The planar cubic Cayley graphs of connectivity 2

Agelos Georgakopoulos*

Technische Universität Graz
Steyrergasse 30, 8010
Graz, Austria

March 2, 2011

Abstract

We classify the planar cubic Cayley graphs of connectivity 2, providing an explicit presentation and embedding for each of them. Combined with [9] this yields a complete description of all planar cubic Cayley graphs.

1 Introduction

This paper is part of a series of related papers on planar Cayley graphs [8, 9, 10, 11]. A major aim of this project is to provide a complete description of the *cubic* planar Cayley graphs, i.e. those in which every vertex is adjacent with precisely three other vertices. This analysis provides surprising new examples, contradicting past conjectures, but also a good insight into planar Cayley graphs in general; see [9].

The aim of the current paper is a classification of the planar cubic Cayley graphs of connectivity 2, yielding an explicit presentation and embedding for each of those graphs. The *connectivity* of a graph G is the smallest cardinality of a set of vertices separating G. Cayley graphs of connectivity 1 are easy to describe. It is very common [6, 7] to consider graphs of connectivity 2 separately from graphs of higher connectivity when studying planar Cayley graphs, and this is so for a good reason: by a classical theorem of Whitney [13, Theorem 11], planar graphs of connectivity at least 3 have a unique, up to homeomorphism, embedding in the sphere S^2, and can thus be analysed taking advantage of this fact.

Our main result is

Theorem 1.1. *Let G be a planar cubic Cayley graph of connectivity 2. Then precisely one of the following is the case:*

  (i) $G \cong Cay \langle a, b \mid b^2, (ab)^n \rangle$, $n \geq 2$;
 (ii) $G \cong Cay \langle a, b \mid b^2, (aba^{-1}b^{-1})^n \rangle$, $n \geq 1$;
(iii) $G \cong Cay \langle a, b \mid b^2, a^4, (a^2b)^n \rangle$, $n \geq 2$;
 (iv) $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (bc)^2, (bcd)^m \rangle$, $m \geq 2$;

(v) $G \cong Cay \langle b,c,d \mid b^2, c^2, d^2, (bc)^{2n}, (cbcd)^m \rangle$, $n, m \geq 2$;

(vi) $G \cong Cay \langle b,c,d \mid b^2, c^2, d^2, (bc)^n, (bd)^m \rangle$, $n, m \geq 2$;

(vii) $G \cong Cay \langle b,c,d \mid b^2, c^2, d^2, (b(cb)^n d)^m \rangle$, $n, m \geq 2$;

(viii) $G \cong Cay \langle b,c,d \mid b^2, c^2, d^2, (bcbd)^m \rangle$, $m \geq 1$;

(ix) $G \cong Cay \langle b,c,d \mid b^2, c^2, d^2, (bc)^n, cd \rangle$, $n \geq 1$ (degenerate cases with redundant generators and G finite).

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

Except for the above presentations we also construct embeddings of these graphs such that the corresponding group action on the graph carries facial walks to facial walks (Corollary 5.1); see also Section 2.3. Tables 1 and 2 summarise some information about these embeddings and some basic properties of the graphs, yielding a more structured presentation of the various possibilities of Theorem 1.1. The interested reader will also find further information, not included in these tables, throughout the course of the proofs. Moreover, in the last section of the paper we point out some interesting corollaries, which are extended in [9] to all cubic planar Cayley graphs.

This paper contributes to the complete classification of the cubic planar Cayley graphs of [9] not only by settling the special case of graphs of connectivity two, but also by providing building blocks for the construction of some of the 3-connected ones. For example, the graphs of type (vi) of Theorem 1.1 and their embeddings that we construct here are used in [9] to produce 3-connected Cayley graphs that have no finite face boundaries, thus contradicting a conjecture of Bonnington and Watkins [3].

In many cases we obtain embeddings of our graphs in the sphere using only the fact that their connectivity is 2, without assuming a priori that the graphs are planar. Thus we obtain the following result, which describes the non-planar cubic Cayley graphs of connectivity 2.

Corollary 1.2. *Let G be a non-planar cubic Cayley graph of connectivity 2. Then one of the following is the case:*

(i) $G \cong Cay \langle a,b \mid b^2, a^n, \ldots \rangle$, $n > 2$; *or*

(ii) $G \cong Cay \langle b,c,d \mid b^2, c^2, d^2, (bc)^k, \ldots \rangle$, $k \geq 3$, G *has no hinge, and every bc cycle contains a pair x, y of vertices separating the graph such that $x^{-1}y$ is an involution.*

2 Definitions

2.1 Cayley graphs and group presentations

We will follow the terminology of [5] for graph-theoretical terms and that of [2] 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

$G = Cay\langle a, b \mid b^2, \ldots\rangle$, G is planar, and $\kappa(G) = 2$		
$a^n = 1$.	$a^n \neq 1$. Then G has a consistent embedding σ in which a preserves spin.	
$\Gamma \cong \langle a, b \mid b^2, a^4, (a^2b)^n\rangle$, $n \geq 2$. Thus G has a consistent embedding in which a reverses spin and b preserves spin.	$\Gamma \cong \langle a, b \mid b^2, (ab)^n\rangle$, $n \geq 2$ and b preserves spin in σ.	$\Gamma \cong \langle a, b \mid b^2, (aba^{-1}b^{-1})^n\rangle$, $n \geq 1$ and b reverses spin in σ.

Table 1: Classification of the cubic planar Cayley graphs with 2 generators and connectivity 2 (types (i) to (iii) of Theorem 1.1).

$G = Cay\langle b, c, d \mid b^2, c^2, d^2, \ldots\rangle$, G is planar, and $\kappa(G) = 2$		
$(bc)^n = 1$ (i.e. G has a 2-coloured cycle).		G has no 2-coloured cycle.
G has no hinge, $\Gamma \cong \langle b^2, c^2, d^2 \mid (bc)^2, (bcd)^m\rangle$, $m \geq 2$, and G has a consistent embedding in which all edges preserve spin.	G has no hinge, $\Gamma \cong \langle b^2, c^2, d^2 \mid (bc)^{2n}, (cbcd)^m\rangle$, $n, m \geq 2$, and G has a consistent embedding in which c preserves spin and b, d reverse spin.	G has no hinge, $\Gamma \cong \langle b^2, c^2, d^2 \mid (b(cb)^n d)^m\rangle$, $n, m \geq 2$, and G has an embedding in which all edges preserve spin.
	G has a hinge, $\Gamma \cong \langle b^2, c^2, d^2 \mid (bc)^n, (bd)^m\rangle$, $n, m \geq 2$, and G has a consistent embedding in which all edges reverse spin.	G has a hinge $\Gamma \cong \langle b^2, c^2, d^2 \mid (bcbd)^n\rangle$, $n \geq 1$, and G has an embedding in which only b preserves spin.

