isometric embedding in H, there exists a constant $C \geqslant 1$ and a bilipschitz embedding of G into H^{+C}, the graph obtained from H by adding C pendant vertices to each vertex of H.

Note that the number of ends of a locally finite graph is preserved under taking quasi-isometries. We will implicitly use this property extensively in the remainder of the thesis.

3 Separations and canonical tree-decompositions

In this section, we introduce the notions of separations and canonical tree-decompositions which will be central in the structure theorems we will present in the next subsections. Many concepts we will present admit analogous formulations with respect to (edge-)cuts. Even though the proofs that will be discussed are based on the structure of vertex separations in graphs, we mention that many known accessibility results were first proved working on the structure of edge-cuts, and based on the theory of cuts introduced by Dicks and Dunwoody in [DD89].

3.1 Separations

As we will use later in this chapter many results of Grohe [Gro16], we use his notations and definitions for all objects related to separations and tangles (which differ from the more conventional ones from [Die17]). A *separation* in a graph $G = (V, E)$ is a triple (Y, S, Z) such that Y, S, Z are pairwise-disjoint subsets of $V(G)$, $V = Y \cup S \cup Z$ and there is no edge between vertices of Y and Z. A separation (Y, S, Z) is *proper* if Y and Z are nonempty. In this case, S is a *separator* of G.

The separation (Y, S, Z) is said to be *tight* if there are some components C_Y, C_Z respectively of $G[Y], G[Z]$ such that $N_G(C_Y) = N_G(C_Z) = S$. The *order* of a separation (Y, S, Z) is $|S|$ and the *order* of a family $\mathcal{N}$ of separations is the supremum of the orders of its separations. In what follows, we will always consider sets of separations of finite order. We will denote with $\mathrm{Sep}_k(G)$ (respectively $\mathrm{Sep}_{<k}(G)$) the set of all separations of G of order k (respectively less than k).

Given a separation (Y, S, Z) and an automorphism γ of a graph G, let $\gamma \cdot (Y, S, Z) := (\gamma \cdot Y, \gamma \cdot S, \gamma \cdot Z)$. If a group Γ acts on a graph G, then for each $k \geqslant 0$, Γ also induces an action on $\mathrm{Sep}_k(G)$.

The following lemma was originally stated in [TW93] for transitive graphs, but the same proof immediately implies that the result also holds for quasi-transitive graphs.

Lemma 3.1 (Corollary 4.3 in [TW93])**.** *Let G be a locally finite graph. Then for every $v \in V(G)$ and $k \geqslant 1$, there is only a finite number of tight separations (Y, S, Z) of order k in G such that $v \in S$. Moreover, for any group Γ acting quasi-transitively on G and any $k \geqslant 1$, there is only a finite number of Γ-orbits of tight separations of order at most k in G.*

3.2 Canonical tree-decompositions

A *tree-decomposition* of a graph G is a pair $(T, \mathcal{V})$ where T is a tree and $\mathcal{V} = (V_t)_{t \in V(T)}$ is a family of subsets V_t of $V(G)$ such that:

- $V(G) = \bigcup_{t \in V(T)} V_t$;
- for every nodes t, t', t'' such that t' is on the unique path of T from t to t'', $V_t \cap V_{t''} \subseteq V_{t'}$;
- every edge $e \in E(G)$ is contained in an induced subgraph $G[V_t]$ for some $t \in V(T)$.

Note that in our definition of tree-decomposition, we allow T to have vertices of infinite degree. The sets V_t for every $t \in V(T)$ are called the *bags* of $(T, \mathcal{V})$, and the induced subgraphs $G[V_t]$ the *parts* of $(T, \mathcal{V})$. The *width* of $(T, \mathcal{V})$ is the supremum of $|V_t| - 1$, for $t \in V(T)$. Note that the width of a tree-decomposition can be infinite. The sets $V_t \cap V_{t'}$ for every $tt' \in E(T)$ are called the *adhesion sets* of $(T, \mathcal{V})$ and the *adhesion* of $(T, \mathcal{V})$ is the supremum of the sizes of its adhesion sets (possibly infinite). We also let $V_\infty(T) \subseteq V(T)$ denote the set of nodes $t \in V(T)$ such that V_t is infinite.

For any subgroup Γ of $\mathrm{Aut}(G)$, we say that a tree-decomposition $(T, \mathcal{V})$ is *canonical with respect to* Γ, or simply Γ-*canonical*, if Γ induces a group action on T such that for every $\gamma \in \Gamma$ and $t \in V(T)$, $\gamma \cdot V_t = V_{\gamma \cdot t}$. By definition of a group action on a graph, $t \mapsto \gamma \cdot t$ is an automorphism of T for any $\gamma \in \Gamma$. In particular, for every $\gamma \in \Gamma$, note that γ sends bags of $(T, \mathcal{V})$ to bags, and adhesion sets to adhesion sets. When $(T, \mathcal{V})$ is $\mathrm{Aut}(G)$-canonical, we simply say that it is *canonical*.

Remark 3.2. If $(T, \mathcal{V})$ is a Γ-canonical tree-decomposition of a graph G, then Γ acts both on G and T, so there are two different notions of a stabilizer of a node $t \in V(T)$: $\Gamma_t = \mathrm{Stab}_\Gamma(t)$ (where we consider the action of Γ on T), and $\mathrm{Stab}_\Gamma(V_t)$ (where we consider the action of Γ on G). Observe that for any $t \in V(T)$ we have $\Gamma_t \subseteq \mathrm{Stab}_\Gamma(V_t)$. The reverse inclusion does not hold in general (when there are adjacent nodes $s, t \in V(T)$ with $V_s = V_t$, automorphisms of T might exchange s and t and thus stabilize $V_s = V_t$ without stabilizing s or t). However, if $t \in V(T)$ is such that $\{t' \in V(T) : V_{t'} = V_t\} = \{t\}$, then $\Gamma_t = \mathrm{Stab}_\Gamma(V_t)$. In particular, if $(T, \mathcal{V})$ has finite adhesion, then every bag V_t with $t \in V_\infty(T)$ appears only once in the decomposition, and thus for each such node $t \in V_\infty(T)$ we have $\Gamma_t = \mathrm{Stab}_\Gamma(V_t)$.

3.3 Edge-separations and torsos

Consider a tree-decomposition $(T, \mathcal{V})$ of a graph G, with $\mathcal{V} = (V_t)_{t \in V(T)}$. Let A be the set of all the orientations of the edges of $E(T)$, i.e. A contains the pairs (t_1, t_2), (t_2, t_1) for every edge $t_1 t_2$ of T. For an arbitrary pair $(t_1, t_2) \in A$, and for each $i \in \{1, 2\}$, let T_i denote the component of $T - \{t_1 t_2\}$ containing t_i. Then the *edge-separation* of G associated to (t_1, t_2) is (Y_1, S, Y_2) with $S := V_{t_1} \cap V_{t_2}$ and $Y_i := \bigcup_{s \in V(T_i)} V_s \setminus S$ for $i \in \{1, 2\}$.

The *torsos* of $(T, \mathcal{V})$ are the graphs $G[\![V_t]\!]$ for $t \in V(T)$, with vertex set V_t and edge set $E(G[V_t])$ together with the edges xy such that x and y belong to a common adhesion set of $(T, \mathcal{V})$.

If Γ acts on G and $\mathcal{N}$ is a family of separations of G, we say that $\mathcal{N}$ is Γ-*invariant* if for every $(Y, S, Z) \in \mathcal{N}$ and $\gamma \in \Gamma$, we have $\gamma \cdot (Y, S, Z) \in \mathcal{N}$. Note that if $(T, \mathcal{V})$ is Γ-canonical, then the associated set of edge-separations is Γ-invariant.

Remark 3.3. If $(T, \mathcal{V})$ is a Γ-canonical tree-decomposition of a locally finite graph G whose edge-separations are tight, with finite bounded order, then by Lemma 3.1 the action of Γ on $E(T)$ must induce a finite number of orbits. In particular, Γ must also act quasi-transitively on $V(T)$.

The *treewidth* of a graph G is the infimum of the width of $(T, \mathcal{V})$, among all tree-decompositions $(T, \mathcal{V})$ of G. Note that adding to a tree-decomposition of bounded width the restriction that it must be canonical can be very costly in the finite case: while it is well known that every cycle graph C_n on n vertices has treewidth 2, the example below shows that in any canonical tree-decomposition of C_n, some bag contains all the nodes of C_n.

Example 3.4. Let C_n be the cycle graph on n elements. Note that the additive group $\mathbb{Z}_n$ acts transitively by rotation on C_n. We let α be a generator of $\mathbb{Z}_n$ of order n. Let $(T, (V_t)_{t \in V(T)})$ be a $\mathbb{Z}_n$-canonical tree-decomposition of C_n. Without loss of generality we may assume that T is finite, by contracting every edge $tt' \in E(T)$ such that $V_t = V_{t'}$. We may also assume that no edge tt' of T is *inverted* by α, i.e. such that $\alpha \cdot (t,t') = (t',t)$, as if it was the case we could subdivide the edge tt' (i.e. add a new vertex t^* between t and t') and let $V_{t^*} := V_t \cap V_{t'}$. If we let T' be the tree of the tree-decomposition obtained after performing such a subdivision, note that the obtained tree-decomposition is still $\mathbb{Z}_n$-canonical as α induces an automorphism of T' that stabilizes the vertex t^* and acts on $V(T)$ the same way that it did before the subdivision. After this operation none of the edges tt^*, t^*t' is inverted by α. It is an easy exercise to prove that if no edge of T is inverted by α, there exists a vertex $t \in V(T)$ stabilized by α, and hence by all the elements of $\mathbb{Z}_n$. Then as $\mathbb{Z}_n$ acts transitively on G, we must have $V_t = V(C_n)$ for such a $t \in V(T)$.

The following two lemmas will be useful when working with canonical tree-decompositions, as they allow to use alternatively torsos or parts of a tree-decompositions according to our purposes.

Lemma 3.5. *Let G be a locally finite Γ-quasi-transitive graph and $(T, (V_t)_{t\in V(T)})$ be a Γ-canonical tree-decomposition of G with finite adhesion whose parts are connected subgraphs of G and such that $E(T)$ admits only finitely many Γ orbits. Then for every $t \in V(T)$, $G[\![V_t]\!]$ is quasi-isometric to $G[V_t]$.*

Proof. We will show that the identity on V_t induces a quasi-isometry between $G[V_t]$ and $G[\![V_t]\!]$. For each $t \in V(T)$, $G[V_t]$ is a subgraph of $G[\![V_t]\!]$, so for every $u, v \in V_t$ we have $d_{G[\![V_t]\!]}(u,v) \leqslant d_{G[V_t]}(u,v)$. We will now show that there exists a constant $C \geqslant 0$ such that for each $t \in V(T), u, v \in V_t$:

$$d_{G[V_t]}(u,v) \leqslant C \cdot d_{G[\![V_t]\!]}(u,v).$$

This will immediately imply that id_{V_t} induces the desired quasi-isometry.

As $E(T)$ has finitely many orbits under the action of the automorphism group of G, up to applying an automorphism from Γ, there are only finitely many pairs $\{u,v\}$ such that u, v lie in a common adhesion set of $(T, \mathcal{V})$. In particular, as G is connected this means that for each $t \in V(T)$, the set of values $D_t := \{d_{G[V_t]}(u,v), \exists s \in N_T(t), u, v \in V_s \cap V_t\}$ admits a maximum. As $V(T)$ has finitely many Γ-orbits, the set of values $\bigcup_{t\in V(T)} D_t$ also admits a maximum $C \in \mathbb{N}$. In particular such a constant C satisfies the above inequality. □

Lemma 3.6. *Let G be a connected locally finite Γ-quasi-transitive graph and $(T, (V_t)_{t\in V(T)})$ be a Γ-canonical tree-decomposition of G with finite adhesion, connected torsos and such that $E(T)$ admits only finitely many Γ orbits. Then there exists a Γ-canonical tree-decomposition $(T, (V'_t)_{t\in V(T)})$ of G with finite adhesion, with the same decomposition tree T and Γ-action of Γ on T, and such that for each $t \in V(T)$, $G[V'_t]$ is connected and quasi-isometric to $G[\![V_t]\!]$.*

Proof. As $E(T)$ has finitely many Γ-orbits, note that the set $\{d_G(u,v) : \exists t \in V(T), uv \in E(G[\![V_t]\!]) \backslash E(G)\}$ is bounded and thus admits a maximum, say $k \in \mathbb{N}$. We let $\mathcal{V}' := (V'_t)_{t\in V(T)}$ be defined by $V'_t := B_k(V_t) = \{v \in V(G) : \exists u \in V_t, d_G(u,v) \leqslant k\}$. We show that $(T, \mathcal{V}')$ is a tree-decomposition of G with the desired properties.

We first check that $(T, \mathcal{V}')$ is a tree-decomposition. As $V_t \subseteq V'_t$ for each $t \in V(T)$, we have $\bigcup_{t\in V(T)} V'_t = V(G)$ and every edge of G must be contained in some part $G[V'_t]$ for some $t \in V(T)$. It remains to show that for all distinct nodes $t, t', t'' \in V(T)$ such that t' lies on the shortest path from t to t'', $V'_t \cap V'_{t''} \subseteq V'_{t'}$. Let $u \in V'_t \cap V'_{t''}$ and $(x, y) \in V_t \times V_{t''}$ be such that $d_G(x, u) \leqslant k$ and $d_G(y, u) \leqslant k$. We consider the path p going from x to y in G obtained by concatenating a shortest path from x to u with a shortest path from u to y. Then as $(T, \mathcal{V})$ is a tree-decomposition of G, p must intersect $V_{t'}$. In particular, u is at distance at most k from any point of p, so we just proved that $u \in V'_{t'}$, implying that $(T, \mathcal{V}')$ is a tree-decomposition. In fact, this argument also implies that $(T, \mathcal{V}')$ has finite adhesion: if $u \in V'_t \cap V'_{t''}$ and t' is chosen to be the neighbor of t on the unique shortest path from t to t'' in T, then any path p obtained as above must intersect $V_t \cap V_{t'}$. In particular it implies that $u \in B_k(V_t \cap V_{t'})$, which must be finite of bounded size as G has bounded degree and $V_t \cap V_{t'}$ is finite.

Note that the action of Γ on T when T is seen as the decomposition tree of $(T, \mathcal{V})$ immediately extends to the same action on T when T is considered as the decomposition tree of $(T, \mathcal{V}')$, hence $(T, \mathcal{V}')$ is still Γ-canonical and $E(T)$ still has finitely many Γ-orbits with respect to this action. Moreover for every $t \in V(T)$ and every $uv \in E(G^+[V_t]) \setminus E(G)$, our choice of k ensures that there exists a path from u to v in $G[V'_t]$, so the parts of $(T, \mathcal{V}')$ are connected.

Let $t \in V(T)$, and fix any projection $\pi : V'_t \to V_t$ such that $\pi|_{V_t} = \mathrm{id}_{V_t}$ and such that for each $v \in V'_t$, $d_G(\pi(v), v) = d_G(V_t, v) = \min\{d_G(v, u) : u \in V_t\}$. We show that π defines a quasi-isometry between $G[V'_t]$ and $G[\![V_t]\!]$. First, note that for each pair of vertices $x, y \in V_t$ such that there exists a path p from x to y in G which only intersects V_t in $\{x, y\}$, the vertices x and y must belong to a common adhesion set of $(T, \mathcal{V})$. In particular, for all vertices $u, v \in V'_t$, if p is a path of length d in $G[V'_t]$ from u to v, then there is a path p' of length at most $d + 2k$ between $\pi(u)$ and $\pi(v)$ in $G[V'_t]$, and we obtain a path q in $G[\![V_t]\!]$ of length at most $d + 2k$ between $\pi(u)$ and $\pi(v)$ after replacing every subpath of p' having only its two endpoints in V_t by an edge of $G[\![V_t]\!]$. In particular if we choose p to be a shortest path from u to v in $G[V'_t]$, we get

$$d_{G[\![V_t]\!]}(\pi(u), \pi(v)) \leqslant d_{G[V'_t]}(u, v) + 2k.$$

We now consider a path p from $\pi(u)$ to $\pi(v)$ in $G[\![V_t]\!]$ of length d. Then by the choice of k and π, there exists a path from u to v in $G[V'_t]$ of length at most $dk + 2k$, showing that

$$d_{G[V'_t]}(u, v) \leqslant k d_{G[\![V_t]\!]}(\pi(u), \pi(v)) + 2k.$$

As π is surjective onto V_t, it follows that it is a quasi-isometry between $G[V'_t]$ and $G[\![V_t]\!]$. □

3.4 Nested sets of separations

We define an order $\leqslant_{\mathrm{RS}}$ on the set of separations of a graph G as follows. For any two separations $(Y, S, Z), (Y', S', Z')$, we write $(Y, S, Z) \leqslant_{\mathrm{RS}} (Y', S', Z')$ if and only if $Y' \subseteq Y$ and $Z \subseteq Z'$. This order corresponds to the (inverse of the) one introduced in [RS91], and we will distinguish it with another order $\leqslant_{\mathrm{G}}$ introduced in [Gro16] that we will use later which is similar but does not admit exactly the same definition. Intuitively, $(Y, S, Z) \leqslant_{\mathrm{RS}} (Y', S', Z')$

means that (Y, S, Z) points towards a "Z-direction" in a more accurate way than (Y', S', Z') does.

Two separations $(Y, S, Z), (Y', S', Z')$ of a graph G are said to be *nested* if (Y, S, Z) is comparable either with (Y', S', Z') or with (Z', S', Y') with respect to the order $\leqslant_{\mathrm{RS}}$. A set of separations $\mathcal{N}$ of G is *nested* if all its separations are pairwise nested. We say that $\mathcal{N}$ is *symmetric* if for every $(Y, S, Z) \in \mathcal{N}$, we also have $(Z, S, Y) \in \mathcal{N}$. It is not hard to observe that if $(T, \mathcal{V})$ is a tree-decomposition and $\mathcal{N}$ denotes its set of edge-separations, then $\mathcal{N}$ is symmetric and nested. Moreover, if $(T, \mathcal{V})$ is Γ-canonical, then $\mathcal{N}$ is also Γ-invariant with respect to the action of Γ on the set of separations of G.

Extending a known result [CDHS11, Theorem 4.8] for finite graphs, Carmesin, Hamann and Miraftab proved in [CHM22, Theorem 3.2] that symmetry and nestedness together with a third property are sufficient conditions to obtain a tree-decomposition from a nested set of separations.

According to the notation from [CHM22], we say that a set of separations $\mathcal{N}$ has *finite intervals* if for every pair of separations $(Y_1, S_1, Z_1), (Y_2, S_2, Z_2) \in \mathcal{N}$ with $(Y_1, S_1, Z_1) \leqslant_{\mathrm{RS}} (Y_2, S_2, Z_2)$ there are only finitely many separations $(Y, S, Z) \in \mathcal{N}$ such that:

$$(Y_1, S_1, Z_1) \leqslant_{\mathrm{RS}} (Y, S, Z) \leqslant_{\mathrm{RS}} (Y_2, S_2, Z_2).$$

Theorem 3.7 (Theorem 3.2 in [CHM22])**.** *Let $\mathcal{N}$ be a nested set of separations with finite intervals in an arbitrary graph G. Then there exists a tree-decomposition $(T, \mathcal{V})$ of G such that the edge-separations of $(T, \mathcal{V})$ are exactly the separations from $\mathcal{N}$ and the correspondence is one-to-one. Moreover, if $\mathcal{N}$ is Γ-invariant with respect to some group Γ acting on G, then $(T, \mathcal{V})$ is Γ-canonical.*

Lemma 3.8. *If G is connected, locally finite and $\mathcal{N}$ is a nested set of separations in G such that for every $(Y, S, Z) \in \mathcal{N}$, S has uniformly bounded diameter with respect to the metric d_G, then $\mathcal{N}$ has finite intervals.*

Proof. We let $d(X, Y) := \min\{d(x, y) : x \in X, y \in Y\}$ for every two nonempty sets $X, Y \subseteq V(G)$. Let k be an upper bound on the diameter of the separators S in G for $(Y, S, Z) \in \mathcal{N}$. For every triple $(Y_1, S_1, Z_1), (Y_2, S_2, Z_2), (Y, S, Z) \in \mathcal{N}$ with $(Y_1, S_1, Z_1) \leqslant_{\mathrm{RS}} (Y, S, Z) \leqslant_{\mathrm{RS}} (Y_2, S_2, Z_2)$, note that $d(S, S_1) \leqslant d(S_1, S_2) := k'$, and as G is locally finite, the size of $|S|$ must be finite. Moreover, for every $u \in S_1, v \in S$ we must have $d(u, v) \leqslant k' + k$, hence every separator S of such a separation (Y, S, Z) of $\mathcal{N}$ lives in the ball of radius $k' + k$ around S_1, and as G is locally finite, there are only finitely many such separators. As G is locally finite, for every finite set S, $G - S$ has only finitely many connected components so S is the separator of only finitely many separations of G. It follows that we can only find finitely many such $(Y, S, Z) \in \mathcal{N}$. □

3.5 Separations of order at most 3

If G is not connected, then the tree-decomposition $(T, \mathcal{V})$ where T is a star whose central bag is empty and where we put a bag for each connected component of G can easily be seen to be a canonical tree-decomposition with adhesion 0, as every automorphism of G acts on

T by permuting some branches. If we start from a connected graph G, it is well known that the block cut-tree of G is a canonical tree-decomposition $(T, (V_t)_{t \in V(T)})$ of G whose adhesion sets have size 1 and such that for each $t \in V(T)$, $G[V_t] = G[\![V_t]\!]$ has either size at most 2 or is 2-connected. A similar result holds for separations of order 2. This was proved by Tutte [Tut84] in the finite case, and generalized to infinite graphs in [DSS98]. See also [CK23, Theorem 1.6.1] for a more precise version.

Theorem 3.9 ([DSS98])**.** *Every locally finite graph G has a canonical tree-decomposition of adhesion at most 2 with tight edge-separations, whose torsos are minors of G and are complete graphs of order at most 2, cycles, or 3-connected graphs.*

For separations of order 3, a similar result was obtained by Grohe for finite graphs [Gro16].

Theorem 3.10 ([Gro16])**.** *Every finite graph G has a tree-decomposition of adhesion at most 3 whose torsos are minors of G and are complete graphs on at most 4 vertices or quasi-4-connected graphs.*

In Section 7 we extend Theorem 3.10 to locally finite graphs, while making sure that most of the construction (except the very end) is canonical. More precisely, we will reproduce in Section 6 the main steps of the work of [Gro16] and give the additional arguments to extend them to locally finite graphs. A consequence is that Theorem 3.10 extends to locally finite graphs. A downside is that we cannot require Grohe's decomposition to be canonical, as illustrated by Example 3.11 below, which was introduced in [Gro16] in the finite case.

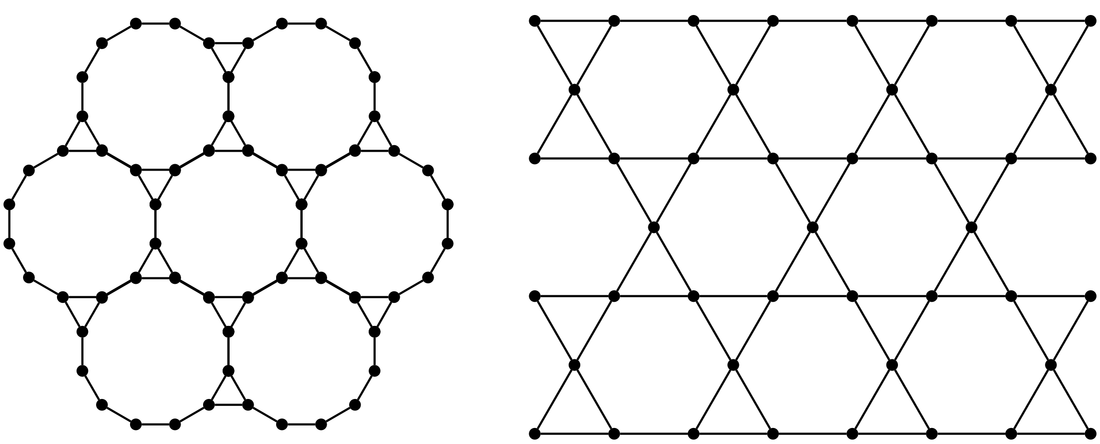

Figure 1.1: Left: a finite section of the 3-connected infinite graph obtained by replacing in the infinite hexagonal planar grid each vertex by a triangle with three vertices of degree 3. Right: a finite section of the quasi-4-connected torso $G[\![V_{z_0}]\!]$ of $(T, \mathcal{V})$. Note that it does not depend of the choice of V_{z_0}.

Example 3.11. Consider the 3-connected infinite planar graph H obtained from the infinite hexagonal planar grid by replacing each vertex by a triangle with three vertices of degree 3 (see Figure 1.1 (left) for a finite part of this graph). We let M be the set of edges connecting pairs of triangles, or equivalently the set of edges that do not belong to any triangle (note that M is a perfect matching). The tree-decomposition $(T, \mathcal{V})$ of H obtained by extending the ideas in [Gro16] to the infinite case has an infinite bag V_{z_0} obtained by selecting one endpoint of each edge of M (which is equivalent to fixing an orientation of each of these edges). The tree T is a star with center z_0, and its other bags are finite. While there are many different choices for V_{z_0}, none of them gives a canonical tree-decomposition. Moreover, it is even possible to choose V_{z_0} such that $(T, \mathcal{V})$ is not Γ-canonical for any subgroup Γ of $\mathrm{Aut}(G)$ with a quasi-transitive action on G. For example, one can fix some special vertex $v_0 \in V(G)$, and let A be the orientation of the edges uv connecting two triangles of G defined by $(u, v) \in A$ if $d_G(u, v_0) < d_G(v, v_0)$. If $d_G(u, v_0) = d_G(v, v_0)$, then we choose an arbitrary orientation of uv to be in A. We let V_{z_0} be the set of tails of the arcs in A. Assume that Γ is a subgroup of $\mathrm{Aut}(G)$ such that the decomposition $(T, \mathcal{V})$ obtained for this choice V_{z_0} is Γ-canonical. Then note that as V_{z_0} is the only infinite bag of $(T, \mathcal{V})$, the only automorphisms $g \in \mathrm{Aut}(G)$ that can induce an automorphism on T must fix V_{z_0}, so they must also preserve the orientation induced by A. As the triangle containing v_0 is the only triangle of G with three incoming edges, Γ must be a subgroup of its stabilizer, hence it cannot act quasi-transitively on G. Indeed one can check more generally that no tree-decomposition of H satisfying the properties of Theorem 3.10 can be canonical. To see this, assume for the sake of contradiction that such a decomposition $(T, \mathcal{V})$ exists. Then one of its edge-separations should be proper of order 3. Note that the only such separations separate a subgraph of a triangle from the rest of the graph. Let (Y, S, Z) be such a separation, such that Z is finite. Then there exists an edge e from M with one endpoint z in Z and the other s in S. Note that there exists an automorphism $\gamma \in \mathrm{Aut}(H)$ exchanging the two endpoints of e. In particular, as $(T, \mathcal{V})$ is canonical, both $\gamma \cdot (Y, S, Z)$ and (Y, S, Z) must be edge-separations of $(T, \mathcal{V})$, which can be seen to be impossible.

This example illustrates the fact that in general it is impossible to obtain a *canonical* tree-decomposition having exactly the properties described in Theorem 3.10.

In a recent work, Carmesin and Kurkofka [CK23] proved that 3-connected graphs can be decomposed into simpler pieces (namely quasi-4-connected graphs, wheels and thickened $K_{3,m}$'s) in a canonical way, suggesting another way to extend the ideas from Theorem 3.9 to 3-connected graphs. The decomposition they obtain is not exactly a tree-decomposition, as they construct an $\mathrm{Aut}(G)$-invariant nested set of *mixed-separations* of order 3 (while Theorem 3.10 only deals with (vertex-)separations of order 3), that is a set of separations of order 3 of the barycentric subdivision of G.

3.6 Combining canonical tree-decompositions

Let $(T, \mathcal{V})$ and $(T, \mathcal{V}')$ be tree-decompositions of two graphs G, G', respectively. We say that $(T, \mathcal{V})$ and $(T', \mathcal{V}')$ are *isomorphic* if there exists an isomorphism φ from G to G', and an isomorphism ψ from T to T' such that for each $t \in V(T)$, we have: $V'_{\psi(t)} = \varphi(V_t)$.

Let G be a graph and let Γ be a group acting on G. Let $(T, \mathcal{V})$, with $\mathcal{V} = (V_t)_{t \in V(T)}$, be a Γ-canonical tree-decomposition of G, and recall that for any $t \in V(T)$, $\Gamma_t = \mathrm{Stab}_\Gamma(t)$ denotes the stabilizer of the node t with respect to the action of Γ on the tree T. For each $t \in V(T)$, let $(T_t, \mathcal{V}_t)$ be a Γ_t-canonical tree-decomposition of $G[\![V_t]\!]$. Our goal will be to refine $(T, \mathcal{V})$ by combining it with the tree-decompositions $(T_t, \mathcal{V}_t)_{t\in V(T)}$. If we want the resulting refined tree-decomposition of G to be canonical, we need to impose a condition on the tree-decompositions $(T_t, \mathcal{V}_t)_{t\in V(T)}$ (namely that they are consistent with the action of Γ on G). This is captured by the following definition. Let $\gamma \cdot (T_t, \mathcal{V}_t) := (T_t, \gamma \cdot \mathcal{V}_t)$ be the tree-decomposition of $G⟦\gamma \cdot V_t⟧ = G⟦V_{\gamma\cdot t}⟧$ with bags $\gamma \cdot \mathcal{V}_t := (\gamma \cdot V_s)_{s\in V(T_t)}$. Observe that $\gamma \cdot (T_t, \mathcal{V}_t)$ is $\Gamma_{\gamma\cdot t}$-canonical. We say that the construction $t \mapsto (T_t, \mathcal{V}_t)$ is Γ-*canonical* if for each $\gamma \in \Gamma$ and $t \in V(T)$, the tree-decompositions $\gamma \cdot (T_t, \mathcal{V}_t)$ and $(T_{\gamma\cdot t}, \mathcal{V}_{\gamma\cdot t})$ are isomorphic. We emphasize here that the first tree-decomposition is indexed by T_t, while the second is indexed by $T_{\gamma\cdot t}$.

The *trivial tree-decomposition* of a graph G consists of a tree T with a single node, whose bag is $V(G)$. Note that the trivial tree-decomposition is canonical.

Lemma 3.12. *Assume that G is locally finite, Γ acts on G and $(T, \mathcal{V})$ is a Γ-canonical tree-decomposition of G with finite adhesion. Let $\{t_i, i \in I_\infty\}$ be a set of representatives of the Γ-orbits of $V_\infty(T)$, indexed by some set I_∞. Assume that for every $i \in I_\infty$ there exists a Γ_{t_i}-canonical tree-decomposition $(T_{t_i}, \mathcal{V}_{t_i})$ of the torso $G⟦V_{t_i}⟧$ of finite adhesion. Then we can find some family $(T_t, \mathcal{V}_t)_{t\in V(T)}$ extending the family $(T_{t_i}, \mathcal{V}_{t_i})_{i\in I_\infty}$ such that each $(T_t, \mathcal{V}_t)$ is a Γ_t-canonical tree-decomposition of $G⟦V_t⟧$, the construction $t \mapsto (T_t, \mathcal{V}_t)$ is Γ-canonical, and for each $t \in V(T) \setminus V_\infty(T)$, $(T_t, \mathcal{V}_t)$ is the trivial tree-decomposition of $G⟦V_t⟧$.*

Proof. First we check that for each $\gamma \in \Gamma$ and every $t \in V_\infty(T)$, $\gamma \cdot \Gamma_t \cdot \gamma^{-1} = \Gamma_{\gamma\cdot t}$. For this we claim that we only need to prove the inclusion $\gamma \cdot \Gamma_t \cdot \gamma^{-1} \subseteq \Gamma_{\gamma\cdot t}$, as the converse then follows from replacing (γ, t) by $(\gamma^{-1}, \gamma \cdot t)$. Let $\gamma' \in \Gamma_t$. Then

$$\gamma \cdot \gamma' \cdot \gamma^{-1} \cdot (\gamma \cdot t) = \gamma \cdot \gamma' \cdot t = \gamma \cdot t,$$

thus $\gamma \cdot \gamma' \cdot \gamma^{-1} \in \Gamma_{\gamma\cdot t}$ and $\gamma \cdot \Gamma_t \cdot \gamma^{-1} \subseteq \Gamma_{\gamma\cdot t}$, as desired.

We complete I_∞ into a set I of representatives of the Γ-orbits of $V(T)$, and for each $i \in I \setminus I_\infty$ we let $(T_{t_i}, \mathcal{V}_{t_i})$ denote the trivial tree-decomposition of $G⟦V_{t_i}⟧$. For each $t \in V(T)$, we let $\gamma \in \Gamma$ and $i \in I$ be such that $t = \gamma \cdot t_i$ and let $(T_t, \mathcal{V}_t) := \gamma \cdot (T_{t_i}, \mathcal{V}_{t_i})$. We check that $(T_t, \mathcal{V}_t)$ is well-defined: for any two $\gamma, \gamma' \in \Gamma$ such that $\gamma \cdot t_i = \gamma' \cdot t_i = t$, we have $\gamma' \cdot \gamma^{-1} \in \Gamma_{t_i}$. As $(T_{t_i}, \mathcal{V}_{t_i})$ is Γ_{t_i}-canonical, $\gamma' \cdot \gamma^{-1} \cdot (T_{t_i}, \mathcal{V}_{t_i}) = (T_{t_i}, \mathcal{V}_{t_i})$ so we have $\gamma \cdot (T_{t_i}, \mathcal{V}_{t_i}) = \gamma' \cdot (T_{t_i}, \mathcal{V}_{t_i})$ and $(T_t, \mathcal{V}_t)$ is well-defined for each $t \in V(T)$. The fact that the construction $t \mapsto (T_t, \mathcal{V}_t)$ is Γ-canonical immediately follows from the definition. Finally, the tree-decomposition $(T_t, \mathcal{V}_t)$ is Γ_t-canonical because if it is not trivial, and $i \in I_\infty$ and $\gamma \in \Gamma$ are such that $\gamma \cdot t_i = t$, then $\Gamma_t = \gamma \cdot \Gamma_{t_i} \cdot \gamma^{-1}$ and thus Γ_t induces a group action on $(T_t, \mathcal{V}_t) = \gamma \cdot (T_{t_i}, \mathcal{V}_{t_i})$. □

Given two tree-decompositions $(T, \mathcal{V}), (T', \mathcal{V}')$ of a graph G, with $\mathcal{V} = (V_t)_{t\in V(T)}$ and $\mathcal{V}' = (V'_t)_{t\in V(T')}$, we say that $(T', \mathcal{V}')$ *refines* $(T, \mathcal{V})$ with respect to some family $(T_t, \mathcal{V}_t)_{t\in V(T)}$ of tree-decompositions if for every $t \in V(T)$, T_t is a subtree of T' such that $V_t = \bigcup_{s\in V(T_t)} V'_s$ and the trees $(T_t)_{t\in V(T)}$ are pairwise vertex-disjoint, cover $V(T')$ and for every edge $uv \in E(T)$, there exist $u' \in V(T_u), v' \in V(T_v)$ such that $u'v' \in E(T')$.

We say that $(T', \mathcal{V}')$ is a *subdivision* of $(T, \mathcal{V})$ if T' is obtained from T after considering a subset $E' \subseteq E(T)$ and doing the following for every edge $tt' \in E'$: we subdivide the edge tt' (by adding a new vertex t^* between t and t'), and we add a corresponding bag

$V_{t^*} := V_t \cap V_{t'}$ in the tree-decomposition. Note that if $(T', \mathcal{V}')$ is a subdivision of $(T, \mathcal{V})$, the two tree-decompositions have the same edge-separations.

The following result from [CHM22] will allow us to construct canonical tree-decompositions inductively:

Proposition 3.13 (Proposition 7.2 in [CHM22])**.** *Assume that G is locally finite, Γ acts on G and $(T, \mathcal{V})$ is a Γ-canonical tree-decomposition of G with finite adhesion, with $\mathcal{V} = (V_t)_{t \in V(T)}$. Assume that for every $t \in V(T)$, there exists a Γ_t-canonical tree-decomposition $(T_t, \mathcal{V}_t)$ of the torso $G[\![V_t]\!]$ of finite adhesion such that the edge-separations induced by $(T_t, \mathcal{V}_t)$ in $G[\![V_t]\!]$ are tight and pairwise distinct, and the construction $t \mapsto (T_t, \mathcal{V}_t)$ is Γ-canonical. Then there exists a Γ-canonical tree-decomposition $(T', \mathcal{V}')$ of G that refines $(T, \mathcal{V})$ with respect to a family $(T'_t, \mathcal{V}'_t)_{t \in V(T)}$ such that for each $t \in V(T)$, $(T'_t, \mathcal{V}'_t)$ is a Γ_t-canonical tree-decomposition of $G[\![V_t]\!]$ which is a subdivision of $(T_t, \mathcal{V}_t)$, and such that every adhesion set of $(T', \mathcal{V}')$ is either an adhesion set of $(T, \mathcal{V})$ or an adhesion set of some $(T_t, \mathcal{V}_t)$ for some $t \in V(T)$. Moreover, the construction $t \mapsto (T'_t, \mathcal{V}'_t)$ is Γ-canonical.*

Remark 3.14. In the original statement of [CHM22, Proposition 7.2], the fact that each tree-decomposition $(T'_t, \mathcal{V}'_t)$ is Γ_t-canonical and that the construction $t \mapsto (T'_t, \mathcal{V}'_t)$ is also Γ-canonical is not stated explicitly, however the authors show it explicitly in the proof.

Hence putting Lemma 3.12 together with Proposition 3.13, we immediately get:

Corollary 3.15. *Assume that G is locally finite, with a group Γ acting on G and $(T, \mathcal{V})$ a Γ-canonical tree-decomposition of G of finite adhesion, where $\mathcal{V} = (V_t)_{t \in V(T)}$. Let $\{t_i : i \in I_\infty\}$ denote a set of representatives of the orbits $V_\infty(T)/\Gamma$ such that for each $i \in I_\infty$, there exists a Γ_{t_i}-canonical tree-decomposition $(T_{t_i}, \mathcal{V}_{t_i})$ of $G[\![V_{t_i}]\!]$ with finite adhesion, such that the edge-separations induced by each $(T_{t_i}, \mathcal{V}_{t_i})$ in $G[\![V_{t_i}]\!]$ are tight and pairwise distinct. Then there exists a Γ-canonical tree-decomposition of G that refines $(T, \mathcal{V})$ with respect to some family $(T'_t, \mathcal{V}'_t)_{t \in V(T)}$ of Γ_t-canonical tree-decompositions of $G[\![V_t]\!]$ such that for each $i \in I_\infty$, $(T'_{t_i}, \mathcal{V}'_{t_i})$ is a subdivision of $(T_{t_i}, \mathcal{V}_{t_i})$, and for every $t \in V(T) \setminus V_\infty(T)$, $(T'_t, \mathcal{V}'_t)$ is the trivial tree-decomposition of $G[\![V_t]\!]$. Moreover, the construction $t \mapsto (T'_t, \mathcal{V}'_t)$ is Γ-canonical.*

Remark 3.16. The proof of Proposition 3.13 given in [CHM22] still holds if we assume more generally that $E(T_{t_i})$ has finitely many Γ_{t_i}-orbits instead of assuming that separations of $(T_{t_i}, \mathcal{V}_{t_i})$ are tight. In particular the same assumption can be done for Corollary 3.15.

A crucial property of a canonical tree-decomposition of a locally finite quasi-transitive graph is that the torsos or parts of the tree-decomposition are themselves quasi-transitive. This is proved in [HLMR22, Proposition 4.5] in the special case where Γ acts transitively on $E(T)$. We give here a more general proof, which is self-contained.

Lemma 3.17. *Let $k \in \mathbb{N}$, let G be a locally finite graph, and let Γ be a group acting quasi-transitively on G. Let $(T, \mathcal{V})$, with $\mathcal{V} = (V_t)_{t \in V(T)}$, be a Γ-canonical tree-decomposition of G of finite adhesion such that $E(T)$ has finitely many Γ-orbits. Then, for any $t \in V(T)$, the group $\Gamma_t := \mathrm{Stab}_\Gamma(V_t)$ induces a quasi-transitive action on $G[V_t]$, and thus also on $G[\![V_t]\!]$.*

Proof. We fix an orientation A of $E(T)$. Note that the number of Γ-orbits of A must be finite. Let $e_1, \dots, e_m \in A$ be representatives of each of the Γ-orbits of A, and let $\overline{e_1}, \dots, \overline{e_m}$ denote their inverse pairs. Fix any node $t \in V(T)$.

We consider an arbitrary vertex $z \in V(G)$, and let $t_0 \in V(T)$ be such that $z \in V_{t_0}$. We define the following subset of $V_t \cap (\Gamma \cdot z)$.

$$\Theta_z := \{y \in V_t \cap (\Gamma \cdot z) : y = \gamma \cdot z \text{ for some } \gamma \text{ such that } \gamma \cdot t_0 = t\}.$$

We first show that Γ_t acts transitively on Θ_z. Let $y, y' \in \Theta_z$ and $\gamma, \gamma' \in \Gamma$ be such that $y = \gamma \cdot z$, $y' = \gamma' \cdot z$ and $\gamma \cdot t_0 = \gamma' \cdot t_0 = t$. Then if we set $\alpha := \gamma' \cdot \gamma^{-1}$, we have

$$\alpha \cdot t = \gamma' \cdot \gamma^{-1} \cdot t = \gamma' \cdot t_0 = t,$$

and thus $\alpha \in \Gamma_t$. As $\alpha \cdot y = y'$, this shows that Γ_t acts transitively on Θ_z.

For $i \in [m]$, we define:

$$\Psi_i := \{y \in V_t \cap V_{t'} : \text{there exists } \gamma \in \Gamma \text{ such that } (t, t') = \gamma \cdot e_i\}$$

$$\Psi_{i+m} := \{y \in V_t \cap V_{t'} : \text{there exists } \gamma \in \Gamma \text{ such that } (t, t') = \gamma \cdot \overline{e_i}\}.$$

We observe that if a vertex of $V_t \cap (\Gamma \cdot z)$ does not lie in Θ_z, it has to lie in one of sets Ψ_i for $i \in [2m]$. To see this, let $y \in V_t \cap (\Gamma \cdot z)$, and $\gamma \in \Gamma$ be such that $y = \gamma \cdot z$. If $y \notin \Theta_z$, then $\gamma \cdot t_0 \neq t$. In this case the unique path in T from t to $\gamma \cdot t_0$ contains at least one edge. Let t' be the neighbor of t on this path. As $V_{\gamma \cdot t_0} = \gamma \cdot V_{t_0}$ and $z \in V_{t_0}$, we have $y \in V_t \cap V_{\gamma \cdot t_0}$. Hence as $(T, \mathcal{V})$ is a tree-decomposition, $y \in V_t \cap V_{t'}$. Thus if we let i be such that $(t, t') = \beta \cdot e_i$ or $(t, t') = \beta \cdot \overline{e_i}$ for some $\beta \in \Gamma$, we obtain that $y \in \Psi_i \cup \Psi_{i+m}$. This shows that V_t is covered by the union of the sets Θ_z, $z \in V(G)$ (there at most $|V(G)/\Gamma|$ such sets), and the sets Ψ_i, $i \in [2m]$.

We now show that Γ_t acts quasi-transitively on each Ψ_i, $i \in [2m]$. Let $i \in [m]$, $y_1, y_2 \in \Psi_i, t_1, t_2 \in V(T)$ and $\beta_1, \beta_2 \in \Gamma$ such that $(t, t_1), (t, t_2) \in E(T)$, $(t, t_1) = \beta_1 \cdot e_i$ and $(t, t_2) = \beta_2 \cdot e_i$. We set $\alpha := \beta_2 \cdot \beta_1^{-1}$ and note that α sends the directed edge $\beta_1 \cdot e_i$ to $\beta_2 \cdot e_i$. Let S_i be the separator of G associated to the edge-separation induced by the edge e_i in $(T, \mathcal{V})$. The previous remark implies that α sends $\beta_1 \cdot S_i$ to $\beta_2 \cdot S_i$ and that $\alpha \in \Gamma_t$. As for every $i \in [m]$, S_i has size at most k, we just proved that the action of Γ_t on Ψ_i induces at most k orbits. The case $i \in \{m+1, \dots, 2m\}$ is exactly the same.

As V_t is covered by the union of the sets Θ_z, $z \in V(G)$ (there are at most $|V(G)/\Gamma|$ such sets, and Γ_t acts transitively on each of these sets), and the sets Ψ_i, $i \in [2m]$ (and the action of Γ_t on each of these sets induces at most k orbits), we have $|V_t/\Gamma_t| \leqslant 2km + |V(G)/\Gamma|$, which implies that Γ_t acts quasi-transitively on $G[V_t]$. As $(T, \mathcal{V})$ is Γ-canonical, for each $\gamma \in \Gamma$ and each edge e lying inside some adhesion set of the tree-decomposition, γ sends e to a pair of vertices in another adhesion set of the tree-decomposition (and this pair of vertices must thus be joined by an edge in the corresponding torso). It follows that any automorphism $\gamma \in \Gamma_t$ of $G[V_t]$ is also an automorphism of the torso $G[\![V_t]\!]$. Hence, Γ_t also acts quasi-transitively on $G[\![V_t]\!]$. □

4 Quasi-transitive graphs of bounded treewidth

We survey in this subsection known characterisations of locally finite quasi-transitive graphs with bounded treewidth.

4.1 Quasi-transitive graphs of bounded pathwidth

A *path-decomposition* of a graph G is a tree-decomposition $(P, \mathcal{V})$ of G where P is a path. The *pathwidth* is the minimum width of a path-decomposition of G. As a warm-up we briefly show the following characterisations of quasi-transitive locally finite graphs with bounded pathwidth. Many results in the same spirit as Lemma 4.1 below can be found in the literature.

Lemma 4.1. *Let G be a connected locally finite quasi-transitive graph. The following are equivalent.*

- (i) *G has bounded pathwidth;*
- (ii) *there exists a subgroup Γ of $\mathrm{Aut}(G)$ and a Γ-canonical path-decomposition $(P, \mathcal{V})$ of bounded width of G such that Γ acts transitively on P;*
- (iii) *G is finite or 2-ended;*
- (iv) *G does not contain the infinite regular tree of degree 3 as a minor.*

Proof. First, observe that $(ii) \Rightarrow (i)$ is trivial.

We first show $(i) \Rightarrow (iii)$. As a quasi-transitive graph has either $0, 1, 2$ or infinitely many ends we just need to show that quasi-transitive graphs with 1 or infinitely many ends have unbounded pathwidth. By [Tho92, Proposition 5.6], if a quasi-transitive graph has only one end, then this end must be thick. By [Hal65], any graph with a thick end must contain a subgraph isomorphic to some subdivision of the hexagonal *half-grid*, i.e., of the intersection of the hexagonal grid, embedded isometrically in $\mathbb{R}^2$, with the upper half-space of $\mathbb{R}^2$. In particular, every one-ended graph has unbounded treewidth (and thus pathwidth). It is also not hard to see that every quasi-transitive graph with infinitely many ends contains a subdivision of the infinite regular tree of degree 3. In particular it must contain any finite forest as a minor and thus cannot have bounded pathwidth [RS83].

Note that in fact we just proved in the previous paragraph that every graph having one or infinitely many ends has the infinite regular tree of degree 3 as a minor, implying also that $(iv) \Rightarrow (iii)$ holds. Moreover, as the infinite regular tree of degree 3 contains all finite trees as minors, it immediately gives the implication $(i) \Rightarrow (iv)$.

The proof of $(iii) \Rightarrow (ii)$ follows from Lemma 4.2 below. Assume that G has 2 ends as if it is finite then the trivial path-decomposition has bounded width and is canonical. Let $\gamma_0 \in \mathrm{Aut}(G), (Y, S, Z)$ be given by Lemma 4.2 such that (Y, S, Z) is a proper separation of G separating the two ends of G and such that $\gamma_0 \cdot (S \cup Z) \subseteq Z$, and set $\Gamma := \langle \gamma_0 \rangle$. We let $(Y_i, S_i, Z_i) := \gamma_0^i \cdot (Y, S, Z)$ for each $i \in \mathbb{Z}$. Then $S_j \cup Z_j \subseteq Z_i$ for all $i < j$ and (Y_i, S_i, Z_i) also separates the two ends of G. For all $i \in \mathbb{Z}$, we also let $V_i := V(G) \setminus (Y_i \cup Z_{i+1})$. Then for each $i \in \mathbb{Z}$, $S_i \cup S_{i+1} \subseteq V_i$, $V_{i+1} = \gamma_0 \cdot V_i$ and as G has two ends and bounded degree, and as the finite set $S_i \cup S_{i+1}$ separates $V_i \setminus (S_i \cup S_{i+1})$ from these two ends, the graph $G_i := G[V_i]$ is finite. Moreover, for each $i \in \mathbb{Z}$, (Y_i, V_i, Z_{i+1}) is a proper separation of G. We let P be the bi-infinite path with vertex set $\mathbb{Z}$ such that $\{i, (i+1)\} \in E(P)$ for each $i \in \mathbb{Z}$. Then $(P, (V_i)_{i \in \mathbb{Z}})$ is a Γ-canonical path-decomposition satisfying the properties of (ii). □

Lemma 4.2. *Let G be a connected locally finite graph with two ends and Γ a group acting quasi-transitively on G. Then there exists a proper separation (Y, S, Z) of finite order separating the two ends of G and an element $\gamma_0 \in \Gamma$ of infinite order such that $\gamma_0 \cdot (S \cup Z) \subseteq Z$.*

Proof. We let $k \geqslant 1$ be the minimum order of a separation (Y, S, Z) separating the two ends of G, i.e., such that S is finite and both Y and Z contain an infinite component, and let (Y, S, Z) be a separation of order k separating the two ends of G. As $G[Z]$ has exactly one infinite connected component C_Z, up to considering the separation $(Y \cup (Z \setminus C_Z), S, C_Z)$, we may assume that $G[Z]$ is connected. We let $\gamma \in \Gamma$ be such that $\gamma \cdot S \subseteq Z$ (as Z is infinite and G is connected Γ-quasi-transitive, such a γ exists), and set $(Y', S', Z') := \gamma \cdot (Y, S, Z)$. If $Z' \subseteq Z$, we set $\gamma_0 := \gamma$, and claim that γ_0 must have infinite order, as an easy induction shows then that $\gamma_0^{i+1} \cdot (S \cup Z) \subseteq \gamma_0^i \cdot Z$ for each $i \geqslant 0$, and that $\gamma_0^i \cdot Z \subsetneq Z$ for each $i \geqslant 1$. If Z' is not included in Z, then as $G[Z]$ is connected (and thus $G[Z']$ also is) we must have $Y \subseteq Z'$ and $Y' \subseteq Z$. We let $\gamma' \in \Gamma$ be such that $\gamma' \cdot S \subseteq Y' \subseteq Z$, and $(Y'', S'', Z'') := \gamma' \cdot (Y, S, Z)$. Again, if $Z'' \subseteq Y'$, then $Z'' \subseteq Z$ and we conclude as before choosing $\gamma_0 := \gamma'$, thus we assume that we do not have $Z'' \subseteq Y'$. As $G[Z'']$ is connected we have $S' \subseteq Z''$ and thus $S' \cup Z' \subseteq Z''$. We thus conclude by choosing $\gamma_0 := \gamma \cdot \gamma'^{-1}$ and (Y', S', Z') playing the role of (Y, S, Z), as γ_0 satisfies $\gamma_0 \cdot (S'' \cup Z'') = S' \cup Z' \subseteq Z''$. Again, the fact that γ_0 has infinite order immediately follows. □

4.2 Quasi-transitive graphs of bounded treewidth

Tree-partition width. A *tree-partition decomposition* of a graph G is a pair $(T, \mathcal{V})$ where T is a tree and $\mathcal{V} = (V_t)_{t \in V(T)}$ is a partition of $V(G)$ whose parts are called *bags* such that for each $e \in E(G)$, the two endpoints of e lie either in a common bag or in two bags V_t, V_s such that $ts \in E(T)$. Again, the *width* of $(T, \mathcal{V})$ is the supremum of the values $|V_t|$ for $t \in V(T)$, and the *tree-partition width* $\mathrm{tpw}(G)$ of G (also called *domino treewidth* or *strong treewidth* of G) is the minimum width of a tree-partition decomposition of G. Again, if Γ acts on G, we say that a tree-partition decomposition is Γ-*canonical* if there exists an action of Γ on T compatible with the one of G, i.e., such that every automorphism $\gamma \in \Gamma$ induces an automorphism $\gamma \in \mathrm{Aut}(T)$ such that $\gamma \cdot V_t = V_{\gamma \cdot t}$ for all $t \in V(T)$.

Quasi-transitive graphs of bounded treewidth. Without always explicitly naming them, bounded treewidth quasi-transitive graphs have attracted a lot of attention and admit many interesting nontrivial characterisations of different flavours. We give a list here of some of them. Another characterisation we will mention in Chapter 2 is that finitely generated groups admitting a Cayley graph of bounded treewidth are exactly the virtually free groups.

Theorem 4.3 (Theorem 7.4 in [HLMR22], [MS83], [Woe89], [TW93], [KL05], [Ham24])**.** *Let G be a connected Γ-quasi-transitive locally finite graph. Then the following are equivalent:*

(i) *G has finite treewidth;*

(ii) *G has finite tree-partition width;*

(iii) *there exists a* Γ*-canonical tree-decomposition of* G *with tight edge-separations and finite width;*

(iii)′ *there exists a* Γ*-canonical tree-decomposition* $(T, \mathcal{V})$ *of* G *with finite width, connected parts and such that* $E(T)$ *has finitely many* Γ*-orbits;*

(iv) *there exists* $k \geqslant 1$ *such that every end of* G *has degree at most* k*;*

(v) *all the ends of* G *are thin;*

(vi) G *is quasi-isometric to a locally finite tree.*

(viii) *Every locally finite graph which is quasi-isometric to* G *excludes a (finite) minor.*

Many different proofs of some equivalences between the different items of Theorem 4.3 exist in the literature, and sometimes the proofs are only given for Cayley graphs. We give a short roadmap on which references can be used to find proofs of these equivalences the way we stated them. The equivalence between (i) and (ii) is proved in [KL05, Theorem 3.4] in the more general case of bounded degree graphs, extending a well known result in the finite case from [DO95]. The implications $(iii)' \Rightarrow (i)$, $(iii) \Rightarrow (i)$ and $(iv) \Rightarrow (v)$ are immediate. The implication $(iii) \Rightarrow (iii)'$ is an immediate consequence of Lemma 3.6. It is not hard to see that if a graph has an end of degree $k \geqslant 1$, then it admits the $k \times k$ grid as a minor, thus if G has ends of arbitrary large degree it has infinite treewidth, so $(i) \Rightarrow (iv)$ holds. The implication $(v) \Rightarrow (iii)$ follows from [HLMR22, Theorem 7.4] and a proof of the implication $(vi) \Rightarrow (ii)$ can be found for example in [Ant11, Theorem 4.7] that holds more generally for locally finite connected graphs. To get the implication $(iii) \Rightarrow (vi)$, observe that if $(T, \mathcal{V})$ is a canonical tree-decomposition of G of bounded width with tight edge-separations, then Lemma 3.1 implies that every vertex v only appears in a finite bounded number of bags V_t. In particular T is locally finite and quasi-transitive, and if we consider a mapping $f : V(G) \to V(T)$ such that for each $v \in V(G), v \in V_{f(v)}$, then it is not hard to check that f defines a quasi-isometry between G and T. The equivalence between (vi) and $(viii)$ is proved in [Ham24] and generalizes a result proved for groups in [Khu23]. We also refer to [Ant11, Theorem 4.7] for further characterisations of locally finite graphs of bounded treewidth. In view of Theorem 4.3, a natural question is the following:

Question 4.4. *If* G *is a locally finite quasi-transitive graph of bounded treewidth, then does* G *admit a* Γ*-canonical tree-partition decomposition of finite width for some group* Γ *acting quasi-transitively on* G*?*

5 Planar quasi-transitive graphs

As mentioned in the introduction, the structure of planar Cayley/transitive/quasi-transitive graphs has attracted a lot of interest. Starting with the work of Maschke [Mas96], a lot of work has been done in an attempt to give a precise description of planar Cayley graphs (see Section 13.2 for the state of the art). We refer in particular to the work of Droms [Dro06], who gave a partial answer to this question, and introduced a general method to decompose planar Cayley graphs in terms of "simpler" planar Cayley graphs. His proof consists in a clever application of Bass-Serre theory concepts. In this section, we will prove a decomposition theorem for locally finite quasi-transitive planar graphs, which is reminiscent of Droms' result. Our approach is slightly different and mainly based on an application of results of Hamann [Ham15, Ham18b] about the structure of cycles in locally finite planar quasi-transitive graphs. Roughly speaking, our main result in this section states that every 3-connected locally finite planar quasi-transitive graph admits a canonical tree-decomposition whose edge-separations correspond to separations associated to cycles in any planar drawing of G, and whose parts are vertex-accumulation-free planar graphs. As a consequence, we give a proof of Corollary 5.10, which was already proved in [HLMR22, Theorem 7.6] using the tree-amalgamation machinery. Our proof offers the advantage to be more explicit, as it does not use as a blackbox result the fact that planar quasi-transitive locally finite graphs are vertex-accessible.

