The vertex-transitive TLF-planar graphs

David Renault

LaBRI – Université Bordeaux I
351, cours de la Libération – 33400 Talence, FRANCE

Abstract

We consider the class of the topologically locally finite (in short TLF) planar vertex-transitive graphs, a class containing in particular all the one-ended planar Cayley graphs and the normal transitive tilings. We characterize these graphs with a finite local representation and a special kind of finite state automaton named *labeling scheme*. As a result, we are able to enumerate and describe all TLF-planar vertex-transitive graphs of any given degree. Also, we are able decide to whether any TLF-planar transitive graph is Cayley or not.

Key words: vertex-transitive, Cayley graph, planar graph, automorphisms group, topologically locally finite,

Introduction

Vertex-transitive graphs – or transitive graphs in short – are graphs whose group of automorphisms acts transitively on their sets of vertices. These graphs possess a regular structure, being structurally the same from any vertex. When such a graph is planar, this regular structure confers symmetry properties to the embedding of the graph : the action of automorphisms on the graph can locally be represented as the action of an isometry of the geometry it is embedded in.

The class $\mathcal{T}$ of the topologically locally finite (in short TLF) planar transitive graphs is a subclass of the class of the transitive planar graphs. These graphs possess a planar embedding such that the set of vertices in this embedding is a locally finite subset of the plane. For example, $\mathcal{T}$ contains the planar Cayley

Email address: `renault@labri.fr` (David Renault).

graphs of the discrete groups of isometries of the plane, be it hyperbolic or Euclidean. We find in $\mathcal{T}$ both tree-like graphs of finite treewidth and one-ended graphs such as the Euclidean grid.

The TLF-planar graphs are related to several fields: first, such graphs represent adapted models of computation for parallel algorithms such as cellular automata [8, 4], and provide examples of structures for interconnection networks [6]. The class $\mathcal{T}$ is also connected to the vertex-transitive tilings of the plane whose set of vertices is topologically locally finite [5]. Second, while many problems in combinatorial group theory are undecidable, the planar vertex-transitive graphs possess structural as well as geometrical properties that allow for a more specific approach. For example, the unique embedding property in the case of the 3-connected planar graphs [14] constrains the structure of the automorphisms of the graph. Finally, putting aside the finite graphs in $\mathcal{T}$, the infinite graphs in $\mathcal{T}$ may factor into finite transitive graphs embedded into compact manifolds, as it is the case for Euclidean graphs [15].

The characterization of the finite planar transitive graphs is a result of Fleischner and Imrich [3]. These graphs turn out to be the complexes associated to the uniform convex polyhedra. The finite non-planar case happens to be much more difficult, and has been thoroughly studied up to 26 vertices by McKay *et al.* [11, 13], but few general results are known. Cayley graphs are natural examples of vertex-transitive graphs, and McKay also studied the problem of the determination of those graphs that were transitive but not Cayley graphs [9, 10]. The problem of enumerating the normal Cayley graphs, has been solved by Chaboud [2] and then extended to the TLF-planar Cayley graphs [12].

In this paper, we give an exhaustive description of the class $\mathcal{T}$ of the TLF-planar vertex-transitive graphs. Our description of the class includes both finite and infinite graphs. This extends the Cayley case [12] in several ways. First, the Cayley case is mainly dedicated to the description of groups which happen to have a planar Cayley graph. This article focuses on the graphs themselves, by describing all their possible groups of automorphisms, and as a consequence all their possible embeddings. Second, we can highlight properties of the transitive graphs that do not hold when we restrict ourselves to Cayley graphs. For example, TLF-planar Cayley graphs can always be represented by the Cayley graph of a discrete group of isometries of the plane. There exist transitive graphs for which this is impossible, independently of the embedding. Finally, $\mathcal{T}$ contains a strictly larger class of graphs than the Cayley case. Such a simple graph as the complex associated to the dodecahedron is an example of transitive but non Cayley planar graphs. More precisely, there exist infinite families of graphs having this property.

In this article, we refine the description of the groups of automorphisms of the graphs and the geometrical properties of their possible planar embeddings given in [12]. We represent these graphs by their geometrical invariants in a structure called a *labeling scheme*, as long as a special kind of finite state automaton called a *border automaton.* We show that there exists a bijection between this representation and the class of the TLF-planar transitive graphs. Our main result (page 22) is:

Theorem 12 (Enumeration) *Given a number $d \geq 2$, it is possible to enumerate all the TLF-planar transitive graphs having internal degree d, along with their labeling schemes.*

Each vertex-transitive graph belonging to the class $\mathcal{T}$ is effectively computable (*i.e.* there exists an algorithm able to build every finite ball of the graph). Associated to our results on Cayley graphs, this allows us to determine which of these graphs are Cayley, and more precisely:

Corollary 13 (Cayley checking) *If Γ is a TLF transitive graph, then it is decidable whether Γ is the Cayley graph of a group or not, and obtain an enumeration and a description of the groups having Γ as a Cayley graph.*

Finally, thanks to the characterization of the embeddings of the graphs in $\mathcal{T}$, it is possible to compute their connectivity and approximate their growth rate, which can be either linear, quadratic or exponential, depending on their local geometrical properties.

1 TLF-planar transitive graphs

A *graph* Γ consists of a pair (V, E), V being a countable set of *vertices* and a set of *edges*, where E is a subset of the pair of elements of V. Each edge corresponds to a pair of vertices (v_1, v_2) called its extremities. An edge with the same extremities (v, v) is called a loop. The graphs that we consider are loopless. An edge (u, v) is said to be *incident* to the vertices u and v. A *labeling* of the graph is an application from the set of edges into a finite set L of labels or colors. The *degree* of a vertex is the number of edges incident to this vertex. A *path* of Γ is a sequence of vertices (v_n) in Γ such that for all n, there exists an edge between v_n and v_{n+1}. A *cycle* is a finite path whose initial and terminal vertices are the same. A *simple* cycle is a cycle where no vertex appears twice.

Remark 1 (*Considerations on the construction of subgraphs)*
We occasionally build subgraphs of Γ by considering a certain subset of vertices and edges (V', E') where $V' \subset V$ and $E' \subset E$, or equivalently by removing from the graph a subset of its vertices and edges. Then, we can consider the remaining set of vertices and edges as a subspace of the graph seen as a metric space, and the connected components of this subspace. These components may

not be graphs themselves, since some edges will not have vertices as their extremities. We can resolve this problem and consider these components as new graphs by adding new vertices to the extremities of these edges.

A graph Γ 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 *connected component* is an equivalence class of vertices for the relation "to be connected". Notice that both definitions are coherent whether Γ is seen as a graph or as a metric space. A *n-separation* is a set of n vertices whose removal separates the graph in two or more connected components not reduced to a single edge. 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 *regular* when all its vertices have the same degree d. The graphs we will be dealing with are connected and regular.

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 graph. When both graphs are labeled, we impose that the morphisms also preserve the labels of the edges. A graph is said to be *vertex-transitive* – or *transitive* in short – 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, then it is vertex-transitive.

Remark 2 (*About the lower degree transitive graphs*)
There exists only one non-trivial connected transitive graph of degree 1*, which corresponds to* K_2*, the graph reduced to a single edge. Transitive graphs of degree* 2 *correspond to cyclic graphs* C_n *where* n *may be infinite. These graphs possess exactly one labeling when* n *is odd and two labelings when* n *is even, these labelings corresponding to the planar Cayley graphs of degree* 2 *associated to the dihedral groups and cyclic groups. In the following, we shall only be interested in connected transitive graphs of degree* $d \geq 3$.

