The planar Cayley graphs are effectively enumerable

Agelos Georgakopoulos*
Mathematics Institute
University of Warwick
CV4 7AL, UK

Matthias Hamann†
Department of Mathematics
University of Hamburg
Bundesstraße 55
20146 Hamburg, Germany

November 30, 2015

Abstract

We show that a group admits a planar, finitely generated Cayley graph if and only if it admits a special kind of group presentation we introduce, called a planar presentation. Planar presentations can be recognised algorithmically. As a consequence, we obtain an effective enumeration of the planar Cayley graphs, yielding in particular an affirmative answer to a question of Droms et al. asking whether the planar groups can be effectively enumerated.

1 Introduction

1.1 Overview

Groups which act discretely on the real plane $\mathbb{R}^2$ by homeomorphisms, called *planar discontinuous groups*, are a classical topic the study of which goes back at least as far as Poincaré [24], and are now fully classified and considered to be well understood. The finite ones were classified by Maschke [20], and important contributions to the infinite case where made by Wilkie [30] and Macbeath [19]. See [18, Prop. III. 5. 4] or [31] for a survey. These groups are closely related to surface groups, which have influenced most of combinatorial group theory [23].

*Supported by EPSRC grant EP/L002787/1, and by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No 639046). The first author would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme 'Random Geometry' where work on this paper was undertaken.

†Both authors have been supported by FWF grant P-19115-N18.

Planar discontinuous groups coincide with the groups admitting a planar modified[1] Cayley complex, and they are Fuchsian when they do not contain orientation-reversing elements [23]. These Cayley complexes correspond exactly to the Cayley graphs that can be embedded into $\mathbb{R}^2$ without accumulation points of vertices.

Planar Cayley graphs that can only be embedded in $\mathbb{R}^2$ with accumulation points of vertices, and their groups, are a bit harder to understand, and still a topic of ongoing research [2, 7, 9, 10, 12, 13, 22]. An important fact is that they are finitely presented and hence accessible [7, 10, 11]. Dunwoody [10, Theorem 3.8] uses this to prove that such a group, or a subgroup of index two of its, is a fundamental group of a graph of groups in which each vertex group is either a planar discontinuous group or a free product of finitely many cyclic groups and all edge groups are finite cyclic groups (possibly trivial).

In this paper we extend the aforementioned correspondence between planar Cayley complexes and accumulation-free planar Cayley graphs by relaxing the notion of planarity of a Cayley complex in such a way that it corresponds exactly to its Cayley graph being planar. Our new notion of '*almost planarity*' of a Cayley complex can be recognised algorithmically (Theorem 7.1). In view of the Adjan-Rabin theorem [1, 25], this is a rather rare case of a decidable, geometric property of Cayley complexes.

The key to this is a concept of *planar group presentation* we introduce; this is a type of group presentation that guarantees the planarity of the corresponding Cayley graph, and conversely, we show that every planar, finitely generated Cayley graph admits such a group presentation.

1.2 Results

The *Cayley complex* X corresponding to a group presentation $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ is the 2-complex obtained from the Cayley graph G of $\mathcal{P}$ by glueing a 2-cell along each closed walk of G induced by a relator $R \in \mathcal{R}$. We say that X is *almost planar*, if it admits a map $\rho : X \to \mathbb{R}^2$ such that the 2-simplices of X are *nested* in the following sense. We say that two 2-simplices of X are nested, if the images of their interiors are either disjoint, or one is contained in the other, or their intersection is the image of a 2-cell bounded by two parallel edges corresponding to an involution $s \in \mathcal{S}$.[2] We call the presentation $\mathcal{P}$ a *planar presentation* if its Cayley complex is almost planar. Our first result is

Theorem 1.1. *Every planar, finitely generated Cayley graph admits a planar presentation.*

The main idea behind this is that if two relators in a presentation induce cycles whose interiors overlap but are not nested, then we could replace a subword of one relator by a subword of the other to produce an equivalent presentation

[1]'modified' in the sense that redundant 2-cells are removed; see [18] for details.

[2]The third option can be dropped by considering the *modified* Cayley complex in the sense of [18], i.e. by representing involutions in $\mathcal{S}$ by single, undirected edges.

with less overlapping; our proof that a presentation with no such overlaps exists is based on the machinery of Dunwoody cuts, cf. [5].

In fact, we prove something much stronger than Theorem 1.1. We introduce a specific type of planar presentation, called a *general planar presentation*, and show that every planar, finitely generated Cayley graph admits such a presentation and, conversely, every general planar presentation has a planar Cayley graph (Theorem 5.8). This converse is the hardest result of this paper. Moreover, as general planar presentations can be recognised algorithmically, we obtain

Corollary 1.2. *The set of planar, finitely generated Cayley graphs can be effectively enumerated.*

By an effective enumeration of an infinite set $\mathcal{G}$ we mean a computer program that outputs elements of this set and nothing else, and every element of $\mathcal{G}$ is in the output after some finite time (repetitions are allowed).

This implies a positive answer to a question of Droms et al. [7, 9], namely whether the groups admitting a planar, finitely generated Cayley graph (called *planar groups*) can be effectively enumerated. This question is motivated by the fact that, as a consequence of the Adjan-Rabin theorem [1, 25], planar groups cannot be recognised by an algorithm (taking a presentation as input), and by the fact that planar discontinuous groups have been effectively enumerated [9]. M. Dunwoody (private communication) informs us that the fact that the planar groups can be effectively enumerated should also follow from his result [10, Theorem 3.8] mentioned above with a little bit of additional work (the main issue here is whether the 'or a subgroup of index two' proviso can be dropped).

We remark that there is a huge variety of planar Cayley graphs: even the 3-regular ones form 37 infinite families [12, 13]. Moreover, as the same group can have many planar Cayley graphs, Corollary 1.2 is stronger than saying that planar groups can be effectively enumerated. The aforementioned classification of the 3-regular planar Cayley graphs cost the first author about 100 pages of work [12, 13]. This task would have been significantly simplified if our results had been available at that time, using a computer aided search based on the algorithm behind Corollary 1.2, which is straightforward to implement.

Our proof method is essentially graph-theoretic, and does not appeal to the theory of planar discontinuous groups. We are optimistic that it can be employed in a wider setup including groups like the Baumslag-Solitar group that act on spaces that generalise the plane.

1.3 Planar presentations

The formal definition of a general planar presentation is given in Section 5. Here, we are going to sketch the most interesting special case of this concept, called a special planar presentation. Such presentations always exist for a 3-connected planar Cayley graph, or more generally, for a Cayley graph that can be embedded in the plane in such a way that its label-preserving automorphisms carry facial paths to facial paths.

We say that $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ is a *special planar presentation*, if it can be endowed with a cyclic ordering σ —from now on called a *spin*— of the symmetrization $\mathcal{S}' = \{s, s^{-1} \mid s \in \mathcal{S}\}$ of its generating set, with the following property. Suppose $W_1 = sUt$ and $W_2 = s'Ut'$, where $s, s', t, t' \in \mathcal{S}'$, are two words, each contained in some rotation of a relator in $\mathcal{R}$ (possibly the same relator), where U is any (possibly trivial) word with letters in $\mathcal{S}'$. Then σ allows us to say whether paths induced by these words W_1, W_2 would *cross* each other or not if we could embed the Cayley graph of $\mathcal{P}$ in the plane in such a way that for every vertex the cyclic ordering of the labels of its incident edges we observe coincides with σ. To make this more precise, we embed a tree consisting of a 'middle' path P with edges labelled by the letters in U, and two leaves attached at each endvertex of P labelled with s, s', t, t' as in Figure 1, where the spin we use at each endvertex of P is the one induced by σ on the corresponding 3-element subset of $\mathcal{S}'$. There are essentially two situations that can arise, both shown in that figure. Naturally, we say that W_1, W_2 *cross* each other in the right-hand situation, and they do not in the left-hand one.

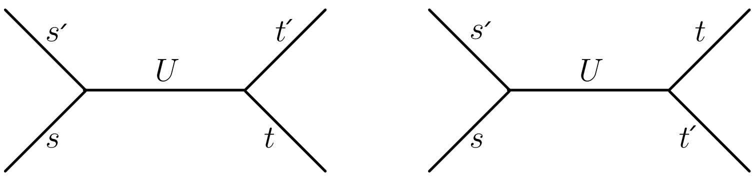

Figure 1: The definition of *crossing*; $W_1 = sUt$ crosses $W_2 = s'Ut'$ in the right, but not in the left.

We then say that $\mathcal{P}$ is a special planar presentation, if there is a spin σ on $\mathcal{S}'$ such that no two words as above cross each other. Note that this is an abstract property of sets of words, and it is defined without reference to the Cayley graph of $\mathcal{P}$; in fact, it can be checked algorithmically. The main essence of this paper is that this is enough to guarantee the planarity of the Cayley graph, and that a converse statement holds.

This generalises an idea from [14], where it was shown that every planar discontinuous group admits a special planar presentation where every relator is *facial*, i.e. it crosses no other word (where we consider words that are not necessarily among our relators).

Our actual definition of a special planar presentation, given in Section 3.1, is in fact a bit more general than the above sketch. Consider for example the Cayley graph of the presentation $\langle a, b \mid a^n, b^2, aba^{-1}b \rangle$. Its Cayley graph is a prism graph with an essentially unique embedding in $\mathbb{R}^2$. Note that the spin of half of its vertices is the reverse of the spin of other half. This is a general phenomenon: every 3-connected Cayley graph has an essentially unique embedding, and in that embedding all vertices have the same spin *up to reflection*. However, for every generator s, either the two end-vertices of all edges labelled s have the same spin, or they always have reverse spins. This yields a classification of generators into *spin-preserving* and *spin-reversing* ones, and our definition of a special planar presentation takes this into account; still, everything can be checked algorithmically.

The situation becomes much more complex however if one wants to consider planar Cayley graphs that are not 3-connected. Such graphs do not always have an embedding with all vertices having the same spin up to reflection; perhaps the simplest such example is the one of Figure 2.

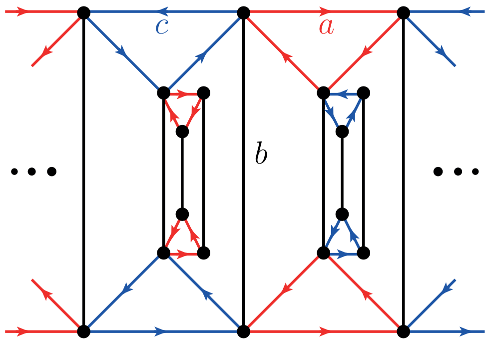

Figure 2: A 2-connected planar Cayley graph from [9], obtained by amalgamating two 6-element groups along an involution, which does not admit a consistent embedding.

In order to capture such Cayley graphs we had to come up with the notion we call a *general planar presentation* (defined in Section 5), which in particular translates, into abstract, algorithmically checkable, properties of words as above, situations as in Figure 2, where a certain generator s with $s^2 = 1$ separates the graph into two parts, and behaves in a spin-preserving way in one part and in a spin-reversing way in the other part. That such general planar presentations always give rise to planar Cayley graphs is the hardest result of this paper, many of its complications arising from the fact that given a general planar presentation with such a 'separating' generator s, it is impossible to predict whether $s = 1$, which would imply that our Cayley graph does not quite have the structure anticipated by the presentation. The situation is complicated further by the fact that separating generators need not be involutions; an example is given in Figure 3.

By considering special planar presentations we also obtain

Theorem 1.3. *The finitely generated Cayley graphs that admit a consistent embedding in $\mathbb{R}^2$ can be effectively enumerated.*

Here, we call an embedding $\rho : G \to \mathbb{R}^2$ *consistent*, if every vertex has the same spin up to reflection, and every generator is either spin-preserving or spin-reversing in the above sense. Equivalently, ρ is consistent, if the label-preserving automorphisms of G carry every facial path with respect to ρ to a facial path.

This paper is structured as follows. After some general definitions, we introduce special planar presentations in Section 3, and show that every 3-connected planar Cayley graph admits such a presentation. Next, we show that the Cayley

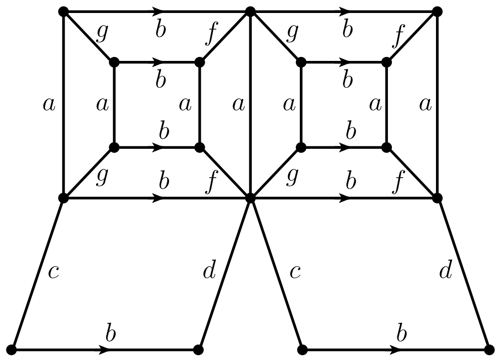

Figure 3: An (infinite) planar Cayley graph, corresponding to the presentation $\langle a, b, c, d, f, g \mid a^2, c^2, d^2, f^2, g^2, (af)^2, (ag)^2, abab^{-1}gbfb^{-1}, cbdb^{-1} \rangle$, with a separating edge b which is not an involution.

graph of every special planar presentation is planar in Section 4. Then we consider the general case, which again we split into two directions: we show that the Cayley graph of every general planar presentation is planar in Section 5, and that every planar Cayley graph admits a general planar presentation in Section 6. In Section 7 we put these facts together to obtain the results mentioned above, and we finish with some open problems in Section 8.

2 Definitions

2.1 Cayley graphs and group presentations

We will follow the terminology of [6] for graph-theoretical terms and that of [4] and [21] for group-theoretical ones. Let us recall the definitions most relevant for this paper.

A *group presentation* $\langle \mathcal{S} \mid \mathcal{R} \rangle$ consists of a set $\mathcal{S}$ of distinct symbols, called the *generators* and a set $\mathcal{R}$ of words with letters in $\mathcal{S} \cup \mathcal{S}^{-1}$, where $\mathcal{S}^{-1}$ is the set of symbols $\{s^{-1} \mid s \in \mathcal{S}\}$, called *relators*. Each such group presentation uniquely determines a group, namely the quotient group $F_{\mathcal{S}}/N$ of the (free) group $F_{\mathcal{S}}$ of words with letters in $\mathcal{S} \cup \mathcal{S}^{-1}$ over the (normal) subgroup $N = N(\mathcal{R})$ generated by all conjugates of elements of $\mathcal{R}$.

The Cayley graph $Cay(\mathcal{P}) = Cay\langle \mathcal{S} \mid \mathcal{R} \rangle$ of a group presentation $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ is an edge-coloured directed graph $G = (V, E)$ constructed as follows. The vertex set of G is the group $\Gamma = F_{\mathcal{S}}/N$ corresponding to $\mathcal{P}$. The set of *colours* we will use is $\mathcal{S}$. For every $g \in \Gamma, s \in \mathcal{S}$ join g to gs by an edge coloured s directed from g to gs. Note that Γ acts on G by multiplication on the left; more precisely, for every $g \in \Gamma$ the mapping from $V(G)$ to $V(G)$ defined by $x \mapsto gx$ is an *automorphism* of G, that is, an automorphism of G that preserves the colours and directions of the edges. In fact, Γ is precisely the group of such automorphisms of G. Any presentation of Γ in which $\mathcal{S}$ is the set of generators will also be called a presentation of $Cay(\mathcal{P})$.

Note that some elements of $\mathcal{S}$ may represent the identity element of Γ, and distinct elements of $\mathcal{S}$ may represent the same element of Γ; therefore, $Cay(\mathcal{P})$ may contain loops and parallel edges of the same colour.

If $s \in \mathcal{S}$ is an *involution*, i.e. $s^2 = 1$, then every vertex of G is incident with a pair of parallel edges coloured s (one in each direction). If s^2 is a relator in $\mathcal{R}$, we will follow the convention of replacing this pair of parallel edges by a single, undirected edge. This convention is common in the literature [18], and is convenient when studying planar Cayley graphs.

If G is a Cayley graph, then we use $\Gamma(G)$ to denote its group.

If R is any (finite or infinite) word with letters in $\mathcal{S} \cup \mathcal{S}^{-1}$, and g is a vertex of $G = Cay\langle \mathcal{S} \mid \mathcal{R} \rangle$, then starting from g and following the edges corresponding to the letters in R in order we obtain a walk W in G. We then say that W is *induced* by R at g, and we will sometimes denote W by gR; note that for a given R there are several walks in G induced by R, one for each starting vertex $g \in V(G)$.

Let $H_1(G)$ denote the first simplicial homology group of G over $\mathbb{Z}$. We will use the following well-known fact which is easy to prove.

Lemma 2.1. *Let $G = Cay\langle \mathcal{S} \mid \mathcal{R} \rangle$ be a Cayley graph. Then the (closed) walks in G induced by relators in $\mathcal{R}$ generate $H_1(G)$.*

2.2 Graph-theoretical concepts

Let $G = (V, E)$ be a connected graph fixed throughout this section. Two paths in G are *independent*, if they do not meet at any vertex except perhaps at common endpoints. If P is a path or cycle we will use $|P|$ to denote the number of vertices in P and $||P||$ to denote the number of edges of P. Let xPy denote the subpath of P between its vertices x and y.

A *hinge* of G is an edge $e = xy$ such that the removal of the pair of vertices x, y disconnects G. A hinge should not be confused with a *bridge*, which is an edge whose removal separates G although its endvertices are not removed.

The set of neighbours of a vertex x is denoted by $N(x)$.