5.1 Cycle nestedness in plane graphs

Recall that a graph G is *planar* if there exists an injective embedding $\varphi : G \to \mathbb{R}^2$. We call such a mapping φ a *planar embedding*, and the pair (G, φ) is called a *plane graph*. Let (G, φ) be a plane graph. By the Jordan curve theorem, for any cycle C of G, $\varphi(C)$ separates $\mathbb{R}^2$ into two disjoint arc-connected regions, one being bounded and called the *interior* of C and the other unbounded called the *exterior* of C. We say that two cycles C, C' in the plane graph (G, φ) are *nested* if $\varphi(C')$ does not intersect both the interior and the exterior regions of C. Note that this definition is symmetric, in the sense that we can exchange the roles of C and C'. Intuitively, it means that $\varphi(C)$ and $\varphi(C')$ do not cross each other. If φ is fixed, we let $V_{\text{int}}(C)$ (respectively $V_{\text{ext}}(C)$) denote the set of vertices $v \in V(G)$ such that $\varphi(v)$ belongs to the interior (respectively exterior) of C. Then $(V_{\text{int}}(C), V(C), V_{\text{ext}}(C))$ is a separation of G, and if C and C' are nested in (G, φ), then $(V_{\text{int}}(C), V(C), V_{\text{ext}}(C))$ and $(V_{\text{int}}(C'), V(C'), V_{\text{ext}}(C'))$ are nested with respect to the definition of nestedness we gave in Section 3. However note that the converse is not true in general as the fact that C and C' are nested might depend of the planar embedding of G we choose.

Recall that by Whitney's theorem, every 3-connected planar graph admits a unique embedding in the 2-dimensional sphere $\mathbb{S}^2$, up to composition with a homeomorphism of $\mathbb{S}^2$. In particular if G is planar 3-connected, then for any cycles C, C' both the unordered pair $\{V_{\text{int}}(C), V_{\text{ext}}(C)\}$ and the property for C and C' to be nested do not depend on the choice of the planar embedding φ of G. In this case, we will then not need to precise the planar embedding of G when talking about nestedness. Note also that if G is 3-connected, then for any pair of cycles C, C' and any automorphism $\gamma \in \text{Aut}(G)$, C and C' are nested if and only if $\gamma \cdot C$ and $\gamma \cdot C'$ are nested.

Cycle space. We say that a set $F \subseteq E(G)$ of edges is *even* if every vertex from $V(G)$ has even degree in the graph $(V(G), F)$. If we identify a cycle with it sets of edges, then the cycles of G are exactly the inclusionwise minimal finite nonempty sets of edges that are even. If $(C_1, \dots, C_k)$ are cycles in G, their $\mathbb{Z}_2$-*sum* $\sum_{i=1}^k C_i$ is the finite subset of $E(G)$ obtained by keeping every edge appearing in an odd number of C_i's. We let $\mathcal{C}(G)$ denote the *cycle space* of G, that is the $\mathbb{Z}_2$-vector space consisting of $\mathbb{Z}_2$-sums of cycles of G. We say that a subset $\mathcal{E}$ of $\mathcal{C}(G)$ *generates* $\mathcal{C}(G)$ if every element of $\mathcal{C}(G)$ can be written as a (finite) $\mathbb{Z}_2$-sum of elements from $\mathcal{E}$.

Remark 5.1. It is well known and not hard to check that elements from $\mathcal{C}(G)$ correspond exactly to the finite even subsets of $E(G)$.

5.2 VAP-free graphs

Given a plane graph (G, φ), an *accumulation point* is a point $x \in \mathbb{R}^2$ that contains infinitely many vertices of G in all its neighborhoods. A planar graph G is *vertex-accumulation-free* or simply *VAP-free* if it admits an embedding in $\mathbb{R}^2$ with no vertex accumulation point, or equivalently an embedding in $\mathbb{S}^2$ with at most one accumulation point. Thomassen [Tho80, Lemma 7.1] gave a proof that countable VAP-free graphs are exactly graphs admitting a planar embedding for which no cycle contains both infinitely many vertices in its interior and exterior regions. In the same paper [Tho80, Theorem 7.4], the author also characterizes 2-connected VAP-free graphs as those satisfying the MacLane planarity criterion. In particular, it implies that 2-connected VAP-free graphs are exactly graphs admitting a planar embedding φ such that the set of finite facial cycles in (G, φ) forms a basis of the cycle space $\mathcal{C}(G)$.

A known result that can be deduced from [Bab97, Lemma 2.3] is that one-ended locally finite planar graphs are VAP-free. We will show in Theorem 5.8 that locally finite VAP-free quasi-transitive graphs form the base class of graphs from which we can inductively build all locally finite quasi-transitive planar graphs. The following is a folklore result about VAP-free graphs.

Proposition 5.2. *If G is a locally finite connected VAP-free graph with at least two ends, then G has bounded treewidth.*

Proof. We let G be a locally finite VAP-free graph and $\varphi : G \to \mathbb{R}^2$ be a VAP-free planar embedding of G.

Assume that G has unbounded treewidth. Then by Theorem 4.3, G has a thick end, thus by a recent strengthening of Halin's grid theorem [GH24], G contains a subdivision H of the infinite hexagonal grid as a subgraph of G. We let ω_0 denote the end of H in G, i.e., the set of rays of G that are equivalent to any ray of H. Let r be a ray in G. We will show that $r \in \omega_0$, which immediately implies that G has a unique end, as desired. As G is connected, we may assume that its first vertex belongs to $V(H)$. Thus every vertex of r is either in $V(H)$ or drawn in a face of $(H, \varphi|_H)$. As φ is a VAP-free embedding, every facial cycle of $(H, \varphi|_H)$ contains only finitely vertices of G in its interior region. In particular, r intersects infinitely many times $V(H)$ so we have $r \in \omega_0$. □

Note that the result of Georgakopoulos and Hamann [GH24] we used in the above proof is an improvement of Halin's grid theorem [Hal65] which states that every locally finite graph with a thick end has a subgraph isomorphic to some subdivision of the hexagonal half-grid.

As the proof from [GH24] is based on the results we will present in Section 7 and on the fact that locally finite planar quasi-transitive graphs are accessible, it is worth mentionning that the above proof of Proposition 5.2 can be adapted to work if we only use Halin's original grid theorem instead.

5.3 Generating families of cycles

For every locally finite graph G and every $i \geqslant 1$, we let $\mathcal{W}_i(G)$ denote the set of closed walks of G that can be generated by the closed walks of length at most i in G (with respect to the definition of generating we gave in Section 2). Similarly we let $\mathcal{C}_i(G)$ denote the subset of $\mathcal{C}(G)$ of cycles that can be written as $\mathbb{Z}_2$-sums of cycles of length at most i.

Theorem 5.3 (Theorem 3.3 in [Ham15])**.** *Let G be a 3-connected locally finite planar graph and Γ be a group acting quasi-transitively on G. Then there exists a nested Γ-invariant set of cycles generating $\mathcal{C}(G)$. Moreover, for any $i \geqslant 0$ there exists a Γ-invariant nested family of cycles $\mathcal{E}_i$ of length at most i generating $\mathcal{C}_i(G)$.*

An analogous result to Theorem 5.3 in the context of closed walks was proved in [Ham18b, Proposition 4.3], namely for every $i \geqslant 0$, every 3-connected quasi-transitive locally finite graph admits an $\mathrm{Aut}(G)$-invariant nested set of cycles generating $\mathcal{W}_i(G)$ (with respect to the definition of generating given in Section 2). In the same paper, the author also proved the following result, which can be seen as a generalization of the result of [Dro06] that finitely generated planar groups are finitely presented.

Theorem 5.4 (Theorem 7.2 in [Ham18b])**.** *Let G be a quasi-transitive planar graph and Γ be a group acting quasi-transitively on G. Then there exists a Γ-invariant set of cycles $\mathcal{E}$ generating $\mathcal{C}(G)$ with finitely many Γ-orbits.*

Again, Theorem 5.4 also admits a similar statement for closed walks [Ham18b, Theorem 5.12]. Despite the fact that the proof of Theorem 5.4 from [Ham18b] is based on Theorem 5.3, the family which is constructed in Theorem 5.4 is not necessarily nested anymore. However we will observe in Corollary 5.6 that combining Theorems 5.3 and 5.4, we can find in the 3-connected case a generating family of cycles which is both nested and has finitely many $\mathrm{Aut}(G)$-orbits. The following is a remark of Matthias Hamann (private communication).

Remark 5.5. Theorem 5.4 was stated in [Ham18b] in a more general way for $\mathbb{Z}$-sums of oriented cycles, i.e., formal sums of oriented cycles with coefficients in $\mathbb{Z}$. More precisely, if we fix an orientation A of $E(G)$ then for each oriented cycle $\overrightarrow{C}$ given by a sequence of oriented edges $(e_1, \ldots, e_k)$, we can associate a vector x in $\{-1, 0, 1\}^A$ with finite support such that for each $(u, v) \in A$,

$$x_{(u,v)} := \begin{cases} 1 & \text{if } e_i = (u, v), \text{ for some } i \in [k], \\ -1 & \text{if } e_i = (v, u), \text{ for some } i \in [k], \text{ and} \\ 0 & \text{otherwise.} \end{cases}$$

In particular the two opposite orientations of a cycle correspond to opposite vectors. If $\overrightarrow{C}$ is an oriented cycle, we let $C \subseteq E(G)$ be its associated unoriented cycle. The *first* *simplicial homology group* $\mathcal{H}_1(G)$ is the $\mathbb{Z}$-submodule of the direct sum $\bigoplus_A \mathbb{Z}$ generated by the (oriented) cycles of G, i.e., the set of finite sums of vectors associated to cycles of G. To simplify notations, we identify oriented cycles of G with their associated vectors. To show that [Ham18b, Theorem 7.2] implies Theorem 5.4 the way we stated it, observe that if $\overrightarrow{\mathcal{E}}$ is a family of $\mathbb{Z}$-vectors associated to oriented cycles that generates $\mathcal{H}_1(G)$, then the associated family $\mathcal{E} := \{C : \overrightarrow{C} \in \overrightarrow{\mathcal{E}}\}$ of unoriented cycles generates the cycle space $\mathcal{C}(G)$ when considering $\mathbb{Z}_2$-sums. For this, we let $C \subseteq E(G)$ be a cycle and we show that C is a $\mathbb{Z}_2$-sum of cycles from $\mathcal{E}$. Let $\overrightarrow{C}$ be any orientation of C and $x \in \mathbb{Z}^A$ be its associated vector. Then there exist oriented cycles $\overrightarrow{C_1}, \dots, \overrightarrow{C_k} \in \overrightarrow{\mathcal{E}}$ with associated vectors $x_1, \dots, x_k \in \mathbb{Z}^A$ such that $x = \sum_{i=1}^k x_i$. Note that there might be repetitions, i.e., the x_i's are not necessarily distinct. We show that C equals to the $\mathbb{Z}_2$-sum $\sum_{i=1}^k C_i$. First, note that for each $(u, v) \in A$ such that $uv \notin C$, we have $x_{(u,v)} = 0$ so the arc (u, v) must appear in the same number of x_i's as the arc (v, u) so uv appears an even number of times in the cycles (with possible repetitions) $C_1, \dots, C_k$ and thus $uv \notin \sum_{i=1}^k C_i$. Now if $(u, v) \in C$ for some $(u, v) \in A$, we must have $x_{(u,v)} \in \{-1, 1\}$. Assume without loss of generality that $x_{(u,v)} = 1$, the other case being symmetric. Then (u, v) must appear exactly $m + 1$ times in the cycles $\overrightarrow{C_i}$, while (v, u) appears m times in the oriented cycles $\overrightarrow{C_i}$ for some $m \in \mathbb{N}$. It means that uv appears in total $2m + 1$ times in the cycles C_i, showing that $uv \in \sum_{i=1}^k C_i$. We thus proved that $C = \sum_{i=1}^k C_i$, as desired.

Again, Theorem 5.4 still holds (see [Ham18b, Theorem 5.12]) if we replace $\mathcal{C}(G)$ with $\mathcal{W}(G)$ and consider the definition of generating family for closed walks we gave in Section 2.

We observe that in the 3-connected case, one can find a generating family $\mathcal{E}$ of cycles combining both the properties of Theorems 5.3 and 5.4.

Corollary 5.6. *Let G be a locally finite 3-connected planar graph and Γ be a group acting quasi-transitively on G. Then there exists a Γ-invariant set of cycles generating $\mathcal{C}(G)$ which is nested and has finitely many Γ-orbits.*

An example of a family satisfying the properties of Corollary 5.6 is given in Figure 1.2 below.

Proof. We let $\mathcal{E}$ be a Γ-invariant family of cycles generating $\mathcal{C}(G)$ with finitely many Γ-orbits given by Theorem 5.4. Then in particular there is a bound $K \geqslant 0$ on the size of the cycles from $\mathcal{E}$. By Theorem 5.3, there exists a nested Γ-invariant family $\mathcal{E}'$ of cycles of lenth at most K in G generating the set $\mathcal{C}_K(G)$. In particular, $\mathcal{E}'$ also generates the whole cycle space $\mathcal{C}(G)$. As G has bounded degree, every vertex $v \in V(G)$ belongs to only finitely many cycles of size at most K. In particular, as Γ acts quasi-transitively on $V(G)$, it implies that Γ also acts quasi-transitively on the set of cycles of size at most K in G. Thus $\mathcal{E}'$ has finitely many Γ-orbits and satisfies the desired properties. □

5.4 Structure of locally finite quasi-transitive planar graphs

If $\mathcal{N}$ is a set of separations, an $\mathcal{N}$-*block* is a maximal set $X \subseteq V(G)$ such that for each $(Y, S, Z) \in \mathcal{N}$, either $X \cap Y = \emptyset$ or $X \cap Z = \emptyset$. If (G, φ) is a plane graph and $\mathcal{E}$ is a set of

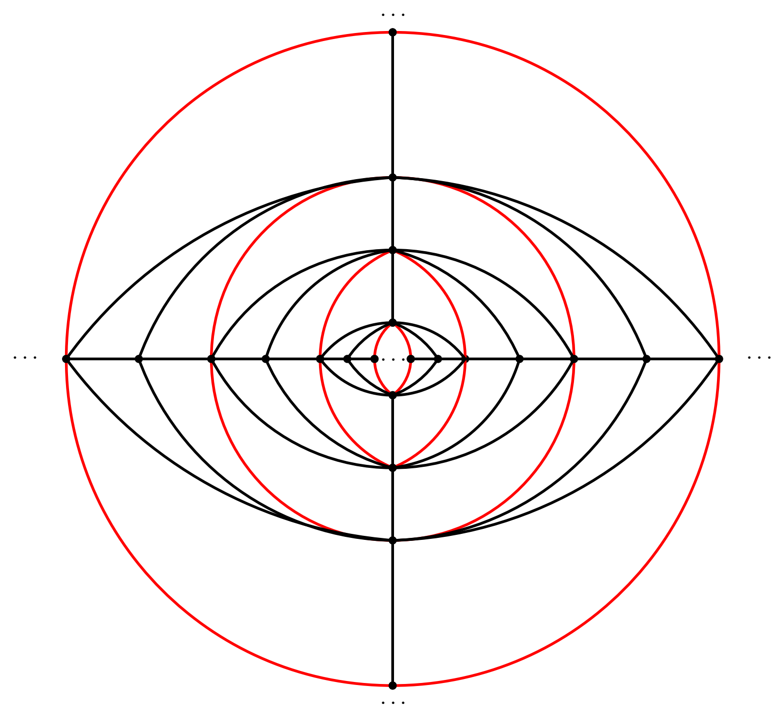

Figure 1.2: A 2-ended locally finite quasi-transitive 3-connected planar graph G. The set of cycles $\mathcal{E}$ formed by the union of the red cycles together with the set of facial cycles of G forms a nested $\mathrm{Aut}(G)$-invariant generating family of the cycle space of G. The $\mathcal{E}$-blocks are the subgraphs obtained after taking two consecutive red cycles, together with the vertices and edges lying between them.

cycles, then an $\mathcal{E}$*-block* of (G, φ) is a set of vertices which is an $\mathcal{N}$-block, where $\mathcal{N}$ denotes the symmetric set of separations induced by $\mathcal{E}$ in (G, φ).

Lemma 5.7. *Let (G, φ) be a 3-connected locally finite plane graph, Γ be a group acting quasi-transitively on G and $\mathcal{E}$ be a Γ-invariant nested family of cycles G of bounded length generating the cycle space $\mathcal{C}(G)$. Then for each $\mathcal{E}$-block X, the family $\mathcal{E}_X := \{C \in \mathcal{E} : V(C) \subseteq X\}$ generates the cycle space $\mathcal{C}(G[X])$.*

Proof. In this proof, we will identify every cycle of G with its even set of edges.

We let C be a cycle of $\mathcal{C}(G[X])$ and $C_1, \dots, C_k \in \mathcal{E}$ be such that C equals to the $\mathbb{Z}_2$-sum $\sum_{i=1}^k C_i$. Choose $C_1, \dots, C_k$ that minimize the number k of cycles from $\mathcal{E}$ required to write C as a $\mathbb{Z}_2$-sum $\sum_{i=1}^k C_i$. As X is an $\mathcal{E}$-block of (G, φ), every cycle from $\mathcal{E}_X$ must be facial in the plane graph $(G[X], \varphi|_{G[X]})$, and C is nested with every cycle from $\mathcal{E}$. We will show that $C_i \in \mathcal{E}_X$ for all $i \in [k]$, implying the desired result.

Assume for a contradiction that $C_i \notin \mathcal{E}_X$ for some $i \in [k]$. As cycles from $\mathcal{E}$ have bounded length, by Lemma 3.8, the set $\mathcal{N}$ of separations induced by $\mathcal{E} \cup \{C\}$ in (G, φ) has finite intervals. We let $(Y_1, S_1, Z_1), (Y_2, S_2, Z_2)$ be two separations respectively induced by C and C_i in the plane graph (G, φ) be such that $(Y_1, S_1, Z_1) \leqslant_{\mathrm{RS}} (Y_2, S_2, Z_2)$. In particular, we have $(Y_2, S_2, Z_2) \in \mathcal{N}$. As $\mathcal{N} \cup \{C\}$ is nested, $\leqslant_{\mathrm{RS}}$ induces a total order on the finitely many separations from $\mathcal{N}$ sandwiched between (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) with respect to $\leqslant_{\mathrm{RS}}$.

In particular, there exists a minimal separation $(Y, S, Z) \in \mathcal{N}$ with respect to $\leqslant_{\mathrm{RS}}$ such that $(Y_1, S_1, Z_1) \leqslant_{\mathrm{RS}} (Y, S, Z) \leqslant_{\mathrm{RS}} (Y_2, S_2, Z_2)$. We let C^* be the associated cycle of $\mathcal{E}$ such that (Y, S, Z) is one of the two symmetric separations induced by C^* in (G, φ). We claim that $C^* \in \mathcal{E}_X$, as if it was not the case, then we could add its vertex set to X and still have some subset of $V(G)$ not separated by cycles from $\mathcal{E}$, and thus contradicting the fact that X is an $\mathcal{E}$-block. We now let D be the cycle associated to the maximal separation $(Y', S', Z') \in \mathcal{E}$ such that $(Y, S, Z) \leqslant_{\mathrm{RS}} (Y', S', Z') \leqslant_{\mathrm{RS}} (Y_2, S_2, Z_2)$ and such that $D \in \mathcal{E}_X$. In particular, by the previous observation D is facial in $(G[X], \varphi|_{G[X]})$ so $\varphi(C_i)$ must be contained in the adhesion of the face Λ of $(G[X], \varphi|_{G[X]})$ which is delimited by D. We let $I \subseteq [k]$ be the set of indices $j \in [k]$ such that $\varphi(C_j)$ is contained in the adhesion of Λ. In particular, $i \in I$ so $I \neq \emptyset$. As $\mathcal{E}$ is nested, for every $j \in [k] \setminus I$, $\varphi(C_j)$ does not intersect Λ. Let C' be the $\mathbb{Z}_2$-sum $\sum_{j \in I} C_j$. Note that the way we defined it, C' is a finite subset of edges of $E(G)$ but not necessarily a cycle of G.

First, note that for each $uv \in E(G) \setminus E(G[X])$, as $uv \notin C$, it must appear in an even number of C_j's. In particular, as we assumed that $I \neq \emptyset$, we must have $|I| \geqslant 2$. Observe that if $\varphi(uv)$ intersects Λ then uv can only appear in cycles C_j such that $j \in I$, and its total number of occurences among the C_j's is even, so $uv \notin C'$. It implies that $C' \subseteq D$. By Remark 5.1, C' is even so as D is a cycle, we have either $C' = \emptyset$ or $C' = D$. According to whether $C' = \emptyset$ or $C' = D$, we consider the decomposition of C as a sum of cycles from $\mathcal{E}$ obtained after either removing the sum $\sum_{j \in I} C_j$ in the decomposition of C or replacing it by the cycle $D \in \mathcal{E}_X$. In both cases it gives a decomposition of C involving at most $k - |I| + 1 < k$ cycles from $\mathcal{E}$, and thus contradicting the minimality of k. □

We are now ready to give the main result of this subsection. This result can be seen as a quasi-transitive version of Droms structure theorem for planar Cayley graphs [Dro06], which states that each planar Cayley graph can be obtained inductively by gluing together (with respect to some specific algebraic operations) planar finite or one-ended Cayley graphs by identifying boundaries of some of their facial cycles.

Theorem 5.8. *Let G be a locally finite 3-connected quasi-transitive planar graph and Γ be a group acting quasi-transitively on G. Then there exists a Γ-canonical tree-decomposition $(T, \mathcal{V})$ of G of finite adhesion whose edge-separations correspond to separations associated to cycles of G and whose parts are 2-connected VAP-free quasi-transitive graphs. Moreover $E(T)$ has finitely many Γ-orbits.*

Proof. We let $\mathcal{E}$ be a nested Γ-invariant family of cycles of G generating $\mathcal{C}(G)$ with finitely many Γ-orbits given by Corollary 5.6. We consider the associated symmetric family $\mathcal{N}$ of separations of G of the form $(V_{\mathrm{int}}(C), V(C), V_{\mathrm{ext}}(C))$ and $(V_{\mathrm{ext}}(C), V(C), V_{\mathrm{int}}(C))$ for each $C \in \mathcal{E}$. As G is 3-connected, our previous remarks imply that $\mathcal{N}$ is a Γ-invariant nested family of separations (with respect to the definition of nestedness given in Section 3). Moreover, as $\mathcal{E}$ has finitely many Γ-orbits, separations in $\mathcal{N}$ must have finite bounded order so Lemma 3.8 implies that $\mathcal{N}$ has finite intervals. We thus can apply Theorem 3.7 and find a Γ-canonical tree-decomposition $(T, \mathcal{V})$ whose edge-separations are in one-to-one correspondence with the different edge-separations of $\mathcal{N}$. In particular, each adhesion set of $(T, \mathcal{V})$ admits a spanning cycle from $\mathcal{N}$ and thus is finite. As $\mathcal{N}$ has finitely many Γ-orbits, Γ acts quasi-transitively on $E(T)$. By Lemma 3.17, every part $G[V_t]$ of $(T,\mathcal{V})$ is Γ_t-quasi-transitive.

As the adhesions of $(T,\mathcal{V})$ are connected, each part $G[V_t]$ must also be connected. Moreover, note that as adhesion sets of $(T,\mathcal{V})$ contain spanning cycles, then for every $t \in V(T)$, $|V_t| \geqslant 3$ and for any three different vertices $u, v, w \in V_t$, any path in G from u to v avoiding w can be modified to a path in $G[V_t]$ from u to v avoiding w. Hence each part of $(T,\mathcal{V})$ is 2-connected.

It remains to show that each part of $(T,\mathcal{V})$ is VAP-free. By [CDHS11, Theorem 4.8][1], parts of $(T,\mathcal{V})$ are either "hubs", i.e., vertex sets of cycles from $\mathcal{E}$, or $\mathcal{N}$-blocks (and equivalently $\mathcal{E}$-blocks). Hubs parts are finite and thus obviously VAP-free. Assume now that $G[V_t]$ is an $\mathcal{E}$-block for some $t \in V(T)$. Then by Lemma 5.7, $\mathcal{E}_{V_t}$ generates the cycle space $\mathcal{C}(G[V_t])$. In particular, note that cycles from $\mathcal{E}_{V_t}$ must be facial in the plane graph $(G[V_t], \varphi|_{G[V_t]})$. The plane graph $(G[V_t], \varphi|_{G[V_t]})$ is thus 2-connected and its cycle space is generated by a family of facial walks, so by [Tho80, Theorem 7.4] it must be a VAP-free graph. □

Corollary 5.9. *For every locally finite 3-connected quasi-transitive planar graph G, and every group Γ acting quasi-transitively on G, there exists a Γ-canonical tree-decomposition $(T,\mathcal{V})$ of G of finite adhesion whose parts are connected and either finite or one-ended and such that $E(T)$ has finitely many Γ-orbits.*

Proof. Let $(T,\mathcal{V})$ be the Γ-canonical tree-decomposition of G given by Theorem 5.8. We let $\{t_i : i \in I_\infty\}$ be a (finite) set of representatives of the orbits $V_\infty(T)/\Gamma$. For each $i \in I_\infty$ such that $G[V_{t_i}]$ has at least 2 ends, Proposition 5.2 implies that $G[V_{t_i}]$ has bounded treewidth. By Lemma 3.5, $G[\![V_{t_i}]\!]$ is quasi-isometric to $G[V_{t_i}]$, so by Theorem 4.3 (vi) it also has bounded treewidth. By Theorem 4.3 $(iii)'$, there exists a Γ_{t_i}-canonical tree-decomposition $(T_{t_i}, \mathcal{V}_{t_i})$ of $G[\![V_{t_i}]\!]$ of finite width whose parts are connected, and such that $E(T_{t_i})$ has finitely many Γ_{t_i}-orbits. Then by Corollary 3.15 together with Remark 3.16, there exists a Γ-canonical tree-decomposition $(T',\mathcal{V}')$ of G refining $(T,\mathcal{V})$ whose torsos are connected with at most one end, and whose adhesion sets are either adhesion sets of $(T,\mathcal{V})$ or adhesion sets of some $(T_t,\mathcal{V}_t)$. In particular, as G is locally finite quasi-transitive, every finite set is the separator of a finite bounded number of separations, hence $E(T')$ must have only finitely many Γ-orbits. Finally, we find a tree-decomposition of G with the desired properties by applying Lemma 3.6 to $(T',\mathcal{V}')$. □

See Figure 1.3 below for an illustration of the tree-decomposition obtained (which turns out to be a path-decomposition in this specific example) when applying the proof with of Theorem 5.8 with respect to the family of cycles from Figure 1.2.

For general planar quasi-transitive graphs, one gets the following

Corollary 5.10. *For every connected planar locally finite graph G, and every group Γ acting quasi-transitively on G, there exists a Γ-canonical tree-decomposition $(T,\mathcal{V})$ of finite adhesion whose parts $G[V_t]$ are connected, either finite or one-ended planar graphs, on which Γ_t acts quasi-transitively for each $t \in V(T)$, and such that $E(T)$ has finitely many Γ-orbits.*

[1]Note that [CDHS11] only deals with finite graphs. However, as explained in [CHM22, Proposition 3.2], the proof of [CDHS11, Proposition 4.8] generalizes to the infinite case when one considers nested sets of separations having finite intervals.

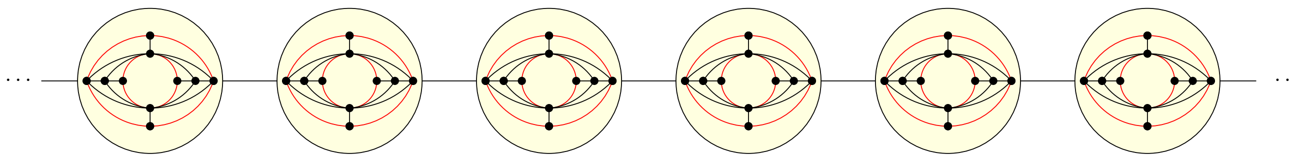

Figure 1.3: The path-decomposition of the graph from Figure 1.2 obtained after applying theorem 3.7 to the nested family of red cycles. The red cycles form the adhesion sets of the path-decomposition.

Proof. We first consider Tutte's canonical tree-decomposition $(T_0, \mathcal{V}_0)$ of G given by Theorem 3.9. We let G^+ be the supergraph obtained from G after adding an edge uv for each pair of vertices u, v belonging to a common adhesion set of $(T_0, \mathcal{V}_0)$. In particular for each $t \in V(T_0)$, $G[\![V_t]\!] = G^+[V_t]$. As the edge-separations of $(T_0, \mathcal{V}_0)$ are tight, Lemma 3.1 implies that for each $v \in V(G)$, there is only a finite bounded number of edges $tt' \in E(T_0)$ such that $v \in V_t \cap V_{t'}$. In particular, for each $v \in V(G)$, there is only a finite bounded number of $t \in V(T_0)$ such that $v \in V_t$. Thus G^+ is also locally finite, and as $(T_0, \mathcal{V}_0)$ is Γ-canonical, Γ also acts quasi-transitively on G^+. We will now show that G^+ is planar. Note that $(T_0, \mathcal{V}_0)$ also corresponds to Tutte's decomposition of G^+, and as every torso $G[\![V_t]\!]$ of $(T_0, \mathcal{V}_0)$ is a minor of G, every torso $G[\![V_t]\!]$ is planar. By a result attributed to Erdös (see for example [Tho80]), a countable graph is planar if and only it excludes $K_{3,3}$ and K_5 as minors. In particular it implies that a countable graph is planar if and only if all its finite subgraphs are planar, so it is enough to check that every finite subgraph of G^+ is planar. As the adhesion sets of $(T_0, \mathcal{V}_0)$ induce complete graphs in G^+ and have size at most 2, note that any finite subgraph of G^+ is obtained after perfoming the following operation a finite number of times: taking two disjoint planar graphs G_1, G_2 with two edges $u_1v_1 \in E(G_1)$ and $u_2v_2 \in E(G_2)$, and gluing them by identifying u_1 with u_2 and v_1 with v_2. Note that such an operation does preserve planarity, hence every finite subgraph of G^+ must be planar and we deduce that G^+ is also planar.

Claim 5.11. If G^+ admits a Γ-canonical tree-decomposition $(T, \mathcal{V})$ of finite adhesion whose parts are connected and finite or one-ended and such that $E(T)$ has finitely many Γ-orbits, then G also admits a Γ-canonical tree-decomposition with the same properties.

Proof of the Claim: Note that $(T, \mathcal{V})$ is also a Γ-canonical tree-decomposition of G. As $E(T_0)$ has finitely many Γ-orbits, note that the set $\{d(u, v) : uv \in E(G^+) \setminus E(G)\}$ must be bounded and thus admits a maximum, say $k_1 \in \mathbb{N}$. As $E(T)$ has finitely many Γ-orbits, the set $\{d_G(u, v) : \exists t \in V(T), uv \in E(G[\![V_t]\!]) \setminus E(G[V_t])\}$ is also bounded, and admits a maximum $k_2 \in \mathbb{N}$. We set $k := \max(k_1, k_2)$ and let $\mathcal{V}' := (V'_t)_{t \in V(T)}$ be defined by $V'_t := B_k(V_t) = \{v \in V(G) : \exists u \in V_t, d_G(u, v) \leqslant k\}$. We claim that the proof of Lemma 3.6 still works here and implies that $(T, \mathcal{V}')$ is a Γ-canonical tree-decomposition of G^+ of finite adhesion.

As G is a subgraph of G^+, $(T, \mathcal{V}')$ is also a Γ-canonical tree-decomposition of G. Finally, note that for every $t \in V(T)$ and every $uv \in E(G^+[V_t]) \setminus E(G)$, as $k \geqslant k_1$ there exists a path from u to v in $G[V'_t]$, so the parts of $(T, \mathcal{V}')$ are connected when considered as a tree-decomposition of G. It remains to show that $G[V'_t]$ has at most one end for each $t \in V(T)$.

We will show that $G[V_t']$ is quasi-isometric to $G^+[V_t]$, which immediately implies the desired result. By Lemma 3.5, $G^+[\![V_t]\!]$ is quasi-isometric to $G^+[V_t]$, hence it is enough to show that $G[V_t']$ is quasi-isometric to $G^+[\![V_t]\!]$.

We claim that the arguments from the proof of Lemma 3.6 can be reproduced almost everywhere. Fix $t \in V(T)$ and let $\pi : V_t' \to V_t$ be such that for all $v \in V_t'$, $d_G(\pi(v), v) = d_G(V_t, v)$. For exactly the same reasons that in the proof of Lemma 3.6, for every $u, v \in V_t'$, if there exists a path of length d from u to v in $G[V_t']$, then there exists a path of length at most $d + 2k$ from $\pi(u)$ to $\pi(v)$ in $G[\![V_t]\!]$, and thus also in $G^+[\![V_t]\!]$ which contains $G[\![V_t]\!]$ as a subgraph, so the inequality

$$d_{G^+[\![V_t]\!]}(\pi(u), \pi(v)) \leqslant d_{G[V_t']}(u, v) + 2k$$

still holds. Conversely if p is a path from $\pi(u)$ to $\pi(v)$ in $G^+[\![V_t]\!]$ of length d, then after replacing every edge of p that belong to $E^+(G) \setminus E(G)$ by a path of length at most k_1 contained in $G[V_t']$ and every edge of $E(G^+[\![V_t]\!]) \setminus E(G^+) = E(G[\![V_t]\!]) \setminus E(G)$ by a path of length at most k_2 in $G[V_t']$, and after extending it to a path from u to v, we obtain a path from $\pi(u)$ to $\pi(v)$ in $G[V_t']$ of length at most $kd + 2k$, showing that

$$d_{G[V_t']}(u, v) \leqslant k \cdot d_{G[\![V_t]\!]}(\pi(u), \pi(v)) + 2k.$$

Hence π defines a quasi-isometry from $G^+[\![V_t]\!]$ to $G[V_t']$, so in particular $G[V_t']$ is also quasi-isometric to $G^+[V_t]$ and thus has at most one end. $\Diamond$

Claim 5.11 allows us to assume without loss of generality that $G^+ = G$, i.e., that for each $t \in V(T_0)$ we have $G[\![V_t]\!] = G[V_t]$. We let $\{t_i : i \in I_\infty\}$ be a (finite) set of representatives of the orbits $V_\infty(T_0)/\Gamma$. For each $i \in I_\infty$, as $G[V_{t_i}]$ is 3-connected, Corollary 5.9 implies that there exists a Γ_{t_i}-canonical tree-decomposition $(T_{t_i}, \mathcal{V}_{t_i})$ of $G[V_{t_i}]$ with finite adhesion whose parts are connected and have at most one end, and such that $E(T_{t_i})$ has finitely many Γ_{t_i}-orbits. In particular, by Corollary 3.15 together with Remark 3.16, there exists a Γ-canonical tree-decomposition $(T, \mathcal{V})$ of G refining $(T_0, \mathcal{V}_0)$, whose parts are connected and either finite or one-ended, and such that $E(T)$ has finitely many Γ-orbits. $\square$

5.5 Quasi-transitive graphs of bounded genus

We observe that if we consider quasi-transitive graphs, the property of being embeddable in a surface of bounded genus is not much more general than the property of being planar. In the finite case, Thomassen [Tho91] and Babai [Bab91] proved simultaneously that for any genus $g \geqslant 3$, there exists only finitely many finite transitive graphs of genus exactly g. For infinite graphs, the following well-known observation can be done (see for example [Lev70]).

Remark 5.12. Let G be a locally finite connected quasi-transitive graph which is embeddable in a surface of finite genus $g \in \mathbb{N}$. Then G must be either finite or planar. Indeed, if we assume that G is not planar, then by Wagner's theorem, G has a minor isomorphic to F, where $F \in \{K_5, K_{3,3}\}$. In particular, note that as F is finite, we can find a model $(M_v)_{v \in V(F)}$ such that M_v is finite for each $v \in V(F)$. By quasi-transitivity of G, we can then find a model of $n \cdot F$ for arbitrary large $n \in \mathbb{N}$, where $n \cdot F$ denotes the finite graph consisting in n

pairwise-disjoint copies of F. Then G is a connected graph with $n \cdot F$ as a minor, so using for example [Mil87, Theorem 1] we deduce that it must have genus at least n.

6 Tangles and structure of 3-connected graphs

Tangles were initially introduced for finite graphs by Robertson and Seymour [RS91] and play a fundamental role in their proof of the Graph Minor Structure Theorem. Intuitively, a *tangle* $\mathcal{T}$ *of order* k is a choice for every separation $(Y, S, Z) \in \mathrm{Sep}_{<k}(G)$ of an "orientation" of (Y, S, Z), i.e., of one of the two symmetric separations (Y, S, Z) and (Z, S, Y) such that the separations of $\mathrm{Sep}_{<k}(G)$ are oriented in a "consistent way", and that they point towards a zone of G which is "highly connected". For example, a tangle of order 2 is exactly an orientation of the edges of the block cut-tree $(T, \mathcal{V})$ of G pointing either towards a vertex of T, i.e., to a block, or towards an end of T, which can be thought as a block of G located at infinity. In general, it is true that for any tree-decomposition $(T, \mathcal{V})$ of G of adhesion at most k with tight edge-separations, any tangle of order $k+1$ orients the edge-separations of $(T, \mathcal{V})$ in such a way, i.e., such that the corresponding orientation of $E(T)$ points either towards a node or towards an end of T.

Tangles of order k are usually presented as a dual object to tree-decompositions, in the sense that if a graph is not well-connected, then it should admit at least one tangle of low order, and an associated tree-decomposition with small adhesion associated to such a tangle. A central result from [RS91] is that every (finite) graph admits a tree-decomposition $(T, \mathcal{V})$ of adhesion at most k distinguishing all its tangles of order $k+1$, i.e., such that every two tangles of order k give a different orientation of $E(T)$. In [CHM22] (see Theorem 6.3), the authors extended this result to every locally finite graph and more importantly they also proved that such a tree-decomposition $(T, \mathcal{V})$ can be chosen to be canonical, extending and unifying many already known results from [CDHS11, DK15, DHL18].

In this section, we present some basic properties of tangles, and we explain and generalize the main ideas from [Gro16] in which the author did an extensive study of the structure of the tangles of order 4 in finite 3-connected graphs. These results will then be used in Section 7 to derive a general structure theorem for minor-excluded locally finite quasi-transitive graphs. More precisely, we give in Section 6.1 the definition of tangles, together with some basic properties, and illustrate them on a specific example in Section 6.2. The main objective of Sections 6.3, 6.4 and 6.5 is to extend results from [Gro16] to locally finite graphs, that we summarized in Theorem 6.26. More precisely, Theorem 6.26 implies that if G is a 3-connected locally finite graph with a unique tangle $\mathcal{T}$ of order 4, then G admits a canonical tree-decomposition $(T, \mathcal{V})$ where T is a star, and all the torsos are finite, except possibly the central torso $G[\![V_{t_0}]\!]$. Moreover, $G[\![V_{t_0}]\!]$ admits an $\mathrm{Aut}(G)$-invariant matching $M \subseteq E(G)$ such that the graph $G[\![V_{t_0}]\!]/M$ obtained from $G[\![V_{t_0}]\!]$ after contracting the edges of M is quasi-4-connected. For our future applications, we will also need to prove in Section 6.6 that if the matching M given by Theorem 6.26 is such that $G[\![V_{t_0}]\!]/M$ is planar, then $G[\![V_{t_0}]\!]$ is also planar. Note that Sections 6 and 7 will be the only parts of the manuscript where we use tangles.

6.1 Tangles

Tangles. We consider here the definition of tangles used by Grohe [Gro16], which slightly differs from the original one from Robertson Seymour [RS91]. We refer to Appendix A from [Gro16] for a correspondence between the two definitions. A *tangle of order* k in a graph G is a subset $\mathcal{T}$ of $\mathrm{Sep}_{<k}(G)$ such that

($\mathcal{T}$1) For all separations $(Y, S, Z) \in \mathrm{Sep}_{<k}(G)$, either $(Y, S, Z) \in \mathcal{T}$ or $(Z, S, Y) \in \mathcal{T}$;

($\mathcal{T}$2) For all separations $(Y_1, S_1, Z_1), (Y_2, S_2, Z_2), (Y_3, S_3, Z_3) \in \mathcal{T}$, either $Z_1 \cap Z_2 \cap Z_3 \neq \emptyset$ or there exists an edge with an endpoint in each Z_i.

Note that ($\mathcal{T}$2) with $(Y_i, S_i, Z_i) = (Y, S, Z)$ for each $i \in [3]$ implies in particular that for every separation $(Y, S, Z) \in \mathcal{T}$, $Z \neq \emptyset$. A tangle in G will be called a G-*tangle*, for brevity. Intuitively, a G-tangle is a consistent orientation of the separations of G, pointing towards a highly connected region of G. More precisely, in finite graphs tangles were initially introduced to generalize the notion of k-*block*. A set of vertices $X \subseteq V(G)$ is a k-*block* for some $k \in \mathbb{N}$ if $|X| \geqslant k$ and if X is an inclusionwise maximal $(k-1)$-*inseparable* set, i.e., for every proper separation $(Y, S, Z) \in \mathrm{Sep}_{<k}(G)$, we have either $X \subseteq Y \cup S$ or $X \subseteq S \cup Z$. When $k = 2$, k-blocks are simply called *blocks*. In general when G is finite there is a one-to-one correspondence between the G-tangles of order 1 and the connected components of G, between the G-tangles of order 2 and the *biconnected* components of G (that is the blocks of G), and between the G-tangles of order 3 and the *triconnected* components of G (that is the 2-blocks of G) [RS91, Gro16].

One direction of this correspondence still holds for general values of k. If $X \subseteq V(G)$ is a k-block for some $k \in \mathbb{N} \setminus \{0\}$, then we define the subset $\mathcal{T}_X$ of $\mathrm{Sep}_{<k}(G)$ by orienting every separation of order at most $k-1$ towards X, i.e., by setting

$$\mathcal{T}_X := \{(Y, S, Z) \in \mathrm{Sep}_{<k} : X \subseteq S \cup Z\}.$$

It was proved in [Gro16, Lemma 3.3] that if X is a k-block such that $|X| > \frac{3}{2} \cdot (k-1)$, then $\mathcal{T}_X$ defines a G-tangle of order k.

In infinite graphs, one can similarly see tangles as a notion generalizing the notion of ends: for each end ω of a graph G and every $k \geqslant \mathbb{N} \setminus \{0\}$, we define the G-tangle $\mathcal{T}_\omega^k$ *of order* k *induced by* ω by

$$\mathcal{T}_\omega^k := \{(Y, S, Z) \in \mathrm{Sep}_{<k}(G) : \omega \text{ lives in a component of } Z\}.$$

The proof that $\mathcal{T}_\omega^k$ is a tangle is an easy exercise.

Projection and lifting. One of the basic properties of tangles is that for any fixed model of a graph H in a graph G, any H-tangle of order k induces a G-tangle of order k. More precisely, if $\mathcal{M} = (M_v)_{v \in V(H)}$ is a model of H in G and (Y, S, Z) is a separation of order less than k in G, then its *projection* with respect to $\mathcal{M}$ is the separation $\pi_{\mathcal{M}}(Y, S, Z) = (Y', S', Z')$ of H of order less than k defined by $Y' := \{v \in V(H) : M_v \subseteq Y\}$, $S' := \{v \in V(H) : M_v \cap S \neq \emptyset\}$ and $Z' := \{v \in V(H) : M_v \subseteq Z\}$.

A proof of the following result can be found in [RS91, (6.1)] or in a more similar version in [Gro16, Lemma 3.11]. Its proof immediately extends to the locally finite case.

Lemma 6.1 (Lemma 3.11 in [Gro16]). *Let G be a locally finite graph. Let $\mathcal{M} = (M_v)_{v\in V(H)}$ be a model of a graph H in G and $\mathcal{T}'$ be an H-tangle of order k. Then the set*

$$\mathcal{T} := \{(Y, S, Z) \in \mathrm{Sep}_{<k}(G) : \pi_{\mathcal{M}}(Y, S, Z) \in \mathcal{T}'\}$$

is a G-tangle of order k, called the lifting *of $\mathcal{T}'$ in G with respect to $\mathcal{M}$.*

Remark 6.2. Assume that $\mathcal{M}$ is a faithful model of H in G with the property that for each $(Y', S', Z') \in \mathrm{Sep}_{<k}(H)$, there exists some $(Y, S, Z) \in \mathrm{Sep}_{<k}(G)$ such that $\pi_{\mathcal{M}}(Y, S, Z) = (Y', S', Z')$ and $S' = S$. Then the function that maps every tangle of order k in H to its lifting in G with respect to $\mathcal{M}$ is injective: if $\mathcal{T}_1' \neq \mathcal{T}_2'$ are two distinct tangles of order k in H, then there exists some $(Y', S', Z') \in \mathcal{T}_1'$ such that $(Z', S', Y') \in \mathcal{T}_2'$. In particular if we consider $(Y, S, Z) \in \mathrm{Sep}_{<k}(G)$ such that $\pi_{\mathcal{M}}(Y, S, Z) = (Y', S', Z')$, we have $(Y, S, Z) \in \mathcal{T}_1$ and $(Z, S, Y) \in \mathcal{T}_2$, where for each $i \in \{1, 2\}$, $\mathcal{T}_i$ denotes the lifting of $\mathcal{T}_i'$ with respect to $\mathcal{M}$. It then follows that $\mathcal{T}_1 \neq \mathcal{T}_2$, as desired. Note that if $(T, \mathcal{V})$ is a tree-decomposition with finitely bounded adhesion and $t \in V(T)$ is such that $G[\![V_t]\!]$ is a faithful minor of G, then any faithful model $\mathcal{M}$ of $G[\![V_t]\!]$ has the property we just described.

If $\mathcal{M} = (M_v)_{v\in V(H)}$ is a model of H in G, and $\mathcal{T}$ is a tangle of G, then $\mathcal{T}' := \{\pi_{\mathcal{M}}(Y, S, Z) : (Y, S, Z) \in \mathcal{T}\}$ is called the *projection* of $\mathcal{T}$. Note that $\mathcal{T}'$ is not a tangle in general. Projecting is the converse operation of lifting in the sense that if $\mathcal{M}$ is faithful, and $\mathcal{T}$ and $\mathcal{T}'$ are tangles of G and H, then $\mathcal{T}$ is the lifting of $\mathcal{T}'$ if and only if $\mathcal{T}'$ is the projection of $\mathcal{T}$.

We now define a partial order $\leqslant_{\mathrm{G}}$ over the set of separations of a graph G by letting for every two separations $(Y, S, Z), (Y', S', Z')$, $(Y, S, Z) \leqslant_{\mathrm{G}} (Y', S', Z')$ if and only if $Y' \subsetneq Y$ or ($Y = Y'$ and $S \subseteq S'$). This order slightly differs from the more conventional order $\leqslant_{\mathrm{RS}}$ defined in Section 3, and used in the literature, and is defined in [Gro16, Subsection 3.2]. We will keep using the order $\leqslant_{\mathrm{G}}$ in this section (and only in this section), in order to stay consistent with Grohe's results.

Region and evasive tangles. A partially ordered set $(X, <)$ is said to be *well-founded* if every strictly decreasing sequence of elements of X is finite. In particular, if $(X, <)$ is well-founded then for every $x \in X$, there exists $y \in X$ which is minimal with respect to $<$ and such that $y \leqslant x$. In the remainder of the manuscript, whenever we consider a minimal separation or a well-founded family of separations, we always implicitly refer to the partial order $\leqslant_{\mathrm{G}}$ defined above.

We will distinguish two types of tangles in infinite graphs:

- the *region tangles*, defined as those which are well-founded (with respect to the order $\leqslant_{\mathrm{G}}$), and
- the *evasive tangles*, which contain some infinite decreasing sequence of separations (with respect to the order $\leqslant_{\mathrm{G}}$).

The tangles we consider in this work will always have order at most 4. Note that if G is 3-connected, an evasive tangle $\mathcal{T}$ of order 4 is exactly a tangle $\mathcal{T}_\omega^4$ induced by an end ω of degree 3. On the other hand, a region tangle is either a tangle of order 4 induced by some end ω of degree at least 4, or a tangle which is not induced by an end. For example, one can check that both graphs in Figure 1.1 have a unique tangle of order 4 which is the tangle induced by their unique end (which is thick), and this tangle is a region tangle in both cases. In particular, for the right graph of Figure 1.1, as it is 4-connected, this tangle is trivial, and only contains separations of the form $(\emptyset, S, V(G) \setminus S)$, for sets S of size at most 3.

Distinguishing tangles in a canonical way. We say that a separation (Y, S, Z) *distinguishes* two tangles $\mathcal{T}, \mathcal{T}'$ if $(Y, S, Z) \in \mathcal{T}$ and $(Z, S, Y) \in \mathcal{T}'$, or vice versa. We say that (Y, S, Z) distinguishes $\mathcal{T}$ and $\mathcal{T}'$ *efficiently* if there is no separation of smaller order distinguishing $\mathcal{T}$ and $\mathcal{T}'$. A tree-decomposition $(T, \mathcal{V})$ *distinguishes* a set of tangles $\mathcal{A}$ if for every two distinct tangles $\mathcal{T}, \mathcal{T}' \in \mathcal{A}$ there exists an edge-separation of $(T, \mathcal{V})$ distinguishing $\mathcal{T}$ and $\mathcal{T}'$. A separation is called *relevant* with respect to $\mathcal{A}$ if it distinguishes at least two tangles of $\mathcal{A}$. A tree-decomposition is *nice* (with respect to $\mathcal{A}$) if all its edge-separations are relevant (with respect to $\mathcal{A}$).

We will need the following result, which extends a central result from [RS91] in the finite case, and which is a canonical version of one of the main results of the grid-minor series in the locally finite case. We will only use it with $k = 4$, but we nevertheless state the result in its most general form.

Theorem 6.3 (Theorem 7.3 in [CHM22]). *Let $k \geqslant 1$ and let G be a locally finite graph. Then there exists a canonical tree-decomposition $(T, \mathcal{V})$ of G that efficiently distinguishes the set $\mathcal{A}_k$ of tangles of order at most k and that is nice with respect to $\mathcal{A}_k$.*

Remark 6.4. The fact that $(T, \mathcal{V})$ is nice in Theorem 6.3 is not explicit in the original statement, however it directly follows from the proof. Moreover, the proof also ensures that the edge-separations of $(T, \mathcal{V})$ are pairwise distinct.

6.2 An example

We give here an example of a one-ended graph that excludes some minor and has infinitely many region tangles of order 4. We show how to distinguish them on this example with a canonical tree-decomposition. As the application of Theorem 6.3 allowing to distinguish all tangles of order 4 will be the very first step of our proof of Theorems 7.1, 7.2 in the next subsection, this example may also be useful to have some intuition on it.

We consider the infinite graph G (a finite section of which is illustrated in Figure 1.4), which is obtained from the infinite triangular grid by adding in each triangular face $f = v_1v_2v_3$ three vertices $\{w_1, w_2, w_3\}$ inducing a $K_{3,3}$ with the vertices of the triangle, and another vertex z connected to each of the w_i's. G has two types of tangles of order 4: one is the tangle $\mathcal{T}_\omega^4$ induced by the unique end ω of G, and all the others are the tangles $\mathcal{T}_f^4$ pointing towards each face $f = v_1v_2v_3$ of the triangular grid; more precisely, $\mathcal{T}_f^4$ has the same set of separations as $\mathcal{T}_\omega^4$ except for $(V(G) \setminus A, \{v_1, v_2, v_3\}, \{w_1, w_2, w_3, z\}) \in \mathcal{T}_f^4$, where

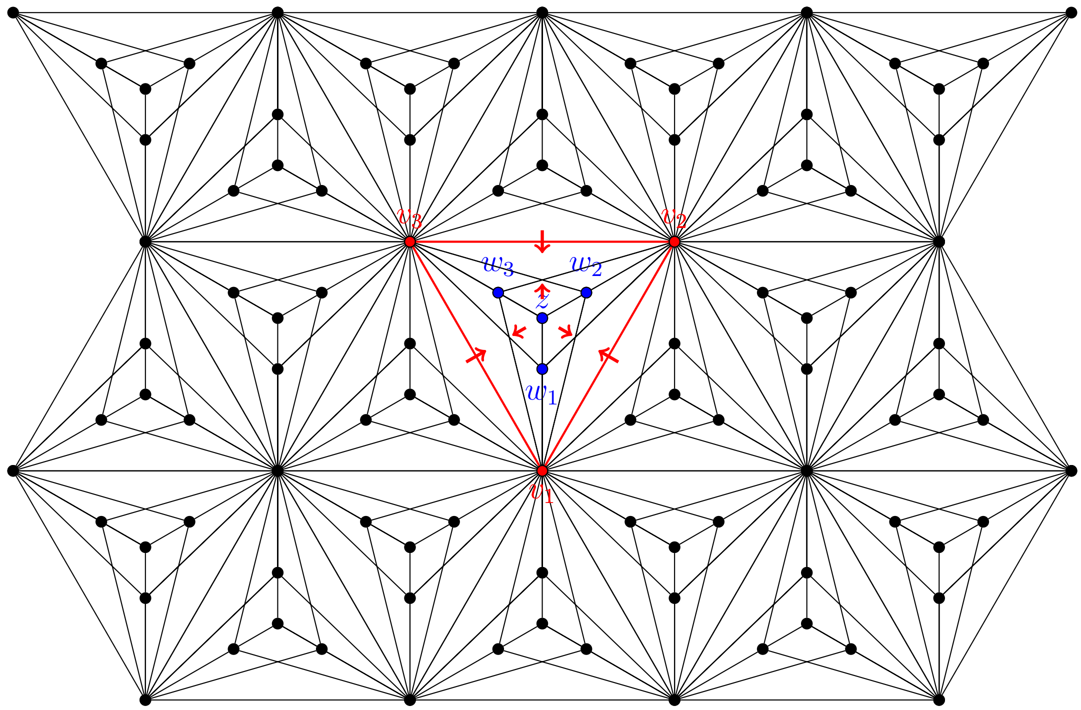

Figure 1.4: A useful example.

