Enumerating planar locally finite Cayley graphs

David Renault*

September 11, 2003

Abstract

We characterize the set of planar locally finite Cayley graphs, and give a finite representation of these graphs by a special kind of finite state automata called labeling schemes. As a result, we are able to enumerate and describe all planar locally finite Cayley graphs of a given degree. This analysis allows us to solve the problem of decision of the locally finite planarity for a word-problem-decidable presentation.

Keywords: vertex-transitive, Cayley graph, planar graph, tiling, labeling scheme

Introduction

Given a group G and a set of generators A, the associated Cayley graph is a natural representation of the action of the generators on G. When a Cayley graph is planar, its embedding possesses symmetry properties induced by the choice of the geometry. The planarity hypothesis also casts a new light on many decidability problems on groups that are unsolvable in the general case.

The problem of enumerating planar Cayley graphs goes back to Maschke [Mas96], who enumerated in 1896 all finite groups having planar Cayley graphs, these groups corresponding to the finite subgroups of the symmetry group of the sphere. Nevertheless, for a given group G, having a planar Cayley graph is a Markov property [LS77], hence there is no algorithm deciding whether a finitely generated group possesses a presentation for which its Cayley graph is planar or not.

Presentations of groups having planar Cayley graphs have already been discussed by Levinson [Lev79], who describes the geometry of his graphs in terms of point and weak-point symmetry and gives necessary conditions on the presentations of these groups. More recent results by Droms et al. [DSS98, Dro01] develop the links between connectivity and planarity inside Cayley graphs, which give hints about the structure of the group. Considering the broader problem of describing the vertex-transitive planar graphs, Imrich and Fleischner [FI79] solved the question in the case of finite graphs.

In this paper, we give an exhaustive description of the class of the locally finite planar Cayley graphs, and answer some decidability problems. Our results are based on Chaboud's work [Cha95], who studied the planar Cayley graphs

*Laboratoire Bordelais de Recherche en Informatique, 351, cours de la Libération, Université Bordeaux I, Talence, France – Email: `renault@labri.fr`

that were also normal tilings. We represent Cayley graphs by finite state automata called *labeling schemes* along with a local geometrical invariant called a *type vector*. We show that there exists a bijection between this representation and the class of locally finite planar Cayley graphs. This extends Chaboud's work by providing a more general class of graphs. We give algorithmic means to describe each of these Cayley graphs, along with their presentations. Our main result (p. 16) is :

Theorem 4.4 (Enumeration)
Given $n \geq 2$, it is possible to effectively enumerate all planar locally finite Cayley graphs having internal degree n, along with one representative presentation.

Each Cayley graph belonging to this class is effectively computable (*i.e.* there exists an algorithm able to build any finite ball of the graph), which leads to algorithms solving the word problem. When the word problem is already decidable in a given presentation, it is possible to decide whether the corresponding Cayley graph is planar locally finite or not. These graphs also correspond to tilings of the plane (seen as spherical, Euclidean or hyperbolic) by regular polygons, and their automorphisms are realized by isometries. Our techniques are similar to those used for groups acting on planar surfaces [Wil66, ZVC80].

Contents

Introduction **1**

1 Definitions **3**

2 Properties of the locally finite planar graphs **4**

3 Rotation and inversion **7**
- 3.1 3-connected case 7
- 3.2 2-separable case 8

4 Labeling schemes **11**
- 4.1 Definitions and properties 11
- 4.2 Characterization 13
- 4.3 Combinatorial approach 16

Discussion **18**

Enumerations of Cayley graphs **19**

1 Definitions

A *directed graph* Γ consists of a pair (V, E), V being the set of *vertices* and $E \subset V \times V$ the set of *edges*. Each edge corresponds to a pair of vertices (s_1, s_2), s_1 being the *initial* vertex and s_2 the *terminal* one. A *labeling* of the graph is an application from the set of edges into a finite set L of labels. The *degree* of a vertex is the number of edges incident to this vertex. A *morphism* from the graph $\Gamma_1 = (V_1, E_1)$ into $\Gamma_2 = (V_2, E_2)$ is an application $\sigma : V_1 \to V_2$ that preserves the edges of the graphs. When both graphs are labeled, we impose that the morphisms also preserve the labels of the edges. A *path* of Γ is a sequence of edges (e_n) in Γ such that $\forall n$, the terminal vertex of e_n and the initial vertex of e_{n+1} are the same. A *cycle* is a finite path whose initial and terminal vertices are the same. A *loop* is a cycle containing a single edge.

Let G be a finitely generated group, and $A = \{a_1, \ldots a_d\}$ a set of generators of G. The *Cayley graph* Γ of G with respect to this set of generators is the directed graph defined as follows : the vertices of Γ are the elements of G, and we draw an edge labeled by a generator a_i from a vertex g_1 to a vertex g_2 whenever $g_1 \cdot a_i = g_2$ in the group G. We shall only be concerned with the case where the graph contains no loops. Therefore, we prevent A from containing the identity of G. We impose that A be stable by inversion in G. Thus, to any edge linking s_1 to s_2 and labeled by a_i, there exists a reciprocal edge from s_2 to s_1 labeled by a_i^{-1}. To avoid redundancy, we reduce both edges into a single undirected edge linking s_1 to s_2.

A graph is said to be *planar* if it can be embedded in the plane, without edges crossing or intersecting other vertices than their extremities. Such an embedding is said to be *locally finite* if its vertices have no accumulation point in the plane – equivalently there exists at most one accumulation point in the sphere. Our embeddings will be considered *tame*, meaning that all edges are $\mathcal{C}^1$ images of $[0; 1]$. A graph is said to be *vertex-transitive* if and only if, given any two vertices $(s_1, s_2) \in \Gamma$, there exists an automorphism of Γ mapping s_1 onto s_2. If Γ is the Cayley graph of a group G, then it is vertex-transitive.

A graph is *connected* if, for every pair of vertices (s_1, s_2) of the graph, there exists a finite path in the graph with extremities s_1 and s_2. A *n-separation* is a set of n vertices whose removal separates the graph in two or more connected components. A *cut-vertex* of Γ is a 1-separation of Γ. A graph is *n-separable* if it contains a n-separation. If it contains no n-separation, it is *(n+1)-connected*. A graph is *uniform* when all its vertices have the same degree d. The graphs we will be dealing with are connected and uniform. There exists only one non-trivial Cayley graph of degree 1, which is the Cayley graph of $\mathbb{Z}/2\mathbb{Z}$, corresponding to a single edge. Cayley graphs of degree 2 corresponds to cyclic groups $\mathbb{Z}/n\mathbb{Z}$ where n might be infinite. Thus, we shall only be interested in graphs of degree $d \geq 3$.

