The planar cubic Cayley graphs

Agelos Georgakopoulos*

Technische Universität Graz
Steyrergasse 30, 8010
Graz, Austria

Abstract

We obtain a complete description of the planar cubic Cayley graphs, providing an explicit presentation and embedding for each of them. This turns out to be a rich class, comprising several infinite families. We obtain counterexamples to conjectures of Mohar, Bonnington and Watkins. Our analysis makes the involved graphs accessible to computation, corroborating a conjecture of Droms.

1 Introduction

1.1 Overview

The study of planar Cayley graphs has a tradition starting in 1896 with Maschke's characterization of the finite ones. Among the infinite planar Cayley graphs, those corresponding to a discontinuous action on the plane have received a lot of attention. Their groups are important in complex analysis, and they are closely related to the surface groups [26, Section 4.10]. These graphs and groups are now well understood due to the work of Macbeath [19], Wilkie [25], and others; see [26] for a survey. The remaining ones are harder to analyse. They have been the subject of more recent work [8, 9, 10, 12], and they are not yet completely classified. For example, we do not know if they can be effectively enumerated [8, 9].

In this paper we study those planar Cayley graphs that are *cubic*, which means that every vertex is adjacent with precisely three other vertices. It turns out that this class is restricted enough to allow for a complete description of all of its elements, while offering enough variety to allow an insight into the general planar Cayley graphs.

Our main result is

Theorem 1.1. *Let G be a planar cubic Cayley graph. Then G has precisely one of the presentations listed in Table 1. Conversely, each of these presentations, with parameters chosen in the specified domains, yields a non-trivial planar cubic Cayley graph.*

*Supported by FWF grant P-19115-N18.

Some of the entries of Table 1 yield counterexamples to a conjecture of Bonnington and Watkins [4] and Bonnington and Mohar [21]; see Section 1.3.

The presentations of Table 1 have a special structure that is related to the embedding of the corresponding Cayley graph, and yield geometric information about the corresponding Cayley complex. The ideas of this paper are used in [16] to prove that every planar Cayley graph admits such a presentation. This solves the aforementioned problem of [8, 9] asking for an effective enumeration; see Section 1.4.

Motivated by Stallings' celebrated theorem, Mohar conjectured that every planar Cayley graph with more than 1 end can be obtained from simpler ones by a glueing operation reminiscent of group amalgamation. Some of the entries of Table 1 disprove Mohar's conjecture, but we will show that the conjecture becomes true in the cubic case after a slight modification. This modified version might be true for all Cayley graphs, not just the planar ones, yielding a refinement of Stallings' theorem. See Section 1.2 for details.

In [15] I asked for a characterization of the locally finite Cayley graphs that admit a *Hamilton circle*, i.e. a homeomorphic image of S^1 in the end-compactification of the graph containing all vertices. Mohar and I conjectured that every 3-connected planar Cayley graph does. As explained in Section 1.5, perhaps the hardest case for this conjecture is the cubic case, and our classification constitutes significant progress in this direction.

1.2 Some examples and Mohar's conjecture

Let us consider some examples. Suppose G is a finite or 1-ended Cayley graph embedded in the plane, that some generator z of G spans a finite cyclic subgroup, and the corresponding cycles of G bound faces in this embedding, see Figure 1 (i). Then, considering the amalgamation product with respect to the subgroup $\langle z \rangle$, and using the generators of G, we obtain a multi-ended Cayley graph G' which can also be embedded in the plane: the face that was bounded by some coset C of $\langle z \rangle$ in G can be used to recursively accommodate the copy K of G glued along C and the further copies sharing a z cycle with K and so on; see Figure 1 (ii).

This kind of amalgamation can be used to produce new planar Cayley graphs from simpler ones, but it cannot yield cubic graphs since the degree of a vertex is increased. To amend this, Mohar [21] proposed the following variant of this operation. For every z cycle C of Figure 1 (ii), rotate one of its sides in such a way that the edges incident with C on either side do not have common endvertices, but appear in an in-out alternating fashion instead; see Figure 1 (iii). It is at first sight not clear why the new graph G'' produced like this is a Cayley graph, but in fact it is, and its group is an overgroup of the group of G, see [14]. We call the operation of Figure 1 a *twist-amalgamation*.

Our next example is slightly more complicated. Let this time G be the Cayley graph of a finite dihedral group shown in the left part of Figure 2. Then, for every 4-cycle C of G bounding a face F, we embed a copy of G in F, with C being glued with a corresponding cycle of the copy. As in our last example, this glueing does not identify vertices with vertices, but rather puts some vertices of one copy of G midway along some edges of the other, see the right part of Figure 2.

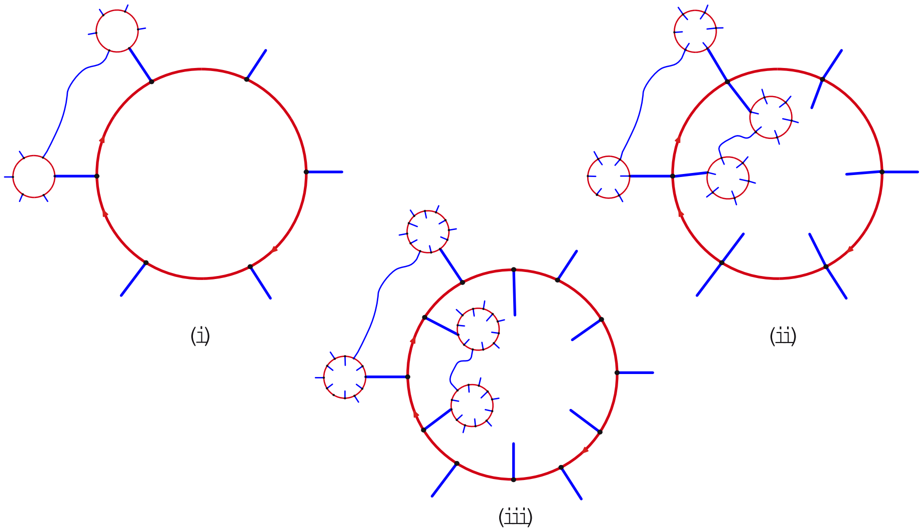

Figure 1: A twist-amalgamation.

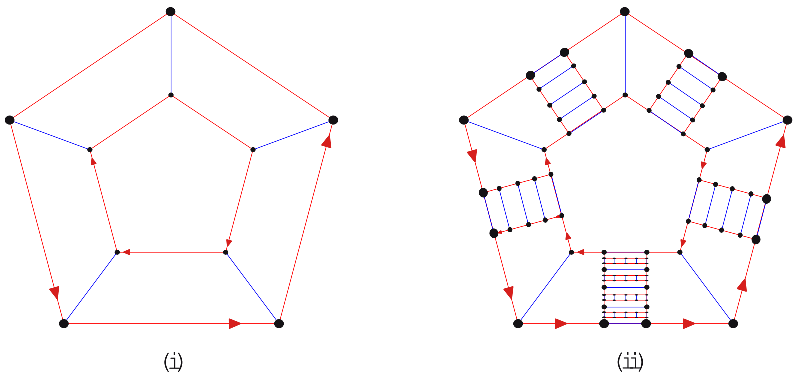

Figure 2: A twist-squeeze-amalgamation.

Unlike the previous example, where every edge of C was subdivided into two, this time every other edge of C is subdivided into three while every other edge is left intact. We recursively repeat this kind of glueing operation for the newly appeared face-bounding 4-cycles. Again, we obtain a new planar Cayley graph. We call the operation of Figure 2 a *twist-squeeze-amalgamation*.

These two examples are special cases of a more general, and more complicated, phenomenon: the hardest task addressed in this paper is to show that for every multi-ended cubic planar Cayley graph G', it is possible to obtain a presentation of G' from one of a finite or 1-ended cubic Cayley graph G embedded in G', by replacing some of the generators by new ones, and replacing each occurrence of an old generator z by a word W_z in the new generators. We call this operation a *word extension of G*. In our two examples it was enough to replace just one generator, and the corresponding word had length two or three. There are many cases where this is enough: the entries (5), (20) and (26) of Table 1 for example were obtained like that. However, there are much harder cases, where one needs to replace all generators by new ones, and there is no upper bound to the length of the replacing words W_z needed. This is the case for the last five entries of Table 1.

In our two examples we saw how a Cayley graph G' with infinitely many ends can be obtained from a finite or 1-ended Cayley graph G by glueing copies of G together. Mohar's aforementioned conjecture [22, Conjecture 3.1] was that every planar Cayley graph graph with more than one end can be obtained in a similar way:

Conjecture 1.1. *Let G be a planar Cayley graph. Then G can be obtained as the tree amalgamation of (subdivisions of) one or more planar Cayley graphs, each of which is either finite or 1-ended. Moreover, the identifying sets in these amalgamations correspond to cosets of finite subgroups.*

We will refrain from repeating the definition of Mohar's *tree amalgamation*, which can be found in [22, Section 2], and contend ourselves with examples. The phrase "subdivisions of" is missing in Mohar's conjecture, but apparently it was intended. In any case, omitting it makes the first sentence false as indicated by the above examples (the first of which is due to Mohar).

The results of this paper corroborate the first sentence of Conjecture 1.1. For example, the graph of Figure 1 (iii) is the tree amalgamation of the graph H obtained from Figure 1 (i) by subdividing every z edge into two, the amalgamations taking place along the facial cycles. Similarly, the graph of Figure 2 (ii) can be obtained from that of Figure 2 (i) by subdividing every directed edge into three, and then amalgamating copies of the resulting graph along the trapezoid cycles.

The second sentence of Conjecture 1.1 is easily shown to be false by the latter example.

1.3 The conjecture of Bonnington and Watkins

Bonnington and Watkins [4] made the following conjecture.

Conjecture 1.2 ([4])**.** *No locally finite vertex-transitive graph of connectivity 3 admits a planar embedding wherein some vertex is incident with more than one infinite face-boundary.*

Mohar [22] disproved this using his aforementioned amalgamation construction to obtain vertex-transitive graphs of connectivity 3 with arbitrarily many infinite face-boundaries incident with each vertex. This motivated Bonnington and Mohar [21] to ask whether there is a planar 3-connected vertex-transitive graph all face-boundaries of which are infinite.

In this paper we show that, surprisingly, such graphs do exist: we construct 3-connected planar cubic Cayley graphs in which no face is bounded by a finite cycle. Note that a 3-connected planar graph has an essentially unique embedding (see Theorem 3.2 below), which makes the existence of such examples more surprising.

In fact, we describe all 3-connected planar cubic Cayley graphs with the property that no face is bounded by a finite cycle: they are precisely the members of the families (34) and (37) of Table 1. See Corollary 9.13 and Corollary 9.19.

1.4 Planar presentations and effective enumeration

Some of the groups appearing in Table 1 were already known and well-studied, namely those admitting a Cayley complex that can be embedded in the plane after removing some redundant simplices, see [12]. These are the 1-ended ones, those of connectivity 1, and some of those appearing in Table 1 with connectivity 2 [13]. This paper is mainly concerned with the remaining ones. It turns out that none of the Cayley graphs of Table 1 from (26) on admits a Cayley complex as above[1] [12].

However, it follows from our results that they admit a Cayley complex X for which there is a mapping $\sigma : X \to \mathbb{R}^2$ such that 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, or their intersection is a 2-simplex bounded by the two parallel edges corresponding to some involution in the generating set. We call such an X an *almost planar* Cayley complex.

Theorem 1.2. *Every cubic planar Cayley graph is the 1-skeleton of an almost planar Cayley complex of the same group.*

A *planar presentation* is a group presentation giving rise to an almost planar Cayley complex. This property can be recognised by an algorithm, see Definition 9.9 for an example or [16] for details.

All presentations in Table 1 are planar. The semi-colons contained in some of them are used to distinguish relators that induce face-boundaries, which appear before the semi-colon, from relators that induce cycles separating the graph in infinite components, and appear after the semi-colon.

Theorem 1.1 yields an effective enumeration of the planar cubic Cayley graphs, corroborating the aforementioned conjecture of Droms et. al. [8, 9]. Moreover, our constructions imply that every planar cubic Cayley graph is effectively computable.

[1] This does not immediately imply that their groups are distinct from the groups of earlier entries, as a group can have various planar Cayley graphs of different nature; in fact, for very small values of the involved parameters it can happen that the corresponding group coincides with one from an earlier entry. For example, as pointed out by M. Dunwoody (personal communication), the group of case (35) for $n = 2$ and $m = 1$ coincides with that of case (8) for $n = 2$ and $m = 4$. However, for larger values of the parameters all groups from (26) on should be distinct from the previous ones; a proof of this fact is in progress.

Using the ideas of this paper Theorem 1.2 is extended in [16] to arbitrary planar Cayley graphs. As planar presentations can be recognised by an algorithm, this settles the general case of the aforementioned conjecture.

1.5 Hamilton circles in Cayley graphs

The following conjecture was motivated by [15]:

Conjecture 1.3 (Georgakopoulos & Mohar (unbublished)). *Every finitely generated 3-connected planar Cayley graph admits a Hamilton circle.*

A *Hamilton circle* of a graph G is a homeomorphic image of the circle S^1 in the end-compactification of G containing all vertices. It has been proved [5] that every 4-connected locally finite VAP-free planar graph has a Hamilton circle, and it is conjectured that the VAP-freeness requirement can be dropped in that theorem; this generalises a classical theorem of Tutte for finite graphs. A Cayley graph of degree 4 or more will either be 4-connected, in which case one can try to apply the above result, or its small separators will give away information about its structure; e.g. we know that a non-4-connected Cayley graph of degree 5 or more cannot be 1-ended [2, Lemma 2.4.].

This means that the cubic case plays an important role for Conjecture 1.3. Thus this paper constitutes an important step towards its proof, as it describes the structure of these graphs in a way that can be exploited to prove hamiltonicity. Indeed, in graphs like the one in Figure 1 (iii) Mohar and I found a way to construct a Hamilton circle of the whole graph using Hamilton circles of the basis graph Figure 1 (i).

1.6 Structure of the proof

This paper is structured as follows. After some definitions and basic facts, we handle the finite and 1-ended case, corresponding to entries (12) to (19) of Table 1, in Section 7; entries (1) and (2) are easy and entries (3) to (11) were handled in [13]; see Section 5. Most of this paper is concerned with entries from (20) on. They are divided naturally into cases according to the number of generators (2 or 3), the existence of cycles avoiding one of the generators, and the spin behaviour (see Section 2.3) in the corresponding embedding; they occupy Sections 8 and 9.

These cases vary considerably in difficulty, in general becoming more difficult as we progress towards the end of Table 1. However, there is a common structure: for every graph G as in the entries from (20) on, we begin by finding a subgroup, typically finite or 1-ended, that has a Cayley graph G_2 topologically embedded in G. This yields a Dunwoody structure tree whose nodes are the copies of G_2 in G, in which two nodes are adjacent if the corresponding copies share a cycle of G. These cycles are typically cycles of minimal length bounding two infinite components of G. We then show that the desired presentation of G can be derived from a presentation of G_2 by translating each generator g of G_2 into a word in the generators of G, which word can be read off the path in G onto which g is mapped when embedding G_2 into G. The presentation of G_2 is obtained from some earlier entry of Table 1, and many of the entries, even multi-ended ones, find themselves embedded in other entries, giving rise to a rich structure.

$\kappa(G) = 1$	1. $\langle a, b \mid b^2, a^n \rangle$, $n \in \{\infty, 2, 3, \ldots\}$ 2. $\langle b, c, d \mid b^2, c^2, d^2, (bc)^n \rangle$, $n \in \{\infty, 1, 2, 3, \ldots\}$
$\kappa(G) = 2$	3. $G \cong Cay \langle a, b \mid b^2, (ab)^n \rangle$, $n \geq 2$ 4. $G \cong Cay \langle a, b \mid b^2, (aba^{-1}b^{-1})^n \rangle$, $n \geq 1$ 5. $G \cong Cay \langle a, b \mid b^2, a^4, (a^2b)^n \rangle$, $n \geq 2$; 6. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bc)^2, (bcd)^m \rangle$, $m \geq 2$ 7. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bc)^{2n}, (cbcd)^m \rangle$, $n, m \geq 2$ 8. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (bd)^m \rangle$, $n, m \geq 2$ 9. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (b(cb)^n d)^m \rangle$, $n \geq 1, m \geq 2$ 10. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bcbd)^m \rangle$, $m \geq 1$ 11. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bc)^n, cd \rangle$, $n \geq 1$
$\kappa(G) = 3$, G is 1-ended or finite, with two generators	12. $G \cong Cay \langle a, b \mid b^2, a^n, (ab)^m \rangle$, $n \geq 3$, $m \geq 2$ 13. $G \cong Cay \langle a, b \mid b^2, a^n, (aba^{-1}b)^m \rangle$, $n \geq 3, m \geq 1$ 14. $G \cong Cay \langle a, b \mid b^2, (a^2b)^m \rangle$, $m \geq 1$ 15. $G \cong Cay \langle a, b \mid b^2, (a^2ba^{-2}b)^m \rangle$, $m \geq 1$
$\kappa(G) = 3$, G is 1-ended or finite, with three generators	16. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bcd)^n \rangle$, $n \geq 1$ 17. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (cbcdbd)^n \rangle$, $n \geq 1$ 18. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (bdcd)^m \rangle$, $n \geq 2, m \geq 1$ 19. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (cd)^m, (db)^p \rangle$, $n, m, p \geq 2$
$\kappa(G) = 3$, G is multi-ended, with two generators	20. $G \cong Cay \langle a, b \mid b^2, (a^2b)^m; a^{2n} \rangle$, $n \geq 3, m \geq 2$ 21. $G \cong Cay \langle a, b \mid b^2, (a^2ba^{-2}b)^m; a^{2n} \rangle$, $n \geq 3, m \geq 1$ 22. $G \cong Cay \langle a, b \mid b^2, a^2ba^{-2}b; (baba^{-1})^n \rangle$, $n \geq 2$ 23. $G \cong Cay \langle a, b \mid b^2, (a^2ba^{-2}b)^m; (baba^{-1})^n \rangle$, $n, m, p \geq 2$ 24. $G \cong Cay \langle a, b \mid b^2, (a^2b)^2; (ab)^{2m} \rangle$, $m \geq 2$ 25. $G \cong Cay \langle a, b \mid b^2, (a^2b)^n; (ab)^{2m} \rangle$, $n \geq 3$, $m \geq 2$
$\kappa(G) = 3$, G is multi-ended, with three generators	26. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bcd)^m; (bc)^n \rangle$, $n \geq 2$, $m \geq 2$ 27. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bcdcbd)^m; (bc)^n \rangle$, $n \geq 2, m \geq 1$ 28. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (cd)^m, (dbcb)^p; (bc)^{2n} \rangle$, $n, m, p \geq 2$ 29. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bcbdcd)^m; (bc)^{2n} \rangle$, $n \geq 2, m \geq 1$ 30. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bcbdcd)^p; (dc)^{2n}, (bc)^{2m} \rangle$, $n, m, p \geq 2$ 31. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bcbdcd)^p; (dbcb)^{2n}, (bc)^{2m} \rangle$, $n, m, p \geq 2$ 32. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bcd)^2; (bcdc)^n \rangle$, $n \geq 2$ 33. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bcd)^k; \mathcal{P} \rangle$, $k \geq 3$, $\mathcal{P}$ is a non-crossing pattern 34. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2; \mathcal{P} \rangle$, $\mathcal{P}$ is a non-regular non-crossing pattern 35. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bdbcdc)^k; (c(bc)^n d)^{2m} \rangle$, $k \geq 2$, $n, m \geq 1, n + m \geq 3$ 36. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bdbcdc)^q; (c(bc)^{n-1}d)^{2m}, (c(bc)^n d)^{2r} \rangle$, $n, r, m, q \geq 2$ 37. $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2; (c(bc)^{n-1}d)^{2m}, (c(bc)^n d)^{2r} \rangle$, $n, r, m \geq 2$

Table 1: Classification of the cubic planar Cayley graphs. All presentations are *planar* in the sense of Section 1.4.

In Section 10 we pose some further problems, one of which seeks for a generalisation of this proof-structure to a general multi-ended Cayley graph.

2 Definitions

2.1 Cayley graphs and group presentations

We will follow the terminology of [7] for graph-theoretical terms and that of [3] for group-theoretical ones. Let us recall the definitions most relevant for this paper.

Let Γ be group and let S be a symmetric generating set of Γ. The Cayley graph $Cay(\Gamma, S)$ of Γ with respect to S is a coloured directed graph $G = (V, E)$ constructed as follows. The vertex set of G is Γ, and the set of colours we will use is S. For every $g \in \Gamma, s \in S$ 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 a *colour-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 colour-automorphisms of G. Any presentation of Γ in which S is the set of generators will also be called a presentation of $Cay(\Gamma, S)$.

If $s \in S$ 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). However, when calculating the degree of a vertex of a Cayley graph we will count only one edge for each such pair, and we will draw only one, undirected, edge in our figures. For example, if S consists of three involutions then we consider the corresponding Cayley graph to be cubic. This convention is common in the literature, and it is necessary if one wants to study the property of being a Cayley graph as a graph-theoretical invariant, like e.g. in Sabidussi's theorem [1, Proposition 3.1].

Given a group presentation $\langle S \mid \mathcal{R} \rangle$ we will use the notation $Cay \langle S \mid \mathcal{R} \rangle$ for the Cayley graph of this group with respect to S.

If $R \in \mathcal{R}$ is any relator in such a presentation and g is a vertex of $G = Cay \langle S \mid \mathcal{R} \rangle$, then starting from g and following the edges corresponding to the letters in R in order we obtain a closed walk W in G. We then say that W is *induced* by R; note that for a given R there are several walks in G induced by R, one for each starting vertex $g \in V(G)$. If R induces a cycle then we say that R is *simple*; note that this does not depend on the choice of the starting vertex g. A presentation $\langle a, b, \dots \mid R_1, R_2, \dots \rangle$ of a group Γ is called *simple*, if R_i is simple for every i. In other words, if for every i no proper subword of any R_i is a relation in Γ.

Define the (finitary) *cycle space* $\mathcal{C}_f(G)$ of a graph $G = (V, E)$ to be the vector space over $\mathbb{Z}_2$ consisting of those subsets of E such that can be written as a sum (modulo 2) of a finite set of circuits, where a set of edges $D \subseteq E$ is called a *circuit* if it is the edge set of a cycle of G. Thus $\mathcal{C}_f(G)$ is isomorphic to the first simplicial homology group of G over $\mathbb{Z}_2$. The *circuit* of a closed walk W is the set of edges traversed by W an even number of times. Note that the direction of the edges is ignored when defining circuits and $\mathcal{C}_f(G)$. The cycle space will be a useful tool in our study of Cayley graphs because of the following well-known fact which is easy to prove.

Lemma 2.1. *Let $G = Cay \langle S \mid R\rangle$ be a Cayley graph of the group Γ. Then the set of circuits of walks in G induced by relators in R generates $\mathcal{C}_f(G)$.*

Conversely, if R' is a set of words, with letters in a set $S \subseteq \Gamma$ generating Γ, such that the set of circuits of cycles of $Cay(G, S)$ induced by R' generates $\mathcal{C}_f(G)$, then $\langle S \mid R'\rangle$ is a presentation of Γ.

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 cycle C of G is *induced* if every edge of G that has both endvertices on C is an edge of C.

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*, a 2-way infinite path is a *double 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.

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 $\mathbb{R}^2 \backslash \sigma(G)$. 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. Note that if G is cubic, which means that $|L| = 3$, then there are only two possible cyclic orders on L, and thus only two possible spins. Call an edge of G *spin-preserving* if its two endvertices have the same spin in σ, and call it *spin-reversing* otherwise. 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 (this definition is natural only if G is cubic; to extend it to the general case, demand that every two vertices have either the same spin, or the spin of the one is obtained by reversing the spin of the other).

It is straightforward to check that σ is consistent if and only if every colour-automorphism of G maps every facial walk to a facial walk.

It follows from Whitney's theorem mentioned in the introduction that if G is 3-connected then its essentially unique embedding must be 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 *VAP-free* for short, if the images of the vertices have no accumulation point in $\mathbb{R}^2$.

3 Known facts

In this section we recall some easy facts about Cayley graphs that we will use later. The reader may choose to skip this section and the next.

We begin with a well-known characterisation of Cayley graphs. Call an edge-colouring of a digraph G *proper*, if no vertex of G has two incoming or two outgoing edges with the same colour.

Theorem 3.1 (Sabidussi's Theorem [23, 1])**.** *A properly edge-coloured digraph is a Cayley graph if and only iffor every $x, y \in V(G)$ there is a colour-automorphism mapping x to y.*

The following classical result was proved by Whitney [24, Theorem 11] for finite graphs and by Imrich [17] for infinite ones.

Theorem 3.2. *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 3.3. *Let G be a 3-connected planar Cayley graph. Then every embedding of G is consistent.*

Finally, we recall the following fact that is reminiscent of MacLane's planarity criterion.

Theorem 3.4 ([13])**.** *Let $\langle S \mid R \rangle$ be a simple presentation and let $G = Cay \langle S \mid R \rangle$ be the corresponding Cayley graph. If no edge of G appears in more than two* *circuits induced by relators in R, then G is planar and has a VAP-free embedding the facial cycles of which are precisely the cycles of G induced by relators in R.*

4 General facts regarding connectivity

Call a relation of a group presentation *cyclic*, if it induces a cycle in the corresponding Cayley graph. An involution is by convention not cyclic. Cyclic relations are useful because they allow us to formulate the following lemma.

Lemma 4.1. *A cubic Cayley graph is 2-connected if and only if each of its generators a is in a cyclic relation.*

Proof. Firstly, note that if G is a cubic Cayley graph that is not 2-connected, then it must have a bridge e. But then the generator a corresponding to e cannot be in any cyclic relation because a cycle cannot contain a bridge.

Conversely, an edge corresponding to a generator a that is in no cyclic relation cannot lie in any cycle; thus it is a bridge, which means that any of its endpoints separates the graph. □

In many occasions we will use some of the graphs of Table 1 as building blocks in order to construct more complicated ones. Our next lemma will be useful in such cases, as it will allow us to deduce the fact that the new graphs are 3-connected from the fact that the building blocks were. Let K, K_1, K_2 be subsets of $V(G)$. We will say that K is *k-connected in G*, if for every vertex set $S \subseteq V(G)$ with $|S| < k$, there is a component of $G - S$ containing $K - S$. Similarly, we will say that K_1 is *k-connected to* K_2 in G, if for every vertex set $S \subseteq V(G)$ with $|S| < k$, there is a path from K_1 to K_2 in $G - S$; in particular, $|K_1|, |K_2| \geq k$.

Lemma 4.2 ([14])**.** *Let G be a graph, and let $\{K_i\}_{i \in I}$ be a family of subsets of $V(G)$ such that the following three assertions hold:*

1. *$\bigcup_I K_i = V(G)$;*
2. *For every $i \in I$, K_i is k-connected in G, and*
3. *For every $i, j \in I$ there is a finite sequence $i = n_0, \ldots, n_r = j$ such that K_{n_m} is k-connected to $K_{n_{m+1}}$ in G for every relevant m.*

Then G is k-connected.

5 Graphs of connectivity 1 or 2

It follows easily from Lemma 4.1 that if G is a cubic Cayley graph of connectivity $\kappa(G) = 1$, then its group $\Gamma(G)$ has one of the following presentations.

1. $\langle a, b \mid b^2, a^n \rangle$, $n \in \{\infty, 2, 3, \ldots\}$
2. $\langle b, c, d \mid b^2, c^2, d^2, (bc)^n \rangle$, $n \in \{\infty, 1, 2, 3, \ldots\}$,

where $n = \infty$ means that the corresponding relator is omitted.

All Cayley graphs corresponding to these presentations are planar; this follows easily from $\kappa(G) = 1$ and the assumption that G is cubic.

The planar cubic Cayley graphs of connectivity 2 were completely analysed in [13], yielding the following classification.

Theorem 5.1. *Let G be a planar cubic Cayley graph of connectivity 2. Then precisely one of the following is the case:*

1. $G \cong Cay\left\langle a, b \mid b^2, (ab)^n \right\rangle$, $n \geq 2$;
2. $G \cong Cay\left\langle a, b \mid b^2, (aba^{-1}b^{-1})^n \right\rangle$, $n \geq 1$;
3. $G \cong Cay\left\langle a, b \mid b^2, a^4, (a^2b)^n \right\rangle, n \geq 2$;
4. $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bc)^2, (bcd)^m \right\rangle$, $m \geq 2$;
5. $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bc)^{2n}, (cbcd)^m \right\rangle$, $n, m \geq 2$;
6. $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (bd)^m \right\rangle$, $n, m \geq 2$;
7. $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (b(cb)^n d)^m \right\rangle, n \geq 1, m \geq 2$;
8. $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bcbd)^m \right\rangle$, $m \geq 1$;
9. $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, cd \right\rangle$, $n \geq 1$ *(degenerate cases with redundant generators).*

Conversely, each of the above presentations, with parameters chosen in the specified domains, yields a planar cubic Cayley graph of connectivity 2.

The above presentations are planar; see [13, Corollary 6.3].

We use Theorem 5.1 to obtain the entries (3)–(11) of Table 1, but we also use these graphs as important building blocks in later constructions of 3-connected Cayley graphs.

The following sections are devoted to the 3-connected case.

6 Crossings of shortest dividing cycles

Given a plane graph G and a cycle C of G, the Jordan curve theorem yields two distinct regions $IN(C), OUT(C)$ of $\mathbb{R}^2 \backslash C$ which we call the *sides* of C. Define the *closed sides* $\overline{IN(C)}, \overline{OUT(C)}$ of C to be the respective union of $IN(C), OUT(C)$ with C.