A graph Γ is said to be *planar* if it can be embedded in the plane, such that no two edges meet in a point other than a common end. By the plane, we mean a simply connected Riemannian surface, homogeneous and isotropic. For our embeddings, we will only consider the three usual geometries : the sphere, the Euclidean and the hyperbolic plane. Our embeddings will be considered *tame*, meaning that all edges are $\mathcal{C}^1$ images of $[0; 1]$. Such an embedding is said to be *topologically locally finite* – in short TLF-planar – if its vertices have no accumulation point in the plane. Equivalently, every compact subset of the plane intersects a finite number of vertices of the embedding. Symmetrically, an embedding is said to be TLF in terms of edges if and only if every compact subset of the plane intersects a finite number of edges of the embedding. The following theorem asserts that a TLF-planar graph always possess such an embedding:

Theorem 1 ([12]) *If the graph Γ is TLF-planar, there exists a tame embedding of the same graph that is TLF in terms of vertices but also of edges.*

We will always suppose that the TLF-planar graphs are embedded in the plane such that their embedding follows Theorem 1. Given a specific embedding of a TLF-planar graph Γ, 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 incident to finitely many vertices of the graph, otherwise it is said to be *infinite*. For TLF-planar graphs, infinite faces are necessarily topologically unbounded in the plane. The *border* of the face $\mathcal{F}$, noted $\partial\mathcal{F}$, is its boundary in topological terms. A face $\mathcal{F}$ is said to be *incident* to a vertex or an edge of the graph if and only if this vertex of edge intersects with $\partial\mathcal{F}$.

In such an embedding, every edge incident to a face is entirely included into the border of this face. Then every edge is incident to exactly two faces of Γ, which it separates. Considering the previous definitions of the faces, the transitivity property of Γ stands out with the following lemma taken from [12]:

Lemma 2 (Intersection of faces) *Let Γ be a vertex-transitive TLF-planar graph. Given two distinct faces of Γ, the intersection of their border, when non-empty, is either a vertex or an edge of Γ.*

Corollary 3 (Preservation of finite faces) *The automorphisms of Γ map the border of every finite face onto the border of another finite face.*

Remark 3 (*On the choice of the embedding)*
The previous statements hold for a particular embedding of Γ that is locally finite in terms of edges and vertices. This embedding may not be unique. For example, if Γ is finite, there exists an embedding of Γ in the sphere, where all faces are topologically bounded. If we select a point inside a face and send this point to the infinity, we obtain another embedding of the graph on a non-compact surface homeomorphic to the Euclidean plane. With the previous definitions, the faces of both embeddings are all finite, and the validity of Corollary 3 is the same for both embeddings.

The *size* of a face of an embedding of Γ corresponds to the number (possibly infinite) of vertices it is incident to. The *type vector* of a vertex of Γ is the sequence of sizes of the faces appearing consecutively around this vertex. It is defined up to rotation and symmetry of the graph. If Γ is transitive and Corollary 3 holds, the type vector is independent of the choice of the vertex, up to permutation of its elements. For example, the type vector of the Euclidean infinite grid is $[4; 4; 4; 4]$ and the type vector of a cyclic graph C_n with n vertices is $[n; n]$.

For a given graph Γ, $\mathrm{Aut}(\Gamma)$ denotes its group of automorphisms, and $\mathrm{Trans}(\Gamma)$ stands for the set of subgroups of $\mathrm{Aut}(\Gamma)$ acting transitively on the set of vertices of Γ. Let G belong to Trans(Γ). G acts on the set of edges of Γ and the set of orbits of edges is finite. Thus a *class* or a *color* of edges is defined as an orbit under the action of G, the set of colors being called E_G. In the same manner, we define classes or colors of finite faces, corresponding to the finite set F_G. This coloring defines a partition of the set of finite faces of the embedding. Infinite faces have a special status since these faces may not be stable by automorphism. Therefore, we request by convention that all of them correspond to a special color in F.

In the following, Γ will be a TLF-planar, connected vertex-transitive graph of finite degree $d \geq 3$. Thus we will speak of vertices, edges and faces of Γ, as defined above. G is a group belonging to Trans(Γ). We will always suppose that the embedding of Γ follows Theorem 1 and that the automorphisms preserve the borders of the finite faces. In the section 2, we analyze these local invariants, in order to obtain a characterization of the graph by its local geometrical properties in section 3. The last section presents some applications of these characterizations.

2 Local geometrical invariants

2.1 Infinite faces and connectivity

Let us give some intuition on the general structure of the graphs in this class. We prove that the finite vertex-transitive planar graphs of degree ≥ 3 are all 3-connected graphs, and in the infinite case, the connectivity of the graphs depends only on the number of infinite faces appearing around each vertex:

Lemma 4 (Connectivity and infinite faces) *If Γ is a TLF-planar transitive graph of degree ≥ 3, let n be the number of infinite faces appearing around a given vertex of Γ. Then, depending on the value of n:*

- $n \geq 2 \Leftrightarrow \Gamma$ *is 1-separable.*
- $n = 1 \Leftrightarrow \Gamma$ *is 2-connected and 2-separable;*
- $n = 0 \Leftrightarrow \Gamma$ *is 3-connected;*

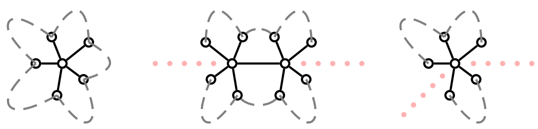

$n = 0$ $n = 1$ $n \geq 2$

Proof :

- ($n \geq 2 \Leftrightarrow \Gamma$ *is* 1-*separable*)
  Given a vertex v of Γ, we consider the set of faces incident to that vertex, and the union of the borders of those faces that are finite. If $n \geq 2$, the union of v and two infinite faces incident to v separates the graph into at least two non-trivial components. Then every vertex of the graph is a cut-

vertex and the graph is 1-separable. On the other hand, if Γ is 1-separable, every vertex must meet at least 2 infinite faces.

- ($n = 1 \Rightarrow \Gamma$ *is* 2*-connected and* 2*-separable*)
  If $n = 1$, then consider an edge of Γ that does not belong to the border of an infinite face. There must exist one, otherwise Γ being of degree at least 3, that would contradict the fact that $n = 1$. The extremities of this edge both meet an infinite face, and these faces are distinct. The removal of these extremities separates Γ, therefore Γ is 2-separable. It is 2-connected because $n < 2$.
- (Γ *is* 2*-connected and* 2*-separable* $\Rightarrow n = 1$)
  Suppose now that Γ is 2-separable and 2-connected, and consider $\{s, t\}$ a 2-separation of Γ. Let Λ be a subgraph separated by $\{s, t\}$. Suppose that we remove Λ from the embedding. The remaining TLF-planar graph possesses a face $\mathcal{F}$, inside which Λ was embedded. Moreover, s and t both belong to the border of $\mathcal{F}$. If $\mathcal{F}$ is finite, embedding Λ inside $\mathcal{F}$ separates the face into at least two subfaces meeting at s and t, therefore contradicting Lemma 2. Therefore $\mathcal{F}$ is infinite. When embedding Λ inside $\mathcal{F}$, there will remain an infinite face in the embedding of Γ. Therefore $n \geq 1$ and since Γ is 2-separable, $n = 1$.

$\square$

If Γ is 1-separable, then every vertex is a cut-vertex. If we cut the graph along its cut-vertices, the remaining components are *2-connected components*. Since the graph is vertex-transitive, the set of components incident to a vertex is independent of the vertex, and finite, because the degree of the graph is finite. These components are TLF-planar graphs, but not necessarily vertex-transitive themselves. They may be finite or infinite. They may be reduced to a single edge. If Γ is at least 2-connected, it is composed of a unique 2-connected component equal to Γ.

2.2 A simple invariant

Consider more closely the implications of Corollary 3. The group G of automorphisms of the graph acts on the set of the finite faces of the embedding. Let us focus on the invariants under this action. The classes or colors of the edges and faces are simple examples of geometrical invariants.

For the sake of clarity, we always mark classes (or colors) of edges with gothic letters $\mathsf{E}_G = \{\mathfrak{a}, \mathfrak{b}, \mathfrak{c} \dots\}$ and classes (or colors) of faces with greek letters $\mathsf{F}_G = \{\alpha, \beta, \gamma \dots\}$. In the remaining of the article, we will suppose that the group G is fixed and therefore drop the letter G.

Lemma 5 (Face separation) *Consider an edge e of Γ belonging to the class $\mathfrak{e} \in \mathsf{E}$ of edges under the action of G. Then the classes of faces separated by e are the same independently of the representative $e \in \mathfrak{e}$.*

Proof : Suppose e is mapped by automorphism on $f \in \mathfrak{e}$. As a result from Theorem 3, the finite faces incident to the edge e are mapped by automorphism onto the finite faces incident to f. And the automorphism is invertible, therefore the finite faces incident to both edges are in bijection. In turn, there is an equal number of infinite faces incident to both edges. This concludes the proof. □

According to Lemma 5, for any class of edges $\mathfrak{e} \in \mathsf{E}$, it is possible to define the *separator* of $\mathfrak{e}$, namely $\mathsf{sep}(\mathfrak{e})$, as the pair of classes of faces separated by $\mathfrak{e}$. This separator is a geometrical invariant under the action of G.