$A := \{v_1, v_2, v_3, w_1, w_2, w_3, z\}$. Note that with respect to our definition, all the tangles of order 4 of G are region tangles. We represented with red arrows the two separations of $\mathcal{T}_f^4$ that are minimal with respect to the order $\leqslant_{\mathrm{G}}$ but which are not minimal separations of $\mathcal{T}_\omega^k$ (for one fixed face f). The three red arrows crossing the red triangle correspond to the minimal separation $(V(G) \setminus A, \{v_1, v_2, v_3\}, \{w_1, w_2, w_3, z\})$ of $\mathcal{T}_f^4$ that points towards the triangular face $v_1v_2v_3$, while the three arrows directed away from z correspond to the minimal separation $(\{z\}, \{w_1, w_2, w_3\}, \{v_1, v_2, v_3\} \cup (V(G) \setminus A))$ of $\mathcal{T}_f^4$. The tree-decomposition $(T, \mathcal{V})$ of Theorem 6.3 distinguishing all the tangles of order 4 is such that T is a star with center $t_0 \in V(T)$ such that $G[\![V_{t_0}]\!]$ is the infinite planar triangular grid. Then T has one vertex t_f for each face $f = v_1v_2v_3$ of $G[\![V_{t_0}]\!]$ and the bag V_{t_f} is finite and contains the 7 vertices $\{v_1, v_2, v_3, w_1, w_2, w_3, z\}$ associated to f. Note that such a tree-decomposition enjoys the properties of Theorems 7.1 and 7.2. However, this is not always the case and we need in general to decompose further some torsos of the tree-decomposition given by Theorem 6.3 in order to obtain such a decomposition.

6.3 Tangles of order 4: orthogonality and crossing-lemma

Torsos of sets We give an alternate definition of torsos in a graph, which we will use in the remainder of this section (and only in this section). For any graph G and $X \subseteq V(G)$, we denote by $G[\![X]\!]$ the graph with vertex set X whose edge set consists of all the pairs uv such that $uv \in E(G)$ or there exists a connected component C of $G - X$ such that $\{u, v\} \subseteq N(C)$. In other words, u and v are adjacent in $G[\![X]\!]$ if and only if there exists some path from u

to v in G which intersects X only in its endpoints. The graphs $G[\![X]\!]$ are called the *torsos* of G. If $(T,\mathcal{V})$ is a tree-decomposition of G with tight edge-separations, then note that the torsos $G[\![V_t]\!]$ coincide exactly with the definition of torsos of a tree-decomposition we gave in Section 3. To prevent any ambiguity between the two definitions, we will ensure that in this section, the tree-decompositions we consider always have tight edge-separations.

Degenerate separations. In this section we introduce some notions from [Gro16] and briefly explain how to extend them in the locally finite case. Unless specified otherwise, we assume in the whole section that the graphs we consider are locally finite and 3-connected.

A separation $(Y,S,Z)\in \mathrm{Sep}_{<4}(G)$ is said to be *degenerate* if

- (Y,S,Z) has order 3,
- $G[S]$ is an independent set, and
- $|Y|=1$.

The following result from [Gro16] immediately generalizes to locally finite graphs.

Lemma 6.5 (Lemma 4.13 and Remark 4.14 in [Gro16])**.** *Let G be a locally finite 3-connected graph, and (Y,S,Z) be a proper separation of order 3. Then $G[\![Z\cup S]\!]$ is a faithful minor of G if and only if (Y,S,Z) is non-degenerate.*

We say that the edge-separations of a tree-decomposition $(T,\mathcal{V})$ are *non-degenerate* if for every $e\in E(T)$, none of the two edge-separations associated to e are degenerate.

Lemma 6.6. *Let G be a locally finite 3-connected graph and let $(T,\mathcal{V})$, with $\mathcal{V}=(V_t)_{t\in V(T)}$, be a tree-decomposition of G whose edge-separations have order 3 and are non-degenerate. Then $G[\![V_t]\!]$ is a faithful minor of G for each $t\in V(T)$.*

Proof. Let $t\in V(T)$, and t' be a neighbor of t in T. Let $(Y_{t'},S_{t'},Z_{t'})$ be the edge-separation of G associated to the (oriented) edge $(t',t)\in E(T)$, that is $S_{t'}=V_t\cap V_{t'}$, $V_{t'}\subseteq Y_{t'}\cup S_{t'}$, and $V_t\subseteq Z_{t'}\cup S_{t'}$. By Lemma 6.5, there is a faithful model $(M_v^{t'})_{v\in(Z_{t'}\cup S_{t'})}$ of $G[\![Z_{t'}\cup S_{t'}]\!]$ in G. As $(M_v^{t'})_{v\in(Z_{t'}\cup S_{t'})}$ is faithful and as $(Y_{t'},S_{t'},Z_{t'})$ is a separation, the only edges of $G[\![Z_{t'}\cup S_{t'}]\!]$ that are not edges of G must be between pairs of vertices of $S_{t'}$, so up to removing vertices from the sets $M_v^{t'}$, we may assume that every $M_v^{t'}$ has size 1, except possibly when $v\in S_{t'}$, in which case the only vertices distinct from v that $M_v^{t'}$ can have must lie in $Y_{t'}$. For every $v\in V_t$, we let:

$$M_v := \bigcup_{\substack{t'\in V(T),\\ tt'\in E(T)}} M_v^{t'}.$$

We show that $(M_v)_{v\in V_t}$ is a faithful model of $G[\![V_t]\!]$ in G. As $(T,\mathcal{V})$ is a tree-decomposition, for every two distinct neighbors t',t'' of t in T, $Y_{t'}\cap Y_{t''}=\emptyset$ so we must have $M_v^{t'}\cap M_v^{t''}=\{v\}$ and $M_v^{t'}\cap M_u^{t''}=\emptyset$ for each distinct vertices $u,v\in V_t$. As $(M_v^{t'})_{v\in(Z_{t'}\cup S_{t'})}$ is a model, we have $M_u^{t'}\cap M_v^{t'}=\emptyset$ for each $u\neq v\in V_t$. It follows that $M_u\cap M_v=\emptyset$ for each distinct $u,v\in V_t$. Now if $uv\in E(G[\![V_t]\!])$ and $uv\notin E(G)$, there must exist some edge-separation $(Y_{t'},S_{t'},Z_{t'})$ such that $u,v\in S_{t'}$ and there exists a path from u to v in $G[S_{t'}\cup Y_{t'}]$. In particular, there must exist $u' \in M_u^{t'}$ and $v' \in M_v^{t'}$ such that $u'v' \in E(G)$. As $u' \in M_u$ and $v' \in M_v$, we proved that $(M_v)_{v \in V_t}$ is a faithful model of $G[\![V_t]\!]$ in G. □

For every tangle $\mathcal{T}$ of a graph G, we denote by $\mathcal{T}_{\min}$ its set of minimal separations (here and in the remainder, minimality of separations is always with respect to the partial order $\leqslant_{\mathrm{G}}$ defined above). If $\mathcal{T}$ has order 4, then we let $\mathcal{T}_{\mathrm{nd}}$ be its set of non-degenerate minimal separations.

Remark 6.7. Let G be locally finite, let $\mathcal{T}$ be a G-tangle of order 4, and let (Y, S, Z) be a degenerate separation of G. Then $(Y, S, Z) \in \mathcal{T}$. This is a direct consequence of [Gro16, Lemma 3.3], which states that if $\mathcal{T}$ is a tangle of order k then for every separation (Y, S, Z) of order $k-1$ such that $|Y \cup S| \leqslant \frac{3}{2}(k-1)$ we have $(Y, S, Z) \in \mathcal{T}$.

For every tangle $\mathcal{T}$ of order 4, we let:

$$X_{\mathcal{T}} := \bigcap_{\substack{(Y,S,Z)\in\mathcal{T},\\ (Y,S,Z) \text{ is non-degenerate}}} (Z \cup S).$$

Note that if $\mathcal{T}$ is an evasive tangle, then $X_{\mathcal{T}}$ is empty. In this case, and because G is 3-connected there exists a unique end ω of degree 3 such that for any finite subset $\mathcal{S}$ of $\mathcal{T}$, the end ω lies in

$$\bigcap_{\substack{(Y,S,Z)\in\mathcal{S},\\ (Y,S,Z) \text{ is non-degenerate}}} (Z \cup S).$$

Remark 6.8. If $(Y, S, Z), (Y', S', Z') \in \mathcal{T}$ are such that $(Y', S', Z') \leqslant_{\mathrm{G}} (Y, S, Z)$ and (Y, S, Z) is non-degenerate, then it is easy to see that (Y', S', Z') is also non-degenerate (recall that G is 3-connected). Also if $\mathcal{T}$ is a region tangle, for every $(Y, S, Z) \in \mathcal{T}$, there exists a separation $(Y', S', Z') \in \mathcal{T}_{\min}$ such that $(Y', S', Z') \leqslant_{\mathrm{G}} (Y, S, Z)$. These observations imply that if $\mathcal{T}$ is a region tangle of order 4 and G is 3-connected and locally finite, then:

$$X_{\mathcal{T}} = \bigcap_{(Y,S,Z)\in\mathcal{T}_{\mathrm{nd}}} (Z \cup S).$$

Crossing and orthogonal separations. Two separations $(Y_1, S_1, Z_1), (Y_2, S_2, Z_2)$ are *orthogonal* if $(Y_1 \cup S_1) \cap (Y_2 \cup S_2) \subseteq S_1 \cap S_2$ (see Figure 1.5a). A set $\mathcal{N}$ of separations is said to be *orthogonal* if its separations are pairwise orthogonal. One can easily show that the set of minimal separations of a (region) tangle of order at most 3 is orthogonal. This does not hold for tangles of order 4, but Grohe [Gro16] proved that for tangles of order 4, minimal separations can only cross in a restricted way. Two separations (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) are *crossing* if $Y_1 \cap Y_2 = S_1 \cap S_2 = \emptyset$ and there is an edge $s_1 s_2 \in E(G)$, with $S_1 \cap Y_2 = \{s_1\}$ and $S_2 \cap Y_1 = \{s_2\}$ (see Figure 1.5b). In this case, we call $s_1 s_2$ the *crossedge* of (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2). We denote by $E^{\times}_{\mathrm{nd}}(\mathcal{T})$ the set of crossedges of $\mathcal{T}_{\mathrm{nd}}$. Lemma 4.16 from [Gro16] generalizes to region tangles of order 4 of locally finite graphs:

Lemma 6.9 (Lemma 4.16 and Corollary 4.20 in [Gro16])**.** *Let G be a locally finite 3-connected graph. Let $\mathcal{T}$ be a region G-tangle of order 4. Then every two distinct minimal separations of $\mathcal{T}$ are either crossing or orthogonal. Moreover, $E^{\times}_{\mathrm{nd}}(\mathcal{T})$ forms a matching in G.*

(a) Orthogonal separations (b) Crossing separations

Figure 1.5: Interaction between minimal separations. The white zones represent empty sets while the grey represent potentially non-empty sets.

In [Gro16], orthogonal sets of separations are presented as the nice case, as they allow to efficiently find quasi-4-connected regions. We show that, up to some additional assumptions, this observation still holds in the locally finite case. We recall that for a tangle $\mathcal{T}$ of order 4, $\mathcal{T}_{\mathrm{nd}}$ denotes its set of minimal non-degenerate separations.

Lemma 6.10. *Let G be a locally finite 3-connected graph. Let $\mathcal{T}$ be a region G-tangle of order 4. Assume that $\mathcal{T}_{\mathrm{nd}}$ is orthogonal. Then $X_{\mathcal{T}} \neq \emptyset$ and the torso $G[\![X_{\mathcal{T}}]\!]$ has size 3 or is a quasi-4-connected minor of G.*

Proof. If every separation of order 3 in G is degenerate, then G is quasi-4-connected and all the separations of $\mathcal{T}_{\mathrm{nd}}$ are non-proper. It follows that $G = G[\![X_{\mathcal{T}}]\!]$ and the desired properties hold.

Assume now that G has a proper non-degenerate separation (Y, S, Z) of order 3. As $\mathcal{T}$ is a region tangle, there is a separation $(Y', S', Z') \in \mathcal{T}_{\mathrm{min}}$ such that $(Y', S', Z') \leqslant_{\mathrm{G}} (Y, S, Z)$. As observed in Remark 6.8, $(Y', S', Z') \in \mathcal{T}_{\mathrm{nd}}$. We claim that $S' \subseteq X_{\mathcal{T}}$ so $X_{\mathcal{T}} \neq \emptyset$: let $(Y_0, S_0, Z_0) \in \mathcal{T}_{\mathrm{nd}} \setminus \{(Y', S', Z')\}$. As (Y_0, S_0, Z_0) and (Y', S', Z') are orthogonal, we must have: $S' \cap Y_0 = \emptyset$ so $S' \subseteq Z_0 \cup S_0$. As $\mathcal{T}$ is a region tangle, the equality $X_{\mathcal{T}} = \bigcap_{(Y,S,Z)\in\mathcal{T}_{\mathrm{nd}}}(Z \cup S)$ holds, and thus we proved that $S' \subseteq X_{\mathcal{T}}$, and so $X_{\mathcal{T}} \neq \emptyset$. Moreover, as G is 3-connected, the separations of $\mathcal{T}_{\mathrm{min}}$ have order 3 so $|X_{\mathcal{T}}| \geqslant |S'| \geqslant 3$.

We now assume that $|X_{\mathcal{T}}| \geqslant 4$ and show that $G[\![X_{\mathcal{T}}]\!]$ is quasi-4-connected. Since G is 3-connected and $|X_{\mathcal{T}}| \geqslant 4$, $G[\![X_{\mathcal{T}}]\!]$ is 3-connected (any proper separation of order at most 2 in $G[\![X_{\mathcal{T}}]\!]$ would induce a proper separation of order at most 2 in G). If $|X_{\mathcal{T}}| = 4$, then $G[\![X_{\mathcal{T}}]\!]$ is clearly also quasi-4-connected so we can assume that $|X_{\mathcal{T}}| \geqslant 5$. Suppose that $G[\![X_{\mathcal{T}}]\!]$ is not 4-connected and let (Y_0, S_0, Z_0) be a proper separation of $G[\![X_{\mathcal{T}}]\!]$ of order at most 3. We will prove that $|Y_0| = 1$ or $|Z_0| = 1$, which will immediately imply that $G[\![X_{\mathcal{T}}]\!]$ is quasi-4-connected. We let Y_1 be the union of all connected components of $G - S_0$ that intersect Y_0 or have a neighbor in Y_0. Let $S_1 := S_0$ and $Z_1 := V(G) \setminus (S_0 \cup Y_1)$. By definition of the torso $G[\![X_{\mathcal{T}}]\!]$, (Y_1, S_1, Z_1) is a proper separation of order at most 3 in G, hence we must have $|S_0| = 3$ as G is 3-connected. Assume first that $(Y_1, S_1, Z_1) \in \mathcal{T}$. If (Y_1, S_1, Z_1) is non-degenerate, then $X_{\mathcal{T}} \cap Y_0 \subseteq X_{\mathcal{T}} \cap Y_1 = \emptyset$ by definition of $X_{\mathcal{T}}$. It follows that $Y_0 = \emptyset$, which contradicts the assumption that (Y_0, S_0, Z_0) is proper. If (Y_1, S_1, Z_1) is degenerate, then $|Y_0| \leqslant |Y_1| = 1$ so as (Y_0, S_0, Z_0) is proper we must have $|Y_0| = 1$. The case $(Z_1, S_1, Z_1) \in \mathcal{T}$ is symmetric. Hence we proved that every separation (Y, S, Z) of $G[\![X_\mathcal{T}]\!]$ of order at most 3 satisfies $|Y| \leqslant 1$ or $|Z| \leqslant 1$ so we are done.

Finally the fact that $G[\![X_\mathcal{T}]\!]$ is a minor of G easily follows from Lemma 6.6: we consider the tree-decomposition $(T, \mathcal{V})$ where T is a star with a central vertex z_0 and one edge z_0z_i for each $(Y_i, S_i, Z_i) \in \mathcal{T}_{\rm nd}$. We let $V_{z_0} := X_\mathcal{T}$ and $V_{z_i} := Y_i \cup S_i$ for each $(Y_i, S_i, Z_i) \in \mathcal{T}_{\rm nd}$. The fact that $(T, \mathcal{V})$ is a tree-decomposition follows from the orthogonality of $\mathcal{T}_{\rm nd}$. Hence by Lemma 6.6, $G[\![X_\mathcal{T}]\!]$ is a minor of G. □

Whenever $\mathcal{T}_{\rm nd}$ is not orthogonal, Lemma 6.10 does not hold anymore and if we want to obtain a canonical tree-decomposition, we will need to consider a larger set, whose torso is not necessarily quasi-4-connected, but can be defined uniquely from the structural properties of $\mathcal{T}$, which will ensure that the resulting decomposition is canonical. For every region tangle $\mathcal{T}$ of order 4 we let:

$$R_\mathcal{T} := \left(\bigcup_{(Y,S,Z) \in \mathcal{T}_{\rm nd}} S \right) \cup \left(\bigcap_{(Y,S,Z) \in \mathcal{T}_{\rm nd}} Z \right).$$

The set $R_\mathcal{T}$ corresponds to the set called $R^{(0)}$ in [Gro16, Section 4.5]. Note that we always have $X_\mathcal{T} \subseteq R_\mathcal{T}$ and that equality holds when $\mathcal{T}_{\rm nd}$ is orthogonal. To illustrate the definition of $R_\mathcal{T}$, it is helpful to go back to the graph G on Figure 1.1 (left). Then G has a unique tangle $\mathcal{T}$ of order 4, the set $\mathcal{T}_{\rm nd}$ is the set of separations (Y, S, Z) of order 3 where Y is a triangular face of G, $S = N(Y)$ and $Z = V(G) \setminus (Y \cup S)$. Hence on this example, the set of crossedges $E^\times_{\rm nd}(G)$ is the set of edges joining two triangular faces, and $R_\mathcal{T} = V(G)$.

While the proof of the following result was originally written for finite graphs, it immediately generalizes to locally finite graphs.

Lemma 6.11 (Lemma 4.32 in [Gro16])**.** *If G is a locally finite 3-connected graph and if $\mathcal{T}$ is a region tangle of order 4 in G, then $G[\![R_\mathcal{T}]\!]$ is a faithful minor of G.*

For each $(Y, S, Z) \in \mathcal{T}_{\rm nd}$, the *fence* $\mathrm{fc}(S)$ of S in G is the union of

- the subset of vertices of S that are not the endpoint of some crossedge of $\mathcal{T}$, and
- the subset of vertices s' such that ss' is a crossedge of $\mathcal{T}$ and $s \in S$.

In particular, as the crossedges form a matching in G (Lemma 6.9), $|\mathrm{fc}(S)| = |S| = 3$ for each $(Y, S, Z) \in \mathcal{T}_{\rm nd}$. A consequence of Lemma 6.9 is the following:

Lemma 6.12. *Let G be a locally finite 3-connected graph and $\mathcal{T}$ be a region G-tangle of order 4. Then G has a tree-decomposition $(T, \mathcal{V})$ of adhesion 3 where $\mathcal{V} = (V_t)_{t \in V(T)}$ and T is a star with central vertex z_0 such that $V_{z_0} = R_\mathcal{T}$. If moreover, $\mathcal{T}$ is the unique G-tangle of order 4, then $(T, \mathcal{V})$ is canonical and every bag except possibly V_{z_0} is finite.*

Proof. We let:

$$V(T) := \{z_0\} \cup \{z_C : C \text{ connected component of } G - R_\mathcal{T}\}$$

where we choose the z_C's to be pairwise distinct nodes. We let T be the star with vertex set $V(T)$ and central vertex z_0, and we define $\mathcal{V} = (V_t)_{t\in V(T)}$ by setting $V_{z_0} := R_{\mathcal{T}}$, and for each connected component C of $G - R_{\mathcal{T}}$: $V_{z_C} := C \cup N(C)$. It is not hard to check that $(T, \mathcal{V})$ is a tree-decomposition of G. By [Gro16, Lemma 4.31] (whose proof extends to the locally finite case), for each component C of $G - R_{\mathcal{T}}$ there exists a unique separator S such that $(Y, S, Z) \in \mathcal{T}_{\mathrm{nd}}$ for some separation (Y, S, Z), $N(C) = \mathrm{fc}(S)$ and $C \subseteq Y$. This implies that $(T, \mathcal{V})$ has adhesion 3, so in particular its edge-separations are tight.

We now prove the second part of Lemma 6.12 and assume that $\mathcal{T}$ is the unique tangle of order 4 of G. Then $\mathcal{T}$ is $\mathrm{Aut}(G)$-invariant, and $(T, \mathcal{V})$ is clearly canonical. If some V_{z_C} is infinite for some $z_C \neq z_0$, then as G is locally finite, $G[V_{z_C}]$ has at least one infinite connected component, and hence there exists some end ω living in $G[V_{z_C}]$. In particular, ω induces some G-tangle $\mathcal{T}_\omega$ of order 4. We let $(Y, S, Z) \in \mathcal{T}_{\mathrm{nd}}$ be the separation given by [Gro16, Lemma 4.31] such that $N(C) = \mathrm{fc}(S)$ and $C \subseteq Y$. Then (Y, S, Z) distinguishes $\mathcal{T}_\omega$ from $\mathcal{T}$, which contradicts the uniqueness of $\mathcal{T}$ in G. □

Note that in the non-orthogonal case, the tree-decomposition from Lemma 6.12 is not the same as the one from [Gro16], as the torso $G[\![R_{\mathcal{T}}]\!]$ associated to the center of the star might not be quasi-4-connected. However, we will prove in Section 6.6 that $G[\![R_{\mathcal{T}}]\!]$ still enjoys the same useful properties as a quasi-transitive quasi-4-connected graph. We note that the crucial ingredient that allows us to obtain a canonical tree-decomposition in the second part of Lemma 6.12 (contrary to Grohe's decomposition) is the assumption that $\mathcal{T}$ is the unique G-tangle of order 4. In particular, one of the most important steps in the proof of the results from Section 7 will be a reduction to the case where graphs have a single tangle of order 4.

6.4 Contracting a single crossedge

In what follows, we let G be a locally finite 3-connected graph, and $\mathcal{T}$ be a region tangle of order 4 in G. Recall that by Lemma 6.9 the set $E^{\times}_{\mathrm{nd}}(\mathcal{T})$ of the crossedges forms a matching. We will see that contracting a crossedge results in a 3-connected graph G' that has a tangle $\mathcal{T}'$ of order 4 induced by $\mathcal{T}$ [Gro16, Subsection 4.5]. More precisely, $\mathcal{T}'$ contains as a subset the projection of $\mathcal{T}$ with respect to the minor G' of G. We give here an overview of the lemmas stated in [Gro16, Subsection 4.5] which all hold when G is locally finite instead of finite, by using the exact same proofs. The only additional property that we need in the locally finite case is that $\mathcal{T}'$ is still a region tangle, which is proved in Lemma 6.16 below.

In the remainder of this subsection, we let (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) be two crossing separations of $\mathcal{T}_{\mathrm{nd}}$ with crossedge s_1s_2. *Contracting* s_1s_2 consists in deleting s_1 and s_2 and adding a new vertex s' whose neighborhood is equal to $N_G(s_1) \cup N_G(s_2) \setminus \{s_1, s_2\}$. We denote by G' the graph obtained after contracting s_1s_2. The *projection* (referred to as *contraction* in [Gro16]) of a set X of vertices of G is defined as

$$X^{\vee} := \begin{cases} X & \text{if } X \cap \{s_1, s_2\} = \emptyset \\ X \setminus \{s_1, s_2\} \cup \{s'\} & \text{if } X \cap \{s_1, s_2\} \neq \emptyset. \end{cases}$$

Given a set X' of vertices of G', the *expansion* $X'_\wedge$ of X' is defined as

$$X'_\wedge := \begin{cases} X' & \text{if } s' \notin X' \\ X' \setminus \{s'\} \cup \{s_1, s_2\} & \text{if } s' \in X'. \end{cases}$$

Observe that for all $X' \subseteq V(G')$, we have $(X'_\wedge)^\vee = X'$ and for all $X \subseteq V(G)$, we have $X \subseteq (X^\vee)_\wedge$ (where the inclusion might be strict). We also define for every $(Y, S, Z) \in \mathrm{Sep}_{<4}(G)$:

$$(Y, S, Z)^\vee := \begin{cases} (Y^\vee \setminus \{s'\}, S^\vee, Z^\vee \setminus \{s'\}) & \text{if } S \cap \{s_1, s_2\} \neq \emptyset \\ (Y^\vee, S^\vee, Z^\vee) & \text{if } S \cap \{s_1, s_2\} = \emptyset. \end{cases}$$

Note that $(Y, S, Z)^\vee$ is exactly the projection $\pi_{\mathcal{M}}(Y, S, Z)$ with respect to the model $\mathcal{M} = (\{v\}_\wedge)_{v \in V(G')}$ of G' in G.

In the context of finite graphs, [Gro16] proves the following lemmas that extend directly to the locally finite case:

Lemma 6.13 (Corollary 4.24 in [Gro16])**.** *The graph G' resulting from the contraction of s_1s_2 is 3-connected.*

Lemma 6.14 (Lemmas 4.26 and 4.27 in [Gro16])**.** *There exists a tangle $\mathcal{T}'$ of order 4 in G' containing the projection of $\mathcal{T}$ with respect to the model $\mathcal{M} = (\{v\}_\wedge)_{v \in V(G')}$.*

Note that the projection of $\mathcal{T}$ with respect to $\mathcal{M}$ is exactly the set $\{(Y, S, Z)^\vee : (Y, S, Z) \in \mathcal{T}\}$. In the remainder of this subsection, we let $\mathcal{T}'$ be the tangle given by Lemma 6.14. In [Gro16], the author gives an explicit definition of $\mathcal{T}'$, but for the sake of clarity we only summarize here the properties of $\mathcal{T}'$ that will be of interest for our purposes.

Note that the inclusion $\{(Y, S, Z)^\vee : (Y, S, Z) \in \mathcal{T}\} \subseteq \mathcal{T}'$ is strict in general, as some separations from $\mathrm{Sep}_{<4}(G')$ might not be projections of separations from $\mathrm{Sep}_{<4}(G)$. The next lemma intuitively states that every separation of $\mathcal{T}'$ is close to an element from $\{(Y, S, Z)^\vee : (Y, S, Z) \in \mathcal{T}\} \subseteq \mathcal{T}'$.

Lemma 6.15 (Definition of $\mathcal{T}'$ and Lemmas 4.23 and 4.25 in [Gro16])**.** *For every separation $(Y', S', Z') \in \mathcal{T}'$ such that $s' \in S'$ and $G'[Z']$ is connected, there exists a separation $(Y, S, Z) \in \mathcal{T}$ such that $S^\vee = S'$ and $Z \setminus S_\vee = Z \setminus \{s_1, s_2\} = Z'$.*

As Lemma 6.15 is not exactly stated this way in [Gro16], we briefly sketch how to obtain it. If $(Y', S', Z') \in \mathcal{T}'$ is such that $s' \in S'$, then by [Gro16, Lemma 4.25] there exists a (unique) connected component C of $G \setminus S'_\vee$ such that the separations (Y'', S'', Z'') of $\mathcal{T}'$ such that $S'' = S'$ are exactly the ones such that $C \subseteq Z''$, and for which every separation $(Y, S, Z) \in \mathcal{T}$ such that $S^\vee = S'$ satisfies $C \subseteq Z$. [Gro16, Lemmas 4.23, 4.25] and the fact that $\mathcal{T}$ is a tangle ensure the existence of a separation $(Y, S, Z) \in \mathcal{T}$ such that $S^\vee = S'$ and $C = Z \setminus \{s_1, s_2\} = Z \setminus S$. In particular by Lemma 6.14 the projection $(Y, S, Z)^\vee = (Y \setminus S', S', C)$ is in $\mathcal{T}'$ so if we assume that $G'[Z']$ is connected, the choice of C imposes $C \subseteq Z'$, thus $Z' = C$. It implies that (Y, S, Z) satisfies the property described in Lemma 6.15.

Lemma 6.16. *$\mathcal{T}'$ is a region tangle.*

Proof. Assume for the sake of contradiction that $\mathcal{T}'$ contains an infinite strictly decreasing sequence of separations $(Y'_n, S'_n, Z'_n)_{n\in\mathbb{N}}$. By Corollary 6.13, G' is 3-connected, so the only possible non-proper separation (Y_n, S_n, Z_n) is for $n = 0$, thus we may assume that all the separations (Y_n, S_n, Z_n) are tight. By Lemma 3.1, there are finitely many integers n for which $s' \in S'_n$. Up to extracting an infinite subsequence, one can assume that for all n, either $s' \in Y'_n$ or $s' \in Z'_n$. If there exists N such that $s' \in Y'_N$ then for all $n \geqslant N$ we must have $s' \in Y_n$ by definition of $\leqslant_{\mathrm{G}}$. Up to extracting another infinite subsequence, we can assume that either $s' \in Y'_n$ for all n or $s' \in Z'_n$ for all n. As a result, and because $\mathcal{T}'$ contains the projection of $\mathcal{T}$ with respect to $\mathcal{M}$, $((Y'_n)_\wedge, S'_n, (Z'_n)_\wedge)_{n\in\mathbb{N}}$ is an infinite decreasing sequence of separations of order 3 in $\mathcal{T}$, contradicting the fact that $\mathcal{T}$ is well-founded. □

We conclude this subsection with the following result relating the degeneracy of minimal separations in G and G'. Its proof is the same as the proof of [Gro16, Lemma 4.28], which directly translates to the locally finite case. To be more precise, we also need the additional assumption that $\mathcal{T}'$ is a region tangle to make the proof work, which is given by Lemma 6.16.

Lemma 6.17 (Lemma 4.28 and Corollary 4.29 in [Gro16])**.** *Either G' is 4-connected and $\mathcal{T}'_{\min} = \{(\emptyset, \emptyset, V(G'))\}$, or*

$$\mathcal{T}'_{\min} = \{(Y,S,Z)^\vee : (Y,S,Z) \in \mathcal{T}_{\min} \text{ and } S^\vee \text{ is a separator of } G'\}.$$

In the latter case, for all $(Y,S,Z) \in \mathcal{T}_{\min}$, (Y,S,Z) is non-degenerate if and only if $(Y,S,Z)^\vee$ is non-degenerate. Moreover, $E^\times_{\mathrm{nd}}(\mathcal{T}') = E^\times_{\mathrm{nd}}(\mathcal{T}) \setminus \{s_1s_2\}$.

6.5 Contracting all the crossedges

In the previous subsection we studied the consequences of contracting a single crossedge in G. However, in our application we will need to contract *all* crossedges of $E^\times_{\mathrm{nd}}(\mathcal{T})$ (which form a matching in G). We now study how this affects G.

Before going further, we will need to introduce some notation, extending the notation from [Gro16] to the infinite case. For convenience we write $M := E^\times_{\mathrm{nd}}(\mathcal{T})$ (and recall that M is a matching in G). For every subset $L \subseteq M$ of crossedges, we let $G^{\backslash L/}$ be the graph obtained from G after contracting each edge $uv \in L$ into a new vertex $s_{u,v}$. Note that the order in which the edges are contracted is irrelevant in the definition of $G^{\backslash L/}$.

We denote $\overline{L} = M \setminus L$. In this section we will also often use the notation $L - L'$ instead of $L \setminus L'$, to avoid any possible confusion when reading superscripts (for instance we will write $G^{\backslash L-L'/}$ instead of $G^{\backslash L\setminus L'/}$).

For every $L \subseteq M$, for every vertex $x \in V(G)$, we let

$$x^{\backslash L/} := \begin{cases} x & \text{if } x \in X \setminus V(L) \\ s_{u,v} & \text{if } x \text{ is the endpoint of a crossedge } uv \in L. \end{cases}$$

For every subset $X \subseteq V(G)$ of vertices, we let $X^{\backslash L/} := \{x^{\backslash L/} : x \in X\}$ be the projection of X to $G^{\backslash L/}$.

Remark 6.18. Note that for every disjoint subsets $K, L \subset M$ and for all $X \in V(G)$, $X^{\backslash K \cup L/} = (X^{\backslash L/})^{\backslash K/} = (X^{\backslash K/})^{\backslash L/}$.

For every $X' \subseteq V(G^{\backslash M/})$ and $L \subseteq M$, let $X'_{/L\backslash}$ denote the maximal set $X \subseteq V(G^{\backslash \overline{L}/})$ such that $X^{\backslash L/} = X'$. In other words $X'_{/L\backslash}$ is the set of vertices obtained after "uncontracting" the edges of L in X. Note that with the notation introduced above we have $G = G^{\backslash \emptyset/}$. Given a separation (Y, S, Z) of G, we define

$$(Y, S, Z)^{\backslash L/} := \left(Y^{\backslash L/} \setminus S^{\backslash L/}, S^{\backslash L/}, Z^{\backslash L/} \setminus S^{\backslash L/}\right).$$

Note that when $L = \{s_1 s_2\}$ consists of a single edge, we recover the notions of the previous subsection; with our previous notation this gives: $x^{\backslash L/} = x^{\vee}$, $X^{\backslash L/} = X^{\vee}$ and $(Y, S, Z)^{\backslash L/} = (Y, S, Z)^{\vee}$.

For each finite subset of crossedges $L \subseteq M$, and every enumeration $(e_1, \dots, e_\ell)$ of the edges of L, we let $\mathcal{T}^{\backslash (e_1, \dots, e_\ell)/}$ denote the tangle of $G^{\backslash L/}$ obtained after iteratively applying Lemma 6.14 to the graphs $G_0 := G, G_1, \dots, G_\ell$ with $G_i := G^{\backslash \{e_1, \dots, e_i\}/}$ for each $i \in [\ell]$.

Lemma 6.19 (Lemma 4.30 (5) in [Gro16])**.** *For every enumeration $(e_1, \dots, e_\ell)$ of a finite set $L \subseteq M$ of crossedges and every permutation σ of $[\ell]$, $\mathcal{T}^{\backslash (e_1, \dots, e_\ell)/} = \mathcal{T}^{\backslash (e_{\sigma(1)}, \dots, e_{\sigma(\ell)})/}$.*

In the remainder of the subsection, for every finite subset $L \subseteq M$, we will denote with $\mathcal{T}^{\backslash L/}$ the unique tangle associated to any enumeration of L given by Lemma 6.19.

Intuitively when G is finite, one of the main properties of $\mathcal{T}^{\backslash L/}$ is that separations of $\mathcal{T}_{\mathrm{nd}}^{\backslash L/}$ are in correspondence with separations of $\mathcal{T}_{\mathrm{nd}}$, and that the only crossing pairs between elements of $\mathcal{T}^{\backslash L/}$ correspond to pairs which were already crossing in $\mathcal{T}$. Thus after each contraction, we reduce the number of crossedges, hence when $L = M$, the family $\mathcal{T}_{\mathrm{nd}}^{\backslash L/}$ must be orthogonal and we can apply results from the previous subsections to the graph $G^{\backslash L/}$. We now show formally how to extend the relevant proofs of [Gro16] to the locally finite case.

In [Gro16], the author proved that if G is finite and 3-connected, for every $L \subseteq M$, the graph $G^{\backslash L/}$ is 3-connected and that $\mathcal{T}^{\backslash L/}$ is a tangle of order 4 induced by $\mathcal{T}$ in $G^{\backslash L/}$. Using the results from the previous subsection, this immediately extends to $G^{\backslash L/}$ and $\mathcal{T}^{\backslash L/}$ when G is locally finite and L is finite, by induction on the size of L.

Theorem 6.20 (Generalization of Lemma 4.30 in [Gro16])**.** *Let $L \subseteq M$ be a finite set of crossedges. Then we have*

1. *$G^{\backslash L/}$ is 3-connected.*

2. *$\mathcal{T}^{\backslash L/}$ is a region tangle of order 4 of $G^{\backslash L/}$ such that*

$$\mathcal{T}_{\min}^{\backslash L/} = \{(Y, S, Z)^{\backslash L/} : (Y, S, Z) \in \mathcal{T}_{\min} \text{ such that } S^{\backslash L/} \text{ is a separator of } G^{\backslash L/}\}$$

   or

$$\mathcal{T}_{\min}^{\backslash L/} = \{(\emptyset, \emptyset, V(G'))\}$$

   if $L = M$ is finite and $G^{\backslash L/}$ is 4-connected.

3. *$E_{\mathrm{nd}}^{\times}(\mathcal{T}^{\backslash L/}) = E_{\mathrm{nd}}^{\times}(\mathcal{T}) \setminus L$.*

4. *$\mathcal{T}^{\setminus L/}$ contains the projection $\{(Y,S,Z)^{\setminus L/} : (Y,S,Z) \in \mathcal{T}\}$ of $\mathcal{T}$ with respect to the model $\mathcal{M} = (\{v\}_{/L\setminus})_{v\in V(G^{\setminus L/})}$.*

We will now extend Theorem 6.20 to the case where $L \subseteq M$ is infinite. Given a set $X \subseteq V(G^{\setminus M/})$, we denote by $M(X) \subseteq M$ the subset of edges of G contracted to a vertex in X.

Lemma 6.21. *The graph $G^{\setminus M/}$ is 3-connected.*

Proof. Assume for the sake of contradiction that $G^{\setminus M/}$ has a separator S of order at most 2. Then the set $L := M(S)$ has size at most 2 and S is a separator of order 2 of $G^{\setminus L/}$. This contradicts Theorem 6.20. □

We now let $L \subseteq M$ be any (not necessarily finite) subset of crossedges and give a general definition of $\mathcal{T}^{\setminus L/}$ extending the previous one. For every $(Y', S', Z') \in \mathrm{Sep}_{<4}(G^{\setminus L/})$, we let $L' := M(S')$. Note that L' is finite, and that $(Y'_{/\overline{L}\setminus}, S', Z'_{/\overline{L}\setminus})$ is a separation of order at most 3 in $G^{\setminus L'/}$. We define $\mathcal{T}^{\setminus L/}$ as the family of separations (Y', S', Z') of $G^{\setminus L/}$ such that $(Y'_{/\overline{L}\setminus}, S', Z'_{/\overline{L}\setminus}) \in \mathcal{T}^{\setminus L'/}$. Note that when L is finite, $(Y', S', Z') = (Y'_{/\overline{L}\setminus}, S', Z'_{/\overline{L}\setminus})^{\setminus L-L'/}$ and thus iterative applications of Lemma 6.14 together with Lemma 6.19 imply that our definition of $\mathcal{T}^{\setminus L/}$ coincides with the one we gave above for finite subsets $L \subseteq M$.

Thanks to Remark 6.18, for all $L \subseteq M$, $\mathcal{T}^{\setminus M/} = (\mathcal{T}^{\setminus L/})^{\setminus \overline{L}/}$ and $G^{\setminus M/} = (G^{\setminus L/})^{\setminus \overline{L}/}$. We say that a set $X \subseteq V(G)$ *hits the edges of L once* if for all $e \in L$, $|X \cap e| = 1$.

Lemma 6.22. *$\mathcal{T}^{\setminus M/}$ is a region tangle of order 4 in $G^{\setminus M/}$ such that $\{(Y,S,Z)^{\setminus M/} : (Y,S,Z) \in \mathcal{T}\} \subseteq \mathcal{T}^{\setminus M/}$.*

Proof. We first prove that $\mathcal{T}^{\setminus M/}$ is a tangle of order 4. To prove ($\mathcal{T}1$), let $(Y', S', Z') \in \mathrm{Sep}_{<4}(G^{\setminus M/})$ and $L := M(S')$. Then L has size at most 3, so by Theorem 6.20, $\mathcal{T}^{\setminus L/}$ is a region tangle of order 4 of $G^{\setminus L/}$. As $(Y,S,Z) := (Y'_{/\overline{L}\setminus}, S', Z'_{/\overline{L}\setminus})$ is a separation of order 3 of $G^{\setminus L/}$ and $\mathcal{T}^{\setminus L/}$ is a tangle of order 4, either $(Y,S,Z) \in \mathcal{T}^{\setminus L/}$ or $(Z,S,Y) \in \mathcal{T}^{\setminus L/}$. By definition of $\mathcal{T}^{\setminus M/}$ we then have either $(Z', S', Y') \in \mathcal{T}^{\setminus M/}$ or $(Y', S', Z') \in \mathcal{T}^{\setminus M/}$, implying that $\mathcal{T}^{\setminus M/}$ satisfies ($\mathcal{T}1$).

To prove ($\mathcal{T}2$), let $(Y'_1, S'_1, Z'_1), (Y'_2, S'_2, Z'_2), (Y'_3, S'_3, Z'_3) \in \mathcal{T}^{\setminus M/}$. Let $L := M(S'_1 \cup S'_2 \cup S'_3)$. Once again L is finite with size at most 9 and for all $i \in \{1,2,3\}$, $(Y_i, S_i, Z_i) := ((Y'_i)_{/\overline{L}\setminus}, S'_i, (Z'_i)_{/\overline{L}\setminus})$ is a separation of order 3 of $G^{\setminus L/}$.

Claim 6.23. For every $i \in \{1,2,3\}$, $(Y_i, S_i, Z_i) \in \mathcal{T}^{\setminus L/}$.

Proof of the Claim: Assume that $i = 1$, the other cases being symmetric. We let $L_1 := M(S'_1)$. Then by definition of $\mathcal{T}^{\setminus M/}$, $(Y''_1, S''_1, Z''_1) := ((Y'_1)_{/\overline{L_1}\setminus}, S'_1, (Z'_1)_{/\overline{L_1}\setminus}) \in \mathcal{T}^{\setminus L_1/}$. Our goal is to show that $(Y_1, S_1, Z_1) = (Y''_1, S''_1, Z''_1)^{\setminus L-L_1/}$. As both L and L_1 are finite and $L_1 \subseteq L$, iterative applications of Lemmas 6.14 and Lemma 6.19 imply $\mathcal{T}^{\setminus L/}$ must contain the projection of $\mathcal{T}^{\setminus L_1/}$ with respect to the model $\mathcal{M} = (\{v\}_{/(L\setminus L_1)\setminus})_{v\in V(G)^{\setminus L/}}$. Thus if we succeed to prove that

$$(Y_1, S_1, Z_1) = (Y''_1, S''_1, Z''_1)^{\setminus L-L_1/}, \tag{1.1}$$

we immediately obtain that $(Y_1, S_1, Z_1) \in \mathcal{T}^{\setminus L/}$, which concludes the claim.

To prove that (1.1) holds, note first that every edge of M is contracted in $G^{\backslash M/}$ so in particular it has its endpoints in exactly one of the three sets $(Y_1')_{/M\backslash}, (S_1')_{/M\backslash}$ and $(Z_1')_{/M\backslash}$. In particular by definition of L_1, the edges of L_1 are all disjoint from $(Y_1')_{/M\backslash}$ and thus $(Y_1')_{/M\backslash} = (Y_1')_{/\overline{L_1}\backslash}$. This implies that $Y_1 = (Y_1')_{/\overline{L}\backslash} = ((Y_1')_{/\overline{L_1}\backslash})^{\backslash L-L_1/} = (Y_1'')^{\backslash L-L_1/}$. As S_1' is disjoint from Y_1' in $G^{\backslash M/}$, it is also disjoint from Y_1 in $G^{\backslash L/}$ so we have $Y_1 = (Y_1'')^{\backslash (L\backslash L_1)/} \backslash S_1''$. Symmetric arguments give $Z_1 = (Z_1'')^{\backslash L-L_1/} \setminus S_1''$, and as $S_1'' = S_1' = S_1$, we get the desired equality. $\Diamond$

By Theorem 6.20, $\mathcal{T}^{\backslash L/}$ is a region tangle of order 4 so Claim 6.23 implies that either $Z_1 \cap Z_2 \cap Z_3 \neq \emptyset$ or there exists an edge of $G^{\backslash L/}$ with both endpoints in $Z_1 \cup Z_2 \cup Z_3$. If $Z_1 \cap Z_2 \cap Z_3 \neq \emptyset$, then $Z_1' \cap Z_2' \cap Z_3' = (Z_1 \cap Z_2 \cap Z_3)^{\backslash \overline{L}/} \neq \emptyset$. Otherwise, there is an edge $e \in E(G^{\backslash L/})$ that has an endpoint in each Z_i, in which case, either $Z_1' \cap Z_2' \cap Z_3' \neq \emptyset$ if $e \in \overline{L}$, or $e^{\backslash \overline{L}/}$ is an edge of $G^{\backslash M/}$ which has an endpoint in each Z_i'. This proves ($\mathcal{T}$2) and shows that $\mathcal{T}^{\backslash M/}$ is a tangle of order 4.

We now prove the inclusion $\{(Y, S, Z)^{\backslash M/} : (Y, S, Z) \in \mathcal{T}\} \subseteq \mathcal{T}^{\backslash M/}$. Let $(Y, S, Z) \in \mathcal{T}$ and $L := M(S)$. By Theorem 6.20 (4), $(Y, S, Z)^{\backslash L/} \in \mathcal{T}^{\backslash L/}$. Write $(Y', S', Z') = (Y, S, Z)^{\backslash M/}$ and note that $(Y'_{/\overline{L}\backslash}, S', Z'_{/\overline{L}\backslash}) = (Y, S, Z)^{\backslash L/}$, thus by definition of $\mathcal{T}^{\backslash M/}$, $(Y, S, Z)^{\backslash M/} \in \mathcal{T}^{\backslash M/}$.

We now prove that $\mathcal{T}^{\backslash M/}$ is a well-founded set. For the sake of contradiction, let $((Y_n', S_n', Z_n'))_{n\in\mathbb{N}}$ be an infinite decreasing sequence of separations of $\mathcal{T}^{\backslash M/}$. The contradiction will follow from the next claim

Claim 6.24. There exists an infinite decreasing sequence $((Y_n'', S_n'', Z_n''))_{n\in\mathbb{N}}$ in $\mathcal{T}^{\backslash M/}$ such that for each $n \geqslant 0$, $(Y_n'', S_n'', Z_n'') = (Y_n, S_n, Z_n)^{\backslash M/}$ for some $(Y_n, S_n, Z_n) \in \mathcal{T}$.

Proof of the Claim: By Lemma 3.1 and because $G^{\backslash M/}$ is 3-connected, for any $n \in \mathbb{N}$, there are finitely many separations (Y_m', S_m', Z_m') such that $S_n' \cap S_m' \neq \emptyset$. Therefore, up to considering a subsequence, we can assume that for all n, $S_n' \subseteq Y_{n+1}'$ and $S_{n+1}' \subseteq Z_n'$. In particular, as by Lemma 6.21 $G^{\backslash M/}$ is 3-connected, S_{n+1}' is included in some connected component C_n of $G^{\backslash M/}[Z_n']$. Then we have

$$(Y_{n+1}', S_{n+1}', Z_{n+1}') \leqslant_{\mathrm{G}} (V(G^{\backslash M/}) \setminus (S_n' \cup C_n), S_n', C_n) \leqslant_{\mathrm{G}} (Y_n', S_n', Z_n'),$$

implying that $(V(G^{\backslash M/}) \setminus (S_n' \cup C_n), S_n', C_n) \in \mathcal{T}^{\backslash M/}$. Hence we may also assume up to replacing (Y_n', S_n', Z_n') with $(V(G^{\backslash M/}) \setminus (S_n' \cup C_n), S_n', C_n)$ that for each $n \geqslant 0$, $G^{\backslash M/}[Z_n']$ is connected.

For each $n \geqslant 0$, we let $L_n := M(S_n')$. Connectedness of $G^{\backslash M/}[Z_n']$ then implies that $G^{\backslash L_n/}[(Z_n')_{/\overline{L_n}\backslash}]$ is connected. Observe that $|L_n| \leqslant 3$ successive applications of Lemma 6.15 imply that there exists some separation $(Y_n, S_n, Z_n) \in \mathcal{T}$ such that $S_n^{\backslash L_n/} = S_n'$ and $Z_n \setminus ((S_n')_{/L_n\backslash}) = (Z_n')_{/\overline{L_n}\backslash}$. We let $(Y_n'', S_n'', Z_n'') := (Y_n, S_n, Z_n)^{\backslash M/}$. By Theorem 6.20 (4), $(Y_n, S_n, Z_n)^{\backslash L_n/} \in \mathcal{T}^{\backslash L_n/}$. Moreover,

$$(Y_n, S_n, Z_n)^{\backslash L_n/} = ((Y_n'')_{/\overline{L_n}\backslash}, S_n'', (Z_n'')_{/\overline{L_n}\backslash}),$$

thus by definition of $\mathcal{T}^{\backslash M/}$, we have $(Y_n'', S_n'', Z_n'') = ((Y_n, S_n, Z_n)^{\backslash L_n/})^{\backslash \overline{L_n}/} \in \mathcal{T}^{\backslash M/}$.

Note that $S''_n = S_n^{\backslash M/} = S_n^{\backslash L_n/} = S'_n$ and

$$Z''_n = Z_n^{\backslash M/} \setminus S'_n = (Z_n \setminus ((S'_n)_{/L_n\backslash}))^{\backslash \overline{L_n}/} = ((Z'_n)_{/\overline{L_n}\backslash})^{\backslash \overline{L_n}/} = Z'_n.$$

We thus deduce that $(Y'_n, S'_n, Z'_n) = (Y''_n, S''_n, Z''_n)$, so $((Y'_n, S'_n, Z'_n))_{n\in\mathbb{N}}$ is an infinite decreasing sequence in $\mathcal{T}^{\backslash M/}$ satisfying the desired properties. ◊

It remains to show how to derive a contradiction from Claim 6.24. For this let $((Y''_n, S''_n, Z''_n))_{n\in\mathbb{N}}$ and $((Y_n, S_n, Z_n))_{n\in\mathbb{N}}$ be as in Claim 6.24. Again, up to considering a subsequence, we can assume that for all n, $S''_n \subsetneq Y''_{n+1}$ and $S''_{n+1} \subsetneq Z''_n$. As $S''_n \cap S''_{n+1} = \emptyset$, the separators S_n and S_{n+1} cannot contain two vertices of a common crossedge of M. Thus, $(S''_n)_{/M\backslash} \subsetneq Y_{n+1}$ and $(S''_{n+1})_{/M\backslash} \subsetneq Z_n$ and hence $(Y_{n+1}, S_{n+1}, Z_{n+1}) <_{\mathrm{G}} (Y_n, S_n, Z_n)$. This proves that $((Y_n, S_n, Z_n))_{n\in\mathbb{N}}$ is an infinite decreasing sequence of separations of G with respect to $\mathcal{T}$, contradicting the fact that $\mathcal{T}$ is a region tangle. □

Lemma 6.25. *Either $G^{\backslash M/}$ is 4-connected and $\mathcal{T}_{\min}^{\backslash M/} = \{(\emptyset, \emptyset, V(G)^{\backslash M/})\}$ or*

$$\mathcal{T}_{\min}^{\backslash M/} = \{(Y, S, Z)^{\backslash M/} : (Y, S, Z) \in \mathcal{T}_{\min} \text{ such that } S^{\backslash M/} \text{ is a separator of } G^{\backslash M/}\}.$$

Finally, we have $E^{\times}_{\mathrm{nd}}(G^{\backslash M/}) = \emptyset$.

Proof. Assume that $G^{\backslash M/}$ is not 4-connected. We first prove the direct inclusion. Let $(Y', S', Z') \in \mathcal{T}_{\min}^{\backslash M/}$, $L := M(S')$ and $(Y_0, S_0, Z_0) := (Y'_{/\overline{L}\backslash}, S', Z'_{/\overline{L}\backslash})$. Then by definition of $\mathcal{T}^{\backslash M/}$, $(Y_0, S_0, Z_0) \in \mathcal{T}^{\backslash L/}$. We prove that (Y_0, S_0, Z_0) is a minimal element of $\mathcal{T}^{\backslash L/}$. Assume for a contradiction that there exists $(Y_1, S_1, Z_1) <_{\mathrm{G}} (Y_0, S_0, Z_0)$ in $\mathcal{T}^{\backslash L/}$. Then $(Y_1, S_1, Z_1)^{\backslash \overline{L}/} \leqslant_{\mathrm{G}} (Y_0, S_0, Z_0)^{\backslash \overline{L}/} = (Y', S', Z')$ and $(Y_1, S_1, Z_1)^{\backslash \overline{L}/} \neq (Y', S', Z')$ because $\overline{L} \cap S' = \emptyset$. Moreover, by Theorem 6.20 (2), we may assume that (Y'_1, S'_1, Z'_1) is minimal and that $(Y_1, S_1, Z_1) = (Y'_1, S'_1, Z'_1)^{\backslash L/}$ for some $(Y'_1, S'_1, Z'_1) \in \mathcal{T}$. Thus as $(Y_1, S_1, Z_1)^{\backslash \overline{L}/} = ((Y'_1, S'_1, Z'_1)^{\backslash L/})^{\backslash \overline{L}/} = (Y'_1, S'_1, Z'_1)^{\backslash M/}$, we must have $(Y_1, S_1, Z_1) \in \mathcal{T}^{\backslash M/}$, contradicting the minimality of (Y_0, S_0, Z_0) in $\mathcal{T}^{\backslash L/}$. Hence $(Y', S', Z') = (Y_0, S_0, Z_0)^{\backslash \overline{L}/}$ for some $(Y_0, S_0, Z_0) \in \mathcal{T}_{\min}^{\backslash L/}$. Again we can apply Theorem 6.20 and write $(Y_0, S_0, Z_0) = (Y, S, Z)^{\backslash L/}$ for some $(Y, S, Z) \in \mathcal{T}_{\min}$. Thus we have:

$$(Y', S', Z') = ((Y, S, Z)^{\backslash L/})^{\backslash \overline{L}/} = (Y, S, Z)^{\backslash M/},$$

so we are done with the direct inclusion.

Conversely, let $(Y_1, S_1, Z_1) \in \mathcal{T}_{\min}$ such that $S_1^{\backslash M/}$ is a separator of $G^{\backslash M/}$. Note that by the previous inclusion and because $\mathcal{T}^{\backslash M/}$ is a region tangle, it is enough to prove that for any $(Y_2, S_2, Z_2) \in \mathcal{T}$ such that $(Y_2, S_2, Z_2)^{\backslash M/} \leqslant_{\mathrm{G}} (Y_1, S_1, Z_1)^{\backslash M/}$, we have $(Y_2, S_2, Z_2)^{\backslash M/} = (Y_1, S_1, Z_1)^{\backslash M/}$. Let $(Y_2, S_2, Z_2) \in \mathcal{T}$ such that $(Y_2, S_2, Z_2)^{\backslash M/} \leqslant_{\mathrm{G}} (Y_1, S_1, Z_1)^{\backslash M/}$, $L := M(S_1 \cup S_2)$ and $\overline{L} := M \setminus L$. For $i \in \{1, 2\}$, the edges in $\overline{L}$ are either contained in Y_i or in Z_i. Note that for each i and $L \subseteq M$:

$$S_i^{\backslash L/} \cup (Z_i^{\backslash L/} \setminus S_i^{\backslash L/}) = (Z_i \cup S_i)^{\backslash L/}$$

and

$$S_i^{\backslash L/} \cup (Y_i^{\backslash L/} \setminus S_i^{\backslash L/}) = (Y_i \cup S_i)^{\backslash L/}.$$

Thus as $(Y_1, S_1, Z_1)^{\backslash M/} \leqslant_{\mathrm{G}} (Y_2, S_2, Z_2)^{\backslash M/}$, we have $(S_1 \cup Z_1)^{\backslash M/} \subseteq (S_2 \cup Z_2)^{\backslash M/}$. By the previous remark that no edge of $\overline{L}$ is contained in Z_i, this implies that $(S_1 \cup Z_1)^{\backslash L/} \subseteq (S_2 \cup Z_2)^{\backslash L/}$. Likewise $(S_1 \cup Y_1)^{\backslash L/} \subseteq (S_2 \cup Y_2)^{\backslash L/}$. As a result, $(Y_1, S_1, Z_1)^{\backslash L/} \preceq (Y_2, S_2, Z_2)^{\backslash L/}$. By Theorem 6.20 (2), $(Y_1, S_1, Z_1)^{\backslash L/} \in \mathcal{T}_{\min}^{\backslash L/}$, thus we have $(Y_1, S_1, Z_1)^{\backslash L/} = (Y_2, S_2, Z_2)^{\backslash L/}$ and $(Y_1, S_1, Z_1)^{\backslash M/} = (Y_2, S_2, Z_2)^{\backslash M/}$, showing that $(Y_1, S_1, Z_1)^{\backslash M/} \in \mathcal{T}_{\min}^{\backslash M/}$.

