J. Eur. Math. Soc. 9, 801–840

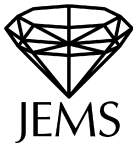

Ivan Dynnikov · Bert Wiest

On the complexity of braids

Received June 20, 2005 and in revised form August 16, 2006

Abstract. We define a measure of "complexity" of a braid which is natural with respect to both an algebraic and a geometric point of view. Algebraically, we modify the standard notion of the length of a braid by introducing generators Δ_{ij}, which are Garside-like half-twists involving strings i through j, and by counting powered generators Δ_{ij}^k as $\log(|k|+1)$ instead of simply $|k|$. The geometrical complexity is some natural measure of the amount of distortion of the n times punctured disk caused by a homeomorphism. Our main result is that the two notions of complexity are comparable. This gives rise to a new combinatorial model for the Teichmüller space of an $n+1$ times punctured sphere. We also show how to recover a braid from its curve diagram in polynomial time. The key rôle in the proofs is played by a technique introduced by Agol, Hass, and Thurston.

Keywords. Braid, curve diagram, complexity, lamination, train track

Contents

I. Dynnikov: Department of Mechanics and Mathematics, Moscow State University, Moscow 119992 GSP-2, Russia; LIFR MIIP, Bol'shoy Vlasyevsky 11, Moscow 119002, Russia;
e-mail: dynnikov@mech.math.msu.su

B. Wiest: IRMAR (UMR 6625 du CNRS), Université de Rennes 1, Campus de Beaulieu, 35042 Rennes Cedex, France; e-mail: bertold.wiest@univ-rennes1.fr

Mathematics Subject Classification (2000): 20F36, 20F65

Introduction

The n-string braid group B_n can be defined in many different ways; we shall mainly be interested in the following two:

Firstly, it has a finite presentation with generating set consisting of Artin's half-twists $\sigma_1, \ldots, \sigma_{n-1}$. (The relations are $\sigma_i\sigma_j = \sigma_j\sigma_i$ for pairs (i, j) such that $|i - j| > 1$, and $\sigma_i\sigma_j\sigma_i = \sigma_j\sigma_i\sigma_j$ for $|i - j| = 1$.)

Secondly, let us denote by D_n a closed disk in the complex plane, centred at the origin, with n punctures lined up on the real line. Then we identify the braid group B_n with the mapping class group of D_n in the standard way [4, 10]. The idea of interpreting braids as isotopy classes of boundary-fixing homeomorphisms of an n times punctured disk D_n is as old as the braid groups themselves. Indeed, Artin [1, 2], who introduced the braid groups, solved the word problem in B_n by showing that a braid $\beta \in B_n$ is uniquely characterized by the images of generators of the fundamental group of D_n under the homeomorphism associated with β.

The aim of this paper is to clarify some aspects of the relation between these two points of view on the braid group. Specifically, we show that there is a notion of *complexity* of a braid which is natural in both frameworks, and which has an added advantage of allowing algorithmically efficient transitions between the two perspectives.

The algebraic point of view is the following: as a generating set, we use the set of all half-twists Δ_{ij} $(1 \leq i < j \leq n)$ involving strands number i through j. Now, the Δ-*length*—our new algebraic notion of complexity—of a braid is simply the shortest possible length of any word representing the braid, but with one modification: a power Δ_{ij}^k $(k \in \mathbb{Z})$ shall not count as having Δ-length $|k|$, as would be usual, but as having Δ-length $\log_2(|k| + 1)$.

The geometric measure of the "complexity" of a braid is as follows. Let E be a set of properly embedded disjoint simple arcs on the disk D_n separating all the punctures (we shall be using the set shown in Fig. 1, but other diagrams would work just as well). Then a *curve diagram* of a braid β is the image $\beta \cdot E$ of E under the homeomorphism β. A simple measure for the complexity of a braid would be the number of intersections of $\beta \cdot E$ with the real line (minimized over the isotopy class of $\beta \cdot E$). However, since in a random braid word this quantity tends to grow exponentially with the length, it is actually more natural not use this quantity itself, but rather its logarithm as a measure of the complexity of β.

The main result of this paper is that the two measures of complexity we have just defined are comparable, in the sense that their ratio is bounded from below and from above by positive constants depending only on n.

In fact, our proof of this result also yields a new, algorithmically very efficient way to calculate a *canonical representative word* for any element of the braid group B_n.

Until recently the most algorithmically efficient treatments of the braid groups were based on purely algebraic ideas that used a presentation of B_n by generators and relations. The approach developed by Garside [15], Thurston [12, Chapter 9], and Birman–Ko–Lee [5] yields an algorithm for finding a canonical form of a braid $\beta \in B_n$ given initially by a word in standard generators. If the input word is of length ℓ, the algorithm in [5] requires at most $O(\ell^2 n)$ operations. Note that Artin's original algorithm is actually very

inefficient in comparison, since the images of the generators of $\pi_1(D_n)$ under the action of the given braid may be of length $O(\exp(\text{const} \cdot \ell))$.

The new algorithm described here, by contrast, is geometrical in flavor and nevertheless efficient. The first crucial idea is that the curve diagram of a braid can be computed efficiently. Indeed, it was pointed out in [11], [10, Chapter 8] that a curve diagram $D = \beta \cdot E$ can be naturally encoded by a $2n$-tuple of integers $\eta(D) = (a_1, b_1, \ldots, a_n, b_n) \in \mathbb{Z}^{2n}$. Moreover, this vector can be computed efficiently: the algorithm constructed there computes the vector $\eta(\beta \cdot E)$ associated to a braid $\beta \in B_n$ of length ℓ in time $O(\ell^2 + n)$ (thus solving the braid recognition problem in time $O(\ell^2)$, but without producing any kind of a canonical word representing the given braid).

Having constructed the curve diagram of β, we then look for an algorithm that, given the curve diagram $\beta \cdot E$ (or its associated vector $\eta(\beta \cdot E) \in \mathbb{Z}^{2n}$) reconstructs a canonical word representing the braid β. Moreover, we want this algorithm to be efficient, and the output braid word to be about as short as possible, in the sense that the Δ-length of the output word surpasses the minimal possible Δ-length, among all representatives of β, by only a linear factor. The centrepiece of this paper is an algorithm which achieves just that.

The idea underlying our algorithm is simply to successively "untangle" the curve diagram $\beta \cdot E$. That is, given a curve diagram D, one can act on it by a generator Δ_{ij}^k of the braid group so as to simplify the diagram (or equivalently, such that the vector $\eta(\Delta_{ij}^k \cdot D)$ is shorter than the vector $\eta(D)$, in an appropriate metric). This process can then be repeated until the diagram E is reached. The braid word one has spelled out during the untangling process is then a representative of β^{-1}. Such untangling is always possible, but it is usually not unique, and the difficulty is to do the untangling in an efficient manner. Our main tool for doing so is a technique introduced by Agol, Hass, and Thurston [3].

A note on history: the term "curve diagram" was introduced in [14], but the basic idea is much older: for instance, it is very explicit in Mosher [22], and indeed it is arguably already present in Artin's original work. The fact that curve diagrams are efficiently computable was popularized by one of us (I.D.) at the *Braids* Colloquium in Toulouse, June 2000, and published later in [11], [10, Chapter 8]. It was also independently discovered by Malyutin [18] in slightly different settings. We are not aware of any literature prior to that. However, the fact that curve diagrams are determined by their intersection numbers with a finite number of curves was well-known before: see e.g. [13, Exposé 6].

The paper is organized as follows. In Section 1 we introduce two measures of complexity of a braid, one geometric and one algebraic, and formulate our main result, that they are comparable. In Section 2.1 we introduce laminations—an important instrument of our constructions. In Section 2.2 we explain a certain special case of Agol, Hass, and Thurston's algorithm for counting the orbits of a collection of isometries of subintervals of $\mathbb{Z}$. In Sections 2.3, 2.4 we show how this technique can be combined with the idea of relaxing integral laminations in order to prove the main theorem. In Section 3.1 we rephrase the main result in terms of quasi-isometries: we introduce two metrics on the braid group B_n corresponding to the two measures of complexity, and prove that they are quasi-isometric. In Section 3.2 we prove that the metric space constructed in the previous section embeds quasi-isometrically in the Teichmüller space of the $n+1$ times punctured sphere, and is in fact quasi-isometric to its so-called thick part. In Section 3.3 we show

that our untangling procedure provides an efficient algorithm for finding σ-consistent braid word representatives. In Section 3.4 we discuss the complexity of our algorithms. At the end of the paper, we discuss some perspectives for further research.

1. Motivating example and statement of the result

As we said in the Introduction, we shall regard braids from B_n as self-homeomorphisms of the punctured disk D_n, which are viewed up to homotopy. The boundary ∂D_n is supposed to be fixed under all homeomorphisms that we consider.

We denote by E the union of $n-1$ arcs in D_n which are shown in Fig. 1 on the left. If β is a braid, then we let $\beta \cdot E$ be the union of arcs obtained from E by the action of β, and we call this the *curve diagram* of β—this is only defined up to isotopies fixing the boundary and the punctures. We recall, however, that by using such an isotopy the curve diagram of any braid can be made *tight* with respect to the horizontal diameter of D_n, meaning that each connected component of $\beta \cdot E$ and the real axis $\mathbb{R}$ are transverse to each other, and there are no puncture-free bigons enclosed by them. Each braid has a unique curve diagram which is tight with respect to the horizontal diameter in the sense that any two such diagrams are related by an isotopy of D_n which preserves the real axis. Details can be found e.g. in [14, 10].

Throughout the paper, all curve diagrams we mention will be assumed tight with respect to the axis unless otherwise specified. We define the *norm* of a curve diagram D to be the number of intersections of D with the real axis:

$$\|D\| = \#(D \cap \mathbb{R}).$$

It is intuitively plausible that in order to create a very complicated curve diagram, one needs a very long braid word. Equivalently, in order to obtain the diagram E by untangling a complicated curve diagram, one needs to act on it by a long braid word. However, there is no simple proportionality relation between length and complexity, as the following example demonstrates.

