JÉRÔME LOS

On the forcing relation for surface homeomorphisms

Publications mathématiques de l'I.H.É.S., tome 85 (1997), p. 5-61

http://www.numdam.org/item?id=PMIHES_1997__85__5_0

© Publications mathématiques de l'I.H.É.S., 1997, tous droits réservés.

L'accès aux archives de la revue « Publications mathématiques de l'I.H.É.S. » (<http://www.ihes.fr/IHES/Publications/Publications.html>) implique l'accord avec les conditions générales d'utilisation (<http://www.numdam.org/conditions>). Toute utilisation commerciale ou impression systématique est constitutive d'une infraction pénale. Toute copie ou impression de ce fichier doit contenir la présente mention de copyright.

NUMDAM

Article numérisé dans le cadre du programme
Numérisation de documents anciens mathématiques
<http://www.numdam.org/>

ON THE FORCING RELATION FOR SURFACE HOMEOMORPHISMS

by JÉRÔME LOS

Introduction

In 1964 Sharkovskii [Sha] published a beautiful theorem on the ordering of the periodic orbits for continuous maps of the interval. The idea behind the theorem was that the existence of a single periodic orbit forces the existence of many others (in general infinitely many). In Sharkovskii's theorem a periodic orbit was only characterized by its period which is the weakest invariant one can use. Later on, several authors have refined the Sharkovskii ordering by using the permutation of the points on the line in order to characterize a periodic orbit (see for instance the monograph [ALM]). The key tools for studying this forcing relation was the use of a canonical piecewise linear model and an associated Markov partitions.

More recently the study of an analogous forcing relation was started for periodic orbits of surface homeomorphisms [Bo], [Ma]. Let S be a compact orientable surface and let $\text{Homeo}^+(S)$ be the group of orientation preserving homeomorphisms of S . If $P = \{x, f(x), \dots, f^n(x) = x\}$ is a periodic orbit of f , then we consider $[f|_{S-P}] \in \text{Mod}_0(S - P)$, the isotopy class of f in the complement of P in S . The group $\text{Mod}_0(S - P)$ is known as the *mapping class group* of the punctured surface $S - P$. The *braid type* of P , denoted by $\text{BT}(P, f, S)$, is the conjugacy class of $[f|_{S-P}]$ in $\text{Mod}_0(S - P)$. In the case where the surface S is the disk D^2 , the relationship between the mapping class group $\text{Mod}_0(D^2 - n \text{ points})$ and the n -strands braid group B_n is well known (see [Bi]). This explains the name braid type. We shall denote a braid type as $[\beta]$, assuming that the surface is given. If no confusion is possible we shall omit the brackets.

Definition 0.1. — For a given compact orientable surface S , we say that a braid type $[\beta]$ *forces* a braid type $[\gamma]$, which we denote by: $[\beta] \succ [\gamma]$, if every $f \in \text{Homeo}^+(S)$ which has a periodic orbit of braid type $[\beta]$ has also a periodic orbit of braid type $[\gamma]$.

In the subclass $\text{Homeo}_1^+(S)$ of homeomorphisms which are isotopic to the identity this forcing relation “ $\succ$ ” is a *partial order*. The proof of this result has been given by

Boyland [Bo]. This proof is given via a bifurcation analysis which is odd in this purely topological context. It turns out that the partial ordering result is an easy corollary of our methods (see § 5).

The goal of this paper is to study some structures of the set

$$\mathcal{G}(\beta) = \{ \text{braid types } [\gamma] \text{ such that: } [\beta] \succ [\gamma] \},$$

which we call the *genealogy* of the braid type $[\beta]$. Our aim is to describe some topological structure of the set $\mathcal{G}(\beta)$ (at least “locally”).

In order to do so, we introduce a graph-like description of $\mathcal{G}(\beta)$, i.e. we represent the elements of $\mathcal{G}(\beta)$ as vertices and we define an oriented path $\pi(\gamma, \delta)$ in $\mathcal{G}(\beta)$ as a segment which connects the braid type $[\gamma]$ to the braid type $[\delta]$ if $[\gamma] \succ [\delta]$. By the transitivity of the order relation, if $[\beta] \succ [\gamma]$ and $[\gamma] \succ [\delta]$ then $[\beta] \succ [\delta]$, therefore there are oriented paths $\pi(\beta, \gamma)$, $\pi(\gamma, \delta)$ and $\pi(\beta, \delta)$ in $\mathcal{G}(\beta)$. We impose further that the oriented path $\pi(\beta, \delta)$ is the concatenation of the oriented paths $\pi(\beta, \gamma)$ and $\pi(\gamma, \delta)$. A path $\pi(\gamma, \delta)$ can then be interpreted as a subset of $\mathcal{G}(\beta)$, namely

$$\{ [\mu] \in \mathcal{G}(\beta) \text{ s.t. } [\gamma] \succ [\mu] \succ [\delta] \}.$$

We do not require uniqueness of this path $\pi(\gamma, \delta)$. Let $[\gamma]$ and $[\delta]$ be two braid types in $\mathcal{G}(\beta)$ which are not related by the forcing relation; such a pair is called *unrelated*. Let $B^+(\gamma, \delta) \subset \mathcal{G}(\beta)$ be the subset of braid types which force both $[\gamma]$ and $[\delta]$. For each $[\mu] \in B^+(\gamma, \delta)$ there are (at least) two paths $\pi(\mu, \gamma)$ and $\pi(\mu, \delta)$ which might have a common initial part. If $\pi(\mu, \gamma) \cap B^+(\gamma, \delta)$ were finite then there would be a “last” element $[\mu_N] \in \pi(\mu, \gamma) \cap B^+(\gamma, \delta)$. Such a point would be called a *ramification point*. At such a point there would be (at least) one arriving path and two starting paths. The notion of the *degree* of a ramification point, as the number of distinct paths arriving and starting at this point would then be natural.

Unfortunately the situation is much more complicated than the above naive description. Indeed the set $\pi(\mu, \gamma) \cap B^+(\gamma, \delta)$ is generally infinite. In section 8, we give a definition of a ramification point and its degree which generalize the above intuitive idea. These definitions are in terms of limits of nested paths in $\mathcal{G}(\beta)$. If $[\gamma]$ and $[\delta]$ are unrelated braid types in $\mathcal{G}(\beta)$ we denote by $\rho(\gamma, \delta)$ the set of ramification points corresponding to this pair.

For a given braid type we also define the *entropy function* $h: \mathcal{G}(\beta) \rightarrow \mathbf{R}^+$ by $h[\gamma] = \min \{ h(\varphi); \varphi \in [f_\gamma], [f_\gamma] \in [\gamma] \}$, where $h(\varphi)$ is the topological entropy of the homeomorphism φ in the isotopy class $[f_\gamma]$ (see [FLP] for instance, or [Lo1]). The function h is a well-defined braid type invariant, since the topological entropy is a topological invariant.

In this paper we shall restrict our attention mainly to the case where the surface is a disk D^2 . Therefore a braid type is a conjugacy class of braids and $\mathcal{G}(\beta)$ is partially ordered, since $\text{Homeo}^+(D^2) = \text{Homeo}_1^+(D^2)$. Nevertheless most of the techniques we develop in this paper and most of the arguments are exactly the same for other surfaces.

The key preliminary step, in order to study the set $\mathcal{G}(\beta)$, is to restrict the study from the entire isotopy class $[f|_{S-P}]$ to a single element, namely the canonical representative of the class as defined by the Nielsen-Thurston theorem [Th1] (see [FLP]). This restriction is justified by the theorems of Birman-Kirdwell [BK], Asimov-Franks [AF] and Hall [Hal] which says, with the above terminology, that if $[f|_{S-P}] \in [\beta]$ is a pseudo-Anosov (P.A.) class, then:

- (*) $\mathcal{G}(\beta)$ is exactly the set of braid types $BT(Q, \varphi, S)$, where Q is a periodic orbit of the pseudo-Anosov representative φ in $[f|_{S-P}] \in [\beta]$.

The second preliminary step is to find a Markov partition for the pseudo-Anosov map, as well as an analogue of the piecewise linear model used in dimension one. This second step has been achieved recently by results of Bestvina-Handel [BH1], the author [Lo2] and a little later by Franks-Misiurewicz [FM]. In these papers the central role is played by a special class of maps on embedded graphs which will be called *efficient maps*. These maps on graphs are closely related to the *train tracks maps* as defined by Thurston [Th2] following ideas of Williams [W]. Notice that, since $\mathcal{G}(\beta)$ is partially ordered, there is a topology on $\mathcal{G}(\beta)$, called the *path topology*, which is induced by the partial ordering. The main results of the paper, concerning the structure of $\mathcal{G}(\beta)$, is collected in the following theorem:

Theorem 0.2. — *Let β be a braid type and let $\mathcal{G}(\beta)$ be its genealogy set. Then $\mathcal{G}(\beta)$ is finite if and only if $h[\beta] = 0$. If $\mathcal{G}(\beta)$ is infinite then the isotopy classes in $[\beta]$ have, at least, one pseudo-Anosov component. We assume that $[\beta]$ is a pseudo-Anosov conjugacy class and we denote by φ_β the pseudo-Anosov element in one of the classes. Then there exists a neighborhood of $[\beta]$ in the path topology, $V_{sc_\beta} \subset \mathcal{G}(\beta)$, on which the following properties hold:*

1. *For every $[\gamma] \in \mathcal{G}(\beta)$, there exists $[\delta] \in V_{sc_\beta}$ so that $[\delta] \succ [\gamma]$.*
2. *V_{sc_β} is metrizable.*
3. *For any $[\gamma] \succ [\delta]$, $[\gamma] \neq [\delta]$ in V_{sc_β} , the intersection $\pi(\gamma, \delta) \cap \mathcal{G}(\beta) \subset V_{sc_\beta}$ is infinite.*
4. *If $[\gamma]$ and $[\delta]$ are two unrelated braid types in V_{sc_β} then a ramification point $\rho(\gamma, \delta)$, as a limit of braid types, corresponds to the type of an infinite orbit of the pseudo-Anosov map φ_β . Therefore $\rho(\gamma, \delta)$ does not define a braid types in $\mathcal{G}(\beta)$.*

For a given unrelated pair $([\gamma], [\delta])$ the set of ramification points $\rho(\gamma, \delta)$ is finite.

Let us consider the set $\overline{\mathcal{G}(\beta)} = \mathcal{G}(\beta) \cup \rho(\gamma, \delta)$ which is obtained by adding to $\mathcal{G}(\beta)$ the set of all the ramification points for all the unrelated pairs in $\mathcal{G}(\beta)$.

5. *The partial order “ $\succ$ ” on $\mathcal{G}(\beta)$ induces a partial order “ $\geq$ ” on $\overline{\mathcal{G}(\beta)}$.*
6. *On any non trivial path $\pi(\gamma, \delta)$ in V_{sc_β} there are infinitely many ramification points.*
7. *The degree of a ramification point is finite.*
8. *$\overline{\mathcal{G}(\beta)}$ is “non simply connected”.*

Notice that the first part of the theorem concerning the zero entropy case is well known. Indeed the zero entropy periodic orbits has been studied by several authors (see for instance [Sm], [GST]).

The paper is organized as follow:

The results about efficient maps and train tracks are recalled in section 1. In section 2 we recall the construction of a Markov partition from an efficient map and also the well-known symbolic description of the periodic orbits. We point out the fact that this symbolic description can be given simply in term of the efficient representative. In section 3 we relate the braid types of the elements of the genealogy set to a certain branched surface embedded in a 3-manifold which is a fibered bundle over the circle and whose fiber is the surface S . Each element of the genealogy set gives rise to a knot which is “carried” by the branched surface. We give a symbolic description of all these knots. The sections 1 to 3 are the preliminary steps which are needed in the paper. Most results in these sections were essentially known even if some of them have not been published.

One of the new technical results of the paper, which uses efficient maps, is called the characterization theorem (see section 4). This theorem can be formulated as follows:

Theorem 0.3. — *Let $[\beta]$ be a pseudo-Anosov braid type and let $[\gamma], [\delta]$ be two elements of $\mathcal{G}(\beta)$. There exists an effective criterion which enables one to decide whether $[\gamma]$ and $[\delta]$ are related or not by the forcing relation.*

This characterization theorem enables one to define the subclass V_{sc_β} of braid types in $\mathcal{G}(\beta)$ and to prove that it forms a neighborhood of $[\beta]$ in $\mathcal{G}(\beta)$. A periodic orbit corresponding to a braid type in this subclass is called a *small cancellation* periodic orbit (see section 5). In this subclass of braid types we define some numerical quantities for each such periodic orbit which allow us to reduce the characterization of the forcing relation to a finite collection of inequalities (see section 6). These inequalities make it possible to define a distance function among small cancellation periodic orbits. That the forcing relation can be read from these inequalities follows from the fact that, in the subclass of small cancellation periodic orbits, we can control effectively the construction of an efficient map for the class defined by $[\delta]$ out of an efficient map for the class defined by $[\beta]$ (with $[\beta] > [\delta]$).

CONTENTS

1. Train tracks, efficient maps	9
2. Markov partition, periodic orbits	16
3. From a symbolic code to a braid type	23
4. Characterization of the forcing	29
5. Small cancellation orbits	33
6. Cancellation rectangles	42
7. Periodic orbits with constraints	47
8. The main theorem.....	51

1. Train-tracks, efficient maps

In this section we review some results about train tracks. The notion of train tracks has been used over the last few years within different contexts which are overlapping. Therefore, to avoid confusions, we shall use distinct words in this paper.

For the original definition of a train track on a surface, which is due to Thurston [Th2], we refer the reader to the bibliography on the subject (see for instance [Pe]).

For free group automorphisms several versions of train tracks also appear, for instance in [BH2], [Lu], [Lo3]. These definitions do not coincide with those of the usual train tracks in the case where the free group is a surface group. In this paper we are dealing with punctured surfaces and therefore the confusion is possible. Moreover the three papers mentioned in the introduction [BH], [Lo1], [FM] do not use the same terminology. We chose to follow, when possible, the terminology of [BH1] which is more appropriate to our needs.

1.1. Efficient representative

Throughout the paper, terminology from graph theory will be used; we refer the reader to [St], [BH2], [Lo3].

Let S be a punctured surface and f be an orientation-preserving homeomorphism of S , i.e. $f \in \text{Homeo}^+(S)$. Fix a base point $O \in S$ and a generating set $\{x_1, x_2, \dots, x_n\}$ for the fundamental group $\pi_1(S, O)$. Then we consider an automorphism $f_\pi : \pi_1(S) \rightarrow \pi_1(S)$ induced by f . This automorphism can be seen from two points of view. The first one is algebraic, i.e. we consider the collection $\{f_\pi(x_1), f_\pi(x_2), \dots, f_\pi(x_n)\}$ of the images of the generators under f_π as words in $\{x_1^{\pm 1}, \dots, x_n^{\pm 1}\}$.

The second point of view is more topological. We consider the graph Γ_0 embedded in S , which has a single vertex $*$ and n oriented edges $\{e_1, e_2, \dots, e_n\}$. We identify the vertex $*$ with the base point $O \in S$ and each oriented edge e_i of Γ_0 with the generator x_i of $\pi_1(S, O)$. With this picture, the automorphism f_π is identified with the continuous map $\Psi_0 : \Gamma_0 \rightarrow \Gamma_0$ which fixes the vertex $*$ and maps each edge e_i to an edge path $\Psi_0(e_i)$ which is identified with the word $f_\pi(x_i)$ by replacing $x_i^{\pm 1}$ by the corresponding edge $e_i^{\pm 1}$.

Consider now the whole set $G(S)$ of graphs which are embedded in S and whose fundamental group, with respect to a vertex, is identified with $\pi_1(S)$. For a given graph $\Gamma \in G(S)$ we denote by $E(\Gamma)$ and $V(\Gamma)$ respectively its sets of edges and vertices. The edges $e_j \in E(\Gamma)$ are oriented, from the *initial vertex* $i(e_j) \in V(\Gamma)$ to the *terminal vertex* $t(e_j) \in V(\Gamma)$. Consider now a map $\Psi : \Gamma \rightarrow \Gamma$, where $\Gamma \in G(S)$, which is homotopy equivalent to $\Psi_0 : \Gamma_0 \rightarrow \Gamma_0$ and such that Ψ maps $V(\Gamma)$ to itself (not necessarily one-to-one). We assume further that the image of each edge $e_i \in E(\Gamma)$ under Ψ is locally injective, i.e. no cancelation $(e_i \cdot e_i^{-1})$ occurs in the word $\Psi(e_j)$. Such a map is called a *topological representative of f* and is denoted as the pair (Ψ, Γ) . We also call the pair (Ψ_0, Γ_0) above an *initial representative of f* . Let (Ψ, Γ) be a topological representative

of f and assume that $\text{Card}[E(\Gamma)] = K$. The integer matrix $M(\Psi, \Gamma)$ of size K whose entries are

$$m_{i,j} = \text{number of letters } e_i \text{ and } e_i^{-1} \text{ in } \Psi(e_j)$$

is called the *incidence matrix*. Its largest eigenvalue is called the *growth rate* of the topological representative and is denoted by $\lambda(\Psi, \Gamma)$.

Definition 1.1. — An *efficient representative* for $f \in \text{Homeo}^+(S)$ is a topological representative (Ψ, Γ) of f so that $\Psi^k(e_i)$ is locally injective for all edges $e_i \in E(\Gamma)$ and all $k > 0$.

The first result has been proved by several authors [BH1], [Lo2], [FM]:

Theorem 1.2. — *If S is a punctured surface and $f \in \text{Homeo}^+(S)$ is isotopic to a pseudo-Anosov homeomorphism then there exists an efficient representative (Ψ, Γ) for f . Moreover the growth rate $\lambda(\Psi, \Gamma)$ of an efficient representative is the dilatation factor of the pseudo-Anosov element in the isotopy class $[f]$.*

The above theorem is proved in [Lo2] and [FM] in the case of the punctured disk. The proof of this theorem is constructive, i.e. it is given by a finite algorithm whose initial data is any initial representative (Ψ_0, Γ_0) .

1.2. Efficient representative: the algorithm

Let us recall the elementary operations of the algorithm mentioned above.

Move 1. — Collapsing an invariant forest.

Suppose that the topological representative (Ψ, Γ) has an invariant forest, i.e. an invariant subgraph T whose connected components $\{T_1, \dots, T_r\}$ are contractible. Then we change Γ to Γ' by collapsing each T_i to a vertex of Γ' . The new map $\Phi : \Gamma' \rightarrow \Gamma'$ is obtained from Ψ by supressing all the words $\Psi(e_i)$ for $e_i \in E(T)$ and removing each occurrence of a letter e_i (resp. e_i^{-1}) in the words $\Psi(e_j)$, where $e_j \in E(\Gamma) - E(T)$.

Move 1'. — Collapsing a pretrivial forest.

Assume that the topological representative (Ψ, Γ) admits a forest F whose image, under Ψ^* , is a collection of vertices. Such a forest is called *pretrivial*. In this case we transform the graph and the map by collapsing the forest F .

Move 2. — Valency one isotopy.

For a topological representative (Ψ, Γ) , suppose that Γ has a valency one vertex. We change the map Ψ by removing all occurrences of the letter e_1 , corresponding to the edge ending at the valency one vertex, from all the edge paths $\Psi(e_j)$, $j \neq 1$. Then we change Γ by removing the edge e_1 .

Move 3. — Valency two isotopy.

For a topological representative (Ψ, Γ) , suppose that Γ has a valency two vertex denoted by v . Chose one of the two edges e_1 or e_2 which are incident at v (say e_1). We

change the graph Γ by removing v and declaring the union of the two edges e_1 and e_2 a single edge e . In order to change the map Ψ , assuming that $v = t(e_1) = i(e_2)$, we replace each occurrence of $e_1.e_2$ (resp. $(e_1.e_2)^{-1}$) by e (resp. e^{-1}) in each edge path. If e_1 occurs alone (at the beginning or the end of an edge path) then we replace it by e and if e_2 occurs alone then we remove it. Finally, in order to define the image of the new edge e , i.e. $\Phi(e)$, we just consider the concatenation of the two edge paths $(\Psi(e_1).\Psi(e_2))'$, where the notation $(\dots)'$ refers to the previous operations on the edge paths.

Move 4. — Pulling tight.

Suppose that $\Psi : \Gamma \rightarrow \Gamma$ is not a topological representative because of the occurrence of a cancellation $e_j.e_j^{-1}$. Then we change the map by removing all these cancellations.

Move 5. — Folding.