Given a specific embedding of a locally finite planar graph, a *face* $\mathcal{F}$ is defined as an arc-connected component of the complement of the graph in the plane. $\mathcal{F}$ is said to be *finite* when it is bounded in the plane, otherwise it is said to be *infinite*. The *border* of this face, noted $\partial\mathcal{F}$, is its boundary in topological terms.

2 Properties of the locally finite planar graphs

In the following, Γ will be a locally finite, planar, connected, vertex-transitive graph. Thus we will speak of vertices, edges and even faces of Γ, as defined above. Our major concern when dealing with such a graph is the following : having set aside the possibility of accumulation points of the vertices, which properties remain when we take into account the whole embedding ? The following approach explores the geometric properties of our graphs.

Lemma 2.1 (Edge locally finite)
Let K be a compact set of the plane. It is possible to modify the embedding of Γ such that the number of edges of Γ intersecting K is finite.

Proof : Let K be a compact set of the plane, intersecting infinitely many edges of the embedding of Γ. Since this set is compact, it meets only a finite number of vertices of Γ, and since the graph is of finite degree, only a finite number of edges intersecting K are incident to vertices effectively belonging to K.

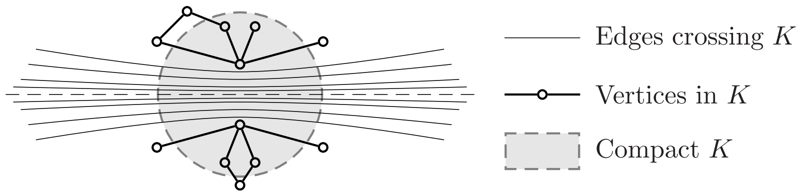

Consider now the connected components of the graph restricted to K (we consider only the edges and vertices entirely belonging to K). If there is just one component, then it is possible to modify the embedding by pushing the accumulating edges outside of the compact, while preserving the locally finite aspect of the graph. If there is more than one component, then consider the subset of edges intersecting K but not incident to vertices of K, each edge corresponding to a couple of vertices (s, t). Since the graph is locally finite, it is possible to extract from this set a series of edges $(s_n, t_n)_{n\in\mathbb{N}}$ such that $d(s_n, K) \to \infty$ and $d(t_n, K) \to \infty$. Thus the graph can not be connected, since there is no way of joining the two components without crossing an infinite number of edges belonging to the series (s_n, t_n), which is impossible. $\square$

Following this lemma, it is possible to construct an *edge-locally finite* embedding of Γ in the following sense : by considering $K_0 \subset \dots \subset K_n \subset \dots$ a series of compacts sets of the plane such that every point of the plane eventually falls into a K_n, and pushing all accumulation points towards infinity. This property allows us to consider that the embedding of Γ contains no accumulation points of the vertices. For such an embedding, the set of points of Γ, adding the point at infinity if Γ is an infinite graph, is a closed subset of the sphere. Faces of Γ are therefore connected open subsets of the sphere. From now on, we will consider that the embedding possesses this property.

Lemma 2.2 (Non-intersection of faces)
The border of a face of Γ has no self-intersection.

Proof : Assume on the contrary that Γ possesses a vertex s belonging at least twice to the border of a given face $\mathcal{F}$. There exists a ball B centered on s such that $B \cap \mathcal{F}$ possessed more than one connected component. And there exists a path drawn in the face connecting one of these components to the other. In concrete terms, the situation looks like the following :

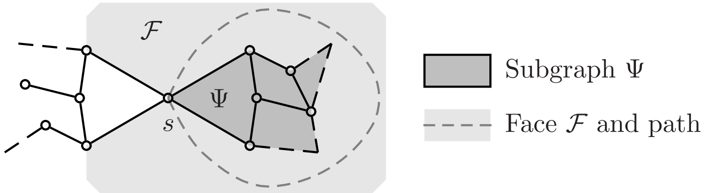

Since Γ is locally finite, we consider the finite subgraph Ψ of Γ contained in one finite connected component of the plane separated by the path drawn inside the face. Call φ an automorphism of Γ mapping s on a vertex of Ψ different from s. Since s must be a cut-vertex of Γ, $\varphi(s)$ is a cut-vertex of $\varphi(\Gamma)$, which separates Ψ into two distinct components, one of which must be finite and entirely inside Ψ, otherwise Ψ would be infinite. This component then corresponds to $\varphi(\Psi)$. But this means that the number of vertices in Ψ is strictly greater than the number of vertices of $\varphi(\Psi)$, which is impossible. $\square$

The preceding lemma implies that around a given vertex of Γ, and a face incident to that vertex, there exists at most two edges incident both to the vertex and to the face. Let's consider more closely the borders of the faces. Our embeddings are considered to be *tame*, meaning that all edges are $\mathcal{C}^1$ images of the unit segment. Tame embeddings possess the following sructure property :

Lemma 2.3 (Border structure)
Let $\mathcal{F}$ be a face of Γ. The border of $\mathcal{F}$ is either a finite cycle of Γ or a two-way infinite path of Γ and its limit point at infinity. In particular, if (s_1, s_2) are any two vertices belonging to this border, then there exists a unique finite path in the border with extremities s_1 and s_2.

Proof : Consider a face $\mathcal{F}$ of Γ and a point $x \in \partial\mathcal{F}$ that is not at infinity. Then x is either on a vertex or on an edge of Γ. We build a non self-intersecting path going through x if x belongs to an edge, or going through the two edges incident to $\mathcal{F}$ if x is a vertex. Thanks to the previous lemma, there exists either a cycle or a two-way infinite path in Γ possessing this property. In the second case, the locally finite property implies that the extremities of the path must join at infinity. In either alternative, this path constitues a simple closed continuous curve of the sphere. According to the Jordan's theorem, it divides the sphere in exactly two connected components, the curve being the complete frontier of both parts. One of these components, for example $\mathcal{H}$, must contain $\mathcal{F}$.

Consider a ball B centered on x such that B only meets the edges incident to x. Therefore, from lemma 2.2, $\mathcal{H} \cap B = \mathcal{F} \cap B$. This implies that around x, the border of $\mathcal{F}$ is the border of $\mathcal{H}$, that is to say the path in Γ going through x. Applying this property to every point belonging to $\partial\mathcal{F}$, the border of $\mathcal{F}$ must entirely contain the edges it is incident to, and any vertex on its border is incident to two edges in the border. This proves the lemma. $\square$

This confirms the following natural intuition : the border of a finite face is composed of a finite number of edges, whereas it is infinite in the case of an infinite face.