Example 1.1. Consider the following two braids: $\alpha = \sigma_2^{-1}\sigma_1$ and $\beta = \sigma_2\sigma_1$. The crucial observation now is that the sequence $\|\alpha^k \cdot E\|$ grows exponentially with k, whereas the sequence $\|\beta^k \cdot E\|$ grows only linearly with k. Indeed, it is an exercise to prove that $\|\alpha^k \cdot E\| = 2(F_{k+2} - 1)$, where $F_0 = 1$, $F_1 = 1$, $F_2 = 2, \ldots$ is the Fibonacci sequence. By contrast, we have $\|\beta^k \cdot E\| = 2[(4k-1)/3] + 4$, where $[x]$ stands for the integral part of x. Thus,

$$\|\alpha^k \cdot E\| \sim \text{const}^k, \qquad \|\beta^k \cdot E\| \sim \text{const} \cdot k.$$

Notice that both braids, α^k and β^k, have the same length $2k$ in Artin's generators σ_i, meaning that the shortest word representing any of them has length $2k$.

The reason why there is a principal difference in the growth of the complexity of curve diagrams $\alpha^k \cdot E$ and $\beta^k \cdot E$ is that α is a so-called pseudo-Anosov braid, whereas

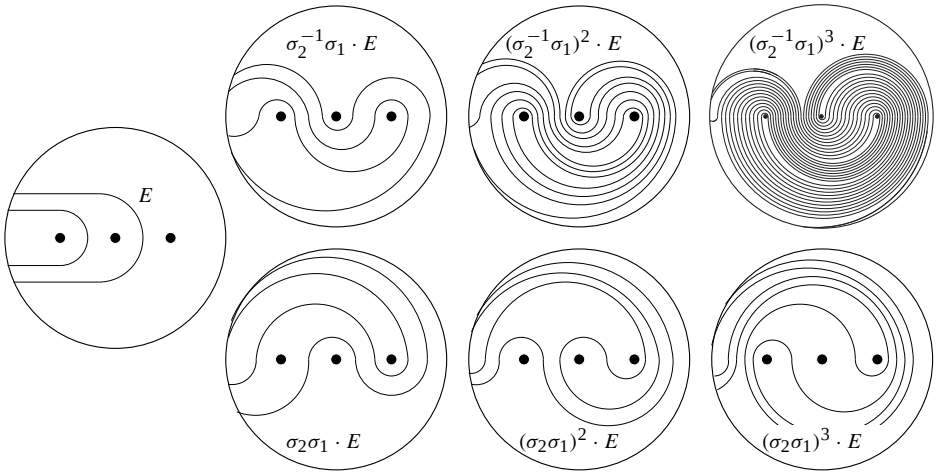

Fig. 1. The diagram E in D_n consists of $n-1$ arcs, each intersecting the real axis once (the case $n=3$ is shown here). The top row shows curve diagrams of the braids α^k, the bottom row of the braids β^k.

$\beta^3 = \Delta^2$, where Δ is the Garside fundamental braid, i.e. the half-twist of all strands at once.

Very roughly speaking, applying repeatedly the same twist Δ entangles the curve diagram much more slowly than applying first one twist, then another one, then again a different one etc.

This example motivates the following modification of the notion of braid length. First, for $1 \le i < j \le n$ let Δ_{ij} denote the half-twist of strands i through j:

$$\Delta_{ij} = (\sigma_i \dots \sigma_{j-1})(\sigma_i \dots \sigma_{j-2}) \dots \sigma_i. \tag{1}$$

This is a generating set of B_n, which contains Artin's standard generators: $\sigma_i = \Delta_{i,i+1}$, and the Garside fundamental braid: $\Delta = \Delta_{1n}$.

Definition 1.1. *By the* Δ-length *of a word* w *of the form*

$$w = \Delta_{i_1 j_1}^{k_1} \dots \Delta_{i_s j_s}^{k_s}, \tag{2}$$

where $k_t \neq 0$ *and* $\Delta_{i_t, j_t} \neq \Delta_{i_{t+1}, j_{t+1}}$ *for all* t*, we shall mean*

$$\ell_\Delta(w) = \sum_{i=1}^{s} \log_2(|k_i| + 1).$$

For a braid $\beta \in B_n$ *we define*

$$\ell_\Delta(\beta) = \min\{\ell_\Delta(w) \mid \textit{the word } w \textit{ represents } \beta\}.$$

Obviously, for any braid β, we have

$$\ell_\Delta(\beta) \le \ell(\beta),$$

where ℓ denotes the ordinary braid length. Note that the Δ-length is in general not an integer.

Definition 1.2. *We define the* complexity *of a braid β as*

$$c(\beta) = \log_2 \|\beta \cdot E\| - \log_2 \|E\|.$$

One of the main results of this paper is the following.

Theorem 1.3. *The complexity and the Δ-length of a braid are comparable. More precisely, there exist constants C_1, C_2 such that*

$$c(\beta) \le C_1 \cdot \ell_\Delta(\beta) \quad \text{and} \quad \ell_\Delta(\beta) \le C_2 \cdot n \cdot c(\beta) \tag{3}$$

for any $\beta \in B_n$.

The constants C_1, C_2 will be given in (4), (5) below. We stress that they are independent of the number of strands n.

Sections 2.1 through 2.4 are devoted to the proof of Theorem 1.3.

2. Proof of the main result

2.1. Integral laminations

Curve diagrams that we consider belong to a wide and very natural class of objects called laminations, which are defined without any reference to the braid groups.

Definition 2.1. *An* integral lamination *in D_n is a non-empty union L of finitely many disjoint simple closed curves and properly embedded arcs in D_n such that*

1) *there are no bigons enclosed by the arcs of L and ∂D_n with no puncture inside;*
2) *no closed component of L bounds a disk with no or just one puncture inside;*
3) *the endpoints of arcs of L are distinct from $\mathbb{R} \cap \partial D_n$ and there is an equal number of them above and below $\mathbb{R}$.*

In particular, conditions 1) and 2) of this definition imply that all curves that can appear as connected components of a curve diagram have the property that they get tangled under the action of some braids. There is only one exception to this rule: a closed curve "parallel" to the boundary of D_n may be present in an integral lamination, even though it does not get tangled under any self-homeomorphism of D_n.

Sometimes we may call integral laminations just laminations for simplicity. The general notion of a lamination will not be needed here.

Integral laminations are regarded modulo isotopy in D_n fixing the boundary. As in the case of curve diagrams, by such an isotopy any integral lamination can be made tight with respect to the axis $\mathbb{R}$.

In what follows, all integral laminations are assumed to be tight with respect to $\mathbb{R}$ unless otherwise specified. In some cases, for technical reasons, we shall consider laminations that are "almost tight" with respect to $\mathbb{R}$. Namely, some laminations L that we consider have exactly one puncture-free bigon enclosed by L and $\mathbb{R}$. However, laminations are always assumed to be transverse to $\mathbb{R}$.

We do not distinguish between two laminations L_1 and L_2 if L_2 is obtained from L_1 by an isotopy of D_n preserving $\mathbb{R}$. In this case we write $L_1 = L_2$. If L_1 and L_2 are related by an isotopy of D_n not necessarily preserving $\mathbb{R}$, we write $L_1 \doteq L_2$. Thus, if both L_1 and L_2 are tight with respect to $\mathbb{R}$, then $L_1 \doteq L_2$ implies $L_1 = L_2$.

The set of isotopy classes of laminations in D_n will be denoted by $\mathcal{L}_n$. The group B_n acts on $\mathcal{L}_n$ in the natural way.

Definition 2.2. *A connected component of an integral lamination is said to be* relaxed *if it intersects the axis $\mathbb{R}$ at most twice. A lamination L is said to be* relaxed *if all connected components of L are relaxed.*

A connected component of a lamination is called even *(respectively,* odd*) if it intersects the axis an even (respectively, odd) number of times.*

If a lamination L has the form $\beta \cdot E$, then each connected component of L is odd. Notice that if a lamination L is very complicated, then it may be quite difficult to decide if all its components are odd, until we have untangled it.

We shall also make use of the following technical definition.

Definition 2.3. *For an integral lamination L in D_n (not necessarily tight with respect to $\mathbb{R}$), by a* closure *of L we shall mean the union $\overline{L}$ of pairwise disjoint simple closed curves in the complex plane such that $\overline{L} \cap D_n$ coincides with L, and $\overline{L} \setminus D_n$ is either empty or consists of arcs intersecting the axis $\mathbb{R}$ exactly once, on the left of D_n. We shall view $\overline{L}$ up to a homeomorphism of the plane preserving the horizontal axis. Clearly, $\overline{L}$ is unique up to such homeomorphisms.*

Curve diagrams as defined in Section 1 are particular cases of laminations. Our basic curve diagram E is an example of a relaxed lamination.

The *norm* of a lamination (not necessarily tight with respect to $\mathbb{R}$) is defined in the same way as for curve diagrams: $\|L\| = \#(L \cap \mathbb{R})$. Notice that $L_1 \doteq L_2$ does not necessarily imply $\|L_1\| = \|L_2\|$, if at least one of the laminations L_1 or L_2 is not tight with respect to the real axis. However, as we have mentioned all laminations L are assumed to be tight if not otherwise specified.

Lemma 2.4. *For any lamination L in D_n, any $1 \le i < j \le n$, and $k \in \mathbb{Z}$,*

$$\|\Delta_{ij}^k \cdot L\| \le (2|k| + 1) \cdot \|L\|.$$

The proof, which is easy, will be left as an exercise to the reader.