This operation has been introduced by Stallings in [St] and Dicks [Di] for automorphisms of free groups. In the surface case we have an additional structure on the graph. Indeed the graph is embedded in an oriented surface S , therefore at each vertex $v \in V(\Gamma)$ there is a cyclic ordering of the incident edges which is induced by the orientation of the surface. Two edges e_1 and e_2 which are incident at a vertex v are said to be *adjacent* if they are consecutive in the cyclic order. Let e_1 and e_2 be two adjacent edges at $v \in V(\Gamma)$ and assume, for simplicity, that $v = i(e_1) = i(e_2)$. The two edges are said to be *tangent under Ψ* if $\Psi(e_1) = M.X$ and $\Psi(e_2) = M.Y$, i.e. in other words that the images of the two edges under Ψ have the same (non trivial) initial edge path. A pair of edges at a vertex v is also called a *turn* in Γ . A turn at v is denoted by the pair $(e_i^{\varepsilon_i}, e_j^{\varepsilon_j})$, where $\varepsilon_j = 1$ if $v = i(e_j)$ and $\varepsilon_j = -1$ otherwise. We say that the Ψ -image of an edge e *crosses* the turn (e_1, e_2) if, with the above orientation, $\Psi(e) = A.(e_1^{-1}.e_2)^{\pm 1}.B$, where A and B are some edge paths. Observe that if the turn (e_1, e_2) is tangent under Ψ and if there is an edge e which crosses (e_1, e_2) , then the second iterate $\Psi^2(e)$ is not locally injective. The same observation holds if an edge crosses a tangent turn under some iterate of Ψ . The goal of the folding operation is to remove these situations which create non local injectivity for some iterates. The transformation is the following:

(i) If the above situation occurs for a topological representative (Ψ, Γ) then we first transform the graph Γ into Γ' by identifying the beginning of the two edges e_1 and e_2 (with the above orientation). This operation creates a new vertex w and a new edge x .

(ii) Then we transform the map (from Ψ to Φ) in the following way (assuming that the orientations are as above). Each edge e_j of Γ is identified with an edge e'_j of Γ' and the new edge is called x . The new map is given by:

$$\Phi(e'_1) = X^*, \quad \Phi(e'_2) = Y^*, \quad \Phi(x) = M^*,$$

and $\Phi(e'_j) = \Psi(e_j)^*$, for all $e'_j \in E(\Gamma') - \{e'_1, e'_2, x\}$.

The notation A^* means that an edge path A in Γ is rewritten in the graph Γ' according to:

$$e_1 \mapsto x.e'_1, \quad e_2 \mapsto x.e'_2, \quad \text{and} \quad e_j \mapsto e'_j, \quad \text{for all other edges.}$$

Notice that the above transformation depends upon the chosen orientation of the edges e_1 , e_2 , x but the adaptations to the other cases are straightforward. Two particular folding operations have to be distinguished:

- If both edge paths X and Y are non trivial then the folding is called *partial*.
- If one edge path X or Y is trivial then the folding is called *absorbing*.

The algorithm which starts at any topological representative and stops at an efficient representative is a sequence of the moves (1)-(5). Notice that, for a given automorphism f_π , there are finitely many efficient representatives. In the free group case, and when the automorphism is irreducible a description of this (finite) set is given in [Lo3]. The surface case is a little different and it is not our aim to describe it here; such a description can be found in [Fe].

1.3. Topological representative: a topological view

Up to this point we have considered a topological representative (Ψ, Γ) of a surface homeomorphism $[\varphi]$ as an algebraic object. The facts that φ is a surface homeomorphism or that Γ is embedded in an oriented surface has not really been used. In this subsection we shall adopt a more topological point of view. This approach will introduce some additional information.

Let Γ be a graph in $G(S)$ and let $N(\Gamma)$ be a regular neighborhood of Γ in S . The neighborhood $N(\Gamma)$ is built by gluing together the following pieces:

- (i) For each vertex $v \in V(\Gamma)$ of valency k we consider a polygon $P(v)$ with k sides.
- (ii) For each edge $e \in E(\Gamma)$ we consider a rectangle $R(e)$.

The neighborhood $N(\Gamma)$ is obtained from the pieces $\{P(v); v \in V(\Gamma)\}$ and $\{R(e); e \in E(\Gamma)\}$ in the obvious way, by gluing together the sides of the polygon $P(v)$ with the corresponding sides of the rectangles $R(e)$ for the edges e which are incident at the vertex v (see figure 1).

FIG. 1. — The neighborhood $N(\Gamma)$

Each edge $e \in E(\Gamma)$ is embedded in $N(\Gamma)$ and intersects exactly two sides of $R(e)$. Let $R(e) = [0, 1] \times [0, 1]$; the edge $e \in E(\Gamma)$ intersects $R(e)$ at the sides $\{0\} \times [0, 1]$

and $\{1\} \times [0, 1]$. The orientation of the surface S induces an orientation of each rectangle; furthermore each edge has a given orientation. We chose the parametrization of $R(e)$ in such a way that the oriented interval $R(e) \cap e$ be $[0, 1] \times 1/2$. This choice fixes the orientation of each interval $\{t\} \times [0, 1]$ which we call a *tie*. Any tie in the rectangle $R(e)$ is denoted by $\theta(e)$.

Let us now consider $f \in \text{Homeo}^+(S)$ and let $f(\Gamma)$ be the image of the graph under f . We now transform the image $f(\Gamma)$ by a sequence of isotopies.

(1) Let f_1 be isotopic to f and satisfy: $f_1(\Gamma) \subset N(\Gamma)$. We assume furthermore that $f_1(N(\Gamma)) \subset N(\Gamma)$. Such a map $f_1 : N(\Gamma) \rightarrow N(\Gamma)$ is an embedding.

(2) Let f_2 be isotopic to f_1 so that the image $f_2(P(v))$ of a polygon $P(v)$ is contained in the interior of a polygon $P(w)$.

(3) Let f_3 be isotopic to f_2 and satisfy, in addition to the properties (1) and (2) above, that the intersection of $f_3(e_i)$ with any tie $\theta(e_j)$ is transverse.

Let us observe that $\{f_3(\text{Int}(e_i)); e_i \in E(\Gamma)\}$ is a set of disjointly embedded paths in $N(\Gamma)$ since Γ is embedded in S and f_3 is a homeomorphism. Observe also that $f_3 : N(\Gamma) \rightarrow N(\Gamma)$ is an embedding which can be extended to a homeomorphism on S by construction. From the definition there is a retraction $\rho : N(\Gamma) \rightarrow \Gamma$ such that $\rho(P(v)) = v$ for all $v \in V(\Gamma)$ and $\rho(R(e)) = e$ for all $e \in E(\Gamma)$. This retraction induces the following commutative diagram:

$$\begin{array}{ccc} N(\Gamma) & \xrightarrow{f_3} & N(\Gamma) \\ \rho \downarrow & & \downarrow \rho \\ \Gamma & \xrightarrow{\psi} & \Gamma \end{array}$$

which defines a continuous map $\psi : \Gamma \rightarrow \Gamma$. From the construction, the pair (ψ, Γ) is a topological representative for the isotopy class $[f]$.

From this topological description we can redefine all the quantities which have been defined algebraically, such as the incidence matrix and the collection of words describing the edge paths $\psi(e_i)$ for $e_i \in E(\Gamma)$. These definitions are straightforward. There is another collection of words which is not directly derived from the previous algebraic description.

Let $\theta(e_i)$ be a given tie in a rectangle $R(e_i)$ and let $\{(\mathcal{Y}_{i_1}^i)^{\varepsilon_1}, \dots, (\mathcal{Y}_{i_r}^i)^{\varepsilon_r}\}$ be the ordered collection of intersection points, with signs, of the tie $\theta(e_i)$ with the image $f_3(\Gamma)$. This collection is such that:

- $\mathcal{Y}_{i_m}^i \in \theta(e_i) \cap f_3(e_{i_m})$, where $e_{i_m} \in E(\Gamma)$;
- $\varepsilon_m = +1$ if the orientations of $\theta(e_i)$ and $f_3(e_{i_m})$ agree and $\varepsilon_m = -1$ otherwise;
- $\mathcal{Y}_{i_m}^i < \mathcal{Y}_{i_{m+1}}^i$, with the ordering given by the orientation of the tie $\theta(e_i)$.

The word $e_{i_1}^{\varepsilon_1} \cdot e_{i_2}^{\varepsilon_2} \dots e_{i_r}^{\varepsilon_r} \equiv \psi^\perp(e_i)$ is called the *transversal word at the edge e_i* of the topological representative (ψ, Γ) (see figure 2).

FIG. 2. — Edge path words — Transversal words

In what follows we shall often use both the topological representative $\psi : \Gamma \rightarrow \Gamma$ and the corresponding embedding $f_3 : N(\Gamma) \rightarrow N(\Gamma)$. For instance we shall assume that a topological representative is represented algebraically by the two collections of words $\{\psi(e_i); e_i \in E(\Gamma)\}$ and $\{\psi^\perp(e_i); e_i \in E(\Gamma)\}$.

1.4. From an efficient representative to a train track

We describe now the construction of a train track map on the surface S starting from an efficient representative (Ψ, Γ) .

We have already defined the notion of tangency under Ψ for a turn at a vertex v . The same notion is defined for any iterate Ψ^k . A *gate* at a vertex v is a subset of $\text{St}(v) = \{e \in E(\Gamma) \mid v = i(e) \text{ or } v = t(e)\}$ whose elements are pairwise tangent under some iterate Ψ^k .

Consider an efficient representative (Ψ, Γ) of f . From Γ and Ψ we construct a graph τ as follow:

(i) For each vertex $v \in V(\Gamma)$ we consider a small disk $\Delta(v)$ centered at v . The circle $\partial \Delta(v)$ intersects Γ at $x_i \in e_i$ for $e_i \in \text{St}(v)$. Let us cut Γ at each point x_i and remove from Γ the segments $e_i \cap \Delta(v)$ for all $e_i \in \text{St}(v)$.

(ii) On the circle $\partial \Delta(v)$, we identify all the points $x_{i_1}, \dots, x_{i_r}$ corresponding to the edges $e_{i_1}, \dots, e_{i_r}$ in the same gate. We assume furthermore that the edges are (locally) smoothly embedded in S so that the edges $e_{i_1}, \dots, e_{i_r}$ have a common tangent space at $x_{i_1} = \dots = x_{i_r}$ which is perpendicular to the circle $\partial \Delta(v)$.

(iii) If an edge e of Γ crosses a turn (e_i, e_{i+1}) , under some iterate Ψ^k , where e_i and e_{i+1} belong to two distinct gates, then the two points x_i, x_{i+1} are connected by an *infinitesimal edge* $\varepsilon_{i, i+1}$. This infinitesimal edge $\varepsilon_{i, i+1}$ is smoothly embedded in $\Delta(v)$ and intersects $\partial \Delta(v)$ perpendicularly at x_i and x_{i+1} (see figure 3).

FIG. 3. — Construction of a train track

Remark. — The algebraic notion of tangency for two adjacent edges at a vertex v coincides, in this construction, with the geometric tangency defined above in (ii).

The resulting graph is denoted by τ . Let us now define a map φ on τ which is induced by the map Ψ on Γ . The graph τ has two types of edges, i.e. the *real edges*, outside the disks $\Delta(v)$, and the *infinitesimal edges*, inside the $\Delta(v)$. The real edges are identified with the edges of Γ . Each real edge e is mapped, under φ , to an edge path $\varphi(e)$ which is obtained (algebraically) from the corresponding edge path $\Psi(e)$ on Γ by inserting, between every two real edges, the appropriate infinitesimal edge of τ . Each infinitesimal edge ε is mapped to another infinitesimal edge by φ . Indeed Ψ maps distinct gates of Γ to distinct gates of Γ . Therefore two distinct gates of Γ at the same vertex which are connected by an infinitesimal edge are mapped to two distinct gates connected by an infinitesimal edge. This defines the permutation of the infinitesimal edges under φ . From the map φ on τ we define an incidence matrix $\mathbf{M}(\varphi, \tau)$ exactly as for a topological representative. We check that this matrix has the following triangular bloc decomposition:

$\begin{pmatrix} \mathbf{N} & \mathbf{A} \\ 0 & \mathbf{M} \end{pmatrix}$. From the above construction, the diagonal bloc $\mathbf{M}$ is the incidence matrix $\mathbf{M}(\Psi, \Gamma)$ of the efficient representative. The other diagonal bloc $\mathbf{N}$ is a permutation matrix corresponding to the permutation of the infinitesimal edges. This observation implies that the largest eigenvalue of $\mathbf{M}(\varphi, \tau)$ is the same than the largest eigenvalue of $\mathbf{M}(\Psi, \Gamma)$. The corresponding eigenvector $\mu(\varphi, \tau)$ is unique (up to scale) and positive.

Proposition 1.3. — *If $f \in \text{Homeo}^+(S)$ belongs to a pseudo-Anosov class then the graph τ defined above is a φ -invariant train track (in the sense of Thurston). Furthermore if we equip each edge e of τ with the weight $\mu(e)$ given by the corresponding entry of the positive eigenvector $\mu(\varphi, \tau)$, then (τ, μ) is an invariant measured train track.*

This formulation is proved in [BH1] and a similar result is given in [Lo2]. The first part of the proposition is proved by checking that the tangencies are mapped to tangencies and that τ is embedded in S . This comes from the fact that (Ψ, Γ) is efficient

and is induced by an embedding (see 1.3). In order to prove the last part of the proposition we have to check that the entries of the eigenvector $\mu(\varphi, \tau)$ satisfy a *switch condition* at each vertex of τ (see figure 4).

FIG. 4. — A switch condition

2. Markov partition, periodic orbits

2.1. Markov partition, foliation

From a measured train track (τ, μ) which is invariant under an homeomorphism f , the construction of an invariant measured foliation and a Markov partition is well known (see [Pe], [FLP]). The idea is to replace each edge e of τ by a rectangle R_e of an appropriate size and with the obvious pair of transverse foliations, i.e. whose leaves are the horizontal (resp. vertical) segments. This operation defines a neighborhood $N(\tau)$ of the train track τ . In order to extend this foliated rectangular partition to a foliation of the entire surface S we have to “fill” the complementary regions. From the construction above and the non reducibility assumption, the connected components of $S - N(\tau)$ are either disks or once punctured disks (with cusped polygonal boundaries). Filling these disks is easy topologically by an isotopy. In fact we have to be more careful since we want to produce not only two foliations but also two measures.

Let us explain this construction more precisely. For a train track map (φ, τ) we have defined the measure $\mu(e)$ on each edge e of τ (see proposition 1.3) which is unique (up to scale). We can associate another measure to each edge e of τ by considering the transposed matrix ${}^t\mathbf{M}(\varphi, \tau)$. Indeed the largest eigenvalue of ${}^t\mathbf{M}(\varphi, \tau)$ is the same as the largest eigenvalue of $\mathbf{M}(\varphi, \tau)$. The corresponding eigenvector $\mu^\perp$ is unique (up to scale) and is non negative. To each edge e of τ we also associate the corresponding entry of $\mu^\perp$ which we call the *transposed measure*. Notice that from the triangular bloc structure of the matrix ${}^t\mathbf{M}(\varphi, \tau)$, the infinitesimal edges have transposed measure zero (this explains the term “infinitesimal”).

Let us go back to the efficient representative (Ψ, Γ) . To each edge we associate a rectangle $R_e = U_e \times S_e$ with the *stable foliation* whose leaves are $y \times S_e$, $y \in U_e$, and the *unstable foliation* whose leaves are $U_e \times x$, $x \in S_e$. Each foliation is equipped

with a *transverse measure*, i.e. the Lebesgue measure on the U_e (resp. S_e) which are locally segments of length $\mu(e)$ (resp. $\mu^\perp(e)$).

Recall that we have considered, around each vertex v of Γ , a disk $\Delta(v)$. The rectangles constructed above are embedded in the surface S so that:

1. Each component of $\partial U_e \times S_e$ of the boundary of R_e is an arc in a circle $\partial \Delta(v)$.
2. Each unstable leaf $U_e \times x$ is isotopic to the edge e , embedded in R_e by an isotopy keeping the endpoints in the $\partial \Delta(v)_e'$.
3. If two adjacent edges e and e' at v belong to the same gate, then $R_e \cap \partial \Delta(v)$ and $R_{e'} \cap \partial \Delta(v)$ intersect at a single point.
4. If $e \neq e'$ then R_e and $R_{e'}$ are disjoint, except for the intersections given by (3) (see figure 5).

FIG. 5. — Construction of a foliation

The above construction defines the rectangles corresponding to the real edges. The same construction holds for the infinitesimal edges $\varepsilon_1, \dots, \varepsilon_r$ inside each disk $\Delta(v)$, except that we have to fix, momentarily, the length of the infinitesimal rectangles to a non zero value. The switch conditions at v , which are satisfied by the entries of the eigenvector $\mu(\varphi, \tau)$, ensure that the *width* of all these rectangles match together.

As we noticed above the union $\bigcup_{e \in \mathbb{E}(\tau)} R_e$ is a neighborhood $N(\tau)$ such that the connected components of $S - N(\tau)$ are disks or once punctured disks. By the previous construction, the (cusped) disk components are contained in some of the disks $\Delta(v)$. Identifying these disks to a point produces a singularity of the foliation. This collapsing enables one to recover the fact that the lengths of the infinitesimal rectangles are zero. Let us call $\mathcal{R}$ the resulting space. The complement $S - \mathcal{R}$ is a collection of once punctured disks. The map φ on τ induces a map $\Phi : \mathcal{R} \rightarrow \mathcal{R}$ which maps each rectangle by an

affine map, stretching the unstable foliation by the factor λ (uniformly) and shrinking the stable foliation by the factor λ^{-1} (uniformly). Let us denote by $\partial\mathcal{R}$ the collection of *boundary curves*. The map Φ permutes the boundary curves. Each component of $\partial\mathcal{R}$ is a cusped polygon. Let us take a sufficiently high iterate, Φ^r such that each side of the polygon is mapped to itself, in an orientation preserving fashion. This iterate has therefore a fixed point on each side of a polygonal boundary curve. Each side of the polygon has a measure (transposed) coming from the foliation. We identify adjacent sides, up to the fixed points of Φ^r , in a measure preserving fashion. Finally all the fixed points, on the sides of a connected component, get identified to a single fixed point p , which we declare to be the puncture which is contained in this punctured disk.

The map Φ induces then a homeomorphism Φ^* on the surface S which is isotopic to the original map f on S . We also obtain the explicit construction of the invariant foliations, as well as the dilatation factor. The map Φ^* is then the pseudo-Anosov representative in the class.

The rectangle decomposition of the surface S coming from the construction is “almost” a *Markov partition* for Φ^* . Indeed for a true Markov partition the entries of the transition matrix consists only of 0 and 1. Whereas in the rectangle decomposition above, the transition matrix may have any non negative integer entries, since this transition matrix is exactly the incidence matrix $M(\Psi, \Gamma)$ of the efficient representative. In order to avoid this formal difficulty we only have to subdivide the rectangles or, which is the same, to subdivide the efficient representative (Ψ, Γ) by adding valency 2 vertices. A “natural” such subdivision is achieved if we declare as a valency 2 vertex each point $x \in \Psi^{-1}(V(\Gamma)) - V(\Gamma)$. This subdivision will be used throughout the paper and the resulting efficient representative is called a *subdivided representative* and is denoted by (Ψ_s, Γ_s) . From this definition it is obvious that all the entries of the incidence matrix $M(\Psi_s, \Gamma_s)$ are 0 or 1.

2.2. Periodic orbits

In the previous section, assuming that $f \in \text{Homeo}^+(S)$ was isotopic to a pseudo-Anosov homeomorphism φ , we described an algorithmic construction of a Markov partition for φ . The description of the periodic orbits is now classical using symbolic dynamics (see [Shu] for instance). The goal of this section is to make this description explicit and to relate the periodic orbits to an underlying efficient representative (Ψ, Γ) . Throughout this section we assume that an efficient representative is given and then also the subdivided representative as defined in the previous section.

The orbits of the pseudo-Anosov homeomorphism $\varphi \in [f]$ are divided into the following two classes:

- (1) The orbit $\mathcal{O}(\varphi, x)$ of a point $x \in S$, under φ is called *singular* if some point $y \in \mathcal{O}(\varphi, x)$ belongs to a stable half-leaf of $\mathcal{F}^s$ which ends at a singular point of the invariant foliations. Such half-leaves are called the *singular leaves*.
- (2) All the other orbits are called *regular*.

Lemma 2.1. — *A point $x \in S$ has a singular orbit under the pseudo-Anosov map φ if and only if its ω -limit set is a periodic orbit of singular points.*

Let us call $\mathcal{L}(s)$ a singular leaf which ends at a singular point s . If $\mathcal{O}(\varphi, x)$ is singular then there exists a point $y = \varphi^m(x)$ such that $y \in \mathcal{L}(s)$ for some singular point s . Since $\mathcal{F}^s$ is invariant under φ then every $\varphi^k(y)$ belongs to a singular leaf. Moreover $\mathcal{F}^s$ is the stable foliation and thus the $\varphi^k(y)$ converge to the periodic orbit $\mathcal{O}(\varphi, s)$. The converse is obvious. $\square$

The singular locus of $\mathcal{F}^s$ consists, by definition, of a collection of singular periodic orbits. All the other periodic orbits are, by lemma 2.1, regular orbits.

From the construction of the invariant foliations, given in the previous section, the singular periodic orbits are given by:

Lemma 2.2. — *The singular points of the invariant foliation $\mathcal{F}^s$ of the P.A. homeomorphism φ are in a one to one correspondance with the set of connected components of $S - N(\tau)$, given by the construction of § 2.1.*

This lemma is just an observation from the previous section. We also observe that the singular periodic orbits are completely determined by the efficient representative (Ψ, Γ) . In the remaining of the paper we shall mainly focus our attention to the regular orbits.

Let us recall that the Markov partition, constructed in § 2.1, have two kinds of rectangles:

- The real rectangles correspond to the real edges of the train track τ .
- The infinitesimal rectangles correspond to the infinitesimal edges of τ .