Lemma 2.4 (Rule of intersection of the faces)
The intersection of two different faces of Γ*, when not empty, is either a single vertex, or an edge along with its extremities.*

Proof : Consider $\mathcal{F}_1$ and $\mathcal{F}_2$ two distinct faces of Γ. Suppose that $\partial\mathcal{F}_1 \cap \partial\mathcal{F}_2$ contains two distinct vertices s_A and s_B of Γ. Following lemma 2.3, these two vertices s_A and s_B can be joined by a path in $\partial\mathcal{F}_1$, and by a path in $\partial\mathcal{F}_2$. If these paths are the same, then they must be reduced to a single edge because the degree of the graph is ≥ 3. Otherwise, we impose that the finite region enclosed by the reunion of this paths does not contain either $\mathcal{F}_1$ or $\mathcal{F}_2$. Concretely, the disposition of the graph is the following :

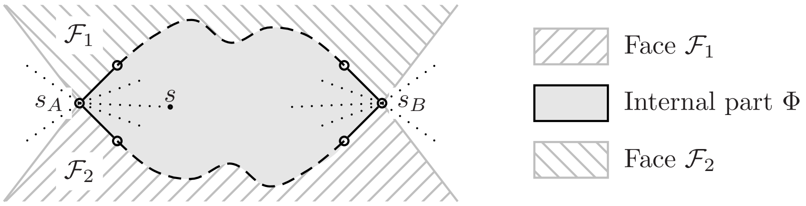

Hence the cycle is a true Jordan curve in the plane. The gray part of the picture corresponds to the finite connected component of the plane separated by this curve. Call Φ the finite subgraph of Γ inside this component. We proceed by induction on the number of vertices in Φ. If Φ contains only two distinct vertices, then the result is true. Consider now a vertex $s \in \Phi$, different from s_A and s_B, and call φ the automorphism of Γ mapping s_A upon s.

Let's examine $\varphi(\Phi)$. $\Phi \subsetneq \varphi(\Phi)$ is impossible, because Φ and $\varphi(\Phi)$ have the same number of vertices. Likewise, $\varphi(\Phi) \subsetneq \Phi$. Suppose *ab absurdo* that $\Phi = \varphi(\Phi)$. There exists a neighbor of s_A in Γ not in Φ, otherwise $\mathcal{F}_1 = \mathcal{F}_2$. This neighbor, is mapped by φ onto a vertex at distance 1 from s, which means a vertex of Φ. Since $\varphi_{|\Phi}$ is a bijection, this is absurd. Therefore, $\varphi(\Phi)$ intersects Φ and $\Gamma\backslash\Phi$.

Take $t \in \Gamma$ a vertex both in $\varphi(\Phi)$ and $\Gamma\backslash\Phi$. Since Φ is connected, consider a simple path in $\varphi(\Phi)$ joining s and t. Since s_A and s_B are the only cut-vertices between Φ and $\Gamma\backslash\Phi$, this path must contain at least one of them. It can't contain both of them, otherwise, when removing Φ from Γ, you obtain one or two connected components, while removing $\varphi(\Phi)$ from Γ, you obtain two or three connected components. Therefore, without loss of generality, if s_B is the only cut-vertex in $\varphi(\Phi)$, then every path from s to t contains s_B. Every path from s_A to $\varphi^{-1}(t)$ contains $\varphi^{-1}(s_B)$. In particular, $\varphi^{-1}(s_B)$ is a cut-vertex for Φ, which is absurd under our hypothesis. By induction, these parts are edges of Γ. Since the degree of Γ is ≥ 3, it is impossible to obtain such a thing. $\square$

Corollary 2.5
The automorphisms of Γ *map finite faces of* Γ *into finite faces of* Γ*.*

Proof : Consider the border of a finite face $\mathcal{F}$. Let φ be an automorphism of Γ mapping $\partial\mathcal{F}$ on a cycle of Γ of the same length. This cycle defines a Jordan curve of the plane. Let Φ be the finite subgraph in the finite connected component delimited by this cycle, we examine $\varphi^{-1}(\Phi)$. This subgraph is connected to $\partial\mathcal{F}$, so it intersects $\mathcal{F}$. Since it can't be connected to another vertex in Γ than the vertices incident to $\mathcal{F}$, it must lie inside another face of Γ, namely $\mathcal{H}$. Following the preceding lemma, the intersection of $\mathcal{F}$ and $\mathcal{H}$ is entirely in the border of $\mathcal{F}$. Hence $\Phi = \emptyset$. $\square$

3 Rotation and inversion

The geometrical properties of planar graphs are often affected by their connectivity. For example, 3-connected planar graphs have a strong geometric property : they possess a unique embedding in the plane[Whi32] or equivalently their dual is uniquely defined. Chaboud [Cha95] studied graphs possessing normality properties, which in fact correspond to 3-connectivity and local finiteness. The connectivity of the graph is directly linked to the number of infinite faces around a vertex :

Lemma 3.1 (Infinite faces)
Γ is 2-separable $\Leftrightarrow$ each vertex is incident to a unique infinite face. Γ is 1-separable $\Leftrightarrow$ each vertex is incident to two or more infinite faces.

In this section, we first recall the lemmas of Chaboud for 3-connected graphs, and then extend his results in the complementary case of the 2-separable graphs.

3.1 3-connected case

In the 3-connected case, all faces of the graph are finite. Consider any vertex s of Γ : as a definition, the generators appearing around s are the labels of the edges of terminal vertex s. The planarity of Γ imposes a cyclic order on these generators.

Lemma 3.2 (Chaboud's rotation lemma)
If Γ is a Cayley graph with respect to a set of generators A, then the generators appear in the same cyclic order around each vertex, up to the direction of rotation.

Proof : According to corollary 2.5, the faces incident to a given vertex s are mapped, by an automorphism φ, onto the faces incident to the image vertex. Consider the set of faces incident to s as a finite set of cycles containing s, intersecting only on the edges incident to s :

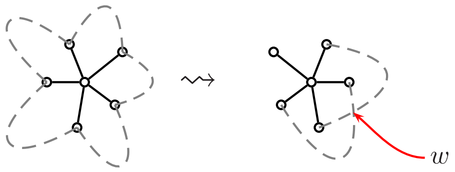

Since faces are mapped onto faces by automorphisms of the graph, then adjacent faces are mapped onto adjacent faces. Therefore, the borders of the images of the faces incident to s can't cross, which implies that the order of rotation of the generators is the same. $\square$

Lemma 3.3 (Chaboud's inversion lemma)
If an edge labeled by a generator a_i joins two vertices of different orientations, then it is the case for every edge labeled by a_i.

Proof : The proof is essentially the same as in the preceding lemma, but instead of considering the paths of the border of the faces incident to a vertex, we consider an edge e labeled by a_1, and the two paths corresponding to the borders of the faces containing e.