Proof of the easy part of Theorem 1.3. Let $\beta \in B_n$ be represented by the word (2). For $t = 0, \dots, s$, put $\beta_t = \Delta_{i_{t+1}j_{t+1}}^{k_{t+1}} \dots \Delta_{i_s j_s}^{k_s}$, so that $\beta_s = 1$ and $\beta_0 = \beta$. Then Lemma 2.4 implies

$$\begin{aligned} c(\beta) &= \log_2 \|\beta \cdot E\| - \log_2 \|E\| = \sum_{t=1}^{s} (\log_2 \|\beta_{t-1} \cdot E\| - \log_2 \|\beta_t \cdot E\|) \\ &= \sum_{t=1}^{s} (\log_2 \|\Delta_{i_t j_t}^{k_t} \cdot (\beta_t \cdot E)\| - \log_2 \|\beta_t \cdot E\|) \\ &\le \sum_{t=1}^{s} \log_2(2|k_t| + 1) \le \log_2 3 \cdot \sum_{t=1}^{s} \log_2(|k_t| + 1) = \log_2 3 \cdot \ell_\Delta(w), \end{aligned}$$

which gives the first inequality in (3) with

$$C_1 = \log_2 3. \tag{4}$$

□

The difficult part of Theorem 1.3, i.e. the second inequality in (3), will be a consequence of the following claim.

Theorem 2.5. *For any integral lamination L in D_n, there exists a braid $\beta \in B_n$ such that $\beta \cdot L$ is relaxed and*

$$\ell_\Delta(\beta) \le 9n \cdot \log_2 \|L\|.$$

Remark 2.1. Note that if a lamination does not have the form of a curve diagram, then there may be more than one braid untangling it.

Thus the constant C_2 in (3) can be set to

$$C_2 = 9. \tag{5}$$

Before embarking on the proof of the theorem, we make a remark on how *not* to prove this result. A naive approach could be to try to proceed by analogy with the proof of the easy part of Theorem 1.3, namely, to show that for a non-relaxed integral lamination L there exists a braid β of the form Δ_{ij}^k such that

$$\log_2 \|\beta \cdot L\| \le \log_2 \|L\| - c \cdot \ell_\Delta(\beta)$$

with some positive constant c independent of L, and then apply induction. Unfortunately, this does not work, because no such constant c exists. In other words, the function

$$\psi(L) = \inf_{\beta = \Delta_{ij}^k;\ \|\beta \cdot L\| < \|L\|} \ell_\Delta(\beta) / (\log_2 \|L\| - \log_2 \|\beta \cdot L\|)$$

is unbounded. For instance, for $L = \Delta_{12}^N \Delta_{34}^N \cdot E$, we find that $\psi(L)$ is comparable to $\log N / (\log 2N - \log N) = \log_2 N$, and thus grows without bound as $N \to \infty$. Moreover, changing the definition of the norm $\|\cdot\|$ in any naive way does not help.

Instead of going this way, we shall apply a very powerful technique due to Agol, Hass, and Thurston [3]: we use a certain complexity function which depends not only on the lamination, but also on the history of the untangling process so far.

2.2. The orbit-counting algorithm of Agol, Hass, and Thurston

In this section, we give a brief account of Agol, Hass and Thurston's technique, adapted to the special case which is relevant for our purposes, and rephrased in a language which is more convenient for us.

By $[i, j]$ with $i, j \in \mathbb{Z}$ we denote the sequence

$$i, i+\epsilon, i+2\epsilon, \ldots, j,$$

where $\epsilon = \pm 1 = \operatorname{sign}(j-i)$, and call such a sequence an *interval*. By the *length* of the interval $[i, j]$ we shall mean the number of elements in it, i.e., $|i-j|+1$.

For two intervals $[i, j]$ and $[k, l]$ of equal length, we define the *interval identification* $[i, j] \leftrightarrow [k, l]$ as the following symmetric relation in $\mathbb{Z}$:

$$i + p\epsilon_1 \leftrightarrow k + p\epsilon_2 \quad \text{for all } p = 0, \ldots, |i-j|,$$

where $\epsilon_1 = \operatorname{sign}(j-i)$ and $\epsilon_2 = \operatorname{sign}(l-k)$. (The interval identification $[i, j] \leftrightarrow [k, l]$ is not distinguished from $[j, i] \leftrightarrow [l, k]$ and $[k, l] \leftrightarrow [i, j]$.)

Definition 2.6. *An* interval identification system *(*IIS *for short) S is an interval $[1, N]$, $N > 0$, endowed with a finite collection of interval identifications*

$$[i_t, j_t] \leftrightarrow [k_t, l_t], \quad t = 1, \ldots, r, \tag{6}$$

within it: $i_t, j_t, k_t, l_t \in [1, N]$. The number N is called the norm *of S and denoted $\|S\|$.*

For such an interval identification system S we denote by $\sim_S$ the equivalence relation in $[1, N]$ generated by the aggregate of all identifications (6). The set $[1, N]/\sim_S$ of equivalence classes will be denoted by Ω_S.

The Agol–Hass–Thurston algorithm that we are going to adapt computes the size of Ω_S in time polynomial in $(\log N, r)$. Note that all naive algorithms one may think of immediately consume time linear in N, which is much worse.

The Agol–Hass–Thurston machinery was originally developed for counting the number of connected components of a normal surface given by its Haken coordinates. We shall apply it to simpler geometrical objects and for quite a different purpose.

Namely, we shall consider the IISs that are naturally associated with (the closure of) a lamination L cut by a ray $(-\infty, a) \subset \mathbb{R}$. The connected components of the cut lamination define an identification relation between the intersection points $L \cap (-\infty, a)$. The formal definition is as follows.

Definition 2.7. *Let $\overline{L}$ be the closure of a lamination L in D_n which is not necessarily tight with respect to the real axis. Let $P_1, \ldots, P_M$, $M = \|\overline{L}\|$ be the intersections of $\overline{L}$ with the axis $\mathbb{R}$, numbered from left to right, and let S be an IIS with $N = \|S\| \le M$. We say that L* carries *S if the following holds:*

1) *$i \leftrightarrow j$ in S if and only if the points P_i and P_j are connected by a segment of $\overline{L}$ not passing through a P_t with $t \le N$;*
2) *for any interval identification $[i_t, j_t] \leftrightarrow [k_t, l_t]$ in S, there is a strip R_t in the complex plane bounded by the straight line segments $P_{i_t}P_{j_t}$, $P_{k_t}P_{l_t} \subset \mathbb{R}$ and two segments*

γ_t, γ_t' of $\overline{L}$ with $\partial\gamma_t = \{P_{i_t}, P_{k_t}\}$, $\partial\gamma_t' = \{P_{j_t}, P_{l_t}\}$. (The arcs γ_t and γ_t' are allowed to coincide, in which case R_t is just an arc.) The bases *$P_{i_t}P_{j_t}$, $P_{k_t}P_{l_t}$ of the strip R_t may overlap and even coincide; besides that, R_t must be embedded;*

3) *the strips R_t are disjoint from each other except at the axis $\mathbb{R}$.*

Not every IIS is carried by a lamination. For example, being carried by a lamination imposes the obvious restriction that every integral point in the interval $[1, N]$ is involved in exactly two interval identifications, which is not true in general. Thus, what we consider is a particular case of the Agol–Hass–Thurston algorithm.

For the rest of the paper, we shall never consider IISs or integral laminations in isolation, but always an IIS S, carried by an integral lamination L. Thus in our situation, it will be convenient to use the geometrical language instead of the combinatorial one. In particular, we shall refer to the elements of Ω_S as connected components of $\overline{L}$ rather than equivalence classes for $\sim_S$.

We shall assume that the whole picture of an integral lamination is rescaled so that the points P_t, $t = 1, \dots, N$, coincide with the integral points $1, \dots, N$ on the real axis. Formally, by *rescaling* we mean a homeomorphism of the plane of the form $(x, y) \mapsto (\varphi(x), y)$, where φ is an increasing function. Clearly, a rescaling preserves the combinatorial structure of a lamination.

We shall also speak of a *strip* $[i, j] \leftrightarrow [k, l]$ instead of an interval identification $[i, j] \leftrightarrow [k, l]$. Notice that such a strip carries a little more information than the corresponding interval identification because in the complex plane it can be attached "from above" and "from below" to the intervals $[i, j]$ and $[k, l]$. We shall always assume that this above-below information is included in the structure of the ISS S.

By the *width* of a strip $R = ([i, j] \leftrightarrow [k, l])$ we shall mean the number of connected components of $L \cap (R \setminus \partial_0 R)$, where $\partial_0 R$ stands for the union of the bases of R. Thus the width of R is equal to the length of the bases $[i, j]$, $[k, l]$ of R, i.e. $|i - j| + 1$. (Note that a strip of width one geometrically looks like a strip of zero width.)

Example 2.1. Fig. 2 shows the closure of the lamination $L = (\sigma_2^{-1}\sigma_1)^2 \cdot E$ endowed with the following interval identification systems:

(a) $N = 26$, $\{[1, 6] \leftrightarrow [12, 7]$, $[13, 19] \leftrightarrow [26, 20]$, $[1, 1] \leftrightarrow [26, 26]$, $[2, 2] \leftrightarrow [3, 3]$, $[4, 14] \leftrightarrow [25, 15]\}$;
(b) $N = 25$, $\{[1, 6] \leftrightarrow [12, 7]$, $[14, 19] \leftrightarrow [25, 20]$, $[1, 1] \leftrightarrow [13, 13]$, $[2, 2] \leftrightarrow [3, 3]$, $[4, 14] \leftrightarrow [25, 15]\}$;
(c) $N = 19$, $\{[1, 6] \leftrightarrow [12, 7]$, $[4, 9] \leftrightarrow [14, 19]$, $[1, 1] \leftrightarrow [13, 13]$, $[2, 2] \leftrightarrow [3, 3]$, $[10, 14] \leftrightarrow [19, 15]\}$;
(d) $N = 14$, $\{[1, 6] \leftrightarrow [12, 7]$, $[4, 4] \leftrightarrow [14, 14]$, $[1, 1] \leftrightarrow [13, 13]$, $[2, 2] \leftrightarrow [3, 3]$, $[5, 9] \leftrightarrow [14, 10]\}$.

In each picture, there are five strips in total, two of them of width one, except in (d), where three strips are of width one.

In all these pictures, the strip $[1, 6] \leftrightarrow [12, 7]$ is attached to both bases from above. The strip $[1, 1] \leftrightarrow [13, 13]$ in (b)–(c) is attached to $\{1\}$ from below and to $\{13\}$ from above.