2.3 Edge and Face vectors

Given a vertex $v \in \Gamma$, consider the finite subgraph Λ of Γ composed of all edges incident to v and its planar embedding induced by the embedding of Γ. Select a particular edge e incident to v. An *edge vector* ξ of Γ around v is the vector whose elements describe the classes of edges appearing around v in Λ in the positive direction, starting from e. Similarly, a *face vector* ϕ around v is the vector whose elements describe the classes of faces appearing around v in Λ, starting from the face next to e in the positive direction. As a convention, we always choose the pair (ξ, ϕ) composed of an edge vector and a face vector, to be *locked* as follows: the edge ξ_i separates the faces ϕ_i and ϕ_{i-1}. We decompose these vectors into blocks separated by infinite faces :

$$\xi : [\overbrace{\xi_{k_1^s}, \dots, \xi_{k_1^e-1}, \xi_{k_1^e}}^{\text{Block 1}}, \overbrace{\xi_{k_2^s}, \dots, \xi_{k_2^e-1}, \xi_{k_2^e}}^{\text{Block 2}}, \dots, \overbrace{\xi_{k_t^s}, \dots, \xi_{k_t^e-1}, \xi_{k_t^e}}^{\text{Block t}}]$$
$$\phi : [\underbrace{\phi_{k_1^s}, \dots, \phi_{k_1^e-1}}_{\text{Block 1}}, \infty, \underbrace{\phi_{k_2^s}, \dots, \phi_{k_2^e-1}}_{\text{Block 2}}, \infty, \dots, \underbrace{\phi_{k_t^s}, \dots, \phi_{k_t^e-1}}_{\text{Block t}}, \infty]$$

Here ξ_k and ϕ_k represent respectively the k-th elements of the vectors ξ and ϕ. The i-th block starts at index k_i^s and ends at index k_i^e. If the graph does not contain any infinite face, then the decomposition contains a single block.

Consider the set of all possible edge and face vectors in the embedding of Γ. Define the following operations on this set:

- Rotation and Symmetry: These operations correspond to the usual isometries of the plane acting on the face and edge vectors.
- Rearrangement: The operations consists in rearranging the blocks while preserving the fact that these blocks are separated by infinite faces. More

graphically, given a permutation σ of the blocks :

$$\xi : [\overbrace{\xi_{k^s_{\sigma(1)}},\ldots,\xi_{k^e_{\sigma(1)}}}^{\text{Block } \sigma(1)}, \overbrace{\xi_{k^s_{\sigma(2)}},\ldots,\xi_{k^e_{\sigma(2)}}}^{\text{Block } \sigma(2)}, \ldots, \overbrace{\xi_{k^s_{\sigma(t)}},\ldots,\xi_{k^e_{\sigma(t)}}}^{\text{Block } \sigma(t)}]$$
$$\phi : [\underbrace{\phi_{k^s_{\sigma(1)}},\ldots,\phi_{k^e_{\sigma(1)}-1}}_{\text{Block } \sigma(1)}, \infty, \underbrace{\phi_{k^s_{\sigma(2)}},\ldots,\phi_{k^e_{\sigma(1)}-1}}_{\text{Block } \sigma(2)}, \infty, \ldots, \underbrace{\phi_{k^s_{\sigma(t)}},\ldots,\phi_{k^e_{\sigma(1)}-1}}_{\text{Block } \sigma(t)}, \infty]$$

- Twist: The twist operation describes a symmetry applied on a single 2-connected component around a vertex. For example, the twist applied onto the first component corresponds to the following transformation:

$$\xi : [\overbrace{\xi_{k_1^e},\ldots,\xi_{k_1^s+1},\xi_{k_1^s}}^{\text{Block 1 reversed}}, \overbrace{\xi_{k_2^s},\ldots,\xi_{k_2^e-1},\xi_{k_2^e}}^{\text{Block 2}}, \ldots, \overbrace{\xi_{k_t^s},\ldots,\xi_{k_t^e-1},\xi_{k_t^e}}^{\text{Block t}}]$$
$$\phi : [\underbrace{\phi_{k_1^e-1},\ldots,\phi_{k_1^s}}_{\text{Block 1 reversed}}, \infty, \underbrace{\phi_{k_2^s},\ldots,\phi_{k_2^e-1}}_{\text{Block 2}}, \infty, \ldots, \underbrace{\phi_{k_t^s},\ldots,\phi_{k_t^e-1}}_{\text{Block t}}, \infty]$$

Two pairs of locked vectors (ξ_1, ϕ_i) and (ξ_2, ϕ_2) around v are said to be *isomorphic* if and only if it is possible to transform the first into the second by a sequence of rotations, symmetries, rearrangements and twists. The following lemma states that these operations describe all the possible edge and face vectors in the embedding:

Lemma 6 (Edge and Face vectors) *The edge vector and the face vector of Γ is independent of the choice of the embedding of Γ and of the vertex around which it is chosen, up to isomorphism.*

Proof : This is a direct consequence of Corollary 3. Since finite faces are mapped onto finite faces by the automorphisms of Γ, then the 2-connected components of Γ are mapped onto 2-connected components. Therefore, the only operations that we may apply on the set of edge and face vectors of Γ are rotations and symmetries in the case of 2-connected graphs, and rearrangement of the 2-connected components for 1-separable graphs. These operations are exactly those described by the twists and rearrangements. □

Given an edge vector ξ and a face vector ϕ of Γ that are locked together, it is therefore possible to determine all possible vectors in the class of isomorphism. Therefore we will only consider *the* edge and face vectors of Γ by choosing a representative in this class.

Example 1. Suppose that the graph Γ possesses the set of colors defined by $\mathsf{E} = \{\mathfrak{b}, \mathfrak{r}, \mathfrak{g}\}$ for its edges and $\mathsf{F} = \{\alpha, \beta, \gamma\}$ for its faces. An example of such a graph appears on Figure 9 page 19. We represent the edge and face vectors (ξ, ϕ) of this graph by the following picture :

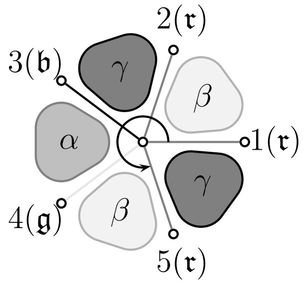

Edge vector ξ : $[\mathfrak{r}, \mathfrak{r}, \mathfrak{b}, \mathfrak{g}, \mathfrak{r}]$

Face vector ϕ : $[\beta, \gamma, \alpha, \beta, \gamma]$

Suppose for our example that the faces colored by β are infinite. Therefore, in the aforementioned decomposition, there are two blocks, one containing the edges numbered $\{1; 5\}$ and the other containing the edges numbered $\{2; 3; 4\}$. Under these hypotheses, we can operate the following transformations onto the pair (ξ, ϕ):

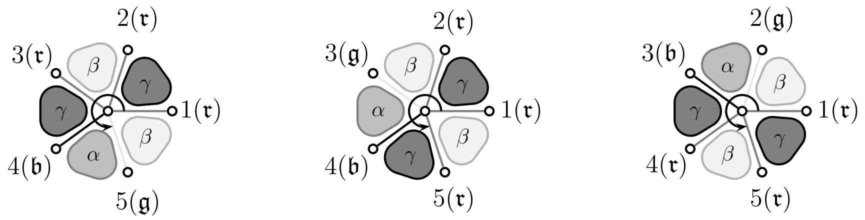

Rotation | Symmetry | Twist

With only two blocks, a rearrangement is the same as a rotation. Moreover, since one of the blocks is stable by symmetry, a twist of this component leaves the pair unchanged. A twist of the other component is the same as a symmetry of the pair.