Call a cycle C of G a *dividing cycle* if both sides of C contain infinitely many vertices, and call C a *shortest dividing cycle* if its length is minimal among all dividing cycles of G. The shortest dividing cycles will play a very important role in this paper: in most of the group presentations we construct, any relator that does not induce a face boundary will induce a shortest dividing cycle. In this section we provide some general facts, that will be useful later, about how pairs of shortest dividing cycles can meet.

We will say that two cycles C, C' of G *cross* each other if none of them is contained in the closure of a side of the other; equivalently, if each of $IN(C), OUT(C)$ meets both $IN(C'), OUT(C')$. It turns out that the ways in which shortest dividing cycles can cross are very restricted.

We will first consider the simplest case, when two shortest dividing cycles C, C' cross only once, that is, when $C - C'$ consists of two components and so does $C' - C$. In this case, $\mathbb{R}^2 \backslash (C \cup C')$ consists of four regions A, B, D, F as in Figure 3.

Since C, C' are dividing cycles, both the inside and the outside of each of them contain infinitely many vertices. This immediately implies that either

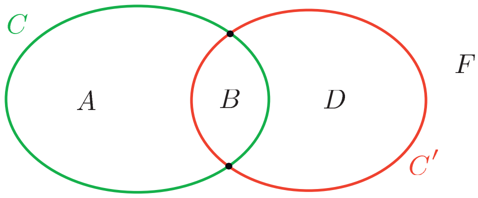

Figure 3: The four regions of $\mathbb{R}^2 \backslash (C \cup C')$.

both A, D are infinite or both B, F are infinite. We may assume without loss of generality that

$$\text{both } A, D \text{ are infinite} \tag{1}$$

since otherwise we could modify the embedding of G so as to fix C but exchange its inside with its outside, which would have the effect of renaming the regions A, B, D, F.

As G is cubic, every time C and C' intersect they must have at least one common edge. This gives rise to the three cases displayed in Figure 4. In two of those cases (upper half of Figure 4) we immediately obtain a contradiction: in each case we obtain two new cycles K, K' (dashed lines) both of which are dividing by (1). Now an easy double counting argument shows that the length at least one of K, K' must be less than $\ell := |C| = |C'|$, the length of a shortest dividing cycle of G: indeed, we have $|K| + |K'| < |C| + |C'|$.

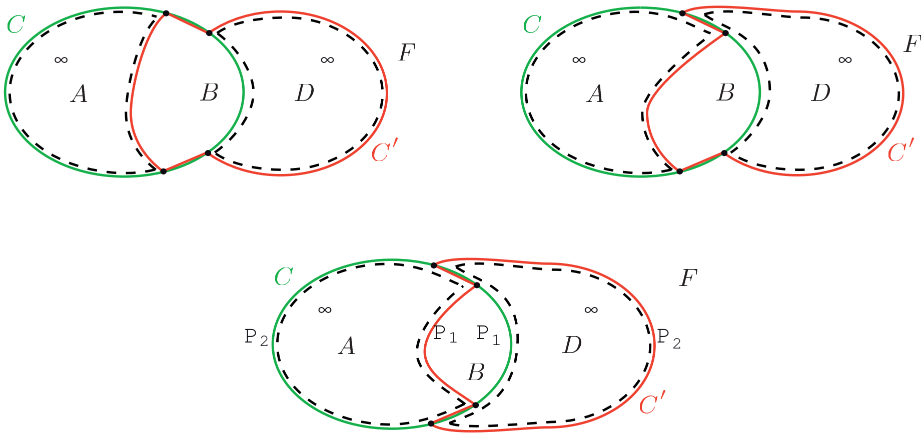

Figure 4: The three ways in which two dividing cycles C, C' can cross. The dashed cycles K, K' are also dividing.

Thus, whenever two shortest dividing cycles C, C' cross only once we must have the situation in the lower half of Figure 4, with each of A and D containing infinitely many vertices. Moreover,

$$\text{at least one of } B, F \text{ is finite,} \tag{2}$$

since otherwise one of the cycles $P_1 \cup P_1'$ and $P_2 \cup P_2'$ is dividing and has length less that ℓ by a double-counting argument as above.

If C, C' cross more than once then more or less the same arguments apply, but we have to be a bit more careful. We define the regions A, B, D, F as follows. We let $A := IN(C) \cap OUT(C')$, $B := IN(C) \cap IN(C')$, $D := OUT(C) \cap IN(C')$, and $F := OUT(C) \cap OUT(C')$. Now each time C, C' cross, the corresponding subpaths must be arranged as in the lower half of Figure 4 as can be shown by repeating the above arguments, except that this time K, K' need not be cycles but could be more complicated closed walks. Still, we can decompose each of them into a finite collection of cycles, one of which will have to be dividing and shorter that ℓ.

Note that $|P_1| = |P'_1|$ must hold in Figure 4 because if, say, $|P_1| < |P'_1|$ then we can replace P'_1 by P_1 in C' to obtain the cycle K' which is shorter than C'. But by (1) K' would then be a dividing cycle, contradicting the fact that C' is shortest possible. Similarly, we obtain $|P_2| = |P'_2|$. Thus, both dotted cycles in the lower half of Figure 4 have the same length as C, and by (1) again they are shortest dividing cycles.

This means that whenever we have two crossing shortest dividing cycles C, C' in a cubic graph G,

there is a shortest dividing cycle K which is the union of a non-trivial subpath of each of C, C' which subpaths witness the fact that C, C' cross. (3)

This will imply the following tool that we will use in many cases.

Corollary 6.1. *Let C, C' be two crossing shortest dividing cycles in a cubic graph G. Let Q be a maximal common subpath of C, C', and let F be a facial subpath of C or C' that has maximal length among all facial subpaths of all shortest dividing cycles of G. Then if Q and F have a common endvertex y, their edges incident with y are distinct.*

Proof. Suppose, to the contrary, that such a path Q finishes with the edge $e = xy$ which is also the last edge of the facial subpath F of C that has maximum length. Since Q was a maximal common subpath of C, C', the edge $f = yz$ of C following Q is different from the edge $f' = yz'$ of C' following Q. By (3), we can combine C and C' into a new shortest dividing cycle K that comprises a subpath P of C ending with e and a subpath P' of C' starting with f'. Recall that the path F is facial, and let H denote its incident face. Since F was a maximal facial subpath of C, H is not incident with the edge $f \in E(C)$. Thus, as G is cubic, H is incident with f'. Now note that K contains f' by construction, as well as the last edge e of F. If K contains all of F, then it contains the facial path $F \cup f'$, which contradicts the maximality of the length of F among all facial subpaths of all shortest dividing cycles. Thus P' must interrupt F. This however yields a contradiction to the embedding. Indeed, since H is a face, it is contained in either $IN(C)$ or $OUT(C)$. Recall also that P' must by definition be contained in $IN(C)$ or $OUT(C)$. But as P' contains the edge f' which is incident with H, both H and P' must be accommodated in the same side of C. Thus P' cannot contain any of the edges incident with F, since these edges lie in the other side of C. This shows that P' cannot interrupt F, and we obtain a contradiction that completes the proof. □

7 The finite and 1-ended cubic planar Cayley graphs

In this section we analyse the cubic planar Cayley graphs that are either finite or infinite but with only 1 end. Many, perhaps all, of these graphs were already known. They appear as entries (12)–(19) of Table 1, but also provide building blocks for many of the more interesting entries.

We begin with some general properties of planar Cayley graphs with at most 1 end.

Theorem 7.1. *Every finite or 1-ended cubic Cayley graph is 3-connected.*

This is proved in [1, Chapter 27, Theorem 3.7.] for finite G and in [2, Lemma 2.4.] for infinite G.

Combining this with Lemma 3.3 easily yields the following well-known fact.

Lemma 7.2. *In an 1-ended plane Cayley graph all face-boundaries are finite.*

Conversely, we have

Lemma 7.3 ([12, Lemma 3.4])**.** *A plane 2-connected graph with no dividing cycle and no infinite face-boundary is either finite or 1-ended.*

Our last lemma is

Lemma 7.4. *Let $G = Cay\langle S \mid \mathcal{R}\rangle$ be finite or 1-ended and planar , and let $\mathcal{R}'$ be a set of relations of $\Gamma(G)$ such that every face boundary of G is induced by some relation in $\mathcal{R}'$. Then $\Gamma(G) \cong \langle S \mid \mathcal{R}'\rangle$.*

Proof. It suffices to show that the edge-set of every cycle C of G is a sum of edge-sets of finite face-boundaries. This is indeed the case, for as G is at most 1-ended, there must be a side A of C containing only finitely many vertices, and so $E(C)$ is the sum of the edge-sets of the face-boundaries lying in A. □

We can now proceed with the main results of this section.

Theorem 7.5. *Let $G = Cay\left\langle a, b \mid b^2, \ldots\right\rangle$ be planar and finite or 1-ended. Then precisely one of the following is the case:*

1. *$G \cong Cay\left\langle a, b \mid b^2, a^n, (ab)^m\right\rangle$, $n \geq 3$, $m \geq 2$ (and a, b preserve spin);*
2. *$G \cong Cay\left\langle a, b \mid b^2, a^n, (aba^{-1}b)^m\right\rangle$, $n \geq 3, m \geq 1$ (and only a preserves spin);*
3. *$G \cong Cay\left\langle a, b \mid b^2, (a^2b)^m\right\rangle$, $m \geq 2$ (and only b preserves spin);*
4. *$G \cong Cay\left\langle a, b \mid b^2, (a^2ba^{-2}b)^m\right\rangle$, $m \geq 1$ (and a, b reverse spin);*
5. *$G \cong Cay\left\langle a, b \mid b^2, a^2, (ab)^n\right\rangle$, $n \geq 2$ or $G \cong Cay\left\langle a, b \mid b^2, ab\right\rangle$ (degenerate cases in which G is not cubic).*

All presentations above are planar.

Conversely, each of the presentations (i)–(iv), with parameters chosen in the specified domains, yields a planar, finite or 1-ended, non-trivial cubic Cayley graph.

Proof. For the forward implication, let $G = Cay\left\langle a, b \mid b^2, \dots\right\rangle$ be planar, with at most one end. G is 3-connected by Theorem 7.1, unless a is an involution too in which case we have one of the degenerate cases of (v). In all non-degenerate cases, the (essentially unique) embedding is consistent with respect to spin by Lemma 3.3. Thus we have precisely one of the following cases.

Case (i): both a, b preserve spin.

Since a preserves spin the walk $P := \mathbb{0}, a, a^2, a^3, \dots$ is facial. By Lemma 7.2 P cannot be a ray, so it spans a finite cycle C of length n say. Similarly, the walk $Q := \mathbb{0}, a, ab, aba, abab, \dots$ is facial because b also preserves spin, and by the above argument it must span a finite cycle C' with edges alternating between a and b. Note that both C, C' are face boundaries. Moreover, by Theorem 3.2 every face-boundary of G is a translate of one of C, C'. By Lemma 7.4 this means that the relations $a^n, (ab)^m$, inducing C and C' respectively, combined with b^2 yield a presentation of $\Gamma(G)$. Thus $G \cong Cay\left\langle a, b \mid b^2, a^n, (ab)^m\right\rangle$ with $m \geq 2$ as claimed.

Case (ii): a preserves spin and b reverses spin.

As in Case (i) we conclude that G has finite a-coloured faces, induced by a relation a^n. As b reverses spin now, a walk Q as above is not facial any more, and instead $\mathbb{0}, a, ab, aba^{-1}, aba^{-1}b, \dots$ is facial. By similar arguments we obtain the desired presentation $G \cong Cay\left\langle a, b \mid b^2, a^n, (aba^{-1}b)^m\right\rangle$, $n \geq 3, m \geq 1$.

Case (iii): b preserves spin and a reverses spin.

The walk $\mathbb{0}, a, aa, aab, aaba, \dots$ is now facial, and spans a finite cycle C of length $3m$ for some $m \geq 2$. It is also straightforward to check that every facial path of G is of that form, in other words, every face boundary is a translate of C. Again by Lemma 7.4 we obtain the desired presentation $G \cong Cay\left\langle a, b \mid b^2, (a^2b)^m\right\rangle$, $m \geq 1$.

Case (iv): both a, b reverse spin.

This case is similar to the previous one, except that the a edges on a facial walk do not all have the same direction now, but instead their directions alternate after each b edge. We thus obtain $G \cong Cay\left\langle a, b \mid b^2, (a^2ba^{-2}b)^m\right\rangle$, $m \geq 1$.

For the converse assertion, given any presentation of the form (i), let us show that the Cayley graph G is planar. We begin with an auxiliary plane graph H, namely, the graph of the regular tiling of the sphere, euclidean plane, or hyperbolic plane with n m-gons meeting at every vertex. The existence of H is well-known and not hard to prove [20]. Note that H is a vertex-transitive graph: for any two vertices $x, y \in V(H)$ it is straightforward to inductively construct an isomorphism between the balls of radius r around these vertices. To obtain G from H, replace every vertex x of H with a cycle C_x of length n, and join each edge incident with x to a distinct vertex of C_x, keeping the cyclic ordering, so that the graph remains planar. Then, assign to each edge of C_x the label a, and direct it in such a way that C_x is oriented clockwise. Moreover, assign to every other edge, coming from an edge of H, the label b. Let G be the resulting coloured graph. We claim that G is the Cayley graph corresponding to presentation (i). Indeed, the fact that G is a Cayley graph follows easily from Sabidussi's Theorem 3.1 and the fact that H was vertex-transitive. The fact that G has the desired presentation now follows from the forward implication which we have already proved since, by construction, all edges of G preserve spin.

Given any presentation of the form (ii), it is now easy to construct an embedding of the corresponding Cayley graph G: we can start with a graph G' of type (i) with parameters $n, m' := 2m$, and then reverse the orientation of 'every other' a-cycle to obtain G. More precisely, let $G' \cong Cay \left\langle a, b \mid b^2, a^n, (ab)^{2m} \right\rangle$, and define a bipartition $\{X, Y\}$ of the a-cycles of G' by letting X (resp. Y) be the set of those a-cycles that can be reached from a fixed vertex $o \in V(G)$ by a path containing an even (resp. odd) number of b-labelled edges. To see that this is indeed a bipartition, note that every relation in the group of $\left\langle a, b \mid b^2, a^n, (ab)^{2m} \right\rangle$ contains an even number of appearances of the letter b. Note moreover, that any two a-cycles that are connected by a b edge lie in distinct classes of this bipartition. Thus, if we reverse the orientation of every a-cycle lying in X we obtain a plane graph G in which only the a edges preserve spin. Again, we can check that G is a Cayley graph using Sabidussi's theorem, and we apply the forward implication to show that G has the presentation (ii).

We handle case (iii) similarly to case (i), except that now the auxiliary graph H is obtained by contracting the b-edges instead of the a-labelled ones. To achieve the desired orientation, replace every b-edge by a 2-cycle, and orient all these cycles clockwise; then make sure that for every face boundary of type $(a^2b)^m$ all edges are oriented in the same direction.

A graph of type (iv) can be obtained by one of type (iii) by reversing the orientation of every other a-path, similarly to the above reduction of type (ii) to type (i).

Finally, the Cayley graphs of type (v) are finite and easy to construct.

It follows easily from our construction that all our presentations are planar. □

We now proceed with the case where G is generated by three generators.

Theorem 7.6. *Let $G = Cay \left\langle b, c, d \mid b^2, c^2, d^2, \ldots \right\rangle$ be planar and finite or 1-ended. Then precisely one of the following is the case:*

1. *$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bcd)^n \right\rangle$, $n \geq 1$ (all colours preserve spin);*
2. *$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (cbcdbd)^n \right\rangle$, $n \geq 1$ (only c, d preserve spin);*
3. *$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (bdcd)^m \right\rangle$, $n \geq 2, m \geq 1$ (only d preserves spin);*
4. *$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (cd)^m, (db)^p \right\rangle$, $n, m, p \geq 2$ (all colours reverse spin);*
5. *$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, cd \right\rangle$, $n \geq 1$ (degenerate, non-3-connected cases with redundant generators).*

All presentations above are planar.

Conversely, each of the above presentations, with parameters chosen in the specified domains, yields a planar, finite or 1-ended, non-trivial cubic Cayley graph.

Proof. By the arguments of Theorem 7.5 G uniquely embeds in the sphere, unless we are in the degenerate case (v), and so the colours behave consistently with respect to spin. We thus have the following non-degenerate cases. Recall that by Lemma 7.2 all face boundaries of G are finite; we are going to tacitly make use of this fact in all cases.

Case (i): all colours preserve spin.

Consider the path $\mathbb{0}, b, bc$, which is facial for some face boundary F since it only has two edges. Since c preserves spin, the next edge on F must be coloured d. Similarly, the edge after that must be coloured b since d preserves spin. Continuing like that we conclude that the edges of F follow the pattern $bcdbcdb\ldots$, in other words, F can be induced by the relation $(bcd)^n$. Since all edges preserve spin, any two faces of G that share an edge $e = xy$ can be mapped to each other by the automorphism of G exchanging x with y, where we are using the fact that all edges correspond to involutions. Thus all face boundaries of G have the same form as F. Similarly to Theorem 7.5, we can now apply Lemma 7.4 to conclude that $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bcd)^n \right\rangle$, $n \geq 1$.

Case (ii): precisely two colours, c, d say, preserve spin.

In this case, at every vertex of G the situation looks locally like Figure 5, and it is straightforward to check that every face boundary is of the form $(cbcdbd)^n$. Indeed, as c, d preserve spin, every facial walk of the form cd is a subwalk of a facial walk of the form $bcdb$. Moreover, as b reverses spin, every facial walk of the form cb is a subwalk of a facial walk of the form cbc and every facial walk of the form db is a subwalk of a facial walk of the form dbd. Furthermore, there is no facial walk of the form bcb or bdb. Combining these facts one obtains that every face boundary is indeed of the form $(cbcdbd)^n$. Thus Lemma 7.4 yields $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (cbcdbd)^n \right\rangle$, $n \geq 1$.

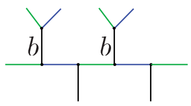

Figure 5: The local situation around any vertex in case (ii).

Case (iii): precisely one colour, d say, preserves spin.

Then a walk alternating in b, c is facial, and so G has two coloured faces, induced by a relation $(bc)^n$ for some $n \geq 2$. It is straightforward to check, by observing the spin behaviour, that every facial walk containing a d-edge is of the form $bdcdbd\ldots$. Thus, by Lemma 7.4 again, $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (bdcd)^m \right\rangle$, $n \geq 2, m \geq 1$.

Case (iv): all colours reverse spin.

Similarly to Case (iii) it follows that any two-coloured walk is facial, and so at every vertex we have three different kinds of incident face boundaries, each a two-coloured cycle. By Lemma 7.4 we obtain $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (cd)^m, (db)^p \right\rangle$, $n, m, p \geq 2$.

The converse assertion can again be proved using explicit constructions similarly to what we did in the proof of Theorem 7.5, but these constructions become more complicated now and so we will follow a different, and interesting, approach that makes use of Theorem 3.4. Given any presentation of one of these four forms, we will check that the presentation is simple. Moreover, it is straightforward to check that each of these presentations satisfies the conditions of Theorem 3.4, with every edge of G appearing in precisely two circuits induced by the specified relators. Thus, by that theorem, the corresponding Cayley graph is planar, and has a VAP-free embedding σ in which the finite face boundaries are induced by those relators, and each edge lies in precisely two such boundaries. This means that all face boundaries in σ are finite. Moreover,

σ has no dividing cycles because it is VAP-free. Thus we can apply Lemma 7.3 which yields that G is 1-ended if it is infinite.

It only remains to check that our presentations are simple indeed. To see that (i) is simple, note that even if we set $d = 1$ we can prove no subword of $(bc)^n$ to be a relation since $(bc)^n$ is simple in the subgroup generated by b, c. Similarly, consider the subgroups of (ii) generated by b, d and b, c to see that the presentation is simple. To see that (iv) is simple, consider the subgroup F consisting of all elements that can be presented by a word of even length in the letters b, c, d; this is indeed a subgroup since every relation of the original group contains an even number of letters. Note that F is generated by bc, cd, db, and in fact one of these generators is redundant. Thus $F \cong \langle a', b' \mid (a')^n, (b')^m, (a'b')^q \rangle$, where $a' := bc, b' := cd$. Note that if we impose $b'^2 = 1$ then this presentation reduces to (i) of Theorem 7.5, and we implicitly checked there that $(a')^n$ is simple in that group. Thus $(a')^n$ must be simple in F, which means that $(bc)^n$ is simple in our original group Γ as the element bc cannot have a smaller order in Γ than in a subgroup of Γ. By symmetry, all relations in (iv) are simple. Instead of showing that presentation (iii) is also simple, let us rather explicitly construct the desired Cayley graph from one of type (iv). For this, let $G' = Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (cd)^{2m}, (db)^{2m} \right\rangle$. We have already proved that G' has an embedding with all edges reversing spin, and we have observed that the set of elements that can be presented by a word of even length forms a subgroup F, which of course has index 2. Note that any two adjacent vertices of G' lie in distinct left cosets of F. Now let G be the graph obtained from G' by exchanging, for every $f \in F$, the labels of the edges incident with f labelled b and c. Note that every vertex of G is still incident with all three labels by the above observation. It is easy to prove that G is a Cayley graph using Sabidussi's theorem. Moreover, the d-edges now preserve spin. Thus, by the forward implication, G has the desired presentation (iii). Finally, presentation (v) gives rise to a finite Cayley graph which is easy to embed.

It follows easily from our construction that all our presentations are planar. □

8 The planar multi-ended Cayley graphs with 2 generators

Having already characterised the 1-ended cubic planar Cayley graphs in the previous section, we turn our attention to our main object of interest, the planar, 3-connected, multi-ended Cayley graphs. In this section we consider those generated by two generators, one of which must be an involution. We will distinguish two cases according to whether the other generator has finite or infinite order, discussed separately in the following two subsections. We will obtain planar presentations for each of those graphs, as well as explicit constructions of their embeddings. Our results are summarised in Theorems 8.3 and 8.5 below.

8.1 Graphs with monochromatic cycles

In this section we consider the case when $a^m = 1$ for some $m > 2$. It turns out that in this case

$$a \text{ reverses spin,} \tag{4}$$

because of the following lemma.

Lemma 8.1. *Let $G = Cay\langle a, b \mid b^2, \ldots\rangle$ be a 2-connected Cayley graph with a consistent embedding in which a preserves spin and such that a has finite order. Then G has at most one end.*

Proof. We begin by showing that

$$G \text{ does not have a dividing cycle.} \tag{5}$$

Indeed, suppose to the contrary that C is a shortest dividing cycle, and choose C so that it has a facial subpath F that has maximum length among all shortest dividing cycles and all their facial subpaths. We distinguish three cases, all of which will lead to a contradiction.

Case I: one of the end-edges e of F is labelled a.

In this case, we can rotate the finite a-cycle containing e by a colour-automorphism of G to translate C to a further shortest dividing cycle C' that crosses C in such a way that e is the last edge of a common subpath of C and C'; see Figure 6 (left). Indeed, since e is the last edge of F, the edge of F before it must have been labelled b, and the edge of C following e is labelled a too, so that such a rotation is possible. This crossing immediately contradicts Corollary 6.1.

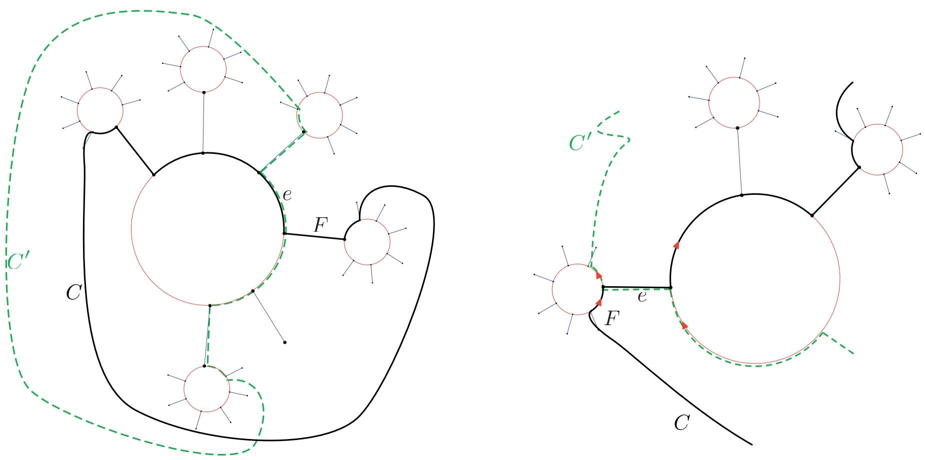

Figure 6: Creating a crossing in Case I (left) and Case II (right).

Case II: both end-edges of F are labelled b, and b reverses spin.

In this case we can apply the colour-automorphism of G that exchanges the endvertices of some end-edge e of F to translate C to a new shortest dividing cycle C' that crosses C; see Figure 6 (right). Again, the crossing we obtain contradicts Corollary 6.1.

Case III: both end-edges of F are labelled b, and b preserves spin.

Let $e = yz$ be an end-edge of F. Since b is an involution, the edge $g = xy$ of F preceding e is labelled a, and so g is not an end-edge of F, which means that F has at least one more edge. The edge $f = wx$ of F preceding g is labelled b again, because otherwise F would not be facial, see Figure 7. We claim that the colour-automorphism T of G that maps z to x translates C to a cycle $C' = T(C)$ that crosses C. To see this, let $\mathcal{F}$ be the face whose boundary contains F, and note that the edge $h = zv$ following F on C is not incident with $\mathcal{F}$, for F is maximally facial. Assume without loss of generality that g is directed from x to y. Then, as b preserves spin, h is directed from v to z. This implies that $T(h) \neq g$, and so C' has an edge, namely $T(h)$, that lies in the side of C not containing $\mathcal{F}$.

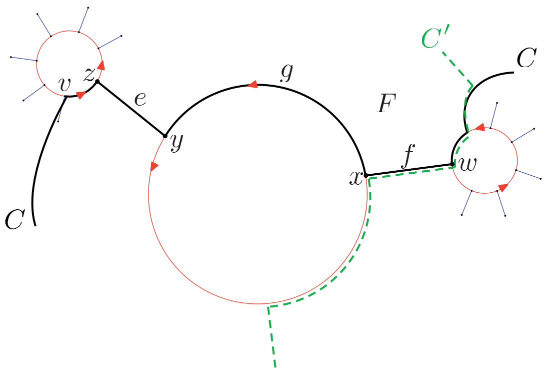

Figure 7: Creating a crossing in Case III.

On the other hand, $T(F)$, which is a subpath of C' since $F \subseteq C$, is facial by Theorem 3.2. Since $T(e) = f$ and we are assuming that both a and b preserve spin, it follows that T 'maps' $\mathcal{F}$ to itself, more precisely, that $T(F)$ is contained in the boundary of $\mathcal{F}$. As F is also contained in the boundary of $\mathcal{F}$, and $|F| = |T(F)|$, this means that if we start at x and walk around $\mathcal{F}$ in the direction of w, then $T(F)$ spans more edges of the boundary of $\mathcal{F}$ than F does. Since $T(F) \subseteq C'$, this proves that C' has an edge in the side of C that contains $\mathcal{F}$. Combined with our earlier observation that C' also meets the other side of C proves that C' crosses C, and in fact so that one of their common subpaths ends with f. As f is also the final edge of $T(F)$, a facial subpath of maximum length, this contradicts Corollary 6.1 again.

Thus, in all three cases we obtained a contradiction, and so we have established (5). We can now exploit this fact to prove our next claim.

$$\text{Every face boundary of } G \text{ is finite.} \tag{6}$$

To see this, let C be a cycle containing an edge e coloured b; such a cycle exists because G is 2-connected. Now as C cannot be dividing by (5), one of its sides contains only finitely many vertices. Thus, all faces contained in that side have a finite boundary. This means that e itself is on a finite face boundary F. Now consider one of the a-cycles D incident with e. Rotating along D by colour-automorphisms of G we can map F to the other face incident with e. Since we can also map e to any other b edge by a colour-automorphismof G, Theorem 3.2

now implies that every face boundary containing a b edge is finite. As we are assuming that a preserves spin, any face boundary not containing a b edge is a finite a-cycle of length m, the order of a. These two observations together prove (6).

We can now apply Lemma 7.3, using (5) and (6), to prove that G has at most one end. $\square$

It follows from (4) that the order m of a is even, since the b edges incident with any a-cycle C must alternate between the two sides of C (Figure 8).

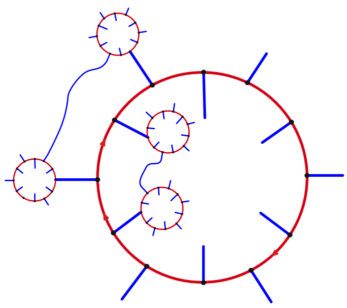

Figure 8: An a-cycle and some of its translates. Every a-edge reverses spin.

Now consider the subgroup Γ_2 of Γ spanned by a^2 and b. We claim that Γ_2 is a proper subgroup of Γ; in fact, that for every a-cycle C of Γ spanned by a, at most half of the vertices of C lie in Γ_2. To see this, note that if x, y are two elements of Γ_2, then there is an x–y path P in G the a edges of which can be decomposed into incident pairs. Now using Figure 8 it is easy to see that whenever such a path P meets an a-cycle C of G, the two edges of P incident with C lie in the same side of C. In other words, P cannot cross any a-cycle C of G. This fact easily implies our claim.