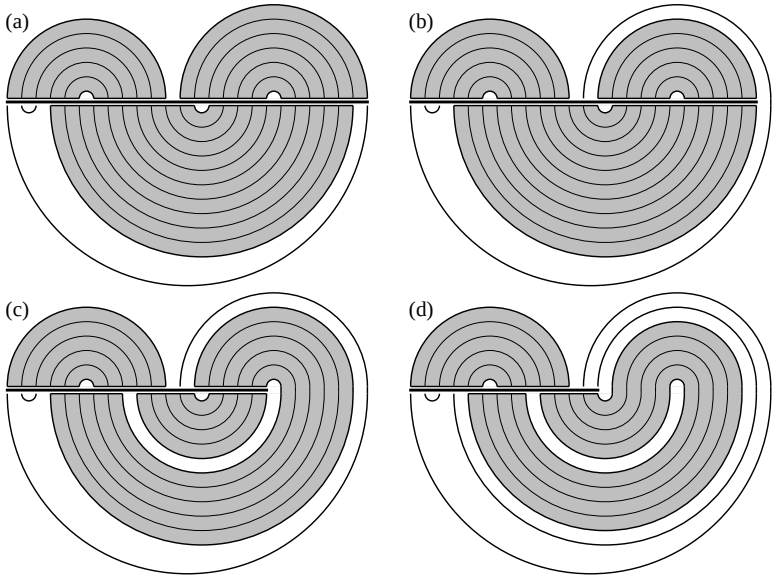

Fig. 2. Four IISs carried by the lamination $(\sigma_2^{-1}\sigma_1)^2 \cdot E$. The sequence (a)→(b)→(c)→(d) is obtained from (a) by successive transmissions.

As this example demonstrates, a lamination may carry many different interval identification systems. If L carries S with given $\|S\| = N$, then the relation $\sim_S$ is completely defined by L. This might suggest that S is uniquely defined by L and N, but this is not necessarily so. Indeed, the structure of S assumes fixing a collection of interval identifications, and therefore one can genuinely change an IIS by replacing an interval identification $[i, j] \leftrightarrow [k, l]$ of width at least two by two interval identifications $[i, p] \leftrightarrow [k, q]$, $[p + \epsilon, j] \leftrightarrow [q + \epsilon', l]$, where $p \in [i, j - \epsilon]$, $\epsilon = \text{sign}(j - i)$, $\epsilon' = \text{sign}(l - k)$. Geometrically, this means that some strip R_t has been cut into two parallel strips. This transformation truly changes the IIS while leaving both N and the underlying lamination L invariant.

The idea of Agol–Hass–Thurston's orbit counting algorithm is to successively simplify an IIS by so-called *transmissions*. In a sense, this algorithm is a generalization of the well-known Euclid algorithm for finding the greatest common divisor of integers.

In our specific situation it works as follows. The input is an interval identification system S which is carried by an integral lamination L. We define a connected component counter, which we set initially to zero. At the "rightmost point" $N = \|S\|$ of the interval $[1, N]$ there are exactly two strips attached, one from below and the other from above. Suppose $R_t = ([i_t, N] \leftrightarrow [k_t, l_t])$ is the wider of the two (no matter which if they happen to have the same width).

It may even happen that those two strips coincide, which means that $k_t = i_t, l_t = N$. In this case, we advance the connected component counter by $N - i_t + 1$ and remove the strip R_t from S. At the same time, we replace N by $i_t - 1$. Intuitively, after this operation

some components of the lamination L are no longer "covered" by any strip of S. This operation is called *removing an annulus*. It is illustrated in Fig. 3(a).

Let us assume now that $l_t \neq N$. In this case we perform a *transmission*—a certain transformation which we are going to explain next. Consider the set X of all subintervals of $[i_t, N]$ that are bases of strips different from R_t. These strips will be referred so as the *denominators* of the transmission, whereas the strip R_t will be called the *numerator*. Let $m \in [i_t, N]$ be the left endpoint of the leftmost interval from X. We cut the strip R_t into a collection of parallel strips so that the base $[i_t, N]$ is cut precisely into subintervals from X, and one more interval $[i_t, m-1]$ provided that $m > i_t$. Then all pairs of strips whose bases have just become matched are stuck together into longer strips, and N is set to $m-1$. This operation, which is called transmission, is illustrated in Fig. 3(b), and examples are given in Fig. 2.

Thus, as a result of a transmission, the numerator and all the denominators are replaced with strips that are all, except at most one, obtained from the denominators by

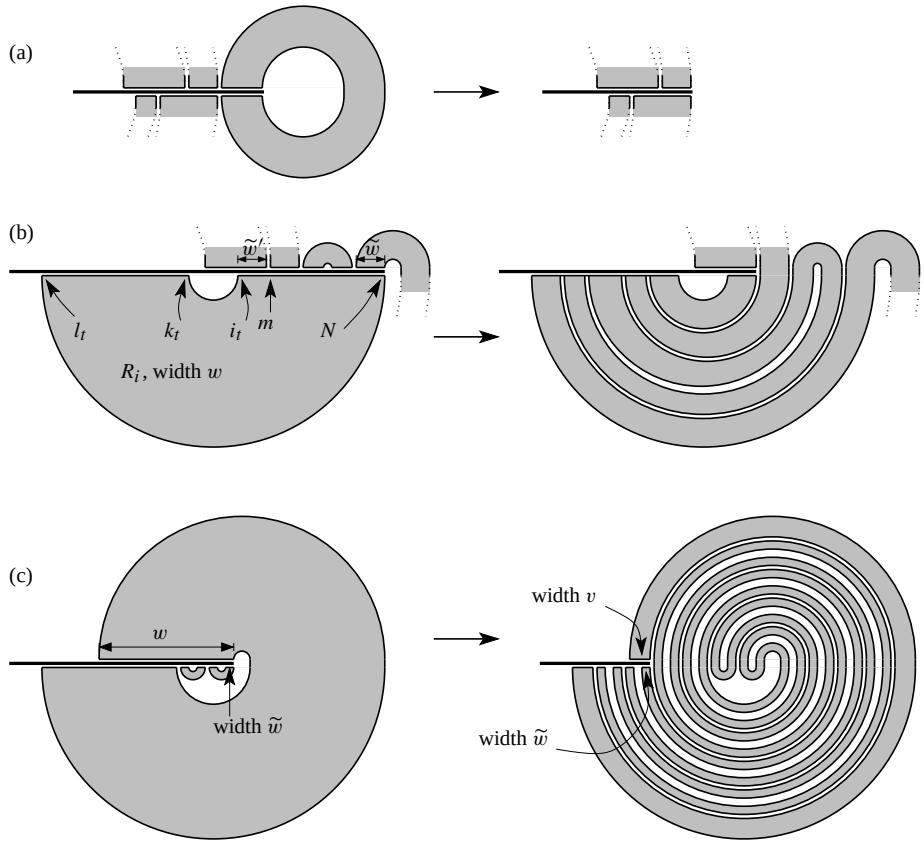

Fig. 3. (a) Removing an annulus; (b) a non-spiralling transmission; (c) a twice-spiralling transmission.

attaching connected components of the cut numerator R_t. If $m > i_t$, then there is one more strip of width $m - i_t + 1$, which we call the *remainder* of the transmission.

There is one exception to the above rule: if the two bases of R_t overlap, i.e. we have $m - 1 = l_t > i_t$, then we can perform the previously explained transmission $d = \left[\frac{N-i_t+1}{N-m+1}\right]$ times at once. This is called a *d-times spiralling transmission*, and it is illustrated in Fig. 3(c). The denominators and the remainder of such a transmission are defined similarly to those in the non-spiralling case. In particular, the width of the remainder equals $(N - i_t + 1) - d(N - m + 1)$.

It is obvious that under a transmission, the number of elements in Ω_S does not change, the norm $N = \|S\|$ decreases, and the number r of interval identifications in S does not increase. Under an annulus removal operation, both N and r decrease, and $|\Omega_S|$ decreases by the value added to the counter. So, after finitely many operations described above, we end up with an empty IIS, and then the connected component counter indicates the number of connected components of the initial IIS.

It is not at all obvious, however, that this procedure is efficient—in particular, that only a relatively small number of transmissions is performed in the process.

Definition 2.8. *The* AHT-complexity $c_{\mathrm{AHT}}(S)$ *of a non-trivial IIS*

$$S = \{[i_t, j_t] \leftrightarrow [k_t, l_t] \mid t = 1, \ldots, r\}$$

is

$$c_{\mathrm{AHT}}(S) = r + \sum_{t=1}^{r} \log_2(|i_t - j_t| + 1) - \frac{1}{2}\log_2 \widetilde{w},$$

where $\widetilde{w}$ is the width of the narrower strip attached to $N = \|S\|$. If S is the trivial ISS, we put $c_{\mathrm{AHT}}(S) = 0$.

Remark 2.2. In this definition, the summand $-\frac{1}{2}\log_2 \widetilde{w}$ looks quite artificial, and it was not present in the original definition by Agol, Hass, and Thurston. Introducing it allows us to prove a better estimate for the simplification effect of a transmission in our specific case of laminations. Note that what we subtract is just one half of one of the summands in the preceding sum, so we count the logarithm of the width of one selected strip with weight one half, whereas all the logarithms of other widths are counted with weight one.

Lemma 2.9. (a) *Performing a non-spiralling transmission on S reduces $c_{\mathrm{AHT}}(S)$ by at least one.*

(b) *Performing a d-times spiralling transmission on S reduces $c_{\mathrm{AHT}}(S)$ by at least* $\log_2(d + 1)$.

Proof. (a) The value $\widetilde{w}$ in Definition 2.8 is the width of the "rightmost" denominator of the transmission to be applied to S. Let $\widetilde{w}'$ be the width of the rightmost denominator in the next transmission step. In addition, let w be the width of the numerator of the transmission.