Example 2. Edge vectors and face vectors in general do not determine a vertex-transitive graph in a unique way. It is quite possible to obtain non-isomorphic graphs possessing the same edge and face vector. For example, consider the graphs on Figure 1. Both graphs correspond to the planar tiling of the hyperbolic plane with decagons; both are vertex-transitive and face-transitive graphs, and they have the same edge vectors. Nevertheless, the borders of the faces differ: for the graph on the left, it corresponds to $(\mathfrak{rgrbrbgbrb})$ and $(\mathfrak{rgbrb})^2$ on the right, where $\mathfrak{r}$, $\mathfrak{g}$ and $\mathfrak{b}$ respectively stand for the three different classes of edges.

For an accurate description of the graph, some complementary informations are therefore needed. Following the intuitions in the case of TLF-planar Cayley graphs [12], these informations are likely to come from local invariants linked to the classes of edges of Γ.

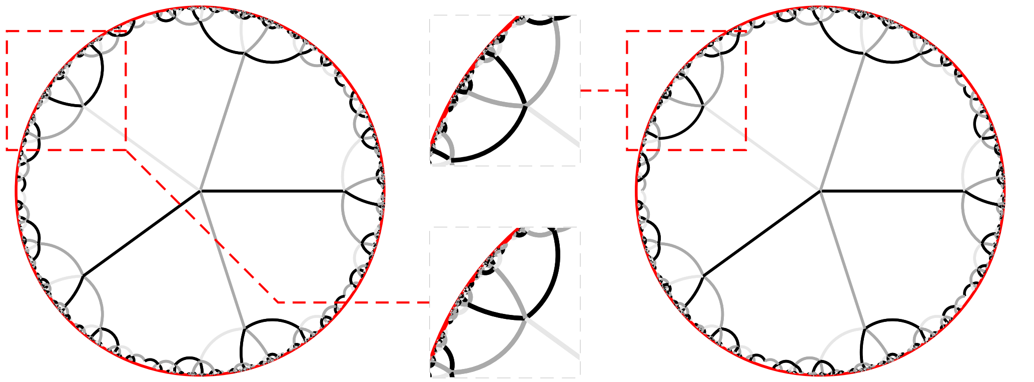

Fig. 1. Non-isomorphic transitive graphs with the same edge and type vectors. The extremities of the lighter edge are non-isomorphic (see detail), which, considering the fact that both graphs are 3-connected, implies the non-isomorphism.

2.4 Edge neighborhoods

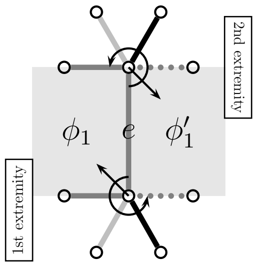

Let e be an edge of Γ, and T the finite tree composed of the edge e and all edges incident to both extremities of e. An *edge neighborhood* for e is a structure composed of the two edge vectors representing the classes of edges appearing around each extremity of e in the same direction, where the first element of each vector corresponds to e, and two face vectors representing the classes of faces appearing around each extremity, such that each corresponding edge and face vectors be locked together. Technically, we represent an edge neighborhood η_e by the following structure:

$$\eta_e = \left\{ \overbrace{\overbrace{[\xi_1, \ldots, \xi_d]}^{\text{Edge vector}}, \overbrace{[\phi_1, \ldots, \phi_d]}^{\text{Face vector}}}^{\text{1st extremity}}, \overbrace{\overbrace{[\xi'_1, \ldots, \xi'_d]}^{\text{Edge vector}}, \overbrace{[\phi'_1, \ldots \phi'_d]}^{\text{Face vector}}}^{\text{2nd extremity}} \right\}$$

$$\text{for } \{\xi_i, \xi'_i\} \subset \mathsf{E}, \quad \text{and } \{\phi_i, \phi'_i\} \subset \mathsf{F}$$

where $\xi_1 = \xi'_1 = e$, $\phi_1 = \phi'_d$ and $\phi_d = \phi'_1$. The vectors $[\xi_1, \ldots, \xi_d]$ and $[\phi_1, \ldots, \phi_d]$ are respectively the edge and face vectors of the first extremity, and the vectors $[\xi'_1, \ldots, \xi'_d]$ and $[\phi'_1, \ldots \phi'_d]$ correspond to the second extremity.

The *color* of an edge neighborhood corresponds to the class of edges of e. The *separator* of η_e, noted $\mathsf{sep}(\eta_e)$, correspond to the pair of classes of faces separated by e, here (ϕ_1, ϕ'_1). An edge neighborhood η_e colored by $\mathfrak{e}$ is said to be *coherent* with a pair of vectors (ξ, ϕ) if and only if both edge vectors and face vectors at each extremity of e are isomorphic to (ξ, ϕ).

Consider the set of edge neighborhoods of the same color of Γ. As for edge and face vectors, it is possible to define operations on this set:

- Inversion and Symmetry: Inversion is the operation exchanging both extremities of the edge neighborhood, and corresponds to an exchange of the edge vectors and of the faces. Symmetry corresponds to the operation of symmetry (defined on the edge and face vectors) applied to each extremity of the edge neighborhood, while preserving the central edge.
- Twist or Rearrangement of an extremity: Let (ξ, ϕ) be the edge and face vectors associated to one extremity of an edge neighborhood η_e. Any twist or rearrangement of (ξ, ϕ) that stabilizes the central edge leaves the faces separated by η_e unchanged and extends naturally on the edge neighborhood.

Two edge neighborhoods η_1 and η_2 of the same color $\mathfrak{e} \in \mathsf{E}$ are said to be *isomorphic* if and only if it is possible to transform η_1 into η_2 by a sequence of inversions, symmetries, twists and rearrangements of any extremity. As was the case for the edge and face vectors, these operations describe all the possible edge neighborhoods in the embedding. We therefore select a single representative for each class of edges in Γ.

Lemma 7 (Edge neighborhood) *The edge neighborhood colored by $\mathfrak{e} \in \mathsf{E}$ of Γ is independent of the choice of the embedding of Γ and of the edge it is referring to, up to isomorphism.*

Proof : The separator of a class of edge is independent of the edge (Lemma 5). Since finite faces are mapped onto finite faces by automorphisms of Γ, the 2-connected components attached to the extremity of a class of edge are mapped onto 2-connected components by automorphism. Therefore, the automorphisms mapping an edge onto another edge correspond to a rearrangement of composition of either natural transformations of the plane preserving this edge but exchanging its extremities (inversion and symmetry) or automorphisms leaving the edge and its extremities stable (rearrangements and twists). $\square$

Let (ξ, ϕ) be an edge and face vector. Two edges $e, f \in \xi$ labeled by the color $\mathfrak{e} \in \mathsf{E}$ are said to be equivalent, namely $e \sim f$, if and only if there exists an isomorphism of (ξ, ϕ) mapping e onto f. Consider the set of edges in ξ colored by $\mathfrak{e}$. Then $\mathsf{eq}_{\xi,\phi}(\mathfrak{e})$ is the number of classes of equivalence inside this set with regard to $\sim$. The uniqueness of the edge neighborhood implies that $\mathsf{eq}_{\xi,\phi}(\mathfrak{e})$ is at least one and at most two. As a matter of fact, each class of equivalence must correspond to an extremity of the edge neighborhood colored by $\mathfrak{e}$, and this neighborhood only has two extremities.

Example 3. Let us define edge neighborhoods coherent with the pair of edge and face vectors described in example 1. We represent the edge neighborhood associated to the edge colored by $\mathfrak{r}$ by the following picture:

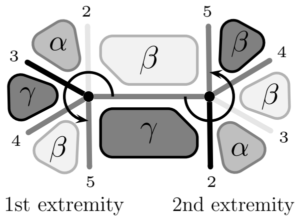

$$\text{1st extremity}\ : [\mathfrak{r}, \mathfrak{g}, \mathfrak{b}, \mathfrak{r}, \mathfrak{r}]\ ,\ [\beta, \alpha, \gamma, \beta, \gamma]$$
$$\text{2nd extremity}\ : [\mathfrak{r}, \mathfrak{b}, \mathfrak{g}, \mathfrak{r}, \mathfrak{r}]\ ,\ [\gamma, \alpha, \beta, \gamma, \beta]$$

The fact that the faces colored by β are infinite allows for the transformation of this labeling scheme by twists of either extremity.