By the same token, given an embedding σ of G, we can modify G and σ to obtain a Cayley graph G_2 of Γ_2, with respect to the generating set $\{a^2, b\}$, and an embedding σ_2 of G_2 as follows. For every a-cycle C of G that contains a vertex in Γ_2, delete all vertices and edges in the side of C that does not meet Γ_2; such a side exists by the above argument. Let G_2' be the graph obtained after doing so for every such cycle. Then, suppress all vertices of G_2' that now have degree two; that is, replace any a-labelled path xPy of length two whose middle vertex now has no incident b edge by a single x-y edge directed the same way as P, bearing a new label z (corresponding to the generator a^2), to obtain G_2.

We will soon see that G_2 uniquely determines G. But let us first look at G_2 more closely. To begin with, using Lemma 4.1 it is easy to show that

$$G_2 \text{ is 2-connected.} \tag{7}$$

Note that the embedding σ_2 is by construction a consistent embedding of G_2, but the z-labelled edges now preserve spin. Applying Lemma 8.1 to G_2 thus implies that

$$G_2 \text{ has at most one end.} \tag{8}$$

It now follows from our characterization of such graphs in Section 7 that,

$$\begin{array}{l} \text{either } G_2 \cong Cay\left\langle z, b \mid b^2, z^n, (zb)^m \right\rangle,\ n \geq 3,\ m \geq 2 \\ \text{or } G_2 \cong Cay\left\langle z, b \mid b^2, z^n, (zbz^{-1}b)^m \right\rangle,\ n \geq 3, m \geq 1. \end{array} \tag{9}$$

Indeed, since G has a monochromatic cycle, it must belong to one of the types (i), (ii) or (v) of Theorem 7.5. However, the third type can immediately be eliminated, as it would either imply that a pair of vertices of G_2 adjacent by a z edge disconnects G, which cannot be the case as G is assumed to be 3-connected, or it would imply that G is a graph on four vertices. (If we drop the assumption that G be 3-connected though then such a graph does exist, and it is unique: it is described in [13, Figure 5].)

We are going to use the presentation (9) of G_2 to obtain a presentation of G. To achieve this, we are going to find some relations of Γ such that the set of cycles induced by these relations generates $\mathcal{C}_f(G)$. In fact, these relations are just the ones that appear in the presentation (9). Intuitively, one way to prove this is as follows. Given an arbitrary cycle C of G, consider the finitely many translates $G_1, G_2, \ldots G_{k_C}$ of G_2' in G that intersect C, and observe that C can be written as a sum of cycles of the G_i. Now as any cycle in G_2 (and G_2') is a sum of cycles induced by the relators in (9), our claim follows by Lemma 2.1. We are going to use a similar argument in several occasions throughout this paper, and rather than repeating the argumentation each time, we are going to use the following more abstract to obtain a rigorous proof of the fact that G has the claimed presentation.

Lemma 8.2 ([14]). *Let G be any graph, and let $\mathcal{X}$ be a set of subgraphs of G with the following properties:*

1. $\bigcup_{H \in \mathcal{X}} H = G$,
2. *no edge of G lies in infinitely many elements of $\mathcal{X}$; and*
3. *there is a tree $T(\mathcal{X}, E_T)$ on $\mathcal{X}$ such that for every edge $e \in E_T$, joining H_i to H_j say, there is a common cycle F_i of H_i and H_j such that F_i separates $\bigcup_{H \in V(T_1)} H$ from $\bigcup_{H \in V(T_2)} H$ where T_1, T_2 are the two components of $T - e$.*

Then, for every choice $(\mathcal{F}_H)_{H \in \mathcal{X}}$ of a generating set $\mathcal{F}_H$ of $\mathcal{C}_f(H)$ for each element H of $\mathcal{X}$, the union $\bigcup_{\mathcal{X}} \mathcal{F}_H$ generates $\mathcal{C}_f(G)$.

Now in order to apply this to our particular Cayley graph G, consider for every vertex $x \in V(G)$ the left coset of Γ_2 in Γ containing x, and let $H^x \subseteq G$ be the corresponding copy of G_2'; in other words, H^x is the image of G_2' under the colour-automorphism x. We are now going organize these copies into an auxiliary tree T satisfying (iii). The vertex set of T is the set of graphs H^x, for all vertices x of G. Two vertices H^x, H^y of T are joined by an edge if and only if they share an a-coloured cycle of G. It follows easily from the definition of G_2' that T is indeed a tree, and that it satisfies (iii). The interested reader will be able to check that T is in fact a Dunwoody structure tree [6] of G.

It is clear that this set $\{H^x \mid x \in V(G)\}$ satisfies (i) and (ii) of Theorem 8.2. Before we apply Theorem 8.2, it remains to define the generating sets $\mathcal{F}_x$. For this, let $\mathcal{F}$ be the set of cycles of G induced by the relators in the presentation

(9) —for which have distinguished two cases— after replacing each appearance of the letter z by the word a^2. Note that all these cycles are contained in G_2'. Moreover, $\mathcal{F}$ generates $\mathcal{C}_f(G_2')$ by Lemma 2.1 since (9) is a presentation of Γ_2.

Thus, defining $\mathcal{F}_x$ to be the image of $\mathcal{F}$ under the automorphism of G that maps G_2' to its copy H^x, we meet the requirement that $\mathcal{F}_x$ generate $\mathcal{C}_f(H^x)$ for every x, and Theorem 8.2 yields that $\bigcup \mathcal{F}_x$ generates $\mathcal{C}_f(G)$. By the second sentence of Lemma 2.1 and the definition of $\mathcal{F}$ it now follows that either $G \cong Cay\left\langle a, b \mid b^2, a^{2n}, (a^2b)^m \right\rangle$, $n \geq 3$, $m \geq 2$ or $G \cong Cay\left\langle a, b \mid b^2, a^{2n}, (a^2ba^{-2}b)^m \right\rangle$, $n \geq 3, m \geq 1$. The first (respectively second) case occurs if b preserves (resp. reverses) spin, as can be seen by applying Theorem 7.5 to G_2. Thus we have

Theorem 8.3. *Let $G = Cay\left\langle a, b \mid b^2, \ldots \right\rangle$ be a 3-connected Cayley graph with more than one end and suppose that a has finite order. Then a reverses spin. If b preserves spin then*
$G \cong Cay\left\langle a, b \mid b^2, (a^2b)^m; a^{2n} \right\rangle, n \geq 3, m \geq 2.$
If b reverses spin then
$G \cong Cay\left\langle a, b \mid b^2, (a^2ba^{-2}b)^m; a^{2n} \right\rangle$, $n \geq 3, m \geq 1.$

In both cases, the presentation is planar.

Moreover, G is the Mohar amalgamation of $G_2 \cong Cay\left\langle z, b \mid b^2, z^n, (zb)^m \right\rangle$ or $G_2 \cong Cay\left\langle z, b \mid b^2, z^n, (zbz^{-1}b)^m \right\rangle$ with itself.

Conversely, each of these presentations, with parameters chosen in the specified domains, yields a Cayley graph as above.

Here, a *Mohar amalgamation* is the operation of Figure 1 (iii) described in the Introduction.

Proof. The forward implication was proved in the above discussion. It remains to prove the converse and the fact that these graphs can be obtained by the claimed Mohar amalgamations. We will prove both these assertions simultaneously. For this, given one of the above presentations, consider the auxiliary presentation $\Gamma_2 = \left\langle z, b \mid b^2, z^n, (zb)^m \right\rangle$ or $\Gamma_2 = \left\langle z, b \mid b^2, z^n, (zbz^{-1}b)^m \right\rangle$ obtained by replacing a^2 by z throughout. Applying Theorem 7.5 (i) or (ii) to this presentation shows that the corresponding Cayley graph G_2 is finite or 1-ended, and has an embedding in which all monochromatic cycles induced by z^n bound faces. The other relation also induces facial cycles but we will not use this fact. Now construct a graph G as the Mohar amalgamation of G_2 with itself with respect to the z-monochromatic cycles, orienting the pasted discs in such a way that all b edges preserve (respectively, reverse) spin if the presentation we started with was of the first (resp. second) kind.

It follows easily from Sabidussi's theorem that the plane edge-coloured graph G just constructed is a Cayley graph. Moreover, G has infinitely many ends: by construction, any z-coloured cycle separates two infinite components. We claim that G is 3-connected. To prove this, suppose $\{x, y\} \subset V(G)$ separates G. It is easy to see that by the construction of G the vertices of any z-coloured cycle cannot be separated by $\{x, y\}$. Moreover, as G_2 is itself 3-connected by Theorem 7.1, no two z-coloured cycles that lie in a common copy of G_2 can be separated by $\{x, y\}$. But for any two z-coloured cycles Z, Z' of G there is by construction a finite sequence $Z_1 = C_1, C_2, \ldots, C_k = Z'$ such that C_i, C_{i+1} are z-coloured cycles lying in a common copy of G_2. Applying Lemma 4.2 using these two facts yields that G is 3-connected as claimed.

We can now apply the forward implication of Theorem 8.3 to this graph G; we thus obtain that the corresponding group has the desired presentation, namely the one we used for the construction of G. $\square$

8.2 Graphs without monochromatic cycles

We now consider the case when a has infinite order, and so G has no monochromatic cycles. Instead, the a edges span double rays in G. Also in this case we will be able to prove that

$$a \text{ reverses spin,} \tag{10}$$

because of the following lemma which is similar to Lemma 8.1.

Lemma 8.4. *Let $G = Cay\langle a, b \mid b^2, \ldots \rangle$ be a 3-connected planar Cayley graph in which a has infinite order. Then a reverses spin.*

Proof. Suppose, to the contrary, that a preserves spin. We will show that G must have a dividing cycle. For this, pick two vertices x, y that lie in the same a-coloured double ray R of G. As G is 3-connected, there are three independent x–y paths P_1, P_2, P_3 by Menger's theorem [7, Theorem 3.3.1]. By an easy topological argument, there must be a pair of those paths, say P_1, P_2, whose union is a cycle C such that some side of C contains a tail of R and the other side of C contains P_3, see Figure 9. We may assume without loss of generality that P_3 is not a single b edge, for we are allowed to choose x and y far apart. Thus the side of C containing P_3 contains at least one vertex z. Now as all a edges preserve spin, the a-coloured double ray R' incident with z is facial, and so it cannot exit the cycle C. This means that C is dividing, since one of its sides contains R' and the other contains R.

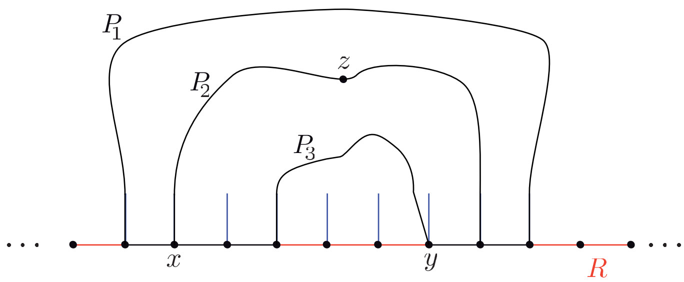

Figure 9: Finding a dividing cycle for the proof of Lemma 8.4.

Now imitating the proof of Lemma 8.1 we can obtain a contradiction to the fact that G has a dividing cycle. Thus a must reverse spin. $\square$

Using (10) and Theorem 3.2 it follows easily that any two face boundaries can be mapped to each other by a colour-automorphism of G. Thus all faces of G have the same size N. This implies that G must have a dividing cycle: if N is infinite, then any induced cycle is dividing. If N is finite, then the existence of a dividing cycle follows immediately from Lemma 7.3.

We now turn our attention to the shortest dividing cycles of G. We will be able to describe these cycles precisely, but in order to do so we have to start with a more modest task, namely to prove that

no shortest dividing cycle of G has an a-labelled subpath comprising more than two edges. (11)

To show this, let p be the maximum length of an a-labelled subpath of a shortest dividing cycle of G, and let P be a an instance of such a path with $||P|| = p$. We need to prove that $p \leq 2$. We distinguish two cases according to the parity of p.

If p is even, and at least 4, then we can shift a shortest dividing cycle C containing P by two edges of the a-coloured double ray containing P to obtain a translate C' of C that crosses C as in Figure 10. It follows from our discussion in Section 6 that one of the regions, D, say resulting from this crossing contains only finitely many vertices, while the regions B and F must each contain infinitely many vertices. Moreover, the paths $xC'y$ and xCy in Figure 10 must have equal lengths, for otherwise the shortest of them provides a shortcut for the C or C', contradicting the minimality of the latter. Thus we can replace xCy by $xC'y$ in C to obtain a new shortest dividing cycle C'' with an a-labelled subpath containing P and two more a-edges. This contradicts the maximality of P.

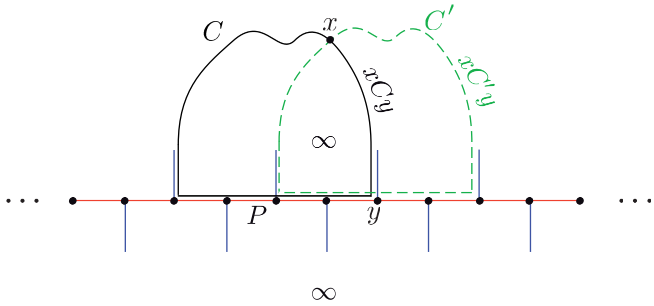

Figure 10: A crossing in the case that p is even.

If p is odd, and at least 3, then we can again shift a shortest dividing cycle C containing P along the a-coloured double ray containing P, this time shifting only by one edge, to obtain a translate C' of C that crosses C as in Figure 11. By the same arguments, one of the regions A, D must be finite, and replacing a subpath of C' for a subpath of C we obtain a new shortest dividing cycle with a longer a coloured subpath than P.

Thus in both cases we obtained a contradiction to the maximality of P, which proves our claim (11).

Next, we prove that

no shortest dividing cycle of G has a facial subpath comprising more than three edges. (12)

Indeed, let F be a maximal facial subpath of the shortest dividing cycle C and suppose that $||F|| \geq 4$. Then F does not contain a bab subpath because a reverses spin, so F must contain an aa subpath. It then follows from (11) that any such subpath of C lies within a $baab$ subpath of C. Since any $baab$ path

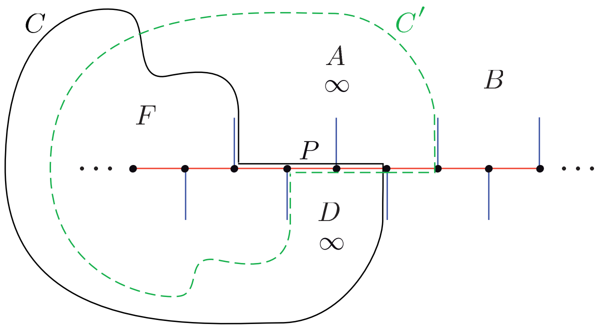

Figure 11: A crossing in the case where p is odd.

is facial in our case, F has to contain a $baab$ subpath too. We distinguish two cases, according to the colour of the last edge e of F. If that colour is a then we have the situation on the left half of Figure 12, while if it is b then we have the situation on the right half of Figure 12; here we are using the fact that F cannot finish with an aa subpath followed by an a-edge on C because of (11).

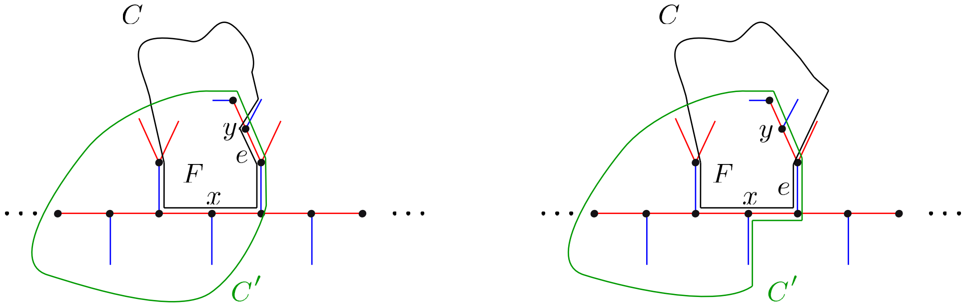

Figure 12: Proof that a facial subpath of length four forces a crossing.

In both cases, the colour-automorphism of G mapping x to y translates C to some other shortest dividing cycle C' that intersects C. We would like to show that C' crosses C. Note that this colour-automorphism maps F to a facial subpath F' of C' that is incident with the same face $\mathcal{F}$ as F was. It is now easy to see that F' must meet both sides of C, for otherwise C' has more edges along the boundary of $\mathcal{F}$ than C has, which would contradict the maximality of F. This proves that C' crosses C indeed. As we could have chosen F to have maximum length among all facial subpaths of all shortest dividing cycles of G without loss of generality, we can apply Corollary 6.1 to this crossing to obtain a contradiction that proves (12).

Combining (12) with (11) easily implies that a shortest dividing cycle cannot even contain an a-labelled subpath comprising more than one edge, which means that

The colours of the edges of every shortest dividing cycle of G alternate between a and b. (13)

8.2.1 Tidy cycles

Now consider a b edge e of a shortest dividing cycle C. We will say that e is *tidy* in C if the two edges incident with e that do not lie in $E(C)$ lie in the same side of C. We will say that C is *tidy* if all its b edges are tidy in C. The reason why we are interested in tidy shortest dividing cycles is the following proposition, which will help us obtain the desired presentation of G:

A shortest dividing cycle is tidy if and only if it is induced by a word of the form $(ab)^n$ and b preserves spin, or it is induced by a word of the form $(aba^{-1}b)^n$ and b reverses spin. (14)

Indeed, this follows immediately from the fact that a reverses spin and the definition of *tidy*.

In order to be able to exploit this fact we need to show that

G has a tidy shortest dividing cycle. (15)

In fact, we will show that every shortest dividing cycle is tidy, unless every face of G has size 6. For this, we first have to show the following.

If G has a face of size greater than 6 then no two shortest dividing cycles of G cross. (16)

To prove this, note first that since a reverses spin, any two faces of G can be mapped to each other, and so every face has size greater than 6 in this case.

Now suppose that two shortest dividing cycles C, C' cross. Then by (2) there is a subpath P of C and a subpath P' of C' such that $P \cup P'$ is a cycle K bounding a finite region B. We will show that such a region contradicts Euler's formula $n - m + f = 2$ for the sphere. To see this, let H be the finite plane subgraph of G spanned by K and all vertices in B. Note that for a cubic finite graph J, using the fact that $|E(J)| = \frac{3}{2}|V(J)|$ and that every edge lies in precisely two faces, Euler's formula can be rewritten as

(Euler's formula for a cubic graph) $$\sum_{k\geq 3} c_k|F_k| = 12, \tag{17}$$

where $|F_k|$ is the number of k-gonal faces of J, and $c_k := 6 - k$ is the *curvature* of each k-gonal face. This means in particular that a cubic plane graph must have some faces of size k less that 6.

Our graph H almost contradicts (17) since all faces of G have size greater that 6, except that it has some vertices of degree two on its boundary K. To amend these degrees, consider the graph H' obtained from two copies of H by joining corresponding vertices of degree two by an edge, and note that H' is cubic. Consider an embedding of H in the sphere such that the two copies of H occupy two disjoint discs D_1, D_2, and the newly added edges and their incident faces lie in an annulus Z that joins these discs. Now note that all faces within these discs still have size greater than 6, contributing a negative curvature to (17), but Z can contain 4-gons. Still, we will show that the number of 4-gons is not enough to balance the deficit in curvature.

To begin with, note that every face in Z has even size. Moreover, since a reverses spin and each of P, P' in the construction of H was a, b alternating by (13), it is easy to see that an a-edge at the boundary of D_1 or D_2 cannot be

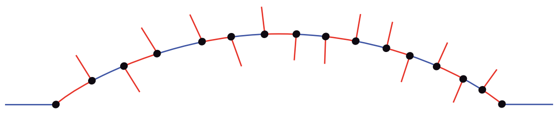

Figure 13:

incident with a 4-gon in Z unless it was one of the four end-edges of P and P' (Figure 13). On the other hand, a b edge at the boundary of D_1 or D_2 can be incident with a 4-gon in Z, however, the fact that a reverses spin easily implies that if e, f are two b edges incident with a 4-gon and both lying on P, say, then there must be a b edge between them on P that is incident with an 8-gon in Z, and the same holds for P'. These two observations together imply that the number $|F_4|$ of 4-gons in Z is bounded from above by $|F_k| \leq 4+2+|F_8| = 6+|F_8|$. But as $c_4+c_8 = 0$, this means that the total curvature contributed to (17) by the faces in Z is at most $6f_4 = 12$, and as D_1, D_2 only contain faces of size greater than 6 each of which contributes a negative curvature, we obtain a contradiction to (17). This proves (16).

This argument also explains why we have to treat the case when every face is a hexagon separately.

We now return to the proof of (15). In fact, we are going to show something stronger: in the case where G has a face of size greater than 6, if a shortest dividing cycle is untidy then it must cross some other shortest dividing cycle, and so by (16) every shortest dividing cycle is tidy.

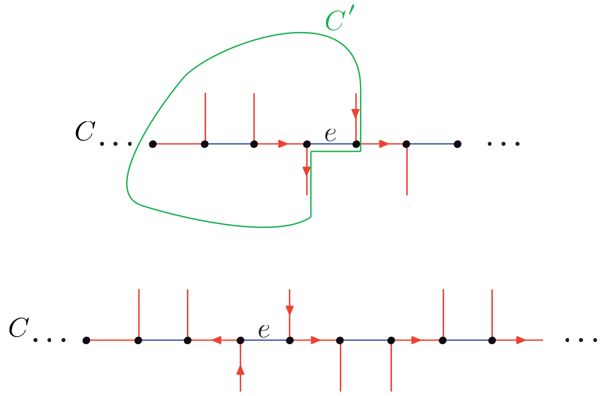

Figure 14: Finding a crossing at an untidy edge e.

For this, suppose there is a shortest dividing cycle C with an untidy b edge e. If b reverses spin then exchanging the endvertices of e by a colour-automorphism translates C to a cycle C' that crosses C and we are done (Figure 14, top). If b reverses spin, then note that the two a-edges of C incident with e both point towards e, or both point away from e (Figure 14, bottom). Now suppose we walk around C starting at e. Every time we come to a tidy b edge, its two incident a-edges in C point in the same direction. Thus, as the two a-edges in C incident with e point in opposite directions, before we arrive at e again we must visit a further untidy edge e' the incident edges of which point away from

e' if the incident edges of e point towards e and the other way round. Now translating e' to e maps C to a cycle C' that again crosses C. Thus in both cases applying (16) we prove that

If G has a face of size greater than 6 then every shortest dividing cycle of G is tidy. (18)

In particular, we proved (15) in the case that G has a face of size greater than 6.

The hexagonal grid case

It remains to consider the case when every face of G has size 6, since, easily, no face can have a size smaller than 6 when a reverses spin and has infinite order. In this case, it is easy to see that any two a-coloured double rays that are joined by a b edge are joined by infinitely many b edges that together with these double rays form an infinite strip of hexagons. Using this fact and (13) it is easy to check that the subgroup of Γ spanned by a has finite index n, in other words, $V(G)$ is spanned by finitely many a-coloured double rays.

Let us first consider the case when b reverses spin, and so all the a double rays point in the same direction, see Figure 15. Now let C be any cycle in G such that the colours of the edges of C alternate between a and b, the directions of the a-edges being arbitrary. For instance, C could be a shortest dividing cycle by (13). We claim that

if w is a word with letters a, a^{-1}, b that induces C then the letter a appears as often as the letter a^{-1} in w. (19)

To see this, note first that translating C by a^2 we obtain a cycle C' whose union with C bounds a strip of $|C|/2$ many hexagons (Figure 15). We call any of the translates of C by $a^{2i}, i \in \mathbb{N}$ a *level*. Enumerate the levels by the integers so that neighbouring levels are assigned consecutive numbers, and these numbers increase whenever we apply a^2.

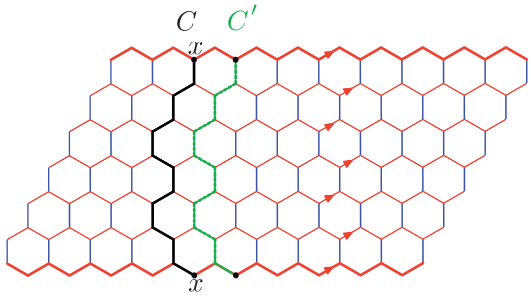

Figure 15: The case when every face is a hexagon and b reverses spin, and a shortest dividing cycle C. The top and bottom rays coincide, and so do the two vertices marked x.

It is straightforward to check, using the spin behaviour of the edges and the structure of C, that for any side A of C, all (a-coloured) edges that have precisely one endvertex in C are directed the same way, that is, either they are all directed from C into A or the other way round. Moreover, if these edges are directed from C say, the edges incident with C on its other side are directed towards C. Now pick a vertex z of C and consider the path P starting at z and induced by w. We may assume without loss of generality that the successor y of z on P is not in C, for otherwise we could have started reading w at y. Note that now any appearance of a or a^{-1} in w forces a change of level, while any appearance of b leaves us in the same level. Moreover, it follows from the aforementioned property of the direction of the edges incident with C that any appearance of a increases the level by one, while any appearance of a^{-1} decreases it. But as w induces a cycle, P must return to its initial vertex, and so (19) follows.

Now let C be any shortest dividing cycle of G. Using (19) we will now modify C into a tidy shortest dividing cycle. For this, note that we can replace any path of the form aba^{-1} in C by $a^{-1}ba$ (and the other way round), to obtain another shortest dividing cycle; indeed, $aba^{-1}a^{-1}ba$ is a relation, as it induces a boundary of a hexagonal face. Applying this operation several times, and using (19), we can reshuffle the letters in a word w inducing C to obtain a new word w' that still induces a shortest dividing cycle and has the form $w' = (aba^{-1}b)^n$. By (14) any cycle induced by w' is tidy, and so we achieved our aim to show (15) in case b reverses spin.

If b preserves spin instead, then the same arguments still apply with the following slight modification. In this case, adjacent a double rays point in different directions. Partition their set into two equal subsets none of which contains two adjacent double rays, and call the elements of one the two subsets the *even* double rays, and call the remaining ones *odd*. Now colour every edge that lies in an odd double ray with a new colour o and reverse its direction. Pretending that we cannot distinguish the colour a from o, and using the fact that every fourth edge on any shortest dividing cycle must be coloured a and every fourth edge must be coloured o, we can apply the above arguments to obtain a shortest dividing cycle C of the form $(abo^{-1}b)^n$. As o^{-1} is an alias for a, (14) implies again that C is tidy, so we have proved (15) in this case too.

Thus (15) holds in all cases.

8.2.2 Structure and presentations

It now follows from (14) that every a-edge of G lies in a unique tidy shortest dividing cycle. Similarly, every b edge of G lies in precisely two tidy shortest dividing cycles. In other words, if we represent the involution b by two parallel edges, then the tidy shortest dividing cycles form a decomposition of $E(G)$ into edge-disjoint cycles.

Moreover,

the tidy shortest dividing cycles incident with (the b edges of) a given tidy shortest dividing cycle C lie alternately in its inside and outside. (20)

This allows us to obtain a presentation, as well as a precise description of the embedding of G by methods similar to those of Section 8.1, where we had dividing a-coloured cycles and the b edges incident with any of them lied alternately in its inside and outside.

So similarly to what we did there, we now define the subgroup Γ_2 of Γ to be the subgroup spanned by b, $a^{-1}ba$ and aba^{-1} if b reverses spin, or the subgroup spanned by b and aba (and the inverse $a^{-1}ba^{-1}$ of aba) if b preserves spin. We claim that Γ_2 is a proper subgroup of Γ, and will prove this by showing that if C is a tidy shortest dividing cycle containing the identity then only one of the sides of C meets Γ_2. To see this, note that by (20) any path in G composed as a concatenation of subpaths induced by the words b, aba or b, aba^{-1} and their inverses generating Γ_2 can never cross from the inside of a tidy cycle to its outside. Thus, only vertices that lie in one of the sides of C, or any other tidy shortest dividing cycle, meet Γ_2.

Given an embedding σ of G we can, still similarly to what we did in Section 8.1, modify G and σ to obtain a Cayley graph G_2 of Γ_2, with respect to the above generating set and an embedding σ_2 of G_2 as follows. For every tidy shortest dividing cycle C of G that contains a vertex in Γ_2, delete all vertices and edges in the side of C that does not meet Γ_2; such a side exists by the above argument. Let G_2' be the graph obtained after doing so for every such cycle. Then, suppress all vertices of G_2' that now have degree two to obtain G_2.

It is easy to prove that G_2 is 2-connected. Indeed, for every vertex $x \in V(G_2)$, the two tidy shortest dividing cycles of G incident with x form a subgraph of G that contains the neighbourhood $N(x)$ of x and is connected even after removing x, which means that x cannot disconnect G_2. With a little bit more effort we can even prove that

G_2 is 3-connected unless $V(G_2)$ is contained in a tidy shortest dividing cycle of G. (21)

For this, suppose that $G_2 - \{x, y\}$ is disconnected. If the vertices x, y lie in no common tidy shortest dividing cycle of G, then by the above argument their neighbourhoods are connected, a contradiction. So let C be an shortest dividing cycle of G containing both vertices.