Suppose that the remainder of the transmission is not trivial. Then its width is exactly $\widetilde{w}'$, and we have $w \geq \widetilde{w} + \widetilde{w}'$. The transmission causes the following change of the AHT-complexity:

$$c_{\mathrm{AHT}}^{\mathrm{old}} - c_{\mathrm{AHT}}^{\mathrm{new}} = \log_2 w - \frac{1}{2}\log_2 \widetilde{w} - \frac{1}{2}\log_2 \widetilde{w}' = \frac{1}{2}\log_2 \frac{w^2}{\widetilde{w}\widetilde{w}'} \geq 1,$$

which follows from the fact that $(a+b)^2 \geq 4ab$ for all $a, b > 0$.

If the remainder is trivial, then the new IIS has a smaller number of interval identifications, so we have

$$c_{\mathrm{AHT}}^{\mathrm{old}} - c_{\mathrm{AHT}}^{\mathrm{new}} \geq 1 + \log_2 w - \frac{1}{2}\log_2 \widetilde{w} + \frac{1}{2}\log_2 \widetilde{w}' \geq 1.$$

(b) Let $w, \widetilde{w}, \widetilde{w}'$ be as before. Suppose that the remainder is non-trivial and let v be its width. Then the two strips attached to the rightmost point of the IIS obtained after the transmission have widths $\widetilde{w}$ and v.

If $\widetilde{w} \leq v$, then $\widetilde{w}' = \widetilde{w}$, and we have

$$c_{\mathrm{AHT}}^{\mathrm{old}} - c_{\mathrm{AHT}}^{\mathrm{new}} = \log_2 w - \log_2 v \geq \log_2(d+1),$$

because $w/v \geq d+1$.

If $\widetilde{w} \geq v$, then $\widetilde{w}' = v$, and we have

$$c_{\mathrm{AHT}}^{\mathrm{old}} - c_{\mathrm{AHT}}^{\mathrm{new}} = \log_2 w - \frac{1}{2}\log_2 \widetilde{w} - \frac{1}{2}\log_2 v \geq \frac{1}{2}\log_2 \frac{(d\widetilde{w}+v)^2}{\widetilde{w}v} \geq \log_2(d+1),$$

since $w \geq d\widetilde{w} + v$ and $(na+b)^2/ab > (n+1)^2$ for all $a \geq b$, $n \geq 1$.

Finally, if there is no remainder, then we have

$$c_{\mathrm{AHT}}^{\mathrm{old}} - c_{\mathrm{AHT}}^{\mathrm{new}} = 1 + \log_2 w - \frac{1}{2}\log_2 \widetilde{w} + \frac{1}{2}\log_2 \widetilde{w}' \geq \log_2 d + 1 \geq \log_2(d+1),$$

since $w \geq d\widetilde{w}$. □

2.3. Relaxing integral laminations

In this section we prove the following claim, which is a "weaker version" of Theorem 2.5.

Theorem 2.10. *For any integral lamination L in D_n, there exists a braid $\beta \in B_n$ such that $\beta \cdot L$ is either relaxed or contains a relaxed even component, and*

$$\ell_\Delta(\beta) \leq 8n^2 \cdot (\log_2 \|L\| + 1).$$

Remark 2.3. For a lamination L that is not a curve diagram, Theorem 2.10 asserts that by a braid of Δ-length $O(\log_2 \|L\|)$ we can "partially" untangle L so that an even component is revealed. It is actually possible to untangle such an L completely by a braid of the indicated Δ-length. The proof of this fact requires more technical details, which we prefer to postpone until the next section.

The proofs of Theorems 2.5 and 2.10 follow essentially the same scheme, but the argument for Theorem 2.10 is more "straightforward". So, in this section, we explain

the main principle that allows us to prove an inequality of the form $\ell_\Delta(\beta) \leq \text{const}(n) \cdot \log_2 \|L\|$, whereas the next section contains details that allow making the const(n) grow as slowly as $O(n)$.

The basic idea is this: we think of our IIS as being made of a rigid horizontal line and a number of rubber-rectangles attached to it, and after each transmission we allow the picture to "relax".

More rigorously, by *relaxing* a lamination L we mean applying a braid β so that the lamination gets simpler, i.e., so as to have $\|\beta \cdot L\| < \|L\|$. For any curve diagram distinct from E, there may be many braids of the form Δ_{ij}^k that relax it, and it is very easy to find at least one of them. However, recursively applying relaxing braids of the form Δ_{ij}^k in a naive way until a relaxed lamination is reached may result in an untangling braid word of length $O(\|L\|)$, because it can be only guaranteed that each relaxation reduces the norm of L by at least some additive constant. The use of the AHT algorithm allows one to make a choice of a relaxation at each step of the untangling process so that the Δ-length of the untangling braid word is of order $O(\log \|L\|)$.

Our algorithm then works as follows: for a given integral lamination L, we construct an IIS S_0 such that L carries S_0 and $\|S_0\| = \|L\|$. Then the construction of the previous section yields a sequence

$$S_0 \overset{d_1}{\longmapsto} S_1 \overset{d_2}{\longmapsto} S_2 \overset{d_3}{\longmapsto} \cdots \overset{d_p}{\longmapsto} S_p = \emptyset, \tag{7}$$

where by $\overset{d}{\longmapsto}$ with $d \geq 2$ we denote a d-times spiralling transmission, by $\overset{1}{\longmapsto}$ a once-spiralling transmission or a non-spiralling transmission, and by $\overset{0}{\longmapsto}$ the removal of an annulus. Lemma 2.9 implies

$$\sum_{i=1}^{p} \log_2(d_i + 1) \leq c_{\text{AHT}}(S). \tag{8}$$

Moreover, for any $i = 1, \ldots, p$ the IIS S_i is still carried by the lamination L.

We put $L_0 = L$, and subsequently find laminations $L_1, \ldots, L_p$ such that, for any $i = 1, \ldots, p$, the following holds:

1) L_i carries S_i;
2) either $L_i = L_{i-1}$ or L_i is obtained from L_{i-1} by a relaxation, $L_i \doteq \beta_i \cdot L_{i-1}$;
3) L_i is the simplest lamination satisfying 1) and 2) (in the sense that it has the minimal norm).

It is not required here that all the laminations L_i are tight with respect to $\mathbb{R}$. We remark, however, that they will be not far from being tight, and most of them will actually be tight. Indeed, the reader who just wishes to understand the principle of the algorithm may safely forget about non-tight laminations.

If L has only odd components, then we end up with

$$L_p = \beta \cdot L = E,$$

where $\beta = \beta_p \beta_{p-1} \ldots \beta_1$. The desired relation between $\ell_\Delta(\beta)$ and $\|L\|$ is then obtained by estimating $c_{\text{AHT}}(L)$ and $\ell_\Delta(\beta_i)$ for all i.

Example 2.2. Consider again the closure $\overline{L}$ of the curve diagram $L = (\sigma_2^{-1}\sigma_1)^2 \cdot E$. Fig. 2(a) shows the corresponding IIS S_0, and from the picture, one can guess the general rule for choosing S_0.

The systems of strips shown in Fig. 2(b), (c), (d) correspond to the lamination-IIS pairs (L, S_1), (L, S_2), (L, S_3), respectively. One can easily see that no braid will simplify L if we require that the lamination obtained still carry S_1 or S_2. This is because the strips in Fig. 2(b), (c) are relaxed, i.e. embedded in the plane "in the optimal way" with respect to the number of intersections with the real axis.

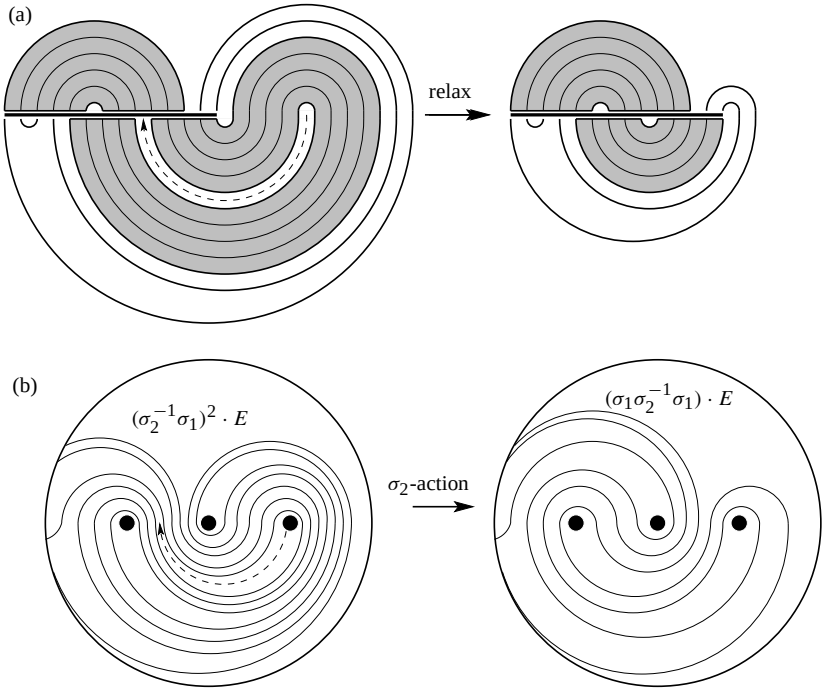

Fig. 4. Part (a) is a continuation of Fig. 2 (the left hand side is the same as Fig. 2(d)). It gives an example of a "relaxation". Part (b) shows the corresponding relaxation of the curve diagram of $(\sigma_1\sigma_2^{-1})^2$.

In Fig. 2(d), by contrast, one of the strips has two "unnecessary" intersections with $\mathbb{R}$, which can be cancelled if we apply the half-twist σ_2. This is illustrated in Fig. 4. Thus, in this example, we have $L_1 = L_2 = L$, $L_3 = \sigma_2 \cdot L$.

Now we give a formal description of the construction. Recall that we assume the whole picture of the lamination to be rescaled so that the closure $\overline{L}$ intersects the real axis in the points $1, \ldots, N = \|\overline{L}\|$.

First of all, we need to define S_0. In order to do so, we cut the lamination $\overline{L}$ along the whole real axis, thus obtaining an identification relation $\leftrightarrow$ on the interval $[1, N]$, where

$N = \|\overline{L}\|$. Then we collect each maximal family of parallel arcs of the cut lamination into a single strip of S_0. In other words, we choose S_0 carried by L so that $\|S_0\| = \|\overline{L}\|$ and S_0 has the minimal possible number of interval identifications (strips).