G is called *k-connected* if $G - X$ is connected for every set $X \subseteq V$ with $|X| < k$. Note that if G is k-connected then it is also $(k-1)$-connected. The *connectivity* $\kappa(G)$ of G is the greatest integer k such that G is k-connected.

A 1-way infinite path is called a *ray*. Two rays are equivalent if no finite set of vertices separates them. The corresponding equivalence classes of rays are the *ends* of G. A graph is *multi-ended* if it has more than one end. Note that given any two finitely generated presentations of the same group, the corresponding Cayley graphs have the same number of ends. Thus this number, which is known to be one of $0, 1, 2, \infty$, is an invariant of finitely generated groups.

A *double ray* is a directed 2-way infinite path.

The set of all finite sums of (finite) cycles forms a vector space over $\mathbb{F}_2$, the *cycle space* of G.

2.3 Embeddings in the plane

An *embedding* of a graph G will always mean a topological embedding of the corresponding 1-complex in the euclidean plane $\mathbb{R}^2$; in simpler words, an embedding is a drawing in the plane with no two edges crossing.

A *face* of an embedding $\sigma : G \to \mathbb{R}^2$ is a component of $\mathbb{R}^2 \setminus \sigma(G)$. The *boundary* of a face F is the set of vertices and edges of G that are mapped by σ to the closure of F. The *size* of F is the number of edges in its boundary. Note that if F has finite size then its boundary is a cycle of G.

A walk in G is called *facial* with respect to σ if it is contained in the boundary of some face of σ.

An embedding of a Cayley graph is called *consistent* if, intuitively, it embeds every vertex in a similar way in the sense that the group action carries faces to faces. Let us make this more precise. Given an embedding σ of a Cayley graph G with generating set S, we consider for every vertex x of G the embedding of the edges incident with x, and define the *spin* of x to be the cyclic order of the set $L := \{xy^{-1} \mid y \in N(x)\}$ in which xy_1^{-1} is a successor of xy_2^{-1} whenever the edge xy_2 comes immediately after the edge xy_1 as we move clockwise around x. Note that the set L is the same for every vertex of G, and depends only on S and on our convention on whether to draw one or two edges per vertex for involutions. This allows us to compare spins of different vertices. Call an edge of G *spin-preserving* if its two endvertices have the same spin in σ, and call it *spin-reversing* if the spin of one of its endvertices is the reverse of the spin of its other endvertex. Call a colour in S *consistent* if all edges bearing that colour are spin-preserving or all edges bearing that colour are spin-reversing in σ. Finally, call the embedding σ *consistent* if every colour is consistent in σ. Note that if σ is consistent, then there are only two types of spin in σ, and they are the reverse of each other.

The following classical result was proved by Whitney [29, Theorem 11] for finite graphs and by Imrich [17] for infinite ones.

Theorem 2.2. *Let G be a 3-connected graph embedded in the sphere. Then every automorphism of G maps each facial path to a facial path.*

This implies in particular that if σ is an embedding of the 3-connected Cayley graph G, then the cyclic ordering of the colours of the edges around any vertex of G is the same up to orientation. In other words, at most two spins are allowed in σ. Moreover, if two vertices x, y of G that are adjacent by an edge, bearing a colour b say, have distinct spins, then any two vertices x', y' adjacent by a b-edge also have distinct spins. We just proved

Lemma 2.3. *Let G be a 3-connected planar Cayley graph. Then every embedding of G is consistent.*

Cayley graphs of connectivity 2 do not always admit a consistent embedding [9]. However, in the cubic case they do; see [13].

An embedding is *Vertex-Accumulation-Point-free*, or *accumulation-free* for short, if the images of the vertices have no accumulation point in $\mathbb{R}^2$.

A *crossing* of a path X by a path or walk Y in a plane graph is a subwalk $Q = e\mathring{Q}f$ of Y where the end-edges e, f of Q are incident with X on opposite sides of X (but not contained in X) and (the image of) $\mathring{Q}$ is contained in X (Figure 4). Note that if Q is a crossing of X by Y, then X contains a crossing $Q' = g\mathring{Q}h$ of Y by X, which we will call the *dual crossing* of Q.

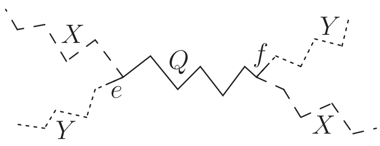

Figure 4: A crossing of X by Y.

3 3-connected planar Cayley graphs admit special planar presentations

3.1 Planar presentations — the special case

We now give the crucial definition of our paper, that of a (special) planar presentation. The intuition behind it comes from the notion of a consistent embedding given above: a planar presentation is a group presentation endowed with some additional data (forming what we will call an embedded presentation) which, once we have proven planarity of the corresponding Cayley graph G, will describe the local structure of a consistent embedding of G, that is, the spin and the information of which generators preserve or reflect it.

Given a group presentation $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$, we will distinguish between two types of generators s: those for which we have s^2 as a relator in $\mathcal{R}$ and the rest. The reasons for this distinction will become clear later. Generators t for which the relation t^2 is provable but not explicitly part of the presentation might exist, but do not cause us any concerns. Given a group presentation $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$, we thus let $\mathcal{I} = \mathcal{I}(\mathcal{P})$ denote the set of elements $s \in \mathcal{S}$ such that $\mathcal{R}$ contains the relator s^2 or s^{-2}.

Let $\mathcal{S}' = \mathcal{S} \cup (\mathcal{S} \setminus \mathcal{I})^{-1}$. For example, if $\mathcal{P} = \langle a, b, c \mid a^2, b^2 \rangle$, then $\mathcal{S}' = \{a, b, c, c^{-1}\}$.

A *spin* on $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ is a cyclic ordering of $\mathcal{S}'$ (to be thought of as the cycling ordering of the edges that we expect to see around each vertex of our Cayley graph once we have proved that it is planar)

An *embedded presentation* is a triple $\mathcal{P}, \sigma, \tau$ where $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ is a group presentation, σ is a spin on $\mathcal{P}$, and τ is a function from $\mathcal{S}$ to $\{0, 1\}$ (encoding the information of whether each generator is spin-preserving or spin-reversing).

To every embedded presentation $\mathcal{P}, \sigma, \tau$ we can associate a tree $\mathbb{T}$ with an accumulation-free embedding in $\mathbb{R}^2$. As a graph, we let $\mathbb{T}$ be $Cay \langle \mathcal{S} \mid s^2, s \in \mathcal{I} \rangle$. Easily, we can embed $\mathbb{T}$ in $\mathbb{R}^2$ in such a way that for every vertex v of $\mathbb{T}$, one of the two cyclic orderings of the colours of the edges of v inherited by the embedding coincides with σ and moreover, for every two adjacent vertices v, w of $\mathbb{T}$, the clockwise cyclic ordering of the colours of the edges of v coincides with that of w if and only if $\tau(a) = 0$ where a is the colour of the v–w edge. (If $\tau(a) = 1$, then the clockwise ordering of v coincides with the anti-clockwise ordering of w.)

Given a word W, we let W^∞ be the 2-way infinite word obtained by concatenating infinitely many copies of W. We say that two words $W, Z \in \mathcal{R}$ *cross*, if there is a 2-way infinite path R of $\mathbb{T}$ induced by W^∞ and a 2-way infinite path L induced by Z^∞ such that L meets both components of $\mathbb{R}^2 \setminus R$.

For example, consider the presentation $\mathcal{P} = \langle n, e, s, w \mid n^2, e^2, s^2, w^2 \rangle$, the spin n, e, s, w, n (read 'north, east, south, west'), and τ identically 0. Then any word containing ns as a subword crosses any word containing ew. The word $nesw$ however crosses no other word, and indeed adding that word to the above presentation yields a planar Cayley graph: the square grid.

Definition 3.1. A *(special) planar presentation* is an embedded presentation $(\mathcal{P}, \sigma, \tau)$ such that

(sP1) no two relators $W, Z \in \mathcal{R}$ cross, and

(sP2) for every relator R, the number of occurrences of letters s in R with $\tau(s) = 1$ (i.e. spin-reversing letters) is even; here, the symbol s^n counts as $|n|$ occurrences of s.

Requirement (sP2) is necessary, as the spin of the starting vertex of a cycle must coincide with that of the last vertex.

The following lemma will later allow us to assume without loss of generality that no relator of $\mathcal{P}$ is a sub-word of a rotation of another relator.

Lemma 3.2. *Let $(\mathcal{P} = \langle \mathcal{S}, \mathcal{R} \rangle, \sigma, \tau)$ be a special planar presentation. Then there is a special planar presentation $(\mathcal{P}' = \langle \mathcal{S}, \mathcal{R}' \rangle, \sigma, \tau)$ such that $\mathcal{P}$ and $\mathcal{P}'$ yield the same Cayley graph, and no element of $\mathcal{R}'$ is a proper subword of another element of $\mathcal{R}'$.*

Proof. We will perform induction on the total length of the words in $\mathcal{R}$.

Suppose that $\mathcal{R}$ contains a word W and a (rotation of a) proper superword $R = WR_2$ of W. To begin with, we may assume that the double rays oW^∞ and oR^∞ induced by them do not coincide. For if this is the case, then $W = U^n$ and $R = U^m$ for some common subword U, and it is easy to modify $\mathcal{R}$ to avoid this situation by replacing R, W with an appropriate power of U.

We may further assume that the double rays oW^∞ and oR^∞ are 'as close as possible' to each other in the following sense.

Let $\mathbb{T}$ be the plane tree corresponding to $\mathcal{P}, \sigma, \tau$ as defined before Definition 3.1. Given three double rays P, S, T in $\mathbb{T}$ which are pairwise non-crossing, we say that S lies *between* P, T if $\bigcap\{P, S, T\} \neq \emptyset$ and P and T lie in distinct components of $\mathbb{R}^2 \setminus S$.

Note that if we fix P and T, then there can only be finitely many double rays S induced by a word in $\mathcal{R}$ lying between P, T because $\mathcal{R}$ is finite.

We say that the double rays P, T in $\mathbb{T}$ are *neighbours*, if no double ray induced by a word in $\mathcal{R}$ lies between them.

Now if $\mathcal{R}$ contains a word W and a proper superword $R = WR_2$ of W, we may assume that the double rays oW^∞ and oR^∞ induced by them are neighbours, because any periodic double ray lying between them contains the path oW and is therefore induced by a superword or subword of W.

Our aim is to replace the word R by its subword R_2 to obtain an equivalent presentation that is closer to satisfying our assertion than $\mathcal{P}$ is. Therefor, we need to show that R_2 does not cross any of our relators. This will be a consequence of

$$\text{Any word with letters in } \mathcal{S} \text{ crossing } R_2 \text{ (with respect to the spin data } \sigma, \tau\text{) also crosses } R \text{ or } W. \tag{1}$$

Let a, z denote the first and last letter of R respectively. Let y denote the last letter of W, and d the first letter of R_2. Note that W starts with a too, and R_2 ends with z. We have $a \neq y^{-1}$ because W is reduced, and $a \neq z^{-1}, d \neq y^{-1}$ because W is reduced.

Furthermore, we may assume that $y \neq z$: for otherwise we can rotate both W and R by moving the letter $y = z$ from the end to the beginning, extending the intersection of the two words; here we used the fact that R^∞ cannot coincide with W^∞ as we noted above.

Thus a, z^{-1}, y^{-1} are all distinct; let us assume that they appear in σ in that order. Our first task is to decide the relative position of d in σ with respect to those letters. It is still possible that d coincides with z^{-1} or a.

Recall that σ is a cyclic ordering on our letters $\mathcal{S} \cup \mathcal{S}'$. We use the notation $\sigma(l, m)$ to denote those letters coming after l and before m in σ. If we want to include l or m we use the notation $\sigma[l, m)$ or $\sigma(l, m]$.

If $d \in \sigma(a, z^{-1})$, then R would cross itself as its rotations contain both za and yd as subwords, which is impossible by (sP1).

If $d \in \sigma(y^{-1}, a)$, then R would cross W as can be seen in Figure 5 by observing the double ray whose two ends are marked $o(R_2W)^{+\infty}$, $o(R_2W)^{-\infty}$, which double ray is induced by R; here we used the fact that the vertex w at which the path oW ends has the same spin as o in the embedding of $\mathbb{T}$ we used to define crossings, because W satisfies (sP2).

If $d = a$, then we can apply one of the two above arguments to the first vertex at which the rays oR^∞ and oW^∞ split to prove that R crosses either itself or W, which is again a contradiction.

These facts combined prove that $d \in \sigma[z^{-1}, y^{-1})$. We have now gathered enough information about σ to allow us to prove (1).

Suppose that R_2 crosses some word X, i.e. there are two crossing double rays P, T where P is induced by R_2 and T is induced by X.

Let $Q = e\mathring{Q}f \subseteq T$ be a crossing of P by T as defined in Section 2.3 (Figure 4), and let $Q' = g\mathring{Q}h \subseteq P$ be its dual crossing.

Let us assume first that $d \neq z^{-1}$; the case $d = z^{-1}$ will be similar.

We can translate P and/or T by some automorphism of $\mathbb{T}$ so as to ensure that the edge e is incident with the path oR_2 induced by R_2 at o (Figure 5), and in fact e is not incident with the last vertex of that path (but may be incident with o).

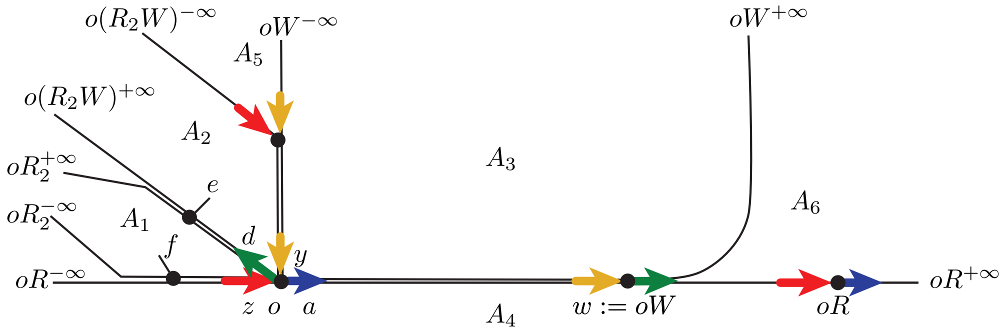

Figure 5: The four double rays $oR^{\infty}, o(R_2W)^{\infty}, oR_2^{\infty}$ and oW^{∞} and the regions A_i they define. Here, the notation $oR^{+\infty}$ means the ray starting at o induced by the 1-way infinite word obtained by repeating R; $oR^{-\infty}$ is defined similarly by repeating the word R^{-1} instead.

Let us fix the embedding ρ of $\mathbb{T}$ in the plane complying with σ, τ as described before Definition 3.1. We use the four double rays $oR^{\infty}, o(R_2W)^{\infty}, oR_2^{\infty}$ and oW^{∞} to divide the plane into regions A_i as shown in Figure 5. Then the edge e must lie in one of these regions, and in each case we obtain a contradiction as follows.

To begin with, note that exactly one of e, f lies in A_1, and we may assume without loss of generality that f does, so e does not lie in A_1.

Case 1: If e lies in A_2, then T crosses not only P but also the double ray $o(R_2W)^{\infty}$. But since R_2W is a rotation of the word $R = WR_2$, (1) is proved in this case.

Case 2: If e lies in A_3, (which can only happen if e is incident with o), then T crosses oW^{∞} and again (1) holds.

Case 3: If e lies in A_4, then T crosses oR^{∞}, and again (1) is proved.

Case 4: If $e = oa$, then as T cannot enter the regions A_3, A_4 for the aforementioned reasons, it has to contain the whole path oW and then continue in the closure of the region A_6. But then some rotation of the relator X inducing T is either a subword or a superword of W, and moreover T lies between oW^{∞} and oR^{∞}. This however contradicts our choice of R, W to be neighbours.

Case 5: If $e = oy^{-1}$, then we apply a similar argument: as T cannot enter the regions A_2, A_3, it has to contain the whole path oW^{-1} and then continue in the closure of the region A_5. Then T lies between oW^{∞} and $o(R_2W)^{\infty}$. As the latter can be induced by R, this again contradicts our choice of R, W as neighbours.

These contradictions complete the proof of (1) in the case where $d \neq z^{-1}$. The case $d = z^{-1}$ is very similar: the only difference is that the region A_1 is a bit smaller in Figure 5.

We now claim that R_2 crosses none of the words in $\mathcal{R} \cup \{R_2\}$.

Indeed, if R_2 crosses a word in $\mathcal{R}$, then applying (1) to that word we obtain a contradiction to the fact that $\mathcal{P}$ was a special planar presentation.

If R_2 crosses itself, then applying (1) again we deduce that R_2 crosses R or W, which are elements of $\mathcal{R}$. But then we are in the previous situation, which cannot occur (here we used the fact that crossing is a symmetric relation).

Keeping the spin data σ, τ, it is clear that R_2 satisfies (sP2) as both W and R satisfied that property.