Lemma 2.3. — *Each point of a regular orbit is contained in a real rectangle.*

From the construction of $\mathcal{F}^s$, the singular points of the invariant foliations which are not punctures are obtained by collapsing the infinitesimal edges. Then these singular leaves come from the infinitesimal rectangles. By definition a regular orbit has no points on a singular leaf and therefore all its points are contained in real rectangles. $\square$

Let us consider the pair $(\mathcal{F}^s, \mathcal{F}^u)$ of invariant foliations of the pseudo-Anosov homeomorphism φ which is constructed from an efficient representative (Ψ, Γ) . We define the following equivalence relation.

Definition 2.4. — Two points x and y in S are *equivalent with respect to $\mathcal{F}^s$* , a relation denoted by $x \sim_s y$, if x and y belong to the same leaf of $\mathcal{F}^s$.

Lemma 2.5. — *Let S^* be the complement, in S , of the singular leaves of $\mathcal{F}^s$. Then there is a well-defined projection $\Pi_\Gamma : S^* \rightarrow \Gamma - V(\Gamma)$ given by $y = \Pi_\Gamma(x) \in \Gamma - V(\Gamma)$ if $x \sim_s y$.*

If x does not belong to a singular stable leaf, then, by lemma 2.3, x belongs to a real rectangle R_e for $e \in E(\Gamma)$. Then the stable leaf through x intersects the real edge e at a point $y \in \Gamma - V(\Gamma)$, proving the lemma. $\square$

Lemma 2.6. — *Let φ be a pseudo-Anosov homeomorphism on the surface S and let (Ψ, Γ) be an efficient representative of φ . If $\mathcal{O}(x, \varphi) = \{x, x_1, \dots, x_r\}$ is a regular periodic orbit of the pseudo-Anosov map φ , then $\Pi_\Gamma(\mathcal{O}(x, \varphi))$ is a collection of $r + 1$ distinct points on $\Gamma - V(\Gamma)$ which are permuted under the map Ψ .*

For proving this result we first observe that two distinct points x_i and x_j of $\mathcal{O}(x, \varphi)$ belong to distinct stable leaves. From lemma 2.5, the projection Π_Γ is well-defined for each point of $\mathcal{O}(x, \varphi)$. The above observation implies that $\Pi_\Gamma(x_i) \neq \Pi_\Gamma(x_j)$. $\square$

Let us now associate to a regular orbit $\mathcal{O}(x, \varphi)$ an ordered sequence of symbols, as usual in symbolic dynamics. In order to be more specific let (Ψ, Γ) be an efficient representative and let $E(\Gamma) = \{e_1, \dots, e_k\}$. We define the set of symbols $\Upsilon(\Gamma) = \{\varepsilon_1, \dots, \varepsilon_k\}$ by the one-to-one correspondence:

$$\varepsilon_i \in \Upsilon(\Gamma) \mapsto e_i \in E(\Gamma).$$

Let $x \in S$ be a point whose orbit, under φ , is regular. We define the *trajectory* $\Pi(x)$ to be the ordered sequence of symbols:

$$\Pi(x) = \dots \varepsilon_{i_0} \varepsilon_{i_1} \dots \varepsilon_{i_r} \dots,$$

where, for all $j \in \mathbf{Z}$, $\varepsilon_{i_j} = \varepsilon_m$ if $\Pi_\Gamma(f^j(x)) \in e_m$.

A sequence $\{\varepsilon_{i_j}\}_{j \in \mathbf{Z}}$ is *admissible* for φ if there exists a point $x \in S$ whose orbit is regular and such that $\{\varepsilon_{i_j}\}_{j \in \mathbf{Z}} = \Pi(x)$.

From the classical results in symbolic dynamics we have:

Lemma 2.7. — *Let $x \in S$ be a point whose orbit $\mathcal{O}(x, \varphi)$ is regular and let (Ψ, Γ) be an efficient representative for φ . Then the following properties are satisfied:*

1. *A symbolic code $\{\varepsilon_{i_j}\}_{j \in \mathbf{Z}}$ is admissible for φ if the (i_j, i_{j+1}) entry of the incidence matrix $M(\Psi, \Gamma)$ is not zero, for all pairs (i_j, i_{j+1}) .*
2. *A symbolic code $\{\varepsilon_{i_j}\}_{j \in \mathbf{Z}}$ is periodic if there exists a positive integer n such that $\varepsilon_{i_{j+n}} = \varepsilon_{i_j}$ for all $j \in \mathbf{Z}$. For a given admissible periodic code $\{\varepsilon_{i_j}\}_{j \in \mathbf{Z}}$, there exists one and only one regular periodic orbit $\mathcal{O}(x, \varphi)$ such that $\{\varepsilon_{i_j}\}_{j \in \mathbf{Z}} = \Pi(x)$ (up to a cyclic permutation).*

Recall that the efficient representative (Ψ, Γ) we consider in this section is a subdivided representative, i.e. its incidence matrix $M(\Psi, \Gamma)$ has only entries 0 or 1. Therefore the classical results of symbolic dynamics applies here (see for instance [Shu], Chap. 10). The uniqueness statement (2) above is not completely classical. Indeed, in general for a Markov partition, non uniqueness may occur if two rectangles have a common boundary. This ambiguity is removed here since we consider open rectangles obtained from the open edges of $\Gamma - V(\Gamma)$. $\square$

At this point we have obtained that each regular periodic orbit $\mathcal{O}(x, \varphi)$ of the pseudo-Anosov map φ defines a unique periodic symbolic code or an ordered collection of points $\Pi_\Gamma(\mathcal{O}(x, \varphi))$ on the graph Γ . The ordering of the points of $\Pi_\Gamma(\mathcal{O}(x, \varphi))$ is induced by the map $\Psi : \Gamma \rightarrow \Gamma$.

There is another classical tool from symbolic dynamics which enables one to keep track of the periodic orbits (see for instance [BGM]). If M is a $k \times k$ non-negative integer matrix we define the *Markov Graph* $\mathcal{MG}(M)$ as the oriented graph with k vertices $\{w_1, \dots, w_k\}$ and with r arrows from w_i to w_j if the entry $m_{i,j}$ of the matrix M is r . From this definition some properties are obvious, such as:

- $\mathcal{MG}(M)$ is transitive if and only if the matrix M is irreducible.
- If M is the transition matrix of a certain Markov partition of a map f , then the periodic orbits of f are in one-to-one correspondence with the closed oriented paths in $\mathcal{MG}(M)$.

2.3. The case of the disk

In this section we consider the particular case where the surface is the disk D^2 . Let $P = \mathcal{O}(x, f)$ be a periodic orbit of $f \in \text{Homeo}^+(D^2)$ and let $[f|_{D^2-P}]$ be the isotopy class of f relative to P . We noticed, in the introduction, the well-known relationship between $[f|_{D^2-P}]$ and an element β in the braid group B_n , where $n = \text{Card}(P)$. In fact the isotopy class $[f|_{D^2-P}]$ defines a braid β modulo the center Z_n of B_n (see [Bi] for more details). Conversely a class $\{\beta\}$ in B_n/Z_n defines a unique isotopy class $[\varphi_\beta] = [f|_{D^2-n \text{ points}}]$. Recall also that a braid $\beta \in B_n$ induces a permutation $s(\beta) \in S_n$. If P is a periodic orbit then the class $[f|_{D^2-P}]$ defines a braid β whose induced permutation $s(\beta) \in S_n$ has a single cycle. From a link point of view, if $s(\beta)$ has a single cycle then the closure $\bar{\beta}$ of the braid β is a knot in S^3 .

As in the previous sections we assume that the class $[f|_{D^2-P}]$ is pseudo-Anosov, we shall say in this case that the corresponding braid β is pseudo-Anosov. From theorem 1.2, there exists an efficient representative for φ_β . Among the efficient representatives of a given homeomorphism, some are easier to work with. From [Lo2] we obtain:

Lemma 2.8. — *Let φ_β be a pseudo-Anosov element in $\text{Homeo}^+(D^2 - n \text{ points})$ which is given by a braid $\beta \in B_n$ whose induced permutation $s(\beta) \in S_n$ has a single cycle. Then there exists an efficient representative $(\Psi_\beta, \Gamma_\beta)$ for φ_β satisfying the following properties:*

1. Γ_β has n edges $B(\Gamma_\beta) = \{e_1, \dots, e_n\} \subset E(\Gamma_\beta)$ whose terminal and initial vertices coincide. Therefore each $e_i \in B(\Gamma_\beta)$ is a simple closed curve in D^2 which bounds a once punctured disk of $D^2 - n \text{ points}$. These edges are called *boundary edges*.
2. Each vertex $v_j = i(e_j) = t(e_j)$, $e_j \in B(\Gamma_\beta)$, has valency three and is called a *boundary vertex*.
3. The complement $T = \Gamma_\beta - B(\Gamma_\beta)$ of the boundary edges is a tree embedded in D^2 .
4. The boundary edges of $B(\Gamma_\beta)$ are permuted under Ψ_β in a single cycle. This permutation, with a suitable labeling, is the permutation $s(\beta)$.
5. All the vertices $v \in V(\Gamma_\beta)$ are permuted under Ψ_β . This permutation has, at least, two cycles one of which is the cycle of the boundary vertices.

An efficient representative satisfying the properties (1)-(5) above is called a *standard representative*, similarly a graph satisfying the properties (1)-(3) above is also called *standard*.

From the property (4) the incidence matrix $M(\Psi_\beta, \Gamma_\beta)$ has a triangular bloc structure:

$$\begin{pmatrix} M_0(\Psi_\beta, \Gamma_\beta) & X \\ 0 & A \end{pmatrix}$$

where the diagonal bloc A corresponds to the boundary edges. Therefore it is a permutation matrix. The other bloc $M_0(\Psi_\beta, \Gamma_\beta)$, corresponding to the edges of the tree T , is a Perron-Frobenius matrix.

From property (2) there is a unique edge t_i in the tree T which belongs to $\text{St}(v_i)$ for each boundary vertex v_i . The edges $\{t_1, \dots, t_n\}$ are called the *terminal edges* of Γ_β . The other edges are called *middle edges*.

Notice that from the triangular bloc structure discussed above, the boundary edges have zero transposed measure. Therefore from the construction of the invariant foliations described in section 2.1, the boundary edges, which bound n punctured disks, are collapsed to the punctures. This gives rise to n distinct one-prong singularities of the foliations (see [FLP]).

This observation implies that for the symbolic description of the periodic orbits given in the previous section we can consider the boundary edges as infinitesimal edges. Therefore the set of symbols we need to describe the regular orbits are in a one-to-one correspondence with the set of edges of the tree T . Another way of stating this property is:

Lemma 2.9. — *Let $\beta \in B_n$ be a pseudo-Anosov braid and let φ_β be the pseudo-Anosov element in the class $[\varphi_\beta]$. Let $(\Psi_\beta, \Gamma_\beta)$ be a standard representative of φ_β . With the above notation the regular periodic orbits of φ_β are in one-to-one correspondence with the closed oriented paths in the Markov graph $\mathcal{MG}(M_0(\Psi_\beta, \Gamma_\beta))$.*

The next property will be used later on. It depends on the particular form of the standard representatives.

Lemma 2.10. — *Let $\beta \in B_n$ be a pseudo-Anosov braid whose induced permutation has a single cycle and let $(\Psi_\beta, \Gamma_\beta)$ be a standard representative for the pseudo-Anosov map φ_β defined by β . For every terminal edge $\{t_1, \dots, t_n\}$ of Γ_β , which is oriented in such a way that the initial vertex $i(t_j)$ is the boundary vertex v_j , the transversal words have the following form:*

$$\Psi_\beta^\perp(t_j) = A_j \cdot t_{\sigma(j)} \cdot A_j^{-1},$$

where σ is the permutation of the n boundary edges (i.e. the permutation $s(\beta)$) and A_j is a word.

The proof of this lemma is clear from the definitions of the transversal words (see § 1.3) and of the standard representatives (see figure 6).

FIG. 6. — Transversal word for a standard representative

3. From a symbolic code to a braid type

The goal of this section is to express explicitly the braid type of a regular periodic orbit, once we know an admissible periodic code with respect to a given subdivided representative. The braid type is obtained via a branched surface which we construct below.

The standard definitions and results about branched surfaces can be found for instance in [W] and [Oe]. Recall that a branched surface B is a finite 2 dimensional CW complex whose set of non manifold locus $\text{bl}(B)$, called the *branch locus*, is a graph. The connected components of $B - \text{bl}(B)$ are called the *sectors* of B . In general some restrictions are imposed to a branched surface. We say that a branched surface B is *embedded* in a 3-manifold M^3 if the sectors are disjointly embedded in M^3 and the connected component of the branch locus are disjoint embedded graphs. The usual definition imposes also that a branched surface has a well-defined tangent bundle (even at the branch locus). In what follows we shall use another restriction:

Definition 3.1. — A *branched surface with bands* in a 3-manifold M^3 is a branched surface which is embedded in M^3 and such that the closure of each sector is homeomorphic to a rectangle $[0, 1] \times [0, 1]$, called a *band*. Furthermore each band has two distinguished sides $\{0\} \times [0, 1]$ and $\{1\} \times [0, 1]$ which are respectively the *top* and the *bottom* of the band. At each point of the branch locus the tangent plane has two sides. On one side k bands are identified by their bottom sides. On the other side only one band starts from its top side (see figure 7).

An oriented path p , embedded in a band $b_i \simeq [0, 1] \times [0, 1]$ is *properly oriented* if p connects $\{0\} \times \{u\}$ to $\{1\} \times \{v\}$ and is transverse to each leaf $\{x\} \times [0, 1]$,

$x \in (0, 1)$. An oriented knot K in M^3 is *carried* by a branched surface with bands B if there is an ambient isotopy in M^3 which transforms K to a knot K_1 embedded in B , and such that the restriction of K_1 to each band is a collection (maybe empty) of properly oriented paths. The main result of this section is:

Theorem 3.2. — *Let β be a braid in B_n such that the corresponding isotopy class $[\varphi_\beta]$ is pseudo-Anosov. The closure $\bar{\beta}$ of β is supposed to be a knot in the solid torus $\mathbf{T}$ and let $M_\beta^3 = \mathbf{T} - \bar{\beta}$. Then there exists a branched surface with bands B_β in M_β^3 which satisfies:*

1. *The branch locus $\text{bl}(B_\beta)$ consists of disjoint intervals $I_1, \dots, I_N$ which are embedded in a single fiber $D^2 \times \{1\}$.*
2. *Each band of B_β is embedded in M_β^3 transversally to each fiber $D^2 \times \{\theta\}$, $\theta \in S^1$.*
3. *The knot $\bar{\beta}$ is contained in the boundary of B_β .*
4. *Let $[\gamma]$ be an element in the genealogy set $\mathcal{G}(\beta)$; then there exists a knot K in M_β^3 which is carried by B_β such that $K = \bar{\delta}$, with $\delta \in [\gamma]$.*
5. *Conversely let K be a knot in M_β^3 which is carried by B_β ; then K is a closed braid $K = \bar{\delta}$ such that $[\delta] \in \mathcal{G}(\beta)$.*

In the remainder of this section we shall give a constructive proof of this theorem. The construction relates the branched surface B_β to a subdivided standard representative (Ψ_s, Γ_s) of φ_β . Moreover we define a symbolic coding of the knots carried by B_β and we prove that this coding is in one-to-one correspondence with the symbolic coding of the periodic orbits, relative to a standard subdivided representative (Ψ_s, Γ_s) .

The above theorem is formulated in the case where the surface is a disk and the corresponding 3-manifold is the solid torus. But the arguments are exactly the same for any surface and the corresponding fiber bundle over the circle. We also assume that the closed braid $\bar{\beta}$ has one component. This is another simplification of the formulation which is sufficient for our needs but is not crucial for the result.

Let us consider the pseudo-Anosov element φ_β in $[f_\beta]$ and let (Ψ, Γ) be a standard representative for φ_β . Let us fix a suspension Φ_β^t of φ_β , in the solid torus $\mathbf{T}$ so that if $\{x_1, \dots, x_n\}$ are the marked points in $D^2 \times \{1\}$ then their orbits, under Φ_β^t , form the closed braid $\bar{\beta}$ in $\mathbf{T}$. We can visualize the 3-manifold M_β^3 as the solid cylinder $D^2 \times [0, 1]$ minus the strands of the braid β , with the identification of $D^2 \times \{0\}$ and $D^2 \times \{1\}$ by the identity map. We parametrize the flow Φ_β^t in such a way that $\Phi_\beta^1(x_0) \in D^2 \times \{1\}$ for all $x_0 \in D^2 \times \{0\}$. The graph Γ is embedded in the disk $D^2 \times \{0\}$ and we choose Φ_β^t so that $\Phi_\beta^1(\Gamma) = f_\beta(\Gamma) \subset N(\Gamma)$ in $D^2 \times \{0\} = D^2 \times \{1\}$, where f_β is the embedding given by the efficient representative (Ψ, Γ) (see § 1.3). Then the set

$$B_\Gamma = \{ \Phi_\beta^t(\Gamma), t \in [0, 1] \}$$

is a two dimensional CW complex in $(D^2 \times [0, 1] - \beta)$. Its branch locus is the suspension of the vertices of Γ , i.e. $\gamma_\Psi = \{ \Phi_\beta^t(V(\Gamma)); t \in [0, 1] \}$.

For a standard representative (Ψ, Γ) , the vertices $V(\Gamma)$ are permuted under Ψ

by lemma 2.8. Then the set γ_Ψ is a geometric braid in $(D^2 \times [0, 1] - \beta)$. The two-complex B_Γ has a boundary component in each boundary disk $D^2 \times \{0\}$ and $D^2 \times \{1\}$. Moreover the strands of β are enclosed in the suspension of the boundary edges $\{\Phi_\beta^t(B(\Gamma)); t \in [0, 1]\}$, which form a collection of braided cylinders. If we collapse each boundary edge to the enclosed puncture, then we get $\Gamma_0 = \Gamma - B(\Gamma)$, which is a tree whose end-points are the punctures. In this case $B(\Gamma_0) = \{\Phi_\beta^t(\Gamma_0); t \in [0, 1]\}$ has the strands of β as boundary components. If $V'(\Gamma_0)$ are the vertices of Γ_0 which are not end-points then $\gamma_0 := \{\Phi_\beta^t(V'(\Gamma_0)); t \in [0, 1]\}$ is a braid. In the disk $D^2 \times \{1\}$ we can apply the equivalence relation 2.4. Since $\Phi_\beta^1(\Gamma) = f_\beta(\Gamma)$, $\Phi_\beta^1(\Gamma_0) \sim_s \Gamma_0$. Therefore we can identify the image of the graph Γ_0 in the disk $D^2 \times \{1\}$ with the graph Γ_0 in the disk $D^2 \times \{0\}$, under the relation $\sim_s$.

Lemma 3.3. — *Let $\bar{B}_\Gamma$ be obtained by:*

- (i) *Identifying $\Phi_\beta^1(\Gamma_0)$ with Γ_0 , under the equivalence relation $\sim_s$.*
- (ii) *Identifying $D^2 \times \{0\}$ with $D^2 \times \{1\}$.*

Then $\bar{B}_\Gamma$ is a two-complex embedded in M_β^3 which satisfies:

- (1) *The knot $\bar{\beta}$ is the boundary of $\bar{B}_\Gamma$.*
- (2) *The branch locus of $\bar{B}_\Gamma$ is contained in*

— the closed braid $\bar{\gamma}_0$, and

— the graph Γ_0 which is embedded in a single fiber $D^2 \times \{0\} \simeq D^2 \times \{1\}$.

The proof of the above lemma is clear from the above observations. The construction of $\bar{B}_\Gamma$ is essentially the same than the original construction pioneered by Williams in [W]. There are nevertheless some slight differences with the usual branched surfaces which can be found in the litterature. For instance, $\bar{B}_\Gamma$ has no well-defined tangent bundle. The branch locus are not generic, i.e. there are generally more than three 2-cells arriving at a branching line. At this point we do not yet have the branched surface B_β of theorem 3.2.

In order to complete the construction we consider the following:

Definition 3.4. — *Let $\bar{B}_\Gamma$ be the 2-complex of lemma 3.3 and let B_β be the open branched surface obtained from $\bar{B}_\Gamma$ by the following operations:*

- (1) *Cut open $\bar{B}_\Gamma$ along each component of the closed braid $\bar{\gamma}_0$ (see lemma 3.3).*
- (2) *Cut open $\bar{B}_\Gamma$ along the finite collection of arcs:*

$$\mathcal{A}_\beta = \{\Phi_\beta^t(x); t \in [0, 1) \text{ and } x \in \Psi^{-1}(V(\Gamma))\}.$$

- (3) *At each point x in the interior of the branch locus, i.e. in*

$$\Gamma_0 - \Psi^{-1}(V(\Gamma_0)) \subset \bar{B}_\Gamma \cap D^2 \times \{1\},$$

we define a tangent plane generated by a vector tangent to the edge e which contains the point $x \in D^2 \times \{1\} \cap \Gamma_0$ and a vector induced by the flow Φ_β^t at this point (transverse to the fiber).

We shall now prove that B_β is the branched surface we are looking for. In order to help the intuition the figure 7 shows the construction of the branched surface, starting from the braid.

(1) B_β is a branched surface with bands.