We now prove that $E^{\times}_{\mathrm{nd}}(G^{\backslash M/}) = \emptyset$. Assume that there are two crossing non-degenerate minimal 3-separations (Y_1'', S_1'', Z_1'') and (Y_2'', S_2'', Z_2'') in $\mathcal{T}^{\backslash M/}$, let $L = M(S_1'' \cup S_2'')$. For $i \in \{1, 2\}$, all crossedges of $G^{\backslash L/}$ lie in Y_i'' or in Z_i'', hence $(Y_i', S_i', Z_i') = ((Y_i'')_{/\overline{L}\backslash}, S_i'', (Z_i'')_{/\overline{L}\backslash})$ is the only 3-separation of $\mathrm{Sep}_{<4}(G^{\backslash L/})$ such that $(Y_i', S_i', Z_i')^{\backslash \overline{L}/} = (Y_i'', S_i'', Z_i'')$. Since $(Y_i'', S_i'', Z_i'') \in \mathcal{T}_{\min}^{\backslash M/}$, from what we just proved, we must have $(Y_i', S_i', Z_i') \in \mathcal{T}_{\min}^{\backslash M/}$. Note that the separations (Y_i', S_i', Z_i') are non-degenerate in $G^{\backslash L/}$. Furthermore, as $L = M(S_1'' \cup S_2'')$, (Y_1', S_1', Z_1') and (Y_2', S_2', Z_2') must be also crossing in $G^{\backslash L/}$, but this contradicts $E^{\times}_{\mathrm{nd}}(T^{\backslash L/}) = E^{\times}_{\mathrm{nd}}(\mathcal{T}) \setminus L$ (third item of Theorem 6.20). □

For each $L \subseteq M$, we let $R^{\backslash L/} := R_{\mathcal{T}}^{\backslash L/}$. Note that for each $L \subseteq M$, $(R_{\mathcal{T}^{\backslash L/}})_{/L\backslash} = R_{\mathcal{T}}$. Thus together with Lemma 6.10, this immediately gives the following, which is the locally finite extension of one of the main results from [Gro16]:

Theorem 6.26. *Let G be a locally finite 3-connected graph, and let $\mathcal{T}$ be a region tangle of order 4 in G. Let $M := E^{\times}_{\mathrm{nd}}(\mathcal{T})$ be the set of crossedges between non-degenerate minimal separations of $\mathcal{T}$. Then the graph $G^{\backslash M/}[\![R^{\backslash M/}]\!]$ is a quasi-4-connected minor of G.*

In order to obtain a proof of Theorem 3.10 in the locally finite case, one can either reuse the arguments from [Gro16, Section 5] (without canonicity), or equivalently adapt our proof from Section 7.4, which basically consists in applying first Theorem 6.3 and then refining the obtained tree-decomposition by decomposing further any infinite torso using Theorem 6.26.

6.6 Planarity after uncontracting crossedges

In general, if $G^{\backslash M/}[\![R^{\backslash M/}]\!]$ is quasi-4-connected, then $G[\![R_{\mathcal{T}}]\!]$ may not be quasi-4-connected anymore. To circumvent this and find a quasi-4-connected region in G, it is proved in [Gro16] that for every subset X' of $R_{\mathcal{T}}$ obtained by deleting one endpoint of each edge of M, the graph $G[\![X']\!]$ is isomorphic to $G^{\backslash M/}[\![R^{\backslash M/}]\!]$. However we cannot choose such a subset X' canonically in general, as illustrated in Example 3.11. Despite the fact that uncontracting the edges of M does not preserve the quasi-4-connectivity of torsos, we now prove that at least planarity is preserved by this operation.

Proposition 6.27. *If $G^{\backslash M/}[\![R^{\backslash M/}]\!]$ is planar, then so is $G[\![R_{\mathcal{T}}]\!]$.*

This is obtained by combining the following two lemmas:

Lemma 6.28. *For every subset $L \subseteq M$, we denote $\overline{L} := M \setminus L$. Assume that for every finite subset $L \subseteq M$, $G^{\backslash \overline{L}/}[\![R_{\mathcal{T}}^{\backslash \overline{L}/}]\!]$ is planar. Then $G[\![R_{\mathcal{T}}]\!]$ is also planar.*

Proof. Assume for the sake of contradiction that G enjoys the properties described above but that $G[\![R_{\mathcal{T}}]\!]$ is not planar. Then by Wagner's theorem [Wag37], $G[\![R_{\mathcal{T}}]\!]$ admits F as a minor, for some $F \in \{K_5, K_{3,3}\}$. Note that we can find a model $(V_v)_{v\in V(F)}$ of F such that each set V_v is finite. Then $X := \bigcup_{v\in V(F)} V_v$ is a finite subset of $V(G)$ and as G is locally finite, there are only finitely many edges in $M\left(X^{\setminus M/}\right)$ (recall that for each subset $X' \subseteq V(G^{\setminus M/})$, $M(X')$ is the set of crossedges of M that contract to a vertex in X'). We let $L := M\left(X^{\setminus M/}\right)$ denote this finite set of edges and note that the sets V_v are also subsets of $V(G^{\setminus L/})$. It follows that $(V_v)_{v\in V(F)}$ is also a model of F in $G^{\setminus \overline{L}/}[\![R^{\setminus \overline{L}/}]\!]$, a contradiction. □

Lemma 6.29 (Planar contraction of a single crossedge). *Let G be locally finite and 3-connected, and $\mathcal{T}$ be a region tangle of order 4 in G. Let (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) be two minimal non-degenerate crossing separations of $\mathcal{T}$. Let s_1s_2 be the corresponding crossedge and G' be the graph obtained from G after contracting s_1s_2. Let $(Y_i', S_i', Z_i') := (Y_i, S_i, Z_i)^\vee$ be the projection of (Y_i, S_i, Z_i) to G' for each $i \in \{1, 2\}$. Let $R := R_{\mathcal{T}} \subseteq V(G)$ and $R' := R^\vee$. If $G'[\![R']\!]$ is planar, then so is $G[\![R]\!]$.*

Proof. We let $H := G[\![R]\!]$ and $H' := G'[\![R']\!]$ and for $i \in \{1, 2\}$, we write $S_i = \{s_i, t_i, r_i\}$ such that s_1s_2 is the crossedge between (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2). Since (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) are crossing, the edge s_1s_2 belongs to $E(G[R]) \subseteq E(H)$. Note that in particular we have $6 = |S_1 \cup S_2| \leqslant |R|$.

Claim 6.30. The neighborhood of s_1 in H is:

$$N_H(s_1) = \{s_2\} \cup (\mathrm{fc}(S_2) \setminus \{s_1\}),$$

and $\mathrm{fc}(S_2)$ is a triangle in H.

Proof of the Claim: Note that by definition of the torso, the projection $(Y_i \cap R, S_i \cap R, Z_i \cap R)$ of (Y_i, S_i, Z_i) to H is a separation of H. Hence the only possible neighbors of s_1 in H must lie in $(R \cap Z_1 \cap Y_2) \cup \{t_2, r_2\}$ (see Figure 1.5b).

Note that as G is 3-connected, H must also be 3-connected: this comes from the fact that $|V(H)| \geqslant 6$ and from the observation that any separator $S \subseteq R = V(H)$ of H is also a separator of G. Thus in particular every vertex of H has degree at least 3.

Then, as $|\{s_2\} \cup (\mathrm{fc}(S_2) \setminus \{s_1\})| = 3$, it is enough to prove that $N_H(s_1) \subseteq \{s_2\} \cup (\mathrm{fc}(S_2) \setminus \{s_1\})$ as equality will be immediately implied as $d_H(s_1) \geqslant 3$. For this we let $t \in N_H(s_1) \setminus \{s_2\}$. We distinguish two cases:

- If $t \in S_2$, then without loss of generality let $t = t_2$. First note that if t is not an endpoint of some crossedge then $t \in S_2 \cap \mathrm{fc}(S_2)$ and there is nothing to prove. Thus we assume that there exists a crossedge t_2s_3 incident to t_2 for some s_3 and we prove that this case implies a contradiction, which will imply the desired inclusion. As $t_2 \neq s_2$ and $E^\times_{\mathrm{nd}}(\mathcal{T})$ is a matching, we have $s_3 \neq s_1$ and there exists $(Y_3, S_3, Z_3) \in \mathcal{T}_{\mathrm{nd}}$ that crosses (Y_2, S_2, Z_2) via the crossedge t_2s_3. As (Y_1, S_1, Z_1) and (Y_2, S_2, Z_2) cross, we have $s_1 \in Y_2$. As (Y_2, S_2, Z_2) and (Y_3, S_3, Z_3) cross, we have $t_2 \in Y_3$ and $S_3 \setminus \{s_3\} \subseteq Z_2$. As we assumed that $s_1t_2 \in E(H)$, $s_1 \neq s_3$ and as $t_2 \in Y_3$, we must have $s_1 \in Y_3 \cup (S_3 \setminus \{s_3\})$. This implies a contradiction as $Y_2 \cap Y_3 = \emptyset$ and $Y_2 \cap (S_3 \setminus \{s_3\}) = \emptyset$.

- If $t \notin S_2$, then we must have: $t \in R \cap Y_2 \cap Z_1$ and by definition of R, as $t \notin \bigcap_{(Y,S,Z)\in\mathcal{T}_{\mathrm{nd}}} Z$, this means that $t \in S_3$ for some $(Y_3, S_3, Z_3) \in \mathcal{T}_{\mathrm{nd}} \setminus \{(Y_1, S_1, Z_1), (Y_2, S_2, Z_2)\}$. By Lemma 6.9, (Y_2, S_2, Z_2) and (Y_3, S_3, Z_3) are either orthogonal or crossing. If we were in the former case, then we should have $S_3 \cap Y_2 = \emptyset$, which is impossible as $t \in S_3 \cap Y_2$. Hence (Y_2, S_2, Z_2) and (Y_3, S_3, Z_3) are crossing, and if s_3 denotes the endpoint of the crossedge between S_2 and S_3, as $S_3 \cap Y_2 = \{s_3\}$ we must have $t = s_3$ so we are done as $s_3 \in \mathrm{fc}(S_2) \setminus \{s_1\}$.

The fact that $\mathrm{fc}(S_2)$ forms a clique follows from [Gro16, Lemma 4.33]. ◊

Note that by symmetry, Claim 6.30 also implies that we have $N_H(s_2) = \{s_1\} \cup (\mathrm{fc}(S_1) \setminus \{s_2\})$ and that $\mathrm{fc}(S_1)$ is a clique.

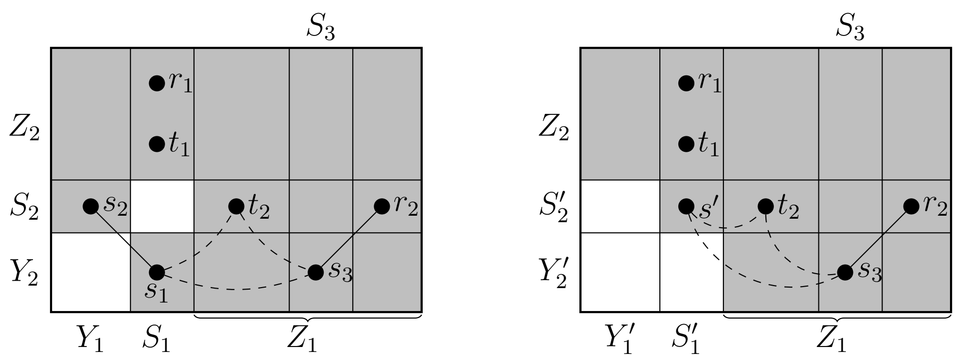

Figure 1.6: Left: The graph G when S_2 is incident to exactly 2 crossedges. Here t_2 is part of no crossedge and S_2 and S_3 are crossing via the crossedge r_2s_3. Hence, $\mathrm{fc}(S_2) = \{s_1, t_2, s_3\}$. Right: The graph G' obtained after contracting the crossedge s_1s_2. The dashed edges are edges that appear in H and H' respectively. The situation is identical when S_2 is incident to 3 crossedges, but harder to illustrate in 2 dimensions.

Recall that by Wagner's theorem [Wag37], a graph is planar if and only if it is K_5 and $K_{3,3}$-minor free. Hence, it is enough to prove that if H contains K_5 or $K_{3,3}$ as a minor, then so does H'. We write $\mathrm{fc}(S_i) = \{s_{3-i}, u_i, v_i\}$ for $i \in \{1, 2\}$, and we recall that the vertex of H' resulting from the contraction of s_1 and s_2 is denoted by s'.

Claim 6.31. If H contains a K_5-minor, then so does H'.

Proof of the Claim: Let $(V_1, \dots V_5)$ be a model of K_5 in H. Let $V'_1, \dots V'_5$ be the projection of the sets V_i to H'. If s_1 and s_2 are in the same set V_i, then $(V'_1, \dots V'_5)$ is also a model of K_5 in H', so we can assume that the vertices s_1 and s_2 belong to distinct sets V_i, say $s_1 \in V_1$ and $s_2 \in V_2$. As by Claim 6.30, s_1 has degree 3 in H, we have $V_1 \neq \{s_1\}$, so V_1 contains one neighbor of s_1 distinct of s_2, say u_2. Since u_2 and v_2 are adjacent in H, the edge u_2s' in H' has an endpoint in $V'_1 \setminus \{s'\}$ and an endpoint in V'_2. Moreover as $u_2v_2 \in E(H')$, the set $V'_1 \setminus \{s'\}$ is connected in H'. Thus $(V'_1 \setminus \{s'\}, V'_2, V'_3, V'_4, V'_5)$ is a model of K_5 in H', as desired. ◊

Claim 6.32. If H contains a $K_{3,3}$-minor, then H' contains a $K_{3,3}$-minor or a K_5-minor.

Proof of the Claim: Let $(V_1, \dots V_6)$ be a model of $K_{3,3}$ in H, such that V_i is adjacent to V_j if i and j have different parities. Let $V'_1, \dots V'_6$ be their projection to H'. If s_1 and s_2 are in the same set V_i, then $(V'_1, \dots V'_6)$ is also a model of $K_{3,3}$ in H', so we can assume that the vertices s_1 and s_2 belong to distinct sets V_i, say $s_1 \in V_1$ and $s_2 \notin V_1$.

If $u_2 \in V_1$, then the edges $s'u_2$ and u_2v_2 in H' ensure that $V'_1 \setminus \{s'\}$ remains connected and that $(V'_1 \setminus \{s'\}, V'_2, \dots V'_6)$ is a model of $K_{3,3}$ in H'. Thus we can assume that $u_2 \notin V_1$ and similarly $v_2 \notin V_1$. Since V_1 is connected, we must have $V_1 = \{s_1\}$.

As s_1 has degree three and V_1 is adjacent to V_2, V_4 and V_6, this implies that s_2, u_2 and v_2 must belong to different sets V_{2i}, say $s_2 \in V_2$, $u_2 \in V_4$ and $v_2 \in V_6$. By applying the same reasoning as for s_2, we obtain $V_2 = \{s_2\}$, $u_1 \in V_3$ and $v_1 \in V_5$. But then $(\{s'\}, V'_3, V'_4, V'_5, V'_6)$ is a model of K_5 in H'. ◊

This concludes the proof of Lemma 6.29. □

Proof of Proposition 6.27. Assume that $G^{\setminus M/}[\![R^{\setminus M/}]\!]$ is planar and for every $L \subseteq M$, set $\overline{L} := M \setminus L$. Then, using Lemma 6.29, we can easily prove by induction on $|L| \in \mathbb{N}$ that for any *finite* set $L \subseteq M$, $G^{\setminus \overline{L}/}[\![R_{\mathcal{T}}^{\setminus \overline{L}/}]\!]$ is planar. In order to be able to use induction, we also need to observe that for the contraction of a single crossedge, the equality $R_{\mathcal{T}'} = R_{\mathcal{T}}^{\vee}$ holds. This is proved in [Gro16, Section 4.5] and can be deduced from item (2) of Theorem 6.20. We thus conclude by Lemma 6.28 that $G[\![R_{\mathcal{T}}]\!]$ is also planar. □

7 Quasi-transitive graphs excluding a minor

In this section, we prove a general structure theorem for locally finite quasi-transitive graphs avoiding some countable graph as a minor. We will then see in Section 8 a number of applications of this result. We start by giving some motivation and context.

7.1 Introduction

A central result in modern graph theory is the Graph Minor Structure Theorem of Robertson and Seymour [RS03], extended to infinite graphs by Kříž and Thomas [KT90]. This theorem states that any graph G avoiding a fixed finite minor has a tree-decomposition, such that each torso is close to being embeddable on a surface of bounded genus. A result of the same type was also proved by Diestel and Thomas [DT99] for graphs excluding a countable graph as a minor. However the tree-decompositions given by these results are not canonical in general, and thus as we saw in the previous subsections, when applied to a quasi-transitive graph we cannot expect their torsos to be still quasi-transitive in general. In this section, we prove a canonical version of the Graph Structure Theorem for locally finite quasi-transitive graphs that exclude a minor. The additional hypothesis of quasi-transitivity has the advantage of making the structure theorem much cleaner: instead of being almost embeddable on a surface of bounded genus, each torso of the tree-decomposition is now simply finite or planar. Intuitively, this is not very surprising as we already observed (see Remark 5.12) that every quasi-transitive graph embeddable in a surface of bounded genus must be either finite or planar, thus if some canonical version of the graph minor structure theorem existed, it should satisfy the properties of Theorem 7.1.

Theorem 7.1 (see Theorem 7.3)**.** *Every locally finite quasi-transitive graph avoiding the countable clique as a minor has a canonical tree-decomposition of finite adhesion whose torsos are finite or planar.*

The tree-decomposition in Theorem 7.1 will be obtained by refining the tree-decomposition obtained in the following more detailed version of the result, which might be useful for applications.

Theorem 7.2 (see Theorem 7.5)**.** *Every locally finite quasi-transitive graph G avoiding the countable clique as a minor has a canonical tree-decomposition with adhesion at most 3 in which each torso is a minor of G, and is planar or has bounded treewidth.*

Interestingly, the proof does not use the original structure theorem of Robertson and Seymour [RS03] (which is a deep result proved in a series of 16 papers) or its extension to infinite graphs by Kříž and Thomas [KT90]. Instead, we rely mainly on the series of results and tools of Grohe [Gro16] we presented in Section 6, together with a result of Thomassen [Tho92] on locally finite quasi-4-connected graphs. Our proof also crucially relies on a recent result of Carmesin, Hamann, and Miraftab [CHM22], which shows that there exists a canonical tree-decomposition that distinguishes all tangles of a given order (in our case, of order 4).

We already mentioned the result of Thomassen [Tho92, Proposition 5.6] that if a locally finite quasi-transitive graph has only one end, then this end must be thick. At some point of our proof, we also need to show the stronger result (see Proposition 7.9), of independent interest, that for any $k \geqslant 1$, a locally finite quasi-transitive graph cannot have only one end of degree k.

Overview of the proof of Theorems 7.1 and 7.2. Consider a locally finite quasi-transitive graph G that excludes the countable clique K_∞ as a minor. Thomassen [Tho92] proved that if G is quasi-4-connected, then G is planar or has finite treewidth (see Corollary 7.8), which implies Theorem 7.2 in this case, with a trivial tree-decomposition consisting of a single node.

To deal with the more general case, the general strategy is the same that in Section 5, in the sense that we show the existence of a canonical tree-decomposition of G whose torsos are minors of G (and thus still K_∞-minor free) and connected enough to apply Thomassen's result on them. We refine Tutte's decomposition using the results from 6. However, as illustrated by Example 3.11, one cannot hope to find a version of Theorem 3.10 where we furthermore ask the decomposition to be canonical. We instead show that we can find a canonical tree-decomposition of G whose torsos are minors of G and "look like quasi-4-connected" graphs.

For this, we proceed in two steps. First, we use Theorem 6.3 to find a canonical tree-decomposition of any 3-connected graph G that distinguishes all its tangles of order 4. Using this result, we show that we can assume that the graph under consideration admits a unique tangle $\mathcal{T}$ of order 4. We then show using results from Section 6 that G has a canonical tree-decomposition of adhesion 3 which is a star and whose torsos are all minors of G and finite, except for the torso H associated to the center of the star, which has the following property: there exists a matching $M \subseteq E(H)$ which is invariant under the action of the automorphism group of G and such that the graph $H' := H/M$ obtained after the contraction of the edges of M is quasi-transitive, locally finite, and quasi-4-connected. In particular, the aforementioned result of Thomassen [Tho92] then implies that H' is planar or has bounded treewidth. Using Proposition 6.27 we then observe that even if H itself is not necessarily quasi-4-connected, it is still planar or has bounded treewidth, which is enough to conclude the proof of Theorem 7.2. The final step to prove Theorem 7.1 consists in refining the tree-decomposition by decomposing again the torsos of infinite treewidth using the canonical decompositions given Theorem 4.3 (*iii*).

If we apply our proof to the left graph H from Example 3.11, then we will simply obtain the trivial tree-decomposition with a single planar part. More precisely, on this example, G has a single tangle of order 4, namely the one induced by its unique end so the canonical tree-decomposition given by Theorem 6.3 is the trivial tree-decomposition, with a single torso $H := G$. Then, the aforementioned $\mathrm{Aut}(H)$-invariant matching M given by Section 6 is exactly the set of edges with endpoints in two different triangles, so in particular the graph $H' := H/M$ obtained after contracting each edge of M corresponds exactly the the right graph from Example 3.11, and is thus planar 4-connected.

Related work. We mentioned in Section 3 that Carmesin and Kurkofka [CK23] recently worked on decompositions of 3-connected graphs with an approach that differs from the one of Grohe. They obtained a canonical decomposition (which is not exactly a tree-decomposition of G) into basic pieces consisting in quasi-4-connected graphs, wheels or thickenings of $K_{3,m}$ for $m \geqslant 0$. In particular, it could be possible to apply their methods instead of those from Section 6 to get another decomposition theorem similar to Theorem 7.1. However it is not clear to us whether this work amounts to consider again the $\mathrm{Aut}(G)$-invariant matching from Section 6, or if it could lead to another proof based on completely different ideas.

7.2 Structure of K_∞-minor free quasi-transitive graphs

Our main result in this section is the following more precise version of Theorem 7.1.

Theorem 7.3. *Let G be a locally finite graph excluding K_∞ as a minor and let Γ be a group with a quasi-transitive action on G. Then there is an integer k such that G admits a Γ-canonical tree-decomposition $(T, \mathcal{V})$ of finite adhesion, with $\mathcal{V} = (V_t)_{t\in V(T)}$, whose torsos $G[\![V_t]\!]$ either have size at most k or are Γ_t-quasi-transitive 3-connected planar minors of G. Moreover, the edge-separations of $(T, \mathcal{V})$ are tight.*

Remark 7.4. A natural question is whether we can bound the maximum size k of a finite bag in Theorem 7.3 by a function of the forbidden minor, when G excludes some finite minor instead of the countable clique K_∞. By taking the free product of the cyclic groups $\mathbb{Z}_k$ and $\mathbb{Z}$, with their natural sets of generators, we obtain a 4-regular Cayley graph consisting of cycles of length k arranged in a tree-like way (see Chapter 2 for a definition of Cayley graphs). This graph has no K_4 minor, but in any canonical tree-decomposition, each cycle C_k has to be entirely contained in a bag, and thus there is no bound on the size of a bag as a function of the forbidden minor in Theorem 7.3. We can replace $\mathbb{Z}_k$ in this construction by the toroidal grid $\mathbb{Z}_k \times \mathbb{Z}_k$, and obtain a Cayley graph with no K_8-minor, such that the bags in any (non-necessarily canonical) tree-decomposition of finitely bounded adhesion are arbitrarily large.

We will also prove the following version of Theorem 7.2 at the same time.

Theorem 7.5. *Let G be a locally finite graph excluding K_∞ as a minor and let Γ be a group with a quasi-transitive action on G. Then there is an integer k such that G admits a Γ-canonical tree-decomposition $(T, \mathcal{V})$, with $\mathcal{V} = (V_t)_{t \in V(T)}$, of adhesion at most 3, and whose torsos $G[\![V_t]\!]$ are Γ_t-quasi-transitive minors of G which are either planar or have treewidth at most k. The edge-separations of $(T, \mathcal{V})$ are all non-degenerate.*

Remark 7.6. If we carefully consider the proof of Theorem 7.5, we can check that if G has only one end (as in the example of Figure 1.4), then the tree-decomposition we obtain has adhesion 3 and consists of a star with one infinite bag associated to its central vertex z_0, and finite bags on its branches. In particular, $G[\![V_{z_0}]\!]$ cannot have bounded treewidth, as otherwise it would have more than one end, hence it must be planar. Thus, Theorem 7.5 implies that every one-ended locally finite quasi-transitive graph that excludes a minor can be obtained from a one-ended quasi-transitive planar graph by attaching in a canonical way some finite graphs on it along separators of order at most 3.

7.3 Tools

Our proof of Theorems 7.3 and 7.5 mainly consists in an application of Theorem 6.26 together with the following result of Thomassen:

Theorem 7.7 (Theorem 4.1 in [Tho92])**.** *Let G be a locally finite, quasi-transitive, quasi-4-connected graph G. If G has a thick end, then G is either planar or admits the countable clique K_∞ as a minor.*

A direct consequence of Theorem 7.7 is the following, which will be our base case in what follows.

Corollary 7.8. *Let G be a quasi-transitive, quasi-4-connected, locally finite graph which excludes the countable clique K_∞ as a minor. Then G is planar or has finite treewidth.*

Proof. Assume that G is non-planar. As G is K_∞-minor free, by Theorem 7.7, all its ends have have finite degree. Then by Theorem 4.3, G has finite treewidth. □

Thomassen proved that if a quasi-transitive graph has only one end, then this end must be thick [Tho92, Proposition 5.6]. We prove the following generalization, which might be of independent interest.

Proposition 7.9. *Let $k \geqslant 1$ be an integer, and let G be a locally finite quasi-transitive graph. Then G cannot have exactly one end of degree exactly k.*

Proof. Assume without loss of generality that G is connected, since otherwise each component of G is also quasi-transitive locally finite, and we can restrict ourselves to a single component containing an end of degree exactly k. Let Γ be a group acting quasi-transitively on G. Assume that G has an end ω of degree exactly k for some integer $k \geqslant 1$. As explained in [TW93, Section 4], an application of Menger's theorem implies that there exists an infinite sequence of sets $S_0, S_1, \ldots$ of size k such that for each $i \geqslant 0$, S_{i+1} belongs to the component G_i of $G_{i-1} - S_i$ where ω lives (where we set $G_0 := G$), and such that there exist k vertex-disjoint paths $P_{1,i}, \ldots, P_{k,i}$ from the k vertices of S_i to the k vertices of S_{i+1}. By concatenating these paths, we obtain k vertex-disjoint rays in G living in ω. As G is connected and locally finite, note that up to extracting a subsequence of $(S_i)_{i\geqslant 0}$, we may assume that the k paths $P_{1,i}, \ldots, P_{k,i}$ are in the same component of $G - (S_i \cup S_{i+1})$. Hence if we set $(Y_i, S_i, Z_i) := (G - (G_i \cup S_i), S_i, G_i)$ for each $i \geqslant 1$, (Y_i, S_i, Z_i) is a tight separation such that for each $i \geqslant 1$, $(Y_{i+1}, S_{i+1}, Z_{i+1}) \leqslant_{\mathrm{RS}} (Y_i, S_i, Z_i)$. Hence by Lemma 3.1, as there are only finitely many Γ-orbits of tight separations of size k, there exist $i < j$ and $\gamma \in \Gamma$ such that $\gamma \cdot (Y_i, S_i, Z_i) = (Y_j, S_j, Z_j)$. Assume without loss of generality that $(i, j) = (0, 1)$. Note that by definition of $\leqslant_{\mathrm{RS}}$, the action of γ preserves the order $\leqslant_{\mathrm{RS}}$, i.e. for each $(Y, S, Z) \leqslant_{\mathrm{RS}} (Y', S', Z')$, we must have $\gamma \cdot (Y, S, Z) \leqslant_{\mathrm{RS}} \gamma \cdot (Y', S', Z')$. We now consider the sequence of separations $(Y'_i, S'_i, Z'_i)_{i\geqslant 0}$ defined for each $i \geqslant 0$ by: $(Y'_i, S'_i, Z'_i) := \gamma^i \cdot (Y_0, S_0, Z_0)$. Then the sequence $(Y'_i, S'_i, Z'_i)_{i\geqslant 0}$ is strictly decreasing according to $\leqslant_{\mathrm{RS}}$. Recall that there exist k vertex-disjoint paths from S_0 to S_1 that extend to k disjoint rays belonging to ω. Then for each $i \geqslant 0$, there exist k vertex-disjoint paths from S'_i to S'_{i+1} such that their concatenations consists in k vertex-disjoint rays that belong to some end ω' of degree exactly k (the fact that the end has degree at most k follows from the fact that all the sets S'_i are separators of size k in G). If $\omega' \neq \omega$ then we are done, so we assume that $\omega' = \omega$. Now, observe that the sequence $(Y''_i, S''_i, Z''_i)_{i\geqslant 0}$ defined for each $i \geqslant 0$ by $(Y''_i, S''_i, Z''_i) := (Y_0, S_0, Z_0) \cdot \gamma^{-i}$ also satisfies that for each $i \geqslant 0$, there exists k vertex-disjoint paths $P''_{j,i} := P_{j,0} \cdot \gamma^{-i}$ for $j \in [k]$ from S''_{i+1} to S''_i. If we consider the k vertex-disjoint rays obtained from the concatenation of the paths $P''_{j,k}$, these rays must belong to the same end ω'' as for each i, the paths $P''_{j,k}$ are in the same component of $G - (S''_i \cup S''_{i+1})$. The end ω'' must have degree exactly k as each (Y''_i, S''_i, Z''_i) is a separation of order k. Moreover the sequence $(Y''_i, S''_i, Z''_i)_{i\geqslant 0}$ is strictly increasing according to $\leqslant_{\mathrm{RS}}$, hence ω and ω'' cannot live in the same component of $G - S_0$. Thus we found an end ω'' distinct from ω of degree k. □

Proposition 7.9 and its proof are reminiscent of Halin's classification of the different types of action an automorphism of a quasi-transitive locally finite graph G can have on the ends of G [Hal73, Theorem 9]. However it is not clear for us whether Proposition 7.9 can be seen as an immediate corollary of Halin's work.

7.4 Proof of Theorems 7.3 and 7.5

Let G be a locally-finite quasi-transitive graph excluding K_∞ as a minor and let Γ be a group inducing a quasi-transitive action on G. Let $(T, \mathcal{V})$, with $\mathcal{V} = (V_t)_{t\in V(T)}$, be a Γ-canonical tree-decomposition of adhesion at most 2 obtained by applying Theorem 3.9 to G. By Lemma 3.17, for each $t \in V_t$, Γ_t acts quasi-transitively on $G_t := G[\![V_t]\!]$. Moreover, as G_t

is a minor of G, it must also exclude K_∞ as a minor.

As the edge-separations of $(T, \mathcal{V})$ are tight, Lemma 3.1 implies that $E(T)$ has only finitely many Γ-orbits, and thus $V(T)$ also has only finitely many Γ-orbits. We let $t_1, \ldots, t_m$ be representatives of the orbits of $V(T)/\Gamma$. For each finite torso G_{t_i} of $(T, \mathcal{V})$, we define $(\widetilde{T}_{t_i}, \widetilde{\mathcal{V}}_{t_i})$ as the trivial tree-decomposition of G_{t_i} (in which the tree $\widetilde{T}_{t_i}$ contains a single node). For each infinite, 3-connected torso G_{t_i} of $(T, \mathcal{V})$, we let $(\widetilde{T}_{t_i}, \widetilde{\mathcal{V}}_{t_i})$ be a Γ_{t_i}-canonical tree-decomposition of G_{t_i} obtained by applying Theorem 6.3 to G_{t_i}, i.e. $(\widetilde{T}_{t_i}, \widetilde{\mathcal{V}}_{t_i})$ distinguishes efficiently all the tangles of G_{t_i} of order 4. By Remark 6.4, the edge-separations of $(\widetilde{T}_{t_i}, \widetilde{\mathcal{V}}_{t_i})$ in G_{t_i} are all distinct. By Remark 6.7, the edge-separations of $(\widetilde{T}_{t_i}, \widetilde{\mathcal{V}}_{t_i})$ in G_{t_i} are non-degenerate. Hence by Lemma 6.6, the torsos of $(\widetilde{T}_{t_i}, \widetilde{\mathcal{V}}_{t_i})$ are minors of G_{t_i}. We now use Corollary 3.15 and find a refinement $(T_1, \mathcal{V}_1)$ of $(T, \mathcal{V})$ with respect to some family $(T_t, \mathcal{V}_t)_{t \in V(T)}$ of Γ_t-canonical tree-decompositions of G_t such that the construction $t \mapsto (T_t, \mathcal{V}_t)_{t \in V(T)}$ is Γ-canonical and such that for each $i \in I$, $(T_{t_i}, \mathcal{V}_{t_i})$ is a subdivision of $(\widetilde{T}_{t_i}, \widetilde{\mathcal{V}}_{t_i})$. Since the construction $t \mapsto (T_t, \mathcal{V}_t)_{t \in V(T)}$ is Γ-canonical, for each $t \in V(T_1)$ the decomposition $(T_t, \mathcal{V}_t)$ is Γ_t-canonical and efficiently distinguishes the tangles of order 4 of G_t (by a slight abuse of notation, we keep denoting by G_t the torso of the tree-decomposition $(T_1, \mathcal{V}_1)$ associated to the node $t \in V(T_1)$). Note that by construction, the adhesion sets of $(T_1, \mathcal{V}_1)$ have size at most 3 and all the edge-separations are tight. Moreover, the torsos of each tree-decomposition $(T_t, \mathcal{V}_t)$ are minors of G_t for each $t \in V(T)$, and as the torsos of $(T, \mathcal{V})$ are minors of G, we also have that the torsos of $(T_1, \mathcal{V}_1)$ are minors of G. In particular, they also exclude K_∞ as a minor. Moreover, by Lemma 3.17, for each $t \in V(T_1)$, Γ_t acts quasi-transitively on G_t. By Lemma 3.1, since all edge-separations of $(T_1, \mathcal{V}_1)$ are tight and have order at most 3, the graph G_t is locally finite for each $t \in V(T_1)$.

Claim 7.10. For each $t \in V(T_1)$ such that G_t is infinite, G_t is 3-connected and has a unique tangle $\mathcal{T}_t$ of order 4. Moreover $\mathcal{T}_t$ is a Γ_t-invariant region tangle and every end of G_t has degree at least 4.

Proof of the Claim: Consider a node $t \in V(T_1)$ such that G_t is infinite. As all torsos are cycles, subgraphs of complete graphs of size at most 3, or 3-connected, G_t itself is 3-connected. Since G_t is connected and infinite, it contains some end ω. Let $\mathcal{T}_t := \{(Y, S, Z), |S| \leqslant 3 \text{ and } \omega \text{ lives in } Z\}$ be defined in G_t. Note that $\mathcal{T}_t$ is a tangle of order 4 in G_t. As G_t is a minor of G, by Lemma 6.1 every tangle $\mathcal{T}'$ of order 4 in G_t induces a tangle $\mathcal{T}$ of order 4 in G, and by Remark 6.2 this mapping is injective. Moreover, note that if (Y, S, Z) is an edge-separation of $(T_1, \mathcal{V}_1)$ such that $V_t \subseteq Z \cup S$, then if $\mathcal{M}$ is any faithful model of G_t in G, the projection $(Y', S', Z') := \pi_{\mathcal{M}}(Y, S, Z)$ is such that $Y' = \emptyset$. Thus $(Y', S', Z') \in \mathcal{T}'$, hence $(Y, S, Z) \in \mathcal{T}$. This means that every edge-separation of $(T_1, \mathcal{V}_1)$ is oriented toward t by $\mathcal{T}$. Hence if G_t admits two distinct tangles $\mathcal{T}_1', \mathcal{T}_2'$ of order 4, the two associated tangles $\mathcal{T}_1, \mathcal{T}_2$ given by Lemma 6.1 must be distinct and not distinguished by $(T_1, \mathcal{V}_1)$, a contradiction. This proves the existence and uniqueness of a tangle $\mathcal{T}_t$ of order 4 in G_t.

Note that as Γ_t acts on G_t and $\mathcal{T}_t$ is the unique tangle of order 4 in G_t, the tangle $\mathcal{T}_t$ is Γ_t-invariant (as a family of separations).

We can also observe that if the end ω in G_t has degree at most 3, then by Proposition 7.9, G_t has another end ω' of degree at most 3 and the construction of $\mathcal{T}_t$ using the end ω' instead of ω yields a different tangle of order 4, which contradicts the uniqueness of $\mathcal{T}_t$. So every end of G_t has degree at least 4.

It remains to prove that $\mathcal{T}_t$ is a region tangle. If not we can find an infinite decreasing sequence of separations of order 3 in G_t, and this sequence defines an end of degree 3 in G_t, which contradicts the fact that every end of G_t has degree at least 4. $\Diamond$

We will need to decompose further the infinite torsos of the tree-decomposition $(T_1, \mathcal{V}_1)$. Let $t \in V(T_1)$ be such that G_t is infinite, and let $\mathcal{T}_t$ be the region tangle of order 4 in G_t given by Claim 7.10. We let $M_t := E^{\times}_{\mathrm{nd}}(\mathcal{T}_t)$ denote the set of crossedges of $\mathcal{T}_t$, $(T'_t, \mathcal{V}'_t)$ be the Γ_t-canonical tree-decomposition of G_t given by Lemma 6.12, and $z_0 \in V(T_t)$ be the center of the star T'_t. By Lemmas 3.17 and 6.11, the graph $H := G_t[\![V'_{z_0}]\!]$ is a Γ_t-quasi-transitive faithful minor of G_t, thus it must also exclude K_∞ as a minor.

Now we observe that Γ_t induces a quasi-transitive group action on $H^{\setminus M_t/}$: for each $w \in V(H^{\setminus M_t/})$ and every $\gamma \in \Gamma_t$, we set:

$$\gamma \cdot w := \begin{cases} s_{\gamma\cdot u, \gamma\cdot v} & \text{if } w = s_{u,v}, \text{ for some } \{u,v\} \in M_t, \text{ and} \\ \gamma \cdot w & \text{otherwise,} \end{cases}$$

where we recall that the notation $s_{u,v}$, for $\{u,v\} \in M_t$, is introduced at the beginning of Section 6.5. As M_t is Γ_t-invariant, we easily see that the mapping γ defines a bijection over $V(H^{\setminus M_t/})$. We let the reader check that it gives a graph isomorphism of $H^{\setminus M_t/}$. Note that the number of Γ_t-orbits of $V(H^{\setminus M_t/})$ is at most the number of Γ_t-orbits of $V(H)$, hence it must be finite.

As $H^{\setminus M_t/}$ is a minor of H, it also excludes the countable clique K_∞ as a minor. It follows from Theorem 6.26 that $H^{\setminus M_t/}$ is quasi-4-connected. Hence, by Corollary 7.8, $H^{\setminus M_t/}$ either has finite treewidth or it is planar. It is not hard to observe that the treewidth of H is at most twice the treewidth of $H^{\setminus M_t/}$ so in particular if we are in the first case, H has also bounded treewidth. In the second case, Proposition 6.27 implies that H is also planar. In both cases, we obtain that $(T'_t, \mathcal{V}'_t)$ is a Γ_t-canonical tree-decomposition of G_t with non-degenerate edge-separations, adhesion 3 and where each torso is a minor of G_t and has either bounded treewidth or is planar. Eventually we can use Proposition 3.13 together with Lemma 3.12 as we did before to find a tree-decomposition $(T^*, \mathcal{V}^*)$ of G with the properties of Theorem 7.5.

We now explain how to derive Theorem 7.3: every torso $G[\![V_t]\!]$ of $(T^*, \mathcal{V}^*)$ which is neither finite nor planar must have bounded treewidth, hence by Theorem 4.3 it must admit a Γ_t-canonical tree-decomposition where each torso has bounded width. Exactly as before we can apply Corollary 3.15 to find a refinement of $(T^*, \mathcal{V}^*)$ with the properties of Theorem 7.3. □

8 Applications of Theorems 7.1 and 7.2

In this section we present some graphical applications of Theorem 7.1. Results and proofs from Sections 8.1, 8.2 and 8.3 come from the paper [EGLD23], while the content of Section 8.4 comes from the paper [EG24a]. We will also give in Sections 14 and 15 other applications of Theorem 7.1, more group-oriented.

8.1 The Hadwiger number of quasi-transitive graphs

A first consequence of Theorem 7.2, is a result on the Hadwiger number of locally finite quasi-transitive graphs. The *Hadwiger number* of a graph G is the supremum of the sizes of all finite complete minors in G. We say that a graph G *attains its Hadwiger number* if the supremum above is attained, that is if it is either finite, or G contains an infinite clique minor. Thomassen [Tho92] proved that every locally finite quasi-transitive 4-connected graph attains its Hadwiger number, and suggested that the 4-connectedness assumption might be unnecessary. We prove in Theorem 8.1 that this is indeed the case.

In fact we prove a stronger statement, namely that every locally finite quasi-transitive graph avoiding the countable clique as a minor also avoids a *singly-crossing* finite graph as a minor. We say that a graph H is *singly-crossing* if H can be embedded in the plane with a single edge-crossing. It was observed by Paul Seymour that Theorem 7.5 bears striking similarities with a structure theorem of Robertson and Seymour [RS93] related to the exclusion of a singly-crossing graph as a minor. Their theorem states that if H is singly-crossing, then there is a constant k_H such that any graph excluding H as a minor has a tree-decomposition with adhesion at most 3 in which all torsos are planar or have treewidth at most k_H. On the other hand, for any integer k there is a finite singly-crossing graph H_k such that any graph with a tree-decomposition with adhesion at most 3 in which all torsos are planar or have treewidth at most k must exclude H_k as a minor (this can be seen by taking H_k to be a 4-connected triangulation of a sufficiently large grid, and adding an edge between two non-adjacent vertices lying on incident faces). Using this last observation, the following result is now an immediate consequence of Theorem 7.2.

Theorem 8.1. *For every locally finite quasi-transitive graph G avoiding the countable clique K_∞ as a minor, there is a finite singly-crossing graph H such that G is H-minor-free. In particular there is an integer k such that G is K_k-minor-free.*

Note that in this application we have not used explicitly the property that the underlying tree-decomposition was canonical, but it is used implicitly in the sense that this is what garantees that the treewidth of the torsos is uniformly bounded in Theorem 7.5.

8.2 Accessibility

Recall that a graph G is vertex-accessible if there is an integer k such that any two distinct ends ω_1, ω_2 in G, can be separated by a set of at most k vertices in G. We already mentioned Dunwoody's result [Dun09] (see also [Ham18b, Ham18a] and/or equivalently Corollary 5.6 for an alternate approach) that locally finite quasi-transitive planar graphs are vertex-accessible. Here we extend the result to locally finite quasi-transitive graphs excluding the countable clique K_∞ (and not necessarily K_5 and $K_{3,3}$) as a minor, and in particular to locally finite quasi-transitive graphs from any proper minor-closed family.

Theorem 8.2. *Every locally finite quasi-transitive graph avoiding the countable clique K_∞ as a minor is vertex-accessible.*

Proof. Let G be a locally finite graph avoiding the countable clique K_∞ as a minor, with a group Γ acting quasi-transitively on G. Let $(T, \mathcal{V})$, with $\mathcal{V} = (V_t)_{t \in V(T)}$, be a Γ-canonical tree-decomposition of G of adhesion at most 3 obtained by applying Theorem 7.5 to G. In particular all torsos are quasi-transitive minors of G, and all non-planar torsos have bounded treewidth. Observe that any two ends living in different parts of the tree-decomposition are separated by the separator of size at most 3 of an edge-separation of the tree-decomposition. Consider any node $t \in V(T)$. If $G[\![V_t]\!]$ is planar then since it is quasi-transitive (by Lemma 3.17) and locally finite, $G[\![V_t]\!]$ is vertex-accessible [Dun09, Theorem 3.8] and thus there is an integer k_t such that all pairs of ends lying in $G[\![V_t]\!]$ can be separated by a set of at most k_t vertices. If $G[\![V_t]\!]$ has bounded treewidth then by Theorem 4.3 there is a integer k_t such that all ends of $G[\![V_t]\!]$ have degree at most k_t, and thus all pairs of ends lying in $G[\![V_t]\!]$ can be separated by a set of at most k_t vertices. As $V(T)/\Gamma$ is finite; there is only a finite number of possible values for the integers k_t, $t \in V(T)$, and thus their maximum k is well-defined. We have proved that every pair of ends in G can be separated by a set of at most $\max\{k, 3\}$ vertices, which concludes the proof. □

8.3 Generating closed walks in K_∞-minor-free quasi-transitive graphs

We recall that a set of closed walks $\mathcal{W}$ *generates* another set of closed walks $\mathcal{W}'$ if every element of $\mathcal{W}'$ can be obtained from elements of $\mathcal{W}$ by adding and deleting spurs, and performing sums, reflections and rotations (see Section 2).

The next result generalizes [Ham18b, Theorem 5.12] (which corresponds to the closed walk version of Theorem 5.4 we discussed earlier in Section 5) to the case where G is a quasi-transitive locally finite planar graph excluding K_∞ as a minor. We reuse some of the arguments of the proof of [Ham18b, Proposition 5.9] and combine them with our structure theorem to extend the result to graphs excluding the countable clique K_∞ as a minor.

Theorem 8.3. *Let G be a locally finite graph excluding the countable clique K_∞ as a minor and let Γ be a group acting quasi-transitively on G. Then the set of closed walks of G admits a Γ-invariant generating set with finitely many Γ-orbits.*

Proof. We consider a Γ-canonical tree-decomposition $(T, \mathcal{V})$, with $\mathcal{V} = (V_t)_{t \in V(T)}$, given by Theorem 7.3. We let A denote the set of pairs $\{x, y\}$ of vertices of G for which there exists an edge-separation (Y, S, Z) of $(T, \mathcal{V})$ such that $x, y \in S$ and $xy \notin E(G)$. By Remark 3.3, as the edge-separations associated to $(T, \mathcal{V})$ are tight, $E(T)/\Gamma$ is finite. As $(T, \mathcal{V})$ has finitely bounded adhesion, this implies that there is a finite number of Γ-orbits of A. We let $\{x_1, y_1\}, \ldots, \{x_\ell, y_\ell\}$ be representatives of these orbits. For each $j \in [\ell]$ we let P_j be a path from x_j to y_j (which always exists, since the edge-separations are tight). For each $\{x, y\} \in A$, we consider the representative $\{x_j, y_j\}$ in the Γ-orbit of $\{x, y\}$, and we define $f(x, y)$ as the image of the path P_j under an automorphism that maps $\{x_j, y_j\}$ to $\{x, y\}$. Note that $f(x, y)$ is an (x, y)-path in G.

We let G^+ be the graph obtained from G by adding all possible edges xy such that $xy \in E(G[\![V_t]\!])$ for some $t \in V(T)$. In other words the edge-set of G^+ is exactly $E(G) \uplus A$. For each walk W in G^+, we define the walk $f(W)$ in G as the walk obtained from W by replacing every edge (x,y) of W such that $\{x,y\} \in A$ by $f(x,y)$ (this definition extends the definition of f above, which applied to walks (x,y) of length 1 in G^+). For each set of walks $\mathcal{S} \subseteq \mathcal{W}(G^+)$, we let $f(\mathcal{S}) := \{f(W) \in \mathcal{W}(G), W \in \mathcal{S}\}$.

Claim 8.4. For every $W \in \mathcal{W}(G^+)$, if $W_1, \ldots, W_k \in \mathcal{W}(G^+)$ generate W in $\mathcal{W}(G^+)$, then $f(W)$ is generated by $f(W_1), \ldots, f(W_k)$ in $\mathcal{W}(G)$.

Proof of the Claim: Let $W \in \mathcal{W}(G^+)$ be generated by $W_1, \ldots, W_k \in \mathcal{W}(G^+)$. We prove by induction on the number of operations needed to generate W from $W_1, \ldots, W_k$ that $f(W)$ is generated by the closed walks $f(W_1), \ldots, f(W_k)$.

If $W = W_i$ for some $i \in [k]$, then the result is immediate. Assume that W is obtained from some closed walk W' after performing a rotation on W', and that W' is generated by $W_1, \ldots, W_k$. We write $W = (v_1, \ldots, v_r)$. If $v_1 = v_2$ or $v_1v_2 \in E(G)$, then $f(W)$ is obtained after performing a single rotation on $f(W')$. If $v_1v_2 \in E(G^+) \setminus E(G) = A$, then $f(W)$ is obtained after performing $|f(v_1, v_2)|$ rotations to $f(W')$. In any case if we assume by the induction hypothesis that $f(W')$ is generated by $f(W_1), \ldots, f(W_k)$, we are immediately done. The case where W is obtained after performing a reflection on W' is even simpler.

Now assume that W is the concatenation of two walks $W', W'' \in \mathcal{W}(G^+)$ for which the induction hypothesis holds. Then we observe by definition of f that $f(W) = f(W') \cdot f(W'')$. Then again we conclude by the induction hypothesis that $f(W)$ is generated by $f(W_1), \ldots, f(W_k)$.

Assume now that $W = (v_1, \ldots, v_{i-1}, v_i, v_{i+1}, \ldots, v_r)$ is obtained from $W' = (v_1, \ldots, v_{i-1}, v_{i+2}, \ldots, v_r)$ after adding the spur v_i between v_{i-1} and v_{i+2} with $v_{i-1} = v_{i+1}$. We let $x := v_{i-1}$ and $y := v_i$ and distinguish two cases:

- If $xy \in E(G)$, then we observe that by definition of f, $f(W)$ is obtained from $f(W')$ after adding the same spur so we are done using the induction hypothesis on W'.
- If $xy \in A$, then $f(W)$ must be of the form $U_1 \cdot f(x,y) \cdot f(x,y)^{-1} \cdot U_2$, where $U_1 \cdot U_2 = f(W')$. This means that $f(W)$ can be obtained from $f(W') = U_1 \cdot U_2$ by adding $|f(x,y)|$ spurs, hence the induction hypothesis on W' implies that $f(W)$ is generated by $f(W_1), \ldots, f(W_k)$.

Finally assume that $W = (v_1, \ldots, v_{i-1}, v_{i+2}, \ldots, v_r)$ is obtained from $W' = (v_1, \ldots, v_{i-1}, v_i, v_{i+1}, v_{i+2}, \ldots, v_r)$ after deleting the spur v_i. Again we let $x := v_{i-1} = v_{i+1}$ and $y := v_i$ and distinguish two cases:

- If $xy \in E(G)$, then as above, $f(W)$ is obtained from $f(W')$ after the removal of a spur and we are immediately done by applying the induction hypothesis on W'.
- If $xy \in A$, then we claim that $f(W)$ is generated by $f(W')$ as $f(W) = U_1 \cdot U_2$, with $U_1 := f((v_1, \ldots, x))$ and $U_2 := f((x, \ldots, v_r))$, and $f(W') = U_1 \cdot f(x,y) \cdot f(x,y)^{-1} \cdot U_2$. This shows that $f(W)$ is obtained from $f(W')$ after deleting $|f(x,y)|$ spurs and we can conclude by the induction hypothesis applied to W' that $f(W)$ is generated by $f(W_1), \ldots, f(W_k)$.

This concludes the proof of Claim 8.4. $\Diamond$

Claim 8.5. $\mathcal{W}(G)$ is generated by $\bigcup_{t\in V(T)} f(\mathcal{W}(G[\![V_t]\!]))$.

Proof of the Claim: Let $W \in \mathcal{W}(G)$. First, note that W can be generated in G^+ by closed walks of $\bigcup_{t\in V(T)} \mathcal{W}(G[\![V_t]\!])$. This comes from the following observation: fix any edge t_1t_2 in T, with associated separation (Y, S, Z) in G. Then any closed walk W in G can be written as the sum of closed walks in $G^+[Y \cup S]$ and $G^+[S \cup Z]$, followed by the removal of spurs corresponding to the edges of the adhesion S. Thus we proved that $\mathcal{W}(G)$ is generated by $\bigcup_{t\in V(T)} \mathcal{W}(G[\![V_t]\!])$ in G^+.

Now observe that as for each $W \in \mathcal{W}(G)$, $f(W) = W$, Claim 8.4 implies that $\mathcal{W}(G)$ is generated by $\bigcup_{t\in V(T)} f(\mathcal{W}(G[\![V_t]\!]))$ in G. $\Diamond$

As G is locally finite, note that for every pair $\{x, y\} \in A$, there are only finitely many paths of the form $\gamma \cdot P_j$ for some $(j, \gamma) \in [\ell] \times \Gamma$ having x and y as endpoints. For each $\{x, y\} \in A$, we let $\mathcal{P}_{x,y}$ denote the set of all such paths and $\mathcal{C}_{x,y}$ denote the set of all closed walks of the form $P \cdot P'^{-1}$ with $P, P' \in \mathcal{P}_{x,y}$. Then $\mathcal{C}_{x,y}$ is finite for each $\{x, y\} \in A$ and the set

$$\mathcal{C} := \bigcup_{\{x,y\}\in A} \mathcal{C}_{x,y}$$

is a Γ-invariant subset of $\mathcal{W}(G)$ with a finite number of Γ-orbits. We also consider the set of closed walks $\mathcal{C}'$ of $\mathcal{W}(G^+)$ of the form xPy for each $\{x, y\} \in A$ and $P \in \mathcal{P}_{x,y}$. Note that $f(\mathcal{C}') \subseteq \mathcal{C}$.

By Remark 3.3, $V(T)/\Gamma$ is finite. As for every $t \in V(T)$, $G[\![V_t]\!]$ is either finite or Γ_t-quasi-transitive planar, by [Ham18b, Theorem 25] the set $\mathcal{W}(G[\![V_t]\!])$ of closed walks of $G[\![V_t]\!] = G^+[V_t]$ has a generating set of cycles with finitely many Γ_t-orbits. We consider representatives $t_1, \dots, t_m$ of each of the finitely many orbits $V(T)/\Gamma$, and for each $i \in [m]$, we let $\mathcal{W}_i$ be a finite set of closed walks of $G[\![V_{t_i}]\!] = G^+[V_{t_i}]$ such that $\Gamma_{t_i} \cdot \mathcal{W}_i$ generates $\mathcal{W}(G[\![V_{t_i}]\!])$.

Claim 8.6. The set

$$\left(\bigcup_{i=1}^{m} f(\mathcal{W}_i)\right) \cdot \Gamma \cup \mathcal{C}$$

generates $\bigcup_{t\in V(T)} f(\mathcal{W}(G[\![V_t]\!]))$ in $\mathcal{W}(G)$.

Proof of the Claim: First, note that for each $i \in [m]$, Claim 8.4 implies that $f(\Gamma_{t_i} \cdot \mathcal{W}_i)$ generates $f(\mathcal{W}(G[\![V_{t_i}]\!]))$.