It could happen that $V(C) = V(G_2)$, which is the case when G has hexagonal faces (for example, when it is the graph of Figure 15). In this case G_2 is not 3-connected: it is a finite cycle with parallel edges. Now assume that this is not the case, which means that some b edge e of $C \cap G_2$ is incident with a further tidy shortest dividing cycle $C_e \neq C$ that contains vertices that lie in $G_2 \backslash C$. But then, given two b-edges $e, e' \in C \cap G_2$ it is straightforward to check that there is a colour-automorphism g of G that rotates C, fixing it set-wise, and maps e to e', and so every b edge in $C \cap G_2$ has this property.

We claim that if e, e' are consecutive b edges of $C \cap G_2$ then there is a path in $G - C$ connecting their incident cycles $C_e, C_{e'}$. To see this, note that as G is 3-connected, removing e and its endvertices p, q does not disconnect G, and so there is a C_e–C path P in $G - \{p, q\}$. Let f be the b edge in C incident with the endpoint of P. If $f = e'$ then P is the path we were looking for and we are done. If $f \neq e'$, then consider the colour-automorphism g of G that rotates C and maps e to e', and let $P' := gP$ be the image of P under g. Since the endvertices of both e and e' lie in G_2, it follows by the construction of G_2 that P and P' lie in the same side of C. Thus, as g rotates C, P' must meet P at some vertex z, say. Now combining the subpaths of P, P' from their starting point up to z we obtain the desired path joining C_e to $C_{e'}$.

In fact, we can now see that e and e' do not have to be consecutive in the above assertion: For if $e = e_1, e_2, \ldots, e_k = e'$ is a sequence of consecutive b edges of $C \cap G_2$, then combining the C_{e_i}–$C_{e_{i+1}}$ paths we just constructed with subpaths of the cycles C_{e_i} we can construct a C_e–$C_{e'}$ path in $G - C$.

We will now use this kind of path to complete our proof of (21). For this, let $C_x \neq C$ be the other tidy shortest dividing cycle of G containing x, and define C_y similarly for y. Note that each of $C_x \cap G_2, C_y \cap G_2$ is connected to C in $G_2 - \{x, y\}$. Thus, if $\{x, y\}$ separates G_2 then it has to separate C into two subarcs P_1, P_2 that lie in distinct components of $G_2 - \{x, y\}$. Pick vertices $v \in P_1 \cap G_2$ and $w \in P_1 \cap G_2$, and let e_v, e_w be the b edges containing v, w respectively. By our last observation there is a C_{e_v}–C_{e_w} path P in $G - C$. If we could transform P into a C_{e_v}–C_{e_w} path P in $G_2 - C$ we would be done, since such a path would contradict the fact that x, y separate P_1 from P_2. But this is easy to do: for every tidy shortest dividing cycle $C' \neq C_{e_v}, C_{e_w}$ of G visited by P, if P enters the side A of C' that does not meet G_2 then it has to exit that side again revisiting C', and we can replace the subpath of P that lies in A by a subarc of C' with the same endpoints. Doing so for every such cycle C', we transform P into a path P' in the auxiliary graph G'_2 (see the definition of G_2), and it is straightforward to transform P' into a v–w path in $G_2 - C$. This completes the proof of (21).

We now consider separately the cases when b preserves or reverses spin in G.

If b reverses spin, then recall that Γ_2 was spanned by b, $c := a^{-1}ba$ and $d := aba^{-1}$, and note that all these generators are involutions. Note moreover that, by construction, σ_2 is a consistent embedding, and that all edges reverse spin in σ_2. Graphs of this kind are characterised in Section 9.1.1 below, and it turns out (Corollary 9.3) that if G_2 is 3-connected then it is finite or 1-ended, so we can use our characterisation of those graphs from Section 7. We obtain that in this case $G_2 \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (cd)^m, (db)^p \right\rangle,\ n, m, p \geq 2$, i.e. possibility (iv) of Theorem 7.6.

If Γ_2 is not 3-connected, then by (21) we have $c = d$ and $G_2 \cong Cay\left\langle b, c \mid b^2, c^2, (bc)^n \right\rangle$, where $n \geq 2$ for if $n = 1$ then the endvertices of any b edge would separate G.

If b preserves spin, then recall that Γ_2 was spanned by b and $a^* := aba$. Note that a^* is not necessarily an involution, and that it preserves spin in σ_2 by construction. We distinguish two cases, according to whether the order n of a^* is finite or infinite.

If n is finite, then G_2 is one of the graphs we have already handled: by Lemma 8.1, G_2 has at most one end, and so it belongs to type (i) or (v) of Theorem 7.5. Thus, if $a^{*2} = 1$ (in which case G has hexagonal faces) then $G_2 \cong Cay\left\langle a^*, b \mid b^2, a^{*2}, (a^*b)^m \right\rangle$, and if $a^{*2} \neq 1$ then $G_2 \cong Cay\left\langle a^*, b \mid b^2, a^{*n}, (a^*b)^m \right\rangle$, $n \geq 3$, $m \geq 2$.

If n is infinite, then in particular $a^{*2} \neq 1$ and so G_2 is 3-connected by (21). By Lemma 8.4 G_2 has at most one end, so we can apply Theorem 7.5 again, but this time the conclusion is that no such graph G_2 exists. Thus n must be finite.

In all cases, we have succeeded in finding a planar presentation of G_2. Similarly to what we did in Section 8.1, we will now use this presentations and apply Theorem 8.2 to obtain a planar presentation of G.

In order to apply Theorem 8.2 we will, similarly to the proof of Theorem 8.3, let $\mathcal{X}$ be the set of all images of the graph G'_2 defined above by colour-automorphisms of G. Again, we define an auxiliary tree T with vertex set $\mathcal{X}$, this time joining two vertices with an edge whenever they share a tidy cycle. Note that (iii) of Theorem 8.2 is satisfied, this time F_i being a tidy cycle. We let again $\mathcal{F}$ be the set of cycles of G induced by the relators in the presentation of Γ_2 obtained above after replacing the auxiliary letters c, d and a^* by the corresponding words. All these cycles are contained in G_2'. Moreover, $\mathcal{F}$ generates $\mathcal{C}_f(G_2')$ by Lemma 2.1. Thus, defining $\mathcal{F}_H$ to be the image of $\mathcal{F}$ under the automorphism of G that maps G_2' to its copy $H \in \mathcal{X}$, we meet the requirement that $\mathcal{F}_H$ generate $\mathcal{C}_f(H)$, and Theorem 8.2 yields that $\bigcup_{H\in\mathcal{X}} \mathcal{F}_H$ generates $\mathcal{C}_f(G)$. Using the second sentence of Lemma 2.1 and the definition of $\mathcal{F}$ we thus obtain presentations of Γ as follows.

If b reverses spin and G_2 is 3-connected then the presentation $G_2 \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (cd)^m, (db)^p\right\rangle$ obtained above translates into $G \cong Cay\left\langle a, b \mid b^2, (baba^{-1})^n, (a^2ba^{-2}b)^m, (baba^{-1})^p\right\rangle$. Note however that the second and fourth relations both induce the tidy shortest dividing cycles of G, and so $n = p$ and the two relations coincide, so one of them can be dropped. If b reverses spin and G_2 is not 3-connected then similarly we obtain $G \cong Cay\left\langle a, b \mid b^2, a^2ba^{-2}b, (baba^{-1})^n\right\rangle$, where the relation $a^2ba^{-2}b$ is tantamount to $c = d$, and the relation c^2 was dropped as it translates into the trivial $a^{-1}baa^{-1}ba = 1$.

If b preserves spin, then the presentations $G_2 \cong Cay\left\langle a^*, b \mid b^2, a^{*2}, (a^*b)^m\right\rangle$ and $G_2 \cong Cay\left\langle a^*, b \mid b^2, a^{*n}, (a^*b)^m\right\rangle$ obtained above easily translate into $G \cong Cay\left\langle a, b \mid b^2, (a^2b)^2, (abab)^m\right\rangle$ and $G \cong Cay\left\langle a, b \mid b^2, (a^2b)^n, (abab)^m\right\rangle$ respectively (note that $(aba)^n = (a^2b)^n$). We thus have

Theorem 8.5. *Let $G = Cay\left\langle a, b \mid b^2, \ldots\right\rangle$ be a 3-connected planar Cayley graph with more than one end and suppose that a has infinite order. Then a reverses spin. If b reverses spin then either*

1. *$G \cong Cay\left\langle a, b \mid b^2, a^2ba^{-2}b; (baba^{-1})^n\right\rangle$, $n \geq 2$, which is the case when G has hexagonal faces (and G_2 is not 3-connected), or*
2. *$G \cong Cay\left\langle a, b \mid b^2, (a^2ba^{-2}b)^m; (baba^{-1})^n\right\rangle$, $n, m, p \geq 2$.*

If b preserves spin then either

3. *$G \cong Cay\left\langle a, b \mid b^2, (a^2b)^2; (ab)^{2m}\right\rangle$, $m \geq 2$, which is the case when G has hexagonal faces, or*
4. *$G \cong Cay\left\langle a, b \mid b^2, (a^2b)^n; (ab)^{2m}\right\rangle$, $n \geq 3$, $m \geq 2$.*

In all cases, the presentation is planar.

Conversely, each of the above presentations, with parameters chosen in the specified domains, yields a planar, 3-connected Cayley graph with more than one end.

Proof. The forward implication was proved in the above discussion.

For the converse implication we follow the approach the proof of Theorem 8.3. Given a presentation $\mathcal{P}$ as in (i)–(iv), we construct a Cayley graph G as follows.

If $\mathcal{P}$ is of type (ii), then we begin by constructing the auxiliary Cayley graph $G_2 \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (cd)^m, (db)^n\right\rangle$, where the parameters n, m here coincide with those in $\mathcal{P}$. By Theorem 7.6 (iv), G_2 has an embedding σ in which, in particular, the 2-coloured cycles induced by $(bc)^n$ bound faces. Let G be the graph obtained from the twist-squeeze-amalgamation (as defined in the Introduction, recall Figure 2) of G_2 with respect to those cycles.

Similarly, if $\mathcal{P}$ is of type (iv), we let $G_2 \cong Cay\left\langle a^*, b \mid b^2, a^{*n}, (a^*b)^m \right\rangle$, and apply Theorem 7.5 (i) to obtain an embedding σ of G_2 in which the 2-coloured cycles induced by $(a^*b)^n$ bound faces. Let G be the graph obtained from the twist-squeeze-amalgamation of G_2 with respect to those cycles. It follows easily from Sabidussi's Theorem 3.1 that G is a Cayley graph; see [14] for details.

By Lemma 4.2 G is 3-connected in both above cases; this can be shown by arguments similar to those of the proof of Theorem 8.3.

Next we claim that a has infinite order in G. To see this note that, by construction, any a-coloured component of G meets infinitely many of the cycles along which the twist-squeeze-amalgamation took place.

If $\mathcal{P}$ is of type (i) or (iii) instead, then our task is easier. Although we could again follow the same approach, starting with a finite graph G_2, it is simpler to construct G directly as in Figure 15: choose the number of parallel monochromatic double rays to be twice the parameter n or m in $\mathcal{P}$, and direct all edges of every second monochromatic double ray the other way if $\mathcal{P}$ is of type (iii).

This completes the construction of G in all cases. The fact that G has indeed the desired presentation $\mathcal{P}$ now follows from the forward implication which we have already proved. □

It follows easily from the proof of Theorem 8.5 that G has precisely two ends if it is of type (i) or (iii), and it has infinitely many ends if it is of one of the other two types.

We have now completed our analysis of the case where G has 2 generators a, b. Let us remark the following, which is perhaps interesting in view of our discussion in Section 1.3.

Corollary 8.6. *Let $G = Cay\left\langle a, b \mid b^2, \ldots \right\rangle$ be a 3-connected planar Cayley graph. Then every face of G has a finite boundary.*

Proof. If G has only one end then this follows from our results of Section 7 (in fact, this is known and holds no matter what the vertex degree is, see [18]).

If G has more than one end, then recall that all our presentations contained a relator inducing a face-boundary. As a always reverses spin in this case (see (4) and Lemma 8.4), and our embeddings are consistent, it follows easily that any two face-boundaries of G can be mapped to each other by a colour-automorphism. Thus all face-boundaries are induced by that relator, and so they are finite. □

9 The planar multi-ended Cayley graphs generated by 3 involutions

Having already described all planar cubic Cayley graphs on two generators, we proceed to the planar cubic Cayley graphs on three generators. Recall that it only remains to describe those that are 3-connected and multi-ended. We divide them into two subclasses, those that have 2-coloured cycles and those that do not, since different arguments are needed in these two cases.

9.1 Graphs with 2-coloured cycles

We have to distinguish three further subcases, according to how the 2-coloured cycles behave with respect to spin.

9.1.1 2-coloured cycles in two spin-reversing colours

We start this section by pointing out that certain choices of spin behaviour cannot give rise to Cayley graphs of the type we are studying. We will make use of these results in subsequent subsections, where we characterise the choices that do give rise to Cayley graphs.

Lemma 9.1. *Let $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, \ldots \rangle$ be planar, multi-ended and 3-connected. If $(bc)^n = 1$ for some $n \in \mathbb{N}$, and both b, c reverse spin, then G has a dividing cycle.*

Proof. Note that under these assumptions, b, c span finite cycles that bound faces of G, see Figure 16. If the remaining faces are also finite then we are done by Lemma 7.3.

If there is an infinite face F, then it must be incident with one of the bc cycles C, so let $e = zy$ be an edge of C incident with F, and assume without loss of generality that e is coloured c, see Figure 16. Let $v = zb$ be the vertex of $C - y$ adjacent with z. Since G is 3-connected, there is a path P in $G - \{z, v\}$ from y to the neighbour x of z outside C. Then $yPxzy$ is a cycle D. Note that there is a side A of D containing F, and so A is infinite.

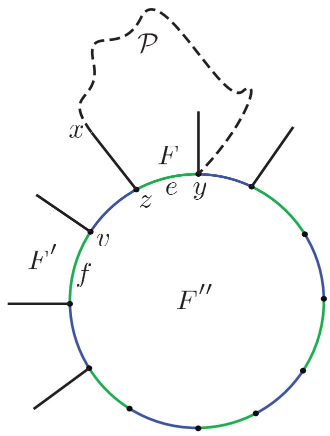

Figure 16: Finding a dividing cycle in the proof of Lemma 9.1.

We would like to show that $G - A$ is also infinite, which would mean that D is dividing. To show this, consider the colour-automorphism g of G mapping z to v. Then g maps e to some other c edge f of C, and it maps F to some infinite face F' incident with f. Note that D cannot contain f by its construction. Thus F' and the finite face F'' bounded by C lie in the same side of D. But that side cannot coincide with A, because as $e \in E(C)$, C separates F'' from $F \subseteq A$. Thus both sides of D are infinite, since they contain the distinct infinite faces F, F'.

□

Lemma 9.2. *Let $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, \dots \right\rangle$ be planar, multi-ended and 3-connected. Then at least one of the colours, d say, preserves spin. Moreover, if both b, c reverse spin then either $(bd)^n = 1$ or $(cd)^n = 1$ holds for some $n \geq 2$.*

Proof. If at least two of the colours preserve spin then there is nothing to show, so suppose G is such a graph in which two colours reverse spin. We distinguish two cases. If two of the colours, b and c say, span finite cycles, that is, if $(bc)^n = 1$ holds, then we can assume that both b, c reverse spin for otherwise our claim is already proved. We can then apply Lemma 9.1, which yields that G has a dividing cycle. If no two of the colours span finite cycles, then we can imitate Lemma 8.4 (see Figure 9) to prove that G has a dividing cycle. Thus G has a dividing cycle in both cases. Moreover, we can assume in both cases that b and c reverse spin.

Choose a shortest dividing cycle C with a facial subpath F of maximum length among all shortest dividing cycles of G. We distinguish three cases according to how F ends, which are very similar to the cases in Lemma 8.1: either F ends with an edge coloured b or c, or both end-edges of F are coloured d and d reverses spin, or both end-edges of F are coloured d and d preserves spin. In the first two cases it is easy to obtain a shortest dividing cycle C' that crosses C similarly to the first two cases of Lemma 8.1 (see Figure 6): consider the colour-automorphism of G exchanging the endvertices of the last edge e of F. But this crossing contradicts Corollary 6.1. Thus the third case must occur; in particular, we have shown the first part of our claim, asserting that d must preserve spin.

Now suppose that the second part of our claim is false, that is, d-edges are in no 2-coloured cycles. Then C must be 3-coloured. Now let e be the last edge of F, and recall that e is coloured d. Assume without loss of generality that the colour of the edge f of F preceding e is c. Note that the edge h following e on C must be coloured c too, since e is the last edge of a facial path and e preserves spin (Figure 17). We can use this fact to show that C has no subpath P of the form db: for then we could map the d-edge of P to e by a colour-automorphism of G to obtain a translate C' of C that crosses C. Indeed, the image of the b edge g of P would then lie in the side of C containing the face $\mathcal{F}$ incident with F. Moreover, C' would have to leave F with an edge that lies in the other side of C, because C' is not allowed to have a path longer than F incident with $\mathcal{F}$ and $g \in E(C')$ is incident with F (Figure 17). But this kind of crossing contradicts Corollary 6.1, which proves our claim that C has no subpath of the form db.

As C must be 3-coloured, this implies that C has a subpath of the form $dc(bc)^k d$ for some $k \geq 1$. Now exchanging the endvertices of the first c edge of such a subpath by a colour-automorphism of G yields again a crossing. We can now use an argument similar to the the proof of (11) for p even (Figure 10) to obtain a contradiction: choosing C so as to maximise the length of a subpath $P \subseteq C$ of the form $dc(bc)^k d$, and considering a crossing as above, (3) implies the existence of a shortest dividing cycle with a longer path of this kind. This completes the proof of the second part of our claim. □

With Lemma 9.2 we immediately obtain

Corollary 9.3. *There is no planar 3-connected multi-ended Cayley graph of the form $Cay\left\langle b, c, d \mid b^2, c^2, d^2, \dots \right\rangle$ in which all edges reverse spin.*

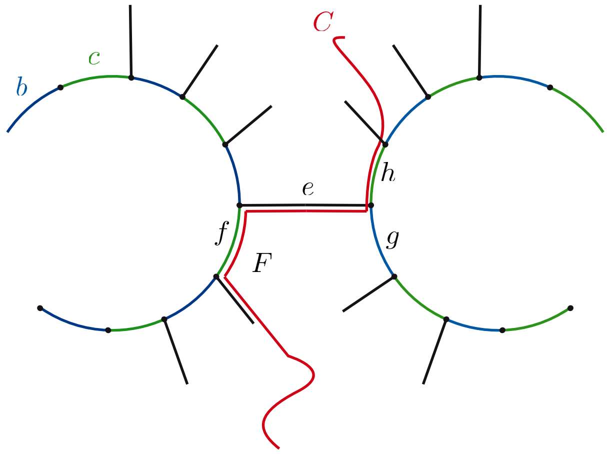

Figure 17: The facial path F in the proof of Lemma 9.2.

9.1.2 2-coloured cycles in two spin-preserving colours

In this section we study the planar 3-connected Cayley graphs of the form $G = Cay\langle b, c, d \mid b^2, c^2, d^2, \dots \rangle$ that have a cycle induced by a relation of the form $(bc)^n$ and both b and c preserve spin.

It is not hard to check that a graph of this kind must be infinite, because every such cycle has a translate of itself in each of its sides, and, by the same argument, every such cycle is dividing (Figure 18); in particular, G is multi-ended.

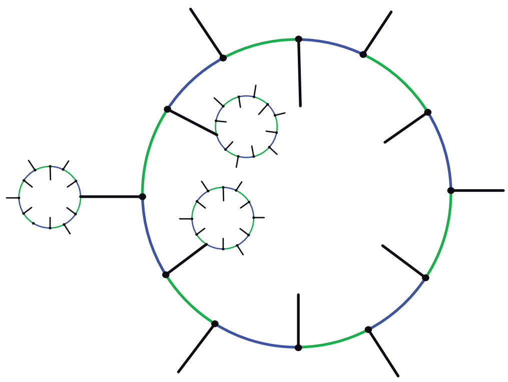

Figure 18: A 2-coloured cycle with spin-preserving edges, and some of its translates.

We will again consider a subgroup Γ_2 of $\Gamma(G)$ as we did in Sections 8.1 and 8.2.2: this time we let Γ_2 be the subgroup of $\Gamma(G)$ generated by $bc, (bc)^{-1}, d$. Let $a := bc$, and note that $a^n = 1$. Again, we will define the auxiliary subgraph $G_2' \subseteq G$ and use it to obtain a Cayley graph G_2 of Γ_2 with an embedding induced by that of G using a construction and arguments very similar to those of Section 8.1.

To begin with, note that if x, y are two elements of Γ_2, then there is an x–y path P in G the b and c edges of which can be decomposed into incident pairs. Thus, since b and c preserve spin, whenever such a path P meets a bc cycle C of G, the two edges of P incident with C lie in the same side of C (Figure 18). In other words, P cannot cross any bc cycle C of G. Now given the embedding σ of G, we can modify G and σ to obtain a Cayley graph G_2 of Γ_2, with respect to to the generating set $\{a, d\}$, and an embedding σ_2 of G_2 as follows. For every bc cycle C of G that contains a vertex in Γ_2, delete all vertices and edges in the side of C that does not meet Γ_2. Let G_2' be the graph obtained after doing so for every such cycle. Then, suppress all vertices of G_2' that now have degree two; that is, replace any bc path xPy of length two whose middle vertex now has no incident d-edge by a single x-y edge, directed the same way as P and bearing the colour a, to obtain the Cayley graph G_2 of Γ_2.

Using Lemma 4.1 it is easy to see that G_2 is 2-connected since G was.

We are now in the fortunate situation of having obtained a Cayley graph of a type that we have already handled: G_2 is generated by two elements, and has monochromatic cycles induced by the relation a^n. Moreover, by the construction of the embedding σ_2, the a-edges preserve spin. Thus we can apply Lemma 8.1, which yields that G_2 has at most one end. So G_2 is one of the graphs in Theorem 7.5, and as a preserves spin cases (iii) and (iv) can be eliminated. The degenerate case (v) of Theorem 7.5 cannot occur, because it would imply that G is not 3-connected as opposite vertices of a bc cycle would separate in that case.

Having obtained a planar presentation of Γ_2, we can now use the same method as in 8.1, namely to apply Theorem 8.2, to yield a planar presentation of Γ.

Theorem 9.4. *Let $G = Cay\left\langle b, c, d \mid b^2, c^2, d^2, \dots\right\rangle$ be a planar 3-connected Cayley graph with more than one end, and suppose that bc has a finite order n and both b, c preserve spin. If d preserves spin then*

$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bcd)^m; (bc)^n\right\rangle$, $n \geq 3$, $m \geq 2$.

If d reverses spin then

$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bcdcbd)^m; (bc)^n\right\rangle$, $n \geq 3, m \geq 1$.

In both cases, the presentation is planar.

Moreover, G is the Mohar amalgamation of $G_2 \cong Cay\left\langle a, b \mid b^2, a^n, (ab)^m\right\rangle$ or $G_2 \cong Cay\left\langle a, b \mid b^2, a^n, (aba^{-1}b)^m\right\rangle$ with itself along the a coloured cycles.

Conversely, each of these presentations, with parameters chosen in the specified domains, yields a Cayley graph as above.

Proof. For the forward implication we apply Theorem 8.2 as in the previous sections. This time the common cycles giving rise to the edges of the auxiliary tree T on the copies of G_2' are the 2-coloured cycles induced by $(bc)^n$. Recall that we obtained a presentation of G_2 in the above discussion from Theorem 7.5: we have $G_2 \cong Cay\left\langle a, d \mid d^2, a^n, (ad)^m\right\rangle$, $n \geq 3$, $m \geq 2$ if d preserves spin and $G_2 \cong Cay\left\langle a, d \mid d^2, a^n, (ada^{-1}d)^m\right\rangle$, $n \geq 3, m \geq 1$ if d reverses spin, with $a = bc$ in both cases. Applying Theorem 8.2, and replacing a back yields the desired planar presentations.

To prove the converse implication, given one of these presentations of the first type we construct the Mohar amalgamation G of $G_2 := Cay\left\langle a, b \mid b^2, a^n, (ab)^m\right\rangle$ with itself along the a coloured cycles —see Introduction. By Sabidussi's theorem G is a Cayley graph; see [14] for details. Lemma 4.2 yields that G is 3-connected since, by Theorem 7.1, G_2 is 3-connected. We can thus apply the forward implication to prove that the presentation we started with is indeed a presentation of G. If we are given a presentations of the second type instead, then we proceed similarly, except that we now let $G_2 := Cay \left\langle a, b \mid b^2, a^n, (aba^{-1}b)^m \right\rangle$. ☐

9.1.3 2-coloured cycles with mixed spin behaviour

In this section we study the planar, infinite, 3-connected Cayley graphs of the form $G = Cay \left\langle b, c, d \mid b^2, c^2, d^2, \ldots \right\rangle$ that have a cycle of the form $(bc)^r$ and precisely one of b, c preserves spin. Let us assume that b preserves spin while c reverses spin (Figure 19).

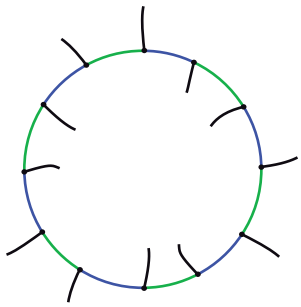

Figure 19: A 2-coloured cycle with mixed spin behaviour.

We proceed similarly to Section 9.1.2: let Γ_2 be the subgroup of $\Gamma(G)$ generated by bcb, c, d. Let $b^* := bcb$, and note that $b^{*2} = 1$ and $(b^*c)^n = 1$ where $n = r/2$ (r must be even). Define the auxiliary subgraph $G_2' \subseteq G$ and use it to obtain a Cayley graph G_2 of Γ_2 with an embedding σ_2 induced by σ similarly to what we did in Section 9.1.2 and Section 8.1: for every bc cycle meeting Γ_2 delete all vertices in its side not meeting Γ_2, then suppress vertices of degree 2. Note that in this case we replace paths of length 3 by edges when suppressing, while in Section 9.1.2 the corresponding paths had length 2.

By the construction of σ_2 the new edges, coloured b^*, reverse spin while the c and d-edges retain their spin behaviour of σ. This means that the new b^*c cycles have all their edges reversing spin. Thus G_2 is the kind of graph we studied in Section 9.1.1 or Section 7 if it is 3-connected. Let us check that this is indeed the case.

Proposition 9.5. *G_2 is 3-connected.*

Proof. Using Lemma 4.1 and the fact that $(b^*c)^n$ is a relation in Γ_2 easily implies that G_2 is 2-connected since G is. Recall that G_2 has 2-coloured cycles of the form $(b^*c)^n$, both b^*, c reverse spin in σ_2, so that every such cycle is a face boundary in σ_2. But by the results of [13], if $\kappa(G_2) = 2$ then G_2 belongs to one of the types (iv), (v), (vi) or (ix) of Theorem 5.1, and of those types, only (vi) and (ix) can have a finite 2-coloured face boundary; see [13, Observation 5.8].

If G_2 belongs to type (ix), which means that it is a finite cycle with some additional parallel edges, then there are three cases to be considered. If $n = 1$, which means that $b^* = c$ and b^*, c span 2-cycles, then any two vertices separating G_2 also separate G, contrary to our assumption that the latter is 3-connected.

If $b^* = d$ or $c = d$ instead, then it is straightforward to check that G must be finite, which we are also assuming is not the case

Thus G_2 belongs to type (vi), and so either $G_2 \cong Cay\left\langle b^*, c, d \mid {b^*}^2, c^2, d^2, (b^*c)^n, (b^*d)^m \right\rangle$, $n, m \geq 2$ or $G_2 \cong Cay\left\langle b^*, c, d \mid {b^*}^2, c^2, d^2, (b^*c)^n, (cd)^m \right\rangle$, $n, m \geq 2$. In the former case the b^* edges are hinges and in the latter the c edges are hinges. We claim that the endvertices of such a hinge also separate G.

For this, let $\{x, y\}$ be the endpoints of a hinge of G_2, let C be the b^*c cycle of G_2 containing x, y, and let K be the component of $G_2 - \{x, y\}$ that does not meet C; such a component exists because $\{x, y\}$ cannot separate C as xy is an edge of C. If $G - \{x, y\}$ has a K–C path P, then P can be modified into a K–C path in $G_2 - \{x, y\}$ as follows: for every bc cycle D met by P that is disjoint from C, note that both endpoints of P lie in the same side of D, for G_2 cannot meet both sides of any such cycle by the construction of G_2. This means that if P enters the side A of D not meeting G_2, then it must exit that side again. Thus, we can replace a subpath of P that has endvertices v, w on D and whose interior lies in A by a v-w subarc of D to obtain a path that does not meet the 'wrong' side A of D. Similarly, if P meets the side of C not containing K, then we replace the part of P in that side by a subarc of C, this time being careful enough to pick that subarc that does not contain x and y. Performing such a modification recursively as long as P meets both sides of a bc cycle D, we modify P into a path P' with the same endvertices that meets at most one of the sides of any bc cycle. It follows easily that P' is a path in G'_2. Moreover, P' does not meet x, y by construction. But then P' contradicts the fact that $\{x, y\}$ separates C from K in G_2. This completes the proof of our claim that G_2 is 3-connected.

□

Now G_2 might be finite or 1-ended, in which case we can use our classification of Section 7, or multi-ended, in which case we can apply Lemma 9.2, which yields that d must preserve spin. Thus, in the case where d reverses spin, we obtain the following classification.