Lemma 2.11. *We have*

$$c_{\mathrm{AHT}}(S_0) \leq (2n-1)(\log_2 \|L\| + 1).$$

Proof. The width of any strip of any IIS carried by L is obviously no larger than the norm of L (unless L is already relaxed). Therefore, it suffices to show that there are at most $2n-1$ strips in S_0. This can be done in numerous ways, e.g., as follows.

On $S^2 = \mathbb{C} \cup \infty$, take a foliation $\mathcal{F}$ with singularities such that:

1) all connected components of $\overline{L}$ are leaves of $\mathcal{F}$;
2) $\mathcal{F}$ is transverse to the segment $I = (1, N) \subset \mathbb{R}$ except at the punctures;
3) all the singularities of $\mathcal{F}$ are simple (see Fig. 5) and the number of them is minimal provided that there is a singularity of type 1 at ∞.

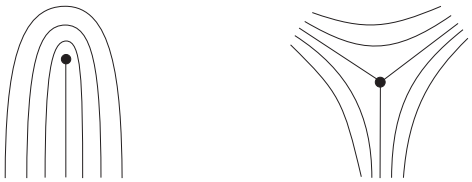

Fig. 5. Singularities of $\mathcal{F}$.

We denote by k_1 and k_2 the number of singularities of type 1 and 2 shown in Fig. 5. Then the singularities of type 1 may occur only at punctures, at the leftmost point of D_n, and at infinity, so we have $k_1 \leq n+2$. Moreover, we have $k_2 = k_1 - 4$, which follows from the Poincaré–Hopf formula $\chi(S^2) = \frac{1}{2}k_1 - \frac{1}{2}k_2$.

If a region complementary to $I \cup \overline{L}$, and not containing ∞, contains $s \geq 1$ singularities of type 2 then it is adjacent to $s+2 \leq 3s$ strips of (L, S) lying on the same side of $\mathbb{R}$. If the outermost region, in which ∞ lies, contains $s \geq 0$ singularities of type 2, then it is adjacent to $s+3 \leq 3s+3$ strips. The sides of strips that are not part of the boundaries of the just mentioned regions are in one-to-one correspondence with singularities of type 1 lying in I. Thus, for the number r of strips, we have

$$2r \leq 3(k_1 - 4) + 3 + (k_1 - 1) \leq 4n - 2,$$

which completes the proof of the lemma. □

Now we proceed with describing the untangling process. Each pair (L, S), where L is a lamination carrying the IIS S, defines a collection of strips. This motivates the following notation:

Definition 2.12. (a) *A* strip system *is a pair* (L, S)*, where* L *is a lamination and* S *is an IIS carried by* L*. (In a strip system* (L, S) *the lamination* L *is not necessarily assumed to be tight with respect to* $\mathbb{R}$*.)*

(b) *For a strip* R_t *of a strip system, we call the number of connected components in* $R_t \setminus \mathbb{R}$ *the* length *of* R_t*.*

(c) *A strip system* (L, S) *is said to be* relaxed *if all strips in it are of length* ≤ 2 *and all the connected components of* L *that are not covered by strips are relaxed.*

Clearly, a strip system is not relaxed if and only if there is a deformation of the complex plane that preserves the bases of the strips and makes at least one of them shorter. The idea of such a deformation is to make the strips tight with respect to the axis $\mathbb{R}$. However, it may be impossible to find such a deformation if we require punctures to be fixed during the deformation. So, in a sense, some punctures provide an obstruction to the relaxation. All the obstructing punctures are located on the right of the right base of the remainder.

Definition 2.13. *Let* α *be a semicircular arc in the complex plane such that* $\partial\alpha = \{P, Q\} \subset \mathbb{R}$*, where* P *is a puncture and* Q *is not. By* sliding the puncture P *along* α *we mean a homeomorphism* φ *of the complex plane such that:*

1) φ *is identical outside a small neighbourhood* U *of* α*;*
2) U *does not contain any other puncture except* P*;*
3) φ *takes* P *to* Q*.*

Viewed up to rescaling, each sliding-a-puncture operation gives rise to a braid. The crucial observation now is that the corresponding braid can be decomposed into two or fewer Δs:

$$\Delta_{ij}^{\epsilon}\Delta_{i,j\pm1}^{-\epsilon},$$

where $\epsilon = \pm 1$ and, by definition, $\Delta_{ii} = 1$. Such a braid is called semicircular in [26].

It is also important to note that in some cases, we can slide a few punctures simultaneously by applying a braid of Δ-length ≤ 2. This occurs if we slide the punctures along

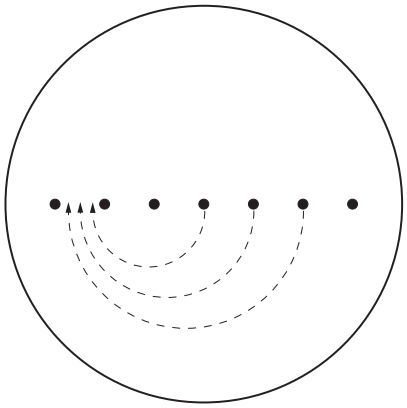

Fig. 6. The sliding of these three punctures is represented by $\Delta_{26}\Delta_{23}^{-1}$ and by $\Delta_{56}^{-1}\Delta_{26}$.

"parallel" arcs, and there are no other punctures between the moving ones and between their destinations (see Fig. 6). The corresponding braid can be represented in the form $\Delta_{ij}^{\epsilon}\Delta_{ij'}^{-\epsilon}$ and in the form $\Delta_{ij}^{\epsilon}\Delta_{i'j}^{-\epsilon}$.

We are now ready to describe the untangling process completely. For a given lamination L, we start by finding the initial IIS S_0. We set $L_0 = L$.

Then, for each $i = 1, 2, \ldots$ we do the following. First, we examine the strip system (L_{i-1}, S_{i-1}), which is relaxed by construction. One of the following situations must occur:

1) $S_{i-1} = \emptyset$, which means that L_{i-1} is relaxed. In this case, we terminate the procedure.
2) An annulus removal operation applies to S_{i-1}. The strip to be removed contains only relaxed components of L_{i-1}. If at least one of them is even, we terminate the procedure. If all of them are odd, we remove the annulus, set $L_i = L_{i-1}$, $\beta_i = 1$, and proceed as before. One can actually see that at this point the lamination is already untangled, so, after removing a few annuli the process will be terminated.

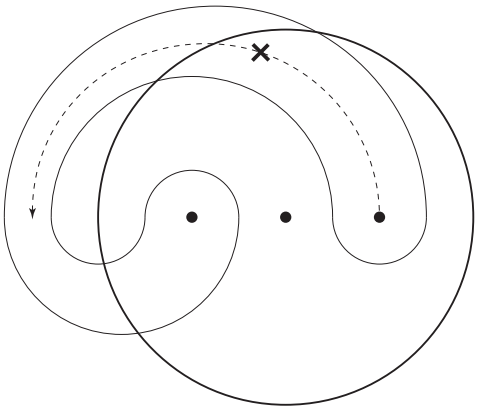

Fig. 7. Punctures should not be slid out of D_n.

3) A transmission $S_{i-1} \overset{d_i}{\longmapsto} S_i$ applies to S_{i-1}. If the strip system (L_{i-1}, S_i) is still relaxed, we set $L_i = L_{i-1}$, $\beta_i = 1$ and proceed as before. If not, we slide the punctures that obstruct the relaxation, along arcs parallel to the arcs of L until they reach the bases of some strips. After that we deform L_{i-1}, keeping the new positions of punctures fixed, so as to reduce the number of intersections with $\mathbb{R}$ on the right of $N_i = \|S_i\|$ as much as possible. (In most cases this just means to make the lamination tight with respect to $\mathbb{R}$ for the new positions of punctures. However, an example of a situation where this is *not* the case is given in Fig. 12.) This replaces L_{i-1} by $L_i \doteq \beta_i \cdot L_{i-1}$, where the braid β_i is obtained by combining all the slidings. The strip system (L_i, S_i) is now relaxed, and we proceed as before. However, there is an exception: if L_{i-1} contains an even non-closed component, then we may not be able to slide some obstructing puncture so as to let the strip system get relaxed (see Fig. 7).

In this case, we apply the slidings until the even component of L gets relaxed in the sense of Definition 2.2 (although the strip system is not yet relaxed in the sense of Definition 2.12) and terminate the procedure. One last note: sometimes there may be more than one way to slide an obstructing puncture. If so, we choose the way that allows moving the puncture farther to the right.

The key ingredient of the proof of Theorem 2.10 is the following bound on the lengths of the braids β_i.