We first let $i \in [m]$ and show that $\Gamma_{t_i} \cdot f(\mathcal{W}_i) \cup \mathcal{C}$ generates $f(\Gamma_{t_i} \cdot \mathcal{W}_i)$. We let $W_i \in \mathcal{W}_i$ and $\gamma \in \Gamma_{t_i}$. One has to be careful as in general the walks $f(\gamma \cdot W_i)$ and $\gamma \cdot f(W_i)$ are not the same. Nevertheless we show that $f(\gamma \cdot W_i)$ is generated by $\gamma \cdot f(W_i)$ and by the walks of $\mathcal{C}$, which is enough to conclude. Indeed, it is not hard to see that $f(\gamma \cdot W_i)$ is generated in $\mathcal{W}(G^+)$ by $\gamma \cdot f(W_i)$ and by the walks $xPy \in \mathcal{C}'$ for each pair $\{x, y\} \in A$ of consecutive vertices of W_i, where $P \in \mathcal{P}_{x,y}$. Thus by Claim 8.4, $f(\gamma \cdot W_i) = f(f(\gamma \cdot W_i))$ is generated by $\gamma \cdot f(W_i) = f(\gamma \cdot f(W_i))$ and by the walks of $f(\mathcal{C}') \subseteq \mathcal{C}$.

To conclude with the proof of the claim we let $t \in V(T)$, and $(i, \gamma) \in [m] \times \Gamma$ be such that $t = \gamma \cdot t_i$. We let $W \in \mathcal{W}(G[\![V_t]\!])$. Then there exist $\gamma \in \Gamma$ and $W' \in \mathcal{W}(G[\![V_{t_i}]\!])$ such that $W = \gamma \cdot W'$. The exact same arguments as in the previous paragraph also apply to prove that $f(W)$ is generated by $\gamma \cdot f(W')$ together with the walks of $\mathcal{C}$. As we just proved above that $f(W')$ can be generated by finitely elements from $\Gamma_{t_i} \cdot f(\mathcal{W}_i) \cup \mathcal{C}$, we conclude that $f(W)$ can be generated by finitely elements from

$$\gamma \cdot (\Gamma_{t_i} \cdot f(\mathcal{W}_i) \cup \mathcal{C}) \cup \mathcal{C} \subseteq (\Gamma \cdot f(\mathcal{W}_i)) \cup \mathcal{C},$$

as desired. ◊

Combining Claims 8.5 and 8.6, we obtain that $\mathcal{W}(G)$ is generated by $\Gamma \cdot (\bigcup_{i=1}^{m} f(\mathcal{W}_i)) \cup \mathcal{C}$ in $\mathcal{W}(G)$. As $\mathcal{C}$ has a finite number of Γ-orbits, this concludes the proof of Theorem 8.3. □

8.4 Quasi-isometry to planar graphs

In a recent work, MacManus [Mac23, Corollary C] proved a structure theorem for quasi-transitive locally finite graphs that are quasi-isometric to a planar graph, namely that any such graph admits a canonical tree-decomposition of finite adhesion and tight edge-separations such that each torso is one-ended and quasi-isometric to a complete Riemannian plane. In particular it implies that such graphs must be accessible, which is in fact the main result from [Mac23].

Another consequence of Theorem 7.1 is that K_∞-minor-free locally finite quasi-transitive graphs form a (proper) subclass of the class of locally finite quasi-transitive graphs that are quasi-isometric to a planar graph.

Theorem 8.7. *Every locally finite quasi-transitive K_∞-minor free graph is quasi-isometric to a planar graph of bounded degree.*

We note that the result which allows us to construct the quasi-isometry using the canonical tree-decomposition was also proved independently by MacManus [Mac23] in a slightly different form (the "if" direction in his Corollary C). Our proof is very similar to his.

Proof. We assume that G is connected, since otherwise we can consider each connected component separately. By Theorem 7.1, G has a canonical tree-decomposition $(T, \mathcal{V})$ whose torsos $G[\![V_t]\!]$, $t \in V(T)$, are either planar or finite and whose adhesion sets have bounded size. For each $u \in V(G)$, we let T_u be the subtree of T with vertex set $\{t \in V(T), u \in V_t\}$. Note that as the edge-separations of $(T, \mathcal{V})$ are tight, Lemma 3.1 implies that T_u is finite for each u.

We let G' be the graph constructed as follow: for each $t \in V(T)$, we let V'_t be a copy of V_t and G'_t (with vertex set V'_t) be a copy of $G[\![V_t]\!]$ if V_t is infinite, or a spanning tree of $G[\![V_t]\!]$ if V_t is finite. For each $u \in V(G)$ and $t \in V(T_u)$, we let $u^{(t)}$ denote the copy of u in V'_t. We let $V(G') := \biguplus_{t \in V(T)} V'_t$. Now for every edge $st \in E(T)$, we choose an arbitrary vertex $u_{st} \in V_t \cap V_s$ (such a vertex exists, since G is assumed to be connected). We let:

$$E(G') := \left(\biguplus_{t \in V(T)} E(G'_t) \right) \uplus \{u_{st}^{(s)} u_{st}^{(t)}, st \in E(T)\}.$$

We also let T' be the 1-subdivision of T (the graph obtained from T by replacing each edge $e = st$ by a two-edge path s, t_e, t). Finally, for each $e = st \in E(T)$, we set

$V'_{t_e} := \{u_{st}^{(s)}, u_{st}^{(t)}\}$ and $\mathcal{V}' := (V'_t)_{t\in V(T')}$. We observe that by definition, $(T', \mathcal{V}')$ is a tree-decomposition of G' whose adhesion sets all have size 1. In particular, for every $t \in V(T')$, $G'[V'_t] = G'[\![V'_t]\!]$. We also note that by the definition of G' and Lemma 3.1, G' has bounded degree.

Claim 8.8. For every graph G, if G has a tree-decomposition $(T, \mathcal{V})$ such that every torso is planar and adhesion sets have size at most 1, then G is planar.

Proof of the Claim: If G contains K_5 or $K_{3,3}$ as a minor, then some torso of $(T, \mathcal{V})$ must also contain K_5 or $K_{3,3}$ as a minor, which is a contradiction. The result then follows from Wagner's theorem [Wag37] stating that any graph excluding K_5 and $K_{3,3}$ is planar. ◊

We now construct a quasi-isometry f from G to G'. For each $u \in V(G)$, we choose some $t_u \in V(T_u)$ and set $f(u) := u^{(t_u)}$. We also let $A_1 := \max\{\mathrm{diam}_G(V_t), V_t \text{ is finite}\}$, $A_2 := \max(1, \max\{|V_t|, V_t \text{ is finite}\})$ and $B := \max\{\mathrm{diam}_T(T_u), u \in V(T)\}$, which all exist by Lemma 3.1, as the edge-separations of $(T, \mathcal{V})$ are tight. We note that for each $t \in V(T)$ such that V_t is finite, since $G'[V'_t] = G'[\![V'_t]\!]$ is connected, its diameter is at most A_2.

We claim that the exact same arguments that in the proof of Lemma 3.5 to show that there exists a constant $C \geqslant 0$ such that for each $t \in V(T), u, v \in V_t$:

$$d_G(u, v) \leqslant C \cdot d_{G[\![V_t]\!]}(u, v),$$

the only difference being that the parts $G[V_t]$ of $(T, \mathcal{V})$ are not necessarily connected thus we have to consider distances in the connected graph G instead of $G[V_t]$.

We now show that there is a constant $\alpha > 0$ such that for every $u, v \in V(G)$ and every $f(u)f(v)$-path P' in G', there exists a uv-path P of size at most $\alpha \cdot |P'|$ in G. By taking P' to be a shortest path from $f(u)$ to $f(v)$ in G', this will imply in particular that $d_G(u, v) \leqslant \alpha \cdot d_{G'}(f(u), f(v))$.

Claim 8.9. For every $u, v \in V(G)$ and $t, s \in V(T)$ such that $u^{(t)}v^{(s)} \in E(G')$ we have

$$d_G(u, v) \leqslant \alpha := \max(A_1, C).$$

Proof of the Claim: Assume first that $s = t$. If V_t is finite, then $d_G(u, v) \leqslant A_1$. If V_t is infinite we must have $uv \in E(G[\![V_t]\!])$, and thus $d_G(u, v) \leqslant C$.

Assume now that $s \neq t$. Then by definition of G', we must have $st \in E(T)$ and $u = v$, and thus $d_G(u, v) = 0$. ◊

We now show that there exists a constant $\beta > 0$ such that for every $u, v \in V(G)$ and every uv-path P in G, there exists a $f(u)f(v)$-path P' of size at most $\beta|P|$ in G'. This directly implies that $d_{G'}(f(u), f(v)) \leqslant \beta \cdot d_G(u, v)$.

Claim 8.10. For every $u, v \in V(G)$ and $t, s \in V(T)$ such that $uv \in E(G)$, $u \in V_t$ and $v \in V_s$ we have:

$$d_{G'}(u^{(t)}, v^{(s)}) \leqslant \beta := (4A_2 + 2)B + A_2.$$

Proof of the Claim: First note that if V_t is finite, then for each $u, v \in V_t$ we have:

$$d_{G'}(u^{(t)}, v^{(t)}) = d_{G'[V'_t]}(u^{(t)}, v^{(t)}) \leqslant A_2.$$

If V_t is infinite, then for each $u, v \in V_t$ such that $uv \in E(G)$, we have $u^{(t)}v^{(t)} \in E(G')$ and thus $d_{G'}(u^{(t)}, v^{(t)}) \leqslant 1$. Since $A_2 \geqslant 1$, it follows that for each $t \in V(T)$ and $u, v \in V_t$ such that $uv \in E(G)$, we have

$$d_{G'}(u^{(t)}, v^{(t)}) \leqslant A_2. \tag{1.2}$$

Now let $u \in V(G)$ and $s, t \in V(T_u)$. We let $(s = t_0, t_1, \dots, t = t_\ell)$ be the shortest st-path in T. Note that it is also a path in T_u, hence $\ell \leqslant B$. Recall that in the construction of G', we have chosen for each edge $st \in E(T)$ a vertex $u_{st} \in V_s \cap V_t$ and we have added an edge in G' between $u_{st}^{(s)} \in V_s'$ and $u_{st}^{(t)} \in V_t'$. For each $i \in [\ell]$, we write $x_i := u_{t_{i-1}t_i} \in V_{t_{i-1}} \cap V_{t_i}$ for the sake of readability. Note that for each $i \in [\ell]$, x_i might be equal to u and that both u and x_i lie in the adhesion set $V_{t_{i-1}} \cap V_{t_i}$. This implies that $u^{(t_i)}$ and $x_i^{(t_i)}$ are adjacent in G' if V_{t_i} is infinite, and $d_{G'}(u^{(t_i)}, x_i^{(t_i)}) \leqslant A_2$ otherwise. So $d_{G'}(u^{(t_i)}, x_i^{(t_i)}) \leqslant A_2$ in both cases, and similarly $d_{G'}(u^{(t_{i-1})}, x_i^{(t_{i-1})}) \leqslant A_2$. It follows that for each $i \in [\ell]$, we have

$$d_{G'}(u^{(t_{i-1})}, u^{(t_i)}) \leqslant d_{G'}(u^{(t_{i-1})}, x_i^{(t_{i-1})}) + d_{G'}(x_i^{(t_{i-1})}, x_i^{(t_i)}) + d_{G'}(x_i^{(t_i)}, u^{(t_i)}) \leqslant 2A_2 + 1.$$

This implies that for every $u \in V(G)$ and $s, t \in V(T_u)$

$$d_{G'}(u^{(s)}, u^{(t)}) \leqslant (2A_2 + 1)B. \tag{1.3}$$

To conclude the proof of the claim, let $uv \in E(G)$. As $(T, \mathcal{V})$ is a tree-decomposition, there exists some $t \in V(T)$ such that $u, v \in V_t$. Then:

$$d_{G'}(f(u), f(v)) \leqslant d_{G'}(u^{(t_u)}, u^{(t)}) + d_{G'}(u^{(t)}, v^{(t)}) + d_{G'}(v^{(t)}, v^{(t_v)}),$$

thus by inequalities (1.2) and (1.3) we obtain $d_{G'}(f(u), f(v)) \leqslant (4A_2 + 2)B + A_2$. ◇

To prove that f is a quasi-isometry, it remains to prove that each $y \in V(G')$ is at bounded distance in G' from $f(V(G))$. For this, let $y \in V(G')$ and $t \in V(T), u \in V(G)$ be such that $y = u^{(t)}$. Then by inequality (1.3), $d_{G'}(y, f(u)) \leqslant (2A_2 + 1)B$ so f is indeed a quasi-isometry. This concludes the proof of Theorem 8.7. □

Note that the converse direction of Theorem 8.7 is wrong in general: consider the graph G obtained from the infinite square grid after adding the two diagonals in each square face. The graph G is locally finite, transitive, quasi-isometric to the infinite square grid and it is not hard to construct a model of K_∞ in G (note that it also follows from Theorem 7.7 that K_∞ is a minor of G).

We also observe that the hypothesis that G is quasi-transitive cannot be dropped in Theorem 8.7: consider the graph G obtained after taking for each k a copy $H_k := K_5^{(k)}$ of the k-subdivision of the complete graph K_5. If we want G to be connected, we might also add in G an infinite path intersecting exactly once each graph H_k. Then G has bounded degree, excludes K_6 as a minor and it can be shown that it is not quasi-isometric to a planar graph. A simple argument for this is that G contains K_5 as an *asymptotic minor* (see Section 9 for a definition), implying that every graph quasi-isometric to G must also contains K_5 as an asymptotic minor. In particular, G cannot be quasi-isometric to a planar graph.

It was proved by MacManus [Mac23] that if a finitely generated group has a Cayley graph which is quasi-isometric to a planar graph, then it is quasi-isometric to a planar

Cayley graph. A natural question is also to ask whether we can require the planar graph of bounded degree in Theorem 8.7 to be quasi-transitive, or even a Cayley graph. Our current proof does not preserve symmetries, as we do a number of non-canonical choices for the images of the vertices.

9 Beyond minor-exclusion

We present in this section structural graph properties, generalizing in different ways the property of excluding K_∞ as a minor. Results and proofs from this section mainly come from the paper [EG24a].

9.1 Quasi-isometries and asymptotic minors

We already saw in Theorem 8.7 that locally finite K_∞-minor free quasi-transitive graphs are quasi-isometric to some planar graph. This result can be seen as a special case of a more general recent conjecture of Georgakopoulos and Papasoglou [GP23] (see Conjecture 9.1 below). Before we state it, we first need to introduce some terminology.

We say that a graph H is a *k-fat minor* of a graph G if there exists a family of connected subsets $(M_v)_{v\in V(H)}$ of $V(G)$ such that

1. for each $u \neq v \in V(H)$, $d_G(M_u, M_v) \geqslant k$;
2. for each $e = uv \in E(H)$ there is a path P_e whose two endpoints lie in M_u and M_v and internal vertices are not in $\bigcup_{v\in V(H)} M_v$, and
3. for every $e \neq e' \in E(H)$, $d_G(P_e, P_{e'}) \geqslant k$ and for every $e = uv \in E(H)$ and $w \notin \{u, v\}$, $d_G(P_e, M_w) \geqslant k$.

A graph H is an *asymptotic minor* of G if for every $k \geqslant 0$, H is a k-fat minor of G. It was observed in [GP23] that for every finite graph H, the property of having H as an asymptotic minor is preserved under taking quasi-isometries. In particular, by Wagner theorem, it implies that every graph quasi-isometric to a planar graph must exclude K_5 and $K_{3,3}$ as an asymptotic minor. The converse implication was conjectured in general locally finite graphs in [GP23, Conjecture 9.1]. In the special case of quasi-transitive locally finite graphs, the authors conjectured the following.

Conjecture 9.1 (Conjecture 9.3 in [GP23])**.** *If G is locally finite, vertex-transitive and excludes some finite graph H as an asymptotic minor, then G is quasi-isometric to a planar graph.*

If G excludes a finite graph H as a minor, then G also excludes H as an asymptotic minor, and by the previous remark, every graph quasi-isometric to G also excludes H as an asymptotic minor. In particular MacManus [Mac23] made the following conjecture, which weakens Conjecture 9.1.

Conjecture 9.2 (Conjecture 8.2 in [Mac23])**.** *If G is a connected locally finite quasi-transitive graph, then G is quasi-isometric to a graph excluding some finite graph as a minor if and only if it is quasi-isometric to some planar graph of bounded degree.*

9.2 Local crossing number

A graph is *k-planar* if it has a drawing in the plane in which each edge is involved in at most k crossings (note that with this terminology, being planar is the same as being 0-planar). The *local crossing number* of a graph G, denoted by $\mathrm{lcr}(G)$, is the infimum integer k such that G is k-planar.

Georgakopoulos and Papasoglu raised the following problem.

Conjecture 9.3 (Problem 9.4 in [GP23]). *For any quasi-transitive graph G of bounded degree, G is quasi-isometric to a planar graph if and only if G has finite local crossing number.*

In this section we prove that the property of having finite local crossing number is preserved under taking quasi-isometries.

Theorem 9.4. *Let G be a graph of bounded degree which is quasi-isometric to a graph H of finite local crossing number. Then G also has finite local crossing number.*

In the particular case $k = 0$, we immediately obtain the "only if" direction of Conjecture 9.3 (we recently learned from Agelos Georgakopoulos that he also proved the case $k = 0$ independently). In particular, in the case $k = 0$, we obtain the following immediate consequence of Theorem 8.7.

Corollary 9.5. *Every locally finite quasi-transitive graph G which is K_∞-minor-free has finite local crossing number.*

The assumption that G is locally finite is necessary, as shown by the graph obtained from the square grid by adding a universal vertex (this graph is K_6-minor free, but is not k-planar for any $k < \infty$). The assumption that G is quasi-transitive is also crucial: consider for each integer ℓ a graph G_ℓ obtained from the square grid by adding an edge between two vertices at distance ℓ in the grid (if G_ℓ is k-planar then $k = \Omega(\ell)$), and take the disjoint union of all graphs G_ℓ, $\ell \in \mathbb{N}$.

We now prove Theorem 8.7. Note that in general, the property of being locally finite, or even of having countably many vertices is not preserved under quasi-isometry. The next lemma will be useful to make sure that we can restrict ourselves to locally finite graphs in the remainder of the proof.

Lemma 9.6. *Let G be a graph of bounded degree which is quasi-isometric to a graph H. Then G is quasi-isometric to a subgraph H' of H of bounded degree.*

Proof. We let $f : V(G) \to V(H)$ and $A \geqslant 1$ be such that for each $x, x' \in V(G)$:

$$\frac{1}{A} \cdot d_G(x, x') - A \leqslant d_H(f(x), f(x')) \leqslant A \cdot d_G(x, x') + A,$$

and such that the A-neighborhood of $f(V(G))$ covers H. We also let $\Delta \in \mathbb{N}$ denote the maximum degree of G. Note that for each $xy \in E(G)$, there exists a $f(x)f(y)$-path P_{xy} in

H such that $|P_{xy}| \leqslant A^2$. We let H' be the subgraph of H given by the union of all such paths P_{xy}.

We first observe that H' is quasi-isometric to G, and that f gives the corresponding quasi-isometric embedding. Note that for every $z \in V(H')$, by construction there must be some edge $xy \in E(G)$ such that $z \in P_{xy}$. In particular, $d_{H'}(z, f(x)) \leqslant A^2$. Note that by construction we clearly have $d_{H'}(f(x), f(y)) \leqslant A^2 d_G(x, y)$ for each $x, y \in V(G)$, and as $d_{H'}(f(x), f(y)) \geqslant d_H(f(x), f(y))$, f induces indeed a quasi-isometric embedding between G and H'.

Now we show that H' has bounded degree. Let $z \in V(H')$ and $xy \in E(G)$ such that $z \in V(P_{xy})$. Then $d_{H'}(z, f(x)) \leqslant A^2$ so $X := \{x \in V(G), z \in V(P_{xy})\}$ has diameter at most $2A^2$ in G. Note that as G degree at most Δ, we have $|X| \leqslant \Delta^{2A^2}$. In particular it implies that H' has degree at most Δ^{2A^2}. □

Given a graph G and an integer $k \geqslant 1$, the *k-th power* of G, denoted by G^k, is the graph with the same vertex set as G in which two vertices are adjacent if and only if they are at distance at most k in G. The *k-blow up* of G, denoted by $G \boxtimes K_k$, is the graph obtained from G by replacing each vertex u by a copy C_u of the complete graph K_k, and by adding all edges between pairs C_u, C_v if and only if u and v are adjacent in G (so that each edge of G is replaced by a complete bipartite graph $K_{k,k}$ in $G \boxtimes K_k$). Quasi-isometries of bounded degree graphs are related to graph powers and blow-ups by the following lemma.

Lemma 9.7. *Let H be a graph, and let G be a graph of degree at most $\Delta \in \mathbb{N}$ which is quasi-isometric to H. Then there is an integer k such that G is a subgraph of $H^k \boxtimes K_k$.*

Proof. We let $A \geqslant 1$ and $f : V(G) \to V(H)$ be such that for each $x, x' \in V(G)$:

$$\frac{1}{A} \cdot d_G(x, x') - A \leqslant d_H(f(x), f(x')) \leqslant A \cdot d_G(x, x') + A,$$

and such that the A-neighborhood of $f(V(G))$ covers H. Note that for each $x, x' \in V(G)$ such that $f(x) = f(x') = y$ we must have $d_G(x, x') \leqslant A^2$, hence

$$|f^{-1}(y)| \leqslant B := \Delta^{A^2}$$

for every $y \in V(H)$. We now show that G is a subgraph of $H' := H^{2A} \boxtimes K_B$, which implies the lemma for $k := \max(2A, B)$.

As in the definition of a blow-up, for each $v \in V(H)$ we denote by C_v the associated clique of size B in H'. For every $v \in V(H)$ we fix an arbitrary injection $g_v : f^{-1}(v) \to C_v$, and define an injective mapping $g : V(G) \to V(H')$ by letting $g(x) := g_{f(x)}(x)$ for each $x \in V(G)$. In other words every two vertices of G having the same image v by f are sent by g in the same clique C_v in H'. By construction g is injective, so we just need to check that it defines a graph homomorphism to conclude that G is a subgraph of H'. Let $xy \in E(G)$. Then $d_H(f(x), f(y)) \leqslant 2A$ so in particular every vertex in $V_{f(x)}$ is at distance at most $2A$ to every vertex in $V_{f(y)}$ in H'. In particular, this means that $g(x)g(y) \in E(H')$, as desired. □

The next observation will allow us to slightly simplify the statement of Lemma 9.7.

Observation 9.8. *For every graph H of bounded maximum degree, and every integer k, there exists a graph G of bounded maximum degree such that $\mathrm{lcr}(G) = \mathrm{lcr}(H)$ and $H^k \boxtimes K_k$ is a subgraph of G^{k+2}.*

Proof. Let G be the graph obtained from H by attaching to each vertex k pendant vertices of degree 1. Note that the graph $H^k \boxtimes K_k$ is a subgraph of G^{k+2}. To see this, one can bijectively map each clique C_v of H for $v \in V(H)$ to the k pendant vertices we attached to v in G, and observe that it gives an isomorphism between the graph induced by these vertices in G^3, and $H \boxtimes K_k$. Since adding pendant vertices does not change the local crossing number, $\mathrm{lcr}(G) = \mathrm{lcr}(H)$. Moreover G has bounded maximum degree if and only if H has bounded maximum degree. □

We can now combine the results above to deduce the following corollary.

Corollary 9.9. *If a graph G with bounded degree is quasi-isometric to a graph of bounded local crossing number, then there exists a planar graph H of bounded maximum degree and an integer k, such that G is a subgraph of H^k.*

Proof. By Lemmas 9.6 and 9.7, there is a graph H_1 of bounded local crossing number and maximum degree and an integer ℓ such that G is a subgraph of $H_1^\ell \boxtimes K_\ell$. By Observation 9.8, there is a graph H_2 of bounded local crossing number and maximum degree such that G is a subgraph of $H_2^{\ell+2}$. Observe that every s-planar graph F_1 is a subgraph of F_2^{s+1}, where F_2 is the planar graph obtained from F_1 by placing a new vertex at each crossing (and note that if F_1 has bounded degree, then F_2 also has bounded degree). It follows that there is a planar graph H of bounded degree and an integer k, such that G is a subgraph of H^k. □

We now prove that bounded powers of planar graphs of bounded degree are ℓ-planar for some ℓ. This was proved for finite graphs in [DMW23, Lemma 12].

Lemma 9.10 ([DMW23])**.** *Let H be a finite planar graph of maximum degree at most Δ and let G be a subgraph of H^k, for some integer k. Then G is ℓ-planar, for $\ell := 2k(k+1)\Delta^k$.*

However, we need a version of Lemma 9.10 for infinite locally finite graphs. The first option is to simply follow the proof of [DMW23], which starts with a planar drawing of G, and adds for any path P of length at most k, an edge between the endpoints of P, drawn in a close neighborhood around P. However in the locally finite case this approach requires that the original planar drawing has the property that every edge has a small neighborhood which does not intersect any other vertices or edges of the graph. Such a drawing always exists but it requires a little bit of work. So instead, we chose to extend Lemma 9.10 to infinite locally finite graphs using a simple compactness argument.

Lemma 9.11. *Let G be a locally finite graph. If there is an integer ℓ such that all finite induced subgraphs of G are ℓ-planar, then G is also ℓ-planar.*

Proof. We first observe that any ℓ-planar embedding of a graph H can be described combinatorially, by considering the planar graph H^+ obtained from H by replacing all crossings by new vertices. The corresponding planar embedding of H^+ can be completely described (up to homeomorphism) by its rotation system (the clockwise cyclic ordering of the neighbors around each vertex), and there are only finitely many such rotation systems if H (and thus H^+) is finite.

We are now ready to prove the lemma. We can assume that G is connected, since otherwise we can consider each connected component independently. Since G is locally finite and connected, it is countable and we can write $V(G) = \{v_1, v_2, \ldots\}$. We define a rooted tree T as follows. The root of T is the unique ℓ-planar embedding of $G[\{v_1\}]$, up to homeomorphism. For every $k \geqslant 1$, and any ℓ-planar embedding of $G_k := G[\{v_1, \ldots, v_k\}]$ we add a node in the tree and connect it to the node corresponding to the resulting ℓ-planar embedding of $G_{k-1} = G_k - v_k$ (the embedding obtained by deleting v_k in the embedding of G_k). The resulting tree T is infinite (since every graph G_k is ℓ-planar by assumption), and locally finite (since every graph G_k has only finitely many different ℓ-planar embeddings, up to homeomorphism). By König's infinity lemma [Kön27] (or by repeated applications of the pigeonhole principle), T has an infinite path starting at the root. This infinite path corresponds to a sequence of ℓ-planar embeddings of G_k, $k \geqslant 0$, with the property that for every $k \geqslant 0$, the ℓ-planar embedding of G_k can be obtained from the ℓ-planar embedding of G_{k+1} by deleting v_{k+1} (and all edges incident to v_{k+1}). By taking the union of all the ℓ-planar embeddings of G_k, $k \geqslant 0$, we thus obtain an ℓ-planar embedding of G, as desired. □

We obtain the following as a direct consequence.

Corollary 9.12. *Let H be a locally finite planar graph of maximum degree at most Δ and let G be a subgraph of H^k, for some integer k. Then G is ℓ-planar, for $\ell := 2k(k+1)\Delta^k$.*

Proof. Let X be a finite subset of $V(G) \subseteq V(H)$ and for any pair $x, x' \in X$ with $d_H(x, x') \leqslant k$, consider a path $P_{x,x'}$ of length at most k between x and x' in H. Let Y be the union of X and the vertex sets of all the paths $P_{x,x'}$ defined above. Then $H[Y]$ is a finite planar graph, and $G[X]$, the finite subgraph of G induced by X, is a subgraph of $H[Y]^k$. By Lemma 9.10, $G[X]$ is ℓ-planar with $\ell := 2k(k+1)\Delta^k$. Since this holds for any finite set X, it follows from Lemma 9.11 that G itself is ℓ-planar, as desired. □

Theorem 9.4 is now a direct consequence of Corollary 9.9 and Corollary 9.12. By combining Theorems 8.7 and 9.4, we then immediately deduce Corollary 9.5.

9.3 Assouad-Nagata dimension

Let (X, d) be a metric space, and let $\mathcal{U}$ be a family of subsets of X. We say that $\mathcal{U}$ is *D-bounded* if each set $U \in \mathcal{U}$ has diameter at most D. We say that $\mathcal{U}$ is *r-disjoint* if for any a, b belonging to different elements of $\mathcal{U}$ we have $d(a, b) > r$.

We say that $D_X : \mathbb{R}^+ \to \mathbb{R}^+$ is an *n-dimensional control function* for (X, d) if for any $r > 0$, (X, d) has a cover $\mathcal{U} = \bigcup_{i=1}^{n+1} \mathcal{U}_i$, such that each $\mathcal{U}_i$ is r-disjoint and each element of $\mathcal{U}$ is $D_X(r)$-bounded. A control function D_X for a metric space X is said to be a *dilation* if there is a constant $c > 0$ such that $D_X(r) \leqslant cr$, for any $r > 0$.

The *asymptotic dimension* of (X, d), introduced by Gromov in [Gro93], is the least integer n such that (X, d) has an n-dimensional control function. If no such integer n

exists, then the asymptotic dimension is infinite. The *Assouad-Nagata dimension* of (X, d), introduced by Assouad in [Ass82], is the least n such that (X, d) has an n-dimensional control function which is a dilation. Clearly the asymptotic dimension is at most the Assouad-Nagata dimension.

It was proved in [BBE+20] that every *bounded degree* graph excluding a minor has asymptotic dimension at most 2, and that any planar graph has asymptotic dimension at most 2. This was improved in [BBE+23], where it was shown that any graph excluding a minor has asymptotic dimension at most 2, and any planar graph has Assouad-Nagata dimension at most 2. This was finally extended by Liu in [Liu23], who proved that any graph avoiding a minor has Assouad-Nagata dimension at most 2 (a different proof was then given by Distel in [Dis23]). All the results on graphs excluding a minor mentioned above (even for bounded degree graphs) crucially rely on the original Graph minor structure theorem of Robertson and Seymour [RS03] we mentioned at the beginning of Section 7.

Using the invariance of Assouad-Nagata dimension under bilipschitz embedding [LS05], we give in this section a short proof of the fact that minor-excluded quasi-transitive graphs have Assouad-Nagata dimension at most 2. An immediate consequence is that *minor-excluded finitely generated groups* (see Chapter 2 for a definition) have Assouad-Nagata dimension at most 2, and thus asymptotic dimension at most 2 (which was originally conjectured by Ostrovskii and Rosenthal in [OR15]). We only use Theorem 8.7 and a few simple tools from [BBE+20, BBE+23] based on the work of Brodskiy, Dydak, Levin and Mitra [BDLM08]. In particular we give a short proof of the fact that planar graphs of bounded degree have Assouad-Nagata dimension at most 2. Crucially, our proof for quasi-transitive graphs excluding only uses Theorem 7.1, and thus *does not rely* on the Graph minor structure theorem of Robertson and Seymour.

Recall that if some graph G admits a quasi-isometric embedding in a graph H, then there is a constant $C \geqslant 1$ and a bilipschitz embedding of G into H^{+C}, the graph obtained from H by adding C pendant vertices to each vertex of H. In particular, if H is planar with bounded degree, then so is H^{+C}. Hence, Theorem 8.7 has the following immediate consequence.

Corollary 9.13. *Every locally finite quasi-transitive K_∞-minor free graph has a bilipschitz embedding in a planar graph of bounded degree.*

We first prove that planar graphs of bounded degree have Assouad-Nagata dimension at most 2 (a stronger version of this result, without the bounded degree assumption, was proved in [BBE+23]).

We need the following result, first proved in [BST12] in a slightly different form. Another proof can be found in [BBE+20, BBE+23] based on a result of Ding and Oporowski [DO95] which states that every graph G of treewidth at most t and maximum degree at most Δ is a subgraph of the strong product of a tree with a complete graph on $24t\Delta$ vertices, i.e., G has tree-partition-width at most $24t\Delta$. The proofs of all these results are fairly short.

Theorem 9.14 ([BST12]). *If a graph G has bounded degree and bounded treewidth, then G has Assouad-Nagata dimension at most 1.*

A *layering* of a graph G is a partition of the vertex set of G into sets $L_0, L_1, \ldots$, called *layers*, so that any pair of adjacent vertices in G either lies in the same layer or in consecutive layers (i.e. layers L_i, L_{i+1} for some $i \geqslant 1$). A simple example of layering is given by a *BFS-layering* of G, obtained by choosing one root vertex v_C in each connected component C of G, and then defining L_i (for all $i \geqslant 0$) as the set of vertices of G at distance exactly i from one of the vertices v_C.

It was proved by Bodlaender [Bod88] that planar graphs of bounded diameter have bounded treewidth. This directly implies the following.

Lemma 9.15. *For any BFS-layering of a planar graph G, and any integer k, the subgraph of G induced by k consecutive layers of $\mathcal{L}$ has bounded treewidth.*

Proof. Let $L_0, L_1, \ldots$ be a BFS-layering of G. Consider the planar subgraph H of G induced by k consecutive layers $L_i, L_{i+1}, \ldots, L_{i+k-1}$. If $i = 0$, then H is a disjoint union of graphs of radius at most k, and thus has bounded treewidth by Bodlaender's result [Bod88]. Assume now that $i \geqslant 1$, and let H^+ be the supergraph of H obtained by adding a vertex r that dominates all the vertices of L_i. Note that H^+ can be obtained from the subgraph of G induced by the layers $L_0, L_1, \ldots, L_{i+k-1}$, by contracting all layers $L_0, L_1, \ldots, L_{i-1}$ (which induce a connected subgraph of G, by definition of a BFS-layering) into a single vertex. Thus H^+ is a minor of G, which implies that H^+ is planar. Moreover, it follows from the definition of a BFS-layering that each vertex of H^+ lies at distance at most k from r, and thus H^+ has diameter at most $2k$. By Bodlaender's result [Bod88], H^+ has bounded treewidth, and thus H (as a subgraph of H^+) also has bounded treewidth. □

The next result appears as Theorem 4.3 in [BBE+23], and is a simple application of the main result in [BDLM08] (which has a nice and short combinatorial proof).

Theorem 9.16 ([BBE+23])**.** *If a graph G has a layering $\mathcal{L} = (L_0, L_1, \ldots)$ such that for any integer k, the disjoint union of all subgraphs of G induced by k consecutive layers of $\mathcal{L}$ has Assouad-Nagata dimension at most n, then G has Assouad-Nagata dimension at most $n+1$.*

We immediately deduce the following.

Corollary 9.17. *Every planar graph G of bounded degree has Assouad-Nagata dimension at most 2.*

Proof. Consider a BFS-layering $\mathcal{L}$ of G. By Lemma 9.15, for any $k \geqslant 1$, the disjoint union of all subgraphs of G induced by k consecutive layers of $\mathcal{L}$ has bounded treewidth. As G has bounded degree, this disjoint union of subgraphs of G also has bounded degree, and thus by Theorem 9.14 it has Assouad-Nagata dimension at most 1. Hence, it follows from Theorem 9.16 that G itself has Assouad-Nagata dimension at most 2. □

We are now ready to prove the main result of this section. Recall that a stronger version (without the quasi-transitivity assumption) was proved in [Liu23, Dis23], but the version below has a reasonably simple proof that does not rely on the Robertson-Seymour graph minor structure theorem.

Theorem 9.18. *Every locally finite quasi-transitive K_∞-minor free graph has Assouad-Nagata dimension at most 2.*

Proof. By Corollary 9.13, every locally finite quasi-transitive K_∞-minor free graph G has a bilipschitz embedding in some planar graph H of bounded degree, which has Assouad-Nagata dimension at most 2. Since Assouad-Nagata dimension is invariant under bilipschitz embedding [LS05], G has Assouad-Nagata dimension at most 2. □

9.4 Open problems

We present in this section directions and problems about the structure of locally finite quasi-transitive graphs described by properties generalizing minor-exclusion.

Quasi-isometries to quasi-transitive planar graphs. It was proved by MacManus [Mac23] that if a finitely generated group has a Cayley graph which is quasi-isometric to a planar graph, then it is quasi-isometric to a planar Cayley graph (see Section 12 for a definition of a Cayley graph). In the same spirit, it is natural to ask whether we can require the planar graph of bounded degree in Theorem 8.7 to be quasi-transitive, or even a Cayley graph. Our current proof does not preserve symmetries, as we do a number of non-canonical choices for the images of the vertices. As remarked by a referee from our paper [EG24a], the stronger question above, whether we can require the planar graph in Theorem 8.7 to be a Cayley graph, is a special case of a Problem of Woess [Woe91, Problem 1], which asked whether every transitive graph is quasi-isometric to a Cayley graph. This turned out to have a negative answer [EFW12] in general, but the question restricted to (quasi-)transitive graphs excluding a minor might still have a positive answer.

A finite list of obstructions. In view of MacManus' characterisation [Mac23, Corollary C] of quasi-transitive graphs which are quasi-isometric to a planar graph, it would be interesting to also find a characterisation in terms of obstructions. Recall the conjecture from [GP23] that a graph is quasi-isometric to a planar graph if and only if it does not contain K_5 or $K_{3,3}$ as an asymptotic minor.

Examples of quasi-transitive graphs that are not quasi isometric to any planar graph include Cayley graphs of a group of Assouad-Nagata dimension at least 3, for instance any grid in dimension 3. This rules out any generalization of Theorem 8.7 using classes of polynomial growth or expansion. This example also shows that we cannot extend Theorem 8.7 to all families of bounded queue-number or stack-number.

Here is perhaps a more interesting example. The *strong product* $G \boxtimes H$ of two graphs G and H has vertex set $V(G) \times V(H)$, and two distinct vertices (u, x) and (v, y) are adjacent if and only if ($u = v$ or $uv \in E(G)$) and ($x = y$ or $xy \in E(H)$). Consider the strong product $T \boxtimes P$ of the infinite binary tree T and the infinite 2-way path P. Using Theorems 9.14 and Lemma 9.16, this graph has Assouad-Nagata dimension at most 2. As it contains a quasi-isometric copy of a 2-dimensional grid, the Assouad-Nagata dimension of $T \boxtimes P$ is indeed equal to 2. On the other hand, we observe that for any integer k, $T \boxtimes P$ contains the complete bipartite graph $K_{k,k}$ as an asymptotic minor. To see this, remark that T contains an infinite k-claw (the graph obtained by gluing k infinite 1-way paths at their starting vertex) as an asymptotic minor (obtained by contracting a subtree of T with k leaves into a single vertex, and pruning the additional branches). The strong product of this infinite k-claw with P consists of k copies of an infinite 2-dimensional grid (restricted to the upper half-plane, say), glued on a common infinite path π. On this path we can select k vertices, arbitrarily far apart, and on each infinite grid we can select a single vertex, arbitrarily far from π, and connect it to the k vertices of π via disjoint paths. By taking a ball of sufficiently large radius around each of the $2k$ vertices, we obtain $K_{k,k}$ as an r-fat minor for arbitrarily large r, and thus $K_{k,k}$ as an asymptotic minor. This is illustrated for $k = 3$ in Figure 1.7. Since containing a graph H as an asymptotic minor is invariant under quasi-isometry, any graph G which is quasi-isometric to $T \boxtimes P$ also contains every $K_{k,k}$ as an asymptotic minor, and thus G cannot be planar.

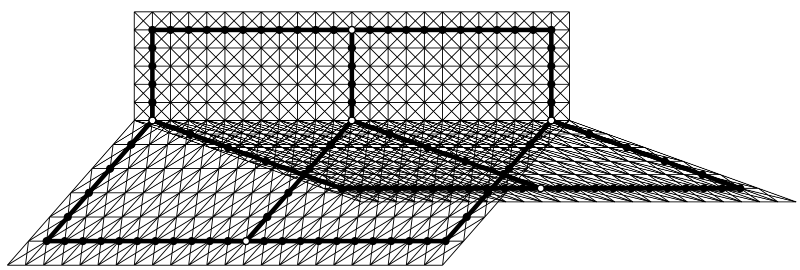

Figure 1.7: A fat $K_{3,3}$-minor in $T \boxtimes P$.

Recall that any graph excluding a minor has Assouad-Nagata dimension at most 2 [Liu23, Dis23]. It is natural to wonder whether some sort of converse holds, that is whether any graph of Assoud-Nagata dimension at most 2 is quasi-isometric to a graph excluding a minor (this would be a natural extension of Theorem 8.7). The example above shows that this is false, even for vertex-transitive graphs.

As explained above, it was conjectured in [GP23] that graphs that are quasi-isometric to a planar graph can be characterized by a finite list of forbidden asymptotic minors. A natural question is whether this can be replaced by a finite list of forbidden quasi-isometrically embedded subgraphs, at least in the case of quasi-transitive graphs. We do not have a good candidate for such a finite list, but it should contain at least the two examples mentioned above: 3-dimensional grids and the product of the binary tree with a path. One difficulty is that no such list is even known (or conjectured to exist) for Assouad-Nagata dimension at most 2, or asymptotic dimension at most 2.

k-planar graphs. First, we observe that Conjecture 9.3 reduces to the case $k = 1$, i.e., that it is equivalent to the following.

Conjecture 9.19. *Every quasi-transitive 1-planar graph of bounded degree is quasi-isometric to a planar graph.*

To see that the case $\mathrm{lcr} \geqslant 2$ reduces to the case $\mathrm{lcr} = 1$, observe that for every k-planar graph G with $k \geqslant 2$, its $(k-1)$-subdivision $G^{(k-1)}$ (the graph obtained from G by replacing every edge by a path on k edges) is locally finite, quasi-transitive, quasi-isometric to G, and 1-planar. To see the last point, consider any embedding of G in $\mathbb{R}^2$ in which every edge is involved in at most k crossings and assume without loss of generality that the crossing points between every two edges are all pairwise distinct. Then for every edge $e \in E(G)$, one can add the $k-1$ corresponding vertices of $G^{(k-1)}$ subdividing e in the drawing by putting at least one vertex on each of the curves connecting two consecutive crossing points of e with other edges.

We note that Conjecture 9.19 (and thus Conjecture 9.3) would be a direct consequence of the following.

Conjecture 9.20. *Let G be a quasi-transitive 1-planar graph of bounded degree. Then there is an integer k and an embedding of G in the plane with at most 1 crossing per edge such that for every pair of crossing edges uv, xy in G, we have $d_G(u, x) \leqslant k$.*

In an early version of the paper [EG24a] we were conjecturing something stronger, namely that for any embedding of G with at most 1 crossing per edge, there is an integer k such that all pairs of crossing edges lie at distance at most k in G. But this is false (as shown by the two-way infinite path, drawn in such a way that it self-intersects at more and more distant points).

In this section, we have mainly considered graphs with finite local crossing number. A natural generalization is the following: a graph is $(<\omega)$-*planar* if it has a drawing in the plane in which each edge is involved in finitely many crossings. In the paper [EG24a], we also asked if for quasi-transitive locally finite graphs, ω-planarity is equivalent to the property of having finite local crossing number. It was observed since by Kolja Knauer (personal communication) that the answer is no, as *any* infinite locally finite graph G is $(<\omega)$-planar. To see this, consider an ordering $v_1, v_2, \dots$ of $V(G)$, embed each vertex v_i in the coordinate (i, i^2) in the plane, and each edge v_iv_j as a segment joining v_i and v_j. Note that by convexity of the function $x \mapsto x^2$, every edge crossing an edge v_iv_j must have an endpoint v_k with $i < k < j$. As for every pair $i < j$ there are only finitely many such vertices v_k and each of them has finite degree, the edge v_iv_j is crossed by only finitely many other edges.

As there exist quasi-transitive graphs of bounded degree that have unbounded local crossing number (the 3-dimensional grid for instance, see [DEW17]), the paragraph above implies that our initial question has a negative answer.

Twin-width. Twin-width is a graph parameter that was discovered by Edouard Bonnet, Eun Jung Kim, Stéphan Thomassé and Rémi Watrigant and introduced in [BKTW22] for finite graphs, generalizing and unifying many results in the domain of fixed parameter algorithm, and also having interesting and deep connections with other fields such as model theory, enumerative combinatorics or algebra (see the twin-width series of papers [BKTW22, BGK+22, BGK+21, BGOdM+22, BGdMT23, BKRT22, BGTT22, BCK+22]). In particular, in [BGTT22], the authors gave a definition of twin-width for infinite graphs of bounded degree, and proved that the property of having bounded twin-width is invariant under taking quasi-isometries. Even though there exist locally finite quasi-transitive graphs (in fact Cayley graphs) with infinite twin-width [BGTT22], the known proof of the existence of such graphs is highly non-trivial and relies on a construction of Osajda [Osa20] of Gromov monster groups (see Section 17 for more details about the construction of such groups). It was also proved in [BGTT22, Proposition 4.2] that finite quasi-transitive graphs admitting an hyperbolic metric have bounded twin-width. Hence the class of locally finite quasi-transitive graphs of bounded twin-width is very rich and interesting from a metric perspective.

The definition of twin-width for finite graph crucially depends on the notion of contraction sequences (see [BKTW22] for a definition): a finite graph has twin-width at most $d \in \mathbb{N}$ if there exists a *contraction sequence* of width at most d. Many other equivalent witnesses of bounded twin-width exist both in finite graphs and locally finite graphs. In particular, in analogy to tree-decompositions of bounded width, a notion of *twin-decomposition* with an underlying tree structure was introduced in [BNdM$^+$24, BGdMT23] as another certificate that a finite graph has bounded twin-width. However, none of the known certificates for bounded twin-width is canonical, in the sense that the automorphisms of the initial graph do not induce as automorphisms of the structures of the certificates of bounded twin-width.

If we want to obtain decomposition results in the spirit Theorem 4.3 (*iii*) for graphs of bounded twin-width, or to decide if such graphs satisfy some of the properties we studied in the previous sections like vertex-accessibility, it is likely that such a canonical certificate of twin-width boundedness would be helpful. Colin Geniet asked explicitly this question during the workshop LoGAlg 2023 in Warsaw.

10 Symmetric proper colorings

We conclude Chapter 1 with a discussion about questions we asked at the end of the paper [EGLD23] related to proper colorings in locally finite quasi-transitive graphs. This provides a perfect transition with Chapter 2, as these problems are related in some way with the study of aperiodicity in subshifts of finite types of finitely generated groups (see Sections 14, 15 and 16). The results and proofs presented in this section come from a joint work in preparation with Tara Abrishami and Louis Esperet.

Recall that a crucial component of the proof of Theorem 7.1 was to adapt a strategy of Grohe [Gro16] to obtain a canonical tree-decomposition up to the contraction of a matching (see also Theorem 6.26). At one point we tried to figure out whether the following was true (in the end it turned out that we did not need to answer these questions, but we believe they might be of independent interest).

Problem 10.1. *Let G be a locally finite quasi-transitive graph. Is there a proper vertex-coloring of G with a finite number of colors such that the colored graph G itself is quasi-transitive (where automorphisms have to preserve the colors of the vertices)?*

Problem 10.2. *Let G be a locally finite quasi-transitive graph. Is there an orientation of the edges of G such that the oriented graph G itself is quasi-transitive (where automorphisms have to preserve the orientation of the edges)?*

We call any coloring c of a graph G such that the colored graph (G, c) is quasi-transitive a *strongly periodic coloring.*

Note that we can also consider the variant of Problem 10.1 with respect to proper edge-colorings. As for any locally finite Γ-quasi-transitive graph, Γ also induces a quasi-transitive group action on $E(G)$, the line graph of G must also be locally finite quasi-transitive, hence this variant correponds to a special case of Problem 10.1.

An example showing that Problem 10.1 has a negative answer was recently constructed by Hamann and Möller (personal communication). It is mainly based on the existence of infinite simple groups which are finitely generated. Together with Tara Abrishami and Louis Esperet, we then observed that a variant of this example could be used to provide a negative answer to Problem 10.2 as well. This second variant was also discovered independently from our work by Norin and Przytycki (personal communication), who furthermore showed that the examples can be chosen to be Cayley graphs (and thus vertex-transitive), rather than merely quasi-transitive.

It remains an interesting problem to understand for which graphs the problems might still have a positive answer. Note that a proper vertex-coloring satisfying the properties of Problem 10.1 also implies a positive answer to Problem 10.2, as seen by choosing a total order on the colors and then orienting each edge from the extremity with the smaller color to the extremity with the larger color (any color-preserving automorphism is then also orientation-preserving). Another simple observation is that Problem 10.1 admits a positive answer when considering bipartite connected graph.

Symmetric proper colorings of 2-ended graphs. We give there a simple proof that Problems 10.1 and 10.2 admit positive answers for 2-ended locally finite quasi-transitive graphs. More precisely, we even show the existence of a strongly periodic proper vertex-coloring using exactly $\chi(G)$ colors. The main motivation for considering this case comes from symbolic dynamics and a question of Carroll and Penland [CP15] (see Conjecture 15.7) and we will present in Section 16 a proof of a more general result on 2-ended quasi-transitive graphs based on a reduction to tiling problems in groups.

Theorem 10.3. *Let G be a connected locally finite Γ-quasi-transitive graph with 2 ends. Then, G has a proper vertex-coloring with $\chi(G)$ colors and such that there exists a cyclic subgroup $\Gamma' \subseteq \Gamma$ such that Γ' preserves vertex colors and acts quasi-transitively on G.*

Proof. We let (Y, S, Z) and γ_0 be given by Lemma 4.2. For each $i \in \mathbb{Z}$, let $(Y_i, S_i, Z_i) := \gamma_0^i \cdot (Y, S, Z)$. Then $S_j \cup Z_j \subseteq Z_i$ for all $i < j$ and (Y_i, S_i, Z_i) also separates the two ends of G. Let $c : V(G) \to [\chi(G)]$ be a proper vertex-coloring of G. By the pigeonhole principle there exists $i < j$ such that $c(\gamma_0^i \cdot x) = c(\gamma_0^j \cdot x)$ for all $x \in S$. Up to replacing γ_0 by γ_0^{j-i}, we may assume that $i = 0$ and $j = 1$, i.e., that $c(\gamma_0 \cdot x) = c(x)$ for all $x \in S$.

For all $i \in \mathbb{Z}$, we let $V_i := V(G) \setminus (Y_i \cup S_{i+1} \cup Z_{i+1})$. Then for each $i \in \mathbb{Z}$, $S_i \subseteq V_i$, $V_{i+1} = \gamma_0 \cdot V_i$ and as G has two ends and bounded degree, the graph $G_i := G[V_i]$ is finite. Moreover, note that $\{V_i : i \in \mathbb{Z}\}$ is a partition of $V(G)$.

We now define a vertex-coloring $\widetilde{c} : V(G) \to [\chi(G)]$ by setting for each $i \in \mathbb{Z}$ and $v \in V_i$, $\widetilde{c}(v) := c(\gamma_0^{-i} \cdot v)$. In other words, the vertex-coloring $\widetilde{c}$ is obtained after repeating periodically $c|_{V_0}$ on each G_i. First, note that $\widetilde{c}$ is well-defined on $V(G)$ as for every $v \in V(G)$

there exists a unique $i \in \mathbb{Z}$ such that $v \in V_i$ and then $\gamma_0^{-i} \cdot v \in V_0$. Moreover, by definition the action of γ_0 on G preserves the colors of $\widetilde{c}$, i.e., $\widetilde{c}(\gamma_0 \cdot v) = \widetilde{c}(v)$ for all $v \in V(G)$.

We show that $\widetilde{c}$ a proper vertex-coloring. Let $uv \in E(G)$ and $i, j \in \mathbb{Z}$ be such that $u \in V_i$ and $v \in V_j$. For each $i \in \mathbb{Z}$, note that $(Y_i, V_i, S_{i+1} \cup Z_{i+1})$ is a separation such that $V_{i-1} \subseteq Y_i$ and $V_{i+1} \subseteq S_{i+1} \cup Z_{i+1}$ so we must have $|i - j| \leqslant 1$. If $i = j$, then $\gamma_0^{-i} \cdot u$ and $\gamma_0^{-i} \cdot v$ are adjacent in G_0 and thus $\widetilde{c}(u) = c(\gamma_0^{-i} \cdot u) \neq c(\gamma_0^{-i} \cdot v) = \widetilde{c}(v)$. Assume that $j = i + 1$, the other case being symmetric. Then $\gamma_0^{-i} \cdot u \in V_0$ and $\gamma_0^{-i} \cdot v$ is adjacent to $\gamma_0^{-i} \cdot u$. Moreover $\gamma_0^{-i} \cdot v = \gamma_0 \cdot (\gamma_0^{-j} \cdot v) \in V_1$. Thus we must have $\gamma_0^{-i} \cdot u \in Y_1$ and $\gamma_0^{-i} \cdot v \notin Y_1$ so $\gamma_0^{-i} \cdot v \in S_1$ and $\gamma_0^{-j} \cdot v \in S_0$. Then by the choice of γ_0, $c(\gamma_0^{-j} \cdot v) = c(\gamma_0^{-i} \cdot v)$. As c is a proper vertex-coloring and by definition of $\widetilde{c}$ we then have $\widetilde{c}(u) = c(\gamma_0^{-i} \cdot u) \neq c(\gamma_0^{-i} \cdot v) = c(\gamma_0^{-j} \cdot v) = \widetilde{c}(v)$, implying that $\widetilde{c}$ is a proper vertex-coloring of G.

As the sets V_i are finite and cover $V(G)$, and that $\gamma_0 \cdot V_i = V_{i+1}$ for each $i \in \mathbb{Z}$, the subgroup Γ' of Γ generated by γ_0 induces a quasi-transitive action on $V(G)$ so we can conclude. □

By our previous remark, Theorem 10.3 implies in particular that for any 2-ended connected locally finite quasi-transitive graphs G, there exists an orientation of the edges of G such that the oriented graph G is quasi-transitive, giving a positive answer in this case to Problem 10.2. We also deduce a positive answer to the edge-coloring version of Problem 10.1:

Corollary 10.4. *If G is locally finite with 2-ends and Γ acts quasi-transitively on G, then there exists a proper edge-coloring of G with $\chi'(G)$ colors and a cyclic subgroup Γ' of Γ such that Γ' preserves edge colors and acts quasi-transitively on G.*

Proof. We let $L(G)$ be the line-graph of G. Then $L(G)$ is also locally finite, connected, and as G is quasi-transitive of bounded degree, there are finitely many Γ-orbits of $E(G)$, so $L(G)$ is also Γ-quasi-transitive. Eventually, we claim that $L(G)$ is quasi-isometric to G and thus also 2-ended as the number of ends of a graph is a quasi-isometric invariant. An easy way to see this is to consider the barycentric subdivision $G^{(1)}$ of G which is quasi-isometric to G, and to observe that the mapping from $L(G)$ to $G^{(1)}$ that sends every vertex of $L(G)$ to the corresponding vertex of $G^{(1)}$ obtained after subdividing the associated edge from G defines a quasi-isometry. We thus conclude applying Theorem 10.3 to $L(G)$. □

We conclude this section by mentioning that all the examples we presented until now of locally finite quasi-transitive graphs G for which there exists a strongly periodic proper vertex-coloring (respectively edge-coloring) using finitely many colors, admit a strongly periodic proper vertex-coloring (respectively edge-coloring) using exactly $\chi(G)$ (respectively $\chi'(G)$) colors. However, there also exist constructions of locally finite quasi-transitive graphs with chromatic number 3, admitting a strongly periodic proper vertex-coloring using 4 colors, but no strongly periodic vertex coloring with 3 colors [AEGH24].

Chapter 2

Finitely generated groups

Notation: In this chapter we will mainly use capital latin letters to denote graphs, and capital greek letters to denote groups, except in Sections 14 and 15, where we will work more on groups than graphs, and thus denote groups with capital latin letters.

For every $n \in \mathbb{N}$, we will use the notation $[n]$ to denote the set of integers $\{1, \dots, n\}$. For every two sets X, Y, we will denote with Y^X the set of mappings from X to Y, and for every $f \in Y^X$ and every $X' \subseteq X$, we will denote with $f|_{X'} : X' \to Y$ the restriction of f on X'.