Table 2: Classification of the infinite cubic planar Cayley graphs with 3 generators and connectivity 2 (types (iv) to (viii) of Theorem 1.1).

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.

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, for simplicity we will draw only one, undirected, edge in this case.

We call a graph *cubic*, if each of its vertices is adjacent with precisely three other vertices.

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 $\langle a, b, \ldots \mid R_1, R_2, \ldots \rangle$ 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 *double ray* is a two-way infinite path in a graph.

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.

2.3 Embeddings into 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 [6]. However, in the cubic case they do (see Corollary 5.1), and a significant part of this paper is concerned with constructing these embeddings.

An embedding is *Vertex-Accumulation-Point-free*, or *VAP-free* for short, if the images of the vertices have no accumulation point in $\mathbb{R}^2$. The following fact will allow us to easily deduce that certain Cayley graphs are planar looking only at the corresponding presentations.

Corollary 2.2 ([8, Section 4]). *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.

3 Cayley graphs with 2 generators

In the rest of the paper G will always denote a Cayley graph and $\Gamma(G)$, or just Γ, will denote the corresponding group. In this section we consider cubic Cayley graphs on two generators a, b where b is an involution.

The following lemma will play an important role.

Lemma 3.1. *Let $G = Cay \langle a, b \mid b^2, \dots \rangle$ be a Cayley graph with $\kappa(G) = 2$, and let P be a shortest path whose endvertices x, y separate G. Then $x^{-1}y$ is an involution, and it equals b or a^i for some $i \in \mathbb{Z}$. Moreover, x sends three independent paths to y.*

Proof. We begin by proving that

$$\text{The component } C \text{ of } G - \{x, y\} \text{ meeting } P \text{ sends four edges to } \{x, y\}. \tag{1}$$

Indeed, suppose that x sends only one edge xx' to C. Then P must contain the edge xx' and, easily, $\{x', y\}$ also separates G. This contradicts the choice of P, since the path $P - x$ is shorter and its endvertices also separate G. Thus x must send two edges to C, and by the same argument y must send another two edges to C, which proves (1).

Our next aim is to prove that

$$\text{There are 3 independent } x\text{–}y \text{ paths in } G. \tag{2}$$

To see this, note that no vertex of C separates $N(x)$ from $N(y)$ in C. Indeed, if v is such a vertex, then $\{x, v\}$ also separates G. But v must lie on P since P connects $N(x)$ to $N(y)$ in C. This contradicts the choice of P, since xPv is a shorter candidate. By Menger's theorem [5], this implies that there are two independent x–y paths P_1, P_2 through C. Note that by (1) $\{x, y\}$ sends only two edges to $G - C$. If these edges are incident with two distinct components of $G - \{x, y\}$, then any of these edges is a bridge of G, contradicting the fact that G is 2-connected. Thus these edges are incident with the same component $C' \neq C$ of $G - \{x, y\}$, and there is an x–y path P_3 in $G - C$. Combined with the paths P_1, P_2 this proves (2).

By (1) each of x, y has a unique neighbour in $G - C$; denote these neighbours by x', y' respectively. Next, we claim that

$$P \text{ is monochromatic.} \tag{3}$$

Indeed, if P contains both colours a, b, then it contains an edge e that has the same colour as the edge xx'. Thus we can translate e to xx' by a colour-automorphism w of G. Let $P' := w(P)$, and distinguish the following two cases: either the two endvertices of P' are in distinct components of $G - \{x, y\}$, or not. In the first case, we obtain a contradiction to (2) since $\{x, y\}$ separates the endvertices of P'. In the second case, we distinguish two subcases: either P' contains the edge yy', or one of the endvertices of P' is x and the other lies in the component C' of $G - \{x, y\}$; no other alternative exists because P' contains xx' and so if x is not an endpoint then P' meets both C and C', and

the edges xx', yy' separate C from C'. Now in the first of these subcases, we obtain a subpath P'' of P' with endvertices x, y that meets C'. This contradicts the minimality of P if $P'' \neq P'$ or, proposition (1) if $P'' = P'$. In the second subcase, as xx' is the only x-C' edge incident with x and the endvertices x, v of P' separate G, it is easy to see that the vertices x', v also separate G. Since both these vertices lie on $P' = w(P)$, we obtain a contradiction to the choice of P since $x'P'v$ is a shorter candidate. Thus we obtained a contradiction in all of these cases, and so (3) must hold.

Now if the colour of P is b, then as b is an involution P consists of a single b edge and we obtain $z = x^{-1}y = b$.

Suppose that every edge of P is labelled a instead. If any of the edges xx', yy' is also labelled a then we can repeat the above argument to obtain a contradiction. Thus both xx', yy' are labelled b in this case. This means that the coset D of the subgroup spanned by a containing x, and thus also y since P is labelled a, is contained in $\overline{C} := C \cup \{x, y\}$. Now consider the element $z := x^{-1}y$ of $\Gamma(G)$; note that yz lies on $D \subseteq C$. Recall that there is an x–y path P_3 in $G - C$ (see (2)), and note that any such P_3 must begin with the edge xx' and finish with $y'y$. Choose P_3 so as to minimize its length. Now consider its translate P_3z starting at y instead of x. By the choice of z this path begins with the edge yy' and terminates in $\overline{C}$. Since xx', yy' are the only edges connecting C' to the rest of the graph, and since y' lies in C', the path P_3z must also contain xx'. Then P_3z has to terminate at x, because otherwise the minimality of $|P_3|$ is contradicted by the subpath of P_3z from y to x. This means that $z = x^{-1}y$ is an involution, as it exchanges x with y. By the choice of z we have $z = a^{\pm||P||}$ in this case as desired.

$\square$

Theorem 3.2. *Let $G = Cay \langle a, b \mid b^2, \ldots \rangle$ be a Cayley graph with $\kappa(G) = 2$ in which a has infinite order. Then either $\Gamma(G) \cong \langle a, b \mid b^2, (ab)^n \rangle$ with $n \geq 2$, or $\Gamma(G) \cong \langle a, b \mid b^2, (aba^{-1}b^{-1})^n \rangle$ with $n \geq 1$. Thus G is planar, and has a consistent embedding in which a preserves spin.*

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

Proof. Let P be a shortest path whose endvertices x, y separate G. By Lemma 3.1 $x^{-1}z$ is an involution, and it equals b or a^i for some $i \in \mathbb{Z}$. But since a has infinite order we have $a^{2i} \neq 1$, and so only the former can be the case. This proves that

Every b edge is a hinge. $\hspace{4cm}$ (4)

We now have enough information about the structure of G to enable as to work out a presentation of $\Gamma(G)$. We start by proving that

There is a simple presentation $\langle a, b \mid R_1, R_2, \ldots \rangle$ of $\Gamma(G)$ such that $a^2 \not\preceq R_i$ holds for every i, $\hspace{3cm}$ (5)

where we write $w \preceq R$ if any relator obtained R from R by rotation and inversion contains the word w, and write $w \not\preceq R$ otherwise.