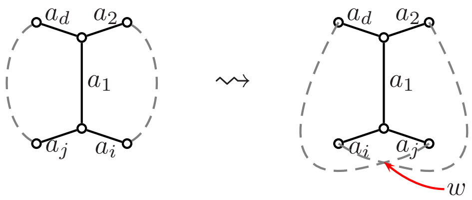

Since these paths do not cross before the automorphism, they can't cross after having been mapped. Therefore if the generator labeling e inverses the direction of rotation of the generator, it is also the case of all edges labeled by a_1. $\square$

A generator a_i is said to be *direct* if it always joins vertices of the same orientation. It is said to be *indirect* in the other case. As a simple corollary, we have that :

Lemma 3.4 (Type Vector)
Let s be a vertex of Γ, and $\{\mathcal{F}_1, \dots, \mathcal{F}_d\}$ the set of faces of Γ incident to s, ordered in cyclic order around s. Let $\mathcal{V} = [k_1, \dots, k_d]$ where k_i is the number of edges incident to $\mathcal{F}_i$. Then $\mathcal{V}$ is independent of the vertex s, up to a cyclic permutation or reversal of the vector.

This vector is called the *type vector* of the graph. Not all possible type vectors correspond to a locally finite graph Γ : [GS87] describes the special case of the Euclidean plane, where among the 21 possible type vectors for tilings by regular polygons, only 12 correspond to vertex-transitive graphs, whereas the remaining 9 can't possibly describe a vertex-transitive graph. Moreover, to a given type vector may correspond multiple non-isomorphic vertex-transitive graphs, or even Cayley graphs, as is the case for $[4, 4, 4, 6]$:

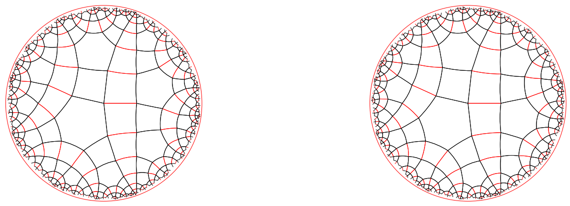

3.2 2-separable case

In the case of connectivity lower than 3, infinite faces appear. For instance, if we consider the Cayley graphs of free groups (infinite regular trees), one can easily devise embeddings that violate the preceding lemmas. A graph Γ is said to possess a *proper labeling* or to be *weakly point-symmetric* [Lev79] if, for a specific embedding, it verifies the rotation lemma and the inversion lemma. The previous subsection shows that 3-connected graphs must have a proper labeling, but this is not general. Our purpose in this section is to prove that, nevertheless, there always exists an embedding of the graph that follows the lemmas of rotation and inversion.

Let Γ be a 2-separable, locally finite, planar and connected vertex-transitive graph. If any vertex of Γ is a cut-vertex, then every vertex is a cut-vertex. If we slit open the graph along its cut-vertices, the remaining connected components are called *2-connected components* or CC_2. Since the graph is vertex-transitive, the set of components incident to a vertex is independent of the vertex, and finite. These components are not necessarily vertex-transitive themselves, and they may be finite or infinite. By hypothesis, 2-connected components are locally finite graphs.

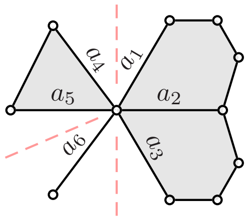

Suppose now that Γ is a Cayley graph. Two generators a_i and a_j are said to be equivalent ($a_i \sim a_j$) when the pair of edges incident to a vertex and labeled by a_i and a_j belong to the same CC_2. This is an equivalence relation on the set of generators A. If a generator is alone in its class, then the CC_2 consists only of a single edge. Let's consider an equivalence class $\mathcal{C}$ for this relation. If s is a vertex of Γ, it is possible to join the two vertices $s \cdot a_i$ and $s \cdot a_j$ by a path remaining in the component corresponding to $\mathcal{C}$ and avoiding the vertex s. The number of classes is equal to the number of CC_2 appearing around a vertex. If Γ is 2-connected, then it possesses a unique CC_2 equals to Γ.

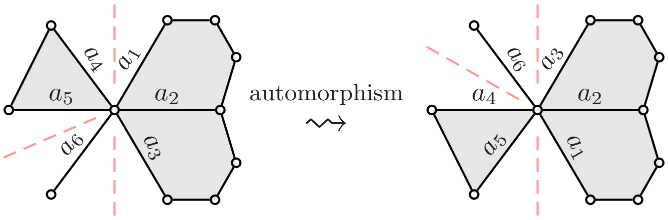

As in the 3-connected case, the order of the generators around a vertex in a given class is independent of the vertex, allowing for the two directions of rotation. The automorphisms of Γ may only modify (a) the cyclic order of the CC_2 appearing around a given vertex and (b) apply a symmetry to a given CC_2. Let's consider the set of CC_2 around a vertex, namely $\{\mathcal{C}_1, \ldots, \mathcal{C}_p\}$. This vertex defines a cyclic order on the set of generators. Our purpose consists in building an embedding of the graph such that the generators appear in the same cyclic order around each vertex.

Example 3.1 – The following graph is an example of 1-separable Cayley graph of degree 5. Its sole 2-connected component, up to isomorphism, is composed of two triangles joined by an edge. Notice that two of these components appear around each vertex, and that they are not equivalent : one component has three edges incident to the vertex, while the other has only two. The automorphisms of the graph may modify the ordering of the generators around a given vertex by twisting these components. A possible presentation for this group is :

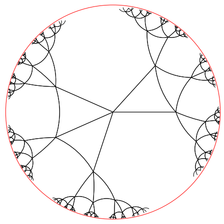

$$G = \left\langle a_1, \ldots, a_5 \;\middle|\; a_4^2, a_2a_3, a_1a_5, a_3a_4a_1 \right\rangle$$

Proposition 3.1 (Proper labeling)
There exists an embedding of Γ *that obeys the inversion and rotation lemmas. Therefore the graph* Γ *possesses a proper labeling of its edges.*

Proof : Consider a vertex s of Γ, and define with this vertex an ordering of the generators of the group. Let's build an embedding of Γ by induction : at the n-th step of the induction, the graph Γ_n is a locally finite graph, isomorphic to a subgraph of Γ, such that every vertex in Γ at distance less than n from s belongs to this subgraph. We also impose that the distance in the plane between a vertex and s be superior to the distance in Γ.

Let's choose for s a point in the plane. Let $\mathcal{C} = \{\mathcal{C}_1, \ldots, \mathcal{C}_p\}$ be the set of CC_2 around a vertex in Γ. By symmetry and rearrangement of the CC_2, it is possible to define an order on $\mathcal{C}$ which is coherent with the order of the generators in each CC_2. Each CC_2 being locally finite and independent in Γ of the others, it is possible to embed each one of them such that the distance condition be verified, and that the order of the generators around s be correct, up to a cyclic permutation. We call this graph Γ_1.