11 Introduction

Many notions introduced in Chapter 1 have their roots in geometric group theory. One of the best examples is the notion of accessibility, introduced first for finitely generated groups by Wall [Wal71], and which is closely related to the celebrated Stallings' ends theorem. This theorem, proved in [Sta68] for torsion-free groups and extended in [Ber68] for groups with torsion, states that every finitely generated group with at least two ends admits a non-trivial splitting over a finite subgroup (see Section 12 for a definition of splitting). According to Wall's definition [Wal71], a finitely generated group is called *accessible* if it admits a finite sequence of splittings over finite subgroups given by Stallings' theorem, such that the base groups involved in the sequence all have either 0 or 1 end. Thomassen and Woess [TW93] introduced the purely graph theoretic definition of (vertex-)accessibility for locally finite graphs that we gave in Section 2, and proved that a finitely generated group is accessible in the sense of Wall if and only one (and thus all) of its associated locally finite Cayley graphs is vertex-accessible. Similarly, many other notions and results mentioned in Chapter 1 can be seen as graphical counterparts of concepts and results from Bass-Serre theory: Woess [Woe89] proved that the finitely generated groups admitting a free group of finite index (these groups are also called *virtually free*) are exactly the groups whose associated locally finite Cayley graphs have bounded treewidth, and more recently, based on a result of Carmesin, Hamann and Miraftab [CHM22], Hamann, Lehner, Miraftab and Rühmann proved a graph theoretic version of Stallings' theorem for locally finite quasi-transitive graphs with many ends [HLMR22]. In this context, canonical tree-decompositions play a crucial role and appear to be the graph counterparts to Bass-Serre covering trees.

Our goal in this section is to present further connections between concepts from graph theory and finitely generated groups. After introducing in Section 12 some notions from geometric group theory that will be relevant for us, we will focus in Section 13 on classes of groups that are the direct analogues of the classes of graphs we studied in Chapter 1, and give a survey of the existing structural results for them. Many of these results are the group analogues of the structural graph theoretic results we presented in Chapter 1. We will then introduce in Sections 14 and 15 some key notions from symbolic dynamics of finitely generated groups. In particular, we will focus on two central questions in the theory of tilings of finitely generated groups, and try to convince the reader that they are intrinsically connected with some of the graph theoretic concepts we discussed in Chapter 1. The first is a conjecture of Ballier and Stein [BS18] related to the *domino problem on groups*. The second is a conjecture of Carroll and Penland [CP15] about the periodicity of tilings in groups. In Section 15, we will discuss further different notions of periodicity/aperiodicity that can be studied in groups, and try to characterize when these notions coincide, for some specific classes of groups. In Section 16, we will suggest a way to extend the notions from symbolic dynamics discussed earlier to locally finite quasi-transitive graphs. In particular it suggests that such an approach could allow to reuse results and ideas from symbolic dynamics to tackle some graph-theoretic questions. Eventually, we will present in Section 17 a simplification of a construction of Osajda [Osa20] of Gromov's monster groups. Osajda's construction is probabilistic and based on the Lovász Local Lemma, and we show that replacing this probabilistic step by a purely combinatorial counting argument popularized by Rosenfeld [Ros20] not only simplifies Osajda's proof, but allows to optimize significantly the rank of the monster groups constructed. Many of the results and observations from Sections 13 and 14 are either already known results or come from the paper [EGLD23], a joint work with Louis Esperet and Clément Legrand-Duchesne. The content of Section 15 mainly comes from a joint work with Étienne Moutot and Solène Esnay, which is unpublished yet. Results of Section 16 are unpublished results for which I am the sole author, and the results and proofs from Section 17 come from the paper [EG24b] co-authored with Louis Esperet.

12 Cayley graphs and group presentations

Group presentations. In the remainder of Chapter 2, unless stated otherwise, the groups we will consider will always be assumed to be *finitely generated*. Up to doubling the size of the generating sets, we will always assume that the sets of generators S that we consider are closed under taking inverse, that is that for each $\alpha \in S$ we also have $\alpha^{-1} \in S$. We will denote with 1_Γ the identity element of Γ, and we will call Γ *trivial* when $\Gamma = \{1_\Gamma\}$. For every $\gamma \in \Gamma$, we will also denote with $|\gamma|_S := k \in \mathbb{N}$ the minimum number k of elements $\alpha_1, \dots, \alpha_k \in S$ such that we can write $\gamma = \alpha_1 \cdots \alpha_k$.

Let A be a *finite alphabet*, that is, a finite set of symbols. For each letter $a \in A$, we choose a new symbol $\overline{A}$ not in A, and call it the *formal inverse of* a. We let $\overline{A} := \{\overline{a} : a \in A\}$ denote the set of formal inverses of elements of A, and we set $\overline{\overline{a}} := a$ for each $a \in A$, so that $\bar{\cdot} : A \uplus \overline{A} \to A \uplus \overline{A}, s \mapsto \overline{s}$ is a fixpoint free involution. We let $(A \uplus \overline{A})^*$ denote the set of words of finite length written on the alphabet $A \uplus \overline{A}$. In particular, we let $\varepsilon \in (A \uplus \overline{A})^*$ denote the empty word. We extend $\bar{\cdot}$ on $(A \uplus \overline{A})^*$ by setting $\overline{\varepsilon} := \varepsilon$ and $\overline{a_1 \ldots a_k} := \overline{a_k} \cdot \ldots \cdot \overline{a_1}$ for each $a_1, \ldots, a_k \in A \uplus \overline{A}$ and $k \geqslant 1$. A word of $(A \uplus \overline{A})^*$ is *reduced* if it contains no two consecutive letters which are formal inverse one of the other. For every word $w \in (A \uplus \overline{A})^*$ such that one can write $w = w_1 a \overline{a} w_2$ for some $w_1, w_2 \in (A \uplus \overline{A})^*$ and $a \in A \uplus \overline{A}$, we say that the word $w' := w_1 w_2$ is obtained from w after performing an *elementary reduction*. It is not hard to check and well-known that every word $w \in (A \uplus \overline{A})*$ admits a unique associated *reduced form*, i.e., a reduced word w' obtained from w after performing at most $|w|/2$ elementary reductions. The *free group* on A is the group $(\mathcal{F}(A), \cdot)$ whose elements are the reduced words of $A \uplus \overline{A}$, equipped with the binary operation that associates to each pair (w_1, w_2) of reduced words the reduced form of $w_1 w_2$. A word $w \in (A \uplus \overline{A})^*$ is *cyclically reduced* if it is reduced and if its first and last letters are not inverse one of another. Again, one can extend the previous definitions and define similarly *elementary cyclic reductions* so that every word admits a unique *cyclically reduced form*.

For every group Γ with a finite set S of generators, there is a (unique) canonical surjective group morphism $\pi_S : \mathcal{F}(S) \to \Gamma$ (see for example [DK18, Proposition 7.21]) such that $\pi_S(s) = s$ and $\pi_S(\overline{s}) = s^{-1}$ for all $s \in S$. For convenience, we extend π_S to a mapping $\pi_S : (S \uplus \overline{S})^* \to \Gamma$, by setting for every word $w \in (S \uplus \overline{S})^*$ with reduced form $w' \in \mathcal{F}(S)$, $\pi_S(w) := \pi_S(w')$. For every word $w \in (S \uplus \overline{S})^*$, we say that w *represents* the group element $\pi_S(w)$. The words of $\pi_S^{-1}(1_\Gamma) \subseteq (S \uplus \overline{S})^*$ are called the *relations* of (Γ, S). A *group presentation* of Γ is a pair (S, R), such that S is a finite generating set of Γ, and such that $R \subseteq \mathcal{F}(S)$ is a set of relations of (Γ, S) such that $\ker(\pi_S)$ is normally generated by R, i.e., such that for every word $w \in \mathcal{F}(S)$, such that $\pi_S(w) = 1_\Gamma$, there exist relations $v_1, \ldots, v_k \in R$ and words $w_1, \ldots, w_k$ (not necessarily distinct) such that w is the reduced form of the word

$$(w_1 v_1 \overline{w_1}) \cdot \ldots \cdot (w_k v_k \overline{w_k}).$$

We usually use the notation $\langle S \mid R \rangle$ instead of (S, R) to denote presentations. If R is finite, we say that $\langle S \mid R \rangle$ is a *finite presentation* of Γ, and we call Γ a *finitely presented* group if it admits a finite presentation. We will also often identify elements of $(S \uplus \overline{S})^*$ with their images by π_S. Note that if $\langle S \mid R \rangle$ is a presentation of Γ and if R' is the set of cyclically reduced forms of the words in R, then $\langle S|R' \rangle$ is also a presentation of Γ. Dunwoody [Dun85] proved that every finitely generated and finitely presented group is accessible. A *Tietze transformation* of a presentation $\langle S \mid R \rangle$ is one of the following operations:

- *Adding a relation to* R: we add a relation $r \in (S \uplus \overline{S})^*$ that is already representing the identity element of G;
- *Removing a relation of* R: we remove a redundant relation r from R, i.e., a relation $r \in R$ such that r also corresponds to the identity element of the group described by the presentation $\langle S \mid R \setminus \{r\} \rangle$;
- *Adding a generator to* S: we add a new letter s to S disjoint from $S \uplus \overline{S}$ together with a relation sw for some word $w \in (S \uplus \overline{S})^*$;

- *Removing a generator from* S: if some relation r can be written $r = sr'$, where neither s nor $\overline{s}$ occur in r', we remove r from R and s from S, and replace every occurence of s (respectively of $\overline{s}$) in the other relations by $\overline{r'}$ (respectively by r').

A basic property of group presentations is that applying a Tietze transformation to a presentation does not change the group it represents.

Cayley graphs. Given a finitely generated group Γ and a finite set S of generators in Γ which is closed under taking inverse, the *Cayley graph* $\mathrm{Cay}(\Gamma, S)$ of Γ with respect to S is the graph with vertex set Γ and such that for every $g \in \Gamma$, the neighbors of g in $\mathrm{Cay}(\Gamma, S)$ are the elements of Γ of the form gs for some $s \in S$. Note that there is a natural edge-labelling and orientation of $\mathrm{Cay}(\Gamma, S)$, which consists in orienting each edge $\{g, gs\}$ from g to gs and labelling it with s for all $g \in \Gamma$ and $s \in S$; in particular in this oriented version, if $s \in S$ is such that $s = s^{-1}$, there are two opposite arcs between each pair $\{g, gs\}$. While we will mainly work with unoriented Cayley graphs, we will sometimes implicitly consider their oriented and edge-labelled version. Note that as S is a finite generating set of Γ, $\mathrm{Cay}(\Gamma, S)$ must be connected and locally finite. Moreover, if S and S' are two different finite generating sets of the same group Γ, it is well-known and not hard to check that the identity mapping id_Γ is a quasi-isometry between $\mathrm{Cay}(\Gamma, S)$ and $\mathrm{Cay}(\Gamma, S')$. As the number of ends of a graph is a quasi-isometric invariant, we can define in particular the number of ends of a finitely generated group Γ as the number of ends of any of its locally finite Cayley graphs. We say that two finitely generated groups Γ and Γ' are *quasi-isometric* to each other if some (and thus all) of their locally finite Cayley graphs are quasi-isometric. The following result of Babai already implies a first connection between the algebraic structure of a group and the topological structure of its Cayley graphs.

Theorem 12.1 ([Bab77])**.** *Let* Γ *be a group with a finite set of generators* S*. If* Γ' *is a finitely generated subgroup of* Γ*, then there exists some finite generating set* S' *of* Γ' *such that* $\mathrm{Cay}(\Gamma', S')$ *is a minor of* $\mathrm{Cay}(\Gamma, S)$*.*

If $\langle S \mid R \rangle$ is a pair where S is a finite generating set of Γ and $R \subseteq (S \cup \overline{S})^*$, we also define the *simplified Cayley complex* $\mathrm{Cay}^{(2)}(\langle S \mid R \rangle)$ of Γ with respect to $\langle S \mid R \rangle$ as the 2-dimensional cell-complex with underlying graph $\mathrm{Cay}(\Gamma, S)$, and where for each relation $r \in R$ and each $g \in \Gamma$, if $r = a_1 \cdot \ldots \cdot a_k$ with $a_1, \ldots, a_k \in S \cup \overline{S}$, we add a 2-cell consisting in a disk whose boundary is identified with the closed walk $(g, ga_1, ga_1a_2, \ldots, gr)$. We identify two 2-cells if they have the same boundary, up to rotation. In other words, if $r' = a_i \ldots a_k a_1 \ldots a_{i-1}$ for some $i \in \{2, \ldots, k\}$, then we identify the 2-cells associated to r and r' (if we do not identify such identifications, then the constructed complex is the *Cayley complex* of (Γ, S), which is more commonly studied in the litterature).

Note that left multiplication by elements of Γ induces a natural transitive action on $\mathrm{Cay}(\Gamma, S)$ and $\mathrm{Cay}^{(2)}(\langle S \mid R \rangle)$. Moreover, this action is also *regular*, i.e., for each pair of vertices u, v of $\mathrm{Cay}(\Gamma, S)$ there is at most (and thus by transitivity exactly) one automorphism $g \in \Gamma$ such that $g \cdot u = v$. In the remainder of the manuscript, we will always implicitly equip a Cayley graph with such a regular transitive action. If Γ is a group acting on a graph G, we call the action *quasi-regular* if for each $v \in V(G)$, the stabilizer Γ_v is finite.

Sabidussi [Sab58] proved that a connected (unlabelled) graph is a Cayley graph if and only if there exists some group Γ acting transitively and regularly on G. Relaxing transitivity and regularity, we obtain the following coarse version of Sabidussi's theorem, which is a special case of the Švarc–Milnor lemma [Š55, Mil68].

Lemma 12.2 (Švarc–Milnor lemma for graphs)**.** *If G is a locally finite and Γ is a group with a quasi-transitive and quasi-regular action on G, then Γ is finitely generated and for each finite set of generators S of Γ, $\mathrm{Cay}(\Gamma, S)$ is quasi-isometric to G.*

The following is a folklore result.

Lemma 12.3. *Let Γ be a group and S a finite generating set of Γ. Then for any set of reduced words $R \subseteq (S \cup \overline{S})^*$, the following are equivalent:*

- *(i)* *$\langle S \mid R\rangle$ is a presentation of Γ;*
- *(ii)* *$\mathrm{Cay}^{(2)}(\langle S \mid R\rangle)$ is simply connected;*
- *(iii)* *There exists a Γ-invariant set of closed walks $\mathcal{E}$ in $\mathrm{Cay}(\Gamma, S)$ generating the set of closed walks $\mathcal{W}(\mathrm{Cay}(\Gamma, S))$ such that the words labelling the walks in $\mathcal{E}$ are exactly the words of R (up to applying some rotations on the walks of $\mathcal{E}$).*

Proof. The equivalence $(i) \Leftrightarrow (ii)$ is a well-known result and can be proved using covering spaces together with the Seifert-Van Kampen theorem (see for example the proof of [DK18, Lemma 7.90]).

To prove $(i) \Rightarrow (iii)$, let W be a closed walk in $\mathrm{Cay}(\Gamma, S)$ and w denote its labelling word. Up to removing spurs in W, assume that $w \neq \varepsilon$ is reduced. Then we must have $\pi_S(w) = 1_\Gamma$, so in particular there exists $v_1, \ldots, v_k \in R$ and $w_1, \ldots, w_k$ such that w is the reduced form of the word

$$(w_1 v_1 \overline{w_1}) \cdot \ldots \cdot (w_k v_k \overline{w_k}).$$

Up to adding some spurs, it means that we can write W as a sum $W_1 \cdot \ldots \cdot W_k$ such that for each $i \in [k]$, W_i is a closed walk in $\mathrm{Cay}(\Gamma, S)$ labelled by the words $w_i v_i \overline{w_i}$. In particular, note that for each $i \in [k]$, W_i is obtained from a closed walk labelled by v_i after adding $|w_i|$ spurs, hence the set $\mathcal{E}$ of closed walks in $\mathrm{Cay}(\Gamma, S)$ with labels in R is indeed a generating set of $\mathcal{W}(\mathrm{Cay}(\Gamma, S))$.

The implication $(iii) \Rightarrow (ii)$ is proved in [Ham18b, Proposition 6.1]. □

Graphs of groups. We give here a few basic definitions related to Bass-Serre theory, and refer the reader to [Ser80, Chapter 1] or [DD89, Chapter 1] for more details.

A *graph of groups* consists of a pair $(G, \mathcal{G})$ such that:

- G is a graph (possibly having loops and multi-arcs) with a fixed orientation A of $E(G)$, together with two mappings $\iota, \tau : A \to V(G)$ such that $\iota(e)$ and $\tau(e)$ are the initial and terminal vertices of each edge $e \in E(G)$ according to the orientation of e in A,
- $\mathcal{G}$ is a family of *vertex-groups* Γ_v for each $v \in V(G)$, *edge-groups* Γ_e^-, Γ_e^+ for each edge $e \in E(G)$, and of group isomorphisms $\phi_e : \Gamma_e^- \to \Gamma_e^+$ for each $e \in E(G)$, such that Γ_e^- and Γ_e^+ are respectively subgroups of $\Gamma_{\iota(e)}$ and $\Gamma_{\tau(e)}$ for each $e \in E(G)$.

When Γ_e^- and Γ_e^+ are trivial for some edge $e \in E(G)$, we identify them and simply write Γ_e.

Let T be a spanning tree of G and $\langle S_v | R_v \rangle$ be a presentation of Γ_v for each $v \in V(G)$. The *fundamental group* $\Gamma := \pi_1(G, \mathcal{G})$ of the graph of groups $(G, \mathcal{G})$ is defined as the group having as generators the set

$$S := \left(\biguplus_{v \in V(G)} S_v \right) \uplus \left(\biguplus_{e \in E(G) \setminus E(T)} \{t_e\} \right)$$

and as relations:

- the relations of each R_v;
- for every edge $e \in E(T)$ and every $g \in \Gamma_e^- \subseteq \Gamma_{\iota(e)}$, the relation $g^{-1}\phi_e(g)$;
- for every edge $e \in E(G) \setminus E(T)$ and every $g \in \Gamma_e^-$, the relation $g^{-1} t_e \phi_e(g) t_e^{-1}$.

It can be shown that for a given graph of groups $(G, \mathcal{G})$, the definition of its fundamental group does not depend of the choice of the spanning tree T (see for example [Ser80, Section I.5.1]). Fundamental groups of graphs of groups generalize the operations of amalgamated free products and (multi-)HNN-extensions of groups (see below for a definition of these operations). Another well-known property is that for any $v \in V(G), e \in E(G)$, Γ_v and Γ_e are subgroups of $\pi_1(G, \mathcal{G})$.

We give a few examples of fundamental groups of graphs of groups.

- If G consists of a single vertex v_0 with a loop e, and if we let Γ_{v_0} and Γ_e be trivial groups, then $\pi_1(G, \mathcal{G})$ is isomorphic to the abelian free group $(\mathbb{Z}, +)$. More generally, if $E(G)$ contains multiple v_0-loops with trivial associated edge-groups, then $\pi_1(G, \mathcal{G})$ is isomorphic to the free group $\mathcal{F}(E(G))$.

- Let Γ_1 and Γ_2 be two finitely generated groups with respective presentations $\langle S_1 | R_1 \rangle$ and $\langle S_2 | R_2 \rangle$ (where S_1 and S_2 are distinct alphabets), and G be the graph with two vertices v_1, v_2 together with a single edge $e := v_1 v_2$. If $\Gamma_{v_i} := \Gamma_i$ for each $i \in \{1, 2\}$ and Γ_e is the trivial group, then the group $\pi_1(G, \mathcal{G})$ is exactly the *free product* of Γ_1 and Γ_2, i.e., the group defined by the presentation $\langle S_1 \uplus S_2 | R_1 \uplus R_2 \rangle$.

- Let Γ be a copy of $(\mathbb{Z}, +)$, a a generator of Γ, and G be the graph with a single vertex v_0 together with a single v_0-loop e. We also fix two integers $m, n \geqslant 1$. We let $\Gamma_e^- := \{a^{i \cdot m} : i \in \mathbb{Z}\} \simeq m\mathbb{Z}$, $\Gamma_e^+ := \{a^{i \cdot n} : i \in \mathbb{Z}\} \simeq n\mathbb{Z}$ and set $\phi_e : \Gamma_e^- \to \Gamma_e^+, a^{i \cdot m} \mapsto a^{i \cdot n}$. Then the group $\pi_1(G, \mathcal{G})$ is the *Baumslag-Solitar* group $\mathrm{BS}(m, n)$ with presentation $\langle a, b \mid a^m b a^{-n} b^{-1} \rangle$.

Bass-Serre theory basically establishes a correspondence between actions of groups on trees and their decompositions into a graph of groups. We summarize in the next lemma some useful properties of one direction of this correspondence. If T is a tree and Γ is a group acting on T, we say that Γ acts on T *without inversion* if there is no pair $(uv, g) \in E(T) \times \Gamma$ such that $g \cdot u = v$ and $g \cdot v = u$.

Lemma 12.4 (Theorem 13 in [Ser80])**.** *Let Γ be a group and T be a tree (not necessarily locally finite) such that Γ acts on T without inversion. Then Γ is isomorphic to the fundamental group of a graph of groups $(G, \mathcal{G})$ such that:*

- *$V(G)$ is in bijection with the set $V(T)/\Gamma$ of Γ-orbits of $V(T)$, and $E(G)$ is in bijection with the set $E(T)/\Gamma$ of Γ-orbits of $E(T)$,*
- *For each $v \in V(G)$, the vertex-group Γ_v is isomorphic to some subgroup $\mathrm{Stab}_\Gamma(x)$ of Γ for some $x \in V(T)$,*
- *For each $e \in E(G)$, the edge-group Γ_e is isomorphic to some subgroup $\mathrm{Stab}_\Gamma(e')$ of Γ for some $e' \in E(T)$.*

Moreover, if each edge-group Γ_e is finitely generated, then each vertex-group is also finitely generated.

Group accessibility. As mentioned in the introduction, the original definition of accessibility concerns finitely generated group and was given by Wall [Wal71]. We give here for completeness the basic definitions of HNN-extensions, free products with amalgamations and accessibility, however we will not use them later, hence the reader is free to jump directly to the next subsection.

We start defining HNN-extensions and free products with amalgamations. We refer to [Ser80, Chapter 1.1] for more details. Let Γ be a finitely generated group and $\langle S \mid R\rangle$ be a group presentation of Γ. Let t be a new letter distinct from S. Assume that Λ is a subgroup of Γ and let $\phi : \Lambda \to \Gamma$ be an injective group morphism. The *HNN-extension of* Γ *with respect to* (Λ, ϕ) is the group denoted $\Gamma *_{\Lambda,\phi}$ (or just $\Gamma *_\Lambda$ when the context is clear) with generating set $S \uplus \{t\}$ and defined by the relations from R, together with a relation $twt^{-1}w'^{-1}$ for each element $h \in \Lambda$ and every two reduced words $w, w' \in S^*$ corresponding respectively to h and $\phi(h)$ in Γ. We call Γ a *factor* of the HNN-extension $\Gamma *_\Lambda$.

Now let Γ_1, Γ_2 be two finitely generated groups with respective presentations $\langle S_1 \mid R_1\rangle$, $\langle S_2 \mid R_2\rangle$, and let Λ_1, Λ_2 be respectively subgroups of Γ_1 and Γ_2 and $\phi : \Lambda_1 \to \Lambda_2$ be a group isomorphism. Assume moreover that S_1 and S_2 are disjoint alphabets. The *amalgamated free product of* Γ_1 *and* Γ_2 *with respect to* $(\Lambda_1, \Lambda_2, \varphi)$ is the group denoted $\Gamma_1 *_{\Lambda_1,\Lambda_2,\phi} \Gamma_2$ (or just $\Gamma_1 *_{\Lambda_1} \Gamma_2$ when the context is clear) with generators $S_1 \uplus S_2$ and defined by the relations from $R_1 \uplus R_2$ together with a relation $w\varphi(w')^{-1}$ for each element $h \in \Lambda_1$ and every two words $w \in (S_1 \uplus S_1^{-1})^*, w' \in (S_2 \uplus S_2^{-1})^*$ corresponding respectively to h and $\phi(h)$ in Γ_1 and Γ_2. We call Γ a *factor* of the amalgamated free product $\Gamma_1 *_{\Lambda_1} \Gamma_2$.

A *splitting of* Γ *over a subgroup* Λ *of* Γ is a decomposition of Γ into either an HNN-extension or an amalgamated free product with respect to subgroups isomorphic to Λ. A splitting is *trivial* if Γ equals to one of the base subgroups involved in its decomposition. Stallings' theorem [Sta68] states that every group with at least 2 ends admits a non-trivial splitting over a finite subgroup. We define accessibility inductively as follows: a finitely generated group Γ is *accessible* if:

- Γ has at most one end,

- or Γ admits a non-trivial splitting over a finite subgroup such that the factors of the splittings are accessible.

In other words, Γ is accessible if there exists at least one finite sequence of splittings over finite subgroups starting from Γ which are given by Stallings' theorem. Walls [Wal71] proved that a finitely generated group is accessible if and only if every sequence of splittings given by Stallings' theorem is finite. A graph theoretical version of Stallings' theorem has been recently proved for quasi-transitive locally finite graphs, where *tree-amalgamations* play the role of splittings [HLMR22]. In particular, the authors asked if a characterization similar to the one of Walls holds for vertex-accessible quasi-transitive graphs, with respect to the operation of tree-amalgamation they consider. We briefly mention a few other known properties about accessibility in finitely generated groups. It can be shown using Bass-Serre theory that a group is accessible if and only if it is the fundamental group of a finite graph of groups with finite edge-groups and whose vertex-groups have at most one end. As we mentioned in the introduction, Thomassen and Woess [TW93] proved that a finitely generated group is accessible if and only if one of its Cayley graphs is vertex-accessible. Finally, Dunwoody [Dun85] proved that every finitely presented group is accessible.

13 Bounded treewidth, planar and K_∞-minor-free groups

In this section, we present structural characterisations of finitely generated groups whose Cayley graphs satisfy the properties studied in Sections 4, 5 and 7. The goal of this section is to survey some known structural properties of such groups, which appear to correspond to the group counterpart of the graph theoretical results we presented in Chapter 1. Results presented in Sections 13.1 and 13.2 are well-known results, while results from Section 13.3 are corollaries of the results from Section 7.

13.1 Virtually free groups

A finitely generated group Γ is said to be *virtually free* if it contains a free subgroup of finite index. Similarly, Γ is said to be *virtually cyclic* if it contains a cyclic subgroup of finite index. In particular, note that a virtually cyclic group Γ is infinite if and only if it contains a subgroup of finite index which is isomorphic to $\mathbb{Z}$. In this case, we say that Γ is *virtually* $\mathbb{Z}$. More generally, if Γ and Λ are any two finitely generated groups such that Γ contains a subgroup isomorphic to Λ with finite index in Γ, then we say that Γ is *virtually* Λ. We say that two finitely generated groups Γ, Γ' are *commensurable* if there exists a finitely generated group Λ such that both Γ and Γ' are virtually Λ. It is not hard to prove that if Γ and Γ' are commensurable, then they are quasi-isometric. However the converse is false in general (see [dlH00, IV.44, IV.47, IV.48]).

Note that by Lemma 4.1 and Theorem 4.3, the property for a locally finite graph of having bounded pathwidth/treewidth is a quasi-isometric invariant. We thus say that a finitely generated group Γ has finite pathwidth (respectively treewidth) if at least one (and thus all) of its locally finite Cayley graphs has finite pathwidth (respectively treewidth).

A well-known result in geometric group theory states that finitely generated 2-ended groups are exactly the groups that are virtually $\mathbb{Z}$ (see for example [DK18, Proposition 9.23]), thus by Lemma 4.1 they correspond exactly to the finitely generated groups having one (or equivalently all) locally finite Cayley graph of bounded pathwidth. A result of Woess [Woe89] states that finitely generated groups whose ends are thins (or equivalently that have finite treewidth by Theorem 4.3) correspond exactly to the virtually free groups. Virtually free groups form a rich class of groups and admit many known characterisations of different flavours. We only mention here that virtually free groups correspond exactly to the groups acting on a tree with finite vertex stabilizers [KPS73]. In particular, by Bass-Serre theory, the virtually free groups are exactly the groups that are isomorphic to the fundamental group of a finite graph of groups with finite vertex-groups and edge-groups. We refer to [Ant11] for a (non exhaustive) list of properties equivalent to the property of having bounded treewidth in Cayley graphs.

13.2 Planar groups

If X is a metric space and $\mathrm{Aut}(X)$ denotes the group of isometries of X, we say that a group Γ is a *discontinuous group of isometries of* X if it is a subgroup of $\mathrm{Aut}(X)$ such that for any $x \in X$, there exists an open neighborhood U_x of x such that for every element $g \in \Gamma \setminus \{1_\Gamma\}$, $g \cdot U_x \cap U_x = \emptyset$.

A finitely generated group Γ is said to be *planar* if it admits a finite set S of generators such that $\mathrm{Cay}(\Gamma, S)$ is planar. Note that unlike the property of having bounded-treewidth, being planar is not a quasi-isometric invariant, hence the property of being planar for a group relies on the choice of a special generating set. Nevertheless, planarity is still significant from a group perspective: as mentioned in Section 1, by a result of Maschke [Mas96], the finite planar groups are exactly the countable discontinuous groups of isometries of the 2-dimensional sphere $\mathbb{S}^2$. In the same paper, Maschke also characterized the finite planar groups as those admitting a presentation $\langle S \mid R\rangle$ such that the associated simplified Cayley complex admits an embedding in $\mathbb{S}^2$.

A finitely generated group Γ is said to be *planar discontinuous* if there exists a finite presentation $\langle S \mid R\rangle$ of Γ such that the associated simplified Cayley complex is planar, i.e., embeddable either in $\mathbb{S}^2$ or $\mathbb{R}^2$. Note that the basic definition of planar discontinuous from Zieschang, Volgt and Coldewey [ZVC80] differs from the one we gave, however it is shown in [ZVC80, Theorems 4.13.11, 6.4.7 and Corrolary 4.13.15] that both definitions are equivalent. Planar discontinuous groups form a rich class of groups, and thanks to the works of Zieschang, Volgt, Coldewey [ZVC80] and of Muller and Schupp [MS83], their algebraic structure is well understood. Among other characterisations, one can mention the followings.

Theorem 13.1 ([ZVC80, Bab97, Geo14])**.** *Let Γ be a finitely generated group. Then the following are equivalent:*

(*i*) Γ *is planar discontinuous;*

(*ii*) *there exists a finite set S of generators of Γ such that* $\mathrm{Cay}(\Gamma, S)$ *is a VAP-free planar graph;*

(iii) *Γ is isomorphic to a subgroup of some discontinuous countable group of isometries of either $\mathbb{S}^2$, or $\mathbb{R}^2$ or of the hyperbolic plane $\mathbb{H}^2$.*

Moreover, if Γ has one end, the previous items are equivalent to the following:

(iv) *Γ admits a regular and quasi-transitive group action on a locally finite planar graph.*

We briefly explain how to obtain the equivalences of Theorem 13.1 from the references we gave. The equivalence $(i) \Leftrightarrow (ii)$ is proved in [Geo14], while $(i) \Leftrightarrow (iii)$ is proved in [ZVC80, Theorems 4.13.11, 6.4.7 and Corrolary 4.13.15]. Assume now that Γ has one end. The implication $(ii) \Rightarrow (iv)$ is immediate and holds in general for groups with an arbitrary number of ends, as a group always acts transitively and regularly on any of its Cayley graphs. To show $(iv) \Rightarrow (iii)$, assume that Γ admits a regular and quasi-transitive group action on some locally finite planar graph G. A graph H is a *topological minor* of G if $V(H) \subseteq V(G)$, and if for every edge $uv \in E(H)$, there exists a path P_{uv} from u to v in $G - V(H)$ such that for every two different edges $uv, u'v' \in E(H)$, the paths P_{uv} and $P_{u'v'}$ do not intersect each other, except possibly in their extremities. Then by Lemma 12.2, G is quasi-isometric to any locally finite graph of Γ, so in particular G has one end. By [Bab97, Theorem 4.1], there exists a topological subgraph H of G which is locally finite, one-ended, 3-connected and such that Γ induces a quasi-transitive group action on H. Note that Theorem 4.1 in [Bab97] is stated in the special case where $\Gamma = \mathrm{Aut}(G)$, however its proof still holds if we replace it with any group Γ acting quasi-transitively on G. By [Bab97, Theorem 4.2], H admits an embedding into some metric space X, with $X \in \{\mathbb{S}^2, \mathbb{R}^2, \mathbb{H}^2\}$ and such that every automorphism of X is induced by an isometry of X. In particular, by our remark from Section 5 that every locally finite quasi-transitive one-ended planar graph is VAP-free, H must be VAP-free. Moreover, it follows from the proof of Theorem 4.2 in [Bab97] that the embedding of H in X has no accumulation point, implying that Γ is a discontinuous group of isometries of X, as desired.

As every one-ended locally finite planar quasi-transitive graph is VAP-free, Theorem 13.1 implies in particular that every one-ended planar group is planar discontinuous. The *surface group Γ_g of genus g* for $g \geqslant 0$ is the fundamental group of the closed orientable surface of genus g. For each $g \geqslant 1$, the surface group of genus g can alternatively be defined by the finite presentation

$$\langle a_1, b_1, \ldots, a_g, b_g \mid [a_1, b_1] \ldots [a_g, b_g] \rangle,$$

where for every two letters a, b, we set $[a, b] := ab\overline{a}\overline{b}$. When $g = 1$, Γ_g is isomorphic to $(\mathbb{Z}^2, +)$ and thus planar, while if $g \geqslant 2$, it is well-known that Γ_g is a planar discontinuous group and admits a VAP-free planar Cayley graph that embeds isometrically into the hyperbolic plane $\mathbb{H}^2$. Moreover, for every $g \geqslant 2$, Γ_2 contains Γ_g as a subgroup of finite index. We say that a group is a *virtually surface* group if there exists $g \in \mathbb{N}$ such that Γ is virtually Γ_g. The following result states that planar discontinuous groups are algebraically close from surface groups.

Theorem 13.2 ([BN46], [Fox52]). *Every planar group with at most one end is a virtually surface group.*

Putting all of these results together, we hope that the reader is convinced that virtually surface groups are well-known from a structural perspective. To describe more generally the structure of planar groups which are not necessarily planar discontinuous, one can refer to the work of Droms [Dro06], who found an explicit inductive way of decomposing any finitely generated planar group into free products with amalgamations and HNN-extensions, with special contraints taking into account the planar embeddings of the Cayley graphs involved. Droms' process of decomposition always ends after a finite number of steps, and the decomposition it gives has to be seen as the group counterpart of Theorem 5.8.

To complete the picture on planar groups, one should also mention the class of finitely generated groups Γ that admit some locally finite Cayley graph $\mathrm{Cay}(\Gamma, S)$ with a planar embedding which is *covariant*, i.e., such that every element of Γ induces an automorphism of $\mathrm{Cay}(\Gamma, S)$ that maps every facial path to a facial path. These groups form a proper subclass of the class of planar groups, and a proper superclass of the class of planar discontinuous groups, and can be characterized as groups admitting a faithful, proper discontinuous and co-compact group action by homeomorphisms on the sphere $\mathbb{S}^2$, the plane $\mathbb{R}^2$, the open annulus, or the Cantor sphere [Geo20, Theorem 1.1 and Proposition 10.1].

13.3 Minor-excluded groups

We say that a finitely generated group Γ is K_∞*-minor-free* if there exists a finite set of generators S such that $\mathrm{Cay}(\Gamma, S)$ is K_∞-minor free. As an immediate corollary of Theorem 7.1, we show that K_∞-minor-free groups have a very specific structure.

Theorem 13.3. *Let Γ be a finitely generated K_∞-minor-free group. Then Γ is isomorphic to the fundamental group of a finite graph of groups $(G_0, \mathcal{G})$ such that:*

- *Every vertex group is finitely generated and either finite or one-ended planar,*
- *Every edge-group is finite.*

Proof. Let S be a finite set of generators such that the graph $G := \mathrm{Cay}(\Gamma, S)$ is K_∞-minor free. By Theorem 7.3, there exists a Γ-canonical tree-decomposition $(T, (V_t)_{t\in V(T)})$ of G of finite adhesion, whose torsos are minors of G which are either 3-connected planar quasi-transitive graphs or finite, and with tight edge-separations, implying that $E(T)$ has finitely many Γ-orbits. We let G^+ be the supergraph of G with vertex set $V(G)$ and edge-set $E(G^+) := \bigcup_{t\in V(T)} E(G[\![V_t]\!])$. Then, $(T, \mathcal{V})$ is also a tree-decomposition of G^+. As $E(T)$ has finitely many Γ-orbits, for every vertex $v \in V(G)$, the set $\{t \in V(T) : v \in V_t\}$ must be finite, so G^+ is also locally finite. Moreover, as $(T, \mathcal{V})$ is Γ-canonical, Γ still defines a regular transitive action on G^+, and we have for each $t \in V(T)$, $G^+[\![V_t]\!] = G^+[V_t] = G[\![V_t]\!]$. For each $t \in V(T)$ such that $G[\![V_t]\!]$ has at least 2 ends, there exists by Corollary 5.9 a Γ_t-canonical tree-decomposition $(T_t, \mathcal{V}_t)$ of $G[\![V_t]\!]$ whose parts are planar, connected with at most one end, and such that $E(T_t)$ has finitely many Γ_t-orbits. Thus applying Corollary 3.15, there exists a Γ-canonical tree-decomposition $(T', (V'_t)_{t\in V(T')})$ of G^+ with finite adhesion refining $(T, (V_t)_{t\in V(T)})$, such that $E(T')$ has finitely many Γ-orbits and whose parts are connected, planar with at most one end. As the action of Γ is regular on G^+, for each $t \in V(T')$, Γ_t also induces a regular group action on $G^+[V'_t]$, and by Lemma 3.17, this action is also quasi-transitive. Thus for each $t \in V(T')$ such that V_t is finite, Γ_t must be finite and, for the same reasons, for each $e \in E(T')$, Γ_e is also finite. Moreover, for each $t \in V(T')$ such that $G^+[V_t']$ is infinite, as the graph $G^+[V_t']$ is planar, quasi-transitive with one end, the implication $(iv) \Rightarrow (i)$ from Theorem 13.1 implies that Γ_t is a one-ended planar group. Thus we immediately obtain the desired result after applying Lemma 12.4 to the action of Γ on T'. □

Note that Theorem 13.3 implies that K_∞-minor-free groups are accessible, which can alternatively be proved using Theorem 8.2 together with Thomassen-Woess' characterisation of accessible groups [TW93]. In fact we obtain as a direct consequence of Theorem 8.3 the more general result that K_∞-minor-free groups are finitely generated, generalizing a result of Droms [Dro06] who proved that planar groups are finitely presented.

Corollary 13.4. *Let Γ be a finitely generated K_∞-minor-free group. Then Γ is finitely presented.*

Proof. We let $G := \mathrm{Cay}(\Gamma, S)$ be a locally finite Cayley graph of Γ which is K_∞-minor-free, and consider the left action of Γ on G. Let $\mathcal{W}$ denote a finite set of representatives of the Γ-orbits of the Γ-invariant generating set of closed walks of G obtained after applying Theorem 8.3 to G. Note that for every $W \in \mathcal{W}$, if r_W denotes the sequence of labels of W with respect to S, then the set of closed walks of G labelled by r_W is exactly the orbit $\Gamma \cdot W$ of W. Thus by Lemma 12.3, $\langle S \mid r_W, W \in \mathcal{W}\rangle$ is a finite presentation of Γ. □

14 Subshifts of finite type and the domino problem

In this section as well as in Section 15, we will adopt notations and vocabulary from symbolic dynamics (see [ABJ18, Chapter 9]). In particular, as we will manipulate more groups than graphs, we will change our convention and denote groups and group elements with letters from the latin alphabet. Theorem 14.7 from Section 14 comes from the paper [EGLD23]. The proof we give there is shorter and based directly on Theorem 13.3.

14.1 Subshifts of finite type

Configurations, patterns, SFTs. Let G be a finitely presented group and A be a finite alphabet. We call elements of A^G *configurations* of G. For any finite subset F of G, we call any coloring $p \in A^F$ a *pattern* of G, and denote with $\mathrm{supp}(p) := F$ the *support of* p. For every configuration $x \in A^G$ and every $h \in G$, we define its *h-translation* $h \cdot x \in A^G$ by setting $h \cdot x(g) := x(h^{-1} \cdot g)$ for each $g \in G$. We say that x *avoids* a pattern $p \in A^F$ if for every $h \in G$, $(h \cdot x)|_F \neq p$. For any finite set of patterns $\mathcal{F}$, we let $X_{\mathcal{F}}$ be the set of configurations that avoid every pattern from $\mathcal{F}$. A *subshift of finite type* (or simply SFT) is a set $X \subseteq A^G$ for which there exists some finite set of forbidden patterns $\mathcal{F}$ such that $X = X_{\mathcal{F}}$.

Remark 14.1. If $\mathcal{F}$ is a set of patterns which is not necessarily finite, then the set $X = X_{\mathcal{F}} \subseteq A^G$ of configurations of G avoiding every pattern from $\mathcal{F}$ is simply called a *subshift*. Note that subshifts exactly correspond to subsets of A^G which are G-invariant, with respect to the action of G by translation on A^G we defined above.

For every $s \in G$, an *s-pattern* is a pattern with support $\{1_G, s\}$. For short, when p is an s-pattern we will describe it with the triple $(p(1_G), s, p(s))$. If S is a finite generating set of G, a subset $X \subseteq A^G$ is a *nearest-neighbor SFT with respect to* S if there exists a set $\mathcal{F}$ of forbidden patterns such that $X = X_{\mathcal{F}}$ and such that for every $p \in \mathcal{F}$, there exists some $s \in S$ such that p is an s-pattern. In other words, a nearest-neighbor SFT with respect to S can be seen as the set of vertex-colorings of $\mathrm{Cay}(G, S)$ that avoid some forbidden patterns $\mathcal{F}$, where each pattern $p \in \mathcal{F}$ is a coloring of the two endpoints of one of the (labelled) edge $\{1_G, s\}$ in $\mathrm{Cay}(G, S)$. Note that in particular, as S is finite every nearest-neighbor SFT with respect to S is also an SFT.

A basic property of an SFT $X = X_{\mathcal{F}}$, is that it is G-invariant, with respect to the action of G on the configurations A^G we defined earlier: for every $x \in X$ and $h \in G$, if x avoids the patterns from $\mathcal{F}$, then so does $h \cdot x$, thus $x \in X$ if and only if $h \cdot x \in X$.

Example 14.2. If $A := [k]$ for some fixed $k \in \mathbb{N} \setminus \{0\}$, note that the set X_{col} of proper k-colorings of $\mathrm{Cay}(G, S)$ is a nearest-neighbor SFT with respect to S, defined by the set of forbidden patterns

$$\mathcal{F} := \{(c, s, c) : c \in A, s \in S\}.$$

We also let X'_{col} be the set of proper k-edge-colorings of $\mathrm{Cay}(G, S)$, and consider the alphabet $B := A^S$ consisting of the injective colorings of the open neighborhood of 1_G in $\mathrm{Cay}(G, S)$ with k-colors. Then X'_{col} is described by the nearest-neighbor SFT $X_{\mathcal{F}'}$ on the alphabet B where

$$\mathcal{F}' := \{(c_1, s, c_2) : c_1(s) \neq c_2(s^{-1})\}.$$

Indeed, for any proper k-edge-coloring $c : E(\mathrm{Cay}(G, S)) \to A$, we can associate a unique configuration $x \in X'_{\mathcal{F}'}$ defined as follows: for each $g \in G$ and $s \in S$, we let $x(g)(s) := c(\{g, g \cdot s\})$. Then as c is a proper coloring, $x(g) : S \to A$ is injective so $x(g) \in B$, and as c is defined on the unordered pairs of adjacent vertices, $x \in X_{\mathcal{F}'}$. Reciprocally, every configuration $x \in X'_{\mathrm{col}}$ uniquely defines a proper k-edge-coloring c of $\mathrm{Cay}(G, S)$ defined on each edge $\{g, g \cdot s\}$ by $c(\{g, g \cdot s\}) := x(g)(s)$. The definition of $\mathcal{F}'$ implies that for each $g \in G, s \in S$, $c(\{g \cdot s, (g \cdot s) \cdot s^{-1}\}) = x(g \cdot s)(s^{-1}) = x(g)(s) = c(\{g, g \cdot s\})$, hence c is well-defined.

Example 14.3. One of the most famous problems from symbolic dynamics is the *Wang tiling problem*. A *Wang tile* is an axis-parallel square with edges of size 1 drawn the Euclidian plane, whose four edges are given colors from a finite alphabet A (see Figure 2.1). In particular, we insist that a tile is embedded in the plane, so its four edges correspond to four different directions, and we do not consider tiles up to rotation. In the Wang tiling problem, we are given as input a finite set $\mathcal{W}$ of Wang tiles, and we must decide if there exists a *valid tiling* of the infinite square grid $\mathbb{G}$ using only the tiles of $\mathcal{W}$, i.e., a mapping $c : V(\mathbb{G}) \to \mathcal{W}$ such that every two tiles sharing a common edge must agree on the color of their common edge. If we let $G := (\mathbb{Z}^2, +)$ and fix a the canonical set of generators $S := \{a, b, a^{-1}, b^{-1}\}$, with $a := (1, 0), b := (0, 1)$, then a Wang tile is exactly a local coloring $c : S \to A$ of the open neighborhood of 1_G in $\mathrm{Cay}(G, S)$. In particular, if we fix a finite set $\mathcal{W} \subseteq A^S$ of Wang tiles, let $B := \mathcal{W}$ and consider the set of forbidden patterns

$$\mathcal{F} := \{(c_1, s, c_2) : c_1(s) \neq c_2(s^{-1})\},$$

then valid Wang tilings correspond exactly to configurations of the nearest-neighbor SFT $X_{\mathcal{F}}$.

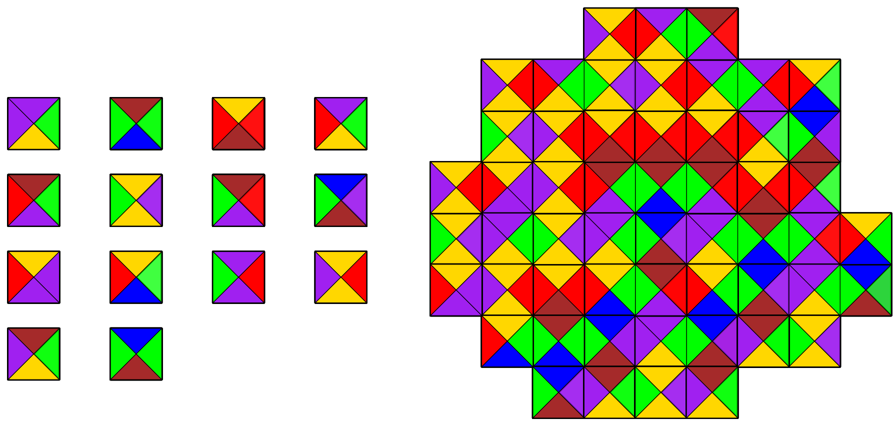

Figure 2.1: Left: a finite set of Wang tiles. Right: a partial valid Wang tiling.

A morphism between two SFTs $X \subseteq A^G$ and $Y \subseteq B^G$ is a continuous map $\sigma : X \to Y$ such that for any $x \in X, h \in G$, we have $\sigma(h \cdot x) = h \cdot \sigma(x)$. If σ is a bijection, we call it a *conjugacy* and we then say that X and Y are *conjugate*. The following result will allow us to consider nearest-neighbors SFTs in what follows, up to taking conjugate.

Proposition 14.4 (Proposition 9.3.21 in [ABJ18]). *Every SFT is conjugate to some nearest-neighbor SFT.*

14.2 Domino Problem

Given a finitely generated group G, together with a finite set S of generators, the *domino problem on* (G, S) asks, when given as input a finite set of forbidden patterns $\mathcal{F}$ which is nearest-neighbor with respect to S, if $X_{\mathcal{F}} \neq \emptyset$. For any two finite generating sets S, S', it is known [ABJ18, Proposition 9.3.26] that the domino problem on (G, S) is decidable if and only if the domino problem on (G, S') is decidable. Thus we will only talk about the domino problem on G, without precising the finite set of generators S.

A central result in symbolic dynamics is Berger's theorem [Ber66], which states that the Wang tiling problem that we introduced in the previous subsection is undecidable. In particular, due to the observations from Example 14.3, it immediately implies that the domino problem is undecidable on $(\mathbb{Z}^2, +)$. On the other hand, there is a simple greedy procedure to solve the domino problem on free groups. More generally, the domino problem is also decidable in virtually free groups [ABJ18, Theorem 9.3.37]. A remarkable conjecture of Ballier and Stein [BS18] asserts that these groups are the only ones for which the domino problem is decidable.

Conjecture 14.5 (Domino problem conjecture [BS18]). *A finitely generated group has a decidable domino problem if and only if it is virtually free.*

Conjecture 14.5 has been proved for many group classes. Among others, it holds for polycyclic groups, surface groups and more generally hyperbolic groups [Jea15a, ABM19, Bar23b]. From a graph theoretic perspective, this conjecture seems quite natural: Berger's theorem asserts that the domino problem is undecidable for grids, and this conjecture states that there should be a way to generalize Berger's result to prove that any group whose Cayley graph admits the infinite grid as a minor (and thus has unbounded treewidth) has also undecidable domino problem (recall that by [GH24], any Cayley graph of infinite treewidth admits the infinite grid as a minor). From this point of view, a natural question is whether Conjecture 14.5 holds for K_∞-minor-free groups. We prove in Theorem 14.7 below that this is indeed the case, as an almost immediate consequence of Theorem 13.3. In particular, it generalizes and unifies the already known results we mentioned above from [Ber66, ABM19] stating that the domino problem is undecidable on $\mathbb{Z}^2$ and more generally on surface groups, and implies that Conjecture 14.5 holds for planar groups. To prove this, we will need the following result of Thomassen, which is tightly related to Theorem 7.7, and which will allow us to reduce the K_∞-minor-free case to the planar case.

Theorem 14.6 (Theorem 5.7 in [Tho92])**.** *Let G be a locally finite, transitive, non-planar, one-ended graph. Then G contains the countable clique K_∞ as a minor.*

Theorem 14.7. *Let G be a finitely generated group excluding the countable clique K_∞ as a minor. Then the domino problem is undecidable on G if and only if*

- *G is one-ended, or*
- *G has an infinite number of ends and has a one-ended planar subgroup which is finitely generated.*

In particular, these situations correspond exactly to the cases where G is not virtually free.

Proof. Let G be a finitely generated K_∞-minor-free group and let S be a finite set of generators such that $\mathrm{Cay}(G, S)$ excludes K_∞ as a minor. If G has 0 end, then it is finite and the domino problem is decidable. If G has 2 ends then G is virtually $\mathbb{Z}$ and the domino problem is also decidable. If G is one-ended, then since G is transitive and excludes the countable clique K_∞ as a minor, it follows from Theorem 14.6 that G is (one-ended) planar. By Theorem 13.2, G contains the surface group Γ_g of genus $g \geqslant 1$ as a subgroup. By [ABM19, Corollary 18], the domino problem is undecidable for Γ_g. It is known that for every finitely generated subgroup H of a finitely generated group G, if the domino problem is undecidable for H then it is also undecidable for G [ABJ18, Proposition 9.3.30]. This implies that if G is one-ended, then the domino problem is undecidable for G.

Assume now that G has an infinite number of ends. By Theorem 13.3, G is isomorphic to a finite fundamental graph of groups with finite edge-groups, and whose vertex-groups are planar with at most one end. If G is not virtually free, then by a result of Karrass, Pietrowski and Solitar [KPS73] one of the (finitely presented) vertex-groups G_u (which is a subgroup of G) in is one-ended. By the paragraph above, the domino problem is undecidable for G_u, and since G_u is a subgroup of G, the domino problem is also undecidable for G. □

In [Mac23], MacManus established the following structural characterisation of finitely generated groups admitting at least one (and thus all) Cayley graph that is quasi-isometric to a planar graph.

Theorem 14.8 (Corollary D in [Mac23])**.** *A finitely generated group G admits a Cayley graph that is quasi-isometric to a planar graph if and only if G is virtually a free product of finitely many free groups and surface groups.*

Note that Theorem 14.8 immediately implies that Theorem 14.7 still holds if we replace the hypothesis for G to be K_∞-minor-free by the hypothesis that G admits a Cayley graph that is quasi-isometric to a planar graph, which by Theorem 8.7 is a more general assumption.

15 Aperiodic subshifts of finite type

Let A be a finite alphabet. Then a configuration $x \in A^{\mathbb{Z}}$ is nothing more than a sequence of elements of A indexed by $\mathbb{Z}$, and a natural notion of periodicity emerges: we say that x is periodic if there exists some period $P \in \mathbb{N} \setminus \{0\}$ such that for each $i \in \mathbb{Z}$, $x(i+P) = x$, or in other words with notations from the previous subsection if $P \cdot x = x$. It is an easy exercise to see that any nonempty SFT on $\mathbb{Z}$ admits a periodic configuration, and a consequence is the existence of a very natural algorithm to solve the domino problem on $\mathbb{Z}$. For more general groups than $\mathbb{Z}$, there are different notions of periodicity, and we will see that in general it is not true anymore that nonempty SFTs always admit periodic configurations.

In subsection 15.1, we will introduce the various notions of periodicity and aperiodicity that are usually studied in symbolic dynamics, and give a positive answer to a Conjecture of Carroll and Penland (see Conjecture 15.7 below) in the special case of K_∞-minor free groups. In subsection 15.2, we will study the (non-)equivalence between two different notions of aperiodicity in some restricted classes of groups. The results and proofs from subsection 15.2 come from a joint work in progress with Etienne Moutot and Solène Esnay.

15.1 Periodicity and aperiodicity

Weak and strong periodicity. We use here the definitions of periodicity and aperiodicity from [CP15]. We let G be a finitely generated group and A a finite alphabet. A configuration $x \in A^G$ is *weakly periodic* if $\mathrm{Stab}_G(x) \neq \{1_G\}$, and x is *strongly periodic* if $\mathrm{Stab}_G(x)$ has finite index in G, or equivalently if the orbit $G \cdot x$ of x is finite.

A SFT $X \subseteq A^G$ is *weakly aperiodic* if it does not contain any strongly periodic configuration. Similarly, X is said to be *strongly aperiodic* if it does not contain any weakly periodic configuration.

A finitely generated group G is *weakly periodic* (respectively strongly periodic) if every nonempty SFT $X \subseteq A^G$, contains at least one weakly (respectively strongly) periodic configuration. In other words, a group is weakly (respectively strongly) periodic if and only if it does not admit a strongly (respectively weakly) aperiodic SFT.

Remark 15.1. Let $x \in A^G$. Then x is strongly periodic if and only if G induces a quasi-transitive group action on the colored graph $(\mathrm{Cay}(G, S), x)$. This follows from the fact that the elements of G which induce automorphisms of the colored graph $(\mathrm{Cay}(G,S),x)$ are exactly the elements from $\mathrm{Stab}_G(x)$. Thus the G-orbits of the vertices of the colored graph $(\mathrm{Cay}(G,S),x)$ correspond to the different right-cosets $\{\mathrm{Stab}_G(x)\cdot g : g\in G\}$.

Remark 15.2. If X and Y are two conjugate SFTs on G, then it is not hard to check that X is weakly (respectively strongly) aperiodic if and only if Y is weakly (respectively strongly) aperiodic.