To begin with, let $\langle a, b \mid R_1, R_2, \ldots \rangle$ be any simple presentation of $\Gamma(G)$. Suppose that R_i contains the subword a^2. Consider a closed walk W in G induced by R_i. Thus W traverses two consecutive edges zx, xv labelled a. Now let $e = xy$ be the b edge incident with x. By (4) e is a hinge. If the edges zx, xv are incident with the same component of $G - \{x, y\}$, then it is easy to see that one of the edges incident with y is a bridge, contradicting the 2-connectedness of G (Figure 1). Thus zx, xv are incident with distinct components of $G - \{x, y\}$, which implies that W, being a closed walk, must visit y.
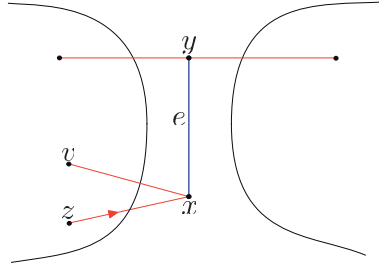

Figure 1: The situation in the proof of (5).

Now as W contains both x, y, we can use the edge $e = xy$ to 'shortcut' W into two walks W_1, W_2, both starting and ending at y, such that the concatenation of the corresponding relations R^1, R^2 yields a word equivalent to R_i (Figure 2). Note that R^1, R^2 are simple since R_i was. Thus we could replace R_i by two relations R^1, R^2 to obtain a simple presentation of $\Gamma(G)$ in which the word a^2 appears less often. We can repeat this procedure as often as needed to replace R_i by a finite family of relations $R^1, \ldots R^k$ in which the words a^2 and a^{-2} do not appear at all. Doing so for each i we obtain a presentation that establishes (5). Note that we did not have to assume that $\Gamma(G)$ is finitely presented; this will follow soon.
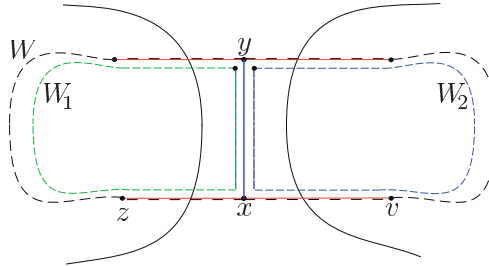

Figure 2: The situation in the proof of (5).

So let $\langle a, b \mid R_1, R_2, \ldots \rangle$ be a presentation of $\Gamma(G)$ as supplied by (5). Next, we claim that

There is no pair $i, j \in \mathbb{N}$ such that $aba \preceq R_i$ and $aba^{-1} \preceq R_j$. \hfill (6)

Indeed, suppose that $ab \preceq R_i$, and consider a cycle K in G induced by to R_i such that both edges in a subpath Q of K induced by ab are incident with the identity $\mathbb{1} \in \Gamma(G)$. Applying (4) again to the b edge e incident with $\mathbb{1}$,

we obtain a separation of G as in Figure 2, with x playing the role of $\mathbb{1}$. Now since K contains the subpath Q containing e, and K is not allowed to visit any vertex twice, the edge f following Q in K is determined: it must be the other a edge incident with the component meeting Q. This means that once we see the subword ab in some relator R_i the following letter (either a or a^{-1}) is determined. This proves (6).

Combining (5) with (6) it follows that once we see the letter a or a^{-1} in some R_i the rest of the word is uniquely determined. In other words, there is, up to rotation of the letters and inversion, only one possible simple relation in $\Gamma(G)$ except for b^2: this relation is either $(ab)^n$ for some $n \geq 2$ or $(aba^{-1}b^{-1})^m$ for some $m \geq 1$.

Let us postpone the rest of the forward implication to prove the converse implication first. Given one of the above group presentations, it is easy to explicitly construct a plane Cayley graph having this presentation: start with an infinite 3-regular tree T, and replace each vertex of T by a cycle of length $2n$ or $4m$, with its edges coloured alternately a and b. Then glue, for every edge uv of T, the two cycles replacing u and v along a b edge. We leave the details to the reader. Note that a preserves spin in any such embedding. Moreover, if $\Gamma(G) \cong \langle a, b \mid b^2, (ab)^n \rangle$ then every b edge preserves spin, while if $\Gamma(G) \cong \langle a, b \mid b^2, (aba^{-1}b^{-1})^m \rangle$ then every b edge reverses spin. It is also clear that each of the above presentations corresponds to a group in which a has infinite order.

Combining what we have proved so far we can now finish off the forward implication: since G must have one of the above presentations, it also has a consistent embedding as constructed above.

$\square$

Next, we consider the case when the order of a is finite.

Theorem 3.3. *Let $G = Cay \langle a, b \mid b^2, a^n, \ldots \rangle$ be a planar Cayley graph with $\kappa(G) = 2$. Then $G \cong Cay \langle a, b \mid b^2, a^4, (a^2b)^n \rangle, n \geq 2$. Thus G has a consistent embedding in which a reverses spin and b preserves spin.*

Conversely, for every $n \geq 2$ the above presentation yields a planar cubic Cayley graph of connectivity 2 in which a has order 4.

Proof. Since G is planar it has some embedding σ. Let C be an a labelled cycle induced by a^n. We claim that either every edge of C preserves spin in σ or every edge of C reverses spin in σ. For suppose not. Then we can find consecutive edges xy, yv on C such that the b edges xz, yw incident with x and y respectively lie in the same side of C and the b edge vu incident with v lies in the other side (Figure 3).

We claim that $\{x, y\}$ do not separate z from w. For if $\{x, y\}$ separates G, then it can only leave two components behind as G cannot contain a bridge, and each of these components sends two edges to $\{x, y\}$. Note however that both a-labelled edges incident with $\{x, y\}$ are incident with a common component containing $C - \{x, y\}$, and so the other two edges xz, yw must go to a common component as desired.

Using this we can now prove the stronger assertion that even C does not separate z from w. Indeed, let R be a z–w path that does not meet $\{x, y\}$, which exists by the above claim. If R meets C then let Q be the subpath of R from z to the first encounter with C. We can rotate C by one step to map y to
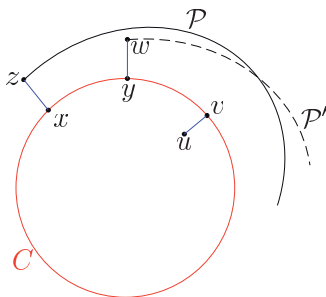
Figure 3: A contradictory situation in the proof of Theorem 3.3.

x —by multiplying G with an element of $\Gamma(G)$— to obtain a translate Q' of Q which, by an easy topological argument, intersects Q. Now combining subpaths of Q and Q' up to such an intersection we can obtain a w–z path that does not meet C.