Since $\mathcal{P}$ satisfied (sP1), and we have just proved that R_2 crosses neither itself nor any other relator in $\mathcal{R}$, this means that the presentation obtained from $\mathcal{P}$ by adding R_2 as a relator is still a special planar presentation. Note that the relator $R = WR_2$ now becomes redundant, and we can remove it. Hence we obtain a special planar presentation of the same Cayley graph in which the relators have strictly smaller total length.

We can repeat this for as long as there are relators in our presentation that are a subword of each other. Since the total length decreases in each step, the process terminates after finitely many steps yielding the desired presentation. $\square$

3.2 The proof

We now prove that planar 3-connected Cayley graphs admit special planar presentations.

Theorem 3.3. *Every planar, locally finite,* 3*-connected Cayley graph admits a special planar presentation.*

In order to prove Theorem 3.3, we need the following general result about generating the first homology group of planar graphs of [16], a generalisation of [15]. For a planar graph G with embedding $\sigma\colon G \to \mathbb{R}^2$, we call two cycles C, D of G *nested* if at most one face of $\mathbb{R}^2 \setminus \sigma(C)$ contains points of $\sigma(D)$. Note that this definition is symmetric in C and D. (It is possible to prove using the 3-connectedness of G that if C, D are not nested, then the corresponding words cross each other in the sense of the previous section.)

Theorem 3.4. [16, Theorem 1.1] *Every planar, locally finite,* 3*-connected graph* G *has an* $\mathrm{Aut}(G)$*-invariant nested set of cycles that generates* $\mathcal{H}_1(G)$*.* $\square$

Proof of Theorem 3.3. Let G be a planar, locally finite, 3-connected Cayley graph, and let $\Gamma := \Gamma(G)$ be its group. By Droms [7, Theorem 5.1], Γ admits a finite presentation $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$. We may replace the generators $\mathcal{S}$ by those finitely many generators that we used to obtain the Cayley graph G, that is, we may assume that $\mathcal{S}$ was used to obtain G. Let $\mathcal{D}$ be a nested $\mathrm{Aut}(G)$-invariant set of cycles in G that generates $\mathcal{H}_1(G)$, which exists by Theorem 3.4, and let $\mathcal{D}' \subseteq \mathcal{D}$ be the subset of those cycles that contain the vertex o. Then the set $\mathcal{R}_{\mathcal{D}'}$ of words corresponding to the cycles in $\mathcal{D}'$ yields a presentation $\mathcal{P}' = \langle \mathcal{S} \mid \mathcal{R}_{\mathcal{D}'} \rangle$ of Γ. Note that a priori the set $\mathcal{D}'$, and hence also $\mathcal{R}_{\mathcal{D}'}$, might be infinite. As

Γ is finitely presented, it is well-known, see e. g. [3], that we can use Tietze-transformations to obtain a finite subset $\mathcal{R}'$ of $\mathcal{R}_{\mathcal{D}'}$ such that $\langle \mathcal{S} \mid \mathcal{R}' \rangle$ is a finite presentation of Γ. We claim that $(\langle \mathcal{S} \mid \mathcal{R}' \rangle, \sigma, \tau)$ is a planar presentation, where σ is the spin of 1 in some embedding ρ of G in $\mathbb{R}^2$, and $\tau(s)$ is 0 for those $s \in \mathcal{S}$ such that the spin of 1 coincides with the spin of the vertex s of G in ρ (and $\tau(s)$ is 1 for every other $s \in \mathcal{S}$).

Indeed, it is easy to check using the nestedness of the finitely many cycles that correspond to the relators $\mathcal{R}'$ and Lemma 2.3 that no two elements of $\mathcal{R}'$ cross. □

We call the Cayley complex X of a presentation $\langle \mathcal{S} \mid \mathcal{R} \rangle$ *almost planar*, if there is a mapping $\sigma\colon X \to \mathbb{R}^2$ such that σ is injective on the 1-skeleton of X, and for every two 2-simplices of X, the images of their interiors under σ are either disjoint or one of these images is contained in the other. (Here, we are using our convention that elements s of $\mathcal{S}$ such that s^2 is a relator in $\mathcal{R}$ give rise to single, undirected edges in X.) Theorem 3.3 has the following consequence, which was conjectured in [12].

Corollary 3.5. *Every planar, locally finite,* 3*-connected planar Cayley graph* G *is the* 1*-skeleton of an almost planar Cayley complex of the group* $\Gamma(G)$ *of* G.

Proof. Since G is planar, there is an embedding $\rho' : G \to \mathbb{R}^2$ by definition. We will extend ρ' to the desired map ρ from the Cayley complex X of $\Gamma(G)$ with respect to the presentation $\langle \mathcal{S} \mid \mathcal{R}' \rangle$ from above. For this, given any 2-cell Y of X with boundary cycle C, we embed Y in the finite component of $\mathbb{R}^2 \setminus C$. It is a straightforward consequence of the nestedness of $\mathcal{D}$ that the resulting map ρ has the desired property. □

4 Planar presentations yield planar Cayley graphs: the consistent case

In this section we prove that every special planar presentation —as defined in Section 3.1— defines a planar Cayley graph with a consistent embedding. This proof contains the fundamental arguments of this paper.

4.1 Fundamental domains

Let again $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ be a presentation, and $G := Cay(\mathcal{P})$ its Cayley graph. Let $T_{\mathcal{S}} = Cay\langle \mathcal{S} \mid \emptyset \rangle$ be the corresponding free tree. Let $N(\mathcal{R})$ denote the normal closure of $\mathcal{R}$ in the group of $\langle \mathcal{S} \mid \emptyset \rangle$, and note that $N(\mathcal{R})$ acts by automorphisms on $T_{\mathcal{S}}$. Then G is, almost by definition, the quotient $T_{\mathcal{S}}/N(\mathcal{R})$ with respect to that action.

In this section we consider all graphs as 1-complexes. The following lemma is folklore; we include a proof for the convenience of the reader.

Lemma 4.1. *$T_{\mathcal{S}}$ admits a connected fundamental domain for the action of $N(\mathcal{R})$.*

Proof. Let D be a maximal subgraph of $T_{\mathcal{S}}$ that is connected and meets each $N(\mathcal{R})$-orbit in at most one point; such a D exists by Zorn's lemma. We claim that D meets every $N(\mathcal{R})$-orbit. For if not, then there exist two adjacent vertices x, xs in $T_{\mathcal{S}}$ (where $s \in \mathcal{S}$) such that xs does not belong to any orbit represented by D but x does. Let $x' = Rx$, where $R \in N(\mathcal{R})$, be the vertex of D that lies in the same orbit as x. Then the vertex Rxs is connected to x', and its orbit is not represented by D. This contradicts the maximality of D, since $D \cup Rxs$ is connected. □

An *open star* is a subspace of a graph consisting of a single vertex and all open half-edges incident with it. A *star* is the union of an open star with some of the midpoints in its closure.

For the connected fundamental domain D provided by Lemma 4.1, we may assume without loss of generality that

$$D \text{ is a union of stars,} \tag{2}$$

since the action of $N(\mathcal{R})$ never identifies two points in the same star.

4.2 Proof of planarity: the consistent case

In this section we prove

Theorem 4.2. *If $(\mathcal{P}, \sigma, \tau)$ is a special planar presentation with countably many generators and relators, then its Cayley graph $Cay(\mathcal{P})$ is planar. Moreover, it admits a consistent embedding, with spin σ and spin-behaviour of generators given by τ.*

For the rest of this section, let us fix $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ as above, and let $G := Cay(\mathcal{P})$. Recall the definition of the embedded tree $\mathbb{T}$ from Section 3.1, and let $o = o_{\mathbb{T}}$ denote the identity of $\mathbb{T}$ seen as a Cayley graph.

Note that it suffices to prove the statement for a finite presentation $\mathcal{P}$; the countably infinite case can then be deduced as follows. If G does not admit a consistent embedding, then by a standard compactness argument there is a finite subgraph $H \subset G$ that does not admit a consistent embedding. It is an easy exercise to show that for some finite $\mathcal{S}' \subset \mathcal{S}$ and $\mathcal{R}' \subset \mathcal{R}$, the Cayley graph $Cay(\langle \mathcal{S}', \mathcal{R}' \rangle)$ also contains H as a subgraph. But $\langle \mathcal{S}', \mathcal{R}' \rangle$ is a finite presentation which is planar with respect to the restriction of σ, τ to $\mathcal{S}', \mathcal{R}'$, leading to a contradiction.

Let D be a connected fundamental domain of $\mathbb{T}$ with respect to the action of $N(\mathcal{R})$, provided by Lemma 4.1. Recall that we may assume that D is a union of stars. Thus the closure $\overline{D}$ of D in $\mathbb{T}$ is the union of D with all midpoints of edges that have exactly one half-edge in D. Moreover, G can be obtained from $\overline{D}$ by identifying pairs of such midpoints: each midpoint m in $\overline{D} \setminus D$ is identified with the unique midpoint m' in D that is $N(\mathcal{R})$-*equivalent* to m, where we call two points or subsets X, Y of $\mathbb{T}$ $N(\mathcal{R})$-equivalent when they lie in the same orbit of the action of $N(\mathcal{R})$ on $\mathbb{T}$. Note that m' might coincide with m, which is the case exactly when it lies on an edge coloured by a generator in the set $\mathcal{I}$ of explicit involutions.

We claim that

$$\text{every two } N(\mathcal{R})\text{-equivalent vertices of } \mathbb{T} \text{ have the same spin.} \tag{3}$$

Indeed, this follows from condition (sP2) of the definition of a planar presentation, according to which every element of $\mathcal{R}$ joins vertices with same spins.

To show that G is planar, and it even admits a consistent embedding, it will suffice to show that these pairs of identified points are *nested* in the embedding of $\overline{D}$ inherited from the embedding of $\mathbb{T}$. Here, we say that two pairs of midpoints x, x' and y, y' in $\overline{D} \setminus D$ are nested, if the x–x' path in $\overline{D}$ does not cross the y–y' path, where we define *crossing* similarly to Section 2.3.

Assuming that such pairs of points are nested, it is easy to prove that G is planar: note that we can cut a closed domain D' of $\mathbb{R}^2$ homeomorphic to a closed disc such that $D' \cap \mathbb{T} = \overline{D}$. Let D'' be a homeomorphic copy of D', and glue D' to D'' by identifying all pairs of corresponding points of their boundaries to obtain a homeomorph S of the sphere. For every pair x, x' of $N(\mathcal{R})$-equivalent points of $\overline{D}$, let X be the x–x' path in $\overline{D}$ and let X'' be its copy in D''. Nestedness implies by definition that these arcs X'' can be continuously deformed into pairwise disjoint arcs. Therefore, the union of $\mathbb{T} \cap \overline{D}$ with all these arcs is an embedding of G on the sphere S (where every midpoint of an edge in $\overline{D}$ became a closed arc). This embedding is consistent because the embedding of $\mathbb{T}$ we started with is.

It thus only remains to prove this nestedness. Our intuition for this is that when both the x–x' path and the y–y' path from above are induced by relators, then these paths cannot cross since that would imply that the corresponding relators cross, which is forbidden. In the next section we will extend this idea to arbitrary pairs of such points, using the fact that the aforementioned paths are cycles of G and cycles of G can be 'proved' using relators.

4.3 Nestedness in $\overline{D}$

Let π denote the canonical covering map from $\mathbb{T}$ to G, and let $o_G := \pi(o_{\mathbb{T}})$ denote the identity element of G. Given a relator W, let W_o denote the closed walk $o_G W$ in G induced by W at o_G. Let $\mathbb{T}_W := \pi^{-1}(W_o)$, and note that $\mathbb{T}_W$ is a union of a set of double-rays of $\mathbb{T}$, which set we denote by $\mathbb{T}[W_o]$, and along each such double-ray we can read the 2-way infinite word W^∞ obtained by repeating W indefinitely; indeed, the only case that could prevent W^∞ from spanning a double-ray is when $W = b$ for some b for which b^2 is a relator, but then W would be a subword of the relator b^2; however, applying Lemma 3.2 we may assume that this is not the case. Equivalently, $\mathbb{T}_W$ is the union of all double rays $N(\mathcal{R})$-equivalent with the double-ray induced by W at $o_{\mathbb{T}}$.

We start our proof of nestedness by noting that in a plane graph, every cycle C separates the graph into two (possibly empty) sides I, O by the Jordan curve theorem, with no edges of the graph joining I to O. Although we have not yet embedded G in the plane, we will be able to show that cycles, or closed walks, of G that are induced by relators enjoy a similar property by exploiting the embedding of $\mathbb{T}$ and the non-crossing property of relators.

4.3.1 Bipartitioning the faces of $\mathbb{T}$

The *dual* graph $\mathbb{T}^*$ of $\mathbb{T}$ is the graph whose vertex set is the set of faces of $\mathbb{T}$, and two faces of $\mathbb{T}$ are joined with an edge e^* of $\mathbb{T}^*$ whenever their boundaries share an edge e of $\mathbb{T}$. Given two faces F, H of $\mathbb{T}$, and an F–H path P_{FH} in $\mathbb{T}^*$, we let $Cr(\mathbb{T}[W_o], P_{FH})$ denote the *number of crossings* of $\mathbb{T}[W_o]$ by P_{FH}; to make this more precise, for a double-ray T in $\mathbb{T}[W_o]$, we write $cr(T, P_{FH})$ for the number of edges e in T such that P_{FH} contains e^*, and we let $Cr(\mathbb{T}[W_o], P_{FH}) := \sum_{T\in\mathbb{T}[W_o]} cr(T, P_{FH})$. We claim that

$$\text{for every two faces } F, H \text{ of } \mathbb{T}, \text{ the parity of the number of crossings } Cr(\mathbb{T}[W_o], P_{FH}) \text{ is independent of the choice of the path } P_{FH}. \tag{4}$$

To see this, note that if C is a cycle in $\mathbb{T}^*$, then $Cr(\mathbb{T}[W_o], C)$ —defined similarly to $Cr(\mathbb{T}[W_o], P_{FH})$— is even because the embedding of $\mathbb{T}$ is accumulation-free and so any ray entering the bounded side of C has to exit it again. This immediately implies (4).

This fact allows us to introduce the following definition

Definition 4.3. Given two faces F, H of $\mathbb{T}$, we write $F \sim H$ if some, and hence every, F–H path P_{FH} in $\mathbb{T}^*$ crosses $\mathbb{T}_W$ an even number of times, i.e. if $Cr(\mathbb{T}[W_o], P_{FH})$ is even.

Note that $\sim$ is an equivalence relation of the set of faces $\mathcal{F}$ of $\mathbb{T}$. Moreover, it uniquely determines an (unordered) bipartition $\{I, O\}$ on $\mathcal{F}$ by choosing one face F and letting $I := \{H \in \mathcal{F} \mid H \sim F\}$ and $O := \mathcal{F} \setminus I$. Note that $\mathbb{T}_W$ decomposes $\mathbb{R}^2$ into regions each of which is in a single side of this bipartition, and crossing $\mathbb{T}_W$ corresponds to alternating between I and O. It will turn out, after we prove that G is planar, that I, O are the lifts of the two sides of the closed walk W_o of G. The following lemma shows that this is plausible and will be needed later.

Lemma 4.4. *The relation $\sim$ is invariant under the action of $N(\mathcal{R})$ on $\mathbb{T}$.*

Proof. It suffices to prove that if F, H are faces in the same orbit of $N(\mathcal{R})$, then $F \sim H$. We may assume that there are vertices x, y in the boundaries of F, H respectively, such that $y = xwRw^{-1}$ for some word w and some relator $R \in \mathcal{R}$: by the definition of the normal closure $N(\mathcal{R})$, if we can prove $F \sim H$ in this case, we can prove $F \sim H$ for every two F, H in the same orbit of $N(\mathcal{R})$.

Since we are free to choose any F–H path in $\mathbb{T}^*$ by (4), let us choose P_{FH} to be one that starts with an edge e^* where e is incident with both x and F, finishes with an edge f^* where f is incident with both y and H, and does not cross the walk from x induced by wRw^{-1} (Figure 6). We need to check that $Cr(\mathbb{T}[W_o], P_{FH})$ is even.

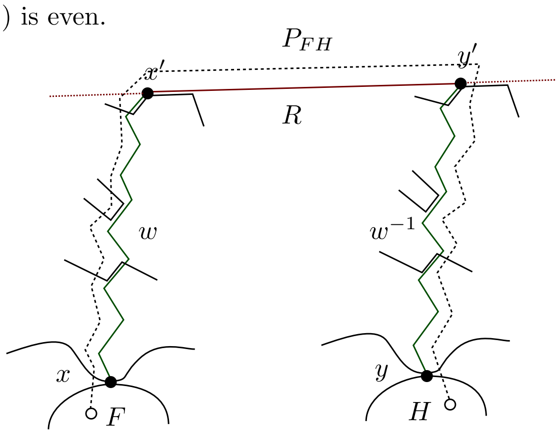

Figure 6: The path P_{FH} (dashed) in the proof of Lemma 4.4 and some elements of $\mathbb{T}[W_o]$ it crosses.

For this, we only need to consider those double-rays $T \in \mathbb{T}[W_o]$ for which $cr(T, P_{FH})$ is odd. We will group those T into pairs, showing that their total contribution is even as desired. For simplicity, we will tacitly assume that $xwRw^{-1}$ crosses any T just once; the general case can be handled with the same arguments.