Theorem 9.6. *Let $G = Cay\left\langle b, c, d \mid b^2, c^2, d^2, \ldots \right\rangle$ be an infinite planar 3-connected Cayley graph such that bc has a finite order n and precisely one of b, c preserves spin (b say) and d reverses spin. Then $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (cd)^m, (dbcb)^p; (bc)^{2n} \right\rangle$, $n, m, p \geq 2$. This presentation is planar.*

Moreover, G is a twist-squeeze-amalgamation of $Cay\left\langle b^*, c, d \mid {b^*}^2, c^2, d^2, (b^*c)^n, (cd)^m, (db^*)^p \right\rangle$ with itself.

Conversely, each of these presentations, with parameters chosen in the specified domains, yields a planar 3-connected multi-ended Cayley graph.

Proof. By the above discussion, G_2 cannot be multi-ended since we are assuming that d reverses spin. Since all edges of G_2 reverse spin, we obtain a presentation of G_2 from Theorem 7.6 (iv) or (v), and by Proposition 9.5 we can exclude (v). Thus we are left with $G_2 \cong Cay\left\langle b^*, c, d \mid {b^*}^2, c^2, d^2, (b^*c)^n, (cd)^m, (db^*)^p \right\rangle$, $n, m, p \geq 2$.

For the forward implication we apply Theorem 8.2 as in the last section. This yields that we can obtain a presentation of G by replacing b^* with bcb in the above presentation of G_2, except that we replace the relation ${b^*}^2$ by b^2. It is easy to see that the asserted presentation is planar.

For the converse implication we proceed as in the proof of Theorem 9.4: let $G_2 := Cay \left\langle b^*, c, d \mid b^{*2}, c^2, d^2, (b^*c)^n, (cd)^m, (db^*)^p \right\rangle$ and let G be the twist-squeeze-amalgamation of G_2 with itself along the b^*c cycles. Again G is a Cayley graph by Sabidussi's theorem; see [14] for details. Lemma 4.2 yields that G is 3-connected since, by Theorem 7.1, G_2 is 3-connected. We can thus apply the forward implication to prove that G has the desired presentation. □

It remains to consider the case when d preserves spin. Again we have to distinguish various cases. The most interesting case is when G_2 is 3-connected and multi-ended, and so applying Lemma 9.2 to G_2 we obtain that G_2 has 2-coloured cycles containing d; we then have to distinguish two subcases according to which of b^*, c participates in those cycles. Let us first consider the case when this is c, and so $(cd)^p$ is a relation for some p. Interestingly, we can now apply Theorem 9.6 to G_2 rather than G: recall that G_2 is 3-connected by Proposition 9.5, and that σ_2 has spin behaviour as in the requirements of that theorem, except that the roles of the letters are now interchanged. Thus, substituting d by b and b^* by d we can apply Theorem 9.6 to obtain a planar presentation of G_2, namely

$$G_2 \cong Cay \left\langle d, c, b^* \mid d^2, c^2, {b^*}^2, (dc)^{2n}, (cb^*)^m, (b^*dcd)^p \right\rangle, \; n, m, p \geq 2. \tag{22}$$

In the second subcase, when $(b^*d)^p$ is a relation rather than $(cd)^p$, we can repeat the same arguments to obtain a similar presentation but with c and b^* interchanged. Thus, in this case we have

$$G_2 \cong Cay \left\langle d, c, b^* \mid d^2, c^2, {b^*}^2, (db^*)^{2n}, (b^*c)^m, (cdb^*d)^p \right\rangle, \; n, m, p \geq 2. \tag{23}$$

Note that these two cases are distinct: if the order of b^*d is finite then the order of dc is infinite and vice-versa. Indeed suppose that the order of b^*d is finite, and recall that d preserves spin while b^*, c reverse spin in G_2. We will show that an infinite $cdcdc\ldots$ walk W, starting at an arbitrary vertex, meets infinitely many of the finite b^*d-cycles. Indeed, suppose there is a last b^*d-cycle C met by W. Then W cannot have met a b^*d-cycle in each side of C, for C separates its sides. But W arrived at C along a c edge, traversed a d-edge of C, and left C by another c edge. Now as d preserves spin, those two c edges lie in different sides of C, and are incident with b^*d-cycles other than C. This contradiction proves our claim. The same argument proves the reverse claim.

Thus we have obtained a presentation of G_2 in the case where d preserves spin too, and can now use this to deduce a presentation of G.

Theorem 9.7. *Let $G = Cay \left\langle b, c, d \mid b^2, c^2, d^2, \ldots \right\rangle$ be a planar 3-connected Cayley graph and suppose that bc has a finite order n, precisely one of b, c preserves spin (b say), and d preserves spin. Then precisely one of the following is the case:*

1. $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bcbdcd)^m; (bc)^{2n} \right\rangle, n \geq 2, m \geq 1.$
2. $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bcbdcd)^p; (dc)^{2n}, (bc)^{2m} \right\rangle, \; n, m, p \geq 2;$
3. $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bcbdcd)^p; (dbcb)^{2n}, (bc)^{2m} \right\rangle, \; n, m, p \geq 2.$

These presentations are planar.

Conversely, each of the above presentations, with parameters chosen in the specified domains, gives rise to a planar 3-connected Cayley graph as above.

Proof. The Cayley graph G_2 (as defined in the beginning of this section) is 3-connected by Proposition 9.5. We distinguish two cases.

Case I: G_2 has at most 1 end.

Recall that both b^*, c reverse spin while d preserves spin in G_2. Thus we are in type (iii) of Theorem 7.6, and so $G_2 \cong Cay\left\langle b^*, c, d \mid b^{*2}, c^2, d^2, (b^*c)^n, (b^*dcd)^m \right\rangle$, $n \geq 2, m \geq 1$. Applying Theorem 8.2 again we obtain

$$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bcbc)^n, (bcbdcd)^m \right\rangle, n \geq 2, m \geq 1,$$

and rewriting $(bcbc)^n$ as $(bc)^{2n}$ we obtain possibility (i) of the statement. Clearly, this is a planar presentation.

Case II: G_2 is multi-ended.

In this case we have already obtained a presentation for G_2 in the above discussion; see (22) and (23). By the same technique as in the first case, and a little bit of rearranging (a relation of the form $(WZ)^n$ is equivalent to $(ZW)^n$), we obtain the claimed presentations (ii) and (iii). It is easy to see that these presentations are planar using the spin behaviour; see Figure 19.

The converse implication can be established as in the proof of Theorem 9.6, by explicitly constructing G as a twist-squeeze-amalgamation of the corresponding G_2. This time we have to apply the converse implication of Theorem 9.6 to obtain the desired G_2. □

We observe the following fact, which follows from Corollary 9.3 and Theorems 9.4, 9.6 and 9.7, and is interesting in view of the forthcoming counterexamples to Conjecture 1.2.

Corollary 9.8. *Let $G = Cay\left\langle b, c, d \mid b^2, c^2, d^2, \ldots \right\rangle$ be a planar 3-connected Cayley graph containing a 2-coloured cycle. Then every face of G is finite.*

9.2 Graphs without 2-coloured cycles

In this section we consider the cubic multi-ended planar Cayley graphs $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, \ldots \right\rangle$ that have no 2-coloured cycles. Similarly to our analysis of the graphs on three generators that do have 2-coloured cycles (Section 9.1), we will have to distinguish cases according to the spin behaviour of the generators.

Recall that our analysis of the graphs with 2-coloured cycles in Section 9.1 was very intimately connected with those cycles: in all non-trivial cases, we used such cycles to split G by finding a subgraph G_2' in which those cycles bound faces. In the current case, the absence of 2-coloured cycles makes our task harder. However, we will still be able to use similar methods. We will be able to find a good substitute for the 2-coloured cycles: namely, the minimal dividing cycles.

9.2.1 All edges preserve spin

In this section we consider a 2-connected multi-ended Cayley graph $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, \ldots \right\rangle$ with no 2-coloured cycles, that has a consistent embedding σ in which all edges preserve spin. We do not demand that G be 3-connected here, because the results of this section are needed for the characterization of graphs of connectivity 2 in [13].

Before we can state our main result of this section we need to define the concept of a non-crossing pattern. Intuitively, a non-crossing pattern is a finite word $\mathcal{P}$ in the letters b, c, d such that whenever $\mathcal{P}$ is a relation of a planar cubic Cayley graph G all vertices of which have the same spin, no cycle of G induced by $\mathcal{P}$ bounds a face, and no two cycles of G induced by $\mathcal{P}$ cross. Making a formal definition out of this intuitive idea is a bit tricky in the absence of a concrete Cayley graph$G$.

For this, let H be a plane graph the edges of which are coloured with the colours b, c, d in such a way that no vertex is incident with more than one edge of the same colour, and let $C \subseteq H$ be a cycle in H. We will say that C *complies* with $\mathcal{P}$ (in H), if one of the words obtained by reading the colours of the edges of C, as we cycle once along C once in a straight manner, is $\mathcal{P}$ and moreover all vertices of C that have degree 3 in H have the same spin. Given two cycles C, R in H both complying with $\mathcal{P}$, we will say that R is a *rotation* of C if $R \cap C$ is a (possibly closed) path. The intuition of this definition is derived from the fact that H can be though of as a Cayley graph in which C, R are induced, starting at the same vertex, by relations that are obtained from each other by rotating the letters.

We can now give the formal definition of a non-crossing pattern.

Definition 9.9. *A non-empty word $\mathcal{P}$ in the letters b, c, d is called a* non-crossing pattern *if it satisfies the following conditions:*

1. *$\mathcal{P}$ contains all three letters b, c, d;*
2. *$\mathcal{P}$contains no consecutive identical letters;*
3. *$\mathcal{P}$ is not of the form $(bcd)^n$ up to rotation and inversion, and*
4. *if C is a cycle complying with $\mathcal{P}$ then no rotation of C crosses D.*

This definition might look somewhat abstract at first sight, but in fact there is an easy algorithm that recognises non-crossing patterns.

It will be easier to understand the necessity of the requirements of Definition 9.9 if one considers the usage of non-crossing patterns in the following theorem: we impose (i) because b, c, d are involutions. With (ii) we prevent 2-coloured cycles, which we have handled in earlier sections. We require (iii) to prevent $\mathcal{P}$ from being a face boundary. Finally, (iv) is the important property of $\mathcal{P}$ from which our results yield their strength.

We can now state the main result of this section, yielding a complete description the corresponding Cayley graphs.

Theorem 9.10. *Let $G = Cay\left\langle b, c, d \mid b^2, c^2, d^2, \ldots \right\rangle$ be a 2-connected multi-ended Cayley graph. Suppose G has a consistent embedding in which all edges preserve spin, and that each of the elements bc, cd, db has infinite order. Then precisely one of the following is the case:*

1. *$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bcd)^2; (bcdc)^n \right\rangle, n \geq 2$ (faces of size 6);*
2. *$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bcd)^k; \mathcal{P} \right\rangle$, $k \geq 3$ (faces of size $3k \geq 9$);*
3. *$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2; \mathcal{P} \right\rangle$, (no finite faces),*

where $\mathcal{P}$ is a non-crossing pattern.

Conversely, for every n or k in the specified domains, and every non-crossing pattern $\mathcal{P}$, the above presentations yield a planar Cayley graph as above.

In the first 2 cases G is always 3-connected. In case (iii) $\kappa(G) = 2$ if $\mathcal{P}$ is regular and $\kappa(G) = 3$ otherwise. If $\kappa(G) = 2$ then G has a hinge if and only if $\mathcal{P}$ is strongly regular.

The above presentations are less explicit than the presentations we have obtained so far because of the presence of $\mathcal{P}$. This is not a shortcoming of our analysis: there is no word with arithmetic parameters capturing all non-crossing patterns, but all non-crossing patterns are needed to make Theorem 9.10 true. However, as we will see in the forthcoming proof, non-crossing patterns have a rather simple structure and they are similar to each other. Since there is an algorithm that recognises them, the set of non-crossing patterns, and thus the set of Cayley graphs described in Theorem 9.10, can be effectively enumerated.

The rest of this section is devoted to the proof of Theorem 9.10, which is completed in page 65. The reader who does not wish to see the details yet could skip the rest of this section, as well as Section 9.2.3 which is similar, and continue with Section 10 in page 79.

Faces and dividing cycles

Let G be a graph as in Theorem 9.10 fixed throughout this section and let Γ be its group. Note that as we are assuming that all colours preserve spin,

$$\text{every facial walk is of the form } \ldots bcdbcd \ldots \text{ or the inverse (Figure 20).} \quad (24)$$

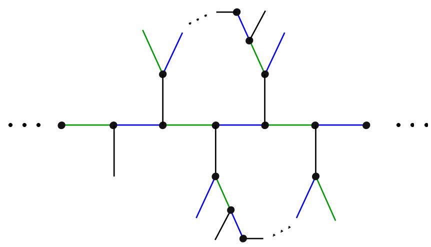

Figure 20: The local situation around a vertex in the case that all edges preserve spin and no 2-coloured cycles exist.

This means that any two face boundaries look locally the same, but in fact more is true: it is easy to check that

Any two face boundaries of σ can be mapped to each other by a colour-automorphism of G. (25)

Indeed, recall that σ is a consistent embedding in which all edges preserve spin. Thus, for every vertex x, the three colour-automorphisms of G exchanging x with its neighbours can be used to map any of the three face boundaries incident with x to each other.

In particular, all faces in σ have the same size, which by (24) is a multiple of 3. It cannot be equal to 3 though, for this would mean that $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, bcd \right\rangle$ and this graph has hinges, which we are assuming is not the case for G. Thus we have proved that

For some $n \in \mathbb{N} \cup \{\infty\}$, every face of σ has size $N = 3n + 6$. (26)

As already mentioned, our analysis of the graphs without 2-coloured cycles will be based on their shortest dividing cycles. We begin by showing that they do exist:

G has a dividing cycle. (27)

Indeed, if all faces in σ are finite, then this follows immediately from Lemma 7.3. If there is an infinite face in σ, then since σ is consistent, every vertex is incident with an infinite face by Lemma 3.2. Let F be an infinite face of G, let $vexe'w$ be a facial walk incident with F comprising two edges e, e', and let f be the third edge incident with x. Since we are assuming that G has no hinge, there is a w–v path P in G that avoids both endvertices x, x' of f. We claim that the cycle $C = vxwPv$ is dividing. Indeed, by its construction C separates F from x', and as F is infinite and x' must be incident with an infinite face too by the above remark, both sides of C contain infinitely many vertices. This completes the proof of (27).

So let C be a shortest dividing cycle of G. Our next claim is that

any maximal 2-coloured subpath of C containing at least 3 edges contains an odd number of edges. (28)

This can be proved by an argument very similar to the one we used in the proof of (11) for the case when p is even.

Using this we are now going to prove that

C has no facial subpath containing more than 3 edges. (29)

To see this, let F be a longest facial subpath of C, and suppose to the contrary that $||F|| > 3$. We may assume without loss of generality that no other shortest dividing cycle C' has a facial subpath longer than $||F||$, for otherwise we could have chosen C' instead of C.

Recall that, by (24), every facial walk is of the form $\dots bcdbcd \dots$ or the inverse (Figure 20). Since all colours behave the same way in this case, we may assume without loss of generality that F starts with a subpath F_4 of the form $bcdb$. Let $\mathcal{F}$ be the face incident with F. Consider the colour-automorphism g of G mapping the fourth vertex y of F_4 to its first vertex x, and let $C' := gC$; see Figure 21. It is easy to check that C' crosses C: indeed, note that C' contains an edge e incident with both x and $\mathcal{F}$, and so C' must leave $\mathcal{F}$ before C as no shortest dividing cycle can have a longer subpath incident with $\mathcal{F}$ than F. This crossing however contradicts Corollary 6.1. This contradiction proves (29).

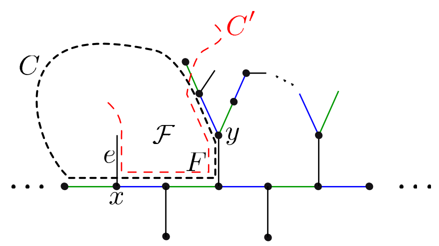

Figure 21:

The dominant colour.

Our next assertion shows that even though in this case G has no 2-coloured cycle, it must have cycles that are not far from being 2-coloured:

There is a colour $a \in \{bcd\}$ such that for every shortest dividing cycle C of G, every other edge of C is coloured a. In particular, $|C|$ is even. (30)

Let P be a maximal 2-coloured subpath of C. Obviously, P contains at least two edges. Moreover, P cannot consist of precisely two edges, because then the subpath P' of C comprising P and its two incident edges would be facial by (24), and this would contradict (29) since $||P'|| = 4$. Thus, $||P||$ is at least 3, and it is odd by (28). This means that the first edge of P has the same colour a as its last edge. Assume without loss of generality that $a = c$.

The path P is a good starting point in our attempt to prove (30): every other edge of P is coloured c since it is 2-coloured. And indeed, we will be able to extend it by adding further 2-coloured subpaths of C, retaining the property that every other edge is coloured c, until exhausting all of C. For this, let e', e be the last two edges of P and let f, f' be the two edges of C succeeding P, appearing in C in that order (Figure 22). Note that as P was chosen to be maximally 2-coloured, e' and f have different colours. This, combined with (29), implies that f' must be coloured c, for otherwise the subpath of C spanned by e', e, f, f' is facial.

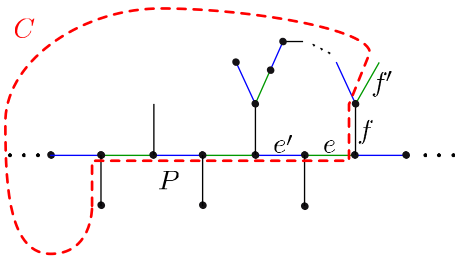

Figure 22: The situation around P in the proof of (30).

Thus e, f, f' span a 2-coloured subpath of C. Let P' be the maximal 2-coloured subpath of C containing these three edges. By (28) $||P'||$ is odd. Moreover, it does not contain e' as f and e' have different colours. Thus P' starts and ends with a c-edge. Consider the path $P \cup P'$, and note that every other edge of this path is coloured c. Now starting with this path instead of P and repeating the above arguments, we find a longer odd subpath of C every other edge of which is coloured c. Continuing like this we prove that every other edge of C bears the same colour. Let us call this colour the *dominant* colour of C. Note that the dominant colour of a cycle is always unique in the current case, since no 2-coloured cycles exist.

We just proved that every shortest dividing cycle of G has a dominant colour. It remains to prove that they all have the same dominant colour. So suppose that C, D are shortest dividing cycles of G with distinct dominant colours c, d respectively, say. Since all cycles are 3-coloured, both C, D contain a b edge, and we may assume that this b edge is the same edge e in both cases, for otherwise we could have considered translates of C, D through e. Now e is surrounded by two c edges in C, and it is surrounded by two d-edges in D. As e preserves spin, this means that C and D cross each other at e. By the remark preceding Corollary 6.1, we can obtain a new shortest dividing cycle C' by combining C and D, following one of them up to e and then switching to the other. But C' then contains a *cbd* subpath (with the edge e in the middle), and can be chosen so that it also contains a *cdc* subpath of C, contradicting the fact that every other edge of C' must bear the same colour. This contradiction proves that the dominant colour is the same for every shortest dividing cycle of G indeed, and completes the proof of (30).

From now on we assume that the dominant colour of the shortest dividing cycles of G is c.

Similarly to the case when G has two generators, we will have to consider the case when G has hexagonal faces separately (see Section 8.2.1). Recall that by (26) every face of σ has size at least 6.

If G has no face of size 6 then no two shortest dividing cycles of G cross. (31)

Indeed, by our discussion in Section 6, in particular by (2), such a crossing gives rise to a finite region bounded by a cycle in $C \cup D$. As this region contains only faces of size larger than 6, we obtain a contradiction to Euler's formula (17) in a way similar to the proof of (16). This proves (31).

The hexagonal grid case.

Let us now consider the case when one, and thus by (26) all, of the faces in σ have size 6. We proceed as in Section 8.2.1 to prove that

If G has a face of size 6 then it has a shortest dividing cycle induced by $(bcdc)^n$. (32)

Indeed, similarly to (19), we can prove that every shortest dividing cycle C contains the same amount of edges from each non-dominant colour b, d. And again, rerouteing C around some of its incident hexagons if needed, in other words, replacing subarcs of the form *bcd* by *dcb*, which we are allowed to do since $bcd(dcb)^{-1}$ is a relation (inducing a hexagonal face), we modify C into a cycle of the same length that is induced by a word of the form $(bcdc)^n$ and is still dividing. This completes the proof of (32).

Note that translates of such a cycle cannot cross, and so we have obtained something similar to (31) for graphs with hexagonal faces.

We already have enough information to finish off the case when the faces of G are hexagonal. It is now not hard to check that the relation $(bcdc)^n$ we just obtained combined with the one inducing the face boundaries, and of course the involution relations for the generators, yield a planar presentation of Γ:

$$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bcd)^2; (bcdc)^n \right\rangle, n \geq 2.$$

This can be proved for example by showing that the underlying graph is isomorphic to one of the graphs in case (i) or (iii) of Theorem 8.5; to see this, look at G through a lens that identifies colours b and d. The details are left to the reader.

Note that if we let $n = 1$ in the above presentation we would obtain a graph in which c edges are hinges, contradicting our assumptions on G.

Conversely, for every $n \geq 2$ one can show that the above presentation corresponds to a planar 3-connected Cayley graph with hexagonal faces by explicitly constructing such a graph: it consists of $2n$ 'parallel' double rays coloured b, d in an alternating fashion, joined by edges coloured c.

Back to the main case: no hexagons.

For the rest of this section we will be assuming that $N > 6$.

The weak colour.

We now turn our attention to the behaviour of the non-dominant colours in the shortest dividing cycles of G: it turns out that, in general, one of them is more 'dominant' than the other. More precisely:

> Suppose that some shortest dividing cycle of G has a bcb subpath. Then no shortest dividing cycle of G has a dcd subpath. (33)

To see this, suppose that some shortest dividing cycle C of G has a bcb subpath and some shortest dividing cycle D of G has a dcd subpath. We may assume without loss of generality that these two subpaths traverse the same c-edge e, for otherwise we can consider a translate of one of the two. As e preserves spin, C and D cross each other at e, contradicting (31).

Proposition (33) implies that at least one of the non-dominant colours $a \in \{b, d\}$ cannot appear too often in any shortest dividing cycle: between any two a-coloured edges in any shortest dividing cycle, there are edges of both other colours. We call this colour the *weak* colour of G, and we call the other non-dominant colour the *semi-dominant* colour of G, unless both non-dominant colours appear with the same frequency, in an alternating fashion, in which case both non-dominant colours are called *weak*. Note that if this is the case, then every shortest dividing cycle is induced by a word of the form $(bcdc)^n$.

The argument of the proof of (33) can be repeated to prove something stronger:

> Suppose that some shortest dividing cycle of G has a bWb subpath, where W is any word in the letters b, c, d. Then no shortest dividing cycle of G has a dWd subpath. (34)

Indeed, such a word W must start and end with the dominant colour c, and it is straightforward to check, using the fact that c preserves spin, that if some shortest dividing cycle C had a bWb subpath and some shortest dividing cycle D had a dWd subpath then some translate of D would cross C, contradicting (31).

Maximal common subpaths and the word Z.

We have already seen that two shortest dividing cycles cannot cross each other. However, they may be tangent, that is, they may have a common subpath provided one is contained in the other. Our next assertion restricts the possible common subpaths.

Let P be a maximal common subpath of two distinct shortest dividing cycles C, D of G, and let P' be another maximal common subpath of two distinct shortest dividing cycles C', D'. Let W_P, W'_P be words inducing P, P' respectively. Then one of these words is a *prefix*, i.e. an initial subword, of the other. (35)

The word 'maximal' here is meant with respect to inclusion: P is a maximal common subpath of C, D if $P \subset C, D$ and for every path Q properly containing P either $Q \not\subset C$ or $Q \not\subset D$ holds.

This proposition implies that there is a unique word $Z = Z(G)$, namely the maximal word for which there are two distinct shortest dividing cycles of G and a common subpath of theirs induced by this word, that governs all the possible intersections of shortest dividing cycles: every maximal common subpath of two shortest dividing cycles is induced by a prefix of Z (or Z itself).

To prove (35), let P, C, D, P', C', D' be as in the assertion, and suppose that none of W_P, W'_P is a prefix of the other. We may assume without loss of generality that P and P' have the same initial vertex x, for otherwise we can translate C', D' by a colour-automorphism of G to achieve this (Figure 23).

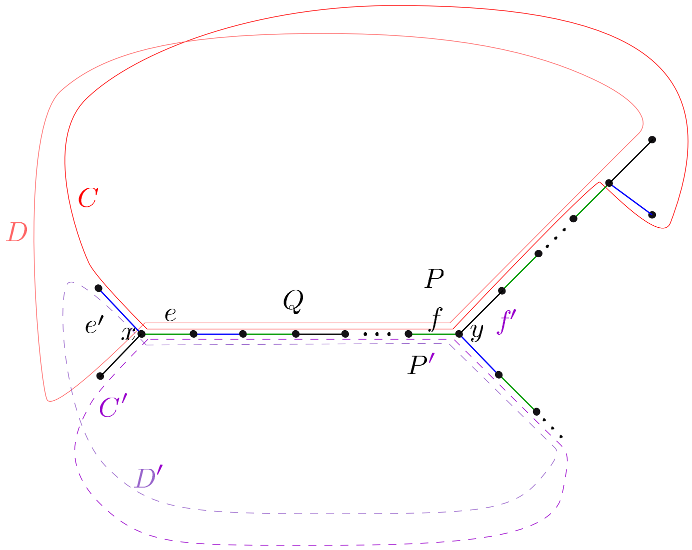

Figure 23: The hypothetical situation when two maximal common subpaths of shortest dividing cycles are not a subpath of each other.

Easily, both P and P' must begin (and end) with an edge e, incident with x, of the dominant colour c. Moreover, one member of each of the pairs C, D and C', D' contains the b edge incident with x and the other member contains the d-edge, for P and P' are maximal common subpaths. Now since we are assuming that neither of W_P, W'_P is a prefix of the other, there is a common vertex $y \in P \cap P'$ after which the two paths split for the first time. Again, the common edge f leading into y must bear the dominant colour c, and so one of the paths P, P' follows the b edge and the other follows the d-edge incident with

y (Figure 23). Consider the common subpath $Q := xPy = xP'y$, and note that Q is a subpath of all four cycles C, D, C', D'. Moreover, for each of the four ways to choose a non-dominant edge e' incident with x and a non-dominant edge f' incident with y, one of these cycles contains $e'Qf'$ as a subpath. In particular, if W_Q is a word inducing Q, then both the words bW_Qb and dW_Qd are induce subpaths of shortest dividing cycles. But this contradicts (34), and so (35) is proved.

Proposition (35) implies that $Z(G)$ is symmetric. Indeed, Z^{-1} also induces a common subpath of shortest dividing cycles, so it must be a prefix of Z. This implies that the element of Γ described by Z is an involution.

Next, we claim that

> For every path P in G induced by Z there are precisely two shortest dividing cycles C, D containing P. Moreover, $C \cap D = P$, that is, C and D have no common edge outside P. (36)

Indeed, by the definition of Z and P there are at least two shortest dividing cycles C, D containing P. If there was a third one C' then, as G is cubic, C' would have an edge f incident with P in common with one of C, D. But this would contradict the maximality of Z as $f \cup P$ would then be a common subpath.

If C and D have a further common edge e outside P, then by the above argument e is not incident with P. This means that we can choose subpaths C', D' of C, D such that $C' \cup D'$ is a cycle shorter that C. Thus this cycle cannot be dividing, and so one of its sides contains only finitely many vertices. This leads to a contradiction as in the proof of (31).

Face incidences and the word A.

Having analysed the ways that shortest dividing cycles can intersect each other let us now see how a shortest dividing cycle can intersect a face boundary. It turns out that there are shortest dividing cycles nicely arranged around each face boundary F, each of them using precisely three consecutive edges of F. Recall that by (24) every facial walk is of the form $\ldots bcdbcd \ldots$ or the inverse.

> For every face boundary F and every bcd subpath P of F there is a shortest dividing cycle C containing P and no other edge of F. (37)

Indeed, the fact that every cycle is 3-coloured and (30) imply together that every shortest dividing cycle contains a bcd subpath. Translating this subpath to P we thus obtain a shortest dividing cycle C containing P. By (29) C contains none of the two edges of F incident with P. It remains to check that C can also not contain an edge e of F not incident with P. But if this was the case then the colour-automorphism g of G mapping the first vertex of P to its last vertex would translate C to a cycle $C' = gC$ that crosses C: indeed, note that g fixes F as F is a concatenation of translates of P, and so C' must also contain a further edge $e' = ge$ on F. An easy topological argument now shows that C and C' cross indeed, contradicting (31). This proves (37).

This motivates us to define the word $A := bcd$, which will play an important role in the sequel.

The word $(AZ)^n$ inducing the shortest dividing cycles.

It turns out that there is a word inducing all the shortest dividing cycles of G:

> For some $n \geq 2$, every shortest dividing cycle of G is induced by the word $(AZ)^n$. In particular, any two shortest dividing cycles can be mapped to each other by a colour-automorphism of G. (38)

To prove this, pick a pair C, C' of shortest dividing cycles that have a common path P induced by Z, and let x be an endvertex of P. Let F be the face-boundary containing the b and d-edge incident with x. We claim that each of C, C' contains a bcd (or dcb) subpath of F incident with x; in other words, C and C' are as in (37) (Figure 24).