So let S be a z–w path that does not meet C. Again, we can rotate C by one step to map y to v. Now this rotation translates S to a w–u path which does not meet C. This contradicts our assumption about the embedding σ, and proves our claim that either every edge of C preserves spin or every edge reverses spin. Translating C to other a cycles of G and using a similar argument we can prove that

> either every a edge preserves spin or every a edge reverses spin in σ. (7)

Let P be a shortest path whose endvertices x, y separate G. Lemma 3.1 implies that either $y = xb$ or $y = xa^i$ and $a^{2i} = 1$. If the former case it is easy to obtain a contradiction: as the edge xy cannot be a bridge, there is a path connecting the two a cycles that contain x and y; but then $\{x, y\}$ cannot separate G since all their neighbours lie in one component. Thus x, y are opposite vertices of some a-labelled cycle C.

Since a is not an involution, for otherwise we are in one of the degenerate cases of type (ix) of Theorem 1.1, there is another pair $\{x', y'\} \neq \{x, y\}$ of opposite vertices on C which also form a separator. Now as x' sends three independent paths to y' by Lemma 3.1, $\{x, y\}$ does not separate x' from y'. Thus, all four a edges incident with x or y are incident with the same component K of $G - \{x, y\}$, which component contains $C - \{x, y\}$. As x sends three independent paths to y, and only two of them can meet K, there is an x–y path L that does not meet C except at its endvertices. By the same argument, there is also an x'–y' path L' that does not meet C except at its endvertices. Now if L' intersects L we easily get a contradiction to the fact L is disconnected from K by $\{x, y\}$ (Figure 4). Thus L' cannot intersect L, and an easy topological argument implies that L' and L lie in different sides of C.

Combined with (7) this implies that every a edge reverses spin in σ, since if a edges preserve spin then all b edges incident with C must lie in one side of C. Moreover, it implies that

$$|C| = 4. \tag{8}$$

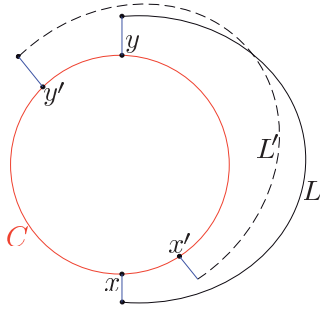Indeed, otherwise there are at least three pairs of opposite vertices on C, and

Figure 4: A contradictory situation in the proof of Theorem 3.3.

so at least two of the corresponding paths L would have to share the same side of C.

Let M be the set of a-labelled cycles visited by P; note that $C \in M$. Easily, for every cycle D in M, only one of the sides of D is met by P: otherwise we could shortcut P using a subarc of D, but P is chosen to have minimum length. Now note that for every vertex v of G there is a $D \in M$ such that v lies either on D or in the side of D not met by P because, by (8) and the fact that every a edge reverses spin, every edge incident with a vertex in $P \cup \bigcup M$ lies in one of those sides. Applying the same argument to the translate P' of P joining xa to ya, and repeating for every a-labelled cycle, it is easy to see that G has the structure of Figure 5, and the presentation $G = Cay \left\langle a, b \mid b^2, a^4, (a^2 b)^m \right\rangle$, where $m := |M| \geq 2$.
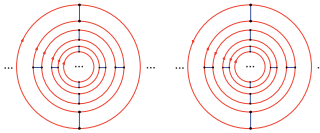

Figure 5: The graph of Theorem 3.3 for $m = 4$.

The embedding claimed in the assertion can be seen in Figure 5. □

4   Cayley graphs with 3 generators

In this section we consider the case when G is defined by three generators b, c, d, all of which are of course involutions since G is cubic. We will distinguish two cases according to whether G has a hinge, i.e. an edge $e = xy$ such that the removal of the pair of vertices x, y disconnects G.

4.1   Graphs with hinges

Theorem 4.1. *Let $G = Cay \left\langle b, c, d \mid b^2, c^2, d^2, \ldots \right\rangle$ be a Cayley graph with $\kappa(G) = 2$ having a hinge. Then either $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (bd)^m \right\rangle$ for $m, n \geq 2$ or $G \cong Cay \left\langle b, c, d \mid b^2, c^2, d^2, (bcbd)^n \right\rangle$ for $n \geq 1$, but not both. In both cases G is planar. In the first case it has an embedding in which every edge*

reverses spin, while in the second it has an embedding in which precisely one of the colours (that of the hinges) preserves spin.

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

Proof. Suppose without loss of generality that every b edge is a hinge, and consider a b edge $e = xy$. We distinguish two cases: either the c edges incident with x, y are incident with the same component of $G - \{x, y\}$ or not; see Figure 6.
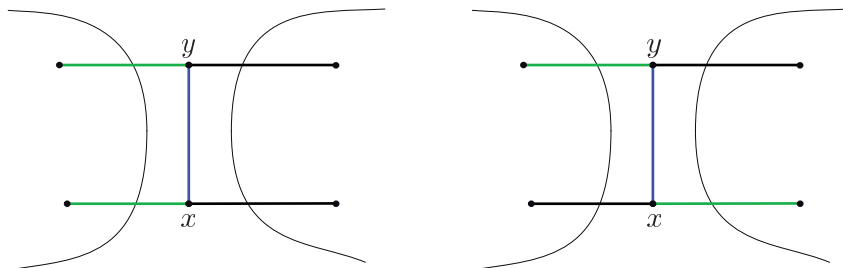

Figure 6: The two cases with respect to the distribution of the edges incident with xy into components.

Case I: In the former case, we can find a simple presentation of Γ in which no relator contains cd or cbd as a subword. Indeed, let $\langle b, c, d \mid R_1, R_2, \ldots \rangle$ be any simple presentation of Γ. Suppose that $cbd \preceq R_i$ for some i, let W_i be a closed walk in G induced by R_i, and let e be a b edge in the middle of a cbd subpath of W_i. Then, as one can easily check using Figure 6, W_i has to visit some endvertex of e twice, contradicting the fact that the presentation is simple.

Now suppose that $cd \preceq R_i$ for some i, and let again W_i be a closed walk in G induced by R_i. Let x be the middle vertex of a cd subpath of W_i, and consider the b edge e incident with x. Since e is a hinge, we have the situation of the left part of Figure 6 again, and we can use e to 'shortcut' W_i into two walks W^1, W^2 such that the concatenation of the corresponding relations R^1, R^2 yields a word equivalent to R, similarly to what we did in Figure 2. Thus we can replace R_i by two relations R^1, R^2 to obtain a presentation of Γ in which the word cd appears less often. Repeating this procedure as often as needed, we obtain, as claimed, a simple presentation $\langle b, c, d \mid R'_1, R'_2, \ldots \rangle$ of Γ in which $cd, cbd \npreceq R'_i$ for every i.