By elementary topological arguments, any such T is either crossed by our walk $xwRw^{-1}$, or it visits x or y and its two rays are separated by $F \cup xwRw^{-1} \cup H$. We will call the former case a crossing of type A, and the latter a crossing of type B, and we will separately show that crossings of each type come in pairs.

Recall the definition of a crossing from Section 2.3. For type A, we will distinguish the following sub-types of crossings $Q = e\mathring{Q}f \subseteq wRw^{-1}$ of such a T by wRw^{-1}:

1. both e, f lie in R;
2. one of e, f lies in w and the other in w^{-1};
3. both e, f lie in w or they both lie in w^{-1}; or
4. exactly one of e, f lies in R (and the other in either w or w^{-1}).

Type 1 cannot occur, as it would imply that the relator R crosses with the relator of T, and we are assuming $\mathcal{P}$ to be a planar presentation. The second one is also impossible, as it would imply that R is contained in a double-ray in $\mathbb{T}[W_o]$, which would in turn imply that one of T, R is a subword of the other, and we are forbidding this too in a planar presentation.

We will define a bijection between those crossings Q that have an end-edge in w and those that have an end-edge in w^{-1}, showing that the total number of crossings is even as desired.

For crossings of type 3, it is easy to bijectively map each crossing Q with e, f in w to a crossing Q' with e, f in w^{-1}: the automorphism of $\mathbb{T}$ mapping the end-vertex x' of w to the starting vertex y' of w^{-1} translates Q to a crossing Q' as desired, as it translates w to w^{-1} and the element of $\mathbb{T}[W_o]$ crossed by w to another element of $\mathbb{T}[W_o]$.

For crossings of type 4 a similar argument applies, but we have to be more careful. Again, given a crossing Q of an element T of $\mathbb{T}[W_o]$ by wRw^{-1} with e, say, in w and f in R, we translate it to a walk Q' by the automorphism of $\mathbb{T}$ mapping x' to y'. We claim that Q' is a crossing of the translate T' of T by wRw^{-1}. This fact is easier to see in Figure 6 than to explain with words, and it follows from the following three facts: a) the double-ray R^∞ obtained by reading the word R indefinitely starting from x' does not cross T by the definition of a planar presentation; b) T' does not contain all of R, since we are assuming that no relator is a sub-word of another relator, and c) x' and y' have the same spin since they are joined by a relator path R.

It remains to consider crossings of type B, i.e. where T visits x or y and its two rays are separated by $F \cup xwRw^{-1} \cup H$. But for any such T, the automorphism of $\mathbb{T}$ mapping x to y or the other way round maps T to another element of $\mathbb{T}[W_o]$ that is crossed by $xwRw^{-1}$ as often as T is, therefore such crossings appear in pairs as well.

Thus we have paired up all crossings, showing that $Cr(\mathbb{T}[W_o], P_{FH})$ is even. □

A further important property of our bipartition is that

$$\text{for every edge } e \text{ of } \mathbb{T}, \text{ the two faces } F, H \text{ of } e \text{ lie in distinct elements of } \{I, O\} \text{ if and only if } e \in \mathbb{T}_W \text{ and } e \text{ lies in an odd number of elements of } \mathbb{T}[W_o]. \tag{5}$$

Indeed, in this case the edge e^* can be chosen as the $F - H$ path P_{FH} in the definition of $\sim$, and $Cr(\mathbb{T}[W_o], P_{FH})$ is just the number of elements of $\mathbb{T}[W_o]$ containing e.

Remark 4.5. *We can define an equivalence relation on the vertices of* $\mathbb{T}$ *similar to our* $\sim$ *for faces: write* $x \sim y$ *if the (unique)* x*–*y *path* P_{xy} *in* $\mathbb{T}$ *crosses* $\mathbb{T}[W_o]$ *an even number of times. Results similar to Lemma 4.4 and* (5) *extend to this relation, but it is more convenient to work with faces.*

4.3.2 Bipartitioning the 'faces' of G

We would like to use the $N(\mathcal{R})$-invariance of the bipartition of $\mathcal{F}$ we defined above (Lemma 4.4) to induce a bipartition on faces of G, but we cannot talk about faces of G before proving that it is planar. However, there is a way around this: for every face F of $\mathbb{T}$, glue a copy of the domain $\overline{F} \subset \mathbb{R}^2$ to G by identifying each point x of ∂F with $\pi(x)$, where π still denotes our covering map from $\mathbb{T}$. If F, F' are equivalent face boundaries, in other words, if $\pi(\partial F) = \pi(\partial F')$, then we identify the corresponding 2-cells glued onto G. These identifications ensure that every edge of G is in the boundary of either exactly one or exactly two 2-cells (but might appear in the boundary of a 2-cell several times); e is in the boundary of only one 2-cell exactly when the two faces of any lift of e to $\mathbb{T}$ are $N(\mathcal{R})$-equivalent. Indeed, the 2-cells we introduced are bounded by walks corresponding to facial 2-way infinite words, and every edge is in exactly two such words. Moreover, the lifts of e map those walks to exactly the boundaries of pairs of incident faces of $\mathbb{T}$, and such pairs are $N(\mathcal{R})$-equivalent for the various decks of the covering.

Let G^2 denote the set of these 2-cells, and let $\overline{G} = G \cup G^2$ denote the 2-complex consisting of G and these 2-cells.

Lemma 4.4 now means that if Z is a closed walk of G induced by a relator, then $\{I, O\}$ induces a bipartition $\pi[I], \pi[O]$ of G^2. Let us from now on denote this bipartition of G^2 by B_Z.

Our next aim is to extend our construction of that bipartition to an arbitrary cycle of G, showing that every cycle has two 'sides'.

To achieve this, given a cycle C of G, we choose a 'proof' P of C; that is, a sequence of closed walks $W_i, 1 \leq i \leq k$, of G induced by rotations of relators such that $C = \sum_{1 \leq i \leq k} W_i$, where $\sum$ denotes addition in the simplicial homology sense. Such a sequence (W_i) exists by Lemma 2.1. For every W_i, let I_{W_i}, O_{W_i} denote the two sides of the bipartition B_{W_i} of G^2 from above; it will not matter which element of B_{W_i} we denote by I_{W_i} and which by O_{W_i}, so we may make an arbitrary choice here.

Define the bipartition $B_C := \{I_C, O_C\}$ of the 2-cells G^2 of $\overline{G}$ by setting $I_C := \triangle_i I_{W_i}$ and $O_C := G^2 \triangle I_C$, where $\triangle$ denotes the symmetric difference. Note that $O_C = \triangle_i O_{W_i}$ when k is odd and $O_C = G^2 \triangle (\triangle_i O_{W_i})$ when k is even, where k is the number of our W_i.

Once we have constructed a planar embedding of G in the plane, it will turn out that B_C corresponds to the bipartition of the faces of G into those lying inside/outside C.[3] To see why this might be true, consider the situation of Figure 7 as an example; we imagine I_C to be the set of 2-cells inside C and O_C the set of 2-cells outside it, or the other way round. Figure 7 is only a help to our imagination since we have not yet proved G to be planar. Our definition of B_C builds on this idea, but has to deal with the fact that we do not yet have an embedding of G in the plane.

The reader who finds this definition surprising might be comforted to know that this is the subtlest idea of the proof of Theorem 4.2, and it took us a long time to come up with. An important point, which explains it to some extent, is that B_C is independent of our above choices of which side to denote by O_{W_i} and which by I_{W_i}, since a bipartition is an undirected pair of sets (the example of Figure 7 might be helpful again). What is more important is that, as we will

[3] It should be said however, that faces of G do not correspond one-to-one to 2-cells of $\overline{G}$; see Section 8.

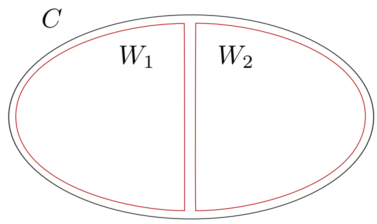

Figure 7: Intuition of the definition of B_C.

see below, B_C is independent of the choice of the proof P. From all we know at the moment however, it is defined as a function of the proof P, so let us denote it by $B_C(P)$.

Our next aim is to show that, in a certain way, $B_C(P)$ behaves like the bipartition of the faces of a plane graph induced by a cycle C: to move between the two sides, one has to cross an edge of C. This is achieved by Lemma 4.7 below, for the proof of which we need the following. A *directed edge* of G is an ordered pair (x, y) such that $xy \in E(G)$. Thus any edge $xy = yx \in E$ corresponds to two directed edges.

Lemma 4.6. *Let e be a directed edge of G, let $W \in \mathcal{R}$ be a relator which is not of the form $b^2 = 1$ for $b \in \mathcal{S}$, and let $o_G W$ be the closed walk of G rooted at o_G induced by W. Then the number of double-rays in $\mathbb{T}[W_o]$ containing e equals the number of times that $o_G W$ traverses $\pi(e)$.*

Proof. If $o_G W$ does not traverse $\pi(e)$ then $\mathbb{T}[W_o]$ avoids e and we are done. So suppose that $o_G W$ does traverse $\pi(e)$. Let $o_G W^\infty$ denote the two-way infinite walk on G obtained by repeating $o_G W$ indefinitely. Let $T \in \mathbb{T}[W_o]$ be the lift of $o_G W^\infty$ to $\mathbb{T}$ (via π^{-1}) sending $\pi(e)$ to e, and note that T is a double-ray containing e. Let Q be the subpath of T that starts with e and finishes when a rotation of the word W is completed. By the definition of $\mathbb{T}[W_o]$, there is a 1–1 correspondence between the elements of $\mathbb{T}[W_o]$ containing e and the directed edges e' in Q that are $N(\mathcal{R})$-equivalent to e: each such element of $\mathbb{T}[W_o]$ can be obtained by translating T by the automorphism of $\mathbb{T}$ sending e' to e.

Now note that $o_G W$ traverses $\pi(e)$ whenever its lift T traverses one of those e'. Combined with the above observations this proves our assertion. □

Lemma 4.7. *For every $e \in E(G)$, the bipartition $B_C(P)$ separates 2-cells of e if and only if $e \in C$.*

Proof. Let I, O be the two elements of $B_C(P)$ as defined above. Then, letting $1_{F \in I}$ denote the indicator function of $F \in I$, we have

$$1_{F \in I} = N_F := |\{W_i \mid F \in I_{W_i}\}| \pmod 2,$$

and similarly

$$1_{H \in I} = N_H := |\{W_i \mid H \in I_{W_i}\}| \pmod 2.$$

But

$$N_F + N_H = |\{W_i \mid W_i \text{ separates } F \text{ from } H\}| \pmod 2$$

by the construction of I, O. We claim that $|\{W_i \mid W_i \text{ separates } F \text{ from } H\}|$ is odd if and only if $e \in E(C)$. Indeed, B_{W_i} separates F from H exactly when W_i traverses e an odd number of times by (5) and Lemma 4.6, and e is in C exactly when there is an odd number of W_i that traverse e an odd number of times.

Since that number is even if $e \notin E(C)$ and odd otherwise, our last congruence yields $N_F + N_H = 1 \pmod 2$ if and only if $e \in E(C)$. Therefore, the previous congruences imply that $1_{F \in I} = 1_{H \in I}$ if $e \notin E(C)$ and $1_{F \in I} \neq 1_{H \in I}$ if $e \in E(C)$, which is our claim. □

Lemma 4.7 implies in particular that $B_C(P)$ is characterised by C alone and is therefore independent of P, since $\overline{G}$ was defined without reference to P. Thus we can denote it by just B_C from now on. In the following, we use again the definition of a crossing from Section 2.3.

Lemma 4.8. *Let C' be a finite path of $\mathbb{T}$ such that $C := \pi(C')$ is a cycle of G, and let $Q = e\mathring{Q}f$ be a crossing of C' in $\mathbb{T}$. Then B_C separates the 2-cells incident with $\pi(e)$ from the 2-cells incident with $\pi(f)$. Moreover, if Q_2 is a path of $\mathbb{T}$ such that $\pi(Q_2)$ is a cycle of G, then Q_2 crosses C' an even number of times.*

Proof. Let F be a face incident with the first edge e of Q, and let H be a face incident with the last edge f of Q. By the definition of a crossing, we can find a finite sequence $(F =)F_1, \dots, F_k(= H)$ of faces of $\mathbb{T}$ such that each F_i shares an edge e_i with F_{i+1} and exactly one of the e_i lies in C': we can visit all faces incident with Q until we reach H. By Lemma 4.7 and Lemma 4.4, B_C separates $\pi(F_1)$ from $\pi(F_k)$. This proves our first assertion.

For the second assertion, note that $\pi(Q_2)$ can be written as a concatenation of subarcs $C_1 D_1 C_2 D_2 \dots C_k = C_1$ where each C_i lifts to a crossing of C' by Q_2 and each D_i avoids C and shares exactly one end-edge with each of C_i and C_{i+1}. We proved above that the 2-cells incident with end-edges of each C_i are separated by B_C. The same arguments imply that the 2-cells incident with end-edges of each D_i are *not* separated by B_C. Since $\pi(Q_2)$ is a cycle, this implies that Q_2 crosses C' an even number of times. □

We can now prove that all pairs of identified points of $\overline{D}$ are nested, completing the proof of Theorem 4.2 started at the beginning of Section 4. Suppose, to the contrary, there are two pairs x, x' and y, y' that are not nested. Let X be the x–x' path in $\overline{D}$, and let Y be the y–y' path. Then X crosses Y exactly once since $\overline{D}$ is a tree, contradicting the last statement of Lemma 4.8 because $\pi(X), \pi(Y)$ are cycles of G by the definition of $\overline{D}$. Thus all such pairs are nested, and our proof is complete.

5 Planar presentations yield planar Cayley graphs: the general case

In this section we extend our definitions and proofs from Section 4 in order to be able to capture planar Cayley graphs without a consistent embedding. As mentioned above, such Cayley graphs are not 3-connected. We start by extending the definition of a planar presentation in Section 5.1, and proceed to prove that such presentations give rise to (not necessarily 3-connected) planar Cayley graphs in Section 5.2.

5.1 Planar presentations — the general case

We now extend the definition of a planar presentation, to capture Cayley graphs with 2-separators that do not admit consistent embeddings.

Let again $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ be a group presentation, and define $\mathcal{S}'$ as above.

A *spin structure* $\mathcal{C}$ on $\mathcal{P}$ consists of a cover $B_1, \ldots, B_k$ of $\mathcal{S}'$ (i.e. $\bigcup_i B_i = \mathcal{S}'$) with the following properties

(S1) for every generator b, the number of B_i's containing b equals the number of B_i's containing b^{-1}, and

(S2) the auxiliary graph X on $\mathcal{C} \cup \mathcal{S}'$ with $s \sim B_i$ whenever $s \in B_i$, is a tree.

The *hinges* of this spin structure are the elements of $\mathcal{S}'$ that have degree at least 2 in X; in other words, $h \in \mathcal{S}'$ is a hinge if $h \in B_i \cap B_j$ for some $i \neq j$. Hinges of a spin structure correspond to edges of our Cayley graph G whose two endvertices separate G.

For example, a, b are the hinges of the presentation

$$\left\langle a, b, c, d, f, g \mid a^2, c^2, d^2, f^2, g^2, (af)^2, (ag)^2, abab^{-1}gbfb^{-1}, cbdb^{-1} \right\rangle$$

given in Figure 3, and b is the only hinge in Figure 2. The tree X of condition (S2) corresponding to the presentation of Figure 3 is shown in Figure 8. Figure 9 shows the corresponding tree X that would result if we amalgamated the above group with two more groups each of which being isomorphic to the subgroup generated by b, c, d along the subgroup spanned by b.

Condition (S2) has the following important consequences:

$$B_i \cap B_j \text{ is either empty or a singleton for every } i \neq j, \tag{6}$$

because if $h, g \in B_i \cap B_j$ then h, g, B_i, B_j span a 4-cycle in X, which cannot happen when X is a tree, and

$$\text{every } B_i \text{ contains at least one hinge unless } k = 1, \text{ i.e. } \mathcal{C} \text{ is the singleton } \{\mathcal{S}'\}, \tag{7}$$

because if each neighbour of B_i in X has degree 1, then B_i and its neighbours form a component of X.

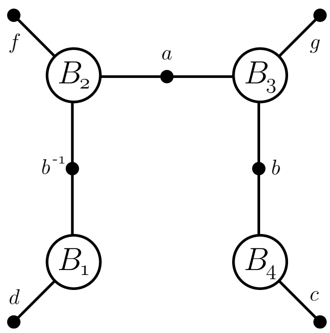

Figure 8: An example: the tree X of condition (S2) corresponding to Figure 3.

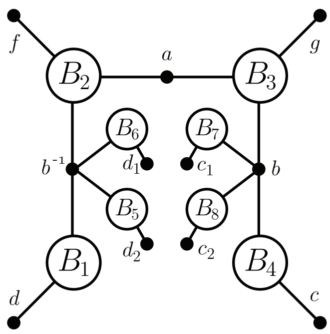

Figure 9: The tree X of condition (S2) corresponding to a variant of Figure 3.