Lift. If G,H are finitely generated groups such that H is a subgroup of G, and if $X := X_{\mathcal{F}} \subseteq A^H$ is an SFT on H with respect to a finite set $\mathcal{F}$ of forbidden patterns, then its *lift* $X^{\uparrow}\subseteq A^G$ is the SFT on G defined by the same set of forbidden patterns $\mathcal{F}$, where we consider $\mathcal{F}$ as a set of patterns of G (recall that $H\subseteq G$). It was observed in [Jea15b] that taking lifts preserves the property of being weakly aperiodic.

Proposition 15.3 (Proposition 1.1 in [Jea15b])**.** *Let G,H be two finitely generated groups such that H is a subgroup of G, and let X be an SFT on H. If X is weakly aperiodic, then $X^{\uparrow}$ is weakly aperiodic.*

Remark 15.4. Note that Proposition 15.3 does not hold anymore in general if we consider strong aperiodicity instead of weak aperiodicity: if H is any finitely generated group that admits a strongly aperiodic SFT $X\subseteq A^H$ (one can take for example $H:=\mathbb{Z}^2$ [Ber66]), and if $G:=H\times(\mathbb{Z}/k\mathbb{Z})$ for some $k\geqslant 2$, then it is not hard to see that $X^{\uparrow}\subseteq A^G$ is not strongly aperiodic anymore: let $x\in X$ and define the configuration $x^{\uparrow}\in A^G$ by setting for each $(g,i)\in G$, $x^{\uparrow}(g,i):=x(g)$. Then by definition of $X^{\uparrow}$, we have $x^{\uparrow}\in X^{\uparrow}$, but for any $i\in(\mathbb{Z}/k\mathbb{Z})\setminus\{0\}$, the element $g:=(1_H,i)\in G\setminus\{1_G\}$ is such that $g\cdot x^{\uparrow}=x^{\uparrow}$, implying that $x^{\uparrow}$ is a weakly periodic configuration.

Barbieri [Bar23a, Proposition 3] recently gave necessary and sufficient conditions for a lift $X^{\uparrow}$ of an SFT to be strongly aperiodic. For our purpose, we only extract the following result:

Theorem 15.5 (Corollary of Proposition 3 in [Bar23a])**.** *Let G,H be finitely generated groups such that H is a subgroup of G and let X be an SFT on H. Assume that there exists some element $g\in G\setminus\{1_G\}$ such that for all $n>0$ and $t\in G$,*

$$tg^nt^{-1}\notin H\setminus\{1_H\}.$$

Then $X^{\uparrow}$ is not strongly aperiodic on G.

Carroll and Penland's conjecture. It is well known that $\mathbb{Z}$ is strongly periodic, i.e., that it admits no weakly aperiodic SFT. Indeed, by Proposition 14.4 and Remark 15.2, it is enough to show that $\mathbb{Z}$ admits no weakly aperiodic nearest-neighbor SFT, which is an easy exercise. We recall that every 2-ended finitely generated group contains $\mathbb{Z}$ as a subgroup of finite index. Using [CP15, Theorem 2] which states that strong periodicity is a commensurability invariant, we immediately get the following result.

Proposition 15.6. *Every 2-ended finitely generated group G is strongly periodic.*

In fact, Carroll and Penland conjectured that a converse of Proposition 15.6 should hold.

Conjecture 15.7 ([CP15])**.** *A finitely generated group is strongly periodic if and only if it is virtually cyclic.*

As for the domino problem conjecture, Conjecture 14.5 is natural from a graph theoretic point of view: Piantadosi [Pia08] proved that non-abelian free groups admit weakly aperiodic SFT, and thus that they are not strongly periodic. In particular, non-abelian free groups are exactly the groups admitting a regular infinite tree as a Cayley graph of degree at least 4. From this perspective, Conjecture 15.7 asks if this aperiodicity property still holds if we relax this hypothesis by simply assuming that G has at least one (and thus all) Cayley graph which contains the infinite regular tree of degree at least 3 as a minor (which by Lemma 4.1 is equivalent to not being virtually cyclic).

Putting together Theorem 15.6 with our remarks from Example 14.2 and Remark 15.1, we immediately obtain the following remark which gives a positive answer to Problem 10.1 for 2-ended Cayley graphs (which we already proved more generally for 2-ended quasi-transitive graphs by Theorem 10.3).

Remark 15.8. For any 2-ended finitely generated group G and every finite generating set S, if $k := \chi(\mathrm{Cay}(G, S))$ (respectively $k := \chi'(\mathrm{Cay}(G, S))$) then there exists a proper k-coloring (respectively k-edge-coloring) c of $\mathrm{Cay}(G, S)$ such that $(\mathrm{Cay}(G, S), c)$ is quasi-transitive.

Again, as for the domino problem conjecture, our previous structural results give a positive answer to Conjecture 15.7 for K_∞-minor-free groups, and more generally for groups with Cayley graphs which are quasi-isometric to a planar graph. We directly prove it for the second class of groups.

Theorem 15.9. *Let G be a finitely generated group and S a finite set of generators such that $\mathrm{Cay}(G, S)$ is quasi-isometric to a planar graph. Then G admits a weakly aperiodic SFT if and only if G has one or infinitely many ends.*

Proof. By Theorem 14.8, G is virtually a free product of finitely many free groups and surface groups. Recall that Carroll and Penland [CP15, Theorem 2] proved that if G, H are two finitely generated commensurable groups, then G admits a weakly aperiodic SFT if and only if H does, hence we may assume without loss of generality that G is just a free product of finitely many free groups and surface groups. Berger [Ber66] proved in his seminal work that $\mathbb{Z}^2$ admits a weakly aperiodic SFT, and Cohen and Goodman-Strauss [CGS17] proved that every surface group of genus $g \geqslant 2$ also has a weakly aperiodic SFT. By Proposition 15.3, if some subgroup of G has a weakly aperiodic SFT, so does G. In particular, if some of the factors involved in the decomposition of G as a free product is an infinite surface group, i.e., a surface group Γ_g with $g \geqslant 1$, then G must admit a weakly aperiodic SFT. Assume that G admits such a subgroup. We show that in this case, G has either one or infinitely many ends. Let H be a subgroup of G isomorphic to a surface group of genus $g \geqslant 1$. By Theorem 12.1, there exists a finite set of generators S' of H such that $\mathrm{Cay}(H, S')$ is a minor of $\mathrm{Cay}(G, S)$. As surface groups are one-ended, by [Tho92, Proposition 5.6] the end of H is thick. In particular, as $\mathrm{Cay}(H, S')$ is a minor of $\mathrm{Cay}(G, S)$, the graph $\mathrm{Cay}(G, S)$ must have at least one thick end. It implies that G cannot have 0 or 2 ends.

To conclude the proof, it remains to prove that the desired result holds if G is a free product of finitely many free groups. In this special case, G must be itself virtually free. In particular, if G is not virtually cyclic, it must contain a non-abelian free group as a subgroup, which by [Pia08] contains a weakly aperiodic SFT. Applying again Proposition 15.3, we conclude that in this case, G also has a weakly aperiodic SFT. □

15.2 Separating weak and strong aperiodicity

Weakly but not strongly aperiodic SFTs. In view of the definitions from the previous subsection, a very natural question is the following: which groups G admit a weakly but not strongly aperiodic SFT? First, observe that such SFTs cannot exist in virtually cyclic groups, as we observed previously that such groups do not even have a weakly aperiodic SFT. A folklore result in tiling theory is that $\mathbb{Z}^2$ also enjoys this property, i.e., that any weakly aperiodic SFT X on $\mathbb{Z}^2$ is also strongly aperiodic. However, this is not true anymore when considering $\mathbb{Z}^d$ for $d \geqslant 3$, which admits a weakly but not strongly aperiodic SFT. Among other, one can also find weakly but not strongly aperiodic SFTs in the following cases:

- If G has infinitely many ends, then on the one hand, by a result of Cohen [Coh17, Theorem 1.5] it cannot have a strongly aperiodic SFT. On the other hand, G has a non-abelian free subgroup H of finite rank (see for example [AMO07, Corollary 1.3]), and by [Pia08], H admits a weakly aperiodic SFT X. Hence by Proposition 15.3, $X^{\uparrow}$ is a weakly but not strongly aperiodic SFT in G.
- Based on a construction of Aubrun and Kari [AK13], Moutot and Esnay [EM22] constructed a weakly but not strongly aperiodic SFT in Baumslag Solitar groups $\mathrm{BS}(m, n)$ for any $m, n \geqslant 1$.
- If $G = \mathbb{Z}^2 \times (\mathbb{Z}/k\mathbb{Z})$ for some $k \geqslant 2$, then G admits a weakly but not strongly aperiodic SFT: let X be a weakly aperiodic SFT on $\mathbb{Z}^2$ given by [Ber66]. We claim that its lift $X^{\uparrow}$ in G is weakly but not strongly aperiodic: by Proposition 15.3, $X^{\uparrow}$ is weakly aperiodic. Moreover, Remark 15.4 immediately implies that $X^{\uparrow}$ is not strongly aperiodic.
- In fact, one can show that if G is virtually $\mathbb{Z}^2$, then it admits a weakly but not strongly aperiodic SFT if and only if it is torsion-free, generalizing the observation from the previous item.

Apart from the few examples we gave, we do not know if there exists any other finitely generated group admitting a weakly but not strongly aperiodic SFT. In a similar fashion to Conjectures 14.5 and 15.7, Nicolas Bitár recently conjectured that these cases should be the only ones for which we can find a weakly but not strongly aperiodic SFT.

Conjecture 15.10 ([Bit24]). *Let G be a finitely generated group. Then G admits a weakly but not strongly aperiodic SFT if and only if it is neither virtually cyclic, nor torsion-free virtually $\mathbb{Z}^2$.*

We prove in the next two theorems that Conjecture 15.10 holds for groups admitting a presentation with one relation and at least 3 generators, and for groups having a Cayley graph which is quasi-isometric to a planar graph which are neither virtually cyclic, nor virtually $\mathbb{Z}^2$.

Theorem 15.11. *Let G be a group which is not virtually cyclic and $\langle S \mid r\rangle$ be a presentation of G such that $|S| \geqslant 3$ and r is a cyclically reduced non-empty word. Then G admits a weakly but not strongly aperiodic SFT.*

Observe that as the group $\mathbb{Z}^2$ admits the presentation $\langle a, b \mid ab\overline{a}\overline{b}\rangle$, we cannot relax the condition $|S| \geqslant 3$ in Theorem 15.11. Finitely generated groups admitting a presentation $\langle S \mid R\rangle$ such that S is finite and $|R| = 1$ are usually called *one-relator groups* in the literature and form a rich class of group which attracted a lot of attention. We will only use a few known results about them and refer the interested reader to [BF95] for further results and questions. In particular, note that for every $g \geqslant 2$, the surface group Γ_g of genus g satisfies the conditions of Theorem 15.11, implying immediately the following.

Corollary 15.12. *For every $g \geqslant 2$, the surface group Γ_g admits a weakly but not strongly aperiodic SFT.*

According to the discussion at the beginning of [CGSR22], the existence of weakly but not strongly aperiodic SFTs in hyperbolic groups should follow from [Gro87, Paragraph 8.4] and [CP93], implying in particular Corollary 15.12.

Combining Corollary 15.12 together with MacManus' structure theorem (Theorem 14.8), we will show that Conjecture 15.10 holds in the special case of K_∞-minor-free groups, and more generally for groups having a Cayley graph which is quasi-isometric to a planar graph.

Theorem 15.13. *Let G be a group which is not virtually H for every $H \in \{\{1_G\}, \mathbb{Z}, \mathbb{Z}^2\}$, and S be a finite set of generators of G such that $\mathrm{Cay}(G, S)$ is quasi-isometric to a planar graph. Then G admits a weakly but not strongly aperiodic SFT.*

Subgroups of finite index. Note that because of the previous example $\mathbb{Z}^2 \times (\mathbb{Z}/k\mathbb{Z})$, the property of admitting a weakly but not strongly aperiodic SFT is not a commensurability invariant in general, as $\mathbb{Z}^2$ does admit such an SFT. Nevertheless, we prove that one direction still holds, namely that if some subgroup of finite index in G admits a weakly but not strongly aperiodic SFT, then it is also the case for G.

Lemma 15.14. *Let G be a finitely generated group and H be a finitely generated subgroup of G of finite index. If H has a weakly but not strongly aperiodic SFT then so does G.*

Note that the SFT we construct in the following proof does not consist in just taking the lift of a weakly but not strongly aperiodic SFT on H. Our construction is quite close to the higher block shift defined in [CP15], but we are not aware of any previous proof of this particular result.

Proof. We let S_H be a finite set of generators for H and $X \subseteq A^H$ be a weakly but not strongly aperiodic SFT of H on some finite alphabet A. By Proposition 14.4, X is conjugate to a nearest-neighbor subshift, so as the property of being weakly or strongly aperiodic is preserved under conjugation, we may assume without loss of generality that X is a nearest-neighbor subshift with respect to S_H, and let $\mathcal{F}_H$ denote its associated set of forbidden patterns. We let $g_1, \dots, g_k \in G$ be representatives of the different right cosets $H \cdot g_1, \dots, H \cdot g_k$ of H with $k := [G : H] \in \mathbb{N} \setminus \{0\}$ and $B := A \times [k]$. Assume that $g_1 = 1_G$. For each $a \in S_H$ and $i \in [k]$, we set $a_i := g_i^{-1} a g_i \in G$ and we consider the finite set $S := \{a_i : a \in S_H, i \in [k]\}$. Recall that for every $a \in G$, an a-pattern is a pattern with support $\{1_G, a\}$. We define a set $\mathcal{F}$ of forbidden patterns in G with support $\{1_G, a_i\}$ for each $a_i \in S$, using the alphabet B, as follows:

1. for each $a \in S_H$ and each a-pattern $p \in \mathcal{F}_H$, we add in $\mathcal{F}$ the a_i-pattern p' of G defined by $p'(1_G) := (p(1_H), i)$ and $p'(a_i) := (p(a), i)$ for each $i \in [k]$;
2. for each $a \in S_H, c_1, c_2 \in A, i \in [k]$ and $j \in [k] \setminus \{i\}$ we add in $\mathcal{F}$ the forbidden a_i-pattern q of G defined by $q(1_G) := (c_1, i)$ and $q(a_i) := (c_2, j)$.

We consider $Y := X_{\mathcal{F}} \subseteq B^G$ the associated SFT and we will show that Y is weakly aperiodic. Intuitively, the forbidden patterns (1) will be useful to “simulate” in each coset the dynamics of X, while the forbidden patterns (2) will allow us to find configurations in Y in which the elements of some coset $H \cdot g_i$ all have their second coordinate colored i (see our remark below).

Let $\pi_1 : B \to A$ and $\pi_2 : B \to [k]$ be the projections on the first and second coordinates of B, i.e., $(c, i) = (\pi_1(c, i), \pi_2(c, i))$ for each $(c, i) \in B$.

First, let us remark that for any $y \in Y$, if there exist $h \in H$ and $i \in [k]$ so that $\pi_2(y(hg_i)) = i$, then $\pi_2(y)$ is constant along the coset $H \cdot g_i$. Indeed, if $\pi_2(y(hg_i)) = i$, the forbidden patterns (2) impose that for each $a \in S_H$, $\pi_2(y(hg_i a_i)) = i$. Moreover, for all $h' \in H, h' g_i a_i = (h'a) g_i$ by definition of a_i. So as S_H is a generating set of H, we can prove by induction on $|h^{-1}h'|_{S_H}$ that $\pi_2(y(h'g_i)) = i$ for all $h' \in H$.

Y is weakly aperiodic. We let $y \in Y$ and show that the G-orbit of y is infinite. Let $i \in [k]$ be such that $y(1_G) = (c, i)$ for some $c \in A$. Then, the translated configuration $y' := {g_i}^{-1} \cdot y \in Y$ is such that $y'(g_i) = (c, i)$. As y' is in the G-orbit of y, it is enough to prove that its G-orbit is infinite so without loss of generality we may assume that $y(g_i) = (c, i)$ for some $c \in A$. In particular, the remark above implies that $\pi_2(y(hg_i)) = i$ for all $h \in H$.

Let us consider the configuration $x^{(y)} \in A^H$ defined for every $h \in H$ by $x^{(y)}(h) := \pi_1(y(hg_i))$.

We show that $x^{(y)} \in X$. For this, let $h \in H, a \in S_H$, and $c_1, c_2 \in A$ be such that $x^{(y)}(h) = c_1$ and $x^{(y)}(ha) = c_2$. We let p be the a-pattern of H defined by $p(1_H) := c_1$ and $p(a) := c_2$, and show that $p \notin \mathcal{F}_H$ (recall that $X = X_{\mathcal{F}_H}$).

Then, as $\pi_2(y(hg_i)) = \pi_2(y(hag_i)) = \pi_2(y(g_i)) = i$, we have $y(hg_i) = (c_1, i)$ and $y(hg_i a_i) = y(hag_i) = (c_2, i)$ so the a_i-pattern p_i of G defined by $p_i(1_G) := (c_1, i)$ and $p_i(a_i) := (c_2, i)$ cannot belong to $\mathcal{F}$. By definition of $\mathcal{F}$, p_i cannot have type (1) so in particular $p \notin \mathcal{F}_H$. We just proved that for all $h \in H$ and $a \in S_H$, $(x^{(y)}(h), a, x^{(y)}(ha)) \notin X$, so as X is nearest-neighbor with respect to S_H, it implies that $x^{(y)} \in X$.

Note that for every $h, h' \in H$,

$$x^{(h \cdot y)}(h') = \pi_1(h \cdot y(h'g_i)) = \pi_1(y(h^{-1}h'g_i)) = x^{(y)}(h^{-1}h') = (h \cdot x^{(y)})(h'),$$

thus $x^{(h \cdot y)} = h \cdot x^{(y)}$ and in particular $h \cdot x^{(y)} \in X$ for each $h \in H$. As X is weakly aperiodic, $x^{(y)}$ has an infinite H-orbit, and the previous equality thus implies that y also has an infinite H-orbit. In particular, y has an infinite G-orbit, showing that Y is weakly aperiodic.

Y is nonempty and not strongly aperiodic. As X is not strongly aperiodic, there exist $x \in X$ and $h_0 \in H \setminus \{1_H\}$ such that $h_0 \cdot x = x$. We define $y \in B^G$ by setting for every $i \in [k]$ and $h \in H$, $y(hg_i) := (x(h), i)$.

First, we show that $y \in Y$. Let $p = ((c_1, i), a_\ell, (c_2, j))$ be any a_ℓ-pattern of y for some $a_\ell \in S$ and let $g \in G$ be such that $(p(1_G), p(a_\ell)) = (y(g), y(ga_\ell))$. We will show that $p \notin \mathcal{F}$.

Assume first that $i \neq j$. Then p is not of type (1) and as for each $a \in S_H$ and $g \in H \cdot g_i$, we have $ga_i \in H \cdot g_i$, we must have $\ell \neq i$ so p is not of type (2) and we indeed get that $p \notin \mathcal{F}$. Assume now that $i = j$ and write $g = hg_i$ for some $h \in H$. Then p is not of type (2) and if $\ell \neq i$, we immediately get $p \notin \mathcal{F}$. We thus assume that $\ell = i = j$. We then have by definition of y that $x(h) = \pi_1(y(hg_i)) = c_1$ and $x(ha) = \pi_1(y(hag_i)) = \pi_1(y(hg_ia_i)) = c_2$ so as $x \in X_H = X_{\mathcal{F}_H}$, the a-pattern $p' := (c_1, a, c_2)$ of H is not in $\mathcal{F}_H$. In particular, we just proved that p is not of type (1), implying that $p \notin \mathcal{F}$.

Finally, let us show that h_0 is a non-trivial period of y. For all $g \in G$, write $g = hg_i$. Then we have

$$h_0 \cdot y(hg_i) = y(h_0^{-1}hg_i) = (x(h_0^{-1}h), i) = (x(h), i) = y(hg_i),$$

implying that $h_0 \cdot y = y$. We thus conclude that X is not strongly aperiodic, as desired. □

Proof of Theorems 15.11 and 15.13. We call a finitely generated group a *one-relator group* if it admits a presentation $\langle S \mid R \rangle$ such that S is finite and $|R| = 1$. To prove Theorem 15.11, we will need a few results about the structure of one-relator groups. In particular, the following result plays a central role in the theory of one-relator groups.

Theorem 15.15 (Freiheitssatz [Mag30]). *Let $G = \langle S \mid r \rangle$ be a one-relator group with r cyclically reduced, such that $s \in S$ appears in r. Then the group generated by $S \setminus \{s\}$ is free of rank $|S| - 1$.*

In the remainder of the proof, we will use the following notation: for every word w written on the alphabet $A := (S \uplus \overline{S})^*$ and every letter $s \in A$, we let $|w|_s \in \mathbb{Z}$ denote the total number of occurences of s in w minus the total number of occurences of $\overline{s}$ in w. Note in particular that for every $w \in A^*$ and every $s \in S$, we have $|w|_s = -|w|_{\overline{s}}$. The next lemma will allow us to conclude when there exists some $s \in S$ such that the unique relation r defining G satisfies $|r|_s = 0$.

Lemma 15.16. *Let G be a group, $S = \{a, b, c, \dots\}$ be a finite set of generators of G of size at least 3 and r be a cyclically reduced word such that $\langle S \mid r \rangle$ is a presentation of G. Assume that $c \in S$ appears in r and $|r|_c = 0$. Then G admits a weakly but not strongly aperiodic SFT.*

Proof. Let B be the alphabet of X, i.e., such that $X \subseteq B^G$. By Theorem 15.15, $\langle a, b \rangle$ is free in G. Thus by [Pia08], it admits a weakly aperiodic SFT X, and Proposition 15.3 implies that $X^\uparrow$ is a weakly aperiodic SFT on G. In the remainder of the proof, we set $A := (S \uplus \overline{S})^*$.

Now, our aim is to prove that $X^\uparrow$ is not strongly aperiodic using Theorem 15.5. We first show that for every $g \in G$, and every word $w \in A^*$ representing g, the value $|w|_c$ does not depend of the choice of w. Indeed, for every two words $w, w' \in A^*$ representing g, the word $\overline{w}w'$ represents the identity element 1_G. If we let w'' denote the cyclically reduced form of $\overline{w}w'$, then there exist $w_1, \dots, w_k \in A^*$ such that

$$w'' = (w_1 r \overline{w_1}) \cdot \ldots \cdot (w_k r \overline{w_k}).$$

In particular, as $|r|_c = 0$, we get $|w''|_c = 0$. Note that for every two words $w_1, w_2 \in A^*$ such that w_2 is obtained after performing an elementary cyclic reduction on w_1, we must have $|w_1|_s = |w_2|_s$ for all $s \in S$. In particular, it implies that $|\overline{w}w'| = |w''|_c = 0$. As $|\overline{w}w'|_c = |\overline{w}|_c + |w'|_c$, we deduce that $|w|_c = -|\overline{w}|_c = |w'|_c$.

Hence for every $g \in G$ and every word $w \in A^*$ representing g, the integer $|g|_c := |w|_c \in \mathbb{Z}$ is well defined. We now let $w_0 := c$, and let g_0 be the group element represented by c. Then for every $t \in G$ and $n > 0$, we have

$$|t g_0^n t^{-1}|_c = n > 0.$$

In particular, as every element g from $\langle a, b \rangle$ satisfies $|g|_c = 0$, we obtain that $t g_0^n t^{-1}$ does not belong to $\langle a, b \rangle$ for every $t \in G, n > 0$. Applying Theorem 15.5, we conclude that $X^\uparrow$ is not strongly aperiodic. □

The next lemma is inspired by the Magnus-Moldavansky rewriting method, a technique used to decompose one-relator groups by means of non-trivial HNN-extensions of simpler one-relator groups. It will be of particular interest for us as it allows to construct group presentations satisfying the properties from Lemma 15.16 starting from any presentation $\langle S \mid R \rangle$ with $|S| \geqslant 3$ and $|R| = 1$.

Lemma 15.17. *Let G be a one-ended group, S be a finite set of generators and $r \in (S \uplus \overline{S})^*$ be a cyclically reduced word such that $\langle S \mid r \rangle$ is a presentation of G. Then there exist some presentation $\langle T \mid r' \rangle$ and some element $t \in T$ such that:*

- $|T| = |S|$,
- *r' is cyclically reduced,*
- *t occurs in r',*
- $|r'|_t = 0$,

- *$\langle T \mid r' \rangle$ is a presentation of G.*

Proof. First, note that as G is one-ended, every element of S must occur at least once in r. If some element $s \in S$ satisfies $|r|_s = 0$, then we conclude choosing $T := S$, wo we may assume that there is no such element. Up to replacing elements of S by their formal inverse from $\overline{S}$, we also assume that for each $s \in S$, $|r|_s > 0$, and write $S = \{s_1, \ldots, s_n\}$ so that

$$0 < |r|_{s_1} \leqslant |r|_{s_2} \leqslant \cdots \leqslant |r|_{s_n}.$$

We set $t_1 := s_1 s_n$ and $t_i := s_i$ for every $i > 1$. We now let r' be the word obtained after rewriting r using the family $T := \{t_1, \ldots, t_n\}$, i.e., after replacing each occurence of s_i (respectively of $\overline{s_i}$) in r by t_i (respectively by $\overline{t_i}$) for $i > 1$, and after replacing each occurence of s_1 (respectively of $\overline{s_1}$) by $t_1 \overline{t_n}$ (respectively by $t_n \overline{t_1}$). Note that $\langle T \mid r' \rangle$ is obtained from $\langle S \mid r \rangle$ after performing the following Tietze transformations:

1 Add new generators $t_1, \ldots, t_n$ together with the relations $t_1 \overline{s_n s_1}$ and $t_i \overline{s_i}$ for each $i > 1$;

2 Remove the letter s_1 and replace its occurences with $t_1 \overline{s_n}$;

3 Remove the letter s_i for each $i > 1$ and replace its occurences with t_i;

4 Remove the redundant relations $t_1 \overline{t_n} t_n \overline{t_1}$ and $t_i \overline{t_i}$ for each $i > 1$.

Hence $\langle T \mid r' \rangle$ is also a presentation of G. Moreover, note that

$$0 < |r'|_{t_n} = |r|_{s_n} - |r|_{s_1} < |r|_{s_n},$$

while for any $i > 1$, $|r'|_{t_i} = |r|_{s_i}$. Note that if the cyclically reduced form of r' is the empty word, then it means that G is a free group, which is impossible as we assumed that G is one-ended. Hence iterating this process at most $n \cdot |r|_{s_n}$ times, we end with a presentation $\langle T \mid r' \rangle$ such that r' is a reduced non-empty word and such that some element $t \in T$ occuring in r' satisfies $|r'|_t = 0$. In particular, $\langle T \mid r' \rangle$ satisfies the desired properties. □

We now prove Theorem 15.11: let G be a group which is not virtually cyclic and let $\langle S \mid r \rangle$ be a presentation of G such that $|S| \geqslant 3$ and such that r is cyclically reduced. Recall that infinitely ended groups admit weakly but not strongly aperiodic SFTs. Hence as G is not virtually cyclic, we immediately get the desired result if G has 0 or infinitely many ends. We assume now that G is one-ended, and let $\langle T \mid r \rangle$ and $t \in T$ be given by Lemma 15.17. Then as $|T| = |S| \geqslant 3$, we can apply Lemma 15.16 to the presentation $\langle T \mid r \rangle$ with t playing the role of c, which gives a weakly but not strongly aperiodic SFT. We thus proved Theorem 15.11.

To prove Theorem 15.13, we let G be a finitely generated group having a Cayley graph which is quasi-isometric to a planar graph, and we assume that G is not virtually H for $H \in \{\{1_G\}, \mathbb{Z}, \mathbb{Z}^2\}$. If G has infinitely many ends, again we already saw that G admits a weakly but not strongly aperiodic SFT. We thus assume that it is not the case, i.e., that G has exactly one end. Then by Theorem 14.8, and because G is one-ended, G is a virtually surface group. In particular, as G is neither finite nor virtually $\mathbb{Z}^2$, G must be virtually Γ_g for some $g \geqslant 2$, and thus in particular, G is also virtually Γ_2. Recall that Γ_2 has the group presentation $\langle a,b,c,d \mid [a,b][c,d]\rangle$, thus in particular, this presentation satisfies the conditions of Theorem 15.11, and there exists a weakly but not strongly aperiodic SFT in Γ_2. Then, by Lemma 15.14, G must also admits a weakly but not strongly aperiodic SFT, which concludes the proof of Theorem 15.13. □

16 SFTs on quasi-transitive graphs

In this section, and until the end of the manuscript, we will adopt again our initial convention, i.e., we will denote graphs with latin capital letters and groups with capital greek letters.

The goal of this section is to propose natural generalizations of the notions and questions from symbolic dynamics we introduced in the previous subsections, to locally finite quasi-transitive graphs. Our hope is that it could help to generalize some results from symbolic dynamics on quasi-transitive graphs. As a first example of application, we give an alternate proof of a generalization of Theorem 10.3, as an immediate corollary of Proposition 15.6.

Subshifts of finite type in quasi-transitive graphs. We start with a definition of SFT for quasi-transitive graphs. Let G be a locally finite graph and Γ be a group with a quasi-transitive action on G. Then Γ also induces a quasi-transitive group action on $E(G)$, and we let $S := \{e_1, \dots, e_k\}$ be a a set of representatives of its different Γ-orbits. We write $e_i = \{u_i, v_i\}$ for each $i \in [k]$ and let A be a finite alphabet. We call a mapping $x : V(G) \to A$ a *configuration* of G and for each $g \in \Gamma$, we define a configuration $g \cdot x \in A^{V(G)}$ by setting $(g \cdot x)(v) := x(g^{-1} \cdot v)$ for each $v \in V(G)$. As before, for every finite set of vertices $F \subseteq V(G)$, a mapping $p : F \to A$ is called a *pattern*, and we denote with $\mathrm{supp}(p) := F$ the *support of* p. We say that a configuration $x : V(G) \to A$ *avoids* p if for every $g \in \Gamma$, $(g \cdot x)|_{\mathrm{supp}(p)} \neq p$. Note that the property for a configuration x to avoid a given pattern p does not depend of the choice of the set S of representatives of $E(G)/\Gamma$. If $\mathcal{F}$ is a finite set of patterns, again we let $X_{\mathcal{F}}$ be the set of configurations of $A^{V(G)}$ avoiding all the patterns from $\mathcal{F}$. We call a subset $X \subseteq A^{V(G)}$ a *subshift of finite type* (or just SFT) of (G,Γ) if there exists some finite set of forbidden patterns $\mathcal{F}$ such that $X = X_{\mathcal{F}}$. If moreover, every pattern from $\mathcal{F}$ has a support of the form $\{u_i, v_i\}$ for some $i \in [k]$, then we say that X is *nearest-neighbor*.

Example 16.1.

- If Γ is a finitely generated group with a finite set of generators S, and if we consider the action of Γ by left-multiplication on $G := \mathrm{Cay}(\Gamma, S)$, then SFTs (respectively nearest-neighbor SFTs) on (G,Γ) correspond exactly to SFTs (respectively nearest-neighbor SFTs) on the group G with respect to S.

- If Γ acts quasi-transitively on a graph G, then for each $k \geqslant 1$, the set X_{col} (respectively X'_{col}) of proper k-colorings (respectively proper k-edge-colorings) of G is a nearest-neighbor SFT on (G,Γ).

Higher block shift. We adapt here a construction known in symbolic dynamics as the *higher block shift* (see for example [CP15] or [ABJ18, Proposition 9.3.21]). We let $A, (G,\Gamma)$ and S be as introduced in the previous paragraph and let $X \subseteq A^{V(G)}$ be an SFT of (G,Γ)

with respect to S. Let $\mathcal{F}$ be a finite set of forbidden patterns such that $X = X_{\mathcal{F}}$. We now explain how to construct a subshift Y on Γ, in the group sense of Section 14.1.

We fix a base vertex $v_0 \in V(G)$ and let $N \geqslant 0$ be such that the supports of all the patterns from $\mathcal{F}$ are included in the ball $B_N(v_0)$ of radius N around v_0 in G. Up to choosing a larger value for N, we also assume that $B_N(v_0)$ contains at least one vertex from each Γ-orbit of $V(G)$. Note that as G is locally finite, $B_N(v_0)$ is finite. We consider the alphabet

$$C := \{\widetilde{p} \in A^{B_N(v_0)} : \widetilde{p}|_{\mathrm{supp}(p)} \neq p, \forall p \in \mathcal{F}\},$$

and let $\mathcal{F}_C$ be the set of patterns p of C^Γ with support $\{1_\Gamma, g\}$ such that there exists $u \in B_N(v_0) \cap B_N(g \cdot v_0)$ for which $(p(1_\Gamma))(u) \neq (p(g))(g^{-1} \cdot u)$ (note that if $u \in B_N(g \cdot v_0)$, then $g^{-1} \cdot u \in B_N(v_0)$). In other words, the patterns from $\mathcal{F}_C$ correspond to the pairs of local colorings of the N-balls around v_0 and $g \cdot v_0$ that do not agree with each other. We let $Y := X_{\mathcal{F}_C} \subseteq C^\Gamma$. Note that in general, there is no reason for $\mathcal{F}_C$ to be finite, so Y is not necessarily a subshift of finite type. The following remark will allow us to find sufficient conditions for $\mathcal{F}_C$ to be finite.

Remark 16.2. Assume that Γ acts quasi-regularly on G. Note that for any two vertices v and v' belonging to the same Γ-orbit, the stabilizers Γ_v and $\Gamma_{v'}$ are conjugate. Thus, as Γ acts quasi-transitively on G, there is a uniform bound on the size of the stabilizers Γ_v. In particular, in this situation the set of automorphisms $g \in \Gamma$ such that $B_N(v_0) \cap B_N(g \cdot v_0) \neq \emptyset$ is finite, so $\mathcal{F}_C$ is finite and Y is an SFT.

Note that Γ might not be finitely generated in general. However, this is not a real issue as the definitions of subshift and SFT we gave in Section 14 still make sense for groups that are not finitely generated.

We now define a mapping $\sigma : X \to C^\Gamma$ as follows. For each $x \in X$, we define $y := \sigma(x) : \Gamma \to C$ by setting $y(g) := (g^{-1} \cdot x)|_{B_N(v_0)}$ for each $g \in \Gamma$, i.e., for all $v \in B_N(v_0)$, $(y(g))(v) = x(g \cdot v)$. In other words, $y(g)$ reproduces the coloring induced by x on the ball of radius N around $g \cdot v_0$ in G. Note that for every $g \in \Gamma$, as $x \in X = X_{\mathcal{F}}$, we indeed have $y(g) \in C$. The next lemma establishes a connection between the properties of the SFT X on (G, Γ) and the properties of the the subshift Y on Γ. It essentially states that X and Y enjoy many common properties.

Lemma 16.3. *σ is an injective Γ-invariant mapping such that $\sigma(X) = Y$.*

Proof. In the remainder of this proof, we let $\mathcal{D} = \{v_1, \dots, v_k\}$ denote a fixed set of representatives of the Γ-orbits of $V(G)$ distinct from $\Gamma \cdot v_0$. By the choice of N, we may assume without loss of generality that $\mathcal{D} \subseteq B_N(v_0)$.

We first check that σ defines a Γ-invariant mapping, with respect to the respective actions of Γ on X and Y. Let $x \in X, h, g \in \Gamma$ and $v \in B_N(v_0)$. We must show that $\sigma(h \cdot x)(g)(v) = (h \cdot \sigma(x))(g)(v)$. This equality holds as by definition of σ we have on the one hand

$$\sigma(h \cdot x)(g)(v) = (h \cdot x)(g \cdot v) = x((h^{-1}g) \cdot v),$$

and on the other hand

$$(h \cdot \sigma(x))(g)(v) = (\sigma(x))(h^{-1}g)(v) = x((h^{-1}g) \cdot v).$$

We now check that $\sigma(x) \in Y = X_{\mathcal{F}_C}$ for each $x \in X$. We fix $x \in X$ and let $y := \sigma(x)$. As σ is Γ-invariant, and as every pattern from $\mathcal{F}_C$ has a support of the form $\{1_\Gamma, g\}$ for some $g \in \Gamma$, it is enough to prove that $y|_{\{1_\Gamma,g\}} \notin \mathcal{F}_C$ for all $g \in \Gamma$. If $g \in \Gamma$ is such that $B_N(v_0) \cap B_N(g \cdot v_0) = \emptyset$, then we clearly have $y|_{\{1_\Gamma,g\}} \notin \mathcal{F}_C$. Let $g \in \Gamma$ be such that there exists some $u \in B_N(v_0) \cap B_N(g \cdot v_0)$. Then we have $(y(1_\Gamma))(u) = x(u) = x(gg^{-1} \cdot u) = (y(g))(g^{-1} \cdot u)$. Thus by definition of $\mathcal{F}_C$, we indeed have $y|_{\{1_\Gamma,g\}} \notin \mathcal{F}_C$, as desired.

We show that σ is injective. Let $x \neq x' \in X$ and $v \in V(G)$ be such that $x(v) \neq x'(v)$. We let $y := \sigma(x)$ and $y' := \sigma(x')$. There exists some $i \in \{0, \dots, k\}$ such that v is in the Γ-orbit of v_i, and by the choice of N recall that $v_i \in B_N(v_0)$. In particular it implies that there exists some $g \in \Gamma$ such that $v \in g \cdot B_N(v_0) = B_N(g \cdot v_0)$, thus $g^{-1} \cdot v \in B_N(v_0)$. We claim that $y(g) \neq y'(g)$. This is because $y(g)(g^{-1} \cdot v) = x(v)$ while $y'(g)(g^{-1} \cdot v) = x'(v)$. In particular, we have $y \neq y'$, implying the injectivity of σ.

We now show that σ induces a surjective mapping onto Y, and we let $y \in Y$. We start by proving the following useful claim.

Claim 16.4. For every $(h, i) \in \Gamma \times \{0, \dots, k\}$ such that $h \cdot v_i \in B_N(v_0)$ and every $g \in \Gamma$, we have

$$y(g)(h \cdot v_i) = y(gh)(v_i).$$

Proof of the Claim: We set $y' := g^{-1} \cdot y$. As $y \in Y$, and as Y is a subshift, we also have $y' \in Y$. Thus, as $Y = X_{\mathcal{F}_C}$, for every h-pattern $p \in \mathcal{F}_C$, we must have $y'|_{\{1_\Gamma,h\}} \neq p$. In particular, as $h \cdot v_i \in B_N(v_0) \cap B_N(h \cdot v_0)$, by definition of $\mathcal{F}_C$, it imposes that:

$$(y'(1_\Gamma))(h \cdot v_i) = (y'(h))(h^{-1}h \cdot v_i).$$

As $y' = g^{-1} \cdot y$, we have on the one hand $(y'(1_\Gamma))(h \cdot v_i) = y(g)(h \cdot v_i)$, and on the other hand $(y'(h))(h^{-1}h \cdot v_i) = y(gh)(v_i)$. This concludes the proof of the claim. $\Diamond$

We now construct a configuration $x \in A^{V(G)}$ from y as follows. For every $v \in V(G)$, by the choice of $\mathcal{D}$ there exists at least one pair $(g, i) \in \Gamma \times \{0, \dots, k\}$ such that $v = g \cdot v_i$. We set $x(v) := y(g)(v_i)$. The following claim implies that the definition of $x(v)$ is independent of the choice of g.

Claim 16.5. Let $g, g' \in \Gamma$ and $i \in \{0, \dots, k\}$ be such that $g \cdot v_i = g' \cdot v_i$. Then

$$y(g)(v_i) = y(g')(v_i).$$

Proof of the Claim: We set $h := g^{-1}g'$. Then $h \cdot v_i = v_i$, and by the choice of N, we have $v_i \in B_N(v_0)$. Thus we can apply Claim 16.4 to the pair (h, i), which gives:

$$y(g)(v_i) = y(g)(h \cdot v_i) = y(gh)(v_i) = y(g')(v_i).$$

This concludes the proof of the claim. $\Diamond$

To prove that σ is surjective onto Y, it remains to prove that $x \in X$ and that $\sigma(x) = y$. It will follow from the next claim.

Claim 16.6. For every $g \in \Gamma$, we have

$$(g^{-1} \cdot x)|_{B_N(v_0)} = y(g).$$

Proof of the Claim: We let $v \in B_N(v_0)$. By the choice of $\mathcal{D}$, there exists $(h, i) \in \Gamma \times \{0, \dots, k\}$ such that $v = h \cdot v_i$. By definition of x and by Claim 16.5, we have $(g^{-1} \cdot x)(v) = x(gh \cdot v_i) = y(gh)(v_i)$. Moreover, as $v = h \cdot v_i \in B_N(v_0)$, Claim 16.4 applies and gives $y(gh)(v_i) = y(g)(h \cdot v_i) = y(g)(v)$. Putting all equalities together, we obtain $(g^{-1} \cdot x)(v) = y(g)(v)$ for every $v \in B_N(v_0)$, so $(g^{-1} \cdot x)|_{B_N(v_0)} = y(g)$. ♢

We now show that $x \in X$, i.e., that for each $g \in \Gamma$ and every $p \in \mathcal{F}$, $(g^{-1} \cdot x)|_{\mathrm{supp}(p)} \neq p$. Recall that by the choice of N, for every $p \in \mathcal{F}$, the support of p is included in $B_N(v_0)$. Moreover, by definition of the alphabet C, for every $p \in \mathcal{F}_C$, $y(g)|_{\mathrm{supp}(p)} \neq p$. Hence, by Claim 16.6, we have $(g^{-1} \cdot x)|_{\mathrm{supp}(p)} \neq p$ for every $p \in \mathcal{F}_C$, implying that $x \in X$, as desired. In particular, Claim 16.6 together with the definition of σ immediately imply that $\sigma(x) = y$. Eventually, we conclude that $\sigma(X) = Y$. □

By Remark 16.2, if we moreover assume that the action of Γ on G is quasi-regular, then Y is also an SFT on Γ. Recall that by Lemma 12.2, in this special case, Γ is finitely generated and its Cayley graphs are quasi-isometric to G. Many of the properties mentioned in Sections 14 and 15 behave well with respect to quasi-isometries. In particular, Cohen [Coh17] proved that if Γ and Γ' are two finitely generated groups that are quasi-isometric, then:

- The domino problem is undecidable on Γ if and only if it is undecidable on Γ';
- If Γ and Γ' are torsion-free, then Γ has a strongly aperiodic SFT if and only if Γ' has a strongly aperiodic SFT;
- Γ has a weakly aperiodic SFT if and only if Γ' has a weakly aperiodic SFT.

From this perspective, Lemma 16.3 suggests that the definition of SFT we introduced in this section is relevant if one is interested in working in generalizations of questions in the spirit of Conjectures 14.5 and 15.7 for quasi-transitive graphs.

SFT of 2-ended quasi-transitive graphs. We end this section with a simple application of Lemma 16.3 on 2-ended quasi-transitive graphs.

Proposition 16.7. *Let G be a connected locally finite 2-ended graph and Γ be a group acting quasi-transitively and quasi-regularly on G. Then G is strongly periodic, i.e., for any SFT X on (G, Γ), there exists a configuration $x \in X$ such that $\Gamma \cdot x$ is finite.*

Proof. By Lemma 12.2, Γ must be finitely generated, and for every finite set of generators S, $\mathrm{Cay}(\Gamma, S)$ is quasi-isometric to G. In particular, Γ has exactly 2 ends. We let $N \geqslant 1, Y$ and $\sigma : X \to Y$ be as in Lemma 16.3. By Remark 16.2, Y is an SFT on Γ, so by Proposition 15.6, Y is not weakly aperiodic and admits a configuration $y \in Y$ such that $\Gamma \cdot y$ is finite. By Lemma 16.3, $\sigma : X \to Y$ is a bijection, so we can consider $x := \sigma^{-1}(y) \in X$. As σ is Γ-invariant, we must have $\Gamma \cdot x = \sigma^{-1}(\Gamma \cdot y)$, implying that $\Gamma \cdot x$ is also finite, as desired. □

We claim that the arguments we used in the proof of Lemma 4.1 still apply to show that for every connected 2-ended locally finite graph G and every group Γ with a quasi-transitive action on G, there must exist a subgroup Γ' of Γ acting quasi-transitively and quasi-regularly on G. We thus immediately obtain alternate proofs of Theorem 10.3 and Corollary 10.4 when applying Proposition 16.7 on the SFTs X_{col} and X'_{col} defined in Example 16.1.

17 Optimization in graphical small cancellation theory

So far in this manuscript, we only worked on quasi-transitive graphs or groups satisfying some "nice properties", and were able to derive structural properties after exploiting the fact that such graphs/groups admit decompositions into pieces with a specific structure. We end Chapter 2 with a specific construction of some monster groups, i.e., of groups that do admit a sophisticated structure, and whose existence often gives rise to counterexamples of deep algebraic or geometric conjectures.

Gromov [Gro03] constructed finitely generated groups whose Cayley graphs contain all graphs from a given infinite sequence of expander graphs of unbounded girth and bounded diameter-to-girth ratio. These so-called *Gromov monster groups* provide examples of finitely generated groups that do not coarsely embed into any Hilbert space, among other interesting properties. If the finite graphs from the sequence used in Gromov's construction admit graphical small cancellation labellings, then one gets similar examples of Cayley graphs containing all the graphs of the family as isometric subgraphs. Osajda [Osa20] recently showed how to obtain such labellings using the probabilistic method. In this section, we will show how to simplify Osajda's approach, decreasing the number of generators of the resulting group significantly. Results and proofs from this section come from the article [EG24b], a joint work with Louis Esperet.

17.1 Introduction

Consider a sequence $\mathcal{G} = (G_n)_{n\geqslant 1}$ of finite bounded degree graphs, whose girth (length of a shortest non-trivial cycle) tends to infinity. We say that the sequence is *dg-bounded* if the ratio between the diameter and the girth of each G_n is bounded by a (uniform) constant, see [AT18]. Consider such a sequence $\mathcal{G}$. Gromov [Gro03] proved that there is a finitely generated group Γ with a finite set of generators S such that the Cayley graph $\mathrm{Cay}(\Gamma, S)$ contains (in a certain metric sense) all the members of $\mathcal{G}$. By choosing $\mathcal{G}$ as a family of suitable expander graphs, this implies that such a group Γ has a number of pathological properties, in particular related to coarse embeddings into Hilbert spaces, or to Guoliang Yu's property A. The construction has also been used very recently to disprove a conjecture on the twin-width of groups and hereditary graph classes [BGTT22]. Gromov [Gro03] introduced the graphical small cancellation condition on the labellings. By the classical small cancellation theory, the existence of labellings of $\mathcal{G} = (G_n)_{n\geqslant 1}$ with the graphical small cancellation condition guarantees that in Gromov's construction each graph G_n embeds *isometrically* in the Cayley graph $\mathrm{Cay}(\Gamma, S)$, which means that the embedding of each G_n in $\mathrm{Cay}(\Gamma, S)$ is distance-preserving and thus in particular the graphs G_n appear as induced subgraphs in $\mathrm{Cay}(\Gamma, S)$. Osajda [Osa20] recently showed, using the probabilistic method, that labellings satisfying Gromov's graphical small cancellation condition do exist, under mild assumptions on $\mathcal{G} = (G_n)_{n\geqslant 1}$.

Given a sequence $\mathcal{G} = (G_n)_{n\geqslant 1}$ of graphs whose edges are labelled with elements from some set S, a *word* in $\mathcal{G}$ is a sequence of labels that can be read along a path of some graph of $\mathcal{G}$. The main idea of graphical small cancellation theory is to assign labels from a finite set S to the edges of all the graphs from the sequence $\mathcal{G} = (G_n)_{n\geqslant 1}$, such that words in each G_n that are sufficiently long compared to the girth of G_n occur only once in all the sequence $\mathcal{G}$ (this will be made more precise in the next section). The labels from S are then used as generators to define the group Γ whose relators are the words labelling the cycles of each G_n. The number of labels (the size of the set S) then gives an upper bound on the minimum number of generators of the group, and thus on the degree of the associated Cayley graph (up to a multiplicative factor of two, if we do not require that S is closed under taking inverses). A natural problem is to minimize this number of generators.

Our purpose in this section is twofold: we present a simplified version of the proof of existence of the labelling of Osajda [Osa20], and significantly decrease the number of generators (and thus the degree of the corresponding Cayley graph). Osajda's proof is based on an application of the Lovász Local Lemma. Instead, we use a self-contained counting argument popularized by Rosenfeld [Ros20], and originally introduced in the field of combinatorics on words in the context of pattern avoidance. This allows us to cleanly handle all the different forbidden patterns at once, instead of sequentially, and greatly reduces the number of labels. We combine this with a significantly simpler (and stronger) analysis of intersecting patterns in order to a obtain a shorter argument that also produces much better bounds.

For the sake of concreteness, if we take $\mathcal{G} = (G_n)_{n\geqslant 1}$ to be the sequence of cubic Ramanujan graphs introduced by Chiu [Chi92], which is likely to offer the best known parameters in terms of degree and diameter-to-girth ratio, our result leads to the existence of a group with 96 generators, whose Cayley graph (of maximum degree 96) contains all the graphs from $\mathcal{G}$ as isometric subgraphs. For the same family, the construction of Osajda [Osa20] uses about 10^{272} generators (although we note that some of the quick optimization steps we perform in Section 17.4 can also be carried directly in Osajda's proof, improving his bound to about 10^{70} generators).

17.2 The graphical small cancellation condition

All the graph we will consider in this section will always be implicitly assumed to be connected. For every graph G we will consider in Section 17, we will always consider that we are given an arbitrary fixed orientation A of the edges of G, and denote with $\vec{G} := (V(G), A)$ the associated directed graph. All the results and properties we will present do not depend on the specific orientation, but the orientation is nevertheless crucial to define the relevant objects that we consider belows. Consider a set $S \uplus \overline{S}$ which is closed under formal inverse. Consider also a labelling $\ell : E(G) \to S \uplus \overline{S}$ of the edges of G by the elements of S. We extend the labelling ℓ to the ordered pairs of adjacent vertices (x, y) in G as follows: if (x, y) is an arc of $\vec{G}$ then $\ell(x, y) := \ell(xy)$ and otherwise $\ell(x, y) := \overline{\ell(xy)}$. The orientation $\vec{G}$ is only used to define this extended labelling ℓ of the the ordered pairs of adjacent vertices, and will not be mentioned elsewhere. We say that the labelling ℓ is *reduced* if for any vertex $v \in V(G)$, and for any pair of distinct neighbors u, w of v in G, $\ell(v, u) \neq \ell(v, w)$. An ℓ-*word* (or simply a *word*, if ℓ is clear from the context) in G is obtained from a path P in G as follows: if

$P = v_1, v_2, \ldots, v_k$, then $\ell(P) := \ell(v_1, v_2) \cdots \ell(v_{k-1}, v_k) \in L^*$ is the ℓ-word associated to P. The *length* of a path is its number of edges. We remark that in this section, we consider paths as either a sequence of vertices, or a sequence of edges, depending on the context, and in particular any path $P = v_1, v_2, \ldots, v_k$ is distinct from its reverse path $\overleftarrow{P} := v_k, v_{k-1}, \ldots, v_1$.

Let $\mathcal{G} = (G_n)_{n\geqslant 1}$ be a sequence of finite graphs. Let λ be a positive real number (for the main application in group theory we need $\lambda \in \left(0, \frac{1}{6}\right]$, but this will not be needed in the full generality of the results presented in this section and the next). Following the terminology of [Osa20], a sequence of labellings $(\ell_n)_{n\geqslant 1}$ of the graphs from $\mathcal{G}$, with labels from some set S as above, is said to satisfy the $C'(\lambda)$-*small cancellation property* if for all $n \geqslant 1$, ℓ_n is a reduced labelling of G_n and no word of length at least $\lambda \cdot \text{girth}(G_n)$ in G_n appears on a different path in $\mathcal{G}$. Small cancellation properties were initially introduced for group presentations, as a convenient tool to construct *word-hyperbolic groups*, see for instance Chapter V in [LS01]. The property $C'(\lambda)$ we use here is defined in the more general context of graphs, and is usually known as *graphical cancellation property* in the literature. In the remainder of Section 17, we will omit the "graphical" term, as there is no risk of confusion with the original small cancellation properties.

Osajda [Osa20] recently proved that under mild assumptions, any sequence of bounded degree dg-bounded graphs of unbounded girth admits small cancellation labellings with a finite number of labels.

Theorem 17.1 ([Osa20]). *Let $\lambda \in \left(0, \frac{1}{6}\right]$ and $A > 0$ be real numbers, and let $\Delta \geqslant 3$ be an integer. Let $\mathcal{G} = (G_n)_{n\geqslant 1}$ be a sequence of finite graphs of maximum degree Δ such that* $\text{girth}(G_n) \to \infty$ *as $n \to \infty$, and* $\text{diam}(G_n) \leqslant A \cdot \text{girth}(G_n)$ *for any $n \geqslant 1$. Assume moreover that $1 < \lfloor \lambda \cdot \text{girth}(G_n) \rfloor < \lfloor \lambda \cdot \text{girth}(G_{n+1}) \rfloor$ for every $n \geqslant 1$. Let*

$$L \geqslant 2e^4 \Delta^{2A/\lambda+2} \cdot (4e^4\Delta)^{8A/\lambda+16}$$

be any even integer. Then $\mathcal{G}$ has a sequence of labellings satisfying the $C'(\lambda)$-small cancellation property, with labels from a set $S \uplus \overline{S}$ of size L.

We note that the bound on A in the statement of Theorem 17.1 is not explicit in [Osa20], but follows from the fact that every finite graph G with a cycle satisfies $\text{girth}(G) \leqslant 2\text{diam}(G) + 1$, and thus if $\text{diam}(G) \leqslant A \cdot \text{girth}(G)$ we must have $A \geqslant \frac{1}{2} - \frac{1}{2\text{girth}(G)} \geqslant \frac{1}{3}$.

The bound on L in Theorem 17.1 has two components: $2e^4\Delta^{2A/\lambda+2}$ comes from a first phase, where Osajda shows how to assign labels in each $G_n \in \mathcal{G}$, so that no word of G_n appears as a word of length at least $\lambda \cdot \text{girth}(G_i)$ in some G_i, with $i < n$. The second component, $(4e^4\Delta)^{8A/\lambda+16}$, comes from a second phase where Osajda shows how to assign labels in each $G_n \in \mathcal{G}$, so that no word of G_n of length at least $\lambda \cdot \text{girth}(G_n)$ appears twice in G_n. This second phase is significantly more involved, which explains the much larger label size. Our main contribution is the following.

- we use a counting argument instead of the Lovász Local Lemma. This allows us to assign labels in a single phase (resulting in an additive combination of the number of labels, instead of a multiplicative one), and optimize the multiplicative constants. Moreover, the resulting proof is completely self-contained.

- we provide a major simplification in the analysis of Osajda's second phase, showing that the long words appearing twice in G_n can be avoided with a number of labels of size comparable to Osajda's first phase.

Using results of Gromov (see [Oll06, Gru15]), Theorem 17.1 leads to the following.

Corollary 17.2. *Let $\lambda, A, \Delta, \mathcal{G}$ be as in Theorem 17.1. Then for any even integer $L \geqslant 2e^4\Delta^{2A/\lambda+2} \cdot (4e^4\Delta)^{8A/\lambda+16}$, there is a group Γ with a set S of L generators such that the corresponding Cayley graph $\mathrm{Cay}(\Gamma, S)$ contains isometric copies of all the graphs from $\mathcal{G}$.*

As alluded to in the introduction, in applications we typically want $\mathcal{G}$ to be a sequence of *expander graphs*. We omit the precise definition here, as it will not be necessary in this section. We only mention that expansion can be defined in several essentially equivalent ways, using isoperimetric inequalities or spectral properties. Families of finite random regular graphs typically have these properties, but constructing explicit families of expander graphs has been an important problem in Mathematics, with major applications in Theoretical Computer Science. The interested reader is referred to the survey [HLW06] for more on expander graphs.