The fact that no R'_i contains the words cd or cbd easily implies that R'_i contains at most two of the letters c, b, d. Thus each $R'_i \neq b^2, c^2, d^2$ is of the form $(bc)^n$ or $(bd)^n$. Since the presentation is simple, it cannot be the case that $(bc)^n, (bc)^l$ with $n \neq l$ are both relators; similarly for $(bd)^n$. Thus Γ has a presentation of the form $\langle b, c, d \mid b^2, c^2, d^2, (bc)^n, (bd)^m \rangle$ where n, m might take the value 0, meaning that the corresponding relator is not present.

Let us now check that $m, n \neq 0$. Indeed, by Lemma 2.1 the set B of cycles in G induced by the relators in the above presentation generate $\mathcal{C}_f(G)$. But if $n = 0$ then no element of B contains a c edge, which implies that no cycle of G contains a c edge. It is an easy graph-theoretical fact that every c edge must be a bridge in this case, which contradicts our assumption that G is 2-connected. Similarly, $m \neq 0$ for otherwise every d edge is a bridge.

We can now apply Corollary 2.2 to obtain an embedding σ. Indeed, the set B of cycles in G induced by the relators in the above presentation form a 2-basis, in which every b edge appears twice and every other edge appears once. It is straightforward to check that every edge reverses spin in σ using the fact that the elements of B are precisely the finite face-boundaries.

Case II: In the second case (Figure 6 right), using very similar arguments as above we obtain a simple presentation $\langle b, c, d \mid R'_1, R'_2, \ldots \rangle$ of Γ in which no relator contains cbc or cd or dbd as a subword. This easily implies that the only possible kind of relator R'_i, except for b^2, c^2, d^2, is $(bcbd)^n$, and again since the presentation is simple only one value for n is allowed. Such a relator with $n > 0$ must exist for otherwise G is not 2-connected.

Similarly to Case I, the set B of cycles in G induced by this relator is a 2-basis of G and applying Corollary 2.2 we obtain a consistent embedding of G in which b edges preserve spin and all other edges reverse spin.

The converse implication can, in both cases, be proved easily by a construction similar to that of Theorem 3.2. $\qquad\square$

4.2 Graphs with no hinges

We continue our analysis with the case when G has no hinge. The following is similar to Lemma 3.1.

Lemma 4.2. *Let $G = Cay \langle b, c, d \mid b^2, c^2, d^2, \ldots \rangle$ be a Cayley graph with $\kappa(G) = 2$ and no hinge, and let P be a shortest path whose endvertices x, y separate G. Then P is two-coloured and $x^{-1}y$ is an involution.*

Proof. By the same arguments as in the proof of assertion (1) of the proof of Lemma 3.1 we can prove that

> The component C of $G - \{x, y\}$ containing P sends four edges to $\{x, y\}$. (9)

Let x', y' be the neighbours of x, y respectively in the other component C' of $G - \{x, y\}$. We claim that none of the edges xx', yy' has the same colour as some edge e of P. For otherwise we can, similarly to the proof of (3) in Lemma 3.1, translate e to that edge to obtain a translate of P that crosses the separator $\{x, y\}$, which easily yields a contradiction to the choice of P. This, combined with the fact that $||P|| > 1$ since G has no hinge, implies that

> P is two-coloured, with the colours b, c say, and both edges xx', yy' bear the third colour d. (10)

Using this it is now easy to prove that $z = x^{-1}y$ is an involution. For if not then we obtain the situation of Figure 7, where P_1 is a shortest x–y path in $G - C$ and P_2 its translate by z. By a similar argument to that of the final part of the proof of Lemma 3.1 we can now prove that vv' is also incident with C', contradicting (9).

$\qquad\square$

4.2.1 Graphs with no hinges and no 2-coloured cycles

We will distinguish two cases according to whether G has a cycle containing only two of the three colours b, c, d or not.
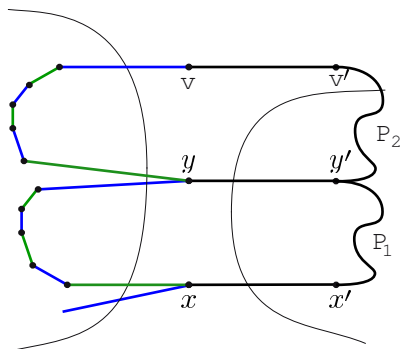
Figure 7: The hypothetical situation when $x^{-1}y$ is not an involution.

Lemma 4.3. *Let $G = Cay \langle b, c, d \mid b^2, c^2, d^2, \ldots \rangle$ be a Cayley graph with $\kappa(G) = 2$ and no 2-coloured cycle and no hinge. Then G is planar and has an embedding in which all edges preserve spin.*

Proof. Let P be a shortest path whose endvertices x, y separate G. Lemma 4.2 yields that P is two-coloured and $z := x^{-1}y$ is an involution. Since z is an involution and G has no 2-coloured cycle, $||P||$ must be odd. Assume without loss of generality that the labels appearing in P are b and c and that $z = b(cb)^n$.

We are now going to construct an explicit embedding σ of G in the sphere. We will construct σ inductively, in ω steps. For this, let $D_0, D_1, D_2, \ldots$ be a —possibly finite— enumeration of the double-rays of G spanned by b and c. Let G_0 be the graph obtained from G after contracting all edges of G not incident to or contained in D_0. We claim that G_0 consists of the double-ray D_0 and an infinite set $\mathcal{P}$ of pairwise disjoint paths of length two with all edges labelled d, each such path joining a vertex x of D_0 to the vertex xw (Figure 8). Indeed, to begin with, D_0 is a subgraph of G_0. Moreover, for every vertex x of D_0 we know that x, xz separate G, and hence also G_0. Furthermore, by (10) one of the components of $G - \{x, y\}$ contains D_0 and the other component C' is incident with the two d edges incident with x and xz. Since C' does not send any other edge to D_0, it is by the definition of G_0 contracted into a single vertex. This proves our claim. Figure 8 shows an embedding σ_0 of G_0. For reasons that will become clear soon, consider σ_0 to be an embedding in the sphere S rather than in the plane. Note that all vertices of D_0 have the same spin in σ_0.

This was the first step of our inductive construction of σ. We now proceed with the remaining steps, in each step i decontracting the double-ray D_i and its incident edges and extending σ_{i-1} into an embedding σ_i that includes D_i.