By cutting along the closed braid $\bar{\gamma}_0$ we obtain a 2-complex whose intersection with each fiber $D^2 \times \{0\}$ of the solid torus is a collection of intervals (see figure 7). Furthermore the intersection: $\mathcal{A}_\beta \cap (D^2 \times \{0\})$ is exactly the subdivision points which gives the subdivided representative (Γ_s, Ψ_s) (see § 2.1). The open edges of Γ_s are a collection $\{I_1, \dots, I_N\}$ of intervals in the fiber $D^2 \times \{0\}$. The suspension of these intervals gives the expected collection of bands which are transverse to each fiber. The branch locus is reduced to the collection of intervals $\{I_1, \dots, I_N\}$ which are embedded in a single fiber. The “smoothing” of definition 3.4-3 defines the tangent space at the branch locus (see figure 7). From the construction, the properties (1) and (2) of theorem 3.2 are satisfied. The property (3) is also clear since the knot $\bar{\beta}$ was the boundary of $\bar{B}_\Gamma$.

It remains to prove the properties (4) and (5). We already observed that the branching intervals are exactly the edges of the graph Γ_s of the subdivided representative. Therefore these intervals are also the real edges of the train track τ which has been constructed in § 2.1.

Let $[\gamma]$ be an element in $\mathcal{G}(\beta)$. From the property (*) of the introduction there exists a periodic orbit Q of the pseudo-Anosov map φ_β such that $BT(Q, \varphi_\beta) = [\gamma]$. Let K_γ be the knot in M_β^3 which is obtained as the suspension of Q , i.e.:

$$K_\gamma = \{ \Phi_\beta^t(x_0); x_0 \in Q \text{ and } t \in \mathbf{R} \}.$$

From the results of section 2.2 each regular periodic orbit of φ_β has a unique symbolic code with respect to the subdivided representative (Γ_s, Ψ_s) . By lemma 2.6 the points of Q have a well-defined projection on Γ_s via Π_Γ . The projection Π_Γ is obtained by identifying the points on the same (regular) leaf of $\mathcal{F}^s$. Let $\mathcal{L}[x_0, y_0]$ be the segment of the stable foliation which connects $x_0 \in Q$ to $y_0 \in \Gamma_s$ and let I_j be the edge of Γ_s which contains y_0 . Let b_j be the band of B_β which is the suspension of the interval I_j . Consider now the arc $\{ \Phi_\beta^t(x_0); t \in [0, 1) \}$. For each $t \in [0, 1)$ the image of the segment $\mathcal{L}[x_0, y_0]$ under Φ_β^t is a segment in the fiber $D^2 \times \{t\}$ which connects the point $\Phi_\beta^t(x_0)$ to a point y_t in the band b_j . At $t = 1$ we are back to the initial fiber $D^2 \times \{0\}$ and $\Phi_\beta^1(x_0) = x_1 \in Q$. Since the points of Q project to distinct points on Γ_s , under Π_Γ , the previous construction can be iterated. If the period of Q is p , then, after p iterations, we are back to the point $x_0 \in Q$ and the corresponding point y_0 in B_β . From the initial knot $K_\gamma = \{ \Phi_\beta^t(x_0); x_0 \in Q \text{ and } t \in [0, p] \}$ we have constructed a knot K on the branched surface B_β given by

$$K = \{ \Phi_\beta^t(\mathcal{L}[x_0, y_0]) \cap B_\beta; t \in [0, p] \}.$$

FIG. 7. — Construction of a branched surface

Let us check that the knot K is carried by B_β . By construction there is an embedded annulus in M_β^3 whose boundaries are K_γ and K . Therefore the two knots are isotopic in M_β^3 . In order to complete the proof of the property (4) of theorem 3.2, we just have to check that the restriction of the knot K with each band of the branched surface is properly oriented. This property is clear since the knot K has an orientation which is induced from the flow and that the bands are also oriented by the flow from the top to the bottom.

In order to prove the last property (5) we first introduce some terminology. Let K be a knot in M_β^3 which is carried by B_β . Then there is a knot K_1 on B_β and an isotopy which transforms K to K_1 . By definition K_1 is transverse to each fiber $D^2 \times \{t\}$, i.e. K_1 is a closed braid $\bar{\gamma}_1$, where $\gamma_1 \in B_p$. Let us fix one point $y_0 \in K_1 \cap D^2 \times \{0\}$ and assume that y_0 belongs to the branching interval I_j . Starting from y_0 , let $b_{j_1}, b_{j_2}, \dots, b_{j_p+1} = b_{j_1}$ be the ordered collection of bands which contain the knot K_1 . The ordering being given by the orientation of the knot.

Now we associate to the knot K_1 the following symbolic code:

$$\Upsilon(K_1) = (\varepsilon_{j_1} \varepsilon_{j_2} \dots \varepsilon_{j_p})^\infty,$$

where the notation $(W)^\infty$ means that the word W is infinitely repeated.

Lemma 3.5. — *Let K be a knot in M_β^3 which is carried by the branched surface B_β and let $\Upsilon(K_1)$ be a corresponding symbolic code defined above. Then $\Upsilon(K_1)$ is an admissible symbolic code with respect to the subdivided representative (Γ_s, Ψ_s) from which B_β has been constructed.*

By construction, the bands $b_1, \dots, b_N$ of B_β are in one-to-one correspondence with the open edges of the graph Γ_s of the subdivided representative. If $(\varepsilon_{j_i} \varepsilon_{j_{i+1}})$ is a pair of consecutive symbols in $\Upsilon(K_1)$, then there exists a band in B_β which connects the branching interval I_{j_i} with $I_{j_{i+1}}$. By construction, this implies that $e_{j_{i+1}} \subset \Psi_s(e_{j_i})$. Therefore the entry (j_i, j_{i+1}) of the incidence matrix $M(\Gamma_s, \Psi_s)$ is not zero. This completes the proof of the lemma by lemma 2.7. $\square$

From lemma 3.5 and 2.7 there exists a unique regular periodic orbit Q of the pseudo-Anosov map φ_β whose symbolic code is $\Upsilon(K_1)$. Let K_0 be the knot in M_β^3 which is the suspension of Q under Φ_β^t . From the proof of the property (4) of theorem 3.2, K_0 is carried by B_β and from the previous arguments this knot is isotopic to K_1 .

Each knot K, K_0, K_1 is a closed braid, respectively $\bar{\gamma}, \bar{\gamma}_0, \bar{\gamma}_1$. By the property (*) of the introduction $[\gamma_0] \in \mathcal{G}(\beta)$. By assumption, $\bar{\gamma}$ and $\bar{\gamma}_1$ are isotopic and by construction $\bar{\gamma}_1$ and $\bar{\gamma}_0$ are also isotopic. Therefore $\bar{\gamma}$ and $\bar{\gamma}_0$ are isotopic in M_β^3 and then in the solid torus $\mathbf{T}$. From a well-known result of Morton [Mo] the braids γ and γ_0 are conjugated in B_p . This completes the proof of theorem 3.2. $\square$

The construction of the branched surface B_β is given here via the suspension flow Φ_β^t . This construction is not algorithmic since the flow is not given by combinatorial data. An algorithmic construction of B_β is in fact possible. Such an algorithm is given in [GL].

4. Characterization of the forcing

The goal of this section is to find an effective criterion which enables one to decide whether two braid types are related or not by the forcing relation. More precisely let $P = \{x_1, \dots, x_p\} \subset S$ and $Q = \{y_1, \dots, y_q\} \subset S$ be two distinct periodic orbits of $f \in \text{Homeo}^+(S)$. From the homeomorphism f and the two periodic orbits P and Q we construct four different isotopy classes, namely:

$$\varphi = [f] \in \text{Mod}_0(S), \quad \gamma = [f|_{S-P}] \in \text{Mod}_0(S - p \text{ points}),$$

$$\delta = [f|_{S-Q}] \in \text{Mod}_0(S - q \text{ points})$$

and $\beta = [f|_{S-P-Q}] \in \text{Mod}_0(S - (p + q) \text{ points}).$

From a 3-manifold point of view these four classes define, via the mapping torus, four distinct manifolds. The relationship between these 3-manifolds is clear. Indeed β defines $M_{\varphi, \beta}^3$ which is the complement of a two-components link in the fibered bundle M_φ^3 . We obtain $M_{\varphi, \gamma}^3$ and $M_{\varphi, \delta}^3$ from $M_{\varphi, \beta}^3$ by forgetting one of the two components (zero surgery). In the same way we define the *forgetting maps*:

$$\begin{array}{ccc} \beta \in \text{Mod}_0(S - (p + q) \text{ points}) & & \\ \downarrow f_P & & \downarrow f_Q \\ \delta \in \text{Mod}_0(S - q \text{ points}) & \gamma \in \text{Mod}_0(S - p \text{ points}) & \end{array}$$

The question we want to address in this section is:

Can we decide whether or not the braid types $[\gamma]$ and $[\delta]$ are related by the forcing relation?

Let $S_n = S - \{n \text{ points}\}$ be a punctured surface and let $\Gamma \in G(S_n)$, i.e. Γ is an embedded graph in S_n whose fundamental group is isomorphic to $\pi_1(S_n)$. Each of the n punctures of S_n is contained in a connected component of $S_n - \Gamma$ which is a once punctured disk $\Delta_1, \dots, \Delta_n$. For a given disk Δ_i we consider the collection $\{e_i^1, \dots, e_i^r\}$ of edges of $E(\Gamma)$ which bound Δ_i . These edges are called the *boundary edges* and the disk Δ_i is called an *r-gon*.

In order to simplify the formulation, let us again restrict to the case of the punctured disk.

Let $\beta \in B_n$ be a braid whose induced permutation $s(\beta)$ has two cycles, i.e. the closed braid $\bar{\beta}$ is a two-components link in S^3 (or in $\mathbf{T}$). The forgetting maps f_P and f_Q defined above induce the following diagram:

$$\delta \in B_q \xleftarrow{\tilde{f}_P} \beta \in B_{p+q} \xrightarrow{f_Q} \gamma \in B_p.$$

Theorem 4.1. — *Let $[\gamma] \in B_p$ and $[\delta] \in B_q$ be two braid types whose induced permutation has a single cycle. Let us assume that γ is a pseudo-Anosov braid; then $[\gamma] \stackrel{reg}{>} [\delta]$ if and only if the following properties hold:*

1. *There exists a braid $\beta \in B_{p+q}$ whose induced permutation has two cycles and such that the two forgetting maps $\tilde{f}_P$ and $\tilde{f}_Q$ defined above give respectively δ and γ , where P and Q are the two collections of punctures which are permuted under $\varphi \in [\varphi_\beta]$.*
2. *There exists an efficient representative $(\Psi_\beta, \Gamma_\beta)$ for φ_β such that each puncture q_i of Q is enclosed in a digon disk Δ_i^2 of $D^2 - P - Q$, whereas each puncture p_j of P is enclosed in a monogon disk Δ_j^1 .*
3. *The incidence matrix $M(\Psi_\beta, \Gamma_\beta)$ has a bloc structure with a permutation matrix bloc corresponding to the boundary edges $B(\Gamma_\beta)$ and a Perron-Frobenius bloc corresponding to the other edges of Γ_β .*

As we noted in section 2.2, we are mainly concerned with the regular periodic orbits of a pseudo-Anosov homeomorphism. For the forcing relation problem these regular orbits, via the property (*) of the introduction, are the whole set $\mathcal{G}(\beta)$ except the finite set of singular periodic orbits given by the singularities of the invariant foliations.

The notation $[\gamma] \stackrel{reg}{>} [\delta]$ used in the theorem means that the periodic orbit of braid type $[\delta]$ is regular with respect to the pseudo-Anosov map φ_γ .

Let us assume that $[\gamma] \stackrel{reg}{>} [\delta]$ and consider the class $[\varphi_\gamma] \in \text{Mod}_0(D^2 - p \text{ points})$ which is supposed to be pseudo-Anosov. Let φ_γ be the pseudo-Anosov element in the class. By the property (*) of the introduction φ_γ has a periodic orbit Q of braid type $[\delta]$. Let us denote by $[\varphi_\beta] \in \text{Mod}_0(D^2 - (p + q) \text{ points})$ the isotopy class obtained from φ_γ by piercing the surface $D^2 - P$ at the points of the periodic orbit Q . By lemma 2.8 the class $[\varphi_\gamma]$ admits an efficient representative $(\Psi_\gamma, \Gamma_\gamma)$ such that each puncture of P is enclosed in a monogon disk. Furthermore the property (3) is satisfied by the incidence matrix $M(\Psi_\gamma, \Gamma_\gamma)$. By lemma 2.6, since the periodic orbit Q of φ_γ is regular, the projection $\Pi_{\Gamma_\gamma}(Q) = \tilde{Q}$ is a well-defined collection of q distinct points which are permuted under Ψ_γ . Furthermore, by the property (4) of lemma 2.8 and the regularity assumption, the points of $\tilde{Q}$ are contained in the interior of the edges of the tree $T_\gamma = \Gamma_\gamma - B(\Gamma_\gamma)$.

The piercing operation of a topological representative.

Let (Ψ, Γ) be a topological representative of a pseudo-Anosov class $[\varphi]$ in $\text{Homeo}^+(S)$ and let $\tilde{Q} = \{x_1, \dots, x_q\}$ be a collection of points on $\Gamma - V(\Gamma)$ which are periodic under Ψ . The subdivision which is obtained by replacing each point x_i of $\tilde{Q}$ by a valency 2 vertex v_i defines a new topological representative (Ψ_q, Γ_q) . If (Ψ, Γ) is efficient then (Ψ_q, Γ_q) is also efficient. The piercing operation is obtained by removing the points of $\tilde{Q}$ from S . We obtain a punctured surface $S - \tilde{Q}$, a new isotopy class $[\varphi|_{S - \tilde{Q}}]$ and a disconnected graph $\Gamma_q - \tilde{Q}$. In order to define a topological representative

for $[\varphi|_S - \tilde{Q}]$ we reconnect the graph by replacing each $v_i \in \tilde{Q}$ by a digon disk Δ_i . This disk Δ_i is bounded by a pair of boundary edges $\{d_i^1, d_i^2\}$. This operation defines a graph $\tilde{\Gamma}_q \in G(S - \tilde{Q})$ from the graph $\Gamma \in G(S)$. We have now to define a map $\tilde{\Psi}_q$ on this graph. Let $D(\tilde{\Gamma}_q) = \bigcup_{i=1}^q \{d_i^1, d_i^2\}$ be the set of boundary edges which have been added by the piercing operation. The edges of $E(\tilde{\Gamma}_q) - D(\tilde{\Gamma}_q)$ are in a one-to-one correspondence with the edges of $E(\Gamma_q)$ and we denote by the same name the corresponding edges. We define the map $\tilde{\Psi}_q$ on $\tilde{\Gamma}_q$ by:

$$\begin{aligned} \tilde{\Psi}_q(e_j) &= (\Psi_q(e_j))^* \quad \text{for } e_j \in E(\tilde{\Gamma}_q) - D(\tilde{\Gamma}_q), \\ \tilde{\Psi}_q(d_i^j) &= d_{\sigma(i)}^j \quad \text{for } d_i^j \in D(\tilde{\Gamma}_q), \end{aligned}$$

where $\sigma(i)$ is the permutation on the indices $\{1, \dots, q\}$ given by the permutation of the vertices $\{v_1, \dots, v_q\}$ induced by the map Ψ_q . In order to define the operation $(X)^*$ used above, let us consider the two edges $\{e_i^1, e_i^2\}$ which are incident at the valency two vertex v_i of Γ_q . The edge path $(X)^*$ on the graph $\tilde{\Gamma}_q$ is obtained from the edge path X on the graph Γ_q by inserting, in X , the appropriate $d_i^j \in D(\tilde{\Gamma}_q)$ each time X crosses the turn $\{e_i^1, e_i^2\}$ in Γ_q .

In order to be more explicit let us first observe that the operation is well defined because Γ and its image under f_3 are embedded in the surface. The map f_3 is the embedding which induces the efficient representative (Ψ, Γ) as defined in § 1.3. The operation is easier to express via the transversal words. Let us assume that $\Psi(x_{i-1}) = x_i$ with $x_{i-1} \in e'$ and $x_i \in e$. The transversal word at the edge e has the form $\Psi^\perp(e) = A \cdot e' \cdot B$. After the subdivision at x_i and x_{i-1} (among the other periodic points) one has $e = \dots e_1 \cdot e_2 \dots$ and $e' = \dots e'_1 \cdot e'_2 \dots$ and the corresponding transversal words are

$$\Psi_q^\perp(e_1) = A^s \cdot e'_1 \cdot B^s \quad \text{and} \quad \Psi_q^\perp(e_2) = A^s \cdot e'_2 \cdot B^s,$$

where A^s and B^s are the words obtained from A and B by the subdivision operation. The piercing operation introduces the additional boundary edges d_i^j . With the above notation and an obvious choice of the orientations for the edges d_i^j one has the following transversal words:

$$\begin{aligned} \tilde{\Psi}_q^\perp(e_1) &= \Psi_q^\perp(e_1) \quad \text{and} \quad \tilde{\Psi}_q^\perp(e_2) = \Psi_q^\perp(e_2), \\ \tilde{\Psi}_q^\perp(d_i^2) &= A^s \cdot d_{i-1}^{r2} \quad \text{and} \quad \tilde{\Psi}_q^\perp(d_i^1) = d_{i-1}^{r1} \cdot B^s. \end{aligned}$$

The adaptation to the other possible orientations is obvious. From the transversal words one easily gets the representative $(\tilde{\Psi}_q, \tilde{\Gamma}_q)$ as explained above.

The next lemma is a direct consequence of the construction.

Lemma 4.2. — *The map $\tilde{\Psi}_q: \tilde{\Gamma}_q \rightarrow \tilde{\Gamma}_q$ is a topological representative of $[\varphi|_S - \tilde{Q}]$. Furthermore if the topological representative (Ψ_q, Γ_q) is efficient then $(\tilde{\Psi}_q, \tilde{\Gamma}_q)$ is also efficient. In addition the incidence matrix $M(\tilde{\Psi}_q, \tilde{\Gamma}_q)$ satisfies the property (3) of theorem 4.1.*

The first statement of the lemma is clear from the construction. The efficiency property is due to the fact that the piercing operation cannot produce non local injectivity. Indeed, none of the turns which have been created by the piercing is tangent under $\tilde{\Psi}_q$. The property of the incidence matrix comes directly from the construction. Indeed, the incidence matrix of the subdivided representative (Ψ_q, Γ_q) satisfies the property (3) of theorem 4.1. We have just added a bloc corresponding to the boundary edges of the digon disks which are permuted, so we obtain a larger permutation matrix bloc. $\square$

Let us now apply the piercing operation to a standard representative $(\Psi_\gamma, \Gamma_\gamma)$ for the braid type $[\gamma]$ with respect to the points $\tilde{Q} = \Pi_{\Gamma_\gamma}(Q)$. By definition the new isotopy class $[\varphi_{\gamma, \tilde{Q}}]$, where $S = D^2 - P$, is a braid type $[\beta]$ which satisfies item (1) of the theorem. The new topological representative $(\tilde{\Psi}_{\gamma_q}, \tilde{\Gamma}_{\gamma_q})$ is efficient by lemma 4.2 and satisfies all the properties of theorem 4.1. This completes the proof of the “only if” part of the theorem.

In order to prove the “if” part of the theorem we consider the class $[\varphi_\beta] \in \text{Mod}_0(D^2 - P - Q)$ and its efficient representative (Ψ, Γ) which satisfies the properties of the theorem. Let $\{\Delta_1^2, \dots, \Delta_q^2\}$ be the collection of digon disks which are permuted under Ψ . The forgetting map is obtained by closing an invariant collection of punctures. At the efficient representative level (Ψ, Γ) the effect of this *closing operation* is exactly the converse of the piercing operation. Indeed, by closing the punctures we obtain a collection of permuted disks $\{\Delta_1^2, \dots, \Delta_q^2\}$ (without punctures). By an isotopy, whose support is $\bigcup_{i=1}^q \Delta_i^2$, we transform each disk Δ_i^2 to a valency two vertex v_i . We obtain a topological representative (Ψ_q, Γ_q) with a collection $\tilde{Q} = \{v_1, \dots, v_q\}$ of valency two vertices which are permuted under Ψ_q . By a sequence of valency 2 isotopies we obtain a topological representative $(\tilde{\Psi}, \tilde{\Gamma})$ for the class $[\gamma]$.

Lemma 4.3. — With the previous notation, if the topological representative (Ψ, Γ) is efficient for the class $[\beta]$ then $(\tilde{\Psi}, \tilde{\Gamma})$ is an efficient representative for the class $[\gamma]$. Furthermore $(\tilde{\Psi}, \tilde{\Gamma})$ admits a periodic orbit $\tilde{Q}$ whose braid type is $[\delta]$.

As for the previous lemma 4.2 the proof of the efficiency is clear since none of the two homotopies above can create non local injectivity.