Lemma 2.14. *For all $i = 1, \dots, p$ we have*

$$\ell_\Delta(\beta_i) \le 4n \cdot \log_2(d_i + 1). \tag{9}$$

Proof. The assertion is non-trivial only in the transmission case, $d_i \ge 1$. We consider the non-spiralling case ($d_i = 1$) first. In Fig. 8 all possible types of obstructing punctures are indicated. In each case, we need to slide such a puncture at most twice. Since the number of punctures to be slid is not larger than n, we have in this case:

$$\ell_\Delta(\beta_i) \le 2 \cdot 2 \cdot n = 4n \cdot \log_2(d_i + 1).$$

Now we turn to the case of a spiralling transmission $S_{i-1} \overset{d_i}{\longmapsto} S_i$. Fig. 9 shows how the obstructing punctures should be slid. There are necessarily some punctures that must be slid $2d$ times. If there are no more obstructing punctures and no punctures at the left base of the numerator, the whole spiral can be untwisted by a braid of the form

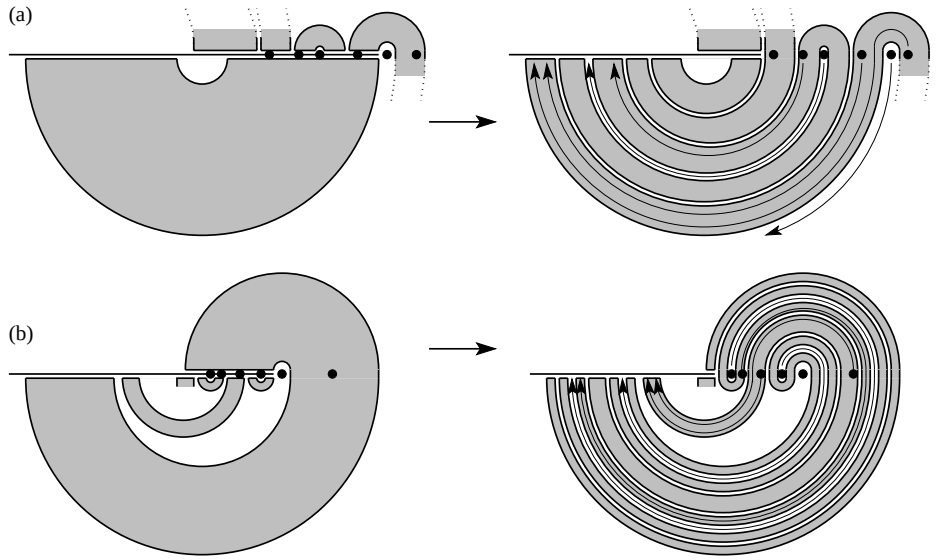

Fig. 8. How to relax after a transmission, in the non-spiralling case.

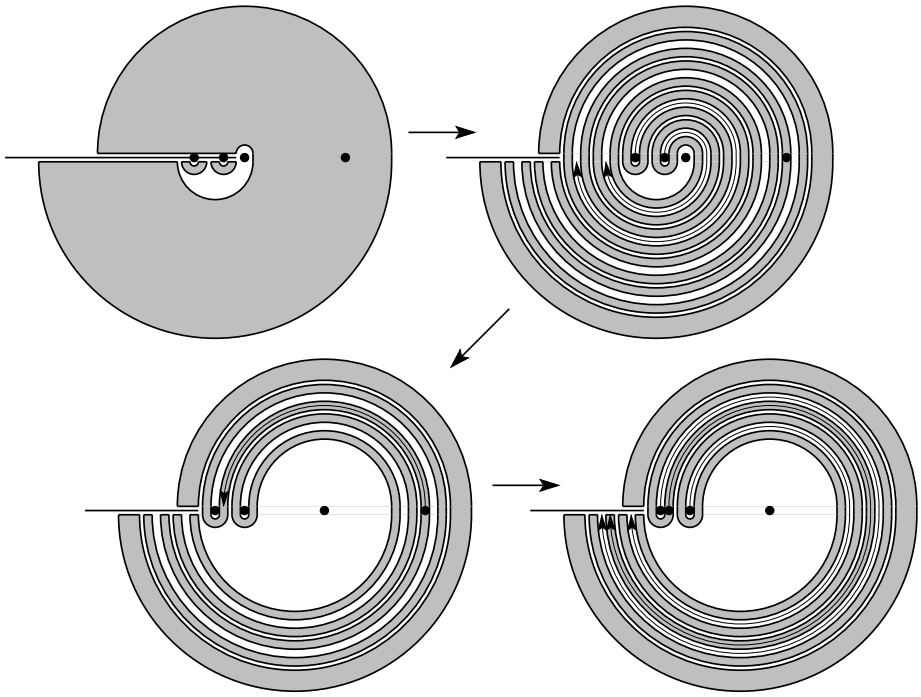

Fig. 9. How to relax after a transmission, in the d-times spiralling case.

$\Delta_{ij}^{\pm 2d}$. There is, however, a complication, if some obstructing punctures are at a smaller distance than $2d$ from their final destination: during the untwisting process, more and more punctures need to be "picked up". For instance, in Fig. 9 we have a twice-spiralling relaxation, but after the first two half-turns of three punctures, a fourth puncture gets picked up.

Let there be obstructing punctures at distances $k_1, \ldots, k_q$ from their final destination, $2d = k_1 > \cdots > k_q \geq 0$, where we regard all punctures within the left base of the numerator also as obstructing (even though they may be at distance zero from their destination). Without loss of generality we may assume that there is exactly one puncture at the centre of the spiral.

In order to deliver all obstructing punctures to their destinations, we first apply the $(k_1 - k_2 - 1)$st power of the half-twist involving the "farthermost" punctures and the central puncture. Now we have to pick up the punctures that are at distance k_2 from their destination. We do this by sliding them back one step. Now they are in a row with the punctures from the first group and all of them are at distance $k_2 + 1$ from their destination. We apply the $(k_2 - k_3)$th power of the half-twist involving all these punctures and the central puncture, then pick up the next portion of punctures, and so forth.

The total number of punctures picked up during the untwisting is at most $n - 2$, and we have $q \leq n - 1$. Picking up each puncture "costs" us two Δs. Thus in this case the

Δ-length of β_i is bounded from above by

$$2(n-2)+\sum_{i=1}^{q}\log_2(k_i-k_{i+1}) \le 2(n-2)+(n-1)\log_2(2d) < 4n\cdot\log_2(d+1),$$

where we set $k_{q+1}=-1$. We leave it to the reader to treat the case when some even component of L gets relaxed during the untwisting of the spiral. □

Proof of Theorem 2.10. This is now an easy corollary to Lemmas 2.11, 2.9, and 2.14: we start with an IIS S_0, carried by the lamination L, of AHT-complexity at most $(2n-1)\cdot(\log_2\|L\|+1)$. Then we apply alternately transmission moves and relaxation moves. During the ith transmission, the AHT-complexity gets reduced by *at least* $\log_2(d_i+1)$, whereas in the subsequent relaxation move a braid of Δ-length *at most* $4n\cdot\log_2(d_i+1)$ is applied. Since the AHT-complexity cannot take negative values, the cycle must stop before a braid word of Δ-length $4n\cdot(2n-1)\cdot(\log_2\|L\|+1)$ has been spelt out. □

Remark 2.4. The untangling process described in this section has been implemented as a `maple`-worksheet by Michel Bonnefont and Erwan Hillion. Their program, which draws pictures of both the curve diagrams and the interval identification systems, is freely available [6].

2.4. Optimizing the untangling procedure

In this section we modify the arguments of the previous section so as to obtain a proof of Theorem 2.5. We use a very similar construction, but modify the definitions of S_0, β_i, and L_i. In order to distinguish from the previous constructions, we add a prime in the notation: S_0', β_i', L_i'. Instead of Lemmas 2.11 and 2.14, we shall get the following estimates:

$$c_{\mathrm{AHT}}(S_0') < 3n\cdot\log_2\|L\|-3n, \tag{10}$$

$$\ell_\Delta(\beta_i') \le 3\cdot\log_2(d_i'+1), \quad i=1,\dots,p', \tag{11}$$

$$\ell_\Delta(\beta_{p'+1}') \le 2n. \tag{12}$$

which together with Lemma 2.9 imply Theorem 2.5. Thus the achievement is to get rid of the factor n in the estimate (11), which is the counterpart of (9), at the expense of enlarging the constant in (10) and getting a more involved construction.

The reason for the factor n appearing in (9) is that we do not know how many punctures we need to slide at each relaxation step, and we estimate the number very roughly by n. The idea is now to move almost all those punctures at once, using the trick indicated in Fig. 6. To this end, we must make sure that there are no punctures in between the destination points, so that the moved punctures do not get shuffled with the others.

Once the new untangling process is described, it is easy, though tiresome, to verify that it works and relations (10), (11), (12) hold. We skip some details of this checking, which contains not much new compared with the previous section. What we do in detail is describing the new rules for relaxing.

First we recall that laminations and their closures that we consider are forbidden to pass through the leftmost point of the disk D_n. The reader might have noticed that the role of this point in our figures is similar to those of punctures. Now it will become even more similar. We call this point the *false puncture* and mark it by $*$ in the figures.

During the untangling process, we shall treat the false puncture almost in the same as a "true" one. Namely, we consider the closure $\overline{L}$ of the initial lamination L as an ordinary lamination in an $(n+1)$-punctured disk D_{n+1} whose punctures are the same as before plus the false puncture.

Denote by ι the inclusion $B_n \to B_{n+1}$ given by $\iota(\sigma_i) = \sigma_{i+1}$. At the ith step of the untangling process, the relaxing braid β'_i will be, in general, a braid from B_{n+1}. However, the resulting braid $\beta'_{p'+1} \dots \beta'_1$ will lie in $\iota(B_n)$: $\beta'_{p'+1} \dots \beta'_1 = \iota(\beta)$. This is achieved by organizing the untangling process so that:

(i) a true puncture is never slid below the false one;
(ii) the false puncture is never slid; the transmission-relaxation procedure is terminated as soon as both bases of the numerator of the transmission to be applied are on the left of $*$, or we get $S'_{p'} = \emptyset$;
(iii) once a true puncture has been moved to the left of $*$, it stays untouched until the final step, when all true punctures that have been slid to the left of $*$ are slid towards the right of $*$ along arcs in the upper half-plane; the additional braid $\beta'_{p'+1}$ does this job.

It is not hard to show (using Lemma 3.10 below) that

$$\ell_\Delta(\beta) \le \sum_{i=1}^{p'+1} \ell_\Delta(\beta'_i). \tag{13}$$

A base of a strip will be called an *A-base* if the strip approaches it from above, and a *B-base* otherwise ('A' stands for 'above' and 'B' for 'below'). To each strip, we associate its *type* that can be either AA, AB, BA, or BB depending on the types of the bases: the first letter indicates the type of the left base, and the second of the right one. If the bases of the strip coincide, it can be thought of as an AB- or BA-strip, this does not matter.

A strip system (L, S) is said to be *almost relaxed* if the length of all its BB-strips is not larger than three, and for all the other strips not larger than two. As before, L is not assumed to be tight with respect to the axis, but all puncture-free bigons enclosed by L and $\mathbb{R}$ must be on the right of $\|S\|$.