More formally, for $i = 1, 2, \ldots$, let G_i be the graph obtained from G by contracting all edges that are not incident to or contained in one of the double-rays $D_0, D_1, \ldots D_i$. Consider the embedding σ_{i-1} of G_{i-1} in the sphere S inherited from the previous step. Let x_i be the (unique) vertex of G_{i-1} into which D_i was contracted. Note that x_i was the middle vertex of some of the d-labelled paths of length two of G_{i-1}. Pick a closed disc X in S into which σ_{i-1} maps x_i and part of its incident d-edges but no other vertex or edge. Then, consider the auxiliary graph G_i', isomorphic to G_0, which is obtained from G just like G_0 except that we keep D_i instead of D_0 and contract all other double-rays. Moreover, let σ_i' be an embedding of G_i' in another copy S' of the sphere similar
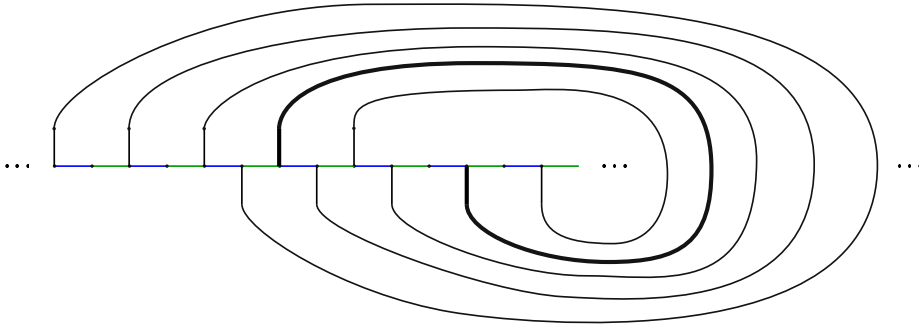
14

Figure 8: Constructing a consistent embedding in the absence of 2-coloured cycles.

to the embedding σ_0 of G_0. Pick a closed disc X' in S' into which σ_{i-1} maps the (unique) vertex of G'_i into which D_0 was contracted, and into which σ_{i-1} maps no other vertex. To obtain the new embedding σ_i, cut both discs X, X' out of their corresponding spheres S, S', and glue the remainders together along their boundaries to obtain a new copy of the sphere S'' in which now G_i is embedded (as a combination of σ_{i-1} and σ'_i). Indeed, note that the boundary of each of the discs X, X' is crossed by precisely two edges, and these are the same d edges e, f of G . Thus we may perform the glueing in such a way that these edges are properly embedded as arcs in S'', half of each edge coming from S and the other half coming from S'.

We would like to obtain an embedding σ of G as a limit of the embeddings σ_i. To achieve this, it is more convenient to think of the embedding σ_i we just constructed as an embedding in S itself rather than some new copy S''. Thus, formally, we define $\sigma_i : G \to S$ so that it coincides with σ_{i-1} in $S \backslash X$ and 'imitates' $\sigma'_i \upharpoonright (S' \backslash X')$ in X. We can now define the embedding σ of G on S as the limit of the σ_i: every point p of G is mapped to the unique point s of S such that $\sigma_i(p) = s$ holds for all but finitely many i.

Next, we claim that σ can be constructed so that every vertex of G has the same spin. For this, suppose that every vertex of G_{i-1} has the same spin ς in σ_{i-1}. Then every vertex of G_{i-1} also has that spin in σ_i by construction. Moreover, all vertices of D_i share the same spin ς' in σ_i, as this was the case in σ'_i. We may assume that $\varsigma = \varsigma'$, since if this is not the case, then we could have flipped the disc $S' \backslash X'$ we replaced X with around before performing the above cutting and glueing operation, to reverse the spins of the vertices of D_i (when flipping that disc we fix the two points on its boundary that meet edges of G). Since there are only two possible types of spin in a cubic graph, we would have then achieved $\varsigma = \varsigma'$. Thus, assuming that every vertex of G_{i-1} has the same spin in σ_{i-1} we deduced that every vertex of G_i has the same spin in σ_i. Since this is the case at the beginning of our construction when $i = 0$, we obtain by induction that every vertex of G has the same spin is σ. $\qquad\square$

We can now plug the embedding σ we just constructed into the following result from [9] to obtain a presentation of the corresponding graphs.

Lemma 4.4 ([9, Theorem 11.10]). *Let $G = Cay \langle b, c, d \mid b^2, c^2, d^2, \ldots \rangle$ be a Cayley graph with $\kappa(G) = 2$, with no 2-coloured cycle, having an embedding in which*

all edges preserve spin. Then $G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (b(cb)^n d)^m \rangle, n \geq 1, m \geq 2.$

Conversely, each of the above presentations, with parameters chosen in the specified domains, yields a planar cubic Cayley graph of connectivity 2 with no two-coloured cycle.

Corollary 4.5. *Let $G = Cay \langle b, c, d \mid b^2, c^2, d^2, \dots \rangle$ be a Cayley graph with $\kappa(G) = 2$ and no 2-coloured cycle and no hinge. Then*
$G \cong Cay \langle b, c, d \mid b^2, c^2, d^2, (b(cb)^n d)^m \rangle, n, m \geq 2.$

Conversely, each of the above presentations, with parameters chosen in the specified domains, yields a planar cubic Cayley graph of connectivity 2 with no hinge and no two-coloured cycle.

Proof. Combining Lemma 4.3 with Theorem 4.4 we obtain a presentation of the desired form, except that Theorem 4.4 allows $n = 1$ as a possibility. However, if $n = 1$ then any b edge would be a hinge as can be easily seen with the help of Figure 8: if $n = 1$, in which case $|z| = 3$, then no path connects the two parts of a bc-double-ray obtained after removing one of the b edges and its endvertices. This contradicts our assumption that G has no hinge. On the other hand, if $n > 1$ then it is easy to check that no edge is a hinge. $\square$

4.2.2 Graphs having 2-coloured cycles and no hinges

We can now proceed with the last result of this section, completing our characterisation of the planar cubic Cayley graphs of connectivity 2. The only remaining case is when G has no hinge but does have a two-coloured cycle.

To begin with, let us first take the finite ones out of the way: by [1, Chapter 27, Theorem 3.7.] finite simple Cayley graphs of degree at least 3 are 3-connected, so in our case G has to have parallel edges in order to have connectivity 2. It follows that if G is finite then

$$\Gamma \cong \langle b, c, d \mid b^2, c^2, d^2, (bc)^n, cd \rangle, n \geq 1.$$

From now on we can assume G to be infinite.

Theorem 4.6. *Let $G = Cay \langle b, c, d \mid b^2, c^2, d^2, \dots \rangle$ be a planar Cayley graph with $\kappa(G) = 2$, with a two-coloured cycle, and no hinge. Then either $\Gamma \cong \langle b, c, d \mid b^2, c^2, d^2, (bc)^2, (bcd)^m \rangle$ with $m \geq 2$, in which case G has a consistent embedding with a single spin, or $\Gamma \cong \langle b, c, d \mid b^2, c^2, d^2, (bc)^{2k}, (cbcd)^m \rangle, k, m \geq 2$ and G has a consistent embedding in which only c preserves spin.*

Conversely, each presentation as above, with parameters in the specified domains, yields a planar Cayley graph of connectivity 2 without hinges.

Proof. Let again P be a shortest path whose endvertices x, y separate G. Lemma 4.2 yields that P is two-coloured and $z := x^{-1}y$ is an involution. We claim that

   the two colours, b, c say, appearing in z span finite cycles.   (11)

Indeed, if b, c span double rays, then it is (at least) one of the other two pairs of colours that span two-coloured cycles. In this case, assume without loss of generality that b, d span cycles. We claim that x, y lie on distinct bd cycles.