It remains to prove that the pseudo-Anosov map φ_γ has a periodic orbit Q of braid type $[\delta]$. From the closing operation the efficient representative (Ψ_q, Γ_q) has a periodic orbit $\tilde{Q}$ of valency two vertices. Under the valency 2 isotopy the points of $\tilde{Q}$ are transformed into a collection of points $\{\tilde{x}_1, \dots, \tilde{x}_q\}$ which are permuted under $\tilde{\Psi}$. This implies the existence of an admissible periodic symbolic code with respect to $(\tilde{\Psi}, \tilde{\Gamma})$. By lemma 2.7, there exists a regular periodic orbit $Q = \{x_1, \dots, x_q\}$ of the pseudo-Anosov map φ_γ . From the construction, the braid type of this periodic orbit is $[\delta]$, which completes the proof of the lemma and of theorem 4.1. $\square$

Theorem 4.1 is easy to understand from an invariant foliation point of view. Indeed let $(\mathcal{F}^s, \mathcal{F}^u)$ be the pair of invariant foliations for the pseudo-Anosov homeomorphism φ_γ . If $Q = \{x_1, \dots, x_q\}$ is a regular periodic orbit of φ_γ , then through each point x_i of Q there is a regular leaf $\mathcal{L}^s(x_i)$ (resp. $\mathcal{L}^u(x_i)$) of $\mathcal{F}^s$ (resp. $\mathcal{F}^u$). Piercing the surface at the points of Q defines a new pair of foliations $(\mathcal{F}_Q^s, \mathcal{F}_Q^u)$ which are invariant under $\varphi_{\gamma|_{S-Q}}$. Each puncture which has been created by the piercing is a 2-prong singularity of the foliations. These 2-prong singularities correspond exactly to the digon disks of the efficient representative discussed above.

5. Small cancellation orbits

At this stage most of the tools we need have been defined and we can start the study of the set $\mathcal{G}(\beta)$. The goal of this section is, roughly speaking, to “localize” the study around the point β . To this end we introduce a class of periodic orbits, with respect to a given efficient representative $(\Psi_\beta, \Gamma_\beta)$ for φ_β . Throughout this section we shall only consider the case of the punctured disk.

By lemma 2.8 we assume that $(\Psi_\beta, \Gamma_\beta)$ is a standard representative for which there is a decomposition: $E(\Gamma_\beta) = B(\Gamma_\beta) \cup T(\Gamma_\beta) \cup L(\Gamma_\beta)$, where

- $B(\Gamma_\beta)$ is the set of boundary edges,
- $T(\Gamma_\beta)$ is the set of terminal edges,
- $L(\Gamma_\beta)$ is the set of middle edges.

Once an efficient representative $(\Psi_\beta, \Gamma_\beta)$ is given, by the results of § 2.2, we obtain a symbolic description of each regular periodic orbit of the pseudo-Anosov map φ_β in term of the subdivided representative $(\Psi_{\beta,s}, \Gamma_{\beta,s})$ (see § 2.1). Recall also that the boundary edges play no role in the symbolic coding of the regular periodic orbits, by lemma 2.9. Let P be a regular periodic orbit of the pseudo-Anosov map φ_β , we say that P *visits* an edge $e_i \in E(\Gamma_\beta)$ if the projection $\tilde{P} = \Pi_{\Gamma_\beta}(P)$ contains a point $\tilde{x} \in e_i$.

Definition 5.1. — Let $\beta \in B_n$ be a pseudo-Anosov braid whose induced permutation has a single cycle and let $(\Psi_\beta, \Gamma_\beta)$ be a standard representative of $\varphi_\beta \in \text{Homeo}^+(D^2 - n \text{ points})$. An orbit P of the pseudo-Anosov map φ_β is called a *small cancellation orbit* with respect to $(\Psi_\beta, \Gamma_\beta)$ if P visits all the terminal edges of Γ_β .

The class of small cancellation orbits (we omit the reference to $(\Psi_\beta, \Gamma_\beta)$ which is fixed) is central in this study. Indeed the idea of a possible localization mentioned above is based on the following:

Theorem 5.2. — Let $\beta \in B_n$ be a pseudo-Anosov braid whose induced permutation has a single cycle and let $(\Psi_\beta, \Gamma_\beta)$ be a standard representative of φ_β . For any $[\gamma] \in \mathcal{G}(\beta)$ there exists a small cancellation periodic orbit P of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$ such that

$$[\beta] \succ [\delta] \succ [\gamma], \quad \text{where } [\delta] = \text{BT}(P, \varphi_\beta).$$

This theorem is the statement of property (1) of the main theorem. In order to fix the notation let us denote the boundary edges $B(\Gamma_\beta)$ by $\{b_1, \dots, b_n\}$ and the terminal edges $T(\Gamma_\beta)$ by $\{e_1, \dots, e_n\}$. We fix the orientation of the terminal edges in such a way that the initial vertex $i(e_j)$ for $j = 1, \dots, n$ is the boundary vertex. The orientation of the other edges $\{e_{n+1}, \dots, e_{n+r}\} = L(\Gamma_\beta)$ is left arbitrary. We also fix the labelling of the terminal edges so that

$$i(e_j) = i(b_j) = t(b_j) \quad \text{for } j = 1, \dots, n$$

and with $\Psi(b_j) = b_{j+1}$ for $j \neq n$ and $\Psi(b_n) = b_1$.

For the subdivided standard representative, each edge e_j of Γ_β is decomposed into a collection $\{e_j^1, \dots, e_j^k\}$ which is ordered by the orientation of the edge. This means, in particular, that for all the terminal edges one has $i(b_j) = i(e_j^1)$. Let us recall, from the definition of a standard representative, that the incidence matrix $M(\Psi_\beta, \Gamma_\beta)$ has a triangular bloc decomposition with a diagonal permutation bloc corresponding to the boundary edges and a Perron-Frobenius bloc $M_0(\Psi_\beta, \Gamma_\beta)$ corresponding to the other edges. By lemma 2.9 the symbolic coding of the regular periodic orbits of φ_β depends only on the bloc $M_0(\Psi_\beta, \Gamma_\beta)$. Recall also that, from the definition of the Markov graph $\mathcal{MG}(M_0)$, there is a one-to-one correspondence between the vertices $\{w_1, \dots, w_{n+r}\}$ of $\mathcal{MG}(M_0)$ and the edges $\{e_1, \dots, e_{n+r}\}$ of Γ_β . The same is true for the subdivided representative. We say that there is a *transition from e_i to e_j under Ψ_β* if there exists an arrow from w_i to w_j in $\mathcal{MG}(M_0)$. Such a transition is denoted by: $e_i \rightarrow e_j$. The difference between the Markov graph of an efficient representative compared with its subdivided representative is that at most one arrow connects two vertices in the last case and any number can occur in the first case.

Lemma 5.3. — *With the above notation, there is a closed path in $\mathcal{MG}(M_0)$ given by the periodic sequence of transitions*

$$(e_1 \rightarrow e_2 \rightarrow \dots \rightarrow e_n)^\infty.$$

For the subdivided representative $(\Psi_{\beta s}, \Gamma_{\beta s})$ the previous closed path is given by the following periodic sequence of transitions:

$$(e_1^1 \rightarrow e_2^1 \rightarrow \dots \rightarrow e_n^1)^\infty.$$

The map $\Psi_\beta: \Gamma_\beta \rightarrow \Gamma_\beta$ is continuous and the boundary edges are permuted under the sequence of transitions $b_1 \rightarrow b_2 \rightarrow \dots \rightarrow b_n \rightarrow b_1$. The continuity of the map Ψ_β completes the proof. $\square$

From the definition of the subdivided representative $(\Psi_{\beta s}, \Gamma_{\beta s})$ the following lemma is obvious:

Lemma 5.4. — *Let $(\Psi_{\beta s}, \Gamma_{\beta s})$ be the subdivided representative of $(\Psi_{\beta}, \Gamma_{\beta})$. If there exists a transition $e_i \rightarrow e_j$ for the representative $(\Psi_{\beta}, \Gamma_{\beta})$ and if e_i is subdivided into $\{e_i^1, \dots, e_i^k\}$ and e_j is subdivided into $\{e_j^1, \dots, e_j^l\}$ then there exists an edge e_i^m of $\Gamma_{\beta s}$ which admits the transitions*

$$e_i^m \rightarrow e_j^r \text{ for } r = 1, \dots, l.$$

In order to start the proof of theorem 5.2 we consider a periodic orbit Q of φ_{β} whose braid type is γ and let $\tilde{Q} = \Pi_{\Gamma_{\beta}}(Q)$ be the projection of Q on the graph Γ_{β} . We assume that Q does not visit all the terminal edges since otherwise the result is obvious.

Case 1.

Assume that the periodic orbit Q visits one and only one terminal edge. The existence of the regular periodic orbit Q implies the existence of a closed path in $\mathcal{MG}(M_0)$ which is given by a periodic sequence of transitions

$$(e_{i_1}^{m_1} \rightarrow e_{i_2}^{m_2} \rightarrow \dots \rightarrow e_{i_q}^{m_q})^{\infty}.$$

Since Q visits one terminal edge, there is one transition such that $e_{i_j}^{m_j}$ is a terminal edge. Let us assume that $i_j = 1$; if it is not the case, then we change the labelling. By lemma 5.4 there exists a transition: $e_{ij-1}^{m_{j-1}} \rightarrow e_{i_j}^1 = e_1^1$ and, by lemma 5.3, there exists a periodic sequence of transitions

$$e_1^1 \rightarrow e_2^1 \rightarrow \dots \rightarrow e_n^1.$$

By lemma 5.4 again there exists a transition $e_n^1 \rightarrow e_1^{m_j} = e_{i_j}^{m_j}$.

Therefore the following periodic sequence of transitions:

$$(*) \quad (e_{i_1}^{m_1} \rightarrow \dots \rightarrow e_{ij-1}^{m_{j-1}} \rightarrow e_1^1 \rightarrow e_2^1 \dots \rightarrow e_n^1 \rightarrow e_{i_j}^{m_j} \dots \rightarrow e_{i_q}^{m_q})^{\infty}$$

is admissible.

Let P be the periodic orbit of φ_{β} whose periodic code is given by the above periodic sequence of transitions. By construction P is a small cancellation orbit. Let us denote by $\tilde{P}$ the projection of P on the graph Γ_{β} . Let us now prove that the braid type $[\delta]$ of P forces the braid type $[\gamma]$ of Q . Let us also denote by R the collection of punctures whose braid type is $[\beta]$.

First step: The piercing operation at $\tilde{P} \cup \tilde{Q}$.

We have defined the piercing operation in § 4 for a single periodic orbit $\tilde{Q}$ of a topological representative (Ψ, Γ) . The same construction obviously works if it is applied to any finite collection of periodic orbits. By rephrasing lemma 4.2 we obtain:

Lemma 5.5. — *Let $(\Psi_{\beta}, \Gamma_{\beta})$ be a standard efficient representative for the class $[\varphi_{\beta}] \subset \text{Homeo}^+(D^2 - R)$. Let $\tilde{P}$ and $\tilde{Q}$ be two distinct regular periodic orbits of periods p (resp. q) for the map Ψ_{β} . Then the piercing operation at the points $\tilde{P} \cup \tilde{Q}$ defines an isotopy*

class $[\varphi_\alpha] \subset \text{Homeo}^+(D^2 - R - \tilde{P} - \tilde{Q})$ and an efficient representative $(\Psi_\alpha, \Gamma_\alpha)$ which satisfies the following properties:

- Each puncture r_i of the collection R is enclosed in a monogon disk Δ_i^R .
- Each puncture p_i of the collection $\tilde{P}$ is enclosed in a digon disk Δ_i^P .
- Each puncture q_i of the collection $\tilde{Q}$ is enclosed in a digon disk Δ_i^Q .

The goal, in the remainder of this section, is to prove that $[\delta] \succ [\gamma]$, where $\delta = \text{BT}(P, \varphi_\beta)$ and $\gamma = \text{BT}(Q, \varphi_\beta)$. To this end we apply theorem 4.1. The first problem is to find a braid η whose closure has two components, one of which corresponding to δ the other one corresponding to γ . Such a braid is easy to obtain from the class $[\varphi_\alpha]$ given by lemma 5.5. Indeed the collection of three periodic orbits $R \cup P \cup Q$ defines a braid $\alpha \in B_{n+p+q}$ whose induced permutation has three cycles or whose closure has three components. Let us denote this braid as $\alpha = \beta \cup_\beta \gamma \cup_\beta \delta$.

In order to compare the braid types $[\gamma]$ and $[\delta]$, using theorem 4.1, we consider the braid $\eta = \gamma \cup_\beta \delta \in B_{p+q}$ which is obtained from the braid α by the forgetting map $\tilde{f}_R$. To this end we apply the closing operation at the punctures of the collection R .

Second step: The closing operation at R .

We start from the efficient representative $(\Psi_\alpha, \Gamma_\alpha)$ given by lemma 5.5. We have already considered a closing operation in § 4, but it was on a collection of permuted digon disks. In this case the closing operation gave an efficient representative. In the present case we consider a closing operation on a collection of permuted monogon disks. As we will observe the effect is drastically different. Indeed, after the closing operation, the resulting topological representative is far from being efficient. Therefore we will have to apply the moves, described in § 1.2, in order to obtain a new efficient representative. The periodic orbit $\tilde{P}$ which induces the collection Δ_i^P of digon disks is a small cancellation periodic orbit. This class has been defined in such a way that these operations can be controlled.

FIG. 8. — Collapsing a monogon

By closing the punctures of the collection R we obtain a collection of permuted monogon disks $\tilde{\Delta}_i^R$ with no punctures.

(1) By a homotopy whose support is $\bigcup_{i=1}^n \tilde{\Delta}_i^R$ we collapse each $\tilde{\Delta}_i^R$ to a point $R_i \in D^2 - \tilde{P} - \tilde{Q}$. After these collapsing operations we obtain a topological representative (Ψ_1, Γ_1) for φ_η with a collection of valency one vertices which are permuted under Ψ_1 (see figure 8).

(2) By a sequence of valency one isotopies (see § 1.2) we remove the above collection of valency one vertices. In order to understand the effects of these valency one isotopies we study them, first on the graph Γ_1 and then on the image $\Psi_1(\Gamma_1)$.

(3) Valency one isotopy on the graph.

Since the periodic orbit P is small cancellation, there is a point of $\tilde{P}$ on each terminal edge of Γ_β and more precisely on each edge e_j^1 of the subdivided graph $\Gamma_{\beta s}$. Therefore the terminal edges of Γ_β give rise to the terminal edges of the graph Γ_α of lemma 5.5 as shown by figure 9.

FIG. 9. — Valency one isotopy on the graph

After the sequence of valency one isotopies we obtain a new graph Γ_2 with one digon disk for each puncture of $\tilde{P}$. Some of these digon disks have a valency 2 vertex (see figure 9). In addition each point of the orbit $\tilde{Q}$ is enclosed in a digon disk.

(4) Valency one isotopy on the image $\Psi_1(\Gamma_1)$.

By construction there are two points of $\tilde{P}$ and one point of $\tilde{Q}$ on the terminal edge e_1 of Γ_β . Each other terminal edge $e_j, j \neq 1$ contains one point of $\tilde{P}$ and no point of $\tilde{Q}$. The situation is shown by figure 10 for the efficient representative $(\Psi_\alpha, \Gamma_\alpha)$.

FIG. 10. — Position of the disks on $(\Psi_\alpha, \Gamma_\alpha)$

The valency one isotopies are obvious for each terminal edge e_j , $j \neq 1$. Indeed from figure 10 we observe that each such valency one isotopy is just a retraction described by figure 9, plus the retraction of the image of each edge which covers the edge between the two disks Δ_j^P and Δ_j^P (see figure 11.a). The valency one isotopy for the terminal edge e_1 is shown by figure 11.b.

FIG. 11. — Valency one isotopy

After this sequence of valency one isotopies we obtain a new topological representative (Ψ_2, Γ_2) . We observe from figure 11 that (Ψ_2, Γ_2) has n valency 2 vertices on some of the disks Δ_i^P . We also observe the existence of a pair of adjacent edges which are tangent under Ψ_2 (see figure 11.b). This tangency creates non local injectivity (see the move 5 in § 1.2).

(5) Folding operation.

The tangency which is shown by figure 11.b implies that the folding operation which has to be applied is an absorbing folding. Indeed there is a boundary edge of a boundary disk Δ_i^P whose image is given by $\Psi_2(b_i^P) = b_{i+1}^P$ and the tangency discussed above is defined between the boundary edge b_i^P and an edge e_j whose image is $\Psi_2(e_j) = b_{i+1}^P \cdot X$, where X is an edge path in Γ_2 . The folding is absorbing since it implies the collapsing of the boundary edge b_i^P . Furthermore since the boundary edges b_j^P are permuted under Ψ_2 and since the edge b_i^P has been collapsed by the absorbing folding, we obtain a pretrivial forest which is therefore collapsed by the move 1' (see § 1.2). These folding and collapsing operations lead to a topological representative (Ψ_3, Γ_3) which is shown by figure 12. The figure shows the result of the sequence of transformations at the terminal edge e_1 (12.a), at the other terminal edges (12.b) and at all other edges (12.c).

(6) Another folding.

After the previous folding plus collapsing operations, all the boundary disks Δ_i^P are monogon disks. Some of them have a 4-valent boundary vertex (see figure 12.a, c).

FIG. 12. — Absorbing folding

The topological representative (Ψ_3, Γ_3) has again a tangency between two edges at such a 4-valent vertex (figure 12.a). Since these 4-valent vertices are permuted under Ψ_3 , this tangency creates non local injectivity. In order to remove these tangencies (for the iterates) we apply a sequence of folding operations at all these valency 4 vertices. These operations are shown by figure 13.

After this sequence of folding operations we obtain a topological representative (Ψ_4, Γ_4) .

Lemma 5.6. — *With the above notation, the topological representative (Ψ_4, Γ_4) of φ_n is efficient.*

Recall that from the algorithm described in § 1.2, the non-local injectivity is only created by some tangencies. The sequence of transformations (1)-(6) started with the efficient representative $(\Psi_\alpha, \Gamma_\alpha)$ given by lemma 5.5. For this representative there are four classes of vertices, namely the vertices on the boundary disks $\Delta_i^R, \Delta_i^P, \Delta_i^Q$ and the other vertices. The boundary disks Δ_i^R have been collapsed by the operations (1), (2), (3). The other transformations (4), (5), (6) have only affected the vertices corresponding to the disks Δ_i^P . After these transformations no tangencies occur at the new

FIG. 13. — Folding again

vertices. Since the other vertices have not been affected by the transformations, no tangencies have been created and therefore (Ψ_4, Γ_4) is efficient. $\square$

The efficient representative (Ψ_4, Γ_4) satisfies all the properties of theorem 4.1 and therefore $[\delta] = \text{BT}(\mathbf{P}, \varphi_\delta) > [\gamma] = \text{BT}(\mathbf{Q}, \varphi_\delta)$, which completes the proof of theorem 5.2 in this case.

The other cases.

In the first case we assumed that the periodic orbit $\mathbf{Q}$ of braid type γ was visiting only one terminal edge. The construction of a small cancellation orbit $\mathbf{P}$ is in fact exactly the same in the case where more than one terminal edge is visited by $\mathbf{Q}$. The proof of the forcing relation via theorem 4.1 is a little bit longer in this case because more

tangencies are created and more folding operations are necessary to remove them. The arguments are the same, namely the elementary operations only affect the disks Δ_i^P which become monogon disks whereas the disks Δ_i^Q are unaffected and thus remain digon disks.

It remains to consider the case where the periodic orbit Q does not visit any terminal edge. In this case, the previous proof shows that every small cancellation periodic orbit P whose braid type is $[\delta]$ forces $[\gamma]$. This completes the proof of the theorem. $\square$

The proof of theorem 5.2 actually gives more information. Indeed, each terminal edge $e_i \in T(\Gamma_\beta)$, $i = 1, \dots, n$, is subdivided into the collection $\{e_i^1, \dots, e_i^l\}$, where we have chosen the edge e_i^1 in such a way that the initial vertex of $i(e_i^1)$ is the boundary vertex v_i of the terminal edge e_i .

Definition 5.7. — Let $\beta \in B_n$ be a pseudo-Anosov braid whose induced permutation has a single cycle and let $(\Psi_\beta, \Gamma_\beta)$ be a standard representative of φ_β . An orbit P of the pseudo-Anosov map $\varphi_\beta \in [\varphi_\beta]$ is called a *very small cancellation orbit* with respect to $(\Psi_\beta, \Gamma_\beta)$ if P visits all the terminal edges e_i^1 , $i = 1, \dots, n$, of the subdivided graph of Γ_β^* .

The proof of theorem 5.2 gives, in fact:

Theorem 5.8. — Let $\beta \in B_n$ be a pseudo-Anosov braid whose induced permutation has a single cycle and let $(\Psi_\beta, \Gamma_\beta)$ be a standard representative of φ_β . For any $[\gamma] \in \mathcal{G}(\beta)$ there exists a very small cancellation periodic orbit P of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$ such that $[\beta] > [\delta] > [\gamma]$, where $[\delta] = \text{BT}(P, \varphi_\beta)$.

From now on we shall mainly consider the braid types corresponding to the very small cancellation periodic orbits.

During the previous proof we had to control the transition between the isotopy classes φ_α and φ_η . In this process some non local injectivity (cancellations) has been “removed”. The name “small cancellation” has been chosen because the amount of such cancellation occurred only at the terminal edges and it was possible to control the variation. Dynamically this transition has two effects. First the topological entropy or the growth rate is strictly decreasing. The second effect is that infinitely many periodic orbits are suppressed. As a corollary of our constructions we obtain the following:

Lemma 5.9. — Let $\beta \in B_n$ be a pseudo-Anosov braid whose induced permutation has a single cycle and let $\mathcal{G}(\beta)$ be its genealogy set. Then for every $[\gamma] \in \mathcal{G}(\beta)$ different from $[\beta]$ the topological entropy function satisfies $h([\gamma]) < h([\beta])$.