A useful family $\mathcal{G}$ for us is the sequence of cubic Ramanujan graphs introduced by Chiu [Chi92]. These graphs are expander graphs (as Ramanujan graphs, they have the best possible spectral expansion), are Δ-regular with $\Delta = 3$, satisfy $\mathrm{girth}(G_n) \to \infty$ as $n \to \infty$ (their girth is logarithmic in their number of vertices) and $\mathrm{diam}(G) \leqslant \frac{3}{2}\,\mathrm{girth}(G) + 5$ for any $G \in \mathcal{G}$. By discarding a bounded number of small graphs in the sequence, this implies that we have $\mathrm{diam}(G) \leqslant (\frac{3}{2} + \epsilon)\mathrm{girth}(G)$ for any $\epsilon > 0$ and any graph G in the sequence, and thus we can take $A \leqslant \frac{3}{2} + \epsilon$ for any $\epsilon > 0$.

17.3 Smaller cancellation labellings

Our main result is the following optimized version of Theorem 17.1.

Theorem 17.3. *Let $\lambda, A, \Delta, \mathcal{G}$ be as in Theorem 17.1, that is $\lambda \in \left(0, \frac{1}{6}\right]$ and $A > 0$ are real numbers, $\Delta \geqslant 3$ is an integer, and $\mathcal{G} = (G_n)_{n\geqslant 1}$ is a sequence of finite graphs of maximum degree Δ such that $\mathrm{girth}(G_n) \to \infty$ as $n \to \infty$, and $\mathrm{diam}(G_n) \leqslant A \cdot \mathrm{girth}(G_n)$ and $1 < \lfloor \lambda \cdot \mathrm{girth}(G_n) \rfloor < \lfloor \lambda \cdot \mathrm{girth}(G_{n+1}) \rfloor$ for every $n \geqslant 1$. Let*

$$L \geqslant 2(\Delta - 1) + 26(\Delta - 1)^{2A/\lambda+2}$$

be any even integer. Then $\mathcal{G}$ has a sequence of labellings satisfying the $C'(\lambda)$-small cancellation property, with labels from a set $S \uplus \overline{S}$ of size L.

We note that the multiplicative constant of 26 in the bound on L can be optimized both for small values of Δ and asymptotically as $\Delta \to \infty$. We have chosen not to do so here for simplicity, and we remark that improving the factor 2 in the exponent of $(\Delta - 1)$ is a more rewarding challenge (see the Section 17.4). When $\Delta \to \infty$, the number L of labels in Theorem 17.3 grows as $O(\Delta^{2A/\lambda+2})$, and we will see in Section 17.4 that this can be easily improved to $O(\Delta^{A/\lambda+2})$. This is to be compared with the bound $O(\Delta^{10A/\lambda+18})$ of Theorem

17.1. In Section 17.4, we will also see several ways to improve the constants significantly when $\Delta = 3$, and the girth of the first graph in the sequence is already quite large.

Similarly as above, we obtain the following corollary.

Corollary 17.4. *Let $\lambda, A, \Delta, \mathcal{G}, L$ be as in Theorem 17.3. Then there is a group Γ with a set S of L generators closed under taking inverse such that the corresponding Cayley graph* $\mathrm{Cay}(\Gamma, S)$ *contains isometric copies of all the graphs from $\mathcal{G}$.*

Using the family of cubic Ramanujan graphs of Chiu [Chi92] mentioned at the end of the previous section, we can apply Corollary 17.4 with $\Delta = 3$, $A = \frac{3}{2}$ and $\lambda = \frac{1}{6}$. Then we obtain a group with a set of $L = 4 + 26 \cdot 2^{20} = 27262980$ generators such that the corresponding Cayley graph contains isometric copies of graphs from an infinite family of expander graphs. We will see in Section 17.4 how to decrease this number of generators to 96.

If instead we apply Corollary 17.2 to the same family $\mathcal{G}$ (and hence with the same parameters $\Delta = 3$, $A = \frac{3}{2}$ and $\lambda = \frac{1}{6}$), the resulting Cayley graph has degree more than 10^{272}.

We now prove our main result.

Proof of Theorem 17.3. Let $\alpha := 2(\Delta - 1)^{2A/\lambda + 2}$, and let

$$L \geqslant 2(\Delta - 1) + 13\alpha = 2(\Delta - 1) + 26(\Delta - 1)^{2A/\lambda + 2}$$

be an even integer. Let $S \uplus \overline{S}$ be a set of L elements, closed under formal inverses (and such that each element $a \in S$ is different from its formal inverse $\bar{a}$). For any $n \geqslant 1$, let $\gamma_n := \lfloor \lambda \cdot \mathrm{girth}(G_n) \rfloor$. In particular $\gamma_n \leqslant \lambda \cdot \mathrm{girth}(G_n) \leqslant \gamma_n + 1$ for any $n \geqslant 1$, and thus

$$\frac{1}{\lambda} \leqslant \frac{\mathrm{girth}(G_n)}{\gamma_n} \leqslant \frac{1}{\lambda} + \frac{1}{\lambda \gamma_n} \leqslant \frac{2}{\lambda}. \tag{2.1}$$

We will sequentially assign labels from $S \uplus \overline{S}$ to the edges of each of the graphs $(G_n)_{n \geqslant 1}$. Assume that for each $i < n$, we have already defined a labelling ℓ_i of the edges of G_i such that the sequence of labellings $(\ell_i)_{i<n}$ satisfies the $C'(\lambda)$-small cancellation property. We now want to define a labelling ℓ_n of G_n so that the sequence $(\ell_i)_{i \leqslant n}$ of labellings of the graphs from $(G_i)_{i \leqslant n}$ still satisfies the $C'(\lambda)$-small cancellation property.

For the proof it will be convenient to consider *partial* labellings of G_n, which are labellings of some subset F of edges of G_n. Equivalently, these are labellings of the edges of $G_n[F]$, the subgraph of G_n induced by the edges of F. We recall that each labelling $\ell(xy)$ of an edge xy yields two labellings $\ell(x, y)$ and $\ell(y, x)$ of the pairs (x, y) and (y, x) by elements of $S \uplus \overline{S}$ that are formal inverse (and that whether $\ell(xy) = \ell(x, y)$ or $\ell(xy) = \ell(y, x)$ depends only on the orientation of the edge xy in some fixed but otherwise arbitrary orientation of the graph under consideration).

Let F be a non-empty subset of $E(G_n)$. We say that a labelling ℓ of $G_n[F]$ with labels from $S \uplus \overline{S}$ is *valid* if it satisfies the following properties:

(a) ℓ is a reduced labelling of $G_n[F]$,

(b) for each $1 \leqslant i < n$, no ℓ_i-word of length at least γ_i in G_i appears as an ℓ-word in $G_n[F]$, and

(c) no ℓ-word of length at least γ_n appears on two different paths of $G_n[F]$.

Let $c(F)$ be the number of valid labellings ℓ of $G_n[F]$ with labels from $S \uplus \overline{S}$ (when F is empty we conveniently define $c(F) := 1$). In the remainder of the proof we will show the following claim, which clearly implies that G_n has a labelling ℓ_n such that the sequence of labellings $(\ell_i)_{i \leqslant n}$ of $(G_i)_{i \leqslant n}$ still satisfies the $C'(\lambda)$-small cancellation property, and thus we can find such labellings in all the graphs from $\mathcal{G}$.

Claim 17.5. For any non-empty $F \subseteq E(G_n)$ and any $e \in F$, $c(F) \geqslant \alpha \cdot c(F \setminus \{e\})$.

We prove Claim 17.5 by induction on $|F|$. Recall that by assumption, $\gamma_i > 1$ for any $i \geqslant 1$, so the properties (a), (b), (c) above are trivially satisfied if F contains a single element e, which is assigned an arbitrary label from $S \uplus \overline{S}$. It follows that $c(\{e\}) = L \geqslant \alpha = \alpha \cdot c(\emptyset)$, as desired. So we can now assume that F contains at least two elements.

Assume that we have proved the claim for any $F' \subseteq E(G_n)$ with $|F'| < |F|$. Consider any edge $xy \in F$. Our goal in the remainder of the proof is to show that $c(F) \geqslant \alpha \cdot c(F \setminus \{xy\})$. Note that by the induction hypothesis, for any subset $F' \subseteq F$ containing xy,

$$c(F \setminus F') \leqslant \alpha^{1-|F'|} \cdot c(F \setminus \{xy\}). \tag{2.2}$$

Let $\mathcal{L}$ denote the set of labellings ℓ of F with labels from $S \uplus \overline{S}$ whose restriction to $F \setminus \{xy\}$ is valid, but such that ℓ itself is not. Then

$$c(F) = L \cdot c(F \setminus \{xy\}) - |\mathcal{L}|. \tag{2.3}$$

Consider first the subset $\mathcal{L}_a \subseteq \mathcal{L}$ of labellings of F that do not satisfy (a) above. Then by definition, for any $\ell \in \mathcal{L}_a$, x has a neighbor z different from y such that $\ell(x, y) = \ell(x, z)$, or y has a neighbor z different from x such that $\ell(y, x) = \ell(y, z)$. By assumption, the labelling ℓ^- of $F \setminus \{xy\}$ obtained from ℓ by discarding the label of xy is valid. Moreover, ℓ can be recovered in a unique way from ℓ^- and the edge xz or yz as above. As there are at most $2(\Delta - 1)$ choices for such an edge incident to xy, we obtain

$$|\mathcal{L}_a| \leqslant 2(\Delta - 1) \cdot c(F \setminus \{xy\}). \tag{2.4}$$

For $1 \leqslant i \leqslant n - 1$, let $\mathcal{L}_i$ be the subset of labellings $\ell \in \mathcal{L}$ of F such that $G_n[F]$ contains a path P containing xy such that $\ell(P)$ coincides with some ℓ_i-word $\ell_i(Q)$ of length γ_i in G_i. Let $\mathcal{L}_n$ be the subset of labellings $\ell \in \mathcal{L} \setminus \mathcal{L}_a$ of F such that $G_n[F]$ contains a path P containing xy such that $\ell(P)$ coincides with some ℓ-word $\ell(Q)$ of length γ_n in $G_n[F]$, for some path Q distinct from P.

For each $1 \leqslant i \leqslant n$ and each labelling $\ell \in \mathcal{L}_i$ as above, let ℓ^- denote the labelling of $F \setminus E(P)$ obtained from ℓ by discarding the labels of the edges of P. Then ℓ^- is a valid labelling of $F \setminus E(P)$. Moreover, if $1 \leqslant i \leqslant n - 1$ or if $i = n$ and P and Q are disjoint, then ℓ^- together with the paths P in G_n and Q in G_i (where each path is viewed as a sequence of edges) are sufficient to recover ℓ in a unique way.

Assume now that $\ell \in \mathcal{L}_n$ (so in particular ℓ is reduced), and the distinct paths P and Q of length γ_n in $G_n[F]$ such that $\ell(P) = \ell(Q)$ are not edge-disjoint. We first observe that $E(P) \cap E(Q)$ is a subpath of P and Q, since otherwise G_n would contain a cycle of length less than $2\gamma_n$, contradicting the assumption that $\mathrm{girth}(G_n) \geqslant \frac{\gamma_n}{\lambda} \geqslant 6\gamma_n$. Let $P = x_0, x_1, \ldots, x_{\gamma_n}$ and $Q = y_0, y_1, \ldots, y_{\gamma_n}$. Then $\ell(x_i, x_{i+1}) = \ell(y_i, y_{i+1})$ for any $0 \leqslant i \leqslant \gamma_n - 1$. Our goal is to show that despite the fact that the edges of $E(P) \cap E(Q)$ have been unlabelled in ℓ^-, we can still recover ℓ from ℓ^-, P and Q.

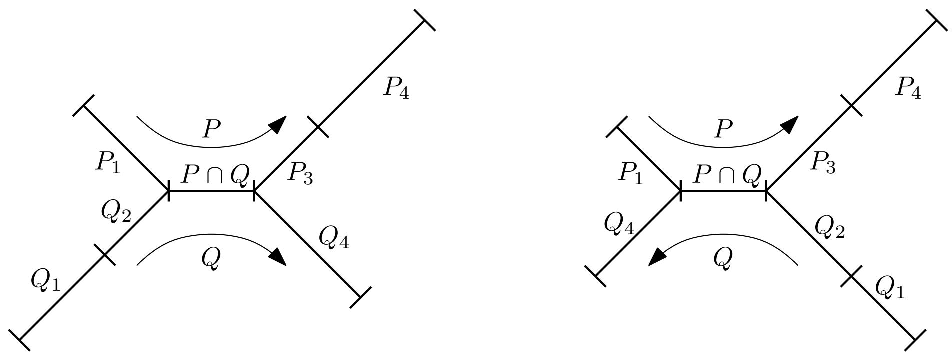

Figure 2.2: Two intersecting paths P and Q.

Assume first that P and Q intersect in the same direction, that is there are integers $0 \leqslant p, q \leqslant \gamma_n - 1$ and $1 \leqslant k \leqslant \gamma_n - 1$ such that $x_{p+i} = y_{q+i}$ for any $0 \leqslant i \leqslant k$. Note that $p \neq q$ since otherwise we would have $x_p = y_p$ and $x_{p+k} = y_{p+k}$ and the fact that $\ell(x_{p-1}, x_p) = \ell(y_{p-1}, y_p)$ or $\ell(x_{p+k}, x_{p+k+1}) = \ell(y_{p+k}, y_{p+k+1})$ would contradict the fact that ℓ is reduced. Up to considering the reverse paths $\overleftarrow{P}$ and $\overleftarrow{Q}$ instead of P and Q, we can assume without loss of generality that $q > p$. Divide P into consecutive subpaths P_1, $P \cap Q$, P_3, and P_4 and divide Q into consecutive subpaths Q_1, Q_2, $P \cap Q$, and Q_4, in such a way that $\ell(P_1) = \ell(Q_1)$ and $\ell(P_4) = \ell(Q_4)$ (see Figure 2.2, left). As P_1 and P_4 are edge-disjoint from $E(P) \cap E(Q)$, both $\ell(P_1)$ and $\ell(P_4)$ can be recovered from ℓ^-. Note that as we assumed that $q > p$, $|Q_2| > 0$, i.e. Q_2 has at least one edge. Let P' be the subpath of P obtained by concatenating $P \cap Q$ and P_3. It remains to explain how to recover $\ell(P')$ from ℓ^-. For this, it suffices to observe that since $\ell(P) = \ell(Q)$, the prefix of $\ell(P')$ of size $|Q_2|$ must be equal to $\ell(Q_2)$. Then the prefix of $\ell(P')$ of size $2|Q_2|$ must be equal to $\ell(Q_2) \cdot \ell(Q_2)$. By iterating this observation, it follows that $\ell(P')$ is a prefix of the word $\ell(Q_2)^\omega$ (the concatenation of an infinite number of copies of $\ell(Q_2)$). Since Q_2 is edge-disjoint from $E(P) \cap E(Q)$, $\ell(P')$ (and thus $\ell(P)$) can be recovered from ℓ^-, P and Q, as desired.

We now assume that P and Q intersect in reverse directions, that is there are integers $0 \leqslant p, q \leqslant \gamma_n$ and $k \geqslant 1$ such that $x_{p+i} = y_{q-i}$ for any $0 \leqslant i \leqslant k$. We say that P and Q *collide* if there is an index i such that either $x_i = y_i$, or $x_i = y_{i+1}$ and $y_i = x_{i+1}$ (think of two particles following the trajectories of P and Q at the same speed). Assume for the sake of contradiction that P and Q collide. If $x_i = y_i$ for some index i, then ℓ is not reduced, which is a contradiction. Otherwise we have $\ell(x_i, x_{i+1}) = \ell(y_i, y_{i+1}) = \ell(x_{i+1}, x_i)$, which contradicts the fact that $\ell(x_i, x_{i+1}) = \overline{\ell(x_{i+1}, x_i)}$ as for each $a \in S$, $\overline{a} \neq a$. So P and Q do not collide, and in particular $p \neq q$. We recall that $\overleftarrow{P}$ and $\overleftarrow{Q}$ denote the paths obtained by reversing

P and Q, respectively. When we use this notation below we also write $\overrightarrow{P}$ and $\overrightarrow{Q}$ instead of P and Q to avoid any confusion. Up to considering $\overleftarrow{P}$ and $\overleftarrow{Q}$ instead of $\overrightarrow{P}$ and $\overrightarrow{Q}$, we can again assume without loss of generality that $q > p$. We divide P into consecutive subpaths P_1, $\overrightarrow{P} \cap \overleftarrow{Q}$, P_3 and P_4 and we divide Q into consecutive subpaths Q_1, Q_2, $\overleftarrow{P} \cap \overrightarrow{Q}$, and Q_4, in such a way that $\ell(P_1) = \ell(Q_1)$, $\ell(P_4) = \ell(Q_4)$ (see Figure 2.2, right). As before, P_1 and P_4 are edge-disjoint from $E(P) \cap E(Q)$, so both $\ell(P_1)$ and $\ell(P_4)$ can be recovered from ℓ^-. As P and Q do not collide, $|Q_2| > |\overrightarrow{P} \cap \overleftarrow{Q}|$, which implies that $\ell(\overrightarrow{P} \cap \overleftarrow{Q})$ is equal to a prefix of $\ell(Q_2)$, and can thus be recovered from ℓ^-. Finally, since $\ell(P) = \ell(Q)$, $\ell(P_3)$ is equal to $\ell(\overleftarrow{P} \cap \overrightarrow{Q})$, which is obtained by reading $\ell(\overrightarrow{P} \cap \overleftarrow{Q})$ backwards. Hence, $\ell(P)$ can be recovered from ℓ^-, P and Q, as desired.

For each $1 \leqslant i \leqslant n$ and each edge e in G_i there are at most $(\Delta-1)^{\gamma_i - 1}$ paths of length γ_i containing e in which e is at a fixed position on the path. Hence, there are at at most $2\gamma_i(\Delta-1)^{\gamma_i-1}$ paths of length γ_i containing e (and in particular at most $2\gamma_i(\Delta-1)^{\gamma_i-1}$ choices for the path P in G_n containing xy when considering a labelling $\ell \in \mathcal{L}_i$). Moreover, each G_i has at most $1 + \Delta + \Delta(\Delta-1) + \cdots + \Delta(\Delta-1)^{\mathrm{diam}(G_i)-1}$ vertices, and thus at most

$$\tfrac{\Delta}{2} \cdot \left(1 + \Delta \tfrac{(\Delta-1)^{\mathrm{diam}(G_i)}-1}{\Delta-2}\right) \leqslant \tfrac{3}{2}(\Delta-1)^{\mathrm{diam}(G_i)+2} \tag{2.5}$$

edges, using $\Delta \geqslant 3$ (the inequality is quite loose here, we have chosen the right-hand side mostly in order to simplify the computation later). It follows that each G_i has at most

$$\tfrac{3}{2}(\Delta-1)^{\mathrm{diam}(G_i)+2} \cdot 2(\Delta-1)^{\gamma_i-1} \leqslant 3(\Delta-1)^{(2A/\lambda+1)\gamma_i+1} \tag{2.6}$$

paths of length γ_i (here the multiplicative factor γ_i disappears since we can count each path from its starting edge). It follows that there are at most $3(\Delta-1)^{(2A/\lambda+1)\gamma_i+1}$ choices for the path Q in G_i when considering a labelling $\ell \in \mathcal{L}_i$. Since $|E(P)| = \gamma_i$, it follows from (2.2) that for each labelling $\ell \in \mathcal{L}_i$, the number of valid labellings ℓ^- of $F \setminus E(P)$ is $c(F \setminus E(P)) \leqslant \alpha^{1-\gamma_i} \cdot c(F \setminus \{xy\})$. As each $\ell \in \mathcal{L}_i$ can be recovered from ℓ^-, P and Q in a unique way, we obtain

$$\begin{aligned} |\mathcal{L}_i| &\leqslant 2\gamma_i(\Delta-1)^{\gamma_i-1} \cdot 3(\Delta-1)^{(2A/\lambda+1)\gamma_i+1} \cdot \alpha^{1-\gamma_i} \cdot c(F \setminus \{xy\}) \\ &\leqslant 6\gamma_i(\Delta-1)^{(2A/\lambda+2)\gamma_i} \cdot \alpha^{1-\gamma_i} \cdot c(F \setminus \{xy\}) \\ &\leqslant 6\gamma_i(\alpha/2)^{\gamma_i} \cdot \alpha^{1-\gamma_i} \cdot c(F \setminus \{xy\}) \\ &\leqslant 6\alpha \cdot \gamma_i(1/2)^{\gamma_i} \cdot c(F \setminus \{xy\}), \end{aligned}$$

where we have used $\alpha = 2(\Delta-1)^{2A/\lambda+2}$ in the third inequality. As a consequence

$$\sum_{i=1}^{n} |\mathcal{L}_i| \leqslant 6\alpha \sum_{i=1}^{n} \gamma_i(1/2)^{\gamma_i} c(F \setminus \{xy\}) \leqslant 12\alpha \cdot c(F \setminus \{xy\}), \tag{2.7}$$

where we have used $\sum_{j=1}^{\infty} j(1/2)^j = 2$. As $\mathcal{L} = \mathcal{L}_a \cup \bigcup_{i=1}^{n} \mathcal{L}_i$, it follows from (2.4) and (2.7) that

$$\begin{aligned} |\mathcal{L}| &\leqslant c(F \setminus \{xy\}) \cdot (2(\Delta-1) + 12\alpha) \\ &\leqslant c(F \setminus \{xy\})(L - \alpha), \end{aligned}$$

by the definition of L. By (2.3), we have

$$\begin{aligned} c(F) &= L \cdot c(F \setminus \{xy\}) - |\mathcal{L}| \\ &\geqslant L \cdot c(F \setminus \{xy\}) - (L-\alpha)c(F \setminus \{xy\}) \\ &\geqslant \alpha \cdot c(F \setminus \{xy\}), \end{aligned}$$

as desired. This completes the proof of Claim 17.5, which concludes the proof of Theorem 17.3. □

17.4 Optimizing the number of generators

So far our goal was to optimize the construction of Osajda [Osa20], while obtaining a result that is comparable to his (i.e., a result with the exact same set of initial assumptions). There are two quick ways to further optimize the number of labels in Theorem 17.3, if we have some control over the family $\mathcal{G}$.

The first way consists in removing all sufficiently small graphs from $\mathcal{G}$ (we have done this already with the cubic Ramanujan graphs of Chiu [Chi92], to argue that A was arbitrarily close to $\frac{3}{2}$ in this case). As the girth of the graphs in $\mathcal{G}$ tends to infinity, the right-hand-side of (2.1) can be replaced by $\frac{1+\epsilon}{\lambda}$ for any $\epsilon > 0$. This allows to replace all instances of $2A/\lambda$ by $(1+\epsilon)A/\lambda$ in the proof, effectively dividing by 2 the exponent of the number of labels in the theorem. Using this observation in the case of the cubic Ramanujan graphs of Chiu [Chi92], with $\lambda = 1/6$, we obtain $\alpha = 2 \cdot 2^{(1+\epsilon)\frac{3}{2}/\frac{1}{6}+2} \leqslant 4097$ for sufficiently small $\epsilon > 0$, and a number of labels $L \geqslant 2 \cdot 2 + 13 \cdot 4097 \approx 53266$ is sufficient.

A more efficient way to decrease the number of labels in the case of families of expander graphs with an explicit description consists in using a more precise bound on the number of edges in a graph $G_n \in \mathcal{G}$, as a function of girth(G_n). In (2.5), we have used that $|E(G_n)| \leqslant \frac{3}{2}(\Delta-1)^{\text{diam}(G_n)+2} \leqslant \frac{3}{2}(\Delta-1)^{A\text{girth}(G_n)+2}$. However, better bounds are known for a number of families $\mathcal{G}$. This is the case for the cubic Ramanujan graphs of Chiu [Chi92] mentioned in the previous section. The graphs G in this class satisfy $|E(G_n)| \leqslant \frac{3}{2} \cdot 2^{(3\text{girth}(G_n)+6)/4}$, which is an improvement over the bound based on the diameter (recall that for these graphs $\Delta = 3$ and A can be made arbitrarily close to $\frac{3}{2}$). Fix any real $\epsilon > 0$, and recall that $\gamma_n = \lfloor \lambda \cdot \text{girth}(G_n) \rfloor$. Using as in the previous paragraph the fact that the girth of the graphs from $\mathcal{G}$ can be made arbitrarily large by discarding a constant number of graphs from the family, we can assume that $\gamma_n\epsilon > \gamma_1\epsilon$ is larger than any fixed constant, and thus $(3\text{girth}(G_n)+6)/4 \leqslant \frac{3+\epsilon}{4\lambda}\gamma_n$ and $|E(G_n)| \leqslant \frac{3}{2} \cdot 2^{\gamma_n \cdot (3+\epsilon)/4\lambda}$, for any $n \geqslant 1$. With $\lambda = 1/6$, we obtain $|E(G_n)| \leqslant \frac{3}{2} \cdot 2^{(9+\epsilon)\gamma_n/2}$, for any $n \geqslant 1$. Substituting this bound in (2.6), we obtain that there are at most $3 \cdot 2^{(11+\epsilon)\gamma_n/2-1}$, paths of length γ_n in G_n. Substituting this bound in the proof of Theorem 17.3, and defining $\alpha := (1+\epsilon)2^{(13+\epsilon)/2}$, we obtain the following.

$$\begin{aligned} |\mathcal{L}_i| &\leqslant 2\gamma_i 2^{\gamma_i - 1} \cdot 3 \cdot 2^{(11+\epsilon)\gamma_i/2-1} \cdot \alpha^{1-\gamma_i} \cdot c(F \setminus \{xy\}) \\ &\leqslant \tfrac{3\alpha}{2}\gamma_i \cdot 2^{(13+\epsilon)\gamma_i/2} \cdot \alpha^{-\gamma_i} \cdot c(F \setminus \{xy\}) \\ &\leqslant \tfrac{3\alpha}{2}\gamma_i \cdot (\tfrac{1}{1+\epsilon})^{\gamma_i} \cdot c(F \setminus \{xy\}), \end{aligned}$$

As $\sum_{j=1}^{\infty} j \cdot (\frac{1}{1+\epsilon})^j$ converges, we can choose again γ_1 sufficiently large so that the truncated sum $\sum_{j=\gamma_1}^{\infty} j \cdot (\frac{1}{1+\epsilon})^j$ is arbitrarily small (say smaller than $\epsilon/(\frac{3\alpha}{2})$). We obtain $\sum_{i=1}^{n} |\mathcal{L}_i| \leqslant \epsilon \cdot c(F \setminus \{xy\})$, and the same computation as in the proof of Claim 17.5 shows that any even number $L \geqslant 2 \cdot 2 + \epsilon + \alpha = \alpha + \epsilon + 4$ of labels is sufficient. Using $\alpha = (1+\epsilon)2^{(13+\epsilon)/2}$, and taking $\epsilon > 0$ sufficiently small, we can obtain that $L = 96$ labels are sufficient.

So, we obtain a group with a set $S \uplus \overline{S}$ of 96 generators whose Cayley graph $\mathrm{Cay}(\Gamma, S)$ contains infinitely many graphs of the sequence of cubic Ramanujan graphs as isometric subgraphs.

17.5 Concluding remarks and questions

The number L of labels in Theorem 17.3 is of order $O(\Delta^{2A/\lambda+2})$, as $\Delta \to \infty$, and the remarks in the previous section improve this bound to $O(\Delta^{A/\lambda+2})$. In typical applications, A is a small constant and the bound becomes $\Delta^{O(1/\lambda)}$. We now observe that this is the right order of magnitude. If G is a Δ-regular graph of girth g, then the ball of radius $g/2$ centered in any vertex induces a tree, and thus for any $\lambda < 1/2$, G contains $\Omega(\Delta^{g/2+\lambda g-1})$ paths of length λg (the ball of radius $g/2$ centered in a vertex contains $\Omega(\Delta^{g/2})$ edges and each of them is the starting point of $\Omega(\Delta^{\lambda g-1})$ paths of length λg). By the $C'(\lambda)$-small cancellation property, all these paths must correspond to different words. As there are at most $L^{\lambda g}$ possible words of length λg, we obtain $L^{\lambda g} = \Omega(\Delta^{g/2+\lambda g-1})$, and thus $L = \Omega(\Delta^{1/2\lambda+1-1/g})$. As the girth of the graphs in our family is unbounded, it follows that $L = \Omega(\Delta^{1/2\lambda+1})$, which shows that the bound in Theorem 17.3 is fairly close to the optimum (up to a small multiplicative factor in the exponent). It remains an interesting problem to close the gap between the upper and lower bounds, both in the case of small degree ($\Delta = 3$) and asymptotically as $\Delta \to \infty$.

It might also be interesting to consider other cancellation properties. For an integer $k \geqslant 1$, a family of labellings $(\ell_n)_{n\geqslant 1}$ of a graph family $\mathcal{G} = (G_n)_{n\geqslant 1}$ satisfies the $C(k+1)$-*small cancellation property* if for any $n \geqslant 1$, ℓ_n is reduced and no cycle C in G_n can be divided into k paths $P_1, \dots, P_k$ such that for each $1 \leqslant i \leqslant k$, the ℓ_n-word associated to P_i appears on a different path in $\mathcal{G}$. This condition is weaker than the $C'(1/k)$-small cancellation property, but nevertheless allows to construct finitely generated groups with interesting properties when $k \geqslant 7$ [Gru15]. A natural problem is to obtain a version of Theorem 17.3 for $C(k)$-small cancellation, with an improved exponent.

We conclude with some algorithmic remarks. Using the constructive proof of the Lovász Local Lemma by Moser and Tardos [MT10], the original proof of existence of the labelling given by Osajda [Osa20] can be turned into an efficient algorithm computing the labels, by which we mean a randomized algorithm, running in polynomial time (in the size of G_n), and computing a $C'(\lambda)$-small cancellation labelling for the sequence of graphs $(G_i)_{1\leqslant i\leqslant n}$. As our main goal was to obtain a simple, self-contained proof of the existence of the labels, we chose to use counting rather than constructive techniques such as the entropy compression method (see [GMP20]). It turns out that our result can also be obtained with this type of techniques, at the cost of a longer and more technical analysis.

Conclusion

The results we presented in this thesis build on a long line of work that started with Maschke's planarity theorem and that aims at providing a better understanding of the geometric and structural properties of Cayley graphs, and more generally of locally finite quasi-transitive graphs. We presented in Chapter 1 general decomposition theorems for planar and minor-excluded locally finite quasi-transitive graphs. We gave graph-theoretic applications, and discussed a number of questions left open after this work.

- As suggested in the recent works of MacManus [Mac23] and of Georgakopoulos and Papasoglou [GP23], the next step of this work would be to obtain similar decomposition theorems for locally finite quasi-transitive graphs satisfying more general properties than excluding a minor or being quasi-isometric to a planar graph. Geometric properties (i.e., properties that are preserved under taking quasi-isometries) are particularly relevant, and in this context, the following classes of locally finite quasi-transitive graphs are natural to study: 1-planar graphs, graphs excluding a fixed graph as an asymptotic minor and graphs that are quasi-isometric to a graph excluding a minor.

- Canonical tree-decompositions, as well as tangles, appeared to be relevant tools to study the structure of minor-excluded locally finite quasi-transitive graphs, and derive properties like accessibility. We also mention a recent result from [DJKK22] going in this direction, that states that a locally finite quasi-transitive graph is accessible if and only if there exists some integer $k \geqslant 1$ such that every two tangles can be distinguished by a separation of order at most k. It suggests that it could be interesting to obtain canonical versions of other structural results in graph theory. In particular, does there exist a canonical notion of twin-width or of twin-decomposition?

- Even though our initial questions about the existence of symmetric proper colorings turned out to have a negative answer in general, we do not have yet a clear idea of which graphs should have such colorings. In particular, do locally finite planar graphs always admit a symmetric proper coloring? If yes, can we find one that uses at most 4 colors?

  We introduced in Chapter 2 some central concepts and problems from symbolic dynamics and presented other applications of the results from Chapter 1. In particular, this suggests that graph theoretic tools could be relevant to tackle Conjectures 14.5 and 15.7. Conversely, we tried to give some evidence that methods and results used in symbolic dynamics could be useful to have some insight about questions from graph theory.

Bibliography

[ABJ18] Nathalie Aubrun, Sebastián Barbieri, and Emmanuel Jeandel. About the Domino Problem for Subshifts on Groups. In V. Berthé and M. Rigo, editors, *Sequences, Groups, and Number Theory*, Trends in Mathematics, pages 331–389. Birkhäuser, Cham, 2018.

[ABM19] Nathalie Aubrun, Sebastián Barbieri, and Étienne Moutot. The domino problem is undecidable on surface groups. In Peter Rossmanith, Pinar Heggernes, and Joost-Pieter Katoen, editors, *44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany*, volume 138 of *LIPIcs*, pages 46:1–46:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[AEGH24] Tara Abrishami, Louis Esperet, Ugo Giocanti, and Matthias Hamann. Symmetric colorings (in preparation), 2024+.

[AK13] Nathalie Aubrun and Jarkko Kari. Tiling problems on Baumslag-Solitar groups. In *Proceedings of the conference on machines, computations and universality 2013, MCU 2013, Zürich, Switzerland, September 9–11, 2013*, pages 35–46. Waterloo: Open Publishing Association (OPA), 2013.

[AMO07] Goulnara Arzhantseva, Ashot Minasyan, and Denis Osin. The SQ-universality and residual properties of relatively hyperbolic groups. *Journal of Algebra*, 315(1):165–177, 2007.

[Ant11] Yago Antolín. On Cayley graphs of virtually free groups. *Groups Complexity Cryptology*, 3(2):301–327, 2011.

[Ass82] Patrice Assouad. Sur la distance de Nagata. *Comptes Rendus de l'Académie des Sciences, Paris, Série I*, 294:31–34, 1982.

[AT18] Goulnara Arzhantseva and Romain Tessera. Admitting a Coarse Embedding is Not Preserved Under Group Extensions. *International Mathematics Research Notices*, 2019(20):6480–6498, 2018.

[Bab77] László Babai. Some applications of graph contractions. *Journal of Graph Theory*, 1(2):125–130, 1977.

[Bab91] László Babai. Vertex-transitive graphs and vertex-transitive maps. *Journal of Graph Theory*, 15(6):587–627, 1991.

[Bab97] László Babai. The growth rate of vertex-transitive planar graphs. In *ACM-SIAM Symposium on Discrete Algorithms*, 1997.

[Bar23a] Sebastián Barbieri. Aperiodic subshifts of finite type on groups which are not finitely generated. *Proceedings of the American Mathematical Society*, 151(09):3839–3843, 2023.

[Bar23b] Laurent Bartholdi. The domino problem for hyperbolic groups, 2023.

[BBE+20] Marthe Bonamy, Nicolas Bousquet, Louis Esperet, Carla Groenland, François Pirot, and Alex Scott. Surfaces have (asymptotic) dimension 2, 2020.

[BBE+23] Marthe Bonamy, Nicolas Bousquet, Louis Esperet, Carla Groenland, Chun-Hung Liu, François Pirot, and Alex Scott. Asymptotic dimension of minor-closed families and Assouad–Nagata dimension of surfaces. *Journal of the European Mathematical Society*, 2023.

[BCK+22] Édouard Bonnet, Dibyayan Chakraborty, Eun Jung Kim, Noleen Köhler, Raul Lopes, and Stéphan Thomassé. Twin-width VIII: delineation and win-wins. In Holger Dell and Jesper Nederlof, editors, *17th International Symposium on Parameterized and Exact Computation, IPEC 2022, September 7-9, 2022, Potsdam, Germany*, volume 249 of *LIPIcs*, pages 9:1–9:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[BDLM08] Nikolay Brodskiy, Jerzy Dydak, Michael Levin, and Atish J. Mitra. A Hurewicz theorem for the Assouad–Nagata dimension. *Journal of the London Mathematical Society*, 77(3):741–756, 2008.

[Ber66] Robert L. Berger. The undecidability of the domino problem. *Memoirs of the American Mathematical Society*, 1966.

[Ber68] Georges M. Bergman. On groups acting on locally finite graphs. *Annals of Mathematics. Second Series*, 88:335–340, 1968.

[BF95] Katalin Bencsáth and Benjamin Fine. Reflections on virtually one-relator groups. In *Groups '93 Galway/St. Andrews'. Proceedings of the international conference, held in Galway, Ireland, August 1-14, 1993. Volume 1*, pages 37–57. Cambridge: Cambridge University Press, 1995.

[BGdMT23] Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, and Stéphan Thomassé. Twin-width V: linear minors, modular counting, and matrix multiplication. In Petra Berenbrink, Patricia Bouyer, Anuj Dawar, and Mamadou Moustapha Kanté, editors, *40th International Symposium on Theoretical Aspects of Computer Science, STACS 2023, March 7-9, 2023, Hamburg, Germany*, volume 254 of *LIPIcs*, pages 15:1–15:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.

[BGK+21] Edouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width III: Max Independent Set, Min Dominating Set, and Coloring. In *48th International Colloquium on Automata, Languages, and Programming (ICALP 2021)*, Glasgow, United Kingdom, July 2021. 33 pages, 6 figures.

[BGK+22] Édouard Bonnet, Colin Geniet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width. II: Small classes. *Combinatorial Theory*, 2(2):42, 2022. Id/No 10.

[BGOdM+22] Édouard Bonnet, Ugo Giocanti, Patrice Ossona de Mendez, Pierre Simon, Stéphan Thomassé, and Szymon Toruńczyk. Twin-width. IV: Ordered graphs and matrices. In *Proceedings of the 54th annual ACM SIGACT symposium on theory of computing, STOC '22, Rome, Italy June 20–24, 2022*, pages 924–937. New York, NY: Association for Computing Machinery (ACM), 2022.

[BGTT22] Édouard Bonnet, Colin Geniet, Romain Tessera, and Stéphan Thomassé. Twin-width VII: groups, 2022.

[Bit24] Nicolás Bitar. Realizability of subgroups by subshifts of finite type, 2024.

[BKRT22] Édouard Bonnet, Eun Jung Kim, Amadeus Reinald, and Stéphan Thomassé. Twin-width VI: the lens of contraction sequences. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 - 12, 2022*, pages 1036–1056. SIAM, 2022.

[BKTW22] Édouard Bonnet, Eun Jung Kim, Stéphan Thomassé, and Rémi Watrigant. Twin-width I: Tractable FO Model Checking. *Journal of the ACM*, 69:1–46, 02 2022.

[BN46] Svend Bundgaard and Jakob Nielsen. Forenklede beviser for nogle sætninger i fladetopologien. *Matematisk Tidsskrift. B*, pages 1–16, 1946.

[BNdM+24] Édouard Bonnet, Jaroslav Nešetřil, Patrice Ossona de Mendez, Sebastian Siebertz, and Stéphan Thomassé. Twin-width and permutations, 2024.

[Bod88] Hans L. Bodlaender. Planar graphs with bounded treewidth. Technical Report RUU-CS-88-14, Department of Computer Science, University of Utrecht, 1988.

[BS18] Alexis Ballier and Maya Stein. The domino problem on groups of polynomial growth. *Groups, Geometry, and Dynamics*, 12(1):93–105, 2018.

[BST12] Itai Benjamini, Oded Schramm, and Ádám Timár. On the separation profile of infinite graphs. *Groups, Geometry, and Dynamics*, 6(4):639–658, 2012.

[CDHS11] Johannes Carmesin, Reinhard Diestel, Fabian Hundertmark, and Maya Stein. Connectivity and tree structure in finite graphs. *Combinatorica*, 34:11–46, 2011.

[CGS17] David B. Cohen and Chaim Goodman-Strauss. Strongly aperiodic subshifts on surface groups. *Groups, Geometry, and Dynamics*, 11(3):1041–1059, 2017.

[CGSR22] David B. Cohen, Chaim Goodman-Strauss, and Yo'Av Rieck. Strongly aperiodic subshifts of finite type on hyperbolic groups. *Ergodic Theory and Dynamical Systems*, 42(9):2740–2783, 2022.

[Chi92] Patrick Chiu. Cubic Ramanujan graphs. *Combinatorica*, 12(3):275–285, 1992.

[CHM22] Johannes Carmesin, Matthias Hamann, and Babak Miraftab. Canonical trees of tree-decompositions. *Journal of Combinatorial Theory, Series B*, 152:1–26, 2022.

[CK23] Johannes Carmesin and Jan Kurkofka. Canonical decompositions of 3-connected graphs. *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1887–1920, 2023.

[Coh17] David Bruce Cohen. The large scale geometry of strongly aperiodic subshifts of finite type. *Advances in Mathematics*, 308:599–626, 2017.

[CP93] Michel Coornaert and Athanase Papadopoulos. *Symbolic dynamics and hyperbolic groups*, volume 1539 of *Lect. Notes Math.* Berlin: Springer-Verlag, 1993.

[CP15] David Carroll and Andrew Penland. Periodic points on shifts of finite type and commensurability invariants of groups. *New York Journal of Mathematics*, 21, 02 2015.

[DD89] Warren Dicks and Martin J. Dunwoody. *Groups Acting on Graphs*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 1989.

[DEW17] Vida Dujmović, David Eppstein, and David R. Wood. Structure of graphs with locally restricted crossings. *SIAM Journal on Discrete Mathematics*, 31(2):805–824, 2017.

[DHL18] Reinhard Diestel, Fabian Hundertmark, and Sahar Lemanczyk. Profiles of separations: in graphs, matroids, and beyond. *Combinatorica*, 39, 11 2018.

[Die17] Reinhard Diestel. *Graph Theory: 5th edition.* Springer Graduate Texts in Mathematics. Springer-Verlag, 2017.

[Dis23] Marc Distel. Proper minor-closed classes of graphs have Assouad-Nagata dimension 2. arXiv e-print 2308.10377, 2023.

[DJKK22] Reinhard Diestel, Raphael W. Jacobs, Paul Knappe, and Jan Kurkofka. Canonical graph decompositions via coverings, 2022.

[DJM93] Reinhard Diestel, H. A. Jung, and Rögnvaldur G. Möller. On vertex transitive graphs of infinite degree. *Archiv der Mathematik*, 60:591–600, 1993.

[DK15] Martin J. Dunwoody and Bernhard Krön. Vertex cuts. *Journal of Graph Theory*, 80(2):136–171, 2015.

[DK18] Cornelia Druţu and Michael Kapovich. *Geometric group theory. With an appendix by Bogdan Nica*, volume 63 of *Colloquium Publications. American Mathematical Society.* Providence, RI: American Mathematical Society (AMS), 2018.

[dlH00] Pierre de la Harpe. *Topics in geometric group theory*. Chicago: The University of Chicago Press, 2000.

[DMW23] Vida Dujmović, Pat Morin, and David R. Wood. Graph product structure for non-minor-closed classes. *Journal of Combinatorial Theory, Series B*, 162:34–67, 2023.

[DO95] Guoli Ding and Bogdan Oporowski. Some results on tree decomposition of graphs. *Journal of Graph Theory*, 20(4):481–499, 1995.

[Dro06] Carl Droms. Infinite-ended groups with planar Cayley graphs. *Journal of Group Theory*, 9(4):487–496, 2006.

[DSS98] Carl Droms, Brigitte Servatius, and Herman Servatius. The structure of locally finite two-connected graphs. *Electronic Journal of Combinatorics*, 2, 01 1998.

[DT99] Reinhard Diestel and Robin Thomas. Excluding a countable clique. *Journal of Combinatorial Theory, Series B*, 76(1):41–67, 1999.

[Dun85] Martin J. Dunwoody. The accessibility of finitely presented groups. *Inventiones Mathematicae*, 81:449–458, 1985.

[Dun09] Martin J. Dunwoody. Planar graphs and covers, 2009.

[EFW12] Alex Eskin, David Fisher, and Kevin Whyte. Coarse differentiation of quasi-isometries. I: Spaces not quasi-isometric to Cayley graphs. *Annals of Mathematics. Second Series*, 176(1):221–260, 2012.

[EG24a] Louis Esperet and Ugo Giocanti. Coarse geometry of quasi-transitive graphs beyond planarity (to appear in the Electronic Journal of Combinatorics), 2024.

[EG24b] Louis Esperet and Ugo Giocanti. Optimization in graphical small cancellation theory. *Discrete Mathematics*, 347(4):7, 2024.

[EGLD23] Louis Esperet, Ugo Giocanti, and Clément Legrand-Duchesne. The structure of quasi-transitive graphs avoiding a minor with applications to the domino problem. In *EUROCOMB 2023, Proceedings of the 12th European Conference on Combinatorics, Graph Theory and Applications*, page 425–432, 2023.

[EGM24] Solène J. Esnay, Ugo Giocanti, and Étienne Moutot. Weakly not strongly aperiodic subshifts of finite type in one-relator and quasi-planar groups (in preparation), 2024+.

[EM22] Solène J. Esnay and Étienne Moutot. Aperiodic SFTs on Baumslag-Solitar groups. *Theoretical Computer Science*, 917:31–50, 2022.

[Fox52] Ralph H. Fox. On Fenchel's Conjecture about F-Groups. *Matematisk Tidsskrift. B*, pages 61–65, 1952.

[Fre44] Hans Freudenthal. Über die Enden diskreter Räume und Gruppen. *Commentarii Mathematici Helvetici*, 17:1–38, 1944.

[Geo14] Agelos Georgakopoulos. Characterising planar Cayley graphs and Cayley complexes in terms of group presentations. *European Journal of Combinatorics*, 36:282–293, 2014.

[Geo20] Agelos Georgakopoulos. On planar Cayley graphs and Kleinian groups. *Transactions of the American Mathematical Society*, 373(7):4649–4684, 2020.

[GH15] Agelos Georgakopoulos and Matthias Hamann. The planar Cayley graphs are effectively enumerable I: Consistently planar graphs. *Combinatorica*, 39:993–1019, 2015.

[GH23] Agelos Georgakopoulos and Matthias Hamann. The planar Cayley graphs are effectively enumerable II. *European Journal of Combinatorics*, 110:103668, 2023.

[GH24] Agelos Georgakopoulos and Matthias Hamann. A full halin grid theorem, 2024.

[Gio24] Ugo Giocanti. A note on the structure of locally finite planar quasi-transitive graphs (in preparation), 2024+.

[GMP20] Daniel Gonçalves, Mickaël Montassier, and Alexandre Pinlou. Acyclic coloring of graphs and entropy compression method. *Discrete Mathematics*, 343(4):111772, 2020.

[GP23] Agelos Georgakopoulos and Panos Papasoglu. Graph minors and metric spaces. arXiv e-print 2305.07456, 2023.

[Gro87] Mikhaïl Gromov. Hyperbolic groups. Essays in group theory, Publ., Math. Sci. Res. Inst. 8, 75-263 (1987)., 1987.

[Gro93] Mikhaïl Gromov. *Geometric group theory. Volume 2: Asymptotic invariants of infinite groups. Proceedings of the symposium held at the Sussex University, Brighton, July 14-19, 1991*, volume 182 of *London Mathematical Society Lecture Note Series*. Cambridge: Cambridge University Press, 1993.

[Gro03] Mikhail Gromov. Random walk in random groups. *Geometric and Functional Analysis*, 13:73–146, 2003.

[Gro16] Martin Grohe. Quasi-4-connected components. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 8:1–8:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.

[Gru15] Dominik Gruber. Groups with graphical $C(6)$ and $C(7)$ small cancellation presentations. *Transactions of the American Mathematical Society*, 367(3):2051–2078, 2015.

[Hal65] Rudolf Halin. Über die maximalzahl fremder unendlicher wege in graphen. *Mathematische Nachrichten*, 30(1-2):63–85, 1965.

[Hal73] Rudolf Halin. Automorphisms and endomorphisms of infinite locally finite graphs. *Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg*, 39:251–283, 1973.

[Ham15] Matthias Hamann. Generating the cycle space of planar graphs. *The Electronic Journal of Combinatorics*, 22(2):research paper p2.34, 8, 2015.

[Ham18a] Matthias Hamann. Accessibility in transitive graphs. *Combinatorica*, 38(4):847–859, 2018.

[Ham18b] Matthias Hamann. Planar transitive graphs. *The Electronic Journal of Combinatorics*, 25(4):research paper p4.8, 18, 2018.

[Ham24] Matthias Hamann. Minor exclusion in quasi-transitive graphs. *Discrete Mathematics*, 347(1):4, 2024. Id/No 113738.

[HLMR22] Matthias Hamann, Florian Lehner, Babak Miraftab, and Tim Rühmann. A Stallings type theorem for quasi-transitive graphs. *Journal of Combinatorial Theory, Series B*, 157:40–69, 2022.

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43:439–561, 2006.

[Jea15a] Emmanuel Jeandel. Aperiodic subshifts on polycyclic groups, 2015.

[Jea15b] Emmanuel Jeandel. Translation-like actions and aperiodic subshifts on groups, 2015.

[Khu23] Ana Khukhro. A characterisation of virtually free groups via minor exclusion. *IMRN. International Mathematics Research Notices*, 2023(15):12967–12976, 2023.

[KL05] Dietrich Kuske and Markus Lohrey. Logical aspects of Cayley-graphs: the group case. *Annals of Pure and Applied Logic*, 131(1):263–286, 2005.

[Kön27] Dénes König. Über eine Schlußweise aus dem Endlichen ins Unendliche. *Acta Litterarum ac Scientiarum. Regiae Universitatis Hungaricae Francisco-Josephinae. Sectio Scientiarum Mathematicarum*, 3:121–130, 1927.

[KPS73] Abraham Karrass, Alfred Pietrowski, and Donald Solitar. Finite and infinite cyclic extensions of free groups. *Journal of the Australian Mathematical Society*, 16(4):458–466, 1973.

[KT90] Igor Kříž and Robin Thomas. Clique-sums, tree-decompositions and compactness. *Discrete Mathematics*, 81(2):177–185, 1990.

[Lev70] Henry Levinson. On the genera of graphs of group presentations. *Annals of the New York Academy of Sciences*, 175(1):277–284, 1970.

[Liu23] Chun-Hung Liu. Assouad-Nagata dimension of minor-closed metrics. arXiv e-print 2308.12273, 2023.

[LS01] Roger C. Lyndon and Paul E. Schupp. *Combinatorial Group Theory*, volume 89 of *Classics in Mathematics*. Springer Berlin Heidelberg, 2001.

[LS05] Urs Lang and Thilo Schlichenmaier. Nagata dimension, quasisymmetric embeddings, and Lipschitz extensions. *IMRN. International Mathematics Research Notices*, 2005(58):3625–3655, 2005.

[Mac67] Angus Macbeath. The classification of non-euclidean plane crystallographic groups. *Canadian Journal of Mathematics*, 19:1192 – 1205, 1967.

[Mac23] Joseph MacManus. Accessibility, planar graphs, and quasi-isometries. arXiv e-print 2310.15242, 2023.

[Mag30] Wilhelm Magnus. Über diskontinuierliche Gruppen mit einer definierenden Relation. (Der Freiheitssatz.). *Journal für die Reine und Angewandte Mathematik*, 163:141–165, 1930.

[Mas96] Heinrich Maschke. The representation of finite groups, especially of the rotation groups of the regular bodies of three-and four-dimensional space, by cayley's color diagrams. *American Journal of Mathematics*, 18(2):156–194, 1896.

[Mil68] John W. Milnor. A note on curvature and fundamental group. *Journal of Differential Geometry*, 2:1–7, 1968.

[Mil87] Gary L. Miller. An additivity theorem for the genus of a graph. *Journal of Combinatorial Theory, Series B*, 43:25–47, 1987.

[MS83] David E. Muller and Paul E. Schupp. Groups, the theory of ends, and context-free languages. *Journal of Computer and System Sciences*, 26(3):295–310, 1983.

[MT10] Robin A. Moser and Gábor Tardos. A constructive proof of the general Lovász Local Lemma. *Journal of the ACM*, 57(2), 2010.

[Oll06] Yann Ollivier. On a small cancellation theorem of Gromov. *Bulletin of the Belgian Mathematical Society - Simon Stevin*, 13(1):75–89, 2006.

[OR15] Mikhail Ostrovskii and David Rosenthal. Metric dimensions of minor excluded graphs and minor exclusion in groups. *International Journal of Algebra and Computation*, 25(04):541–554, 2015.

[Osa20] Damian Osajda. Small cancellation labellings of some infinite graphs and applications. *Acta Mathematica*, 225:159–191, 2020.

[Pia08] Steven T. Piantadosi. Symbolic dynamics on free groups. *Discrete and Continuous Dynamical Systems*, 20(3):725–738, 2008.

[Ros20] Matthieu Rosenfeld. Another approach to non-repetitive colorings of graphs of bounded degree. *Electronic Journal of Combinatorics*, 27(3):3, 2020.

[RS83] Neil Robertson and Paul D. Seymour. Graph minors I excluding a forest. *Journal of Combinatorial Theory, Series B*, 35:39–61, 1983.

[RS91] Neil Robertson and Paul Seymour. Graph minors X. obstructions to tree-decomposition. *Journal of Combinatorial Theory, Series B*, 52(2):153–190, 1991.

[RS93] Neil Robertson and Paul Seymour. Excluding a graph with one crossing. In *Graph structure theory*, pages 669–675. American Mathematical Society, 1993.

[RS03] Neil Robertson and Paul Seymour. Graph minors. XVI. Excluding a non-planar graph. *Journal of Combinatorial Theory, Series B*, 89(1):43–76, 2003.

[Sab58] Gert Sabidussi. On a class of fixed-point-free graphs. *Proceedings of the American Mathematical Society*, 9(5):800–804, 1958.

[Ser80] Jean-Pierre Serre. *Trees*. Springer-Verlag, Berlin, 1980. Translated from the French by John Stillwell.

[Sta68] John R. Stallings. On torsion-free groups with infinitely many ends. *Annals of Mathematics. Second Series*, 88:312–334, 1968.

[Tho80] Carsten Thomassen. Planarity and duality of finite and infinite graphs. *Journal of Combinatorial Theory, Series B*, 29(2):244–271, 1980.

[Tho91] Carsten Thomassen. Tilings of the torus and the Klein bottle and vertex-transitive graphs on a fixed surface. *Transactions of the American Mathematical Society*, 323(2):605–635, 1991.

[Tho92] Carsten Thomassen. The Hadwiger number of infinite vertex-transitive graphs. *Combinatorica*, 12:481–491, 1992.

[Tut84] William T. Tutte. *Graph Theory.* Encyclopedia of Mathematics and its Applications. Cambridge University Press, 1984.

[TW93] Carsten Thomassen and Wolfgang Woess. Vertex-transitive graphs and accessibility. *Journal of Combinatorial Theory, Series B*, 58(2):248–268, 1993.

[Š55] Albert S. Švarc. A volume invariant of coverings (in russian). *Doklady Akademii Nauk SSSR*, 105:32–34, 1955.

[Wag37] Klaus Wagner. Über eine Eigenschaft der ebenen Komplexe. *Mathematische Annalen*, 114:570–590, 1937.

[Wal71] Charles T. C. Wall. Pairs of relative cohomological dimension one. *Journal of Pure and Applied Algebra*, 1:141–154, 1971.

[Wil66] H. C. Wilkie. On non-euclidean crystallographic groups. *Mathematische Zeitschrift*, 91:87–102, 1966.

[Woe89] Wolfgang Woess. Graphs and groups with tree-like properties. *Journal of Combinatorial Theory, Series B*, 47(3):361–371, 1989.

[Woe91] Wolfgang Woess. Topological groups and infinite graphs. *Discrete Mathematics*, 95(1-3):373–384, 1991.

[ZVC80] Heiner Zieschang, Elmar Vogt, and Hans-Dieter Coldewey. *Surfaces and planar discontinuous groups. Revised and expanded transl. from the German by J. Stillwell*, volume 835 of *Lecture Notes in Mathematics.* Springer, Cham, 1980.