Suppose now that we have built Γ_n, and consider the finite set of vertices at distance $n+1$ around which we have not built all the necessary CC_2 yet. Around each vertex in this set, we have only built one CC_2. If this CC_2 is reduced to a single edge, then we choose the order of the generators around the vertex to be the same that the order around the other extremity. Otherwise, the order is already defined by the order of the generators in the CC_2.

Let's consider the face incident to our vertex that does not belong to the component already attached to the vertex. From lemma 2.3, the border of this face is either finite, or a bi-infinite path whose only accumulation point is at the infinity. In each case, it is always possible to embed the missing CC_2 in this face, such that the graph still verifies the distance condition and that the order of the generators around s is correct. After having added a finite number of components, all the vertices at distance $n+1$ are completed, and we obtain Γ_{n+1}.

The limit graph is isomorphic to Γ. Since the distance condition is verified, it is locally finite, and by construction, it obeys the rotation lemma. During the construction, we have chosen the generators alone in their equivalence class to be direct. All the others follow the inversion lemma by an argument similar to Chaboud's lemmas. Hence, this embedding obeys the inversion lemma, as long as the rotation lemma. $\square$

For such a graph, it is therefore possible to extend the notion of type vector that we defined in the 3-connected case. Around a vertex, we may have finite or infinite faces. And considering the embedding provided by the previous proposition, the automorphisms of Γ also map infinite faces onto infinite faces. Therefore we allow infinite faces in type vectors, the type vector of Γ being independent of the chosen vertex. For example, the free group on two elements, with the usual generators, possesses a planar Cayley graph of degree 4, and its type vector is $[\infty, \infty, \infty, \infty]$.

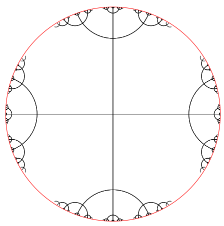

4 Labeling schemes

4.1 Definitions and properties

Let $A = \{a_1, \dots, a_d\}$ a set of generators. A *labeling scheme* on A is given by a 3-tuple (p, σ, τ) :

- a cyclic permutation p on A;
- an involution σ on A ($\sigma = \sigma^{-1}$);
- a function $\tau : A \to \{-1; 1\}$ such that for each $x \in A$, $\tau(\sigma(x)) = \tau(x)$.

The permutation p corresponds to the cyclic ordering of the generators around any vertex of Γ. The involution σ links every generator with its inverse in G. The function τ assigns the value 1 to a generator if it is direct, and -1 if it is indirect. The previous section assures that every locally finite planar Cayley graph possesses a labeling scheme that is coherent with its embedding, *i.e.* describes the order of rotation of its generators around any vertex and the direct/indirect property of the generators. Since we will only be concerned with Cayley graphs, we can choose $p = [1, 2, \dots, d]$ by renaming the generators of the graph.

Using the labeling scheme, it is possible to describe the border of the faces with a finite state automaton, using the general ideas from maps and hypermaps [CM92]. Let's consider a vertex s from Γ around which the generators appear in cyclic order along the clockwise direction. If we follow the edge incident from s labeled by the generator a_i, we can compute the direction of rotation of the generators around the terminal vertex of the edge. Then it is possible to rebuild the border of any face of the graph, be it finite or infinite.

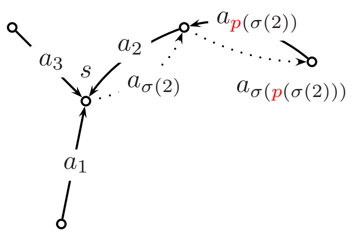

Case were a_2 is direct

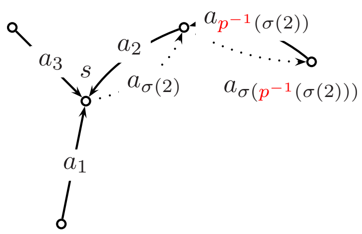

Case were a_2 is indirect

For a more formal point of view, let's consider the set of vertices $C = C^+ \cup C^-$, where C^+ is composed of d vertices $\{1^+, \ldots, d^+\}$ and C^- is composed of $\{1^-, \ldots, d^-\}$. Let s be a vertex of C. We define an action of $\mathbb{Z}$ onto the set of vertices of C by :

1. if $s^+ \in C^+$, then let $t = \sigma(p(s))$.
   If t is direct, then $1 \cdot s^+ = t^+$, otherwise $1 \cdot s^+ = t^-$.

2. if $s^+ \in C^-$, then let $t = \sigma(p^{-1}(s))$.
   If t is direct, then $1 \cdot s^- = t^-$, otherwise $1 \cdot s^- = t^+$.

This action can be naturally extended into an action of monoid. Moreover, it is invertible as a result of the properties of the labeling scheme. Therefore it is an action of the group $\mathbb{Z}$ on a finite set. The orbits of this action correspond to the borders of the faces, described by the generators labeling them, in function of the direction of rotation of the vertices.

It is possible to represent this action with a finite state automaton, each state being a couple (a_i, b) where $a_i \in A$ is a generator and $b \in \{true, false\}$ being a boolean. Our alphabet is composed of two different letters, p for following to the next generator in the current direction of rotation, and i for reversing the generator :

- $(a_i, true)\overrightarrow{p}(a_{(i+1 \bmod d)}, true)$ and $(a_i, false)\overrightarrow{p}(a_{(i-1 \bmod d)}, false)$;
- if $\tau(i) = 1$, $(a_i, b)\overrightarrow{i}(a_{\sigma(i)}, b)$ otherwise $(a_i, b)\overrightarrow{i}(a_{\sigma(i)}, \neg b)$.

Computing the border of a face consists in selecting a starting state in the automaton, and reading the infinite word $(pi)^\omega$, thus obtaining a finite cycle of the automaton.

Example 4.1 – The following example has been taken from [Cha95]. The left picture describes the labeling scheme as defined by Chaboud : the cyclic permutation p corresponds to a cyclic graph of order 8, and the involution σ corresponds to the black arcs linking the vertices of the cycle. A direct generator is pictured as a circle, an indirect one as a square.