From the definitions of the forcing relation and the entropy of a braid type, the inequality $h([\gamma]) \leq h([\beta])$ is clear. Let us prove it is a strict inequality. By theorem 5.8 it is sufficient to consider that $[\gamma]$ is realized as a very small cancellation periodic orbit Q

with respect to a standard representative $(\Psi_\beta, \Gamma_\beta)$. We first apply the piercing operation at the points of $\tilde{Q}$. The new efficient representative has the same growth rate as $(\Psi_\beta, \Gamma_\beta)$. Then we apply the closing operation at the monogon disks corresponding to the periodic orbit R of braid type $[\beta]$. After this closing operation we apply, as above, a sequence of valency one isotopy and the growth rate strictly decreases (see for instance lemma 1.11 of [BH2]). The folding operations that arise possibly after this step will again reduce the growth rate. As a result, either a very small cancellation periodic orbit Q contains some of the extreme points of the terminal edges (boundary vertices), in which case it contains them all and $[\gamma] = [\beta]$, or $h([\gamma]) < h([\beta])$. $\square$

As a direct consequence of this lemma we obtain:

Lemma 5.10. — *For braid types in the disk, the forcing relation is a partial order.*

The reflexivity and the transitivity of the forcing relation is clear from the definition. The only non trivial part is the antisymmetry. Let us assume that $[\beta] \succ [\gamma]$ and $[\gamma] \succ [\beta]$. The first observation is that it is only necessary to consider the case where $[\gamma]$ and $[\beta]$ are pseudo-Anosov braid types. Indeed the finite order case is easy and for the reducible cases the study restricts to the irreducible components (in the sense of the Nielsen-Thurston theorem). If $[\gamma]$ and $[\beta]$ are both pseudo-Anosov, then, by lemma 5.9, one has $[\gamma] = [\beta]$. $\square$

Observe that we have used the assumption that the surface is a disk only in a weak form. The above proof is exactly the same for the braid type of a homeomorphism isotopic to the identity on any surface.

6. Cancellation rectangles

The goal of this section is to define an effective criterion which enables one to compare two very small cancellation periodic orbits. To this end we introduce some numerical quantities which are related to the periodic orbit for a given standard representative.

Let P be a very small cancellation periodic orbit of φ_β with respect to a standard representative $(\Psi_\beta, \Gamma_\beta)$ and let $T(\Gamma_\beta) = \{e_1, \dots, e_n\}$ be the set of terminal edges of Γ_β . Let $\tilde{P} = \Pi_{\Gamma_\beta}(P)$ be the projection of P on the graph Γ_β . For each terminal edge $e_i \in E(\Gamma_\beta)$, $i = 1, \dots, n$, we consider the set of points $\tilde{P}_i = \{p_i^1, \dots, p_i^{m_i}\} = \tilde{P} \cap e_i$. Let us recall that we have chosen the orientation of the terminal edges $\{e_1, \dots, e_n\}$ in such a way that the initial vertex $i(e_i)$ is the boundary vertex v_i of the edge. For the subdivided edges $\{e_i^1, \dots, e_i^{r_i}\}$ we normalize the length of each interval e_i^j to be 1. Therefore the length of a terminal edge e_i is normalized to $|e_i| = r_i$. The position of a point $p_i^j \in e_i$ is given by the distance $l(p_i^j) = d(v_i, p_i^j)$, where d is the usual distance on the real line $\mathbf{R}$ and v_i is the boundary vertex.

On each terminal edge $e_i \in T(\Gamma_\beta)$ two particular points of $\tilde{P}_i$ are important.

Definition 6.1. — Let P be a very small cancellation periodic orbit of the pseudo-Anosov map φ_β with respect to the standard representative $(\Psi_\beta, \Gamma_\beta)$. Let $\tilde{P}$ be the projection of P on Γ_β and set $\tilde{P}_i = \tilde{P} \cap e_i$, where e_i is a terminal edge. Let $p_i^m \in \tilde{P}_i$ be the nearest point to the boundary vertex v_i , i.e. $l(p_i^m) = \inf \{ l(p_i^j); p_i^j \in \tilde{P}_i \}$. Let also $p_i^M \in P_i$ be the point such that $p_i^M = \Psi_\beta^{-1}(p_j^m)$, where p_j^m is the corresponding nearest point to the boundary vertex on the terminal edge e_j .

There is another numerical quantity which can be associated to the points of $\tilde{P}_i$:

Definition 6.2. — Let β , φ_β and $(\Psi_\beta, \Gamma_\beta)$ be as above and let P be a very small cancellation periodic orbit of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$ whose projection on Γ_β is $\tilde{P}$. Let e_i be a terminal edge and let $\tilde{P}_i = \tilde{P} \cap e_i$. For every point $p_i^j \in \tilde{P}_i$ we consider $\Psi_\beta^{-1}(p_i^j) \in e_{ij}$. By lemma 2.10 the transversal word at the edge e_i is of the form $\Psi_\beta^{-1}(e_i) = M_i \cdot e_{\sigma(i)} \cdot M_i^{-1}$. The fact that $\Psi_\beta^{-1}(p_i^j) \in e_{ij}$ implies that:

- (i) either $M_i = B \cdot (e_{ij})^{\pm 1} \cdot A$, where A and B are some words (maybe empty),
- (ii) or $(e_{ij}) = e_{\sigma(i)}$.

The *width* of the point p_i^j on the terminal edge e_i is defined by $w(p_i^j) = |(e_i)^{\pm 1} \cdot A|$ in case (i), where $|X|$ is the geometric length of the word X , and $w(p_i^j) = 0$ in case (ii) (see figure 14).

We define, in the same way, the width of any point $\Psi_\beta(x) \in e_i$ on some terminal edge.

From the previous definition the following properties are obvious:

- (*) $w(p_i^j) = 0 \Leftrightarrow l(\Psi_\beta^{-1}(p_i^j)) < 1$.
- (**) $w(p_i^m) = 0 \Leftrightarrow p_i^m = p_i^M$.

We need one more definition before stating the main result of this section.

Definition 6.3. — Let β , φ_β and $(\Psi_\beta, \Gamma_\beta)$ be as above and let P be a very small cancellation periodic orbit of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$, whose projection on Γ_β is $\tilde{P}$. Let p_i^m and p_i^M be the two special points of $\tilde{P}_i$ given by definition 6.1.

If $p_i^m \neq p_i^M$ then by (**) the width $w(p_i^m)$ is ≥ 1 ; in this case we define the *cancellation rectangle* $R_i(P)$ on the terminal edge e_i to be the set of points $x \in \Gamma_\beta$ such that:

- $\Psi_\beta(x) \in e_i$,
- $l(p_i^m) \leq l(\Psi_\beta(x)) \leq l(p_i^M)$,
- $w(\Psi_\beta(x)) < w(p_i^m)$.

If $p_i^m = p_i^M$, then the cancellation rectangle is degenerate to the point p_i^m .

If P is a very small cancellation periodic orbits of φ_β and if Q is any other periodic orbit, then we write $\tilde{P}_i \triangleleft_i \tilde{Q}_i$ if either $\tilde{Q}_i = \emptyset$, or:

- (i) $l(p_i^m) < l(q_i^m)$ and,
- (ii) no point $q_i^j \in \tilde{Q}_i$ belongs to $R_i(P)$.

Finally we write $P \triangleleft Q$ if $\tilde{P}_i \triangleleft_i \tilde{Q}_i$ for all the terminal edges $\{e_1, \dots, e_n\}$.

FIG. 14. — Cancellation rectangle

We denote by $V\text{sc}(\gamma)$ the set of all the very small cancellation periodic orbits of φ_β whose braid type is $[\gamma]$. More generally we denote by $V(\delta)$ the set of all periodic orbits of φ_β whose braid type is $[\delta]$.

Theorem 6.4. — *Let $\beta \in B_n$ be a pseudo-Anosov braid whose closure is a knot and let φ_β be the pseudo-Anosov map in the isotopy class $[\varphi_\beta]$ defined by β . Let $(\Psi_\beta, \Gamma_\beta)$ be a standard representative of φ_β . If P is a very small cancellation periodic orbit of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$ whose braid type is $[\gamma]$, then $[\beta] \succ [\gamma] \succ [\delta]$ if and only if, for $P \in V\text{sc}(\gamma)$, there exists $Q \in V(\delta)$ so that $P \triangleleft Q$.*

Notice that this result enables one to check the forcing relation very easily since it reduces the problem to a finite number of inequalities.

The strategy for the proof of this theorem is very similar to the one for theorem 5.2. Let us first assume that $[\beta] \succ [\gamma] \succ [\delta]$.

Let us consider any periodic orbit $P \in V\text{sc}(\gamma)$. By the piercing operation at $\tilde{P}$, followed by the closing operation at the monogon disks we obtain a topological representative (Ψ_1, Γ_1) of φ_γ . The algorithm of § 1 applied to (Ψ_1, Γ_1) gives a standard representative $(\Psi_\gamma, \Gamma_\gamma)$. Since $[\gamma] \succ [\delta]$, φ_γ admits a periodic orbit Q' of braid type $[\delta]$. There is a corresponding periodic orbit Q for φ_β .

The piercing operation at the points $\tilde{P} \cup \tilde{Q}$ defines, by lemma 4.2, an efficient representative (Ψ_η, Γ_η) for the braid $\eta = \beta \cup_\beta \gamma \cup_\beta \delta \in B_{n+p+q}$. The closure of η is a three-components link. As in § 4 we consider the collection of boundary disks Δ_P , Δ_Q and Δ_R , where R is the collection of punctures corresponding to the braid β . The forgetting map $\tilde{f}_R$, as defined in § 4, gives: $\tilde{f}_R(\eta) = \xi = \gamma \cup_\beta \delta \in B_{p+q}$. The closing operation of § 4, at the points of R , produces a topological representative for the isotopy class $[\varphi_\xi]$. From the algorithm of § 1, we obtain an efficient representative (Ψ_ξ, Γ_ξ) which satisfies the properties of theorem 4.1, since $[\gamma] \succ [\delta]$. Therefore the efficient representative (Ψ_ξ, Γ_ξ) has a collection of digon disks Δ_Q and a collection of monogon disks Δ_P .

Let us assume that the property $P \triangleleft Q$ fails for some $P \in V \text{sc}(\gamma)$ and every $Q \in V(\delta)$. This implies that either

- a) there exists a terminal edge e_i such that $l(p_i^m) > l(q_i^m)$, or
- b) there exists, for a given terminal edge e_i , a point $q_i^j \in \tilde{Q}_i$ such that $q_i^j \in R_i(P)$.

The two situations are described by figure 15 below.

FIG. 15. — Comparison of the cancellation rectangles

In the situation a), we already used, in the proof of theorem 5.2 (see figure 9), the fact that, after the closing operation and a valency one isotopy on the graph, we obtain a topological representative with a monogon disk Δ_Q^i . This contradicts theorem 4.1.

In the situation b) we obtain, after the closing operation and a valency one isotopy, a topological representative (Ψ_2, Γ_2) which admits an absorbing folding at some boundary disks Δ_Q^i and Δ_P^i (see figure 16). After performing these absorbing folding operations we obtain a topological representative (Ψ_3, Γ_3) for φ_ξ whose boundary disks Δ_Q and Δ_P are all monogon disks. This again contradicts theorem 4.1 and completes the proof of the implication

$$[\beta] \succ [\gamma] \succ [\delta] \Rightarrow \text{for all } P \in V \text{sc}(\gamma), \text{ there exists } Q \in V(\delta) \text{ such that } P \triangleleft Q.$$

Let us now assume that $P \triangleleft Q$ for some $Q \in V(\delta)$.

FIG. 16. — Absorbing foldings

By definition, at each terminal edge $e_i \in T(\Gamma_\beta)$, either $\tilde{Q}_i = \emptyset$ or the inequality $l(p_i^m) < l(q_i^m)$ is satisfied. In the case where $\tilde{Q}_i = \emptyset$ for every terminal edge, $[\gamma] > [\delta]$ by theorem 5.8. Otherwise two cases have to be considered:

- (1) $l(p_i^M) < l(q_i^m)$ or,
- (2) $l(p_i^m) < l(q_i^m) < l(p_i^M)$ and $w(p_i^m) \leq w(q_i^m)$.

In fact the first case has already been studied in the proof of theorem 5.2. Therefore we just have to consider the second case which is shown by figure 17.

FIG. 17. — $P \triangleleft Q$

In this case, after the closing operation and the valency one isotopy we obtain a topological representative (Ψ_2, Γ_2) with one absorbing folding at a boundary disk Δ_P^i . After the sequence of absorbing folding operations at each boundary disk of Δ_P we obtain a topological representative (Ψ_3, Γ_3) which admits a partial folding (see figure 18). After this folding operation we obtain a topological representative (Ψ_4, Γ_4) without tangencies on this edge (see figure 18).

FIG. 18. — (Ψ_4, Γ_4) is efficient

We are now in the same situation as in lemma 5.6. Indeed, after the last partial folding, all the tangencies which have been created after the closing operation are removed. Therefore the topological representative (Ψ_4, Γ_4) of φ_ξ is efficient. Notice that figure 18 illustrates a simple case, where $w(p_i^m) < w(q_i^m)$. There is of course the possibility of having the equality $w(p_i^m) = w(q_i^m)$ for some terminal edge, all other inequalities being as before. In this case there is an additional folding operation leading to the same conclusion.

We observe, as in the proof of theorem 5.2, that the boundary disks Δ_Q have not been affected by the above sequence of elementary moves. Therefore the boundary disks Δ_Q of (Ψ_4, Γ_4) are all digon disks, whereas the boundary disks Δ_P are all monogon disks. The proof of the theorem is complete since the efficient representative (Ψ_4, Γ_4) satisfies the properties of theorem 4.1. $\square$

7. Periodic orbits with constraints

The goal of this section is to show that we can find some sequences of periodic orbits, for a topological representative $\Psi: \Gamma \rightarrow \Gamma$, which satisfies some topological constraints. For instance, in what follows, we will have to prove the existence of infinitely many periodic orbits of Ψ in the complement of some intervals in Γ .

Let (Ψ, Γ) be a topological representative and let (Ψ_s, Γ_s) be the subdivided representative as defined in § 2.1. Let us denote by $\mathcal{MG}(\Psi, \Gamma)$ the Markov graph of the incidence matrix $M(\Psi, \Gamma)$. We fix the notation in such a way that the vertex w_i of $\mathcal{MG}(\Psi, \Gamma)$ corresponds to the edge e_i of the graph Γ . If multi-indices are required in order to define some edges of a graph Γ then the same multi-indices will be used for the corresponding vertices of the Markov graphs.

Definition 7.1. — Let (Ψ', Γ') be a topological representative which is obtained from (Ψ, Γ) by subdivision and let $\pi: w_1 \rightarrow w_2 \dots \rightarrow w_n$ be a path in $\mathcal{MG}(\Psi, \Gamma)$. If the corresponding collection of edges $\{e_1, \dots, e_n\}$ (maybe with repetition) of Γ are subdivided into the edges: $\{(e_1^1, \dots, e_1^{r_1}) \dots (e_n^1, \dots, e_n^{r_n})\}$ of Γ' . Then a collection of vertices $\{w_1^{i_1}, \dots, w_n^{i_n}\}$ of $\mathcal{MG}(\Psi', \Gamma')$, or equivalently a collection of edges $\{e_1^{i_1}, \dots, e_n^{i_n}\}$ of Γ' , is called *coherent with respect to π* if there exists a path $w_1^{i_1} \rightarrow \dots \rightarrow w_n^{i_n}$ in $\mathcal{MG}(\Psi', \Gamma')$.

The main result of this section is:

Lemma 7.2. — Let (Ψ, Γ) be a topological representative whose incidence matrix is Perron-Frobenius. Let $E = \{e_1, \dots, e_n\} \subset E(\Gamma)$ be a proper collection of edges of Γ such that the Markov graph $\mathcal{MG}(\Psi, \Gamma)$ has a closed path $\pi: w_1 \rightarrow w_2 \dots \rightarrow w_n \rightarrow w_1$. Let θ be a collection of periodic orbits of Ψ such that each edge of E contains, at least, one point of θ . We denote by (Ψ', Γ') the topological representative which is obtained from the subdivided representative (Ψ_s, Γ_s) by subdivision at the points of θ . Each edge $e_i \in E$ is subdivided into $\{e_i^1, \dots, e_i^{r_i}\} \subset E(\Gamma')$ ($r_i \geq 2$).

Then, for each coherent choice $\Omega = \{w_1^i, \dots, w_n^i\}$ of vertices in $\mathcal{MG}(\Psi', \Gamma')$, the graph $\mathcal{MG}(\Psi', \Gamma') - \Omega$ is transitive.

Let us start the proof by recalling some well-known facts. By assumption the Markov graph $\mathcal{MG}(\Psi, \Gamma)$ is transitive since the incidence matrix $M(\Psi, \Gamma)$ is Perron-Frobenius. Any subdivision (Ψ', Γ') of the topological representative (Ψ, Γ) defines a topological representative whose incidence matrix is Perron-Frobenius. Therefore the Markov graph $\mathcal{MG}(\Psi', \Gamma')$ of any subdivided representative (Ψ', Γ') is transitive. There is another property which follows from the transitivity:

Closed path property. — If $\pi : w_1 \rightarrow w_2 \dots \rightarrow w_n \rightarrow w_1$ is a closed path in the transitive Markov graph $\mathcal{MG}(\Psi, \Gamma)$, then there exists, at least, one vertex w_i , $i \in \{1, \dots, n\}$, at which more than one arrow arrives and there exists, at least, one vertex w_j , $j \in \{1, \dots, n\}$, at which more than one arrow starts.

Notice that a single closed path cannot be the whole graph $\mathcal{MG}(\Psi, \Gamma)$ since the matrix $M(\Psi, \Gamma)$ is Perron-Frobenius. The closed path property is obvious by the transitivity of the graph.

Let us start the proof by considering one arrow $w_1 \rightarrow w_2$ in the path π . If there is more than one arrow starting from w_1 then we subdivide the corresponding edge e_1 according to the number of starting arrows. This subdivision defines a topological representative (Ψ_1, Γ_1) which corresponds to the subdivision of the edge e_1 of the subdivided representative (Ψ_s, Γ_s) . Let us denote by $e_{1,1}$ the edge of Γ_1 which gives rise to the arrow $w_{1,1} \rightarrow w_2$ in $\mathcal{MG}(\Psi_1, \Gamma_1)$.

Case 1. — There is, at least, another arrow $w_N \rightarrow w_2$ ($N \neq 1$) arriving at w_2 .

(1.i) All the points of $\theta \cap e_2$ are images, under Ψ , of points of $\theta \cap e_{1,1}$.

In this case the edge $e_{1,1}$ of Γ_1 is subdivided into the edges $\{e_{1,1}^1, \dots, e_{1,1}^r\}$ of Γ' and e_2 is subdivided into $\{e_2^1, \dots, e_2^r\}$. By the assumption (1.i) we have $r_2 = r_1 = r$. The graph $\mathcal{MG}(\Psi', \Gamma')$ has therefore r arrows $w_{1,1}^j \rightarrow w_2^j$, $j \in \{1, \dots, r\}$. It also has, at least, r arrows $w_N^j \rightarrow w_2^j$ for $j \in \{1, \dots, r\}$ (see figure 19 (i) and lemma 5.4).

By transitivity there exists a path in $\mathcal{MG}(\Psi', \Gamma')$ which connects w_2^j to w_N for all $j \in \{1, \dots, r\}$.

Property (1.i). — For $j \in \{1, \dots, r\}$, there exists, in $\mathcal{MG}(\Psi', \Gamma')$, a path $w_2^j \rightarrow \dots \rightarrow w_N$, which does not go through any $w_{1,1}^j$.

Indeed this property comes from the fact that only one arrow starts from $w_{1,1}$ (see figure 19).

Claim (1.i). — The property (1.i) implies that each graph $\mathcal{MG}(\Psi', \Gamma') - w_{1,1}^j$, $j \in \{1, \dots, r\}$, is transitive.

Let us remove the vertex $w_{1,1}^j$ from the graph $\mathcal{MG}(\Psi', \Gamma')$ and let us consider a pair x, y of vertices in $\mathcal{G}(\Psi_1, \Gamma_1)$. By transitivity there is a path from x to y in $\mathcal{G}(\Psi_1, \Gamma_1)$.

If this path does not go through $w_{1,1}$ in $\mathcal{MG}(\Psi_1, \Gamma_1)$ then the same path still exists in $\mathcal{MG}(\Psi', \Gamma') - w_{1,1}^i$. If this path goes through $w_{1,1}$ then, after the subdivision, it might go through the arrow $w_{1,1} \rightarrow w_2^j$. If it is not the case, then suppressing the vertex $w_{1,1}^i$ does not affect the path. If the path does go through $w_{1,1} \rightarrow w_2^j$, then we have to find another path. The original path, in $\mathcal{MG}(\Psi', \Gamma')$ can be represented by:

$$x \rightarrow \dots \rightarrow w_{1,1}^i \rightarrow w_2^j \rightarrow \dots \rightarrow y.$$

Since the path arrives at $w_{1,1}^i$ and $\{e_{1,1}^1, \dots, e_{1,1}^r\}$ is obtained from $e_{1,1}$ by subdivision, each arrow arriving at $w_{1,1}^i$ also arrives at all the other $w_{1,1}^i$. By the property (1.i) we can change the original path to the following one:

$$x \rightarrow \dots \rightarrow w_{1,1}^i \rightarrow w_2^i \rightarrow \dots \rightarrow w_N \rightarrow w_2^j \dots \rightarrow y.$$

This new path is not affected by the removal of the vertex $w_{1,1}^i$, which completes the proof of the claim (1.i).