Suppose to the contrary that there is a b, d cycle C containing both x, y. Then x, y separate C into two subarcs P_1, P_2. Note that any colour-automorphism

16

of G mapping x to some other vertex of C must also map y within C. It is now easy to see that, as none of P_1, P_2 can be a single edge since G has no hinge, there is a colour-automorphism g of G mapping x to some vertex x^* of P_1 and y to some vertex y^* of P_2. This, however, easily leads to a contradiction: (10) implies that $\{x, y\}$ separates P_1 from P_2 in G, and so it also separates x^* from y^*. But there is a x^*–y^* path alternating in the colours b and c, namely the translate of P under g, which contradicts the fact that $\{x, y\}$ only sends d-coloured edges to one of the components of $G - \{x, y\}$ by (10). This contradiction proves our claim that x, y lie on distinct bd cycles.

So let C be the bd cycle that contains x. Since $y \notin V(C)$, all vertices of $C - x$ lie in the same component of $G - \{x, y\}$. Again, this contradicts (10), since x sends both a b-coloured and a d-coloured edge to that component. Thus we have proved (11).

Assume from now on that the length of a bc cycle is $2k$. We now distinguish three subcases depending on the length of P.

Case 1: $||P|| = 2$. In this case we have $2k = 4$ since z is an involution. We can now repeat the construction of Lemma 4.3, with the only difference that instead of bc double rays we now have bc cycles of length 4. Instead of Figure 8 we now have the left part of Figure 9. It is easy to check that the embedding σ we obtain must look like the one in the right part of Figure 9. Using this embedding it is straightforward to check that Γ can be represented by $\Gamma \cong \langle b, c, d \mid b^2, c^2, d^2, bcbc, (bcd)^m \rangle$ with $m \geq 2$ (for $m = 1$ we obtain a degenerate case: a group with 4 elements whose Cayley graph is 3-connected).
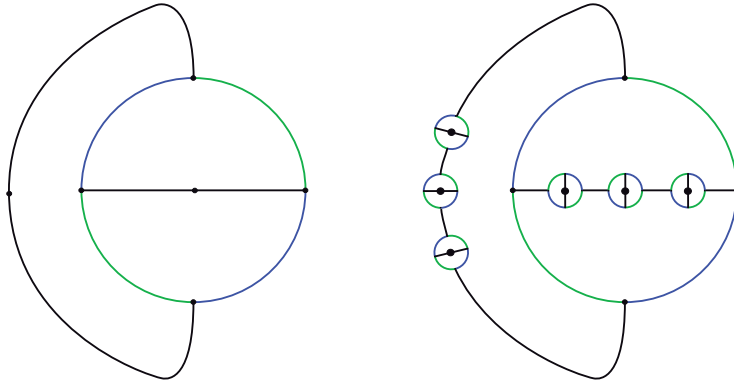

Figure 9: The first step σ_0 in the construction of the embedding for the graphs of Case 1 (left) and a possible later step (right). In this example we chose $m = 4$, i.e. our graph is $Cay \langle b, c, d \mid b^2, c^2, d^2, bcbc, (bcd)^4 \rangle$.

Case 2: $||P|| = 3$. Easily, $2k \geq 6$ in this case. We are again going to repeat the construction of Lemma 4.3, with the bc cycles playing the role of the bc double-rays there, to construct a consistent embedding σ of G. However, this time only one of the colours, c say, will preserve spin, while all b and d edges will reverse spin. The first step σ_0 of this construction is shown in Figure 10 (left). It is not hard to prove that if k is odd then G cannot be planar. For example, if $k = 3$ then G_0 becomes isomorphic to the Kuratowski graph $K_{3,3}$ after suppressing the vertices of degree two.
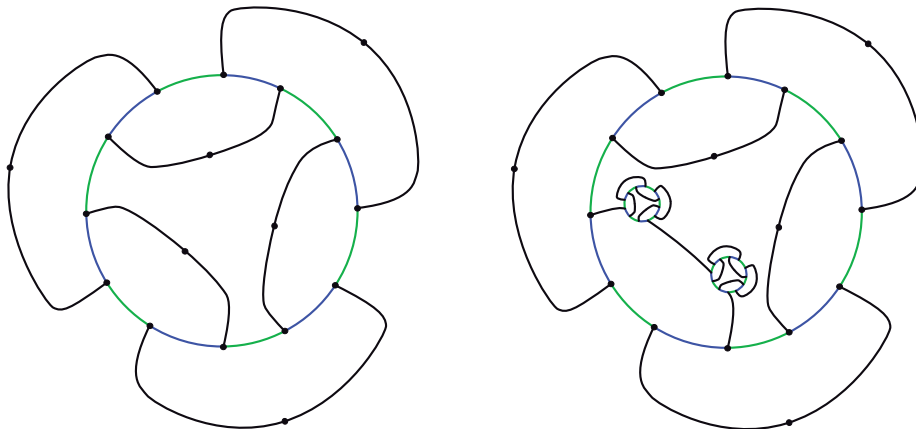
Figure 10: The first step σ_0 in the construction of the embedding for the graphs of Case 2 (left) and a possible later step (right).

Note that, as desired, the c edges preserve spin in this embedding, while the b edges do not. Our construction makes sure that the b and c edges retain this property also in the final embedding σ. We need to make sure that all d edges reverse spin in σ.

Recall that in each step i of our construction we decontract the cycle D_i, which was hidden in some vertex x_i of G_{i-1} that was the middle vertex of a path $P_i = y_i e x_i f z_i$ of length two, where e and f are labelled d. We will make sure by induction that the endpoints of every maximal d-labelled path P, no matter whether P is a path of length two with a contracted vertex in the middle or an original edge of G, have different spins in σ_i. Note that this is the case in the embedding σ_0. Now assuming that it is also true for a later embedding σ_{i-1}, we can easily make sure that it remains true in σ_i: for if it happens to be false, then this can only be so locally at the place where we introduced D_i, and we can flip the newly pasted disc X' around —as we did in the proof of Lemma 4.5— to reverse spins to the desired state. Indeed, there are only two vertices v, w (of D_i) in X' adjacent to some vertex outside X', and these two vertices have different spins because the embedding within X' is by construction similar to the one of left part of Figure 10; now as the neighbours of v, w outside X' have different spins by our induction hypothesis, one of the two possible ways to paste X' into σ_{i-1} (i.e. with or without a flip) preserves our induction hypothesis that endpoints of a maximal d-labelled path have different spins. Thus, in the limit embedding σ, all d edges reverse spin.

We will now use Lemma 2.1 to obtain a presentation of Γ. To begin with, note that by the construction of G, each of the faces containing a cbc subpath in their boundary is induced by a relation of the form $(cbcd)^m$, where m is fixed for all such faces. Let B' be the set of circuits of such face boundaries, in other words, the set of cycles induced by $(cbcd)^m$. Moreover, let B'' be the set of circuits of the bc cycles of G. We claim that $B := B' \cup B''$ generates $\mathcal{C}_f(G)$.