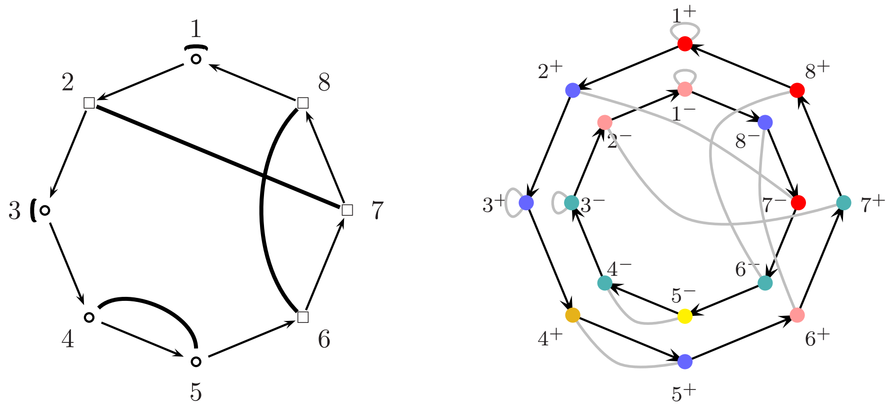

The picture on the right describes the action of $\mathbb{Z}$ on the generators. The action of $1 \in \mathbb{Z}$ on a vertex is computed as follows : follow the oriented black arc starting from the generator, then follow the non-oriented gray arc to obtain the resulting generator. By alternatively following black and gray arcs, one follows the orbit of the generator. —

Different orbits can correspond to same faces. This may happen if $[a_1^{s_1} \ldots a_k^{s_k}]$ and $[a_1^{-s_1} \ldots a_k^{-s_k}]$ are two different orbits, or if an orbit is composed of the inverses of the generators of another orbit. Two such orbits are called *dual orbits*. Dual orbits correspond to borders of the same face, but read in the opposite direction. We say that two generators a_i and a_j *correspond to the same face* if a_j^+ belongs either to the orbit of a_i^+ or to its dual orbit.

For a generator a_i pointing toward the vertex s, let k_i be the smallest positive integer such that $k_i \cdot a_i^+ = a_i^+$. The vector $[k_1, \ldots, k_d]$ is called the *primitive vector* of the labeling scheme. Another type vector $[l_1, \ldots, l_d]$ is a *valid type vector* with regard to this labeling scheme if and only if :

- $\forall i \in [1; d], l_i$ is a multiple of k_i;
- if a_i^+ and a_j^+ correspond to the same face, then $l_i = l_j$.

Example 4.2 – Continuing with the previous example, the orbits under the action of $\mathbb{Z}$ are :

- (1^+ 7^- 8^+) with dual (6^+ 2^- 1^-), of length 3 ;
- (2^+ 3^+ 5^+ 8^-) with dual (7^+ 6^- 4^- 3^-), of length 4;
- (4^+) with dual (5^-) of length 1.

Therefore the primitive type vector corresponding to the labeling scheme is $[3, 4, 4, 1, 3, 4, 3]$. The case of length 1 means that any length is allowed for the face, assuming that it is greater than 3, the smallest possible length for a face to exist. The generators 1, 8 and 6 belong to the same face, as is the case for 2, 3, 5 and 7. All valid type vectors are then of the form $[3p, 4q, 4q, r, 4q, 3p, 4q, 3p]$, where (p, q, r) are positive integers, such that $r \geq 3$. —

We can extend this construction to type vectors with infinite faces. We allow the number ∞ in the type vectors, to stand for infinite faces. If $l_i = \infty$, then the first condition of validity is always true. These type vectors can be seen as limits of finite type vectors, as we shall see in the following sections.

4.2 Characterization

Every locally finite planar Cayley graph possesses a labeling scheme, and a type vector that is valid with regard to this scheme. We will now take an interest in the converse of this result : given a random labeling scheme, and a valid type vector, is it possible to produce a Cayley graph and the corresponding group that have the same labeling scheme and type vector ?

Our goal in this section consists in building specific embeddings of Cayley graphs with particular geometrical properties. Depending on the graph, we select an appropriate geometry : Euclidean, spherical or hyperbolic. As long as our graphs are locally finite, embeddings in spherical geometry will lead to finite graphs. Infinite faces thus can only occur in Euclidean and hyperbolic geometry. The following theorem constructs embeddings that are tilings of the plane by regular polygons. A *regular polygon* is a non-intersecting cyclic graph embedded in the plane such that all edges have the same length and the angle between two consecutive edges is set. This allows for finite or infinite polygons. In the Euclidean plane, an infinite regular polygon corresponds to a half-plane, with border a straight line. In the hyperbolic plane, there exists infinitely many non-isometric infinite polygons[1].

Theorem 4.1 (Existence)
Given a labeling scheme (p, σ, τ), and a valid type vector $[k_1, \dots, k_d]$, there exists a locally finite planar Cayley graph possessing this scheme and type vector. Moreover, all faces of the embedding are regular polygons.

Proof : The proof of this result involves two different parts : first, we show it is possible to glue around a single vertex regular polygons corresponding to the type vector. Then, we build the entire graph by induction, as in proposition 3.1.

Consider any point in the plane. We are going to evaluate the interior angle of a regular polygon of side length l, with k_i sides, and then the total angle corresponding to all the polygons in the type vector must be equal to 2π :

$$\theta_i(l) = \begin{cases} \text{Spherical :} & 2\arcsin\left(\frac{\cos(\pi/k_i)}{\cos(l/2)}\right) \\ \text{Euclidean :} & \frac{(k_i-2)}{k_i}\pi \\ \text{Hyperbolic :} & 2\arcsin\left(\frac{\cos(\pi/k_i)}{\cosh(l/2)}\right) \end{cases} \qquad \text{and} \qquad \sum_{i=1}^{d} \theta_i(l) = 2\pi$$

This equation determines the choice of the geometry : let Σ be the sum of the angles in the Euclidean plane. If $\Sigma = 2\pi$, we choose the Euclidean geometry, and any value is possible for l (the corresponding graphs will be homothetic). If $\Sigma < 2\pi$, we choose the spherical geometry, whereas if $\Sigma > 2\pi$ we select the hyperbolic geometry. In any case, there exists an unique solution of the equation[2] that determines a unique value for the length l, and consequently for the angles θ_i. These values allow us to draw all the regular polygons that correspond to our type vector around a given point of the plane.

We will now build the graph by induction : let ϵ be a vertex of Γ. The graph Γ_n is a planar locally finite graph, isomorphic to a subgraph of Γ, containing all vertices of Γ at distance $\leq n$ from ϵ. The first graph Γ_1 is composed of the glued polygons of the first case. Suppose now that we have built Γ_n. Consider the finite set of vertices of Γ_n at distance n for which all the incident edges still have not been drawn. We shall build the remaining edges with geodesics, such that the length of the geodesic be l, and the angle at the vertex between the edge labeled by a_i and the edge labeled by a_{i+1} be equal to $\theta_i(l)$.

The labeling scheme describes how the edges must be labeled : take a vertex s and an edge labeled by the generator a_i pointing towards s. Depending on the direction of rotation of the generators around s, it is possible to build all edges incident to s, in such a way that the angle between thed edge a_i and the edge a_{i+1} be $\theta_i(l)$. The edge labeled by a_i is seen by its other extremity t as an edge

[1]For example when the relation between the length l of the edge and the angle θ between two edges is $l = \frac{1}{2}\sqrt{2 + 2\cos\theta}$. In the Poincaré model, the vertices of the border are the orbits of a parabolic isometry : they belong to a circle tangent to the circle at infinity.