FIG. 19. — Removing one vertex in a Markov graph

Case (1.ii). — Some points of $\theta \cap e_2$ are the images, under Ψ of points of $\theta \cap e_N$.

In order to simplify the presentation we assume that $\theta \cap e_2$ has only two points, one point being the image of a point in $\theta \cap e_N$, the other point being the image of a point in $\theta \cap e_{1,1}$. In the more general case the arguments are the same. Under this simplified assumption, the edge e_N is subdivided into $\{e_N^1, e_N^2\}$, the edge $e_{1,1}$ is subdivided into $\{e_{1,1}^1, e_{1,1}^2\}$ and the edge e_2 is subdivided into $\{e_2^1, e_2^2, e_2^3\}$. In this case and, with a suitable labeling, the Markov graph has the following arrows (see the figure 19.ii):

$$w_{1,1}^1 \rightarrow w_2^1, w_{1,1}^2 \rightarrow w_2^2, w_{1,1}^3 \rightarrow w_2^3$$

$$w_N^1 \rightarrow w_2^1, w_N^2 \rightarrow w_2^2, w_N^3 \rightarrow w_2^3.$$

In this case there is a property similar to (1.i) as follow:

Property (1.ii). — *There exists a path from w_2^1 to w_N^2 and also a path from w_2^2 and w_2^3 to w_N^1 . Furthermore, as above, there exists such paths which do not go through any $w_{1,1}^j$.*

Claim (1.ii). — *The property (1.ii) implies that the graph $\mathcal{MG}(\Psi', \Gamma') - w_{1,1}^j$ is transitive.*

The proof of the property and the claim are exactly the same as above.

Case 2. — *There is only one arrow arriving at w_2 .*

As above, if there is more than one arrow starting from w_1 , then we subdivide the edge e_1 according to the number of arrows. We denote by $e_{1,1}$ the subdivided edge which gives rise to the arrow $w_1 \rightarrow w_2$. As in the case (1.i) the number of points of θ on $e_{1,1}$ and on e_2 is the same. Therefore the subdivision at the points of θ defines the edges $\{e_{1,1}^1, \dots, e_{1,1}^r\}$ and $\{e_2^1, \dots, e_2^r\}$ and there are r arrows: $w_{1,1}^i \rightarrow w_2^i$, $i \in \{1, \dots, r\}$. At this point it is clear that removing any of the vertices $w_{1,1}^j$ produces a non transitive graph since no arrow arrives at w_2^j . Therefore, if we remove the vertex $w_{1,1}^j$ we have also to remove (at least) the vertex w_2^j .

Let us follow the closed path $\pi: w_1 \rightarrow w_2 \rightarrow \dots \rightarrow w_n \rightarrow w_1$ in $\mathcal{MG}(\Psi', \Gamma')$. From the closed path property there exists at least one vertex w_i in π at which the assumptions of the case 1 are satisfied. In order to simplify the exposition let us assume that the vertex w_i above is the next one in the path, i.e. $w_i = w_3$, otherwise we have to continue the process along the path.

If there is more than one arrow starting from w_2 then, as above, we subdivide the edge e_2 according to the number of arrows. Let us denote by $e_{2,1}$ the edge which gives rise to the arrow $w_2 \rightarrow w_3$. In what follows we are only interested in the subpath $w_{1,1} \rightarrow w_{2,1} \rightarrow w_3$ of the original path π in $\mathcal{G}(\Psi', \Gamma')$. In fact we are only interested in a part of this subpath. Indeed there was one arrow between $w_{1,1}$ and w_2 but, since we have subdivided e_2 into $\{e_{2,1}, \dots, e_{2,l}\}$, there are l arrows starting from $w_{1,1}$, i.e. one arrow to each of the $w_{2,j}$. In order to use the same formulation as above we shall subdivide further the edge $e_{1,1}$, according to the above l arrows, but this is not absolutely necessary. Let us denote by $e_{1,1,1}$ the subdivided edge which gives rise to the arrow $e_{1,1,1} \rightarrow e_{2,1}$. Now we have to subdivide again, at the points of θ . This subdivision gives rise to the following edges: $\{e_{1,1,1}^1, \dots, e_{1,1,1}^m\}$ and $\{e_{2,1}^1, \dots, e_{2,1}^m\}$. The graph $\mathcal{MG}(\Psi', \Gamma')$ has now m arrows $w_{1,1,1}^i \rightarrow w_{2,1}^i$, $i \in \{1, \dots, m\}$. By assumption there is, at least, one arrow $w_N \rightarrow w_3$ for $N \neq 1, 2$. As in the case 1 we have the two possibilities:

(2.i) All the points of $\theta \cap e_3$ are the images under Ψ' of points of $\theta \cap e_{2,1}$.

(2.ii) Some points of $\theta \cap e_3$ are the images under Ψ' of points of $\theta \cap e_N$.

As above a property, similar to (1.i) or (1.ii), is satisfied.

Property 2. — *There exists a path $w_3^i \rightarrow \dots \rightarrow w_N^i$, where w_N^i corresponds to the possible subdivision $e_N \rightarrow \{e_N^1, \dots, e_N^r\}$ of the edge e_N at the points of θ (case 2.ii).*

In this case also this property is sufficient to conclude that, if we remove a coherent sequence of vertices (in this case, any pair $\{w_{1,1,1}^j, w_{2,1}^j\}$), then the resulting graph is transitive.

In order to complete the proof of lemma 7.2 we just have to follow the path $\pi : w_1 \rightarrow \dots \rightarrow w_n \rightarrow w_1$ in the graph $\mathcal{MG}(\Psi, \Gamma)$, arrow after arrow. For each such arrow we apply the arguments of the cases 1 or 2 above, depending upon the occurrence or not of an extra-arrow arriving at the vertex. Finally, after removing any coherent sequence of vertices $\{w_1^{j_1}, w_2^{j_2}, \dots, w_n^{j_n}\}$, we obtain a transitive graph. $\square$

Lemma 7.2 is an example of conditions under which, removing a collection of vertices from a transitive Markov graph preserves the transitivity property. These conditions are sufficient for our needs. It would be interesting to find a general condition under which the transitivity is preserved.

FIG. 20. — Removing a coherent sequence of vertices

8. The main theorem

In this section we apply the results of the previous parts in order to prove each property of the main theorem. Recall that the property (1) has already been proved in § 5.

8.1. A local distance on $\mathcal{G}(\beta)$

The simplest topology we can define on the set $\mathcal{G}(\beta)$ is so that the open sets are union of paths $\pi(\gamma, \delta)$, where γ and δ are two elements of $\mathcal{G}(\beta)$ such that $\gamma \succ \delta$. This topology on the set $\mathcal{G}(\beta)$ is called the *path topology*.

Let us fix, once and for all, a given standard representative $(\Psi_\beta, \Gamma_\beta)$ for the pseudo-Anosov map φ_β induced by a braid $\beta \in B_n$. We also denote, as usual, by $E_T(\Gamma_\beta)$ the set $\{e_1, \dots, e_n\}$ of terminal edges of Γ_β .

Let us denote by $\text{VSC}(\Psi_\beta, \Gamma_\beta)$ the set of very small cancellation periodic orbits of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$. We also set

$$\text{Vsc}(\gamma) = \{ P \in \text{VSC}(\Psi_\beta, \Gamma_\beta); \quad \text{BT}(P, \varphi_\beta) = [\gamma] \},$$

and

$$\text{Vsc}_\beta = \{ [\gamma] \in \mathcal{G}(\beta) \text{ s.t. } \text{Vsc}(\gamma) \neq \emptyset \}.$$

From theorem 5.8, the subset Vsc_β , with the path topology of $\mathcal{G}(\beta)$, forms a neighborhoods of $[\beta]$ in $\mathcal{G}(\beta)$.

In what follows we shall restrict the study of the topology of $\mathcal{G}(\beta)$ to the subset $\text{Vsc}_\beta \subset \mathcal{G}(\beta)$. Let us now define a distance function for two very small cancellation periodic orbits.

Definition 8.1. — Let P and Q be in $\text{VSC}(\Psi_\beta, \Gamma_\beta)$, we define the function $d : \text{VSC}(\Psi_\beta, \Gamma_\beta) \times \text{VSC}(\Psi_\beta, \Gamma_\beta) \rightarrow \mathbf{R}^+$ by

$$d(P, Q) = \sum_{e_i \in E_T(\Gamma_\beta)} |l(p_i^m) - l(q_i^m)| + |\omega(p_i^m) - \omega(q_i^m)|.$$

Let us check that the function d is a distance. The symmetry property is obvious. If we assume that $d(P, Q) = 0$ then for all $e_i \in E_T(\Gamma_\beta)$ $l(p_i^m) = l(q_i^m)$ and $\omega(p_i^m) = \omega(q_i^m)$, which implies that the two periodic orbits P and Q are the same. The triangle inequality is also obvious.

Proposition 8.2. — Let $[\gamma]$ and $[\delta]$ be two elements in $\text{Vsc}_\beta \subset \mathcal{G}(\beta)$; then the function $\Delta : \text{Vsc}_\beta \times \text{Vsc}_\beta \rightarrow \mathbf{R}^+$ defined by

$$\Delta(\gamma, \delta) = \inf \{ d(P_i, Q_j); P_i \in \text{Vsc}(\gamma) \text{ and } Q_j \in \text{Vsc}(\delta) \}$$

is a distance on Vsc_β .

The first observation is that, for a given braid type $[\gamma] \in \mathcal{G}(\beta)$, the set $\text{Vsc}(\gamma)$ is finite. Indeed the set of periodic orbits of a given pseudo-Anosov map and of a given period is finite. The set $\text{Vsc}(\gamma)$ is finite since all its elements have the same period. The function Δ is thus well-defined and each property of a distance is obviously satisfied. The property (2) of the main theorem is proved. $\square$

8.2. Density on a path

The goal of this subsection is to prove the following lemma which implies property (3) of the main theorem.

Lemma 8.3. — Let $\beta \in B_n$ be a pseudo-Anosov braid whose induced permutation has a single cycle and let $(\Psi_\beta, \Gamma_\beta)$ be a standard representative of the pseudo-Anosov element φ_β in the isotopy class defined by β . If Q is a very small cancellation periodic orbit of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$ whose braid type is $\text{BT}(Q, \varphi_\beta) = [\delta]$, there exists infinitely many very small cancellation periodic orbits $\{ P_i, i \in \mathbf{I} \}$ such that

$$\text{for all } i \in \mathbf{I}, \quad [\beta] > [\gamma_i] > [\delta], \quad \text{where } [\gamma_i] = \text{BT}(P_i, \varphi_\beta).$$

In order to prove this lemma we apply theorem 6.4, together with lemma 7.2. Indeed, from theorem 6.4, the relations $[\beta] \succ [\gamma_i] \succ [\delta]$ hold if the periodic orbits P_i and Q satisfy $P_i \triangleleft Q$ for all $i \in \mathbf{I}$ (see definition 6.3). This property means that we have to find infinitely many very small cancellation periodic orbits P_i such that the points of Q avoid the cancellation rectangles $R(P_i)$. A cancellation rectangle is a finite union of disjoint intervals whose boundary points are either periodic points of the map $\Psi_\beta : \Gamma_\beta \rightarrow \Gamma_\beta$ or preimages of such points. We have to check that the assumptions of lemma 7.2 are satisfied with respect to the collections of intervals that we want to remove from the graph Γ_β . To this end we need the following data:

- (1) A Perron-Frobenius matrix.
- (2) A proper collection of edges $E = \{e_1, \dots, e_n\} \subset E(\Gamma_\beta)$, with a closed path: $w_1 \rightarrow w_2 \rightarrow \dots \rightarrow w_n \rightarrow w_1$ in the Markov graph $\mathcal{MG}(\Psi_\beta, \Gamma_\beta)$.
- (3) A collection θ of periodic orbits with at least one point on each edge of E .
- (4) A subdivision with a coherent choice of edges.

The Perron-Frobenius property is satisfied by the bloc of the incidence matrix $M(\Psi_\beta, \Gamma_\beta)$ which corresponds to the edges of the tree $T_\beta = \Gamma_\beta - E(B(\Gamma_\beta))$. This is due to the lemmas 2.8 and 2.9 since φ_β is pseudo-Anosov.

The collection of edges we have to consider is naturally the collection $E_T(\Gamma_\beta)$ of terminal edges of Γ_β . From lemma 5.3, the closed path property is satisfied with respect to this collection of edges.

The collection θ of periodic orbits of lemma 7.2 is the single periodic orbit $\tilde{Q}$ which is the projection on Γ_β of the periodic orbit Q given by assumption of lemma 8.3. Since Q is a small cancellation periodic orbit, property (3) above is satisfied with respect to the set of terminal edges.

It remains to find a coherent choice of edges for the subdivision of the standard representative $(\Psi_\beta, \Gamma_\beta)$ at the points of $\tilde{Q}$. There are two important points of $\tilde{Q}$ on each terminal edge $e_j \in E_T(\Gamma_\beta)$ which have been denoted q_j^m and q_j^M in definition 6.1. These two points might be the same if $\omega(q_j^m) = 0$ (by property (***) of § 6). The points of $\tilde{Q}$ define a subdivision of the representative $(\Psi_\beta, \Gamma_\beta)$. Let us denote the subdivided edges which correspond to the terminal edge e_j by $\{e_j^1, \dots, e_j^{r_j}\}$, where the indices $\{1, \dots, r_j\}$ are ordered according to: $i < k$ if $l_j(x) < l_j(y)$ for $x \in e_j^i$ and $y \in e_j^k$. By construction it is clear that the sequence $E_m = \{e_1^1, \dots, e_n^1\}$ is a coherent choice. If, on the terminal edge e_j , the interval $[q_j^m, q_j^M]$ is not reduced to a point, then this interval is a union of edges e_j^i , $i = 2, \dots, k < r_j$; we denote it by e_j^R . If the interval $[q_j^m, q_j^M]$ is reduced to a point, then we denote by e_j^R the union of the edges e_j^i , $i = 2, \dots, r_j$. Finally we set $E^R = \{e_1^R, \dots, e_n^R\}$. By construction, it is also clear that E^R is a coherent choice of edges. In order to conclude the proof of lemma 8.3, it is sufficient to prove the existence of infinitely many periodic orbits P_i such that the points $(p_i)_j^m$ belong to the intervals of the sequence E_m and the other points of P_i avoid the cancellation rectangle (by theorem 6.4). This property is satisfied by using lemma 7.2 with the

coherent choice E^R . Indeed the graph that is obtained from the initial Markov graph by removing the vertices corresponding to E_R is transitive and not reduced to a single closed path. Therefore it contains infinitely many closed paths and the map Ψ_β admits infinitely many periodic orbits avoiding the edges of E_R . $\square$

8.3. Ramification points

In order to prove the properties (4), (5), (6) and (7) of the main theorem we have first to understand the elementary properties of the ramification points. That is the goal of this part. The basic situation we have to consider is when three braid types $[\beta]$, $[\gamma]$, $[\delta]$ are such that $[\beta] \succ [\gamma]$ and $[\beta] \succ [\delta]$ but $[\gamma]$ and $[\delta]$ are not related by the forcing relation. In this case we say that $[\gamma]$ and $[\delta]$ are *unrelated*.

As above we assume that $\beta \in B_n$ is a pseudo-Anosov braid and we consider a standard representative $(\Psi_\beta, \Gamma_\beta)$ for the pseudo-Anosov element φ_β defined by β . Let us assume that the braid type $[\gamma]$ is realized by a very small cancellation periodic orbit P of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$, i.e. $[\gamma] \in V \text{sc}_\beta$.

Lemma 8.4. — *With the above notation, if a very small cancellation periodic orbit P of braid type $[\gamma]$ is given, there exists infinitely many very small cancellation periodic orbits $\{Q_i, i \in \mathbf{I}\}$ of braid types $\{[\delta_i], i \in \mathbf{I}\}$ such that $[\gamma]$ and $[\delta_i]$ are unrelated for all $i \in \mathbf{I}$.*

In order to prove this lemma we just have, by theorem 6.4, to prove the existence of infinitely many periodic orbits which do not avoid the cancellation rectangles of the periodic orbit P . The existence of such orbits is obvious by the density of the periodic orbits of the map Ψ_β on the graph Γ_β . $\square$

From now on we shall fix a pair $([\gamma], [\delta])$ of unrelated braid types. Furthermore we shall assume that these braid types are realized by two very small cancellation periodic orbits (P, Q) of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$.

With the notation of theorem 6.4, the periodic orbits P and Q and their projections $\tilde{P}$ and $\tilde{Q}$ on Γ_β satisfy the following properties.

There exists a terminal edge $e_i \in E_T(\Gamma_\beta)$ such that either $\tilde{Q} \cap R_i(\tilde{P}) \neq \emptyset$ or $\tilde{P} \cap R_i(\tilde{Q}) \neq \emptyset$. In this case, if we set $Q_i = \tilde{Q} \cap e_i$ and $P_i = \tilde{P} \cap e_i$, then the following inequalities are satisfied (maybe after exchanging P and Q):

$$\text{there exists } q_i \in Q_i \text{ s.t. either } l(p_i^m) < l(q_i) < l(p_i^M) \text{ and } \omega(q_i) < \omega(p_i^m) \\ \text{or } l(q_i^m) < l(p_i^m).$$

For each terminal edge $e_i \in E_T(\Gamma_\beta)$ we consider the following points:

$$r_i^m(P, Q) = \begin{cases} p_i^m & \text{if } l(p_i^m) < l(q_i^m) \\ q_i^m & \text{if } l(q_i^m) < l(p_i^m) \end{cases}$$

and
$$r_i^M(P, Q) = \begin{cases} p_i^M & \text{if } l(p_i^M) < l(q_i^M) \\ q_i^M & \text{if } l(q_i^M) < l(p_i^M). \end{cases}$$

If $r_i^m(P, Q) = p_i^m$, let us consider the set of points $\tilde{Q}_i = Q \cap R_i(P)$ and let $s_i(P, Q)$ be the point $q_i^l \in \tilde{Q}_i$ such that $l(q_i^l)$ is minimal in $\{l(q_i); q_i \in \tilde{Q}_i\}$. Similarly we define the point $v_i(P, Q)$ to be the point $q_i^w \in \tilde{Q}_i$ such that $\omega(q_i^w)$ is minimal in $\{\omega(q_i); q_i \in \tilde{Q}_i\}$. If $r_i^m(P, Q) = q_i^m$, we define in the same way the points $s_i(P, Q)$ and $v_i(P, Q)$ by exchanging the roles of P and Q .

Lemma 8.5. — *With the previous notation there exists infinitely many very small cancellation periodic orbits $\{M_i, i \in \mathbf{I}\}$ of φ_β with respect to $(\Psi_\beta, \Gamma_\beta)$ whose braid types are $\{[\mu_i], i \in \mathbf{I}\}$ such that, for $i \in \mathbf{I}$, $[\mu_i] \succ [\gamma]$ and $[\mu_i] \succ [\delta]$. We denote by $B^+(\gamma, \delta)$ this set of periodic orbits.*

The proof of this lemma is exactly the same as the proof of lemma 8.3. In this case we apply lemma 7.2 with the collection θ of periodic orbits which is the union of the periodic orbits P and Q . By theorem 6.4, a periodic orbit $M \in B^+(\gamma, \delta)$ must satisfy the following inequalities:

$$(*) \quad \begin{cases} l(m_i^m) \leq l(r_i^m(P, Q)) \\ l(m_i^M) \leq l(r_i^M(P, Q)) \\ \omega(m_i^m) \leq \omega(v_i^m(P, Q)) \end{cases}$$

for every terminal edge $e_i \in E_T(\Gamma_\beta)$ (see figure 21). Therefore once again the periodic orbits in $B^+(\gamma, \delta)$ have to avoid a finite collection of intervals in Γ_β and we only have to check that this collection is coherent. This is left as an exercise. $\square$

The same arguments are also used in order to prove the following:

Lemma 8.6. — *With the previous notation we denote by $B^-(\gamma)$ (resp. $B^-(\delta)$) the set of very small cancellation periodic orbits of φ_β which satisfy:*

Each element $N \in B^-(\gamma)$ (resp. $A \in B^-(\delta)$) whose braid type is $[\nu]$ (resp. $[\alpha]$) is such that for all $M_i \in B^+(\gamma, \delta)$ whose braid type is $[\mu_i]$ the following relations hold:

$$[\mu_i] \succ [\nu] \quad \text{and} \quad [\nu] \succ [\gamma] \quad (\text{resp. } [\mu_i] \succ [\alpha] \quad \text{and} \quad [\alpha] \succ [\delta])$$

but $([\nu], [\delta])$ (resp. $([\alpha], [\gamma])$) are unrelated.

Then the sets $B^-(\gamma)$ and $B^-(\delta)$ are infinite.