Example 4. While it could seem that the face vector is superfluous, *i.e.* that the graph could be described simply by its edge neighborhoods without any face vector, consider the particular case of graphs that are both vertex-transitive and edge-transitive, as in the Figure 2. The $[5; 5; 5; 5]$ and $[3; 7; 3; 7]$ graphs – denoted by their type vector – both belong to this class. Both possess the same edge neighborhoods and edge vectors. Yet these two graphs are obviously not isomorphic and their group of automorphisms are distinct, because the first is face-transitive, and the second possesses two different classes of faces.

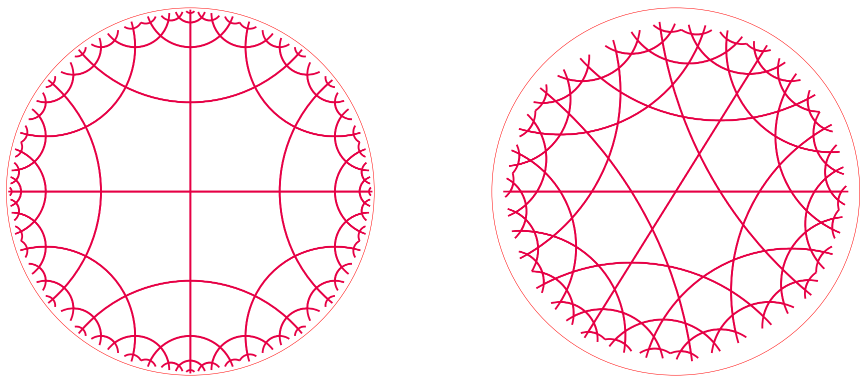

Fig. 2. Non-isomorphic transitive graphs with the same edge neighborhoods, assuming that we remove the face vectors from the edge neighborhoods.

Whitney's Theorem [14] states that the finite planar 3-connected graphs have a unique embedding property *i.e.* their dual is uniquely defined. This property was extended to infinite graphs by Imrich [7]. As a matter of face, when Γ is 3-connected (while remaining transitive and TLF-planar), all faces of the graph are finite, and the graph is obviously composed of a unique 2-connected component. Therefore the classes of isomorphisms of the geometrical invariants described in this section contain neither twists nor rearrangements. This is coherent with the unique embedding property. Notice that this property holds when Γ is at least 2-connected, in the case of transitive TLF-planar graphs. On the other hand, when Γ is 1-separable, these invariants provide an accurate description of the possible embeddings of the graph in the plane.

3 Labeling schemes

Our purpose in this section is to consider the geometrical invariants of Γ and to prove that they are sufficient to give an exact description of the graph. The resulting description of the graph is called a labeling scheme, and extends the notion of labeling scheme for Cayley graphs, detailed in [2, 12].

3.1 Border automaton

Let E and F be two non-intersecting finite sets of colors. A *labeling scheme* of degree d is a 3-tuple (ξ, ϕ, η) possessing the following properties:

(i) $\xi \in \mathsf{E}^d$ is an edge vector and $\phi \in \mathsf{F}^d$ is a face vector;
(ii) for each color $\mathfrak{e} \in \mathsf{E}$, $\mathsf{eq}_{\xi,\phi}(\mathfrak{e})$ does not exceed two;
(iii) for each color $\mathfrak{e}$ in ξ, there exists a unique edge neighborhood $\eta_\mathfrak{e} \in \eta$ of the same color; all edge neighborhoods in η must be coherent with (ξ, ϕ) and if $\mathsf{eq}_{\xi,\phi}(\mathfrak{e}) = 2$, then each class of equivalence must appear on an extremity of $\eta_\mathfrak{e}$.

Given a graph Γ, then ξ and ϕ stand for Γ's edge and face vectors, locked together. The set η stands for the set of edge neighborhoods of Γ. In the following, $\eta_\mathfrak{e}$ stands for the edge neighborhood in η colored by $\mathfrak{e}$. The isomorphisms of labeling schemes are defined as the isomorphisms of the elements of the scheme. The results in the previous sections ensure that for any pair (Γ, G), there exists a labeling scheme (ξ, ϕ, η) corresponding to the coloring of the vertices and faces associated to the group G of the graph Γ, up to isomorphism. Notice that the condition (ii), along with Lemma 7 constrain the number of possible labeling schemes.

Consider a labeling scheme (ξ, ϕ, η). Let e be an element of ξ and $\eta_e \in \eta$ be the edge neighborhood labeled by the color of e. The operation of *gluing* η_e *with* e is possible if and only if there exists an edge neighborhood κ isomorphic to η_e such that one extremity of κ be exactly equal to the pair (ξ, ϕ), with $e \in \xi$ as the central edge of the neighborhood κ.

Lemma 8 (Edge reconstruction) *Let (ξ, ϕ, η) be a labeling scheme, $e \in \xi$ and $\eta_e \in \eta$ colored by e, there exists a unique way to glue η_e with e, up to isomorphism of the other extremity of η_e.*

Proof : Suppose that there exists two different ways to glue η_e onto e. If each extremity of η_e may be glued onto e, this means that all edges colored by e are equivalent with regard to $\sim$. If only one extremity of η_e may be glued onto e, then there exists two classes of equivalence with regard to $\sim$ for the color of e, each of them at one extremity of η_e. In any case, since the number of classes of equivalence for the color of e does not exceed two, that leads to a unique possibility to glue η_e onto e, up to isomorphism. The existence follows from the fact that every class of equivalence of every color appears on an extremity of the associated edge neighborhood. $\square$

The gluing of edge neighborhoods allows the reconstruction of the graph. Let us start from an initial graph Λ composed of all edges incident to a central vertex v. Suppose that these edges are labeled accordingly to the edge vector ξ. Obviously, this planar graph does not include faces for the moment, nevertheless we are expecting to build a face ϕ_i between the edges ξ_i and ξ_{i+1}. By gluing the appropriate edge neighborhood $\eta_{\mathfrak{e}}$ onto the edge ξ_i, we create $d-1$ new edges incident to the other extremity of ξ_i labeled such that the obtained edge neighborhood is isomorphic to $\eta_{\mathfrak{e}}$. The other extremities of these new edges correspond to new vertices of the graph. We will see later how it is possible to close the border of the faces.

Example 5. Consider the labeling scheme defined in Figure 3, where the set of colors are $\mathsf{E} = \{\mathfrak{b}, \mathfrak{r}, \mathfrak{g}\}$ and $\mathsf{F} = \{\alpha, \beta, \gamma\}$. It is based on the examples 1 and 3. We associate to each color in E a unique edge neighborhood. The face β is supposed to be infinite, thus allowing by a twist exchanging the faces colored by β that $\mathsf{eq}_{\xi,\psi}(\mathfrak{r}) = 2$.

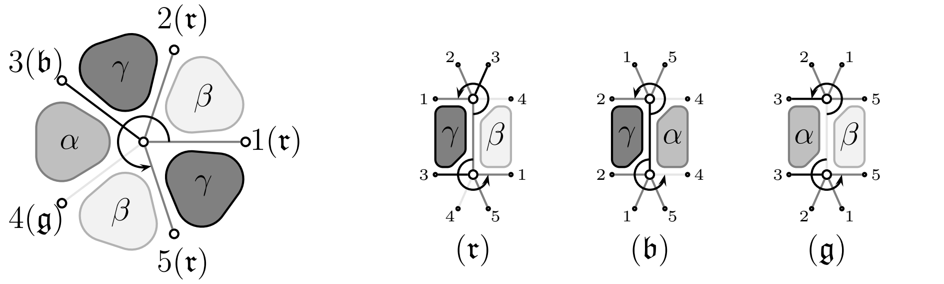

(a) Edge and Face vectors (b) Edge neighborhoods

Fig. 3. A labeling scheme of degree 5 with its components.