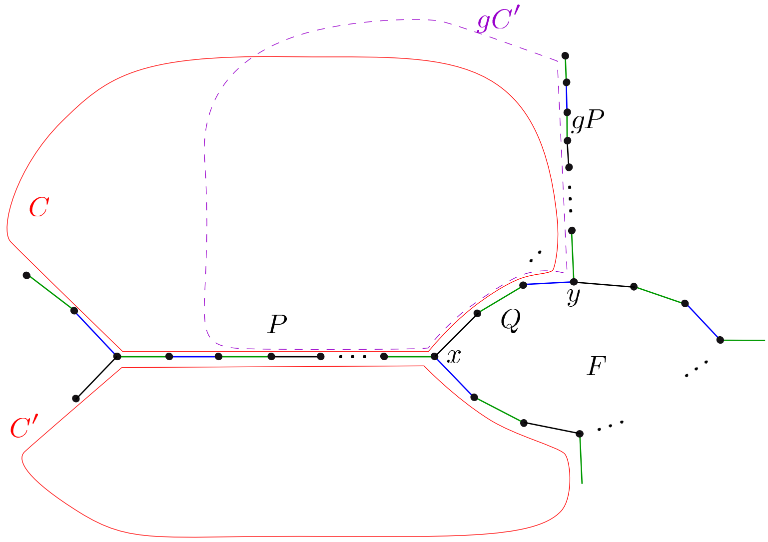

Figure 24: The situation arising in the proof of (38).

To begin with, note that as x was by definition a last common vertex of C and C', one of them, C say, contains the b edge incident with x and the other contains the d-edge. By (29) C cannot have a subpath on F containing more than three edges. So it suffices to show that C does not leave F after having traversed less than three edges. If this is the case though, then the shortest dividing cycle D containing the bcd subpath Q of F starting at x, which cycle is provided by (38), is distinct from C. Thus, D cannot contain all of P because then it would have a common subpath with C properly containing P and this would contradict the choice of P. So D leaves P at some of its interior vertices, which means that D enters a side of one of C, C' not containing F. On the other hand, D also meets the side of each of C, C' that does contain F; see (Figure 24). This means that D crosses one of C, C', contradicting (31).

This proves that, as claimed, C contains the bcd subpath of F incident with x. By the same argument we can prove that C' contains the dcb subpath of F incident with x, but we will not need this. This is a good start for the proof of (38): we just proved that C contains an AZ subpath.

Now consider the colour-automorphism g of G mapping x to the other end-vertex $y = xbcd$ of Q. Repeating the above arguments we see that one of

gC, gC', in fact it must be gC', also contains Q. We claim that one of C, gC' must contain all three paths P, Q, gP. For suppose that C misses part of gP and gC' misses part of P. Then, as C is not allowed to cross gC by (31), it leaves gP entering the side of gC' not containing F. Similarly, gC' leaves P entering the side of C not containing F. But if each of C, gC' meets the side of the other not containing a face then the two cycles must cross each other contradicting (31). This proves our claim that one of C, gC' must contain all three paths P, Q, gP. Note that by the definition of Z and the fact that both these cycles contain Q, this immediately implies that $C = gC'$.

We just proved that C has a $ZAZ^{-1} = ZAZ$ subpath. Now repeating the previous arguments at the other end of P we prove that C has a $AZAZ$ or $A^{-1}ZAZ$ subpath. The latter possibility can however not occur, for it would mean that C crosses C' when leaving P; this can be seen by observing the spin of the end-vertices of P. Thus C has a $(AZ)^2$ subpath. Continuing like this, we prove that C is induced by $(AZ)^n$. Moreover, $n \neq 1$ for otherwise we would easily obtain, with the above arguments, that the faces have size 6 which we are assuming is not the case.

By the same arguments, we can prove that C' is induced by $(A^{-1}Z)^n$, but as $Z = Z^{-1}$ this means that C' is induced by $(AZ)^n$ too.

We started with C, C' being arbitrary shortest dividing cycles having a subpath induced by Z. Thus if we could prove that every shortest dividing cycle of G has such a subpath this would complete the proof of (38). This is indeed the case. For let D be a shortest dividing cycle and let R be a maximum-length subpath of D shared with another shortest dividing cycle $D' \neq D$. Easily, R has at least one edge. By (35), R is an initial subpath of some path P induced by Z which is a maximal common subpath of two shortest dividing cycles C, C'. Now if $|R| \neq |P|$ then $D \neq C, C'$. But D shares an edge e not contained in but incident with P with one of C, C', for these two cycles together use all edges incident with the first vertex of P. This means that a common subpath of D and one of C, C' is eR, contradicting the maximality of R. This proves that every shortest dividing cycle of G has a subpath induced by Z and completes the proof of (38).

The subgroup Γ_2, societies, and the subgrpaph G'_2.

Let a, z be the elements of Γ corresponding to the words A, Z respectively, and let Γ_2 be the subgroup of Γ spanned by a, b. This subgroup, and the above results relating A and Z to shortest dividing cycles, will allow us to follow an approach similar to that of the previous sections, the shortest dividing cycles now playing the role of the monochromatic cycles of Section 8.1 or the 2-coloured cycles of Section 9.1.

As in those sections, we are going to show that G is a union of subdivisions of isomorphic copies of the Cayley graph G_2 of Γ_2 with respect to the generating set $\{a, z\}$. We would like to define this kind of subdivision G'_2 of G_2 similarly to previous sections, by deleting for each shortest dividing cycle C meeting Γ_2 all vertices in one of the sides of C, and then suppress vertices of degree 2 to obtain G_2 from G'_2. However, things are more complicated now and we need some preparatory work before we can show that this operation yields G_2 indeed. They reader may choose to skip this preparatory work and continue reading after (41), perhaps after having a look at the following vital definition.

Definition 9.11. *Let C be a shortest dividing cycle of G, and let x be a vertex of C such that the word $(AZ)^n$ induces C if the starting vertex is x. Then, we call the set of vertices of C that can be reached from x by subarc of C induced by a prefix of $(AZ)^n$ a* society *of C.*

Note that, by (38), every shortest dividing cycle has at least one society. One of the major points of this paper, to be proved in the sequel, is that every shortest dividing cycle has precisely two societies, one corresponding to each of its sides as indicated by our next claim:

For every shortest dividing cycle C of G and every society S of C, one of the sides of C contains all edges incident with C at an element of S. (39)

This can be seen by observing the spin behaviour or by using (31), see Figure 24.

Let us next check that shortest dividing cycles do not separate cosets of Γ_2 (compare this with the behaviour of 2-coloured cycles in earlier sections).

For every shortest dividing cycle C of G and every left coset Δ of Γ_2 in G, at most one of the sides of C contains elements of Δ. (40)

Suppose to the contrary there are elements $x, y \in \Delta$ in distinct sides of C, and let $W = w_1 \dots w_k$ be a word with letters $w_i \in \{A, A^{-1}, Z\}$ inducing a x–y path $P = xPy$ in G. Assume that W has minimum length among such words. Define the corners of P to be its vertices reachable from x by paths induced by prefixes of W, and note that every corner of P lies in Δ by definition. Moreover, by the minimality of W the only corners of P that do not lie on C are x and y.

By the definition of Z and (37) there is for every letter w_i a shortest dividing cycle C_i containing the corresponding subpath P_i of P induced by w_i. Even more, C_i and C_{i+1} have a common subpath Q_i induced by Z for every relevant i, and the endvertices of Q_i are corners of P. It might be the case that $C_i = C_{i+1}$. As x, y lie in distinct sides of C, and no C_i can cross C by (31), C *bounds* C_1 from C_k, i.e. C_1, C_k lie in distinct closed sides of C. This means that there is an i such that either C bounds C_i from C_{i+1}, or $C_{i+1} = C$ and $C_i \neq C$. If the former is the case then Q_i must clearly be a subarc of C.. But then one of C_i, C_{i+1} has a common subpath with C that properly contains Q since G is cubic, and this contradicts the maximality in the definition of Z. If the latter is the case, then Q_i is contained in $C = C_{i+1}$, and so C_i must leave C immediately before and after Q_i. But (39) now implies that P leaves C entering the side from which it approached C, contradicting our assumptions. This proves (40).

Using this we can prove the following observation. A *metaedge* is a path of G induced by one of the words A, Z.

Let $P \neq P'$ be metaedges. If both P, P' have endvertices in Γ_2 then they are independent. (41)

Recall that two paths are called *independent* if their interiors are disjoint.

To prove this, let C, C' be a shortest dividing cycle containing P, P' respectively, which exists by (36) and (37). Suppose first that both P, P' are induced by the word A. Let $\mathcal{F}, \mathcal{F}'$ be the face whose boundary contains P, P' respectively, and note that $\mathcal{F} \neq \mathcal{F}'$ since $P \neq P'$. It is easy to see that C separates $\mathcal{F}$ from $\mathcal{F}'$, because none of $\mathcal{F}, \mathcal{F}'$ can contain the other. But then, C separates the vertices of Γ_2 that lie on the boundary of $\mathcal{F}$ from the vertices of Γ_2 that lie on the boundary of $\mathcal{F}'$, which contradicts (40).

Suppose now that P is induced by the word Z. Let $D \neq C$ be a further shortest dividing cycle containing P, provided by (36). By (31) the closure of one of the sides of C' contains both C, D. Thus, by an easy topological argument, one of C, D, call it K, separates the other from C'. But then both sides of K meet Γ_2 as each of C', C, D contains a vertex in Γ_2 not contained in K. Again, this contradicts (40), and so (41) is established.

By (41) the Cayley graph G_2 of Γ_2 with respect to the generating set $\{a, z\}$ has a topological embedding in G: we can obtain G_2 from G by substituting for every two adjacent vertices x, y of G_2 the x–y path in G induced by A, A^{-1} or Z by an x-y edge labelled a or z accordingly. This yields indeed a topological embedding of G_2 in G since by (41) all these paths are independent. Starting with G_2 and replacing each edge back by the corresponding path induced by A, A^{-1} or Z we obtain the subdivision G_2' of G_2 alluded to earlier. Compared with G_2 the graph G_2' has the advantage that it is a subgraph of G while still capturing the structure of G_2. As in earlier sections, this will come in handy later, when we will try to yield a presentation of Γ from a presentation of Γ_2.

Moreover, for every coset Δ of Γ_2 in Γ we find an isomorphic copy of G_2' in G whose vertices of degree 3 are precisely the elements of Δ: such a copy can be obtained by mapping any vertex of Γ_2 to any vertex in Δ by a colour-automorphism of G. For every such copy define its *corners* to be its vertices of degree 3; in other words, the elements of the corresponding coset.

A planar presentation of Γ_2.

Note that G_2 is a cubic Cayley graph on two generators a, z, and our embedding σ of G induces, when combined with the aforementioned topological embedding of G_2 in G, an embedding σ_2 of G_2 in the sphere. It is straightforward to check that both a, z preserve spin in σ_2, for example using the fact that all vertices of G have the same spin in σ. It is also not hard to see that G_2 is 2-connected: apply Lemma 4.1 using the fact that $(az)^n$ is a relation by (38). Now if a has finite order, then Lemma 8.1 implies that G_2 has at most one end, and so we can obtain a planar presentation of G_2 from Theorem 7.5 (i): $G_2 \cong Cay\left\langle a, z \mid z^2, a^k, (az)^n \right\rangle$, $k \geq 3$, $n \geq 2$. Note that n can be read off (38) in this case.

If a has infinite order, then G_2 cannot be 3-connected by Lemma 8.4. Thus $\kappa(G_2) = 2$, and we can obtain a planar presentation of G_2 from Theorem 5.1 (i): $G_2 \cong Cay\left\langle a, z \mid z^2, (az)^n \right\rangle$, where again n is as in (38). Cases (ii) and (iii) of Theorem 5.1 cannot arise here, because we already know from (38) that az has finite order and this is not the case in these groups; see [13] for more details.

As in earlier sections we are going to plug these presentations into Theorem 8.2 to obtain a presentation of Γ.

Splitting G into copies of G_2'.

Before we can apply Theorem 8.2 we need a couple of further preparatory observations. Call a cycle of a copy of G_2' *basic* if it is a shortest dividing cycle of

G. Also call the corresponding cycle of G_2 basic.

For every pair of distinct copies H, H' of G_2' there is a unique basic cycle C of H that bounds H from H'. (42)

For this, let e be an edge in $E(H') \backslash E(H)$, which must exist if H, H' are distinct, and let P be a (possibly trivial) e–H path in G. Let f be the unique edge in $P \cup e$ incident with a vertex v of H. Then v is not a corner of H because $f \notin E(H)$. Thus v lies in the interior of a path Q of H induced by A or Z. If Q is induced by A, then the unique basic cycle C of H containing Q bounds f from the face boundary F of G containing Q. By (39) H meets the side of C containing $F \backslash Q$ as this side contains Q'. By (40), H does not meet the other side of C which contains f. Applying (40) again but this time on the coset of H', we obtain that H' only meets the side of C that does contain f since P must lie in that side. Thus $H - C$ and $H' - C$ lie in distinct sides of C, and so C is as desired.

In the other case, where Q is induced by Z, a similar argument applies except that now there are two basic cycles of H containing Q, and we have to choose C to be the one bounding f from the other.

The uniqueness of C follows easily from the fact that for every other basic cycle D of H we now know that C bounds D from H'.

Using (42) we can prove the following.

For every copy H of G_2' and every basic cycle C of H there is a copy $H' \neq H$ of G_2' such that $C \subseteq H'$ and C bounds H from H'. (43)

To see this, recall that C bounds a face $\mathcal{F}$ of H by (40). Let P be a maximal subpath of C that is contained in a basic cycle C' of some copy H' of G_2' that lies in $\mathcal{F} \cup C$. To see that such copies exist, note that as C is dividing in G, there is an edge xy of G such that $x \in V(C)$ and y lies in $\mathcal{F}$. Easily, x is not a corner of H for otherwise xy would lie in H. So let M be the coset of G_2 containing x, and let H_M be the corresponding copy of G_2'. Then H_M must be contained in $\mathcal{F} \cup C$ by (40), and it meets C at x. Thus any basic cycle of H_M containing x is a candidate for C'.

We claim that $P = C$ (in particular, P is a closed path). Suppose to the contrary that P has distinct endvertices $v, w \in V(C)$. Consider the copy J of G_2' corresponding to the coset of Γ_2 containing v, and note that, again by (40), J lies in $\mathcal{F} \cup C$ too as it contains an edge e, incident with v, that lies in $\mathcal{F}$ (Figure 25). Moreover J contains the edge $f \in E(C)$ incident with v that does not lie in P.

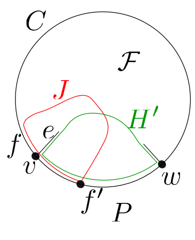

Figure 25: The situation in the proof of (43).

By (42) there is a basic cycle D of J that bounds J from H. Thus, easily, $E(H) \cap E(J) \subseteq D$; in particular, $E(C) \cap E(J) \subseteq D$. Let P' be the maximal common subpath of C and J containing f, and note that $P' \subseteq D$ too. By the maximality of P, and as P' contains the edge $f \notin E(P)$, we have $P \not\subseteq P'$. By an easy topological argument, D now crosses C' contradicting (31). This contradiction proves our claim that $P = C$, which immediately implies (43).

It is straightforward to strengthen (43) to demand that there is a colour-automorphism g of G such that $H' = gH$ and $C = gC$; just use the fact that C' in the above proof is a translate of C by (38). Thus, any society S of C is mapped by g to a further society S' of C, which is distinct from S as, by (39) and the choice of H', it follows that S and S' point into distinct sides of C. Note that C cannot have a third society, as this would have to share a side with one of S, S' easily yielding a contradiction to (31). Thus, every shortest dividing cycle has two *dual* societies, one for each of its sides. This remarkable fact is one of the main ideas of this paper, and it is not peculiar to the current section: dual societies were implicit in all cases we have seen so far. For example, the vertices on the big cycle of Figure 1 (iii) pointing to one of its sides form a society, and the reader will easily spot the dual societies of the cycles of Figure 2 (ii) corresponding the 4-cycles of Figure 2 (i).

Back to our analysis of G, we note that for every vertex x of G there is a coset of Γ_2 containing x, and so there is a copy H^x of G'_2 in G containing x as a corner. We will use these graphs H^x in our application of Theorem 8.2.

For this, let T be the graph with vertex set $\{H^x \mid x \in G\}$ in which two vertices H^x, H^y are joined by an edge if they share a basic cycle bounding H^x from H^y. We claim that T is a tree. It is easy to see that T is acyclic, since any cycle would yield a topological impossibility. To prove that T is connected, suppose it is not and choose vertices H^x, H^y in distinct components of T minimizing the distance from H^x to H^y in G. Let C_0 be the basic cycle of H^x bounding it from H^y as provided by (42). Let H_1 be the other vertex of T, provided by (43), containing C_0. If H_1 lies in the component of T containing H^y then we obtain a contradiction, since H_1 is joined to H^x by an edge of T. If not, then we can repeat the above procedure, replacing H^x by H_1, to obtain H_2 and C_2 bounding H_2 from H^y. Continuing like this, we obtain a sequence $H_0 = H^x, H_1, H_2, \ldots$ of vertices of T all of which lie in the component of H^x, and a sequence of basic cycles C_i bounding H^y from $H_0, \ldots, H_i$. But after finitely many steps we must obtain some C_i that is disjoint from C_0, for all C_i have the same length and G is locally finite. This means that if P is a shortest H^x–H^y path in G, then C_i intersects P at a vertex $z \notin V(H^x)$ for C_i bounds H^y from H^x. But then the distance from H_i to H^y is shorter than the distance from H^x to H^y, which contradicts the choice of the pair H^x, H^y. This contradiction proves that T is a tree as claimed.

We can now apply Theorem 8.2. As in earlier sections, we choose for every H^x the generating set $\mathcal{F}_x$ induced by the presentation of Γ_2 we obtained above. This easily yields the following presentations:

- $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bcd)^k; (bcdZ)^n \right\rangle$, $k \geq 3$, $n \geq 2$ (faces of size $3k \geq 9$);
- $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bcdZ)^n \right\rangle$, $n \geq 2$ (no finite faces).

Indeed, recall that either $G_2 \cong Cay \left\langle a, z \mid z^2, a^k, (az)^n \right\rangle$, $k \geq 3$, $n \geq 2$, which is the case if a has finite order, or a has infinite order and $G_2 \cong Cay \left\langle a, z \mid z^2, (az)^n \right\rangle$. By the above discussion, and replacing bcd back for a, we then obtain $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, Z^2, (bcd)^k, (bcdZ)^n \right\rangle$, $k \geq 3$, $n \geq 2$ in the first case and $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, Z^2, (bcdZ)^n \right\rangle$, $n \geq 2$ in the second. In both cases we can omit the relation Z^2 though because, as mentioned earlier, (35) implies that Z is a symmetric word and so Z^2 can be deduced from the relations b^2, c^2, d^2. Our presentations are planar by (31).

Note that these presentations coincide with (ii) and (iii) of Theorem 9.10. Presentation (i) corresponds to the case when G has hexagonal faces and was obtained in our earlier discussion for that case. This completes the proof of the forward implication of Theorem 9.10.

The converse implication

It remains to prove the converse implication of Theorem 9.10 for presentations of type (ii) and (iii): given a non-crossing pattern $\mathcal{P}$ and a $k \in \{3, 4, \ldots\} \cup \{\infty\}$, we have to show that the Cayley graph $G = Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bcd)^k; \mathcal{P} \right\rangle$ has the desired properties.

As in the forward implication, we can prove the existence of a dominant and a weak colour in $\mathcal{P}$; this time, instead of using Corollary 6.1 like in the proof of (29) we can just use the fact that *cp* is non-crossing by assumption. The existence of the word Z can be proved as above too using this assumption. Thus, similarly to, (38), we can prove that

$$\mathcal{P} \text{ is of the form } (AZ)^n \tag{44}$$

up to rotation and inversion, where again $A = bcd$ and Z is a symmetric word. This means that every cycle C 'induced' by $\mathcal{P}$ has a society in the sense of Definition 9.11. A *pending edge* of C is an edge incident with but not contained in C. Similarly to (39) we have:

For every cycle C complying with the pattern $\mathcal{P}$, and every side A of C, there is a society S of C such that all pending edges of C incident with S lie in A. (45)

Indeed, if s, t are subsequent elements of S whose pending edges lie in distinct sides of C, then it is easy to construct a rotation of C containing the s–t path 'induced' by A or Z.

If $k < \infty$, i.e. if our given presentation is of type (ii), we apply the converse of Theorem 7.5 (i) to obtain a (finite or 1-ended) 3-connected planar Cayley graph G_2 of the group $\Gamma_2 = \left\langle a, z \mid z^2, a^k, (az)^n \right\rangle$, each face of which is induced by a^k or $(az)^n$. If $k = \infty$, i.e. if we are in type (iii), we apply the converse of Theorem 5.1 (i) to obtain the planar 2-connected Cayley graph G_2 of the group $\Gamma_2 = \left\langle a, z \mid z^2, (az)^n \right\rangle$. The embedding of the latter graph is not unique; we choose an embedding in which each cycle induced by $(az)^n$ bounds a face. This graph, and the desired embedding, can be obtained from a 3-regular tree by replacing every vertex with a cycle of length $2n$ alternating in the colours a, z, and letting the cycles replacing two adjacent vertices of the tree share a single a-edge. Note that G_2 has infinite faces too, bounded by a-coloured double rays.

In both cases, replace each a-edge of G_2 by a path of length three with edges coloured bcd respecting the directions of the a-edges, and replace each z edge of G_2 by a path along which the colours of the edges read as the letters in Z. Let G_2' be the resulting edge-coloured graph. Notice the similarity with the G_2' used for the forward implication of Theorem 9.10.

We will use similar ideas as in the proof of the forward implication, except that instead of finding copies H^x of G_2' as subgraphs of a given graph G, we now have to construct them from scratch. In the end we will have to check that their union is a Cayley graph, and that it has the desired properties.

Let $H_0 := G_2'$. As earlier, we call any cycle of G_2' a *basic cycle* if the corresponding cycle of G_2 was induced by $(az)^n$. We will construct a sequence $(H_i)_{i \in \mathbb{N}}$ by inductively glueing a copy of G_2' in each face of H_{i-1} bounded by a basic cycle. In order to be able to do so we will need the following assertion, which is reminiscent of (43).

For every cycle C complying with the pattern $\mathcal{P}$ and every side A of C, there is a society S of C such that all pending edges of C incident with S lie in A. (46)

We will imitate the proof of (43), except that we now have no underlying graph G in which we can look for C', and so we will have to make do with rotations of C. Note that by (44) there is at least one society T of C, so we can assume that the pending edges of C incident with S lie in the wrong side $B \neq A$ of C. We may also assume that every vertex x of C is incident with a pending edge, for otherwise we may attach such an edge, embedding it so that x has the right spin and giving it the colour missing from x, without affecting the compliance of C with $\mathcal{P}$. So let xy be a pending edge of C such that $x \in V(C)$ and y lies in A, which exists by (iii) of Definition 9.9. Let t be a vertex in T such that the colours of the two edges of C incident with t are the same as the colours of the two edges of C incident with x (one of which colours must be the dominant colour c). To see that such a t exists, note that T has by definition two vertices joined by a bcd subpath of C, and as c is the dominant colour, one of these vertices is incident with a b and a c edge and the other is incident with a d and a c edge of C. Let D be the t–x rotation of C, and recall that D does not cross C.

By the construction of D, x is a member of a society of D. Thus we can ask for a maximal subpath P of C that is contained in some rotation D of C such that some society member x of D lies in C and has a pending edge in A. As in the proof of (43) (recall Figure 25) we claim that $P = C$. Suppose, to the contrary, that P has distinct endvertices $v, w \in V(C)$. By the previous arguments, we can construct a rotation D' of C having v as a society vertex. Note that D' contains the edge $f \in E(C)$ incident with v that does not lie in P.

Let P' be the intersection of C and D'. By the maximality of P, and as P' contains the edge $f \notin E(P)$, we have $P \nsubseteq P'$, and so there is an interior vertex z of P at which D' leaves C. It follows that D' crosses D, as can be easily seen by looking at the edges incident with the endpoints of the path $P \cap P' = zPv$. But as D and D' share some edges they are a rotation of each other, and this crossing contradicts the definition of $\mathcal{P}$. This proves that $P = C$ as claimed. Combined with (45), this easily implies (46) as $P = C$ has by definition a society member with a pending edge in A.

Note that a society as in (46) is uniquely determined once we fix the side A: if there were two distinct such societies S, T on C then it would be easy to find

$s, s' \in S$ and $t, t' \in T$ with $s' = sz, t' = tz$ such that s, t, s', t' appear in that order on C, but then the corresponding rotations of C would cross. This means that

for every cycle C complying with $\mathcal{P}$ there are precisely two societies on C, with pending edges on opposite sides. This societies are disjoint. (47)

Back to our construction of H_i, let B be a basic cycle of H_0 one of the sides $\mathcal{F}_B$ of which is a face of H_0. Note that all corners of H_0 have the same spin since this was the case in G_2. Thus B complies with $\mathcal{P}$, and we may apply (47) to obtain two distinct societies of B. One of these societies must be the set of corners of B by the construction of G_2'; we will use the other society S to extend H_0: build an isomorphic (respecting edge-colours) copy H_B of G_2' that has B as a basic cycle too, and in which the vertices in S are corners, embedding H_B in the closure of $\mathcal{F}_B$. Let H_i be the union of H_{i-1} with all these graphs H_B, one for every facial basic cycle B of H_0 as above. Note that every corner of H_1 has the same spin; we use this fact as an induction hypothesis in the construction of $(H_i)_{i \in \mathbb{N}}$. The following steps $H_2, H_3, \dots$ of this construction are similar to the last one: we consider all basic cycles B bounding a face of H_{i-1}, and embed a copy of G_2' in that face attaching it at the yet unused society of B. All basic cycles comply with $\mathcal{P}$ at every step by our induction hypothesis, and so we can always apply (47). Since the newly attached pending edges of any B as above are put in the 'right' side of B, our induction hypothesis is preserved. Note that a basic cycle B considered in step i will not be reconsidered in any subsequent step as it does not bound a face any more.

Let $G := \bigcup_{i \in \mathbb{N}} H_i$. By construction G is planar. To see that it is cubic, each vertex x being incident with all three edge-colours b, c, d, note that if x has an incident pending edge e in H_i, then e lies in a face bounded by a basic cycle B_i of H_i. Thus in the next step, e will either lie in $H_{B_i} \subseteq H_{i+1}$ or inside a face of some basic cycle $B_{i+1} \neq B_i$. It follows easily that e lies in some H_j for otherwise there is an infinite sequence $(B_i)_{i \in \mathbb{N}}$ of nested distinct cycles of the same length with a common vertex x, which cannot be the case in a locally finite graph.

Furthermore, all vertices of G have the same spin since this was the case in each H_i.

We claim that G is a Cayley graph. This will follow from the following two assertions and Sabidussi's theorem.

For every two basic cycles C, D of G, any societies S, T of C, D respectively, and any members $s \in S, t \in T$, there is a colour-automorphism g of G such that $gC = D$ and $gS = T$ and $gs = t$. (48)

To begin with, it easy to construct a colour-preserving isomorphism g' from C to D with $g'S = T$ and $g's = t$. We will extend g' to a colour-automorphism of the whole graph G. By the construction of G there are copies H_C and H_D of G_2' in G in which the members of S and T are corners. Indeed, if C was constructed in step i, and S has pending edges in the side of C that is a face of H_i, then we explicitly constructed H_C in the definition of H_{i+1}. If S has pending edges in the other side, then we take H_C to be the copy of G_2' in the construction of H_i containing C. The same goes to D.

Now let g_0 be a colour-preserving isomorphism from H_C to H_D that coincides with g' on C. Such a colour-preserving isomorphism exists because H_C and H_D are both copies of G'_2, and any basic cycle of G'_2 can be 'rotated' by the action of Γ_2.

Having defined g_0 we proceed inductively, in ω steps, to define extensions $g_1 \subset g_2, \subset, \ldots$, as follows. In each step $i > 0$ consider each new basic cycle B in the domain of g_{i-1}, let $B' := g_{i-1}B$, and note that precisely one of the two societies S_B of B —provided by (47)— has the property that its pending edges lie in the domain of g_{i-1}. Let S^B be the other society of B. Similarly, let $S^{B'}$ be the society of B' whose pending edges do not lie in the range of g_{i-1}. Then let g_i map H_B to $H_{B'}$, extending g_{i-1}, where H_B and $H_{B'}$ are defined as H_C above, but this time with respect to the societies S^B and $S^{B'}$.

Let $g := \bigcup g_i : G \to G$. To see that g is surjective notice the similarity in the definitions of g_i and H_i. Injectivity follows easily by considering the embedding of G and its subgraphs H_B. Thus g is the desired colour-automorphism.

It is easy to see that

for every $v \in V(G)$ there is a basic cycle B of some H_i such that v is in a society of B. (49)

Indeed, let i be the first index such that all three edges of v lie in H_i. Then by the construction of H_i there was a basic cycle B that had v as a society vertex.

Combining (48) with (49), we see that for every $v, w \in V(G)$ there is a colour-automorphism of G mapping v to w. Sabidussi's theorem thus yields that G is a Cayley graph as claimed.

It follows immediately from Lemma 4.1 that G is 2-connected.

Finally, we claim that each pair of colours spans double rays in G, in other words, each of bc, cd, db has infinite order. To see this, let P be a component of G spanned by two colours. Note that for every basic cycle C of G, if P meets C then P crosses C; this is easy to see using the fact that all vertices have the same spin and every other edge of C bears the dominant colour c. Pick a basic cycle C_0 meeting P, and let x_1 be the first vertex of P after leaving C_0. Let C_1 be a basic cycle containing x_1, and let x_2 be the first vertex of P after leaving C_1. Continue like this to define infinite sequences x_i and C_i. By our previous remark, each C_i bounds C_{i-1} from C_{i+1}; thus the C_i are all distinct, and so P is infinite proving our claim.