We are now in a position to define a ramification point for the periodic orbits P and Q of braid type $[\gamma]$ and $[\delta]$ defined above. Let $[\mu_i]$ be the braid type of some $M_i \in B^+(\gamma, \delta)$ and $[\nu_i]$ be the braid type of some $N_i \in B^-(\gamma)$. By definition there is a path $\pi(\mu_i, \nu_i)$ in $\mathcal{G}(\beta)$. By lemma 8.3 there are infinitely many braid types on this path. Let us consider

$$[\mu_{i+1}] \neq [\mu_i] \text{ the braid type of some } M_{i+1} \in B^+(\gamma, \delta) \cap \pi(\mu_i, \nu_i)$$

$$\text{and} \quad [\nu_{i+1}] \neq [\nu_i] \text{ the braid type of some } N_{i+1} \in B^-(\gamma) \cap \pi(\mu_i, \nu_i).$$

This defines a “nested” sequence of paths $\pi(\mu_{i+1}, \nu_{i+1}) \subset \pi(\mu_i, \nu_i)$. By exchanging $B^-(\gamma)$ and $B^-(\delta)$ we also define a nested sequence of paths $\pi(\mu_{i+1}, \alpha_{i+1}) \subset \pi(\mu_i, \alpha_i)$.

FIG. 21. — Periodic orbits in $B^+(P, Q)$ on e_i

Lemma 8.7. — *Any sequence of paths $\pi(\mu_i, \nu_i)$ or $\pi(\mu_i, \alpha_i)$ as defined above does not converge in $\mathcal{G}(\beta)$.*

Let us assume that such a sequence of paths has a limit in $\mathcal{G}(\beta)$. If this limit is a point in $\mathcal{G}(\beta)$, then, from the inequalities (*) above, this point is realized as a braid type of a very small cancellation periodic orbit M of φ_β . It satisfies the following inequalities, for each terminal edge e_i :

$$l(m_i^m) = l(r_i^m(P, Q)), \quad l(m_i^M) = l(r_i^M(P, Q)) \quad \text{and} \quad \omega(m_i^m) \leq \omega(r_i^m(P, Q)).$$

Since ω is an integer function the last inequality has at most finitely many solutions. Therefore there are, at most, finitely many accumulation points of the sequences $\pi(\mu_i, \nu_i)$ and the limit cannot be a path.

Claim. — *An orbit of φ_β which satisfies the above equalities cannot be a periodic orbit.*

Indeed, those equalities correspond, by construction, to points on Γ_β which belongs to two distinct periodic orbits (namely P and Q). Therefore a single periodic orbit of φ_β cannot satisfy these equalities since otherwise two periodic orbits would have some common points as their projections on Γ_β . This is impossible by lemma 2.5. $\square$

Definition 8.8. — Let $(\Psi_\beta, \Gamma_\beta)$ be a standard efficient representative of the pseudo-Anosov homeomorphism φ_β . Let $([\gamma], [\delta])$ be two unrelated braid types in $\mathcal{G}(\beta)$ which are realized by two very small cancellation periodic orbits of φ_β . We denote by $\rho(\gamma, \delta)$ the set of accumulation points of all the possible sequences of paths $\pi(\mu_i, \nu_i)$ and $\pi(\mu_i, \alpha_i)$ as defined by 8.6 and 8.7. A point $\rho \in \rho(\gamma, \delta)$ is called a *ramification point*.

We define the *completion* $\bar{\mathcal{G}}(\beta)$ of $\mathcal{G}(\beta)$ to be the set obtain from $\mathcal{G}(\beta)$ by adding all the possible ramification points.

The completion as defined above is very similar to the completion of the rational numbers by the reals. Strictly speaking a ramification point is the type of some infinite

orbit under the pseudo-Anosov homeomorphism φ_β . They can thus be considered as conjugacy classes of braids with infinitely many strands, i.e. in the infinite group B_∞ . From the definition and the proof of lemma 8.7 we have:

Lemma 8.9. — *For a given pair $([\gamma], [\delta])$ of unrelated braid types in $\mathcal{G}(\beta)$ the set of ramification points $\rho(\gamma, \delta)$ is finite. Moreover each element $\rho \in \rho(\gamma, \delta)$ is the type of an infinite orbit of the pseudo-Anosov homeomorphism φ_β .*

Now let us prove that the partial ordering on the set $\mathcal{G}(\beta)$ induces a partial ordering on the completed set $\bar{\mathcal{G}}(\beta)$.

Definition 8.10. — Let ρ_1 and ρ_2 be two elements of $\bar{\mathcal{G}}(\beta)$. We say that ρ_1 *forces* ρ_2 , which we denote by the symbol $\rho_1 \succeq \rho_2$, if there exist two sequences of paths $\pi(\mu_i, \nu_i)$ and $\pi(\eta_i, \xi_i)$ in $\mathcal{G}(\beta)$ whose limits are ρ_1 and ρ_2 and such that, for all sufficiently large $i \in \mathbb{N}$, $[\mu_i] \succ [\eta_i]$ and $[\nu_i] \succ [\xi_i]$.

Let us prove:

Lemma 8.11. — *The relation “ $\succeq$ ” is a partial order in the set $\bar{\mathcal{G}}(\beta)$.*

Let us first observe that if ρ is an element of $\mathcal{G}(\beta)$ then there exist sequences of paths $\pi(\mu_i, \nu_i)$ whose limit is ρ , the obvious one being the constant sequence. We first check that the previous definition, restricted to $\mathcal{G}(\beta)$, gives the usual partial ordering. This is clear by using the constant sequences. Let us consider some other sequences converging toward ρ_1 and ρ_2 in $\mathcal{G}(\beta)$ such that $\rho_1 \succeq \rho_2$. The sequences of paths $\pi(\mu_i, \nu_i)$ and $\pi(\eta_i, \xi_i)$ are defined in such a way that, for all $i \in \mathbb{N}$, $\mu_i \succ \rho_1 \succ \nu_i$ and $\eta_i \succ \rho_2 \succ \xi_i$. Two cases are possible:

(1) there exists N such that, for all $i \geq N$, $\pi(\mu_i, \nu_i) \cap \pi(\eta_i, \xi_i) = \emptyset$

In this case, from the definition, we have $\nu_i \succ \eta_i$ for all $i \geq N$, which implies that $\rho_1 \succ \rho_2$ and $\rho_1 \neq \rho_2$.

(2) for all N , there exists $i \geq N$ such that $\pi(\mu_i, \nu_i) \cap \pi(\eta_i, \xi_i) \neq \emptyset$.

In this case we have $\rho_1 = \rho_2$. Indeed if it were not the case then we could find a subsequence with empty intersection for large i . We observe, by the same argument, that if two elements ρ_1 and ρ_2 in $\bar{\mathcal{G}}(\beta)$ are distinct and there are two nested sequences $\pi(\mu_i, \nu_i)$ and $\pi(\eta_i, \xi_i)$ converging respectively toward ρ_1 and ρ_2 , then we can find subsequences which are disjoint.

In order to prove that “ $\succeq$ ” is a partial order we check each property. The reflexivity is obvious. The transitivity is a little bit less obvious. We proceed as follows: let $\rho_1 \succeq \rho_2$ and $\rho_2 \succeq \rho_3$ be three elements in $\bar{\mathcal{G}}(\beta)$. Then we can find four sequences of paths

$$\pi(\mu_i^j, \nu_i^j) \rightarrow \rho_j, \quad \text{for } j = 1, 3, \quad \text{and} \quad \pi(\eta_i^j, \xi_i^j) \rightarrow \rho_2 \quad \text{for } j = 1, 3$$

such that, for all sufficiently large $i \in \mathbb{N}$,

$$\begin{aligned} \mu_i^1 > \eta_i^1 \quad \text{and} \quad \nu_i^1 > \xi_i^1 \\ \eta_i^3 > \mu_i^3 \quad \text{and} \quad \xi_i^3 > \nu_i^3. \end{aligned}$$

We assume that ρ_1, ρ_2, ρ_3 are pairwise distinct since otherwise the transitivity is obvious. Then we can assume that, for sufficiently large i , the two pairs of sequences are disjoint. This implies that, for sufficiently large i , $\xi_i^3 > \mu_i^3$ and $\nu_i^1 > \eta_i^1$. By definition of these sequences, we have $\eta_i^1 > \xi_i^3$. This completes the proof of the transitivity since $\mu_i^1 > \eta_i^1 > \xi_i^3 > \mu_i^3$ and $\nu_i^1 > \eta_i^1 > \xi_i^3 > \mu_i^3 > \nu_i^3$.

In order to prove the antisymmetry we have to check that the “ loops ” $\rho_1 \geq \rho_2 \geq \rho_1$ in $\bar{\mathcal{G}}(\beta)$ are not possible except if $\rho_1 = \rho_2$. This is a consequence of the antisymmetry in $\mathcal{G}(\beta)$. $\square$

The properties (4) and (5) of the main theorem are now proved. The property (6) can be formulated as:

Lemma 8.12. — *Let $\rho_1 \geq \rho_2$ be two distinct ramification points in $\bar{\mathcal{G}}(\beta)$; then there exist infinitely many ramification points $\{\rho_i \in \bar{\mathcal{G}}(\beta), i \in \mathbb{N}\}$ such that, for all $i \in \mathbb{N}$, $\rho_1 \geq \rho_i \geq \rho_2$.*

Let us consider two sequences of paths $\pi(\mu_i, \nu_i) \rightarrow \rho_1$ and $\pi(\eta_i, \xi_i) \rightarrow \rho_2$ such that $\mu_i > \eta_i$ and $\nu_i > \xi_i$ for all sufficiently large $i \in \mathbb{N}$. Since ρ_1 and ρ_2 are distinct we can assume that, for sufficiently large i , the two sequences of paths $\pi(\mu_i, \nu_i)$ and $\pi(\eta_i, \xi_i)$ are disjoint. Let $\alpha \in \mathcal{G}(\beta)$ be such that, for a given sufficiently large i , $\nu_i > \alpha$ but (α, η_i) are unrelated. Infinitely many such braid types exist by lemma 8.4. We consider now a ramification point $\rho \in \rho(\alpha, \eta_i)$. By construction such a ramification point satisfies

$$\rho_1 \geq \nu_i \geq \rho \geq \eta_i \geq \rho_2,$$

proving the lemma. $\square$

Definition 8.13. — Let $\rho \in \bar{\mathcal{G}}(\beta)$ be a ramification point. A sequence of paths $\pi(\mu_i, \nu_i)$ whose limit is ρ can be split in two *half-paths* which we denote by: $\pi^+(\mu_i, \rho) \cup \pi^-(\rho, \nu_i)$. We say that two half-paths $\pi^+(\mu_i^1, \rho)$ and $\pi^+(\mu_i^2, \rho)$ are *equivalent* if

$$\begin{aligned} &\text{there exists } K \in \mathbb{N} \text{ such that for all } i > K, \pi^+(\mu_i^1, \rho) \subset \pi^+(\mu_i^2, \rho), \\ &\text{or } \pi^+(\mu_i^2, \rho) \subset \pi^+(\mu_i^1, \rho). \end{aligned}$$

The same definition holds for two half-paths $\pi^-(\rho, \nu_i^1)$ and $\pi^-(\rho, \nu_i^2)$.

It is clear that this relation among half-paths is an equivalence relation.

Definition 8.14. — Let $\rho \in \bar{\mathcal{G}}(\beta)$ be a ramification point. The *degree* of a ramification point is the number of equivalence classes of half-paths $\pi^+(\mu_i, \rho)$ and $\pi^-(\rho, \nu_i)$ whose limit is ρ .

Lemma 8.15. — *The degree of a ramification point in $\bar{\mathcal{G}}(\beta)$ is finite.*

Let us consider $\rho \in \rho(\gamma, \delta)$, where $\gamma = \text{BT}(\text{P}, \varphi_\beta)$ and $\delta = \text{BT}(\text{Q}, \varphi_\beta)$ are such that P and Q are two unrelated very small cancellation periodic orbits of φ_β , relative to a standard representative $(\Psi_\beta, \Gamma_\beta)$. In order to prove this result we have to control the possible ways a sequence of paths $\pi(\mu_i, \nu_i)$ converges toward a ramification point. Recall that a ramification point $\rho \in \rho(\gamma, \delta)$ is given by a finite collection of equalities, relative to the graph Γ_β . These equalities are obtained as the limits of squeezed inequalities (see the proof of lemma 8.5). From the finiteness of the graph there are only finitely many ways to obtain such sequences of squeezed inequalities. This is a sketch of proof, the details are easy from theorem 6.4 and Lemma 7.2 (see also figure 22). $\square$

FIG. 22. — Convergence toward a ramification point

Let us now prove the property (8) of the main theorem.

Lemma 8.16. — *The set $\bar{\mathcal{G}}(\beta)$ is “non simply connected”.*

The proof of lemma 8.16 consists in exhibiting non-trivial loops in $\bar{\mathcal{G}}(\beta)$. These loops are obtained from, at least, two paths of the following form: $\pi(\gamma, \mu) \cup \pi(\mu, \delta)$ and $\pi(\gamma, \nu) \cup \pi(\nu, \delta)$ such that (μ, ν) are unrelated in $\mathcal{G}(\beta)$.

These paths are obtained in the following situation. Let P_1, P_2, Q_1, Q_2 be two pairs of very small cancellation periodic orbits, relative to a standard representative $(\Psi_\beta, \Gamma_\beta)$. Assume that there exists a terminal edge $e_i \in E_T(\Gamma_\beta)$ such that the points of P_1, P_2, Q_1, Q_2 satisfy:

$$\begin{aligned} l(p_{1,i}^m) &< l(q_{1,i}^m) < l(p_{2,i}^m) < l(q_{2,i}^m), \\ \omega(p_{1,i}^m) &= \omega(p_{2,i}^m) < \omega(q_{1,i}^m) \leq \omega(q_{2,i}^m), \\ l(q_{1,i}^M) &< l(q_{2,i}^M). \end{aligned}$$

Under these assumptions, and if we set $\gamma = \text{BT}(\mathbf{P}_1, \varphi_\beta)$, $\mu = \text{BT}(\mathbf{P}_2, \varphi_\beta)$, $\nu = \text{BT}(\mathbf{Q}_1, \varphi_\beta)$, and $\delta = \text{BT}(\mathbf{Q}_2, \varphi_\beta)$, then from theorem 6.4 the following relations are satisfied: $\gamma \succ \mu \succ \delta$ and $\gamma \succ \nu \succ \delta$, but μ and ν are unrelated. If this situation occurs, then the two paths $\pi(\gamma, \mu) \cup \pi(\mu, \delta)$ and $\pi(\gamma, \nu) \cup \pi(\nu, \delta)$ form a non trivial loop. The proof that such a situation occurs is given by the proof of lemma 8.3 and 8.4, as soon as there is a pair of very small cancellation periodic orbits with different ω . Therefore this property occurs for some periodic orbits if there is a terminal edge whose transversal word has length strictly greater than one (and thus greater or equal three by lemma 2.10). This property is always satisfied for a pseudo-Anosov class (this is an easy exercise). $\square$

Remark. — From the definition 8.1 of the distance function between small cancellation periodic orbits, it follows that a cycle as above satisfies $d(\mathbf{P}, \mathbf{Q}) \geq 1$. Thus these cycles cannot be too small.

Some further remarks. — The entropy functions has been defined in the introduction. It is clearly an interesting braid type invariant since, by definition, $h[\varphi]$ is decreasing along the paths. The zero entropy cases are important in this description since they forms the “ends” of the genealogy set. The transitions between the positive and the zero entropy parts of the set $\mathcal{G}([\varphi])$ are still unclear, as well as the possible zero entropy subgraphs which are contained in a given $\mathcal{G}([\varphi])$.

In this paper we have focused our attention on the punctured disks cases, because the formulation is easier. But most of the results can be generalized, without any changes, in all the cases where the notion of terminal edges applies. This is the case if S is any compact surface and $\mathbf{P}$ is a periodic orbit of $f \in \text{Homeo}^+(S)$ such that $\text{BT}(\mathbf{P}, f)$ does not belong to the genealogy set $\mathcal{G}([f])$. This is in particular the case if f is isotopic to the identity and $\mathbf{P}$ is a pseudo-Anosov braid type.

There are also many questions from knot theory which are related to the forcing relations described in this paper. Indeed, in the punctured disk case, the genealogy set $\mathcal{G}(\beta)$ consists of conjugacy classes of braids. Furthermore, the results of section 3 give an explicit construction of all these braids. The partial ordering of $\mathcal{G}(\beta)$ gives a partial ordering for closed braids in the solid torus. It would be interesting to understand if this infinite collection of knots (or links) share some particular properties.

REFERENCES

- [ALM] L. ALSEDA, J. LLIBRE, M. MISIUREWICZ, *Combinatorial dynamics and entropy in dimension one*, World scientific, Advanced series in non linear dynamics, vol. 5 (1993).
- [AF] D. ASIMOV, J. FRANKS, Unremovable closed orbits, *Lect. Notes in Math.*, vol. 1007 (1981); revised version (1989).
- [BH1] M. BESTVINA, M. HANDEL, Train tracks for surface homeomorphisms, *Topology* **34** (1995), 109-140.
- [BH2] M. BESTVINA, M. HANDEL, Train tracks and automorphisms of free groups, *Ann. Math.* **135** (1992), 1-51.
- [Bi] J. BIRMAN, *Braids, links and mapping classgroups*, Ann. of Math. Studies **82** (1974), Princeton Univ. Press.

- [BK] J. BIRMAN, M. KIRDWELL, Fixed points of pseudo-Anosov diffeomorphisms of surfaces, *Adv. in Math.* **46** (1982), 73-98.
- [BGM] L. BLOCK, J. GUCKENHEIMER, M. MISIUREWICZ, L. YOUNG, Periodic points and topological entropy of one-dimensional maps, *Lect. Notes in Math.* **819** (1980).
- [Bo] P. BOYLAND, *Notes on dynamics of surface homeomorphisms*, Preprint, Warwick (1989).
- [Di] W. DICKS, Groups, trees and projective modules, *Lect. Notes in Math.* **790** (1980).
- [FLP] A. FATHI, F. LAUDENBACH, V. POENARU, *Travaux de Thurston sur les surfaces*, Astérisque, vol. 66-67 (1979).
- [Fe] J. FEHRENBACH, *Roots of surface homeomorphisms*, Preprint, INLN (1997).
- [FM] J. FRANKS, M. MISIUREWICZ, Cycles for disk homeomorphisms and thick trees, *Contemp. Math.* **152** (1993), 69-139.
- [GST] J. M. GAMBAUDO, S. VAN STRIEN, C. TRESSER, The periodic orbit structure of orientation-preserving diffeomorphisms of D^2 with topological entropy zero, *Ann. Inst. H. Poincaré (phys. théo.)* **49** (1989), 335-356.
- [GL] F. GAUTIER, J. LOS, *Combinatorial suspension for disk homeomorphisms*, Preprint, INLN, 1996.
- [Hal] T. HALL, Unremovable periodic orbits of homeomorphisms, *Math. Proc. Phil. Soc.* **110** (1991), 523-531.
- [Lo1] J. LOS, Knots, braid index and dynamical type, *Topology* **33** (1994), 257-270.
- [Lo2] J. LOS, Pseudo-Anosov maps and invariant train tracks in the disk: a finite algorithm, *Proc. Lond. Math. Soc.* (3), **66** (1993), 400-430.
- [Lo3] J. LOS, On the conjugacy problem for automorphisms of free groups, *Topology* **35** (1996), 779-806.
- [Lu] M. LUSTIG, *Automorphisms, train tracks and non-simplicial $\mathbf{R}$ -tree actions*, Preprint, Bochum (1993).
- [Ma] T. MATSUOKA, The number and linking of periodic solutions of periodic systems, *Invent. Math.* **20** (1983), 319-340.
- [Mo] H. MORTON, Infinitely many fibered knots having the same Alexander polynomial, *Topology* **17** (1978), 101-104.
- [Oe] U. OERTEL, Incompressible branched surfaces, *Invent. Math.* **76** (1984), 385-410.
- [Pe] R. PENNER (with Harer), *Combinatorics of train tracks*, Ann. Math. Studies, **125** (1992), Princeton Univ. Press.
- [Sha] A. N. SHAROVSKII, On the coexistence of cycles of a continuous map of a line into itself, *Ukr. Math. Z.* **16** (1964), 61-71.
- [Shu] M. SHUB, *Stabilité globale des systèmes dynamiques*, Astérisque, vol. 56 (1978).
- [Sm] J. SMILLIES, Periodic points of surface homeomorphisms with zero entropy, *Erg. Theo. Dyn. Syst.* **3** (1983), 315-334.
- [St] J. R. STALLINGS, Topology of finite graphs, *Invent. Math.* **71** (1983), 551-565.
- [Th1] W. THURSTON, On the geometry and dynamics of diffeomorphisms of surfaces, *Bull. A.M.S.* **19** (1988), 415-431.
- [Th2] W. THURSTON, *On the geometry and topology of three manifolds*, Preprint, Princeton (1977).
- [W] R. S. WILLIAMS, One dimensional non-wandering sets, *Topology* **6** (1967), 473-487.

Institut non linéaire de Nice,
 U.M.R. CNRS, 129,
 Université de Nice - Sophia Antipolis,
 1361, route des Lucioles,
 06560 Valbonne.

Manuscrit reçu le 7 mars 1994.