[2]Except for the following type vectors : $[3, 3, p \geq 5]$, $[3, 4, p \geq 6]$ and $[3, 5, p \geq 9]$. There does not exist a labeling scheme validating any of these type vectors (cf p 17).

labeled by $a_{\sigma(i)}$ pointing towards t. The direction of rotation of the generators around t is inverted if $\tau(i) = 1$, otherwise it remains the same. On the same principle, it is now possible to build the generators around t.

Suppose *ab absurdo* that this construction leads to two edges crossing. Then consider one of the finite faces delimited by these two edges and the remaining part of the graph. If there are no more edges to add into this face, then since our construction fixes the angles between two edges, then this angle is always the same inside the face. This comes from the fact that the construction due to the labeling scheme follows the borders of the faces. This face is a regular polygon, as the length and angles have been chosen in that sense. The only possible way the edges can then cross is either at their extremities, or being in fact the same edge. In each case, the face correspond to a regular polygon and the edges do not really cross. In the case where there still exists edges to add in the interior of the face, we reason by induction, adding the finite set of remaining edges one by one. □

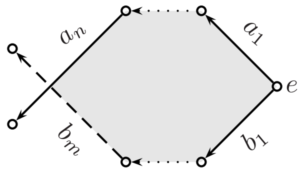

Moreover, this tiling of the plane is vertex-transitive : the construction is independent of the starting vertex, and applying this construction with two different starting vertices creates the automorphism mapping the first vertex on the second. The labeling of the edges entails that the automorphism group is simply transitive, which imposes that the graph is a Cayley graph. It is possible to give a canonical presentation of the underlying group :

- Set of generators : $A = \{a_1, \ldots, a_d\}$ (corresponding to the free group of degree d);
- Relators for the inverses : $a_i a_{\sigma(i)}$ for all $i \in [1; d]$;
- Relators for the border of the face between a_i and a_{i+1} : since for any face, the length of the corresponding orbit $\mathcal{O}_i = a_{j_1} a_{j_2} \ldots a_{j_{p_i}}$ divides the length of the face, the corresponding relator is $\mathcal{O}_i^{n_i}$ where $n_i \cdot p_i = k_i$ is the number in the type vector. For infinite faces, there is no added relator.

Example 4.3 – The type vector $[3, 4, 4, 3, 4, 3, 4, 3]$ is validated by the previous labeling scheme. This leads us to a Cayley graph of degree 8, pictured below. For this graph, $\Sigma = 14/3$, consequently our embedding is built in the Poincaré model of the hyperbolic plane. This graph has no infinite face, therefore it is 3-connected. The arrangement between the squares and the triangles is directed by the labeling scheme. Following our method, it is possible to compute any finite ball of the Cayley graph and embed it in the hyperbolic plane with regular polygons. A presentation of the underlying group is given by :

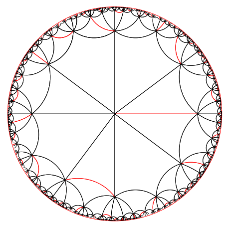

$$G = \left\langle a_1, \ldots, a_8 \;\middle|\; a_1a_6a_2, a_7a_6a_4a_3, a_5^3, a_1^2, a_3^2, a_2a_7, a_6a_8, a_4a_5 \right\rangle$$

Example 4.4 – Finite planar Cayley graphs are embedded in the sphere.

We find among them the Archimedean solids, and graphs such as the one drawn on the right on the sphere, which is called the *snub cube*. This one is of degree 5, and its type vector is $[3, 3, 3, 3, 4]$. Square faces are glued such that no two of them are adjacent. The corresponding group is of order 24, and a possible presentation of this group is given by :

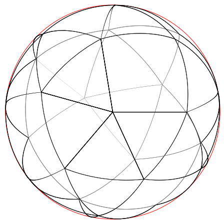

$$G = \left\langle a_1, \ldots, a_5 \;\middle|\; a_1a_2, a_3^2, a_4a_5, a_1^4, a_2a_3a_5, a_4^3 \right\rangle$$

Theorem 4.2 (Coherence)
If Γ is a locally finite Cayley graph that possesses a labeling scheme (p, σ, τ) and a type vector $[k_1, \ldots, k_d]$, then the preceding construction yields a graph isomorphic to Γ.

Proof : Consider Γ our Cayley graph and Θ the graph produced by the algorithm. These two graphs have the same sets of generators, and share in common the relators corresponding to the border of the faces. And every relator in Γ correspond to a cycle in the graph, that is to say a finite product of relators of the faces. Hence the cycles of Γ correspond to cycles of Θ, and *vice versa*. Therefore the two groups are the same, and because of the uniqueness of the construction, the graphs are isomorphic. □

Therefore, for any Cayley graph obeying our conditions, it is possible to embed the graph in a particular geometry of the plane with regular polygons. Moreover, automorphisms of the graph map faces onto isometric faces, even infinite faces. Consequently, every automorphism of Γ extends into an isometry of the plane. Since the automorphisms group of Γ is isomorphic to the group G, then G is a discrete group of isometries of the plane, which sums up to :

Corollary 4.3
Every locally finite planar Cayley graph may be embedded in the plane as a vertex-transitive tiling by regular polygons. Automorphisms of the graph then correspond to isometries of the plane.

This result appears in [Bab97] : since the tiling is made of regular polygons, the growth of the graph is approximatively equivalent to the growth of balls in the geometry. Therefore graphs drawn in the Euclidean space are of quadratic growth, whereas graphs drawn in the hyperbolic space are of exponential growth.

4.3 Combinatorial approach

Consider at first the problem of the description of the set of labeling schemes. Given a labeling scheme, it is simple to compute its primitive vector. The number of different labeling schemes of degree d is bounded by $I(d) * 2^d$, where $I(d)$ is the number of involutions in $\mathfrak{S}_d$. It is therefore possible to express the following corollary :

Theorem 4.4 (Enumeration)
Given a number $n \geq 2$, it is possible to effectively enumerate all planar locally finite Cayley graphs having internal degree n, along with their canonical presentation.

Having enumerated all possible labeling schemes of given degree, it becomes possible to count all possible presentations that lead to locally finite planar Cayley graphs. Another interest is, given a planar vertex-transitive graph that is also locally finite, to test whether this graph is a Cayley graph or not, and if the test is positive, the list of group presentations corresponding to this graph. It is already known that all archimedean tilings are Cayley graphs. Yet the graph of the icosahedron in the sphere and the $[5,5,5,5]$ isohedral tiling of the hyperbolic plane by pentagons are both vertex-transitive graphs that are not Cayley.