To sum up, starting with an arbitrary non-crossing pattern $\mathcal{P}$ and an arbitrary value $k \in \{3, 4, \ldots\} \cup \{\infty\}$, we constructed a plane Cayley graph with the properties required by the forward implication of Theorem 9.10. Thus we can now apply the forward implication, which yields that G has one of the presentations of the assertion. As we constructed G to have faces of size $3k$, and dividing cycles induced by $\mathcal{P}$, it follows that this presentation is indeed the intended one. This completes the proof of the backward implication of Theorem 9.10. It only remains to prove the last statement about $\kappa(G)$.

The connectivity of G

In this section we determine the connectivity of each of the graphs of Theorem 9.10 of type (ii) or (iii); we have already seen that those of type (i) are

3-connected. The most interesting result is that such a graph can be 3-connected even if its faces have infinite size.

We will say that our non-crossing pattern $\mathcal{P}$ is *regular* if it is, up to rotation, of the form $\mathcal{P} \cong (dc(bc)^m)^n$. Note that in this case we have

$$Z = c(bc)^{m-1}, \tag{50}$$

by (38) and the definition of A.

We begin with a basic observation that will be useful later.

$$Z \text{ contains the letter } d \text{ if and only if } \mathcal{P} \text{ is not regular.} \tag{51}$$

Indeed, if $\mathcal{P}$ is regular then any d-edge e uniquely determines a cycle through e induced by $\mathcal{P}$. As Z induces the intersection of two distinct cycles induced by $\mathcal{P}$, the forward implication follows.

For the backward implication, suppose now that $\mathcal{P}$ is not regular, which means that $\mathcal{P}$ has two 2-coloured bc 'subpaths' P, Q of distinct lengths surrounded by d-edges. Now given a d-edge $e = uv$ of G, we can let $\mathcal{P}$ induce two cycles C, D of G by starting at u, say, and reading $\mathcal{P}$ once starting at the beginning of P and once at the beginning of Q. Then C, D are distinct, and they both contain e. The definition of Z and (35) now immediately imply that Z contains the letter d as claimed.

Let N be the size of the faces of G, and recall that $N = 3k$ is finite if G is of type (ii) and it is infinite if G is of type (iii). The main result of this subsection is

Lemma 9.12. *If $N < \infty$ or $\mathcal{P}$ is not regular then G is 3-connected.*

Proof. If for every copy H_C of G_2' as in the construction of G from page 60, all vertices in $V(H_C)\backslash\{x, y\}$ lie in a common component of $G - \{x, y\}$ for every pair of vertices $\{x, y\} \subset V(G)$, then G is 3-connected by Lemma 4.2. Thus we can assume that $V(H_C)\backslash\{x, y\}$ is disconnected in $G - \{x, y\}$ for some H_C and some pair of vertices $\{x, y\}$. This implies in particular that $x, y \in V(H_C)$, for H_C is 2-connected.

Let us first consider the case where both x, y are corners of H_C. Then N cannot be finite, because in this case G_2 in 3-connected, and so no two corners of its subdivision $G_2' \cong H_C$ can disconnect it. Thus we may assume that $N = \infty$, and so $G_2 = Cay\left\langle a, z \mid z^2, (az)^n \right\rangle$. We claim that x, y must lie on a common basic cycle of H_C. For if not, then the vertices of each basic cycle remain connected after removing x, y. Moreover, as H_C is the union of basic cycles, of which any two incident ones D, D' have two corners in common, x, y cannot disconnect D from D' unless $x, y \subset V(D)$. But if x, y leaves every basic cycle connected, and every two adjacent ones are connected to each other, then it leaves their union H_C connected. Thus we can indeed assume that x, y lie in $V(D)$ for some basic cycle D of H_C, and since they are corners of H_C they lie in a common society S of D.

It also follows from the previous argument that either x, y are the common corners of two basic cycles D, D', in which case they are joined by a subpath of D induced by Z, or both components R, Q of $D - \{x, y\}$ contain a vertex of S, in other words, x, y are not neighbouring society vertices. In both cases,

each of R, Q contains the interior of at least one Z-path with endvertices in S. (52)

The following assertion will now easily imply the existence of an R–Q path in the copy H_D of G'_2 containing the dual society $\overline{S}$ of S, provided by (47), as corners.

If $\mathcal{P}$ is not regular then the interior of every Z-path in some society S contains an A-path between vertices of the dual society. (53)

To prove (53), note that if it fails then for some $u, v \in S$ and the u–v path P induced by Z, any vertex $z \in \overline{S} \cap P$ is at distance less that four from one of u, v for A has length three. As every second edge of P bears the dominant colour c, this distance must be precisely two, so assume that $d(u, z) = 2$ (Figure 26). Easily, P must start and end with the dominant colour c by the definition of Z. Similarly, the second edge of P cannot bear the weak colour d, for no basic cycle has a dcd subpath by (33). Thus $z = ucb$. Consider the two basic cycles C, C' containing P, as well as the basic cycle D containing the Z-path starting at z.

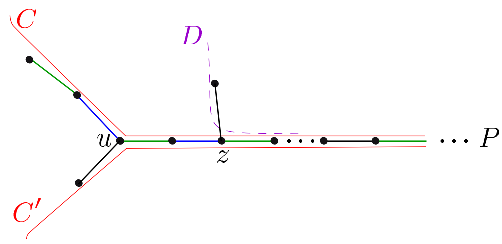

Figure 26: The cycles C, C' and D in the proof of (53).

Now recall that P contains a d-coloured edge by (51), and so we can let e be the d-coloured edge of P that is closest to u. Let B, B', B'' be the maximal bc 2-coloured subpaths of C, C' and D respectively containing z. Note that both B, B' start at e by the definition of the latter, and precisely one of them stops at u. Thus $|B| \neq |B'|$. Note also that B'' starts at z and cannot go past e, so that $|B''| < |B|, |B'|$. This means that $|B|, |B'|, |B''|$ are pairwise distinct, and so one of them contains at least 2 more b edges than some other. But this contradicts (34). This contradiction proves (53).

Let T be the Z-path incident with x, say, and assume without loss of generality that $T \subseteq R \cup \{x\}$, which we may by (52). Then (53) yields a pair of elements v, w of $\overline{S}$ that both lie in the interior of T. Now assume v is closer to x than w on T, and let $v' = vZ$ be the element of $\overline{S}$ joined to v by a Z-path $T' \subset D$. Note that by the choice of v' and (52) we must have $v' \in Q$. This means that the copy H_D of G'_2 containing the dual society $\overline{S}$ as corners contains a P–Q path L in $G - \{x, y\}$: consider the basic cycle D' of H_D containing T' that is not D itself, and let $L = D' - \mathring{T}'$ be the other v-v' subarc of that basic cycle. But this path L easily contradicts our assumption that H_C meets two distinct components of $G - \{x, y\}$.

This contradiction implies that x, y cannot both be corners of H_C, and so we may assume that y, say, is not a corner of H_C from now on.

Note that x, y cannot disconnect any two corners of H_C from each other because, as the reader will easily check, G'_2 cannot be disconnected by removing a vertex and an edge. This implies that x, y lie on a common path B of H_C

induced by Z or A, disconnecting part of B from the rest of H_C. Let D be a basic cycle of H_C containing B. Note that $D-\{x,y\}$ has two components R,Q one of which, R say, is a proper subpath of B.

Now consider the other copy $H_D \supset D$ of G'_2 provided by (43). If H_D has a corner in R, then H_D contains an R–Q path L as above, and such a path easily implies that H_C cannot meet two distinct components of $G-\{x,y\}$ contrary to our assumptions. If, on the other hand, H_D has no corner in R, then H_D contains a further basic cycle $D_1 \neq D$ such that $\{x,y\} \cup R$ is contained in some path B_1 of D_1 induced by Z or A. We claim that $\{x,y\}$ must also disconnect R from $H_D - R$ in G. Indeed, if it does not, then there is a path from R to $(D-R) \subset H_D$ in $G-\{x,y\}$, and this path contradicts our assumption that $\{x,y\}$ disconnects R from $H_C \supset D$ in $G-\{x,y\}$. This proves our claim.

Now repeating these arguments on D_1 and the corresponding copy of G'_2 instead of D and H_C, and iterating, we either obtain a contradiction after finitely many such steps, or an infinite sequence D_i of distinct basic cycles containing R. But as G is locally finite, such a sequence cannot exist and we have a contradiction in any case.

□

Lemma 9.12 implies that once $\mathcal{P}$ is non-regular, G is 3-connected even if all its faces have infinite boundary:

Corollary 9.13. *For every non-regular non-crossing pattern $\mathcal{P}$, the graph $Cay\left\langle b,c,d \mid b^2,c^2,d^2;\mathcal{P}\right\rangle$ is 3-connected and has no finite face boundary.*

Since non-regular non-crossing patterns do exist, take for example $(d\ cbc\ d\ cbcbc)^n$, we obtain examples of planar 3-connected Cayley graphs no face of which is bounded by a cycle, which is interesting in view of our discussion of Section 1.3.

Our next result is that if $\mathcal{P}$ is regular and the faces have infinite size then G is not 3-connected. This means that Lemma 9.12 is best possible.

Lemma 9.14. *Let $G = Cay\left\langle b,c,d \mid b^2,c^2,d^2;(b(cb)^m d)^n\right\rangle$, $m \geq 1$, $n \geq 2$. Then $\kappa(G)=2$. Moreover, G has a hinge (coloured c) if and only if $m=1$.*

Proof. Recall that G is 2-connected by Lemma 4.1. Let us check that G is not 3-connected.

In our proof of the converse implication of Theorem 9.10 we constructed G as the union of the sequence of its subgraphs H_i, and we will base our proof on that sequence. Recall that when $N=\infty$, $G'_2 = H_0$ was obtained from the Cayley graph $G_2 = Cay\left\langle a,z \mid z^2,(az)^n\right\rangle$ after replacing each a-edge by a bcd-path and replacing each z edge by a path induced by $Z = c(bc)^{m-1}$, where we used (50). Consider a basic cycle C of H_0, and two d-edges e,f of C. We claim that these two edges separate G.

To see this, let P,Q be the two components of $C-\{e,f\}$. Note that no basic cycle D of H_0 meets both P and Q, which can be seen by applying (51) to the pair C,D. This means that $\{e,f\}$ separates H_0, as the latter is the union of its basic cycles. Similarly, $\{e,f\}$ separates H_C, the copy of G'_2 that we glued along C in the construction of H_1, into two components containing P and Q respectively. It now follows easily that $\{e,f\}$ separates G into two components containing P and Q respectively, as G can be obtained from $H_0 \cup H_C$ by inductively glueing copies of G'_2 along some basic cycle, and so none of these copies meets both P and Q. Now choosing one endvertex from each of e, f yields a separator of G, which means that G is not 3-connected.

It remains to prove that G has a hinge (coloured c) if and only if $m = 1$. Firstly, it is easy to see that, in every case, no edge coloured b or d is a hinge: for if uv is such an edge, then consider the Z-path P starting at u, and the two basic cycles C, C' containing P (Figure 26). Note that three of the four vertices adgecent with uv lie in $C \cup C'$ and are thus in a common component K of $G - \{u, v\}$. Moreover, the fourth vertex x, which is adjacent with v, is connected to the neighbours of u by a path that avoids u since G is 2-connected. Thus, x also lies in K. This proves that uv is not a hinge as it fails to disconnect its neighbourhood. Let us now consider the case when uv is a c-edge instead. If $m > 1$, then by (50) the path P defined as above has length greater than 1, and the we can apply the same arguments to show that uv is not a hinge. If $m = 1$ however, we have $Z = c$, and so the endvertices of any c edge separates in G'_2 the two basic cycles containing it. This easily implies that any c edge is a hinge of G. □

Lemma 9.14 combined with Lemma 9.12 determine $\kappa(G)$ for each of the graphs G as in Theorem 9.10 (ii), (iii).

9.2.2 All edges reverse spin

If G is 3-connected and all its edges reverse spin, then by Lemma 9.2 G cannot be multi-ended. Thus we can proceed with the next case.

9.2.3 Mixed spin behaviour

The last case we have to consider is that of a planar, multi-ended 3-connected Cayley graph $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, \ldots\right\rangle$ with no 2-coloured cycles, with both spin-preserving and spin-reversing colours. The main result of this section, characterizing these graphs, is the following.

Theorem 9.15. *Let $G = Cay\left\langle b, c, d \mid b^2, c^2, d^2, \ldots\right\rangle$ be a planar 3-connected multi-ended Cayley graph with both spin-preserving and spin-reversing edges. Suppose that all three of bc, cd and db have infinite order. Then precisely one of the following is the case.*

1. $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bdbcdc)^k; (c(bc)^n d)^{2m}\right\rangle, k \geq 2, n, m \geq 1,\ n + m \geq 3$;
2. $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bdbcdc)^q; (c(bc)^{n-1} d)^{2m}, (c(bc)^n d)^{2r}\right\rangle, n, r, m, q \geq 2$;
3. $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2; (c(bc)^{n-1} d)^{2m}, (c(bc)^n d)^{2r}\right\rangle,\ n, r, m \geq 2$.

All these presentations are planar.

Conversely, for every k, n, m in the specified domains the above presentation yields a Cayley graph as above.

The rest of this section is devoted to the proof of Theorem 9.15, so let us fix a Cayley graph G as in its assertion. Our analysis will be similar to that of Section 9.2.1.

It follows from Lemma 9.2 that only one of the colours b, c, d can be spin-reversing. Assume from now on that this colour is d. This implies that

$$\text{every facial walk is of the form } \ldots bdbcdc\ldots. \tag{54}$$

The shortest dividing cycles

As in the previous section our analysis will be based on the shortest dividing cycles of G.

Proposition 9.16. *For every shortest dividing cycle C of G, there is a colour $a \in \{bcd\}$, such that every other edge of C is coloured a. (In particular, $|C|$ is even.)*

As earlier, we will call this colour a the *dominant* colour of C.

Proof. We begin by showing that

$$\text{no facial subpath of a shortest dividing cycle contains two } d\text{-edges.} \tag{55}$$

For let P be a facial subpath of an shortest dividing cycle chosen so as to maximize the number p_d of d-edges in P, and suppose that $p_d \geq 2$. Let $e = uv$ be the last d-edge of P, and let C be a shortest dividing cycle containing P. Then the colour-automorphism of G exchanging u and v maps C to a shortest dividing cycle D, and it is not hard to see, using (54), that D crosses C. However, as in the proof of Corollary 6.1, such a crossing gives rise to a new shortest dividing cycle C' containing a facial subpath $P' \supseteq P$ with more than p_d d-edges, contradicting the choice of P. This proves (55).

Now let C be any shortest dividing cycle of G. Note that as C contains all three colours, it must contain two adjacent edges coloured b and c unless every other edge of C is coloured d. In the latter case our assertion is already proved, so assume from now one that the former is the case. So let P be a maximal bc-coloured subpath of C with $||P|| \geq 2$. If $||P|| = 2$ then both edges e, f of C incident with P are coloured d, and so $e \cup P \cup f$ is, by (54), a facial subpath of C containing two d-edges contradicting (55). Thus $||P|| \geq 3$.

Similarly to the proof of (11) we can now prove that

$$\text{every maximal } bc\text{-coloured subpath of a shortest dividing cycle of } G \text{ has odd length.} \tag{56}$$

It follows that $||P||$ is odd, and so P starts and ends with the same colour, c say. Let x be the last vertex of P, let e be the d-edge of C following P, and let f be the edge of C after that. We claim that the colour of f is c. For if it is b, then the colour-automorphism exchanging the endvertices of e maps C to a cycle C' crossing C. Note that the maximal bc-coloured subpath Q of C' starting at x has odd length by (56). Now let C'' be the translate of C' obtained by mapping x to $w := xcb$ (Figure 27). Note that C'' also crosses C, and by (3) we can replace a subpath of C containing e by a subpath of C'' not containing e to obtain a new shortest dividing cycle D. As each of Q and P had odd length, D has a maximal bc-coloured subpath of even length contradicting (56). This contradiction proves our claim that the colour of f is c.

Our next assertion will allow us to determine the colour of the edge of C following f.

$$\text{No shortest dividing cycle of } G \text{ has a 2-coloured subpath of length at least four containing } d\text{-edges.} \tag{57}$$

To prove this, let Q be a 2-coloured subpath of a shortest dividing cycle containing d-edges that maximises $|Q|$ among all such paths, and suppose that Q

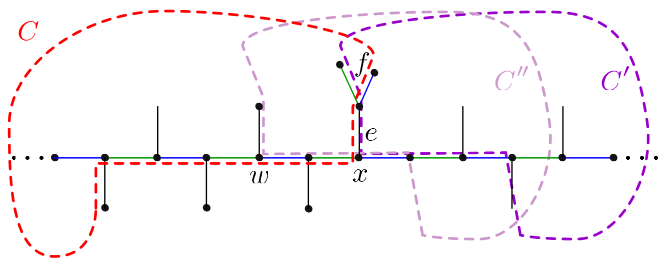

Figure 27: The construction of D in the proof of Proposition 9.16.

contains at least four edges. Let D be a shortest dividing cycle containing Q. Then, exchanging the endvertices of one of the d-edges e of Q by a colour-automorphism of G we obtain one of the two situations depicted in Figure 28. In both cases though, (3) implies the existence of a shortest dividing cycle containing a 2-coloured path that extends Q contradicting its maximality. This contradiction proves (57).

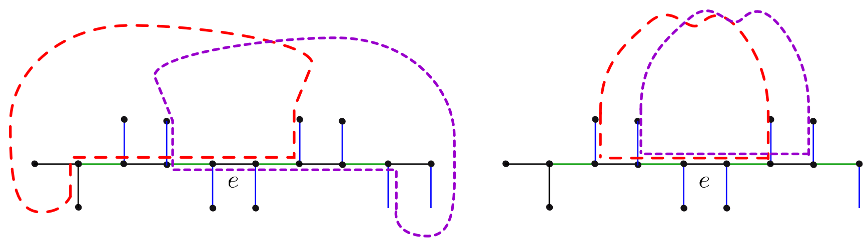

Figure 28: The situation in the proof of (57).

Now as P ends with a c edge, and f is also coloured c, (57) implies that the edge of C following f is coloured b. Let P_1 be the maximal bc-coloured subpath of C containing f. Then $||P_1|| \geq 2$, and by (56) $||P_1||$ is odd. Thus, P_1 starts and ends with a c edge, as was the case for P. This allows us to apply the same arguments to P_1 to show that it is followed by a d-edge and another odd bc-coloured path P_2, and so on. This proves that every other edge of C is coloured c as desired. □

As in Section 9.2.1 it is important to know whether G has a hexagonal face. Fortunately, it does not:

Proposition 9.17. *G has no hexagonal face.*

Proof. Suppose, to the contrary, that G does have a hexagonal face-boundary H. Let C be a shortest dividing cycle of G. We distinguish two cases according to the spin-behaviour of the dominant colour a of C.

If a is the spin-reversing colour d, then our approach is similar to that of Section 8.2.1. Let R be bc-double-ray incident with H. Considering the colour-automorphisms mapping R to itself it is easy to check that every face boundary incident with R must be a hexagon too, as it can be mapped to H. We can proceed as in Section 8.2.1 to prove an assertion similar to (19) there. For this, we superimpose a new dummy colour o to each edge of G coloured b or c, and direct all edges bearing that colour 'to the right', as in Figure 15. We can now repeat the proof of (19) to show that o appears in C the same number of times in each direction, and use this fact to construct a new shortest dividing cycle C' that 'reads' $(odo^{-1}d)^n$. Note however that the way we directed the o edges combined with the fact that d reverses spin implies that such a cycle C' contains only two of the colours b, c, d, contradicting our assumption that G has no 2-coloured cycle.

If a is one of the spin-preserving colours, c say, then as C must contain all three colours, it must have a $bcdc$ subpath. But such a subpath can be shortcut by a db path since, by (54), $bdbcdc$ is the relation inducing the hexagonal face-boundary H. This contradicts the minimality of C. □

Using this fact we can now enrich our knowledge about the shortest dividing cycles. Let C be a shortest dividing cycle of G, and let $\mathcal{T}$ be the set of translates of C by colour-automorphisms of G.

No two elements of $\mathcal{T}$ cross. (58)

Suppose, to the contrary, two elements D, D' of $\mathcal{T}$ cross. We will now use the ideas of Section 8.2.1 to obtain a contradiction from Euler's formula (17). By (2) there is a subpath P of D and a subpath P' of D', with common endvertices u, v, such that $P \cup P'$ is a cycle K bounding a region B containing finitely many vertices. We repeat the construction of H' of Section 8.2.1: let again H be the finite plane subgraph of G spanned by K and all vertices in B. Let H' be the graph obtained from two copies of H by joining corresponding vertices of degree two by an edge, and note that H' is cubic. Consider an embedding of H' in the sphere such that the two copies of H occupy two disjoint discs D_1, D_2, and the newly added edges and their incident faces lie in an annulus Z that joins these discs. Since by Proposition 9.17 all faces within these discs have size greater than 6, contributing a negative curvature to (17), it suffices to show that the total curvature of the faces in Z does not exceed 12 to obtain a contradiction to (17). The faces in Z have even size by construction, so that it suffices to show that Z cannot contain more than six 4-gons. The presence of spin-reversing edges makes this task slightly harder than in Section 8.2.1, and we need a new argument.

We are going to reduce Z to a small auxiliary graph Z' by performing operations that leave the curvature, as well as the spin-behaviour of the colours, invariant. We will end up with only few possibilities for Z', in which it will be easy to count the 4-gons. The first of these operations suppresses a pair of spin-preserving edges on opposite sides of Z. Let e be an edge of Z coloured b or c which is an interior edge of P or P' such that both edges f_1, f_2 of P or P' incident with e bear the other spin-preserving colour (this could be the dominant colour for example). Let e', f_1', f_2' be the corresponding edges on the other side of Z. Note that because of the spin-behaviour, precisely one of the endvertices of e is adjacent to its counterpart in e' by a d-edge g. The situation is thus as in the left part of Figure 29. Now let us delete e, e' and g, and identify f_1 with f_2 and f_1' with f_2' to obtain a new graph Z^*. Note that Z^* has the same 'curvature' as Z: indeed, Z^* has one face less than Z, and precisely 6 edges less than Z in the count of the sum of face-sizes of (17) (g is counted twice because it lies in two face boundaries). Thus, this operation is indeed neutral as far as Euler's formula is concerned.

We can perform a similar operation whenever Z, or any auxiliary graph Z^* obtained after performing the above operation a number of times, has a subpath of the form *cdcdc* or *bdbdb* on its boundary: the right part of Figure 29 shows how to remove this kind of subpath and its counterpart affecting neither the total curvature nor the spin-behaviour of the colours on the boundary of Z.

We now distinguish two cases according to which colour is dominant in C (and thus in any element of $\mathcal{T}$).

If the dominant colour is a spin-preserving one, let us say c, then we can apply the operation of the left half of Figure 29 repeatedly to eliminate from Z each b edge that was an interior edge of P or P'. Then, we can apply the

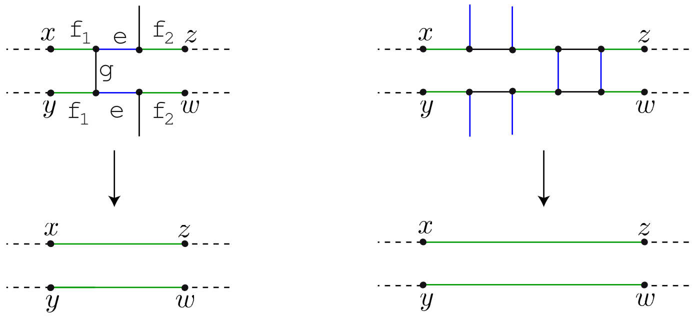

Figure 29: Two operations that simplify Z without changing its curvature.

operation of the right half of Figure 29 repeatedly to eliminate pairs of d-edges, thus leaving an auxiliary graph Z' with very few edges: the vertices u, v (recall that these where the common endpoints of P and P') and their incident c edges split Z' into two parts, each of which contains no b edge that is not incident with u or v, and contains at most one pair of d-edges not incident with u, v. In fact, it must contain precisely one pair of d-edges not incident with u, v, for any attempt to construct such a part without such d-edges leads to a contradiction to the spin-behaviour as displayed in Figure 30. It follows that each of those parts must be one of the three graphs in the upper row of Figure 31. Now as

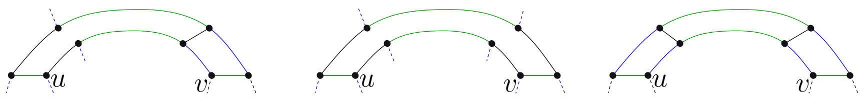

Figure 30: Some failed attempts to create a building block for Z'. Note that the c edges in the middle are not spin-preserving as they should. The dashed edges are not part of Z' and are only a visual aid.

u and v are incident with one edge of each colour, Z' is either the union of the two leftmost graphs of the upper row of Figure 31, or the union of two copies of the rightmost one. Thus Z' is one of the graphs of the bottom row. In both cases, it has precisely six 4-gons and two 6-gons, accounting for a contribution of 12 'curvature' units to Euler's formula (17). To sum up, the contribution of the faces of H' that lie inside Z to the left part of (17) equals 12. The contribution of the remaining faces of H' is strictly negative, since each such face has size larger than 6. This contradicts (17) in the case when the dominant colour is a spin-preserving one.

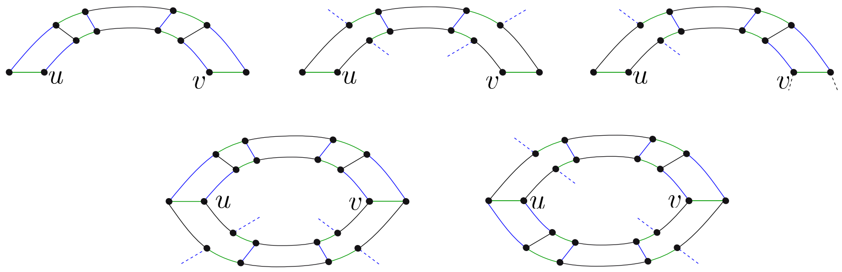

Figure 31: The possible combinations for Z'.

We now turn to the case when the dominant colour of C is d. In this case we can still use the reducing operation in the right half of Figure 29. In addition, we can use two further operations, shown in Figure 32, that can recursively suppress any path of the form *cdbdc*, or *bdcdcdbdb*, and similarly with the roles of b and c interchanged. After performing such operations as often as possible, we are left with a graph Z' which again we think of as the union of two parts joined at u and v. Each of those parts has now at most three pairs of d-edges not incident with u or v, for if it has more then one of the above operations can be applied. We are left with few possibilities, and an easy case study (for which Figure 32 might still be helpful) shows that none of these two parts can contain more than three 4-gons, again leading to a contradiction of (17). This completes the proof of (58).

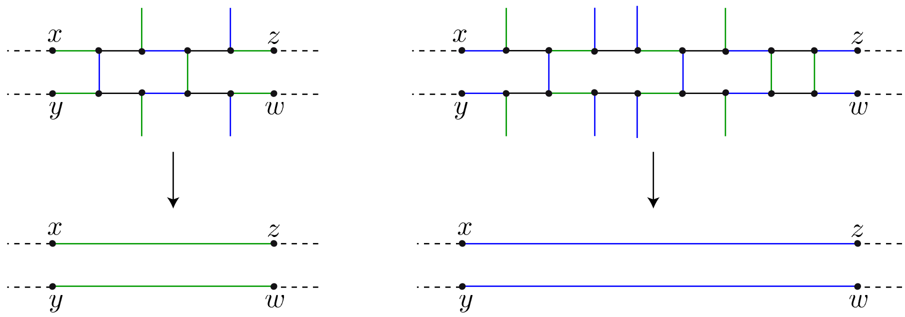

Figure 32: Further reducing operations for the case that d is the dominant colour of C.

We can now apply (58) to exclude d as a dominant colour:

The spin-reversing colour d is not dominant in any shortest dividing cycle of G. (59)

Indeed, if d was the dominant colour of a shortest dividing cycle C, then C would contain a d-edge e incident with both a b and a c edge of C. But then the colour-automorphism of G that exchanges the endvertices of e would map C to a shortest dividing cycle C' that crosses C, contradicting (58).

In fact, we can prove a bit more about dominant colours:

If C, B are shortest dividing cycles of G then they have the same dominant colour. (60)

Indeed, if not, then by (59) we can assume that the dominant colour of C is c and that of B is b. Note that as all cycles must be 3-coloured, C must have a subpath of the form $cdcbc$ and B must have a subpath of the form $bdbcb$. Translating one of those paths to the other as in Figure 33 we obtain a crossing of two shortest dividing cycles. By (3) this implies the existence of a new shortest dividing cycle containing a subpath of the form $dbcd$. But this contradicts Proposition 9.16 and so (60) is proved.

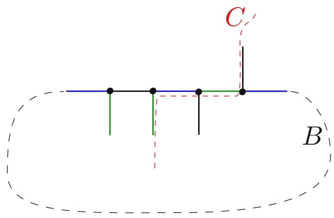

Figure 33: The crossing cycles in the proof of (60).