A *general embedded presentation* is a quintuple $\mathcal{P}, \mathcal{C}, \sigma, \tau, \mu$ as follows; $\mathcal{P}$ is a group presentation and $\mathcal{C}$ a spin structure on $\mathcal{P}$ as above; σ is a function of $i \in \{1, \ldots, k\}$ assigning a spin (i.e. a cyclic ordering) to each $B_i \in \mathcal{C}$; $\tau \colon \mathcal{S} \times \{1, \ldots, k\} \to \{0, 1\}$ encodes the information of whether each generator is spin-preserving or spin-reversing in each B_i it participates in (if $s \in \mathcal{S} \setminus B_i$, then the value of $\tau(s, i)$ will be irrelevant in the sequel); and for every $b \in \mathcal{S}$, and every i for which $b \in B_i$, $\mu(b, i)$ is a B_j such that $b^{-1} \in B_j$, and $\mu(b, i) \neq \mu(b, m)$ for $m \neq i$. This μ encodes the information of which pairs of B_i incident with the two endvertices of a given hinge belong to the same block of G. The use of $\mathcal{S}$ rather than $\mathcal{S}'$ in the definition of μ and τ is intended: the values we assign to each $b \in S$ give us enough information about how to treat b^{-1}.

For the time being, the data σ, τ, μ are abstract objects describing the intended structure and embedding of our Cayley graph given by $\mathcal{P}$. But we will, similarly to what we did in Section 4, indeed prove that if these data satisfy certain conditions, then the Cayley graph is indeed planar and can be embedded in the intended way.

As an example, the presentation $\left\langle \mathcal{S} \mid b^2, a^3, c^3, aba^{-1}b, cbcb \right\rangle$ of the graph of Figure 2 can be endowed with the following data. The spin structure $\mathcal{C}$ consists of two sets $B_1 = \{b, c, c^{-1}\}, B_2 = \{b, a, a^{-1}\}$. We can then let $\sigma(1) = (b, c, c^{-1})$, $\sigma(2) = (b, a^{-1}, a)$ —but any other σ would do in this case as there are only two cyclic orderings of a set of three elements, and they are the reflection of each other— $\tau(b, 1) = 0$, $\tau(b, 2) = 1$ —this is the most interesting aspect of this graph: any b edge is spin-preserving in one of its incident blocks and spin-reversing in the other— and $\mu(b, 1) = B_1$, $\mu(b, 2) = B_2$ —because b stabilises the two components into which it splits the graph.

Our general definition of a planar presentation will be very similar to that of Section 3.1, and still based on the idea of non-crossing relators. One difference is that we have to embed the tree $\mathbb{T} = Cay \left\langle \mathcal{S} \mid s^2, s \in \mathcal{I} \right\rangle$ in $\mathbb{R}^2$ more carefully: rather than demanding every vertex to have the same cyclic ordering of its incident colours in the embedding, which would in general make it impossible to adhere to the spin-behaviour encoded by τ, we embed $\mathbb{T}$ (accumulation-free) in $\mathbb{R}^2$ in such a way that the following two conditions are satisfied. Given a vertex $x \in V(\mathbb{T})$ and $B_i \in \mathcal{C}$, we write $B_i(x)$ for the edges of x with labels in B_i.

(B1) σ is respected, i.e. for every vertex $x \in V(\mathbb{T})$, and every $B_i \in \mathcal{C}$, the cyclic ordering induced on $B_i(x)$ by our embedding coincides with $\sigma(i)$ up to reflection. Moreover, the edges of $B_i(x)$ are consecutive in our embedding.

(B2) τ is respected, i.e. for every edge $e = vw$ of $\mathbb{T}$, and every i such that the label s of e is in $B_i \in \mathcal{C}$, we have $1_{\sigma(i)}(B_i(v)) = 1_{\sigma(j)}(B_j(w))$ if and only if $\tau(s, i) = 0$, where $B_j = \mu(s, i)$ and $1_{\sigma(i)}(B_i(v))$ is 1 if the clockwise cyclic ordering of the colours of the edges of $B_i(v)$ coincides with $\sigma(i)$ and 0 otherwise.

We repeat the definition of *crossing* from Section 3.1 verbatim: given a word W, we let W^∞ be the 2-way infinite word obtained by concatenating infinitely many copies of W. We say that two words $W, Z \in \mathcal{R}$ *cross*, if there is a 2-way infinite path R of $\mathbb{T}$ induced by W^∞ and a 2-way infinite path L induced by Z^∞ such that L meets both components of $\mathbb{R}^2 \setminus R$.

The second and final difference of our generalised definition of a planar presentation compared to that of Section 3.1 will be an additional condition reflecting the idea that in a planar Cayley graph of connectivity 2, we can choose the relators in such a way that each cycle they induce is contained in a block. Recalling that our spin structure $\mathcal{C}$ is intended to capture the decomposition into blocks, the following definition should not be too surprising.

We say that a relator R is *blocked* with respect to $\mathcal{C}$, if it satisfies the following two properties. Firstly, for every two (possibly equal) consecutive letters st appearing in R^∞ or $(R^{-1})^\infty$, there is some $B_i \in \mathcal{C}$ containing both s^{-1}, t. Secondly, for every three consecutive letters sbt, where b is a hinge, appearing in R^∞ or $(R^{-1})^\infty$, if B_i is the unique element of $\mathcal{C}$ containing s^{-1}, b, then $\mu(b, i)$ contains both b^{-1}, t, unless $s = b = t$ and $b^2 \in \mathcal{R}$; here, the existence of such a B_i is guaranteed by the previous requirement, and its uniqueness is a consequence of (6) in the definition of a spin structure.

Definition 5.1. A general *planar presentation* is a general embedded presentation such that

(P1) every relator in $\mathcal{R}$ is blocked with respect to $\mathcal{C}$;

(P2) no two relators $W, Z \in \mathcal{R}$ cross;

(P3) for every relator R, the number of occurrences of letters t in R with $\tau(t, i) = 1$ (i.e. spin-reversing letters), where i is the unique value for which $s^{-1}, t \in B_i$ for the letter s preceding t in R, is even[4]; here, the symbol s^n counts as $|n|$ occurrences of s;

(P4) no relator is a sub-word of a rotation of another relator.

(We could try to omit (P4) by generalising Lemma 3.2.)

Note that a planar presentation as defined in Section 3.1 is a special case of a general one when $\mathcal{C}$ consists of a single set coinciding with $\mathcal{S}'$.

5.2 Proof of planarity: the general case

For a hinge $h \in \mathcal{S}$, we let $\mathcal{C}(h) := \{B_i \in \mathcal{C} \mid h \in B_i\}$ and let $N(h)$ be the cardinality $|\mathcal{C}(h)|$. Note that $|\mathcal{C}(h)| = deg_X(h)$, where the tree X is as in (S2) of the definition of a spin structure.

Every hinge $b = xy \in E(\mathbb{T})$ of $\mathbb{T}$ labelled h naturally splits $\mathbb{T}$ into $N(h)$ subtrees: each of these subtrees contains b, it contains all edges of x with labels in a component of $X - h$ containing some $B_i \in \mathcal{C}(h)$ and no other edges of x, and it contains those edges of y with labels in the component of $X - h^{-1}$ containing $\mu(h, i)$ and no other edges of y; moreover, each such subtree is maximal

[4] The existence and uniqueness of this B_i is a consequence of (P1); see the definition of 'blocked'.

with these properties. Let $Sep_b = \{T_1, T_2, \dots, T_{N(h)}\}$ denote the set of those subtrees, and note that $\bigcap Sep_b = \{b\}$.

Definition 5.2. A *pre-block* of $\mathbb{T}$ is a maximal subtree $A \subseteq \mathbb{T}$ not separated by any Sep_b; that is, for every hinge b of $\mathbb{T}$, A is contained in some element of Sep_b.

Alternatively, we can define a pre-block as a maximal subtree of $\mathbb{T}$ such that for every $x, y \in V(A)$, if we let $s_1 s_2 \dots s_k$ denote the word (with letters in $\mathcal{S}$) read along the x–y path, then s_{j-1}^{-1}, s_j lie in a common element of $\mathcal{C}$ for every $j > 1$, and whenever s_j is a hinge, and $s_{j-1}^{-1}, s_j \in B_i \in \mathcal{C}$, then $s_j^{-1}, s_{j+1} \in \mu(s_j, i)$.

5.3 The embedding ρ of $\mathbb{T}$

Recall that our proof of Theorem 4.2 starts with an embedding of the corresponding tree $\mathbb{T}$ respecting the spin data. In our new setup of a general embedded presentation our spin data give us some restrictions but do not uniquely determine an embedding of $\mathbb{T}$, and in fact we have to be careful with our choices in order for the proof in subsection 5.4 to work.

Recall that our general embedded presentation consists of the data $\mathcal{P}, \mathcal{C}, \sigma, \tau, \mu$. For $B \in \mathcal{C}$ and a vertex $x \in V(\mathbb{T})$, recall that $B_i(x)$ denotes the edges going out of x whose labels are in B. We claim that there is an embedding $\rho : \mathbb{T} \to \mathbb{R}^2$ satisfying all of the following (the first two were also used in the definition of crossing relators in Section 5.1).

(ρ1) σ is respected, i.e. for every vertex $x \in V(\mathbb{T})$, and every $B_i \in \mathcal{C}$, the cyclic ordering induced on $B_i(x)$ by ρ coincides with $\sigma(i)$ up to reflection. Moreover, the edges of $B_i(x)$ are consecutive in the spin of x induced by σ.

(ρ2) τ is respected, i.e. for every edge $e = vw$ of $\mathbb{T}$, and every i such that the label s of e is in $B_i \in \mathcal{C}$, we have $1_{\sigma(i)}(B_i(v)) = 1_{\sigma(j)}(B_j(w))$ if and only if $\tau(s, i) = 0$, where $B_j = \mu(s, i)$ and $1_{\sigma(i)}(B_i(v))$ is 1 if the clockwise cyclic ordering of the colours of the edges of $B_i(v)$ coincides with $\sigma(i)$ and 0 otherwise.

(ρ3) μ is respected: let $b \in E(\mathbb{T})$ be a hinge, and U, W two paths containing b contained in distinct pre-blocks containing b. Then U, W do not cross each other (at b).

(ρ4) If x, y belong to the same $N(\mathcal{R})$-orbit (where $N(\mathcal{R})$ is the normal subgroup generated by $\mathcal{R}$ as in Section 2.1), and b is a hinge at x with label in $h \in \mathcal{I}$, and $h \neq 1$, then the *local spin* at x with respect to b coincides up to reflection with the local spin at y with respect to the corresponding hinge labelled h.

Here, the *local spin* with respect to a generator $h \in \mathcal{S}'$ at a vertex x is the cyclic ordering on $N_X(h)$ induced by the embedding, where X denotes the tree from Section 5.1.

If G is a planar Cayley graph, then the results of Section 6.2 imply that if we embed the universal cover $\mathbb{T}$ of G into $\mathbb{R}^2$ in a way that locally imitates an embedding of G, then all above properties are satisfied.

Properties (ρ1) to (ρ3) are not hard to satisfy: we can embed $\mathbb{T}$ by starting with the star $E(o)$ and then recursively attaching the star $E(v)$ of a new vertex to the subtree embedded so far, and it is always possible to embed $E(v)$ without violating any of (ρ1)–(ρ3). In fact we could have several ways to extend the current embedding to $E(v)$, arising by 'permuting' those $B_i(v), 1 \leq i \leq k$ that do not contain the edge of v embedded before, and by 'reflecting' any such $B_i(v)$. These choices are in direct analogy to the flexibility we have in the embedding of any planar Cayley graph of connectivity 2: permuting the $B_i(v)$ corresponds to 'activating' a hinge b incident with v to exchange the order in which blocks separated by b are embedded. Reflecting a $B_i(v)$ corresponds to flipping such a block around.

These choices mean that (ρ4) will be violated unless we make them carefully. To achieve this, recall from (S2) of Section 5.1 that the auxiliary graph X on $\mathcal{C} \cup \mathcal{S}'$ with $s \sim B_i$ whenever $s \in B_i$, is a tree. Let X^ℓ denote the tree obtained from X by attaching to each vertex v in $\mathcal{S}' \subset V(X)$ a new leaf, which leaf we denote by $\ell(v)$.

Fix an embedding $\chi : X^\ell \to \mathbb{R}^2$ of that tree with the following two properties. Firstly, the spin of any vertex $B \in \mathcal{C}$ of X^ℓ coincides with $\sigma(B)$ up to reflection.

Recall that $N(v) = N_G(v)$ denotes the neighbourhood of v in a graph G. For every hinge $h \in \mathcal{S} \setminus \mathcal{I}$, note that $\mu(h, \cdot)$ defines a bijection between $N_X(h)$ and $N_X(h^{-1})$ by the definition of μ. We extend that bijection to $N_{X^\ell}(h)$ and $N_{X^\ell}(h^{-1})$ by mapping $\ell(h)$ to $\ell(h^{-1})$. The second property we impose on χ is that the spin it induces on $N_{X^\ell}(h)$ coincides up to reflection with the μ-image of that spin induced by χ on $N_{X^\ell}(h^{-1})$, and this holds for every such h.

For an involution hinge $h \in \mathcal{I}$, $\mu(h, \cdot)$ still defines a bijection between $N_X(h)$ and $N_X(h^{-1}) = N_X(h)$, and we do not impose any requirement on χ as we did for $h \in \mathcal{S} \setminus \mathcal{I}$. Instead, we let χ embed $N_{X^\ell}(h)$ with an arbitrary spin $\phi = \phi(h)$, and define

Definition 5.3. The *dual spin* of ϕ is the cyclic ordering on $N_{X^\ell}(h)$ obtained by composing ϕ with $\mu(h, \cdot)$.

To satisfy (ρ4), we will construct ρ in such a way that the local spin with respect to h at every vertex in a given $N(\mathcal{R})$-orbit either always coincides with ϕ or it always coincides with the dual of ϕ. We remark that we cannot construct ρ algorithmically since we cannot predict which vertices of $\mathbb{T}$ are in the same $N(\mathcal{R})$-orbit; we can only prove the existence of such a ρ abstractly.

We think of this χ as providing instructions about how to construct ρ. As an example, if the set $\mathcal{I}$ of involutions in $\mathcal{S}$ is empty, then every vertex of $\mathbb{T}$ will have the same spin up to reflection in ρ, and that spin can be read from χ by contracting all non-leaves of X^ℓ into a single vertex; that vertex has the right spin in the resulting star.

Let $o = x_1, x_2, \ldots$ be an enumeration of $V(\mathbb{T})$ such that $\{x_1, \ldots, x_k\}$ spans a connected subgraph for all k. We will construct ρ by embedding the x_i one at a time as indicated above. To begin with, we embed one edge e_0 incident with $x_1 = o$ in the 0th step. From now on, each step i begins with some vertices being embedded fully, i.e. with all incident edges, and some vertices having exactly one of their edges embedded in the current embedding ρ_{i-1} of some subtree of $\mathbb{T}$. Let j be the smallest index such that x_j has exactly one of its edges e_i embedded in ρ_{i-1}. We may assume without loss of generality that $j = i$ by changing our enumeration.

We extend ρ_{i-1} to ρ_i by embedding the remaining edges incident with x_i. This will be done by the performing the following recursive procedure on X^ℓ to obtain an embedded star S_i with its edges labelled by $\mathcal{S}'$, and then embedding $N_{\mathbb{T}}(x_i)$ with the same spin as S_i.

To begin with, let ℓ be the unique leaf of X^ℓ such that $\ell = \ell(s)$ for the label $s \in \mathcal{S}$ of the edge e_i considered as outgoing from x_i. We distinguish the following cases.

Case 1: If $s \notin \mathcal{I}$, and s is a hinge, then we embed the star $N(s)$ of s in X^ℓ into $\mathbb{R}^2$ so that the spin of s in this embedding coincides with the spin of s in χ up to reflection; there are exactly two possibilities for this —because of reflection— and we choose the unique one guaranteeing (ρ3): unless we are in step $i = 1$, in which case we just embed $N(s)$ with the spin of s in χ without reflection, the other endvertex x of e_i has already been fully embedded, and the local spin with respect to e_i (which now label s^{-1} as seen from x) at x coincides up to reflection with that induced on $N(s^{-1})$ by χ by induction hypothesis. We use the possibility to reflect or not in order to guarantee that the clockwise ordering of the B_i in $N(s)$ coincides with the counterclockwise ordering of the $\mu(s, i)$ induced by the spin of x in the embedding ρ_{i-1}.

Case 2: If $s \notin \mathcal{I}$, and s is not a hinge, then it has exactly two neighbours in $N(s)$ ($\ell(s)$ and the unique $B \in \mathcal{C}$ containing s), and so reflection does not change the spin; we just embed $N(s)$ in the unique possible way.

Case 3: If $s \in \mathcal{I}$, and s is not a hinge, then again we just embed $N(s)$ in the unique possible way.

Case 4: Finally, if $s \in \mathcal{I}$, and s is a hinge, then we follow a similar approach to the $s \notin \mathcal{I}$ case, except that we now do not insist that the spin of s in the embedding of $N(s)$ we produce coincides with the spin of s in χ up to reflection; we just make sure that (ρ3) is satisfied, by embedding $N(s)$ so that the clockwise ordering of the B_i in $N(s)$ coincides with the counterclockwise ordering of the $\mu(s, i)$ induced by the spin of x in the embedding ρ_{i-1}; again this is well-defined unless we are in step $i = 1$, in which case we just embed $N(s)$ with the spin of s in χ.