Let us try to glue the edge neighborhoods consecutively, while following the edges constituting the border of the orange face (cf. Figure 4). We begin by gluing the golden edge neighborhood $\eta_{\mathfrak{g}}$ onto the unique golden $\mathfrak{g}$ edge belonging to the edge vector. Having glued two red edge neighborhoods $\eta_{\mathfrak{r}}$, we can

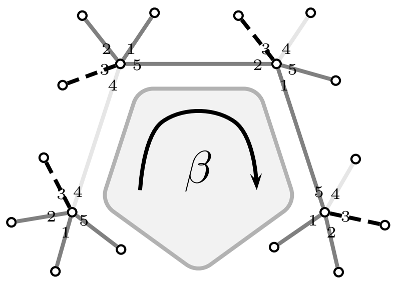

Fig. 4. Successive gluings of the edge neighborhoods along the border of the face colored by β.

ourselves continue the process indefinitely, by gluing red edge neighborhoods along the border of the face. If this labeling scheme corresponds to an existing graph Γ, then this process describes the border of a face colored by β in Γ.

Lemma 9 (Description of the faces) *Consider a graph Γ and its labeling scheme (ξ, ϕ, η). It is possible to build a finite state automaton and an operation onto that automaton allowing to construct the borders of the faces of Γ.*

Proof : Consider an automaton built over the alphabet A. A language $L \subset A^\star$ acts naturally over the set of states of the automaton. Therefore, we consider the following partition of the states of the automaton: two states u and v are equivalent if and only if there exists $l \in L$ such that it is possible to start from the state u, and reach the state v by reading the word l on the automaton. Now we will build an automaton by describing its set of states and the language acting on these states.

A *configuration* of a labeling scheme (ξ, ϕ, η) is a pair $(\mathcal{C}, b)$ where $\mathcal{C}$ is a class of equivalence for $\sim$ in (ξ, ϕ) and $b \in \{+, -\}$ is a direction of rotation. A configuration represents a block of edges (separated by infinite faces) attached to a vertex, an edge inside this component and a direction expressing in which way the block is embedded in the plane. Two configurations are said to be *equivalent* if (i) they correspond to two blocks, (ii) it is possible to map the first block onto the second by an isomorphism of edge and face vectors and (iii) this map sends the edge of the first configuration onto the edge of the second configuration, and maps the direction of rotation accordingly.

For a given labeling scheme, the number of configurations is finite, bounded by $2d$. These configurations define the states of our automaton. Let us define the following relations between configurations:

- The next relation describes whether a configuration comes next to another one in the edge vector, given a direction of rotation: $(\mathcal{C}, b) \xrightarrow{\mathsf{next}} (\mathcal{C}', b)$ if and only if there exist $(\bar{\xi}, \bar{\phi})$ isomorphic to (ξ, ϕ) such that if $\bar{\xi}_i$ (the i^{th} element of the vector ξ) belongs to $\mathcal{C}$, the edge next to $\bar{\xi}_i$ according to the direction b (if b is positive, that is $\bar{\xi}_{i+1}$, otherwise that is $\bar{\xi}_{i-1}$) belongs to $\mathcal{C}'$.

- The inv relation describes how the configuration is modified when we cross the corresponding edge: $(\mathcal{C}, b) \xrightarrow{\mathsf{inv}} (\mathcal{C}', b')$ if and only both configurations correspond to the same color of edge $\mathfrak{e}$, and an edge neighborhood isomorphic to $\eta_{\mathfrak{e}} \in \eta$ has $(\mathcal{C}, b)$ and $(\mathcal{C}', b')$ at its extremities.

Given these two relations, it is possible to build a finite state automaton on the set of configurations. This automaton is called the *border automaton* associated to this labeling scheme. By Lemma 8, it is connected. The inv relation is involutive, while the next relation may have more than one successor, therefore this automaton is non-deterministic in general. Determining the border of a face $\mathcal{F}$ is only a matter of reading the infinite word $(\mathsf{next} \cdot \mathsf{inv})^{\omega}$, with starting state a configuration such that the face $\mathcal{F}$ appears between this configuration and the configuration next to it in the automaton. Such a word is an orbit of the automaton under the action of $(\mathsf{next} \cdot \mathsf{inv})$. □

For a given labeling scheme, we can associate to an element of the face vector – or a face – an orbit of the automaton. The colors of the faces are supposed to distinguish the classes of the faces under the group of automorphism. Two finite faces are said to be *equivalent* if they have the same orbit or orbits that correspond to the same border read in opposite directions. If two faces of the same color are not equivalent, then the labeling scheme is said to be *invalid*. Labeling schemes resulting from TLF-planar graphs are always valid. This property ensures that given two faces of the same color, there exists an automorphism mapping the first onto the second.

The orbits of the automaton can be classified into two categories : the cyclic orbits and the acyclic ones. Cyclic orbits correspond to faces that can be "closed", meaning that their border is periodic. Consider the orbit containing the i-th edge of ξ, with positive direction of rotation. If this orbit is cyclic, then we define k_i as the size of this orbit, otherwise $k_i = \infty$. Let $(a_i), i \in [1; d]$ be a set of formal letters such that $a_i = a_j$ if and only if the i-th and j-th face of ϕ are of the same color. The vector whose elements are the $k_i a_i$ (or simple ∞ if k_i is infinite) is called the *primitive type vector* of (ξ, ϕ, η). A type vector $[l_1, \ldots, l_d]$ is said to be *valid* with regard to that labeling scheme if and only if there exists a valuation of the (a_i) in $\mathbb{N}$ such that $\forall i, l_i = k_i a_i$ and all values in the vector are greater than three.

3.2 Examples of construction

Example 6. Consider the labeling scheme defined in Figure 5, with degree 5, $\mathsf{E} = \{\mathfrak{b}, \mathfrak{r}, \mathfrak{g}\}$ and $\mathsf{F} = \{\alpha, \beta, \gamma\}$. With each color in E we associate a unique edge neighborhood. Let us try to compute the associated border automaton. Even if $|\mathsf{E}| = 3$, there exist 10 different configurations, one for each edge and

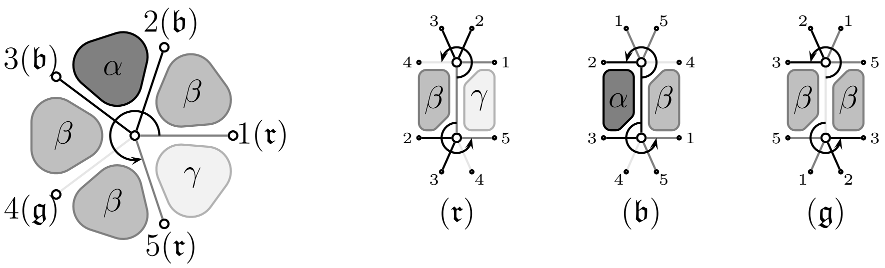

Fig. 5. Another labeling scheme of degree 5

direction of rotation. Since the next operation corresponds to a rotation of the edge vector, we represent the border automaton with two cycles of length 5. On the Figure 6, next edges appear in black, while inv edges appear in dashed gray lines:

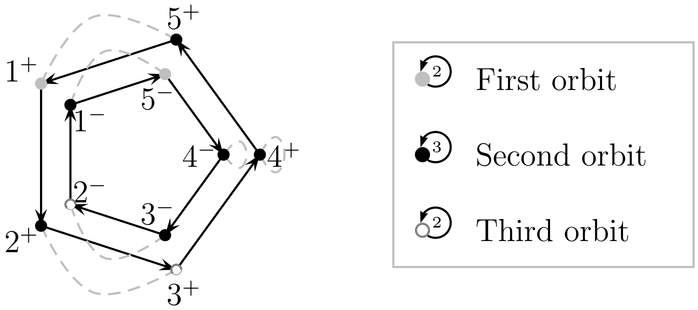

Fig. 6. The border automaton associated to the labeling scheme in Figure 5.

The possible orbits are therefore [1], [2, 4, 5], and [3]. The first one corresponds to faces with red borders $\mathfrak{r}^\star$, and the third one to faces with blue borders $\mathfrak{b}^\star$. The others have border $(\mathfrak{bgr})^\star$. We have three possible classes of borders of the faces, corresponding to each orbit of the automaton and to each color in F. The corresponding primitive type vector is: $[3n, m, 3n, 3n, p]$ for $(n, m, p) \in \mathbb{N}$. On Figure 7, there is an planar embedding in the hyperbolic plane of a graph possessing this labeling scheme and type vector $[3, 4, 3, 3, 5]$. The faces α correspond squares and the faces γ to pentagons, while the faces with 3-colored borders are triangles. Notice that this graph possesses a trivial stabilizer and is therefore a Cayley graph.