In terms of graph theory, it is also possible to enumerate all primitive type vectors, for a given degree and, having removed the redundant ones, to obtain all possible distinct vectors. For example, the different primitive vectors of degree 3 are given by the following :

$$\Big\{[3n,3n,3n],[n,2m,2m],[2n,2m,2p]\Big\}$$

...for values of n, m, p such that the faces of the graph are at least triangles. The set of planar graphs obtained is strictly larger than the set of hyperbolic graphs presented in [GS87] (*cf.* appendix).

Second, we devise algorithms to effectively describe the graph/tiling corresponding to a given labeling scheme. The formulas given in Theorem 4.1 allow us to compute the angles and lengths of the edges of the graph. Then it is possible to draw any finite subgraph in the given geometry. In particular, given a word on the generators, it is possible to compute approximately the position of the corresponding element of the group in the tiling in linear time on the length of the word. This leads to an algorithm to decide the word problem in quadratic time. Our algorithms are used to draw in Postscript language the graphs given in example in annex of this article.

Third, we take interest in the following problem : given a finite presentation of a group G, is it decidable whether this presentation possesses a planar locally finite Cayley graph or not ? We impose that the set of generators A of the presentation be closed under inversion. Such a presentation is said to be *full* if and only if, if l is the length of the longest relator in R, all elements of $u \in A^\star$ of length $\leq l$ such that $u = \epsilon$ in G belong to R. More precisely, the presentation contains all relators of length $\leq l$. In the case of full presentations, the problem is decidable :

Theorem 4.5 (Full presentation $\Rightarrow$ Decidable)
Let $\langle A, R\rangle$ be a full presentation of the group G. Then it is decidable whether this presentation possesses a locally finite planar Cayley graph or not.

Proof : Suppose we possess a full presentation $\langle A, R\rangle$ of a group G. Considering all relators of length 2, we already know the inverse vector σ, which maps a_i on $a_{\sigma(i)} = a_i^{-1}$. Consider n the number of generators in A. The number of labeling schemes possessing n generators, and having σ as an inverse vector is finite, and the word problem is decidable on any of the corresponding graphs. Call $\mathcal{A}$ this set schemes. It remains to find whether any of these "candidates" is the adequate Cayley graph for G. A necessary condition consists in finding whether for a given candidate, the borders of the faces belong to G or not.

We state that in a full presentation, these borders must belong to the presentation for the candidate to be valid. Suppose *ab absurdo* that the border of one of the faces does not belong to the presentation initially. Consider l the size of the longest relator in R. Either l is greater than the size of the largest face in the Cayley graph, and then it is decidable whether the relator belongs to the presentation or not, or l is smaller and it is impossible for the relators in R to generate the same group as the labeling scheme.

It suffices then to check if the remaining relators of the presentation correspond to relators in G, which corresponds to the resolution of the word problem. The exhaustion of all possible schemes leads either to a negative result, or to a labeling scheme corresponding to the initial presentation. □

Corollary 4.6 (Word problem decidable case)
If the word problem is decidable on a presentation of G, then it is decidable whether this presentation corresponds to a locally finite planar Cayley graph or not.

Discussion

The properties of locally finite planar Cayley graphs are rather strong properties. There exists two distinct directions which could enlarge our class or graphs. First, our study is restricted to Cayley graphs, when the most part of graph theory deals with vertex-transitive graphs. There exists graphs that are locally finite, planar, vertex-transitive, but not Cayley graphs of groups. Studying this class of graphs could lead to some intuition about the difference between vertex-transitive and Cayley graphs, would it be only for planar locally finite ones. Secondly, the most obvious property that we could explore is the locally finite one. Our graphs contain only one accumulation point, at most. Yet Levinson showed that the number of accumulation points, if superior to one, might be either two or infinite. The study of the behavior of these graphs in geometrical terms is also of a certain interest in terms of groups and graph theory.

References

[Bab97] László Babai. The growth rate of vertex-transitive planar graphs. In *Proc. 8th Ann. Symp. on Discrete Algorithms*, pages 564–573. New Orleans LA, ACM—SIAM, 1997.

[Cha95] Thomas Chaboud. *Pavages et Graphes de Cayley planaires*. PhD thesis, Ecole Normale Supérieure de Lyon, Février 1995.

[CM92] R. Cori and A. Machi. Maps, hypermaps and their automorphisms : a survey. *Expositiones Mathematicae*, 10:403–467, 1992.

[Dro01] C. Droms. Planar Cayley graphs. (unpublished), February 2001.

[DSS98] C. Droms, B. Servatius, and H. Servatius. Connectivity and planarity of Cayley graphs. *Beiträge zur Algebra and Geometrie*, 39:269–282, 1998.

[FI79] H. Fleischner and W. Imrich. Transitive planar graphs. *Mathematica Slovaca*, 29(2):97–105, 1979.

[GS87] B. Grünbaum and G. C. Shephard. *Tilings and Patterns.* W. H. Freeman Company, 1987.

[Lev79] H. Levinson. Planar Cayley diagrams : Weak point-symmetry. *Congressus Numerantium*, 24:679–697, 1979.

[LS77] R.C. Lyndon and P.E. Schupp. *Combinatorial Group Theory.* Springer-Verlag, 1977.

[Mas96] H. Maschke. The representation of finite groups, especially of the rotation groups of the regular bodies of three and four-dimensional space, by Cayleys color diagrams. *American Journal of Mathematics*, 18:156–194, 1896.

[Whi32] H. Whitney. Congruent graphs and the connectivity of graphs. *American Journal of Math.*, 54:150–168, 1932.

[Wil66] H.C. Wilkie. On non-eulidean crystallographic groups. *Mathematische Zeitschrift*, 91:87–102, 1966.

[ZVC80] H. Zieschang, E. Vogt, and H.D. Coldewey. *Surfaces and planar discontinuous groups.* Number 835 in Lecture Notes in Mathematics. Springer, 1980.

Enumerations of Cayley graphs

In the following sections[3] we enumerate all locally finite planar Cayley graphs of small inner degree (from 2 to 4). Cayley graphs of degree 2 correspond to cyclic groups. For the other degrees, we enumerate all possible labeling schemes, each one corresponding to a family of planar Cayley graphs. For each class, we give the corresponding presentation and type vector, as long as a representative of this class. Some classes may have the same representative in terms of groups or unlabeled graphs, but the corresponding Cayley graphs are distinct. The following table displays the number of such classes :

Degree	Classes	Degree	Classes
1	0	4	26
2	1	5	64
3	8	6	253

[3]The corresponding graphs and labeling schemes appear in an additional file which may be found at `http://www.labri.fr/Perso/~renault/research/pages.ps.gz` ($\approx$ 438 ko.).