Once $N(s)$ is embedded as above, we set $X^\ell_0 := N(s)$ and proceed by the following recursive procedure, which produces embeddings of an increasing sequence $X^\ell_1, \ldots, X^\ell_k (= X^\ell)$ of subtrees of X^ℓ to embed the rest of X^ℓ.

For $j = 1, 2, \ldots$, pick a leaf v_j of X^ℓ_{j-1} which is not a leaf of X^ℓ; if no such leaf exists then $X^\ell_{j-1} = X^\ell$ and we stop. Then we extend the current embedding of X^ℓ_{j-1} by embedding $N(v_j)$ in such a way that the spin of v_j coincides up to reflection with that induced by χ, unless $v_j \in \mathcal{I} \subseteq \mathcal{S}$ and $v_j \neq 1$, in which case we do the following. Let $y_i = x_i v_j$ be the vertex of $\mathbb{T}$ joined to x_i by the edge labelled v_j. If no vertex of $\mathbb{T}$ from the $N(\mathcal{R})$-orbit of x_i or y_i has been embedded yet by ρ_i, then we embed $N(v_j)$ with local spin given by χ. If some vertex of $\mathbb{T}$ from the $N(\mathcal{R})$-orbit of x_i has already been embedded by ρ_i, we embed $N(v_j)$ with same spin up to reflection as we used so far for all $x_j, j < i$, that are $N(\mathcal{R})$-equivalent to x_i; (we make this choice in order to satisfy (ρ4)). Otherwise, we embed $N(v_j)$ with the dual spin —recall Definition 5.3— up to reflection of the spin we used so far for all $x_j, j < i$ that are $N(\mathcal{R})$-equivalent to y_i. Note that these choices ensure that $N(v_j)$ is embedded with the same spin up to reflection —namely, either that induced by χ or its dual— for all vertices in an $N(\mathcal{R})$-orbit, where we use the fact that, as $v_j \neq 1$, x_i and y_i are never in the same orbit.

In all cases, we still have the option of reflecting. If $v_j \in N(s)$, which means that $v_j \in \mathcal{C}$ and v_j contains the label s of e_i, then we have to worry about satisfying (ρ2); but one of the two choices we have due to the option of reflecting will satisfy (ρ2) for $e = e_i$ and $B_i = v_j$ and we make that choice. (If $v_j \notin N(s)$ then we do not worry about μ and τ; the other endvertices of the edges incident with x_i will make sure that this data is respected, just as we were careful above when embedding $N(s)$ for the label s of e_i.)

Let $X^\ell_j := X^\ell_{j-1} \cup N(v_j)$.

The procedure finishes when all of X^ℓ has been embedded. Then, we contract all non-leafs of X^ℓ to obtain the desired embedded star S_i out of that embedding. Finally, we embed $N_{\mathbb{T}}(x_i)$ with the same spin as S_i to extend ρ_{i-1} to ρ_i.

Let $\rho = \bigcup \rho_i$ be the limit of the ρ_i. We claim that ρ satisfies conditions (ρ1)–(ρ4). Indeed, if any of them is violated, then there is a first step in the above procedures violating it. But we designed all steps so that none of those conditions are violated: condition (ρ1) is never violated because we chose χ so that the spin of every $B_i \in \mathcal{C}$ coincides with $\sigma(i)$ up to reflection, which implies that the corresponding edges of x_i appear in that cyclic order up to reflection in S_i, and therefore in ρ, by the construction of the embedded star S_i. Condition (ρ2) is never violated because of the way we embedded $N(v_j)$ for $v_j \in N(s)$ in the construction of S_i. Condition (ρ3) is never violated because of the way we embedded $N(s)$ in the first step of the construction of S_i. Finally, condition (ρ4) is never violated because of the way we embedded $N(v_j)$ for $v_j \in \mathcal{I}$ in the construction of S_i.

In fact, we obtain a slightly stronger property than (ρ4), and this will be useful later:

Condition (ρ4) remains true if we define *local spin* using X^ℓ instead of X. (8)

5.4 Planarity of blocks

A *block* of G is an image $\pi([A])$ under the covering map π, where A denotes a pre-block of $\mathbb{T}$ and $[A] := \{x \in V(\mathbb{T}) \mid x \simeq_N y \text{ for some } y \in A\}$ denotes its $N(\mathcal{R})$-equivalence class.

Note that every block of G is connected: given vertices x, z in a block $K = \pi([A])$, we can find $x', z' \in A$ (and not just in the $N(\mathcal{R})$-orbit of A) with $\pi(x') = x, \pi(z') = z$, and so the x'–z' path P in A yields the x–z path $\pi(P)$ in K.

Lemma 5.4. *Every block of G is planar.*

In fact, we will prove a stronger statement similar to Theorem 4.2, namely, that every block admits an embedding into $\mathbb{R}^2$ respecting σ and τ.

The proof of this follows the lines of our proof of the planarity of G in the consistent case (Theorem 4.2), and we assume that the reader has already understood that proof. Here we will point out the differences.

Let K be a block of G. Let D be a fundamental domain of K in $\mathbb{T}$; that is, D is a subset of $\mathbb{T}$ containing exactly one point from each $N(\mathcal{R})$-orbit O such that $\pi(O) \in K$. The proof of Lemma 4.1 can be repeated to prove that D can be chosen to be connected since K is. Moreover, we may still assume without loss of generality that D is a union of stars as we did in (2). Thus the closure $\overline{D}$ of D in $\mathbb{T}$ is still the union of D with all midpoints of edges that have exactly one half-edge in D, and K can be obtained from $\overline{D}$ by identifying pairs of $N(\mathcal{R})$-equivalent midpoints. As in the proof of Theorem 4.2, we will prove that any two pairs of such $N(\mathcal{R})$-equivalent midpoints are nested in the sense of Section 4.2.

In order to guarantee this nestedness, we will have to embed $\mathbb{T}$ appropriately; in our general setup, $\mathbb{T}$ cannot be embedded consistently as in the case of special planar presentations, and this is why we are now only trying to prove the planarity of a block, and not of all of G at once.

Adhering to the notation of Section 4.2, for a relator W, we still use W_o to denote the closed walk $o_G W$ in G induced by W at o_G, and let $\mathbb{T}_W := \pi^{-1}(W_o)$, which is a union of a set of double-rays of $\mathbb{T}$, which set we denote by $\mathbb{T}[W_o]$.

Recall that in Section 4.2 we introduced a relation $F \sim H$, meaning that any F–H path P_{FH} in $\mathbb{T}^*$ crosses $\mathbb{T}_W$ an even number of times. We will introduce a similar relation now, but we have to refine its definition due to the fact that our embedding is only consistent when restricted to a pre-block.

Recall we have chosen an embedding ρ of $\mathbb{T}$ in Section 5.3. For a pre-block C of $\mathbb{T}$, we define a *super-face* of C to be a face of the embedding $\sigma(C)$ of C inherited by ρ. The super-faces of $\mathbb{T}$ are the super-faces of all of the pre-blocks of $\mathbb{T}$. Note that a super-face can contain several faces of $\mathbb{T}$.

We will define our relation $\sim_K$, or just $\sim$ if K is fixed, on the set of super-faces of pre-clusters contained in $\pi^{-1}(K)$. Given two super-faces F, H lying in pre-clusters contained in $\pi^{-1}(K)$, let $\mathbb{T}[W_o]_K$ denote the subset of $\mathbb{T}[W_o]$ contained in $\pi^{-1}(K)$. Now pick two faces $F' \subseteq F, H' \subseteq H$ contained in the super-faces F, H, and write $F \sim H$ if for each F'–H' path $P_{F'H'}$ in $\mathbb{T}^*$, the number of crossings $Cr(\mathbb{T}[W_o]_K, P_{F'H'})$ of $\mathbb{T}[W_o]_K$ by $P_{F'H'}$ (as defined in Section 4.3.1) is even. Note that $Cr(\mathbb{T}[W_o]_K, P_{F'H'})$ is independent of the choice of $P_{F'H'}$ by the same arguments we used to prove (4). Therefore, it is also independent of the choice of F', H', because if F'' is another face contained in F, then the F'–F'' path of $\mathbb{T}^*$ contained inside F crosses no element of $\mathbb{T}[W_o]_K$, because a super-face of any pre-cluster C in $\pi^{-1}(K)$ meets no element of $\mathbb{T}[W_o]_K$ by the definitions.

5.4.1 The bipartitions $\{I, O\}$

An important part of our planarity proof in the consistent case was that $\sim$ was invariant under the action of $N(\mathcal{R})$ (Lemma 4.4). Below (Lemma 5.6) we prove an analogous statement for the general case, namely that the restriction of $\sim$ to the super-faces of the pre-blocks in $\pi^{-1}(K)$ is $N(\mathcal{R})$-invariant.

The rest of our proof is almost identical to that of Theorem 4.2, except that we are now working with the block K of G rather than the whole graph.

As in Section 4.3.1, the equivalence relation $\sim$, now restricted on the set of super-faces $\mathcal{F}$ of $\pi^{-1}(K)$, uniquely determines a bipartition $\{I, O\}$ on $\mathcal{F}$ by choosing one super-face $F \in \mathcal{F}$ and letting $I := \{H \in \mathcal{F} \mid H \sim F\}$ and $O := \mathcal{F} \setminus I$.

Next, we adapt the material of Section 4.3.1 to our new setup. For every super-face F in $\pi^{-1}(K)$, glue a copy of the domain $\overline{F} \subset \mathbb{R}^2$ to K by identifying each point of ∂F with $\pi(\partial F)$. If F, F' are equivalent face boundaries, in other words, if $\pi(\partial F) = \pi(\partial F')$, then we identify the corresponding 2-cells glued onto K. Let K^2 denote the set of these 2-cells, and let $\overline{K} = K \cup K^2$ denote the 2-complex consisting of K and these 2-cells. Notice the similarity between G^2 and $\overline{G}$ as defined in Section 4.3.2 and K^2 and $\overline{K}$.

Lemma 5.6 now means that if Z is a closed walk of G (here we really mean G and not just K) induced by a relator, then $\{I, O\}$ induces a bipartition $\pi[I], \pi[O]$ of K^2. Let us still denote this bipartition of K^2 by B_Z.

We extend that bipartition to an arbitrary cycle in K just like in Section 4.3.1: given a cycle C of K, we choose a 'proof' P of C; that is, a sequence of closed walks $W_i, 1 \leq i \leq k$ of G induced by rotations of relators such that $C = \sum_{1 \leq i \leq k} W_i$. The existence of such a sequence (W_i) is not affected by the fact that we are focusing on a subgraph K; the W_i are allowed to be arbitrary relators. For every W_i, let I_{W_i}, O_{W_i} denote the two sides of the bipartition B_{W_i} of K^2 from above, and define the bipartition $B_C := \{I_C, O_C\}$ of K^2 by $I_C := \triangle_i I_{W_i}$ and $O_C := G^2 \triangle I_C$.

Lemmas 4.6, 4.7 and 4.8 remain true and can be proved with the same arguments, except that we replace G by K everywhere. As in the finish of the proof of Theorem 4.2, the last lemma says that any two cycles of K cross each other an even number of times, and therefore any two pairs of identified points of $\overline{D}$ are nested.

This completes the proof of Lemma 5.4, except that we still have to prove the two lemmas we used above:

Lemma 5.5. *For $b \in \mathcal{I}$ with $b = 1$, and any relator W in $\mathcal{R}$, the number of elements of $\mathbb{T}[W_o]$ containing any edge e labelled by b is even.*

Proof. Let T be an element of $\mathbb{T}[W_o]$ containing e. The automorphism β of $\mathbb{T}$ exchanging the two endvertices of e maps T to an element T' of $\mathbb{T}[W_o]$ because $b = 1$ and so the two end-vertices of e are $N(\mathcal{R})$-equivalent. Note that $T \neq T'$ even if T, T' contain the same vertices, because they have opposite directions (remember that double-rays are directed by definition). Note that $\beta(T') = T$. Therefore, β establishes a bijection without fixed points on the elements of $\mathbb{T}[W_o]$ containing e, which means that the number of those elements is even. □

Lemma 5.6. *For every block K of G, the restriction of $\sim_K$ to the super-faces of $\pi^{-1}(K)$ is invariant under the action of $N(\mathcal{R})$ on $\mathbb{T}$.*

Proof. We will adapt the proof of Lemma 4.4. Since K is fixed, let us just write $\sim$ instead of $\sim_K$.

We need to prove that if F, H are super-faces of $\pi^{-1}(K)$ in the same orbit of $N(\mathcal{R})$, then $F \sim H$. Again, we may assume that there are vertices x, y in the boundaries of F, H respectively, such that $y = xwRw^{-1}$ for some word w and some relator $R \in \mathcal{R}$: by the definition of the normal closure $N(\mathcal{R})$, if we can prove $F \sim H$ in this case, we can prove $F \sim H$ for every two F, H in the same orbit of $N(\mathcal{R})$.

Let α_{FH} be the automorphism of $\mathbb{T}$ mapping x to y.

Decompose the path $Q := xwRw^{-1}$ into (inclusion-)maximal subpaths contained in a pre-block. Then we can write

$$Q = P_1 \cup P_2 \cup \ldots \cup P_k (= P'_k) \cup P'_{k-1} \cup \ldots \cup P'_1,$$

where the P_i, P'_i are those maximal subpaths, P'_i is $N(\mathcal{R})$-equivalent to P_i for every $i < k$, and P_k contains the subpath of Q induced by R (such a P_k exists because every relator R is blocked). Note that the intersection of any two subsequent P_i or P'_i is either a hinge separating the corresponding pre-blocks, or a single vertex incident with such a hinge.

Since we are free to choose any F–H walk P_{FH} in $\mathbb{T}^*$ to decide whether $F \sim H$, we will choose a convenient one, which we construct now.

Recall that every $P_i, i > 1$ starts and ends at hinges, which we will call h_{i-1}, h_i, separating its pre-block from the pre-blocks containing P_{i-1}, P_{i+1} respectively; here h_{i-1}, h_i may or may not be contained in P_i as end-edges.

Let C_i be the pre-block containing P_i and let C'_i be the pre-block containing P'_i.

Let $\Pi_i, k > i > 1$, be an (inclusion-)minimal path in $\mathbb{T}^*$ joining a super-face incident with h_{i-1} to a super-face incident with h_i —where we say that a super-face F is *incident* with an edge if the boundary of F contains that edge— such that all vertices of Π_i are faces sharing a vertex with P_i, and Π_i does not intersect P_i (at a midpoint of any edge); see Figure 10. Define Π'_i similarly using P'_i instead of P_i. Note that there are exactly two such paths Π_i to choose from, one on either side of P_i; it doesn't matter much which of the two we will choose, but let us make 'the same' choice for both Π_i and Π'_i; more precisely, we ensure that

$$\Pi_i \text{ crosses an edge } e \text{ of } C_i \text{ (incident with } P_i) \text{ if and only if } \Pi'_i \text{ crosses the edge } \alpha_{FH}(e) \text{ of } C'_i. \tag{9}$$

This is possible because ρ embeds C_i the same way as C'_i up to reflection, and Π_i is uniquely determined once we choose which of the two super-faces of C_i incident with h_i we want it to contain; by choosing Π'_i to contain the corresponding super-face incident with h'_i, our claim is satisfied. Note that Π_i does not cross h_i, because if it did we could shorten it.

For $i = 1$ we let Π_1 be a minimal path in $\mathbb{T}^*$ joining F to a super-face incident with h_1, and otherwise be defined similarly to $\Pi_i, k > i > 1$. Define Π'_1 similarly. Finally, let $\Pi_k = \Pi'_k$ be a minimal path in $\mathbb{T}^*$ joining a super-face incident with h_{k-1} to a super-face incident with $\alpha_{FH}(h_{k-1})$ without crossing P_k.

Let $\sqcup_i, k > i \geq 1$ be a path in $\mathbb{T}^*$ joining the last vertex of Π_i to the first vertex of Π_{i+1} such that all vertices of $\sqcup_i$ are faces sharing a vertex with $P_i \cap P_{i+1}$, and define $\sqcup'_i$ similarly for Π'_i, Π'_{i+1}; there are several choices for this $\sqcup_i$, so let us make it uniquely determined: if $P_i \cap P_{i+1}$ is a single vertex, then there are two candidates, and we always choose the one crossing h_i. If $P_i \cap P_{i+1}$ is the hinge h_i, then there are up to four choices, and we choose the one that crosses h_i and is contained in the two super-faces of C_i incident with h_i and in the two super-faces of C_{i+1} incident with h_i.

It follows from the choice of $\sqcup_i$ that it behaves well with respect to elements of $\mathcal{C}$:

$$\text{If } \sqcup_i \text{ meets an edge in } B_i(v) \setminus \{h_i\} \text{ (where } B_i \in \mathcal{C}) \text{ where the vertex } v \text{ is incident with } h_i, \text{ then } \sqcup_i \text{ meets every edge of } B_i(v). \tag{10}$$

A similar but slightly stronger is true for Π_i:

$$\text{If } \Pi_i \text{ meets an edge lying inside some super-face of } C_i, \text{ then } \Pi_i \text{ visits all faces incident with } P_i \text{ inside that super-face.} \tag{11}$$

Indeed, Π_i is by definition a minimal path joining certain super-faces of C_i; therefore, it crosses any super-face either completely or at a single boundary edge.