Example 7. Consider now the case of example 5 page 15, where $\mathsf{E} = \{\mathfrak{b}, \mathfrak{r}, \mathfrak{g}\}$ and $\mathsf{F} = \{\alpha, \beta, \gamma\}$. As in the previous example, we associate to each color in E a unique edge neighborhood. The labeling scheme is valid with regard to our definitions. When we assume that the face β is infinite, we come down with the border automaton appearing on Figure 8.

The first orbit corresponds to the face γ, which is bordered by $(\mathfrak{rrb})^\star$. The second orbit corresponds to the face α, bordered by $(\mathfrak{bg})^\star$. More interesting is the case of the face β. Notice that due to the possible twists, the next operator is not deterministic.

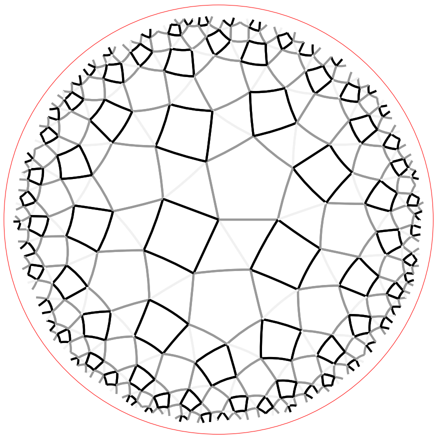

Fig. 7. Example of a 3-connected TLF-planar vertex-transitive graph, with type vector $[5; 3; 4; 3; 3]$, associated to the labeling scheme described in the Example 4.

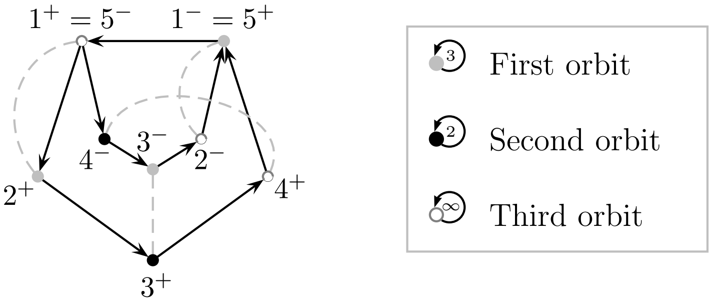

Fig. 8. The border automaton associated to the labeling scheme in Figure 3.

This leads to faces bordered by aperiodic words from $\mathfrak{r}^\omega\mathfrak{g}\mathfrak{r}^\omega$. There does not exist any embedding of the graph can avoid these aperiodic faces (cf. Fig. 9), and all infinite faces either have a unique yellow edge in their border or none at all. The associated primitive type vector is given by $[\infty, 3n, 2m, \infty, 3n]$ for values of $(m, n) \in \mathbb{N}$, which leads to 1-separable graphs such as the one on the right. For Cayley graphs, it is always possible to find an embedding where the border of the faces are periodic [12]. Hence, this set of graphs contains no Cayley graphs, because of this aperiodic border property.

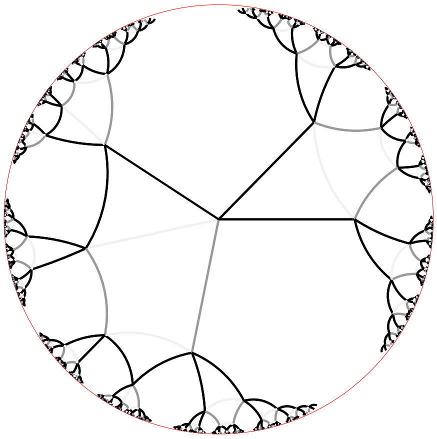

Fig. 9. Example of a 1-separable TLF-planar vertex-transitive graph, with type vector $[3; \infty; 3; 4; \infty]$, associated to the labeling scheme described in the Example 5.

Every TLF-planar vertex-transitive graph possesses a valid labeling scheme and a vector type that is valid regarding to that scheme. Consider now the converse of this result: given a general valid labeling scheme, and a valid type vector, is it possible to produce a vertex-transitive graph that has the same labeling scheme and type vector ?

Our goal in this section consists in building specific embeddings of vertex-transitive planar graphs with particular geometrical properties. Depending on the graph, we select an appropriate geometry: Euclidean, spherical or hyperbolic. As long as our graphs are TLF-planar, 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.

Theorem 10 (Existence) *Given a labeling scheme (ξ, ϕ, η), and a valid type vector $[k_1, \ldots, k_d]$, there exists a TLF-planar vertex-transitive graph possessing this scheme and type vector. Moreover, all faces of the embedding are regular polygons.*

Proof : 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. The following values result from simple trigonometry in the different geometries:

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

The equation (1) 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 (1) that determines a unique value for the length l, and consequently for the angles θ_i [1]. These values allow us to draw all the regular polygons that correspond to our type vector around a given point of the plane.

[1] 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 22).

Let ϵ be a point in the plane. By induction, we build a planar locally finite graph Γ_n, with a central vertex ϵ, such that all vertices at distance $\leq n$ from ϵ have degree d. The graph Γ_1 is composed of the glued polygons of the first case, centered onto ϵ, but restricted to the edges incident to ϵ. The labeling of Γ_1 is chosen isomorphic to (ξ, ϕ), such that the face ϕ_i corresponds to a regular polygon with k_i sides. Suppose now that we have built Γ_n. Consider the finite set of vertices of Γ_n at distance n of degree less than d. We shall build the remaining edges with geodesics, such that the length of the geodesic be l, and the angle at the vertex corresponding to the face ϕ_i be equal to $\theta_i(l)$.

The labeling scheme describes how the edges must be labeled : take a vertex s and an edge e colored by $\mathfrak{e}$ incident to s. According to Lemma 8, there exists a unique way to glue the edge neighborhood $\eta_{\mathfrak{e}}$ onto e, up to isomorphism. The choice of the edge neighborhood is not crucial, because the isomorphisms leave the angles of the faces constant. The construction of Γ_{n+1} is just a matter of gluing the edge neighborhoods on the edges possessing an extremity at distance $n+1$ from ϵ, by choosing an adequate edge neighborhood isomorphic to $\eta_{\mathfrak{e}}$.

This construction can not lead to intersecting edges. In fact; all edges have the same length, and inside a given face, all angles are the same. Therefore the construction of a face ultimately leads to a regular polygon in the plane. This answers the problem of the closure of the faces.

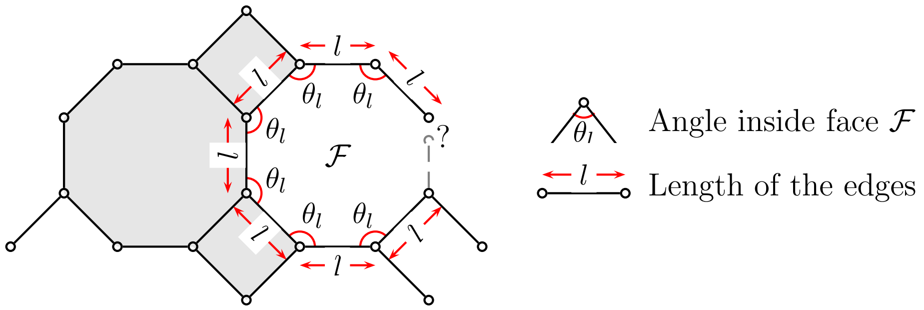

Hence the limit graph Γ is well defined. This graph is vertex-transitive: the construction is independent of the starting vertex, and applying this construction with two different starting vertices creates an automorphism mapping the first vertex on the second. The labeling of the edges entails that the automorphism group is transitive. The automorphisms stabilizing a vertex exactly correspond to the classes of equivalence of the configurations. □

Theorem 11 (Coherence) *If Γ is a TLF-planar vertex-transitive graph that possesses a labeling scheme (ξ, ϕ, η) and a type vector $[k_1, \ldots, k_d]$, then the preceding construction yields a graph isomorphic to Γ.*

Therefore, for any vertex-transitive 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 finite faces onto isometric faces. Consequently, in the case of the 2-connected graphs, every automorphism of Γ extends into an isometry of the plane. This result appears in Babai [1] in the case of 3-connected graphs, and is here extended to the larger case of the TLF-planar graphs. Notice that the case of 1-separable graphs is more complex: even is the faces of the graph are regular polygons, example 7 shows that there exists graphs for which it is impossible to find an embedding where automorphisms may be realized by isometries of the plane.