To see this, let D be any cycle of G. Note that D can be written as a sum (in $\mathcal{C}_f(G)$) $D = \sum_i D_i$ of cycles D_i such that no D_i meets both sides of some bc cycle: indeed, if D meets both sides of the bc cycle C, then we can split D into two subpaths P_1, P_2 with endvertices on C, and combine each P_i with a subarc

18

C' of C with the same endvertices to obtain two cycles $D_1 := P_1 \cup C'$, $D_2 := P_2 \cup C'$ whose sum is D, so that P_1, P_2 together cross C less often than D does. If some of the D_1 still crosses C, we can repeat this operation on D_1 and so on. As D can meet at most finitely many bc cycles, performing this kind of operation a finite number of times we obtain the desired presentation for D. It is now not hard to see that each D_i can be written as a sum of elements of B; see Figure 10. Thus $D = \sum_i D_i$ too can be written as a sum of elements of B. This proves our claim that B generates $\mathcal{C}_f(G)$. Applying Lemma 2.1 we now immediately obtain a presentation: $\Gamma \cong \langle b, c, d \mid b^2, c^2, d^2, (bc)^{2k}, (cbcd)^m \rangle$ with $k \geq 2$ and $m \geq 2$. (Again, for $m = 1$ we have only one bc cycle and G is 3-connected.)

Note that in the embedding we constructed, every vertex is incident with two finite and one infinite face boundary. In [9] we use these Cayley graphs as a building block in order to obtain 3-connected plane Cayley graphs which have no finite face boundaries, contrary to a conjecture of Bonnington and Watkins [3] that none exist. We now state some properties of the graphs we just constructed to be used in [9].

Proposition 4.7. *Let $G = Cay \langle b, c, d \mid b^2, c^2, d^2, (bc)^{2k}, (cbcd)^m \rangle$. Then the following assertions are true.*

(i) *G has a consistent embedding σ in which c preserves spin while b, d reverse spin. In this embedding, each vertex is incident with two faces bounded by a cycle induced by the relator $(cbcd)^m$ and one face that has infinite boundary.*

(ii) *G 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;*

(iii) *If a pair of vertices s, t of a cycle C of G induced by $(cbcd)^m$ separates G, then both s, t are incident with a d-edge of C;*

(iv) *G has no hinge;*

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

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

Proof. The consistent embedding σ required by (i) was constructed above.

To prove that (ii) is true, it suffices to show that if uv is an edge coloured b or c then there are two independent u–v paths none of which is uv itself. This is indeed true, and is easy to check using Figure 10.

To see that (iii) is true, suppose that s is not incident with a d-edge of C, in which case it must be incident with a b-edge of C. Consider the two neighbours $u = sc$ anf $v = sb$ of s on C, and note that if $\{s, t\}$ separates G, then u, v lie in distinct component because of the cycle C. However, it is easy to find a u–v path in G that does not meet C except at u and v (use the bc cycle containing uv). This easily yields a contradiction.

It is easy to check (iv) and (v) using Figure 10.

Finally, note that if two cycles C, D of G induced by the word $(cbcd)^m$ share an edge uv, then the colour of uv is c, and the bc cycle containing uv has a subpath joining C to D without using any of u, v as claimed by (vi). $\qquad\square$

Case 3: $||P|| > 3$. Easily, $2k \geq 8$ in this case. It is now easy to check that G cannot be planar: using one of the b, c cycles C and three translates of P with endpoints on C one can force $K_{3,3}$, one of the Kuratowski graphs, as a topological minor of G.

To prove the converse implication, that every presentation as in Theorem 4.6 yields a planar Cayley graph without hinges, we can use the constructions of Figures 9 and 10 to explicitly construct planar Cayley graphs with embeddings as in the assertion. The fact that these Cayley graphs do indeed have the desired presentation follows from the forward implication, which we have already proved. $\qquad\square$

We point out the following observation, which is a consequence of our analysis, to be used in [9].

Observation 4.8. *No graph of the types (iv) or (v) of Theorem 1.1 has an embedding in which some 2-coloured cycle bounds a face.*

5 Summary and final remarks

Combining our analysis of Sections 3 and 4 immediately yields our main results:

Proof of Theorem 1.1. Any cubic Cayley graph G has either two or three generators. In the former case Theorems 3.2 and 3.3 imply that G is of type (i), (ii) or (iii). In the latter case, Theorems 4.1 4.5 and 4.5 imply that G is of one of the types (iv) to (ix).

We need to check that G cannot belong to more than one of these types. For most of the pairs of types this is obvious; we now consider the remaining pairs. The fact that G cannot be of both type (i) and (ii) follows from assertion 6 of the proof of Theorem 3.2. The graphs of the last two types have no 2-coloured cycles, thus they cannot also belong to some of the types (iv) to (vi). Theorem 4.6 shows that no graph of both types (iv) and (v) can exist.

The converse implication is again obtained by combining the corresponding results in Sections 3 and 4. $\qquad\square$

Similarly, Corollary 1.2 follows easily from propositions 3.2, 4.1, 4.2 and 4.3.

Our analysis of Sections 3 and 4 also implies

Corollary 5.1. *Every cubic planar Cayley graph of connectivity 2 admits a consistent embedding.*

An important property of planar a Cayley graph is whether it admits a VAP-free embedding (as defined in Section 2.3), since this can have important implications for its group-theoretical [8] as well as its graph-theoretical properties [4, 12]. For the graphs studied in this paper we can always decide whether they have this property:

Corollary 5.2. *A cubic planar Cayley graph of connectivity 2 admits a VAP-free embedding if and only if it does not belong to one of the types (iii), (iv), (v), and (vii) of Theorem 1.1.*

The interested reader will be able to prove Corollary 5.2 using our analysis of these graphs.

It is proved in [8] that a Cayley graph admits a VAP-free embedding if and only if it is the 1-skeleton of a Cayley complex that can be embedded in the plane after removing some redundant simplices. Thus, by the above corollary, some of our graphs are not 1-skeletons of any such Cayley complex. However, they are 1-skeletons of an *almost planar* Cayley complex, that is, a Cayley complex that can be mapped to $\mathbb{R}^2$ in such a way that the images of the interiors of any two 2-simplices are either disjoint or one of them 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.

Corollary 5.3. *Every cubic planar Cayley graph of connectivity 2 is the 1-skeleton of an almost planar Cayley complex.*

This fact can be easily seen by considering the embeddings we constructed and the presentations we chose.

Corollary 5.3 is extended in [9] to all cubic planar Cayley graphs by more involved arguments.

Acknowledgement

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

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

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

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

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

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

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

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

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

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

[12] G. Kozma. Percolation, perimetry, planarity. *Rev. Mat. Iberoamericana*, 23(2):671–676, 2007.

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