Finally, we obtain P_{FH} by concatenating all the $\Pi_i, \sqcup_i, \Pi'_i$ and $\sqcup'_i$:

$$P_{FH} := \Pi_1 \cup \sqcup_1 \cup \Pi_2 \ldots \cup \sqcup_{k-1} \cup \Pi_k (= \Pi'_k) \cup \sqcup'_{k-1} \ldots \cup \sqcup'_1 \cup \Pi'_1.$$

We need to check that $Cr(\mathbb{T}[W_o]_K, P_{FH})$ is even. We will do so by showing that the contributions of the Π_i to $Cr(\mathbb{T}[W_o]_K, P_{FH})$ cancel with those of the Π'_i, and the contributions of the $\sqcup_i$ cancel with those of the $\sqcup'_i$.

Let T be an element of $\mathbb{T}[W_o]_K$ with odd $cr(T, P_{FH})$ —which as in Section 4.3.1 denotes the number of crossings of T by P_{FH}; only such T matter. Let $T' := \alpha_{FH}(T)$.

Let us first consider the total number of crossings of such T by the subpaths $\Pi_i, \Pi'_i, i < k$, of P_{FH}.

If T is contained in C_i, then $cr(T, \Pi_i) = cr(T', \Pi'_i)$ by (9).

If T is not contained in C_i, then Π_i crosses T an even number of times (0 or 2): this is easy to see when $T \cap P_i$ is a single vertex v by applying (11) to that vertex. The situation is slightly subtler when $T \cap P_i$ is a hinge g —no other option is possible as distinct pre-blocks intersect at an edge at most by

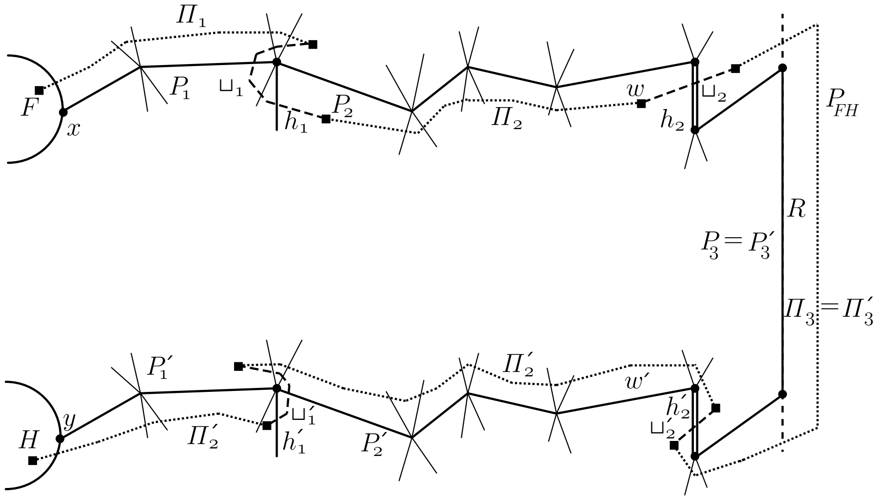

Figure 10: The path P_{FH} (dashed) in the proof of Lemma 5.6 with the paths Π_i, Π'_i, $\sqcup_i$, $\sqcup'_i$.

construction. In this case, we remark that the pre-block D containing T lies in some super-face of C_i by the construction of ρ, and again Π_i must cross all faces incident with g inside that super-face by (11), therefore crossing both edges of T incident with g.

Finally, $\Pi_k = \Pi'_k$ has an even contribution to $Cr(\mathbb{T}[W_o]_K, P_{FH})$ by the proof of Lemma 4.4.

These facts combined show that $\sum_{T \in \mathbb{T}[W_o]_K} cr(T, \bigcup_i \Pi_i)$ is even.

Next, we consider the total number of crossings of such T by the subpaths $\sqcup_i, \sqcup'_i$. Suppose $cr(T, \sqcup_i)$ is odd. Then it must equal 1 as $\sqcup_i$ is too short to cross a double-ray three times, where we used property ($\rho 3$) of our embedding ρ that pre-blocks do not cross each other.

Let v_i be the last vertex of P_i and v'_i the last vertex of P'_i. If the local spin at v_i with respect to h_i coincides up to reflection with the local spin at v'_i with respect to h'_i, then $cr(T, \sqcup_i) = cr(T', \sqcup'_i)$ (here, local spin refers to X^ℓ rather than X; recall (8)). Therefore, the total contribution of the pair T, T' to $Cr(\mathbb{T}[W_o]_K, P_{FH})$ is even and can be ignored.

If those local spins do not coincide up to reflection, then by the choice of ρ ($\rho 4$), the label of h_i is an involution $b \in \mathcal{I}$ with $b = 1$. In this case however, Lemma 5.5 applies, yielding that the set H of elements of $\mathbb{T}[W_o]_K$ containing h_i is even. We claim that $T \in H$ (i.e. $h_i \subset T$): this follows from $cr(T, \sqcup_i) = 1$, the fact that $\sqcup_i$ only contains faces of $\mathbb{T}$ incident with h_i by its construction, and (10). Moreover, (10) also implies that $cr(R, \sqcup_i) = 1$ for every other $R \in H$. But as $|H|$ is even, the total contributions $\sum_{R \in H} cr(R, \sqcup_i)$ of its elements are even and can be ignored as well.

Summing up, we proved that both

$$\sum_{T \in \mathbb{T}[W_o]_K} cr(T, \bigcup_i \Pi_i) \quad \text{and} \quad \sum_{T \in \mathbb{T}[W_o]_K} cr(T, \bigcup_i \sqcup_i)$$

are even. Therefore $Cr(\mathbb{T}[W_o]_K, P_{FH})$ is even as well, since it is the sum of those two sums by definition. □

5.5 From the planarity of blocks to the planarity of G

The main aim of this section is to prove

Lemma 5.7. *Every hinge of G separates its incident blocks.*

Proof. The statement is equivalent to the statement that every cycle of G *crosses* each hinge b an even number of times, where the number of crosses of b by C is the maximum number of edge disjoint subpaths P_i of C such that b separates each P_i into two (possibly trivial, but non-empty) subpaths that lie in distinct blocks.

To prove the latter, let $C = c_0c_1 \dots c_k$ with $c_k = c_0$ be a cycle, and let $L = t_0t_1 \dots t_k$ be a lift of C to $\mathbb{T}$ via π^{-1}. Fix a hinge b. We may assume without loss of generality that c_0 is not a vertex of b. Let $P = w_1R_1w_1^{-1} \dots w_kR_kw_k^{-1}$ be a proof of C in our presentation.

Since $c_0 \notin b$ and since the end vertices of $w_iR_iw_i^{-1}$ are $N(\mathcal{R})$-equivalent to c_0, any crossings of b by P occur inside the subpaths $w_iR_iw_i^{-1}$ and not when switching from w_{i-1} to w_i. We have no crossings of b inside any R_i because our relators are blocked. Moreover, any crossings of b inside a w_i are paired up by crossings of b inside w_i^{-1}. Thus the number of crossings of b by P, and hence by C, is even. □

This, combined with the planarity of blocks we proved in the previous section, easily implies the planarity of G:

Theorem 5.8. *Let G be the Cayley graph of a general planar presentation. Then G is planar.*

Proof. Combining Lemma 5.4 with Lemma 5.7 easily yields that G is planar. Indeed, we can embed G one block at a time: since incident blocks share a hinge only by Lemma 5.7, if we have already embedded a block A meeting a block B at a hinge b, then it is easy to embed B inside one of the two faces (we are free to choose) of the current embedding whose boundary contains b. □

6 Every planar Cayley graph admits a general planar presentation

In this section we prove the converse of Theorem 5.8, namely that every planar Cayley graph admits a general planar presentation.

We start by showing that every planar Cayley graph of connectivity 1 can be extended into a 2-connected one using redundant generators; see Lemma 6.1 below. We then show that every 2-connected planar Cayley graph admits a general planar presentation in Section 6.2. After this is achieved, we deduce that planar Cayley graphs of connectivity 1 admit general planar presentations too, by removing the redundant generators.

6.1 Planar Cayley graphs of connectivity 1

Lemma 6.1. *Every planar, locally finite, Cayley graph of connectivity 1 can be extended into a planar 2-connected, locally finite, Cayley graph by adding redundant generators.*

Proof. We proceed by induction on the number of blocks incident with the vertex o, where a *block* means a maximal 2-connected subgraph in this subsection. Pick two such blocks B, C, an edge from B corresponding to some generator b, and an edge from C corresponding to some generator c. Introduce a new redundant generator x and the relation $x = b^{-1}c$. Clearly, the resulting Cayley graph G' obtained from the original Cayley graph G by adding the generator x has less blocks incident with o than G.

We claim that G' is still planar. If none of b^2 or c^2 is a relator, then this is an easy exercise, based on the observation that G can be embedded in such a way that for every vertex v, the edges labelled b and c emanating from v lie in a common face boundary.

If however b^2, say, is a relator, then it is a bit harder to avoid that the two x edges emanating out of o and ob cross in our embedding. Still, the following observation will help us embed G' in this case (and it is also applicable to the case where none of b^2 or c^2 is a relator). A good example to bear in mind throughout the rest of the proof is where G is the Cayley graph $Cay\left\langle b, c \mid b^2, c^2 \right\rangle$ of the free product of two copies of $\mathbb{Z}/2\mathbb{Z}$, and $x = bc$.

Let H_0 be the graph consisting of a single vertex, and suppose that for every $i \in \mathbb{N}$, the graph H_i is obtained from H_{i-1} by attaching a planar graph P_i to H_{i-1} by identifying some vertex $p_i \in V(P_i)$ with some vertex $h_i \in V(H_{i-1})$, and possibly joining a neighbour p'_i of p_i to a neighbour h'_i of h_i with an edge. Then $\bigcup_{i \geq 0} H_i$ is planar. (12)

To prove this, we first use induction to show that H_i is planar: given an embedding of H_{i-1}, observe that p'_i, p_i lie in a common face F_i since they are neighbours. Likewise, h'_i, h_i lie in a common face of P_i, and we may assume that that face is the outer face by embedding P_i appropriately. We now embed H_i by drawing P_i inside F_i and, if there is a $h'_i - p'_i$ edge in H_i, joining h'_i to p'_i with an arc in F_i that avoids the rest of the graph.

The fact that $\bigcup_{i \geq 0} H_i$ is planar now follows from a standard compactness argument.

To complete our proof, we will show that our G' can be constructed as described in (12).

Indeed, let $\mathcal{H}$ be the set of blocks (i.e. maximal 2-connected subgraphs) of G, and let $H_1, H_2, \dots$ be an enumeration of $\mathcal{H}$ such that for $i > 1$, H_i is incident with some H_j for $j < i$. Then G' has the claimed structure, with the x-edges playing the role of the $h'_i - p'_i$ edges. $\square$

6.2 Cayley graphs of connectivity 2

In this section, we will complete the proof of our main theorem by showing that every locally finite 2-connected planar Cayley graphs admits a general planar presentation.

A *cut* in a graph G is a set of vertices C spanning a connected subgraph of G, such that the *boundary*

$$\partial C := \{x \in V(G) \setminus C \mid x \text{ has a neighbour in } C\}$$

of C is finite and $C \cup \partial C \neq V(G)$. The *order* of C is the cardinality of ∂C.

We call two cuts C, D *nested* if, setting $C^* := V(G) \setminus C$ and $D^* := V(G) \setminus D$, one of the four relations holds:

$$C \subseteq D, \quad C \subseteq D^*, \quad C^* \subseteq D, \quad C^* \subseteq D^*.$$

We call a set of cuts *nested*, if every two of its elements are nested.

Definition 6.2. Given a nested set $\mathcal{C}$ of cuts, a *block* is a maximal subgraph H such that for every cut C, we have either $V(H) \subseteq C \cup \partial C$ or $V(H) \subseteq C^*$ but not both.

To obtain a *torso* of a block H from H we add all edges xy such that $\{x, y\} \subseteq V(H)$ is a boundary of a cut in $\mathcal{C}$.

Tutte [28] showed that every finite 2-connected graph G has an $\mathrm{Aut}(G)$-invariant nested set $\mathcal{C}$ of cuts of order 2 whose torsos are either 3-connected or cycles. This theorem also holds for locally finite graphs, see Droms et al. [8]. Nevertheless, we will refer to it as *Tutte's theorem*. To each such nested set of cuts, there is an associated tree T that admits a bijection from $V(T)$ to the blocks and boundaries of cuts in $\mathcal{C}$ such that, for any $t_1, t_2 \in V(T)$ and any t on the unique t_1–t_2 path in T, the image of t separates the images of t_1 and t_2.[5] We call this tree T the *decomposition tree* of the set of cuts.

A 2-*separator* is the boundary of a cut of order 2. Lemma 6.3 allows us to assume that all 2-separators of G are joined by an edge, i.e. they are hinges in the sense of Section 5.1. Given two Cayley graphs G, H, we call G a *Tietze-supergraph* of H if there are presentations $\langle \mathcal{S}_G \mid \mathcal{R}_G \rangle$ of $\Gamma(G)$ and $\langle \mathcal{S}_H \mid \mathcal{R}_H \rangle$ of $\Gamma(H)$ with $G = Cay\,\langle \mathcal{S}_G \mid \mathcal{R}_G \rangle$ and $H = Cay\,\langle \mathcal{S}_H \mid \mathcal{R}_H \rangle$ and with $\mathcal{S}_G \supseteq \mathcal{S}_H$ and $\mathcal{R}_G \supseteq \mathcal{R}_H$.

[5] Readers that are familiar with tree-decompositions of graphs might notice that this just says that for every nested set of cuts, we find a tree-decomposition of the graph whose parts are the blocks and boundaries of cuts.

Lemma 6.3. *Every planar 2-connected Cayley graph G has a planar Tietze-supergraph H in which every pair of vertices that separates H is connected by an edge. In addition, the new edges are labelled by a new redundant generator. (Moreover, if G is locally finite, then so is H.)*

Proof. To begin with, pick a $\Gamma(G)$-invariant nested set $\mathcal{C}$ of cuts of order 2. This set exists due to Tutte's theorem mentioned above. For every pair of non-adjacent vertices x, y such that one component of $G - \{x, y\}$ lies in $\mathcal{C}$, we add a new redundant generator a and relation $a = x^{-1}y$. Let us show that the nestedness of $\mathcal{C}$ implies that we do not lose planarity.

Note that every 2-separator lies on the boundary of some face. So if we join x_1 and y_1 by a new edge and also want to join x_2 and y_2, then the only reason why we cannot do this is because the edge x_1y_1 separates the face on whose boundary the vertices x_2 and y_2 lie. So, originally, all four vertices x_1, x_2, y_1, y_2 are distinct and lie on a boundary C of some face F in this order (either clockwise or anticlockwise). For $i = 1, 2$, let P_i be an x_i–y_i path whose inner vertices lie in a component of $G - \{x_i, y_i\}$ that avoids x_j and y_j for $j \neq i$. As the two paths P_i lie outside of F, the path P_2 connects a vertex in the inner face of $P_1 + y_1x_1$ to one in its outer face, which is impossible due to the Jordan curve theorem. This proves that we can indeed add the aforementioned redundant generators and relations without losing planarity.

Since every vertex has only finitely many neighbours and every two of them can be separated by only finitely many 2-separators (see e.g. [27, Proposition 4.2]), the resulting Cayley graph G' is still locally finite. □

Call a graph *well-separated* if it is 2-connected and every 2-separator is joined by an edge.

Theorem 6.4. *Every planar locally finite Cayley graph G with $\kappa(G) = 2$ admits a general planar presentation.*

Proof. By Lemma 6.3, we may assume that G is well-separated. Let $\mathcal{C}$ be a $\Gamma(G)$-invariant nested set of cuts of order 2 as in Tutte's Theorem. Let $\mathcal{B}_o$ be the set of blocks (in the sense of Definition 6.2) that contain the vertex o. For $B \in \mathcal{B}_o$, let S_B be the set of those generators $s \in \mathcal{S} \cup \mathcal{S}^{-1}$ such that the edge with label s starting at o lies in B. Then $\mathcal{S} \cup \mathcal{S}^{-1}$ is covered by the set of S_B. We fix an embedding ρ of G in $\mathbb{R}^2$, and endow every S_B with the cyclic order induced by ρ at o. Let $\mathcal{B}'_o \subseteq \mathcal{B}_o$ be maximal such that no two distinct $B, B' \in \mathcal{B}_o$ are of the form $B = g(B')$ for any $g \in \Gamma(G)$. We can apply Theorem 3.4 to each $B \in \mathcal{B}'_o$ to obtain a set $\mathcal{D}_B$ of cycles that generates $\mathcal{H}_1(B)$, and is invariant under the stabiliser of B in $\Gamma(G)$. Then it is easy to see that

$$\mathcal{D} := \bigcup_{\substack{B \in \mathcal{B}'_o \\ g \in \Gamma(G)}} g(\mathcal{D}_B)$$

generates $\mathcal{H}_1(G)$. Let $\mathcal{R}_\mathcal{D}$ be the set of words corresponding to cycles in $\mathcal{D}$. Easily, $\langle \mathcal{S} \mid \mathcal{R}_\mathcal{D} \rangle$ is a presentation of $\Gamma(G)$. Once more, we use Tietze-transformations to obtain a finite subset $\mathcal{R} \subseteq \mathcal{R}_{\mathcal{D}}$ with $\langle \mathcal{S} \mid \mathcal{R}_{\mathcal{D}} \rangle = \langle \mathcal{S} \mid \mathcal{R} \rangle$, which is possible as $\Gamma(G)$ is finitely presented (Droms [7, Theorem 5.1]). To see that the set $\mathcal{C} := \{B_1, \dots, B_n\} := \{S_B \mid B \in \mathcal{B}_o\}$ is a spin structure of $\mathcal{P} := \langle \mathcal{S} \mid \mathcal{R} \rangle$, it remains to show that the graph $\mathcal{T} := (\mathcal{C} \cup \mathcal{S}', \mathcal{E})$, where $xy \in \mathcal{E}$ if and only if $x \in y$ or $y \in x$, is a tree.