A direct consequence of this result pertains to the growth-rate of the TLF-planar graphs. When Γ is 2-separable, its structure is arborescent and its growth-rate is either linear or exponential. On the other hand, when Γ is at least 3-connected, the embedding of Γ given by the Theorem 10 is quasi-isometric to the plane inside which it is defined. Therefore the growth-rate of the graph is either quadratic or exponential.

3.4 Combinatorial approach

Consider the problem of finding all the valid labeling schemes for graphs of degree d. The number of colors in E and F is bounded by d. Therefore the number of possible edge and face vectors is bounded by d^{d-1}, up to rotation, and the number of edge neighborhoods for a given color by $2d$. It is therefore possible to enumerate all possible labeling schemes, and then to extract the valid schemes by computing the border automaton. The particular properties of labeling schemes allow to drastically reduce this theoretical upper bound.

Theorem 12 (Enumeration) *Given a number $d \geq 2$, it is possible to enumerate all TLF-planar transitive graphs having internal degree d, along with their labeling scheme and primitive type vector.*

A possible question lays in determining all possible type vectors for TLF-planar vertex-transitive graphs. We enumerate the graphs of a given degree, and eliminate the redundant ones by testing if they belong to the same class of isomorphism. For example, the possible type vectors for vertex-transitive graphs of degree 3 are:

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

... for values of n, m, p such that the faces of the graphs are at least triangles. The set of planar graphs obtained is presented in appendix, and extends the vertex-transitive graphs presented in [5]. Notice also that this set is strictly larger than the set of possible Cayley graphs of the same degree, containing in particular the graph associated to the dodecahedron, with type vector $[5;5;5]$ and the icosidodecahedron, with type vector $[3;5;3;5]$, both of which are not Cayley graphs. More generally, this is the case of all tilings with type vector $\{[6n\pm1, 6n\pm1, 6n\pm1]\}$. This is an example o an infinite family of TLF-planar transitive graphs which are not Cayley.

More precisely, for a given labeling scheme (ξ, ϕ, η) and valid type vector, we have the possibility to check whether the associated unlabeled graph is Cayley or not, by enumerating all Cayley graphs having the same type vector [12], computing the labeling schemes of these transitive graphs and comparing them to (ξ, ϕ, η):

Corollary 13 (Cayley checking) *If Γ is a TLF-planar transitive graph, then it is decidable whether Γ is the Cayley graph of a group or not, and obtain an enumeration and a description of the groups having Γ as a Cayley graph.*

Another side-result is the possibility to enumerate for a given graph Γ all the groups of automorphisms acting transitively on the set of vertices of Γ. This raises the question on finding graphs which are vertex-transitive but not Cayley, *i.e.* who do not possess a group of automorphism acting simply transitively on its vertices. Consider the stabilizer of a vertex of Γ: it must be a finite group of isometries of the plane fixing the vertex, hence it is either cyclic or dihedral. If this stabilizer does not contain any rotation, then the subgroup of direct isometries of $\mathrm{Aut}(\Gamma)$ acts simply transitively on the vertices of Γ, therefore such graphs must possess rotations in their stabilizers.

The $\mathsf{multiply}(k)$ operation on a labeling scheme (ϕ, ξ, η) of degree d consists in building another labeling scheme (ϕ', ξ', η') of degree $k \times d$ defined by:

- ϕ' and ξ' consist of k copies of ϕ and ξ, *i.e.* $\phi'_i = \phi_{(i \bmod d)}$, and $\xi'_i = \xi_{(i \bmod d)}$;
- each edge neighborhood $\eta'_{\mathfrak{e}}$ is a copy of $\eta_{\mathfrak{e}}$ where to the edge and face vectors at each extremity have been replaced by (ξ', ϕ').

The set of configurations of the automaton remains unchanged by this operation, as is the construction of the border automaton. This operation allows to build labeling schemes with stabilizers containing rotations. Moreover, this operation is invertible: for a labeling scheme whose stabilizer possesses rotations, it is possible to extract a labeling scheme without rotations by an operation of division.

Lemma 14 (Taxonomy of prime degree graphs) *Let Γ be a 3-connected planar vertex-transitive graph of prime degree d. Either Γ is a face and edge-transitive graph with type vector $[k;\ldots;k]$, or it is a Cayley graph.*

Proof : Consider the group of automorphisms of Γ and the stabilizer G_s of a vertex s of Γ. Since the graph can be embedded in the plane by regular polygons, so that the automorphisms map the finite faces onto finite faces, G is a group of isometries of the plane and G_s is a finite group of isometries of the plane leaving s stable. If G_s contains a non-trivial rotation, since the degree of Γ is prime, the graph is edge and face-transitive. On the contrary, if G_s contains no rotation, then it has a single non-trivial element which is a symmetry. Then the subgroup of the direct isometries in G acts simply and transitively on the vertices of Γ. Therefore Γ is a Cayley graph. $\square$

This analysis points out the classes of graphs that are likely not to be Cayley graphs. For the lesser degrees, these graphs are either face-transitive $[k;\ldots;k]$, or edge-transitive $[k_1,k_2,\ldots,k_1,k_2]$. There exists an example of graph of degree 9 belonging to this class that is neither face-transitive nor edge-transitive.

Discussion

The geometrical invariants of TLF-planar graphs allowed here the description of the graphs by their local properties, in the case of vertex-transitive graphs. Nevertheless, there exists different directions where we could extend this study. First, the dual graph of a transitive planar graph, is not transitive, but cofinite. That is to say that there exists only a finite number of orbits of the vertices under the action of the group of automorphisms. Even if these graphs in general do not respect the same geometrical properties as the transitive graphs, it would be interesting to study this family and their embeddings. We could then question the possibility of extending the representation by labeling schemes to cofinite graphs. Second, the key property of local finiteness is rather restrictive: our graphs possess one accumulation point at most. Yet Levinson showed that the number of accumulation points, if superior to one, is either two or infinite. The analysis of the embeddings of planar graphs with two or more accumulation points could lead to a geometrical representation for a larger class of vertex-transitive planar graphs.

References

[1] 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.

[2] T. Chaboud. *Pavages et graphes de Cayley planaires.* PhD thesis, Ecole Normale Supérieure de Lyon, Février 1995.

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

[4] Max Garzon. *Models of massive parallelism: analysis of cellular automata and neural networks*. Springer-Verlag, 1995.

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

[6] M.-C. Heydemann and B. Ducourthial. *Graph symmetry: algebraic methods and applications*, volume 497 of the NATO ASI Series (Series C: Mathematical and Physical Sciences), chapter Cayley graphs and Interconnection networks, pages 167–224. Kluwer Academic Publishers, 1997.

[7] W. Imrich. On Whitney's theorem on the unique embeddability of 3-connected planar graphs. In *Recent advances in graph theory, Proc. Symp. Prague 1974*, pages 303–306, 1975.

[8] J. Mazoyer and M. Delorme. *Cellular automata : a parallel model*, volume 460 of *Mathematics and its applications*. Kluwer, 1998.

[9] B. D. McKay and C. E. Praeger. Vertex-transitive graphs which are not Cayley graphs I. *J. Austral. Math. Soc. (Series A)*, 56:53–63, 1994.

[10] B. D. McKay and C. E. Praeger. Vertex-transitive graphs that are not Cayley graphs II. *J. Graph Theory*, 22(4):321–334, 1996.

[11] B. D. McKay and J. F. Royle. The transitive graphs with at most 26 vertices. *Ars Combin.*, 30:161–176, 1990.

[12] D. Renault. Enumerating planar locally finite Cayley graphs. Submitted. Also available from `arxiv.org`, at "`http://arxiv.org/abs/cs.DM/0309017`", August 2003.

[13] G. F. Royle. Catalogues of combinatorial objects. "`http://www.cs.uwa.edu.au/~gordon/data.html`", 1997. (Last access : 17 December 2004).

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

[15] J. Wolf. *Spaces of constant curvature*. Publish or Perish, Wilmington, 1984. Chap. 2.