By (59) and (60) all shortest dividing cycles of G have the same, spin-preserving, dominant colour. Assume from now one that this colour is c. We can now describe the shortest dividing cycle precisely:

Every shortest dividing cycle of G is induced by the word $(c(bc)^n d)^{2m}$ for fixed $n, m \geq 1$. (61)

Let C be a shortest dividing cycle, and define a *bc-interval* of C to be a maximal subpath of C not containing a d-edge. Note that (61) is equivalent to saying that all bc-intervals of shortest dividing cycles of G have the same length $2n+1$. Suppose, to the contrary, that C has a d-edge e such that the two bc-intervals on either side of e have different lengths. Then, observing the spin behaviour of the edges, it is easy to see that the colour-automorphism of G that exchanges the endvertices of e maps C to a shortest dividing cycle that crosses C (Figure 34), contradicting (58). In particular, C does not contain a dcd subpath, for such a subpath contains a bc-interval of length 1, and C has some bc-interval of length at least 3 as it is 3-coloured. This proves that every shortest dividing cycle C has the desired form $(c(bc)^n d)^m$, with $n \geq 1$, and it only remains to show that n and m cannot vary for a different shortest dividing cycle D. The fact that n cannot vary can be proved with a similar argument, by considering a colour-automorphism that maps a d-edge of D to a d-edge of C (see Figure 34 again). It follows that m cannot vary either, since all shortest dividing cycles have the same length by definition. The fact that C has an even number of d-edges, giving rise to the exponent $2m$, can easily be proved by observing the spin behaviour of the edges.

The subgroup Γ_2 and the subgrpaph G'_2.

Define the words $B := bcdcb$ and $Z := c(bc)^{n-1}$, where n is supplied by (61). Note that every face boundary of G is, by (54), of the form $\ldots dBdBdB \ldots$. Moreover, every shortest dividing cycle is, by (61), of the form $(BZdZ)^m$ (compare this with (38)).

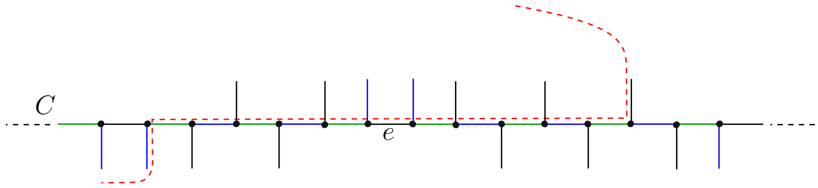

Figure 34: Finding a crossing if two bc-intervals of C have different lengths in the proof of (61).

Let b^* be the element of Γ corresponding to the word B and let z be the element of Γ corresponding to the word Z. Let Γ_2 be the subgroup of Γ generated by $\{b^*, z, d\}$.

As in (40) we still have:

for every shortest dividing cycle C of G and every coset Δ of Γ_2 in G, at most one of the sides of C contains elements of Δ. (62)

To prove this, one can use similar arguments as in the proof of (40), however, there is an easier way: it is easy to check, just by observing the spin behaviour, that no path induced by a word in the letters b^*, z, d can cross a cycle C induced by the word $(BZdZ)^m$, the word inducing the shortest dividing cycles.

Next, we check that

for every path P in G induced by Z there are precisely two shortest dividing cycles C, D containing P. Moreover, $C \cap D = P$, that is, C and D have no common edge outside P. (63)

The first part of this assertion is much easier to prove than the corresponding assertion (36) in the previous section: it follows immediately from (61). The second part can be proved like (36): if C, D have a common edge outside P, then using a subpath of each we can form a cycle K shorter than $|C|$, and using Euler's formula as in (58) we can prove that K is dividing, a contradiction.

The following assertion strengthens (62) and can be proved like (41). A *metaedge* is a path of G induced by one of the words B, Z, d.

Any two metaedges of Γ_2 are independent. (64)

This means that, as in earlier sections, the Cayley graph G_2 of Γ_2 with respect to the generating set $\{b^*, z, d\}$ has a topological embedding in G: we can obtain G_2 from G by substituting, for every two vertices x, y of G_2 that are adjacent by a b^* or z edge, the x–y path in G induced by B, Z with an x-y edge of the corresponding colour b^*, z. This yields indeed a topological embedding of G_2 in G since by (64) all these paths are independent. Starting with G_2 and replacing each b^* or z edge back by the corresponding path induced by B or Z we obtain a subdivision $G_2' \subseteq G$ of G_2 that will be useful later.

Note that every edge of G_2 corresponds to an involution. Moreover, every b^* or d-edge of G_2 is spin-reversing while every z edge is spin-preserving; this can be deduced from the spin behaviour of the original edges of G.

The planar presentation of G

As in earlier sections we will express G as a union of copies of G_2', and then apply Theorem 8.2 in order to deduce a presentation of G from a presentation of G_2.

In this section we define a *society* as in Definition 9.11 of the previous section, except that we now base this definition on the word $(BZdZ)^m$, which induces the shortest dividing cycles in the current case.

To begin with, we claim that

$$\text{every shortest dividing cycle } C \text{ of } G \text{ contains precisely two distinct societies.} \tag{65}$$

Indeed, recall that C is induced by the word $(c(bc)^n d)^{2m}$ by (61). Note that for every two 'consecutive' d-edges d_1, d_2 on C, the two b edges incident with d_1 lie in one side of C, while the b edges incident with d_2 lie in the other side of C. Thus there is a bipartition $\{D_1, D_2\}$ of the set of d-edges of C such that all b edges incident with an element of D_i lie in the same side of C. It is now straightforward to check that the endvertices of all the edges in each of the D_i lie in a common society of C, and these two societies are distinct for $i = 1, 2$.

This allows us to create a structure tree T on the set of left Γ_2 cosets in Γ as we did in Section 9.2.1: join two such cosets with an edge, if the corresponding copies of G_2' share a shortest dividing cycle of G. It follows from (62) that T is acyclic, and from (65) that it is connected. Thus once more, we can apply Theorem 8.2, with the H_i being the vertices of T and the F_i being the shortest dividing cycles of G giving rise to the edges of T. This yields that given any presentation of Γ_2 we can transform it into a presentation of Γ by replacing any occurrence of the letters b^*, z by the corresponding words $B = bcdcb$ and $Z = c(bc)^{n-1}$ and adding the involution relations b^2, c^2.

So let us find a presentation of Γ_2. We distinguish three cases according to the connectivity and number of ends of G_2.

Case I: G_2 is 3-connected and finite or 1-ended. In this case we apply Theorem 7.6 (iii) to G_2, which yields $G_2 \cong Cay\left\langle b^*, z, d \mid {b^*}^2, z^2, d^2, (b^*d)^r, (b^*zdz)^m \right\rangle$, $r \geq 2, m \geq 1$. Note that b^* and d are interchangeable in this presentation, while z, being the only spin-preserving colour, plays a special role.

Substituting b^*, z as suggested above we obtain the following presentation for Γ:

$$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (cbdbcd)^k; (c(bc)^n d)^{2m} \right\rangle, k \geq 2, n, m \geq 1.$$

This is a planar presentation: the relation $(cbdbcd)^k$ corresponds to face-boundaries, and the last relation corresponds to shortest dividing cycles. We will prove in the next subsection that $n + m \geq 3$ must hold in this case.

Case II: G_2 is 3-connected and multi-ended. In this case we can apply Lemma 9.2, which yields that G_2 must have a 2-coloured cycle involving the only spin-preserving colour z. We have already characterized the graphs of this type: Theorem 9.6 yields that $G_2 \cong Cay\left\langle b^*, z, d \mid {b^*}^2, z^2, d^2, (b^*d)^q, (dzb^*z)^m; (zb^*)^{2r} \right\rangle$, $r, m, q \geq 2$, or (exchanging b^* and d in the above presentation and rearranging) $G_2 \cong Cay\left\langle b^*, z, d \mid {b^*}^2, z^2, d^2; (b^*d)^q, (dzb^*z)^m; (zd)^{2r} \right\rangle$, $r, m, q \geq 2$.

Substituting b^*, z as above we obtain in the first case

$$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2; (cbdbcd)^q; (c(bc)^n d)^{2m}, (c(bc)^{n+1} d)^{2r} \right\rangle \text{ with } n \geq 1 \text{ and } r, m, q \geq 2,$$

where we used the fact that $(c(bc)^{n-1}bcdcb)^{2r} = (c(bc)^n dcb)^{2r} = (c(bc)^{n+1}d)^{2r}$. In the second case we obtain

$$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2; (cbdbcd)^q; (c(bc)^n d)^{2m}, (c(bc)^{n-1}d)^{2r} \right\rangle, \text{ with } n, r, m, q \geq 2.$$

In the latter presentation we are demanding $n \geq 2$ because if $n = 1$ then G has 2-coloured cycles contrary to our assumption.

Note that these two presentations are the same, as can be seen by exchanging m with r and n with $n+1$. Thus we omit the first one. We have now obtained possibility (ii) of Theorem 9.15.

Case III: G_2 is not 3-connected.

It follows from Lemma 4.1 that G_2 must be 2-connected. Thus, in this case G_2 is one of the graphs of Theorem 5.1. Since it has three generators, it has to belong to one of the types (iv)–(ix) of that theorem. We will be able to eliminate most of these types as a possibility for G_2, leaving only type (v) as a possibility.

It is made clear in [13] that for every graph of type (iv), all edges participating in the 4-cycles induced by the relation $(bc)^2$ must preserve spin in any embedding. But our G_2 has an embedding in which two of the colours reverse spin, and so G_2 cannot be of this type.

For a graph of type (vi), the colour participating in both relations has the property that any edge e of that colour is a hinge, and it separates the graph in two components each of which sends two edges of the same colour to e. Let us check that G_2 cannot have a hinge $e = uv$. If uv is coloured d or b^*, then note that it is contained in a basic cycle C, and each of its endvertices u, v is incident with a basic cycle $D_u, D_v \neq C$. Now note that $(C \cup D_u \cup D_v) - \{u, v\}$ is connected, which means that if $G_2 - \{u, v\}$ is disconnected, then G_2 was already disconnected before removing $\{u, v\}$, a contradiction. If uv is coloured z instead, then note that the d-edge incident with u and the b^*-edge incident with v lie in a common component of $G_2 - \{u, v\}$ because there is a basic cycle containing these two edges and uv. But this contradicts the property of the separating colour described above. Thus in all cases we obtain a contradiction if G_2 is of type (vi).

It also follows from the analysis in [13] that for every consistent embedding of a graph of type (vii) at least two colours preserve spin, and so again the embedding of G_2 we have implies that G_2 cannot be of that type either.

Suppose now G_2 is of type (viii), which means that $G_2 \cong Cay \left\langle b^*, z, d \mid {b^*}^2, z^2, d^2, (dzb^*z)^m \right\rangle$. Then replacing b^* and z as above, we obtain the following presentation for G: $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (c(bc)^n d)^{2m} \right\rangle, n \geq 2$. Note however, that the latter presentation is identical with that of Theorem 5.1 (vii); thus, by the converse implication of that theorem, G is not 3-connected in this case contradicting our assumption.

If G_2 is of the degenerate type (ix), then there must be a pair of edges of G_2 that have common endvertices. No z-edge can participate in such a pair, because ZB and Zd are both subwords of the word $ZdZB$ inducing the shortest dividing cycles. But if a d and a b^* edge form such a pair, then the corresponding cycle of G bounds a hexagonal face, which cannot be the case by Proposition 9.17. Thus G_2 is not of type (ix) either.

The only possible candidate left is type (v), and this possibility can indeed occur as we will see in the next subsection. In this case we have either

$G_2 \cong Cay \left\langle b^*, z, d \mid {b^*}^2, z^2, d^2, (b^*z)^{2r}, (zb^*zd)^m \right\rangle$, $r, m \geq 2$, or
$G_2 \cong Cay \left\langle b^*, z, d \mid {b^*}^2, z^2, d^2, (dz)^{2r}, (zb^*zd)^m \right\rangle$, $r, m \geq 2$,
depending on which of the two spin reversing colours b^*, d forms 2-coloured cycles with z. Replacing b^* and z as above, we obtain the following two presentations respectively:

$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2; (c(bc)^{n+1}d)^{2r}, (c(bc)^n d)^{2m} \right\rangle$, $n \geq 1$, $r, m \geq 2$, or
$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2; (c(bc)^{n-1}d)^{2r}, (c(bc)^n d)^{2m} \right\rangle$, $n, r, m \geq 2$.

Again, there is no difference between these two presentations except for the naming of the parameters, and we can omit the first one. We have thus obtained possibility (iii) of Theorem 9.15.

This completes the proof of the forward implication of Theorem 9.15.

The converse implication

In this section we show that for every presentation as in Theorem 9.15 the corresponding Cayley graph is planar and 3-connected. Our approach is very similar to that of the proof of Theorem 9.10 (page 58), and it more or less goes through the proof of the forward implication the other way round.

Consider first a Cayley graph of **type (i)**:

$$G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (cbdbcd)^k, (c(bc)^n d)^{2m} \right\rangle, k \geq 2, \; n, m \geq 1.$$

We are going to construct an embedding of G. For this, consider first the auxiliary Cayley graph $G_2 = Cay \left\langle b^*, z, d \mid {b^*}^2, z^2, d^2, (b^*d)^k, (b^*zdz)^m \right\rangle$. Then Theorem 7.6 (iii) yields an embedding σ_2 of G_2 in which only z preserves spin, and G_2 is 3-connected by Theorem 7.1. Modify G_2 into a further auxiliary graph G'_2 by replacing each b^* edge of G_2 by a path of length 5 with edges coloured $bcdcb$ (recall that this was the word B), and replacing each z edge of G_2 by a path of length $2n-1$ with edges coloured $c(bc)^{n-1}$ (the word Z). Note that these words are symmetric, and so it does not matter at which end of those paths we start colouring the new edges.

Note that every cycle of G_2 induced by $(b^*zdz)^m$ has turned into a cycle of G'_2 induced by the word of (61). Moreover, every such cycle bounds a face of G'_2. We call these cycles the *basic cycles* of G'_2. A *corner* of G'_2 is a vertex of degree 3. By (65) every basic cycle B of G'_2 contains precisely two distinct 'societies'. Note that precisely one of these societies S_B consists of corners of G'_2, while the elements of the other society T_B are non-corners. Now for every basic cycle B of G'_2, construct a copy H_B of G'_2 such that $B \subseteq H_B$ and the elements of T_B are corners of H_B, and embed H_B in the face of G'_2 bounded by B. Repeat this inductively ad infinitum for each of the newly appeared basic cycles. Let G be the resulting plane graph and σ its embedding. It follows from Sabidussi's Theorem that G is a Cayley graph: assertions (48) and (49) are still valid, and imply that the colour-automorphisms of G act transitively on its vertices.

We will now prove that G is 3-connected using Lemma 4.2. For this, let $(H_i)_{i \in \mathbb{N}}$ be an enumeration of the copies of G'_2 in G, and for every i let K_i be the set of corners of H_i. The requirements (i) and (ii) of Lemma 4.2 are satisfied for $k = 3$ since G_2 is 3-connected. We will show that the third requirement (iii) is satisfied unless $n = m = 1$. Indeed, if the latter is the case, then we have $Z = c$ and so the z edges of G_2 are also edges of G. Moreover, any basic cycle C is induced by $(b^*zdz)^m$, and so after subdividing it only contains one metaedge.

Note that such a metaedge contains all elements of one of the societies of C. This implies that removing the first and last edge of this metaedge disconnects the two copies of G'_2 that share C in G, which means that G is not 3-connected in this case.

If on the other hand one of n, m is greater than 1, then each basic cycle C contains more than one metaedge, and it follows that C contains at least four edges joining its two societies. Thus requirement (iii) of Lemma 4.2 is satisfied too, and so G is 3-connected in this case.

Next, we claim that G has no 2-coloured cycle unless $n = m = 1$. We begin with showing that bd and cd have infinite order independently of the values of n and m. For this, note that for every b-edge e of G there is a b^* metaedge starting with e, and a basic cycle C_e containing this metaedge. The spin behaviour implies that the two d-edges incident with e lie in distinct sides of C_e. Now given a path P of G the edges $b_1 d_1 b_2 d_2 \ldots$ of which alternate in the colours b, d, consider the cycles C_{b_i} obtained as above, and note that $C_{b_i} \neq C_{b_{i+1}}$; indeed if $C_{b_i} = C_{b_{i+1}}$ then d_i is a *chord* of C_{b_i}, i.e. an edge having both vertices on that cycle, but a basic cycle cannot have a chord since $k > 1$. Note moreover that, by our previous remark about the spin, $C_{b_{i-1}}$ and $C_{b_{i+1}}$ lie in distinct sides of C_{b_i}. This immediately implies that P cannot be a closed path, and so bd has infinite order indeed.

Similarly, note that for every path P induced by $b\ c(bc)^n\ b$ there is a basic cycle C that contains the interior of P and the first and last edge of P lie in distinct sides of C. Adapting the above argument we conclude that bc has infinite order too.

If $n > 1$ then for every dcd-path there is a basic cycle C containing only the middle edge, and the two incident d-edges lie in distinct sides of C, which again allows us to prove that cd has infinite order. If $n = 1$, then note that if G has a finite cd-cycle C, then C is also a cycle of G_2 since $Z = c$ in this case. Moreover, since c preserves spin, some of the b^*-edges incident with C lie in one of its sides and some of them lie in its other side; the situation around C looks like Figure 19 after adapting the colours. If no such edge is a chord of C then, easily, C is a dividing cycle, which contradicts the fact that G_2 is at most 1-ended. If C has a chord $e = uv$, then consider the shortest subarc P of C with endvertices u, v. If $||P|| > 3$, then there is a colour-automorphism G_2 that fixes C, maps u to a vertex u' of P that has an incident b^*-edge in the side of C in which e also lies, and maps v to a vertex v' outside P (use Figure 19 again to see this). But then the edges $uv, u'v'$ must cross, yielding a contradiction.

Note that $||P|| \neq 1$ because $k > 1$. If $||P|| = 3$, and so P is induced by cdc, then $P \cup e$ is a face-boundary of G_2. But every face-boundary of G_2 is induced by one of the relators $(b^*d)^k, (b^*zdz)^m$ in its presentation. Thus m must equal 1 in this case. To sum up, we proved that if $n+m \geq 3$ then cd has infinite order (in G) too. The interested reader will be able to check that this assertion is best possible: if $n = m = 1$ then cd must have finite order in G. We do not need this fact for our proof though since $n = m = 1$ is already forbidden because of the connectivity.

The fact that G has the desired presentation now follows from the forward implication of Theorem 9.15, which we have already proved, since we checked that G has the desired properties.

Consider now a presentation of **type (ii)**:

$$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (bdbcdc)^q; (c(bc)^n d)^{2m}, (c(bc)^{n+1}d)^{2r}\right\rangle, n \geq 1, r, m, q \geq 2$$

We can then construct an embedding of G by the same method, except that we have to start with a different G_2: this time we let $G_2 \cong Cay\left\langle b^*, z, d \mid {b^*}^2, z^2, d^2, (b^*d)^q, (dzb^*z)^m, (zb^*)^{2r}\right\rangle$, and it follows from Theorem 9.6 that G_2 is again planar and 3-connected, and has the desired spin behaviour. Otherwise, the construction remains the same.

The parameter m is now large enough to make sure that the requirements of Lemma 4.2 for $k = 3$ are satisfied in all cases, and so G is 3-connected. Moreover, we can prove that G has no 2-coloured cycle by the same arguments, and our task is made easier by the fact that $m \geq 2$ now.

Finally, consider a presentation of **type (iii)**:

$$G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2, (c(bc)^{n-1}d)^{2m}, (c(bc)^n d)^{2r}\right\rangle, n, r, m \geq 2.$$

We use the same approach again, except that it is now trickier to show that the resulting Cayley graph G is 3-connected. In this case we start our construction letting $G_2 \cong Cay\left\langle b^*, z, d \mid {b^*}^2, z^2, d^2, (b^*z)^{2r}, (zb^*zd)^m\right\rangle$, which corresponds to type (v) of Theorem 5.1 and has connectivity 2. It is proved in [13, Proposition 5.7] that this G_2 has the following properties:

1. G_2 has a consistent embedding σ_2 in which z preserves spin while b^*, d reverse spin (see [13, Figure 10]). In this embedding, each vertex is incident with two faces bounded by a cycle induced by the relator $(zb^*zd)^m$ and one face that has infinite boundary.

2. G_2 cannot be separated by removing two edges e, f unless both e, f are coloured d, and it cannot be separated by removing a vertex and an edge e unless e is coloured d;

3. If a pair of vertices s, t of a cycle C of G_2 induced by $(zb^*zd)^m$ separates G_2, then both s, t are incident with a d-edge of C;

4. G_2 has no hinge;

5. for every cycle C of G_2 induced by the word $(zb^*zd)^m$, and every b^* edge vw of C, there is a v–w path in G_2 meeting C only at v, w, and

6. If two cycles C, D of G_2 induced by the word $(zb^*zd)^m$ share an edge uv, then there is path from C to D in $G_2 - \{u, v\}$.

Construct G and an embedding of its using G_2 and σ_2 (provided by (i) as in the previous cases.

Although $\kappa(G_2) = 2$, we will be able to prove

Proposition 9.18. *G is 3-connected.*

Proof. We will apply Lemma 4.2, with K_i being, as usual, the set of corners of a copy H_i of G_2' in G, and $(H_i)_{i \in \mathbb{N}}$ being an enumeration of these copies.

So let us check that K_i is 3-connected in G. To begin with, note that H_i is the union of basic cycles, that is, cycles induced by the word $(zb^*zd)^m$, because this word contains all colours, and so every edge is in a basic cycle.

For every basic cycle C of H_i, $K_i \cap C$ is 3-connected in G. (66)

Indeed, suppose there are vertices $s, t \in V(G)$ separating two vertices $x, y \in K_i \cap C$. Since C is a cycle, both s, t must lie on C, with each of the two components of $C - \{s, t\}$ containing one of x, y. It follows easily from (ii) that none of s, t can be an interior vertex of a b^* or z metaedge contained in C, and so $s, t \in K_i \cap C$. By (iii) each of s, t is incident with a d-edge of C, and these two d-edges d_s, d_t are distinct by (iv). Note that by the choice of the words Z, B, the b^* metaedge b_s containing d_s is contained in C, and note also that d_t is not contained in b_s as the latter contains only one d-edge. Thus, each of the two components S_1, S_2 of $C - \{s, t\}$ contains one of the endvertices p, q of b_s. Now by (v), there is a path P joining p to q in the copy of G'_2 sharing C with H_i which path has no interior vertex on C. Thus $s, t \notin V(P)$, and P connects the two components of $C - \{s, t\}$. This contradicts our assumption that s, t separate x, y in G, and proves (66).

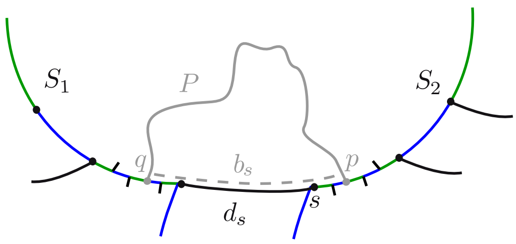

Figure 35: The path P accounting for the 3-connectedness of G.

Next, we claim that

for every $v, w \in K_i$, there is a finite sequence of basic cycles $C_0, \dots C_k$ of H_i such that $v \in C_0$, $w \in C_k$, and C_i shares a metaedge with C_{i+1} for every relevant i. (67)

Indeed, since the word $(zb^*zd)^m$ inducing the basic cycles involves all three colours, any two edges of H_i sharing a vertex lie in a common basic cycle, and so (67) can be proved by induction on the length of a v–w path in G_2.

Now (vi) yields that for every two basic cycles C, D of H_i sharing a metaedge, C is 3-connected to D in H_i. Combining this with (67) and (66) implies that K_i is 3-connected in G: given $v, w \in K_i$, and a sequence C_i as in (67), we can construct a v–w path in G avoiding any fixed pair of vertices by combining paths joining two suitable vertices of C_i for every relevant i; see the proof of Lemma 4.2 for a more detailed exposition of this argument.

Thus we have proved that the set of corners K_i of any copy H_i of G'_2 in G is 3-connected in G. Moreover, for any two copies H_i, H_j of G'_2 sharing a basic cycle C, it is clear that K_i is 3-connected to K_j because C contains more than 2 pairwise disjoint paths joining its two societies, and each of K_i, K_j contains a distinct society of C. We can thus apply Lemma 4.2 to prove that G is 3-connected. □

This completes the proof of the converse implication of Theorem 9.15.

The graphs of the last type (iii) also have the surprising property that we already encountered in Corollary 9.13 that none of their faces is bounded by a cycle:

Corollary 9.19. *For every $n, r, m \geq 2$, the graph* $G \cong Cay\left\langle b, c, d \mid b^2, c^2, d^2; (c(bc)^{n-1}d)^{2m}, (c(bc)^n d)^{2r} \right\rangle$ *is 3-connected and has no finite face boundary.*

Proof. Recall that in order to construct a G as above, we started with the graph G_2 and its embedding σ_2, and inductively glued copies of G_2 inside the basic cycles. By (i) σ_2 had faces with infinite boundary, and these faces were left intact by our construction, except for subdividing the edges in their boundary. Thus G also has at least one face with infinite boundary at every vertex. But as two of the colours b, c preserve spin in G, such a face can be mapped to any other by a colour-automorphism of G, implying that every face has infinite boundary. □

10 Outlook

In Section 1.2 we showed two examples of a *word extension*. This operation was implicit throughout the paper whenever we used a presentation of G_2 to obtain one of G. It would be interesting to study word extensions in greater generality. One aim could be to refine Stallings' theorem into a theorem about all Cayley graphs rather than a theorem about their groups: prove that every multi-ended Cayley graph can be obtained from simpler Cayley graphs by means of certain operations including word extensions. A modest first step in this direction would be to prove that every multi-ended cubic Cayley graph is a word extension of a cubic Cayley graph of a subgroup. To put it in a different way:

Problem 10.1. *Every multi-ended cubic Cayley graph $G = Cay(\Gamma, S)$ contains a subdivision of a cubic Cayley graph G_2 of a proper subgroup of Γ, and G is the union of the translates of G_2 under Γ.*

One of the most important ideas in this paper was the use of societies (Definition 9.11), and realising that they can be found in each of our multi-ended graphs. It would thus be interesting to prove that they appear in all multi-ended planar Cayley graphs, not just the cubic ones. It also seems promising to try to generalise the concept to non-planar Cayley graphs, perhaps using the ideas of [11]. For example, the big cycle of Figure 1 (iii) accommodates two dual societies, each being the set of its vertices sending edges to one of its sides. This graph was the result of a word extension of the graph of Figure 1 (i), where a generator a was replaced using the word a_*^2; see Section 8.1. But we could have used the word a_*^3 instead, in which case each edge of the original cycle would be subdivided into three, we would have three 'dual' societies, and removing the cycle would leave three infinite components. Thus the resulting graph is not planar, but still it can be analysed by our methods.

Theorem 1.1 shows that every cubic planar Cayley graph admits a planar presentation with at most 6 relators. This motivates

Problem 10.2. *Let $f(n) \in \mathbb{N} \cup \{\infty\}$ be the smallest cardinal such that every n-regular planar Cayley graph admits a planar presentation with at most $f(n)$ relators. Is $f(n)$ finite for each n? If yes how fast does it grow with n?*

In this paper we constructed surprising examples of planar 3-connected Cayley graphs in which no face is bounded by a cycle. Thus one can ask

Problem 10.3. *Is there, for every $k \in \mathbb{N}$, a planar k-connected Cayley graph in which no face is bounded by a cycle?*

Acknowledgements

I am very grateful to Bojan Mohar, for triggering my interest in the topic and for later valuable discussions, and to Martin Dunwoody, for discussions leading to improvements in the final version.

References

[1] L. Babai. Automorphism groups, isomorphism, reconstruction. In *Graham, R. L. (ed.) et al., Handbook of combinatorics. Vol. 2. Amsterdam: Elsevier (North-Holland)*, pages 1447–1540. 1995.

[2] L. Babai. The growth rate of vertex-transitive planar graphs. In *Proceedings of the eighth annual ACM-SIAM symposium on Discrete algorithms*, pages 564–573, New Orleans, Louisiana, United States, 1997. Society for Industrial and Applied Mathematics.

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

[4] C. P. Bonnington and M. E. Watkins. Planar embeddings with infinite faces. *Journal of Graph Theory*, 42(4):257–275, 2003.

[5] Q. Cui, J. Wang, and X. Yu. Hamilton circles in infinite planar graphs. *J. Combin. Theory (Series B)*, 99(1):110–138, 2009.

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

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

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

[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 and B. Krön. Vertex cuts. Preprint.

[12] A. Georgakopoulos. A group has a flat cayley complex if and only if it has a VAP-free cayley graph. Preprint 2010.

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

[14] A. Georgakopoulos. Word extensions of groups. In preparation.

[15] A. Georgakopoulos. Infinite hamilton cycles in squares of locally finite graphs. *Advances in Mathematics*, 220:670–705, 2009.

[16] A. Georgakopoulos and M. Hamann. In preparation.

[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] B. Krön. Infinite faces and ends of almost transitive plane graphs. Preprint.

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

[20] K. Mitchell. Constructing semi-regular tilings. *Talk given at the Spring 1995 Meeting of the Seaway Section of the MAA.* `http://people.hws.edu/mitchell/tilings/Part1.html`.

[21] B. Mohar. Personal communication.

[22] B. Mohar. Tree amalgamation of graphs and tessellations of the cantor sphere. *Journal of Combinatorial Theory. Series B*, 96(5):740753, 2006.

[23] G. Sabidussi. On a class of fixed-point-free graphs. *Proc. Am. Math. Soc.*, 9:800–804, 1958.

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

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

[26] 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.