Let us suppose that $\mathcal{T}$ is not a tree. Obviously, $\mathcal{T}$ is connected. So it contains some cycle $S_1 s_1 \dots S_m s_m S_1$ with $S_i \in \mathcal{C}$ and $s_i \in \mathcal{S}'$. For each $i \leq m$, let $B(S_i) \in \mathcal{B}_o$ be such that $S_i = S_{B(S_i)}$. As each element of $\mathcal{B}_o$ is a block, there is some path P_i in S_i connecting the end vertices of s_{i-1} and s_i distinct from o (with $s_0 = s_m$). The concatenation of all these paths P_i is a cycle C in G that crosses all hinges s_i precisely once as $S_i \neq S_{i+1}$ (with $S_{m+1} = S_1$). But this is not possible as each cycle, and hence also C, must lie in a unique block of G.

For $i \leq n$, let $B(i)$ be that element of $\mathcal{B}_o$ with $S_{B(i)} = B_i$. For every hinge $b \in \mathcal{S}$ incident with o and every $i \leq n$ with $b \in B_i$, let $\mu(b,i)$ be that B_j with $b(B(i)) = B(j)$. So we have $b^{-1} \in B_j$. Let $\sigma(i)$ be the spin of B_i at o. To define whether every generator is spin-preserving or spin-reversing in each element of the spin-structure (it participates in), we remember that the blocks —being either 3-connected or cycles— have a unique embedding in the plane. So for $s \in \mathcal{S}$ and $i \leq n$, we define $\tau(s,i)$ to be 0 if s is spin-preserving in $B(i)$ and 1 otherwise. (Note that τ is also defined if $s \notin B_i$.) Clearly, $(\mathcal{P}, \mathcal{C}, \sigma, \mu, \tau)$ is a general embedded presentation.

As every element of $\mathcal{D}$ lies in a unique block, every $R \in \mathcal{R}$ is blocked with respect to $\mathcal{C}$ by definition, and the number of spin-reversing generators in R is even. As $\mathcal{D}$ is nested, it is easy to check that no two relators cross. The fact that no cycle is a subgraph of any other cycle implies that no relator is a sub-word of a rotation of another relator, and hence our general embedded presentation is a general planar presentation. $\square$

With an argument similar to the proof of Corollary 3.5, we obtain:

Corollary 6.5. *Every planar well-separated Cayley graph G with $\kappa(G) = 2$ is the 1-skeleton of an almost planar Cayley complex of $\Gamma(G)$.* $\square$

6.3 Consistent embeddings lead to special planar presentations

In the previous section, we have seen that 2-connected planar Cayley graphs admit general planar presentations. However, if the Cayley graph has a consistent embedding, we obtain a bit more even for 1-connected graphs:

Theorem 6.6. *Every planar Cayley graph with a consistent embedding admits a special planar presentation.*

Proof. Let G be such a graph. First note that, by repeating the arguments of the proof of Lemma 6.3, we can join the two vertices of any 2-separator $\{x,y\}$ by a new edge whenever $xy \notin E(G)$ and $G - \{x,y\}$ has two components C with

$\partial C = \{x, y\}$, while keeping the embedding consistent. So we may assume that every maximal 2-connected subgraph of G is well-separated.

Let $\mathcal{B}$ a set of blocks of the maximal 2-connected subgraphs of G consisting of one block from each $\Gamma(G)$-orbit. As before, Theorem 3.4 gives us for each $B \in \mathcal{B}$ a nested set of cycles $\mathcal{D}_B$, invariant under the stabiliser in $\Gamma(G)$ of B, generating $\mathcal{H}_1(B)$. Let $\mathcal{R}_B$ be the set of words corresponding to the cycles of $\mathcal{D}_B$. As above, Tietze-transformations give us a finite $\mathcal{R} \subseteq \bigcup_{B \in \mathcal{B}} \mathcal{R}_B$ such that $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ is a finite presentation of $\Gamma(G)$, where $\mathcal{S}$ is the generating set of G.

If we let σ be the spin of one fixed vertex x and $\tau(s) = 0$ if the edge from x labelled s is spin-preserving and $\tau(s) = 1$ otherwise, then $(\mathcal{P}, \sigma, \tau)$ is a special planar presentation of $\Gamma(G)$. Indeed, nestedness of the cycles in $\mathcal{D}_B$ implies that the corresponding words are non-crossing, the fact that they are cycles implies that no relator is a subword of any other relator, and the embedding implies that every relator contains an even number of spin-reversing letters. □

7 Conclusions

We now put the above results together to prove the statements of the introduction.

Let $CC(\mathcal{P})$ denote the Cayley complex of a presentation $\mathcal{P}$. Call a map $\rho\colon CC(\mathcal{P}) \to \mathbb{R}^2$ *consistent* if its restriction to $Cay(\mathcal{P})$ is consistent. Call ρ *nested* if it witnesses the fact that $CC(\mathcal{P})$ is almost planar, i.e. if the images under ρ of the interiors of any two 2-cells are either disjoint, or one is contained in the other.

The following might be interesting as it exhibits a geometric property of Cayley complexes which can be decided by an algorithm.

Theorem 7.1. *There is an algorithm that given a presentation $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ decides whether $CC(\mathcal{P})$ admits a nested, consistent map into $\mathbb{R}^2$.*

Proof. We claim that $CC(\mathcal{P})$ admits a nested, consistent map into $\mathbb{R}^2$ if and only if there is a spin σ on $\mathcal{S}$ and a 'spin-behaviour' function τ from $\mathcal{S}$ to $\{0, 1\}$ such that the triple $(\mathcal{P}, \sigma, \tau)$ is a special planar presentation.

To prove the backward direction, note that if $\mathcal{P}, \sigma, \tau$ is a special planar presentation, then $Cay(\mathcal{P})$ admits a consistent embedding ρ into $\mathbb{R}^2$ by Theorem 4.2. Extend this embedding into a map ρ' from $CC(\mathcal{P})$ to $\mathbb{R}^2$ by mapping each 2-cell inside the closed curve to which ρ maps its boundary. Then ρ' is nested because no two words in $\mathcal{R}$ cross each other by the definition of a special planar presentation.

For the forward direction, given such a map $\rho\colon CC(\mathcal{P}) \to \mathbb{R}^2$, we can read the spin data σ, τ from ρ since ρ is consistent. Then $\mathcal{P}, \sigma, \tau$ is an embedded presentation. To prove that it is a special planar presentation it remains to show that no two words in $\mathcal{R}$ cross each other, which follows immediately from the nestedness of ρ. □

By using general planar presentations instead of special ones, Theorem 7.1 can be generalised to yield a further decidable property of Cayley complexes, but instead of maps into $\mathbb{R}^2$ we have to consider maps into larger spaces obtained by glueing copies of $\mathbb{R}^2$ along (possibly closed) bounded simple curves —to which we map the hinges of our Cayley graphs— in a tree like fashion. We leave the details to the interested reader.

Our results do not yet answer the following

Problem 7.2. *Is there an algorithm that given a presentation $\mathcal{P} = \langle \mathcal{S} \mid \mathcal{R} \rangle$ decides whether $CC'(\mathcal{P})$ is planar?*

In this problem $CC'(\mathcal{P})$ denotes the complex obtained from $CC(\mathcal{P})$ by removing redundant 2-cells, that is, if a set of 2-cells have the same boundary, we remove all but one of them. Some authors still call $CC'(\mathcal{P})$ the Cayley complex of $\mathcal{P}$. (In Theorem 7.1 it does not make a difference whether we consider $CC(\mathcal{P})$ or $CC'(\mathcal{P})$.)

We remark that it is not true that $CC(\mathcal{P})$ is planar if and only if $\mathcal{P}$ is a facial presentation in the sense of [14]; the presentation $\mathcal{P} = \left\langle a, b \mid a^2, b^3, ab^{-1} \right\rangle$ if facial, but its Cayley complex consists of a single vertex, two loops, a 2-cell winding twice around a loop, and a 2-cell winding three times around the other loop.

Theorem 7.3. *The Cayley graphs that admit a consistent embedding in the plane are effectively enumerable.*

Proof. We claim that any effective enumeration of the special planar presentations is an effective enumeration of those Cayley graphs. Indeed, this is the case as every Cayley graph with a consistent embedding admits a special planar presentation by Theorem 6.6 and as the Cayley graph of every special planar presentation is planar by Theorem 4.2.

In order to effectively enumerate the special planar presentation, it suffices to produce an enumeration of the embedded presentations, and output those embedded presentations that satisfy the three conditions in the definition of a special planar presentation (Definition 3.1); it is easy to see that these conditions can be checked algorithmically. ☐

Theorem 7.4. *The planar, finitely generated, Cayley graphs are effectively enumerable.*

Proof. Similarly to the proof of Theorem 7.3, we claim that any effective enumeration of the general planar presentations gives rise to an effective enumeration of the planar Cayley graphs. Indeed, according to Lemmas 6.1 and 6.3, every locally finite planar Cayley graph has a locally finite planar Tietze-supergraph obtained by adding redundant generators. This supergraph admits a general planar presentation by Theorem 3.3 and Theorem 6.4 —recall that every special planar presentation induces canonically a general planar presentation— and the Cayley graph of every general planar presentation is planar by Theorem 5.8.

Thus it suffices to start with an enumeration of the general embedded presentations, and output those that satisfy the four conditions of Definition 5.1, which can be checked algorithmically.

For each such output $G = \langle \mathcal{S} \mid \mathcal{R} \rangle$, check for every $s \in \mathcal{S}$ whether s is an *obviously redundant* generator, i.e. there is exactly one relator $W_s \in \mathcal{R}$ in which s appears, and s appears exactly once in W. For every such s found, output the presentation $G' := \langle \mathcal{S} \setminus \{s\} \mid \mathcal{R} \setminus \{W_s\} \rangle$. Then, recursively apply the same check to G', removing any obviously redundant generators of that presentations and so on. In this way we obtain the locally finite planar Cayley graphs of connectivity 1 and those of connectivity 2 that are not well-separated as mentioned earlier according to Lemmas 6.1 and 6.3. □

8 Further remarks

In Section 6 we prove that every planar Cayley graph G admits a planar presentation such that every relator induces a cycle of G (rather than an arbitrary closed walk with repetitions of vertices). It would be interesting if we could strengthen the definition of a planar presentation in such a way that this is always the case in the resulting planar Cayley graph. Some strengthening will be necessary as shown by the example $\mathcal{P} = \left\langle a, b \mid a^2, b^3, ab^{-1} \right\rangle$ from the previous section. This is a planar presentation —even stronger, every relator is facial— but it is easy to see that its group is the group of one element. Our optimism that this may be possible stems from the fact that it was possible in the cubic case [12].

If we could do this, then it would probably help to prove that the planar Cayley graphs are effectively constructible:

Conjecture 8.1. *There is an algorithm that given a general planar presentation $\mathcal{P}$, and $n \in \mathbb{N}$, outputs the ball of radius n in the Cayley graph of $\mathcal{P}$.*

This was proved in [12] in the cubic case.

A further interesting question, also asked in [12], is whether for every $n \in \mathbb{N}$ there is an upper bound $f(n)$, such that every n-regular planar Cayley graph admits a planar presentation with at most $f(n)$ relators. This would strengthen Droms' result [7, Theorem 5.1] that planar groups are finitely presented.

We conclude with a rather unrelated observation. It is known that the fundamental group of a finite graph of groups with residually finite vertex groups and finite edge groups is residually finite [26, II.2.6.12]. Combining this with Dunwoody's result mentioned in the introduction, we obtain the following corollary, to which this paper has no contribution

Corollary 8.2. *Every planar group is residually finite.*

References

[1] S. I. Adjan. The algorithmic unsolvability of checking certain properties of groups. *Dokl. Akad. Nauk SSSR*, 103:533–535, 1955.

[2] G. N. Arzhantseva and P.-A. Cherix. On the Cayley graph of a generic finitely presented group. *Bulletin of the Belgian Mathematical Society - Simon Stevin*, 11(4):589–601, December 2004.

[3] G. Baumslag. *Topics in Combinatorial Group Theory*. Birkhäuser Verlag, Basel, 1993.

[4] O. Bogopolski. *Introduction to Group Theory*. EMS, Zuerich, Switzerland, 2008.

[5] W. Dicks and M. J. Dunwoody. *Groups acting on graphs*. Cambridge University Press, 1989.

[6] Reinhard Diestel. *Graph Theory* (3rd edition). Springer-Verlag, 2005. Electronic edition available at: `http://www.math.uni-hamburg.de/home/diestel/books/graph.theory`.

[7] C. Droms. Infinite-ended groups with planar Cayley graphs. *J. Group Theory*, 9(4):487–496, 2006.

[8] C. Droms, B. Servatius, and H. Servatius. The structure of locally finite two-connected graphs. *Electron. J. of Comb.*, 2:research paper 17, 1995.

[9] C. Droms, B. Servatius, and H. Servatius. Connectivity and planarity of Cayley graphs. *Beitr. Algebra Geom.*, 39(2):269–282, 1998.

[10] M.J. Dunwoody. Planar graphs and covers. Preprint.

[11] M.J. Dunwoody. The accessibility of finitely presented groups. *Invent. Math.*, 81(3):449–457, 1985.

[12] A. Georgakopoulos. The planar cubic cayley graphs. Preprint 2011.

[13] A. Georgakopoulos. The planar cubic cayley graphs of connectivity 2. Preprint 2010.

[14] A. Georgakopoulos. Characterising planar Cayley graphs and Cayley complexes in terms of group presentations. *Europ. J. Comb.*, 36:282–293, 2014.

[15] M. Hamann. Generating the cycle space of planar graphs. submitted, arXiv:1411.6392.

[16] M. Hamann. Planar transitive graphs. submitted, arXiv:1511.08777.

[17] W. Imrich. On Whitney's theorem on the unique embeddability of 3-connected planar graphs. In *Recent Adv. Graph Theory, Proc. Symp. Prague 1974*, pages 303–306. 1975.

[18] Roger C. Lyndon and Paul E. Schupp. *Combinatorial Group Theory.* Springer Science & Business Media, January 2001.

[19] A.M. Macbeath. The classification of non-euclidean plane crystallographic groups. *Can. J. Math.*, 19:1192–1205, 1967.

[20] H. Maschke. The representation of finite groups, especially of rotation groups of three and four dimensional space, by Cayley's color diagrams. *Amer. J. Math*, 18:156–194, 1896.

[21] J. Meier. *Groups, Graphs and Trees.* Cambridge University Press, 2008.

[22] B. Mohar. Tree amalgamation of graphs and tessellations of the cantor sphere. *J. Combin. Theory (Series B)*, 96(5):740–753, 2006.

[23] B. Fine P. Ackermann and G. Rosenberger. On surface groups: Motivating examples in combinatorial group theory. In *Groups St Andrews 2005*, pages 96–129. Cambridge University Press, 2007.

[24] H. Poincaré. Théorie des groupes fuchsiens. *Acta Mathematica*, 1(1):1–62, 1882.

[25] M. O. Rabin. Recursive unsolvability of group theoretic problems. *Annals of Mathematics.*, 67:172–194, 1958.

[26] J.-P. Serre. Arbres, amalgames, SL_2. *Astérisque*, 2(46), 1977.

[27] C. Thomassen and W. Woess. Vertex-transitive graphs and accessibility. *J. Combin. Theory (Series B)*, 58:248–268, 1993.

[28] W.T. Tutte. *Connectivity in graphs.* Mathematical Expositions. 15. Toronto: University of Toronto Press; London: Oxford University Press. IX, 145 p. , 1966.

[29] H. Whitney. Congruent graphs and the connectivity of graphs. *American J. of Mathematics*, 54(1):150–168, 1932.

[30] H.C. Wilkie. On non-Euclidean crystallographic groups. *Math. Z.*, 91:87–102, 1965.

[31] H. Zieschang, E. Vogt, and H.-D. Coldewey. *Surfaces and planar discontinuous groups. Revised and expanded transl. from the German by J. Stillwell.* Lecture Notes in Mathematics 835. Springer-Verlag, 1980.