We define the new untangling procedure $\overline{L} = L'_0 \mapsto L'_1 \mapsto \cdots$ so as to comply with the following rules:

(iv) $S'_0 \overset{d'_1}{\longmapsto} S'_1 \overset{d'_2}{\longmapsto} S'_2 \overset{d'_3}{\longmapsto} \cdots \overset{d'_{p'}}{\longmapsto} S'_{p'}$ is a sequence of transmissions and annulus removal operations;
(v) for any $i = 0, \dots, p'$, the strip system (L'_i, S'_i) is almost relaxed;
(vi) for any $i = 1, \dots, p'$ we have $L'_i \doteq \beta'_i \cdot L'_{i-1}$ with some $\beta'_i \in B_{n+1}$;
(vii) for any $i = 0, \dots, p'$, the interiors of all strips of (L'_i, S'_i) and their A-bases are free of punctures; the false puncture $*$ is not contained in any (A- or B-) base of a strip.

The IIS S_0 from the previous section (which, we recall, has at most $2n-1$ strips) does not in general satisfy condition (vii). This is because some punctures may sit on the A-bases of strips, and a base of a strip may contain $*$. We resolve this by cutting those strips into a few parallel ones. This results in enlarging the number of strips by at most $n+2$, and one can show that the number of strips will be enlarged exactly by $n+2$ only if it was strictly smaller than $2n-1$ before cutting. So, the number r of strips in S_0' is at most $3n$.

One now obtains (10) as follows. Let $w_1, \dots, w_r$ be the widths of the strips of S_0'. If L is not relaxed, then $r \geq 4$. We also have

$$\|L\| \geq \frac{r}{2}, \qquad \sum_{i=1}^{r} w_i \leq \|L\|.$$

This implies

$$\begin{aligned} c_{\mathrm{AHT}}(S_0') &\leq r + \sum_{i=1}^{r} \log_2 w_i = r + \log_2\Big(\prod_{i=1}^{r} w_i\Big) \\ &\leq r + \log_2\left(\frac{\sum_{i=1}^{r} w_i}{r}\right)^r \leq r + r(\log_2 \|L\| - \log_2 r) \\ &= 3n \cdot \log_2 \|L\| - (3n-r)\log_2 \|L\| - r\log_2 \frac{r}{2} \\ &\leq 3n \cdot \log_2 \|L\| - 3n \cdot \log_2 \frac{r}{2} \leq 3n \cdot \log_2 \|L\| - 3n. \end{aligned}$$

Provided that conditions (v), (vii) above are satisfied up to $i = k-1$, we shall explain how to define β_k'. For simplicity, we will assume that, during the untangling process, no two punctures become immediate neighbours so that the lamination does not traverse the interval between them. One can easily show that this is not a loss of generality, since such two punctures can be treated as a single one.

The notion of obstructing puncture was defined somewhat loosely in the previous section. Now we make it more precise. To this end, consider the strip system (L_{k-1}', S_{k-1}') and the transmission $S_{k-1}' \overset{d_k'}{\longmapsto} S_k'$.

Definition 2.15. *An arc $\alpha \subset L_{k-1}'$ will be said to be* essential *if it satisfies the following conditions: α lies in the lower half-plane, and we have $\partial\alpha \subset \mathbb{R}$; the left endpoint of α is located in the right base of the numerator of the transmission $S_{k-1}' \overset{d_k'}{\longmapsto} S_k'$, but not in the right base of the remainder.*

By definition, an essential arc is contained in the numerator or in a denominator of the transmission $S_{k-1}' \overset{d_k'}{\longmapsto} S_k'$. We also remark that it has both its endpoints to the right of $*$. Our untangling process will be organized as follows: at each step, the essential arcs form a family of parallel, concentric semicircles; in particular, it makes sense to talk about an outermost essential arc. The strips of (L_{k-1}', S_k') that require relaxation after the transmission will be exactly those that contain an essential arc of (L_{k-1}', S_{k-1}'). The relaxation is

achieved by "pushing all essential arcs across the real line". So, by *obstructing punctures* we shall mean those punctures that are located between the endpoints of the outermost essential arc.

Depending on the type of the numerator of the transmission, the following cases are possible:

Case AA. In this case BB-denominators of length one with both bases participating in the transmission cannot occur. Indeed, between those bases there must be a puncture, which contradicts the requirement that all A-bases are free of punctures. Thus, any length one BB-denominator has one of its bases further to the left. Such a denominator gives rise to a length two AB- or BA-strip, which does not need to be simplified.

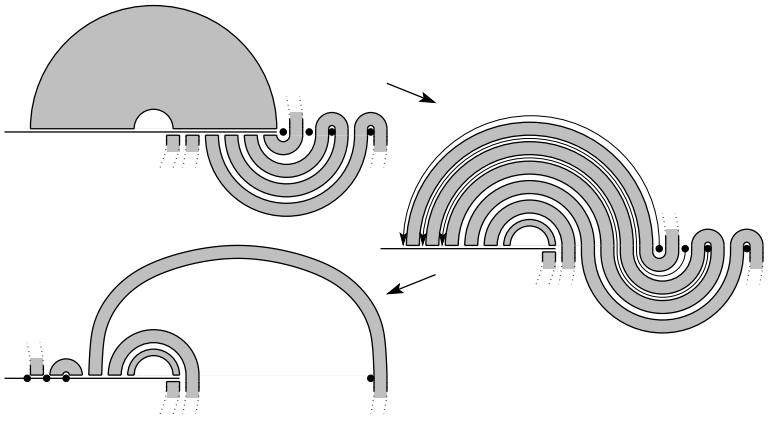

Fig. 10. Case where the numerator is of type AA.

All the other denominators are of AB type and length two, or BB type and length three. The obstructing punctures should be slid along arcs parallel to the denominators toward the right base of the numerator (by one Δ), and then along the numerator toward the left base (two more Δs); see Fig. 10. We make just one exception to this rule: if all denominators are of type BB, then an obstructing puncture positioned on the immediate right of the right base of the numerator (if there is one) does not participate in the second sliding. This is not important for the moment but will be in the proof of Lemma 3.9 below. In this case, we have $\ell_\Delta(\beta'_k) \leq 3$.

Case BB, length one. No relaxation is needed at this point, since every strip that is created during the transmission is of AB type and length two, or of BB type and length three.

Case BB, length three. The obstructing punctures may be inside the right base of the numerator and on the immediate right of that base. They are slid twice along the numerator to the right (see Fig. 11). We have $\ell_\Delta(\beta'_k) \leq 3$.

Case AB, non-spiralling. The obstructing punctures, which are inside and on the immediate right of the B-base of the numerator, are slid twice along the numerator (see Fig. 12). Again, we have $\ell_\Delta(\beta'_k) \leq 3$.

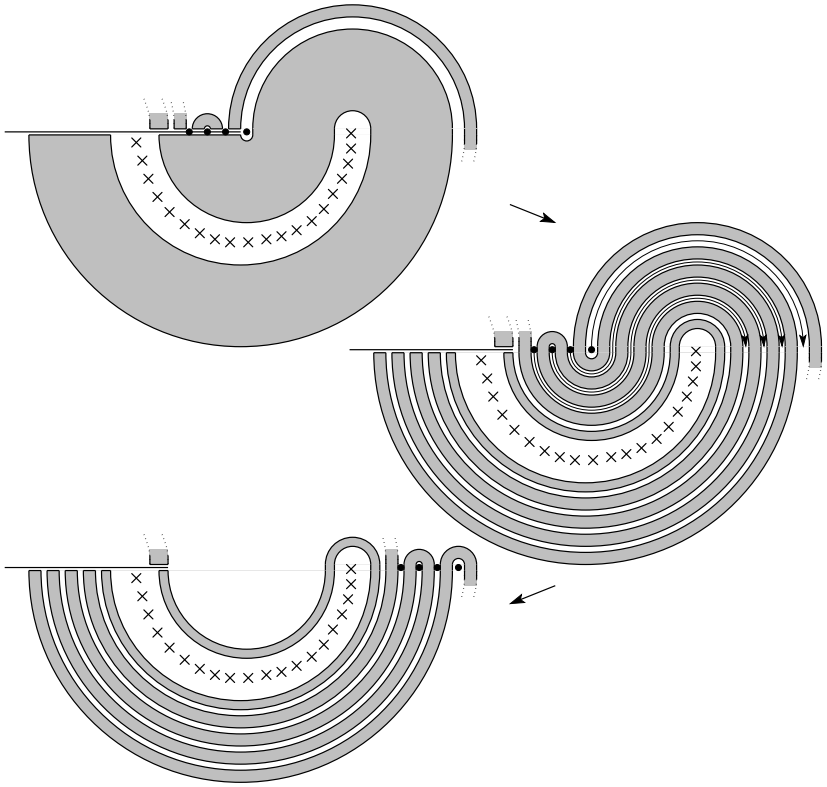

Fig. 11. Case where the numerator is of type BB, length three.

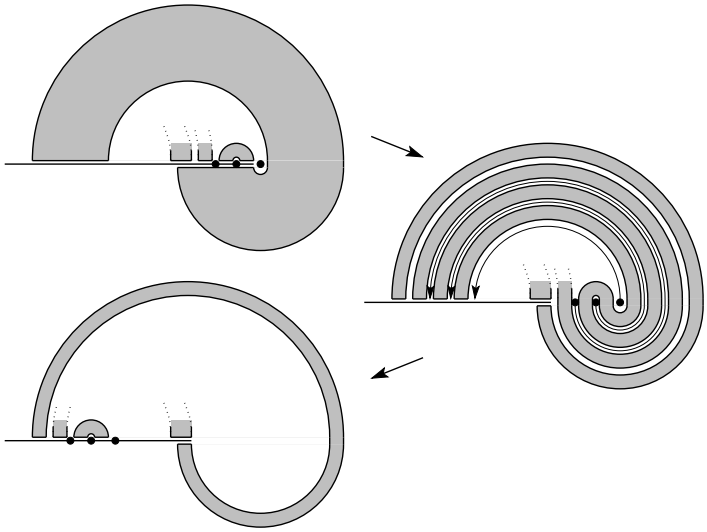

Fig. 12. Case where the numerator is of type AB, non-spiralling.

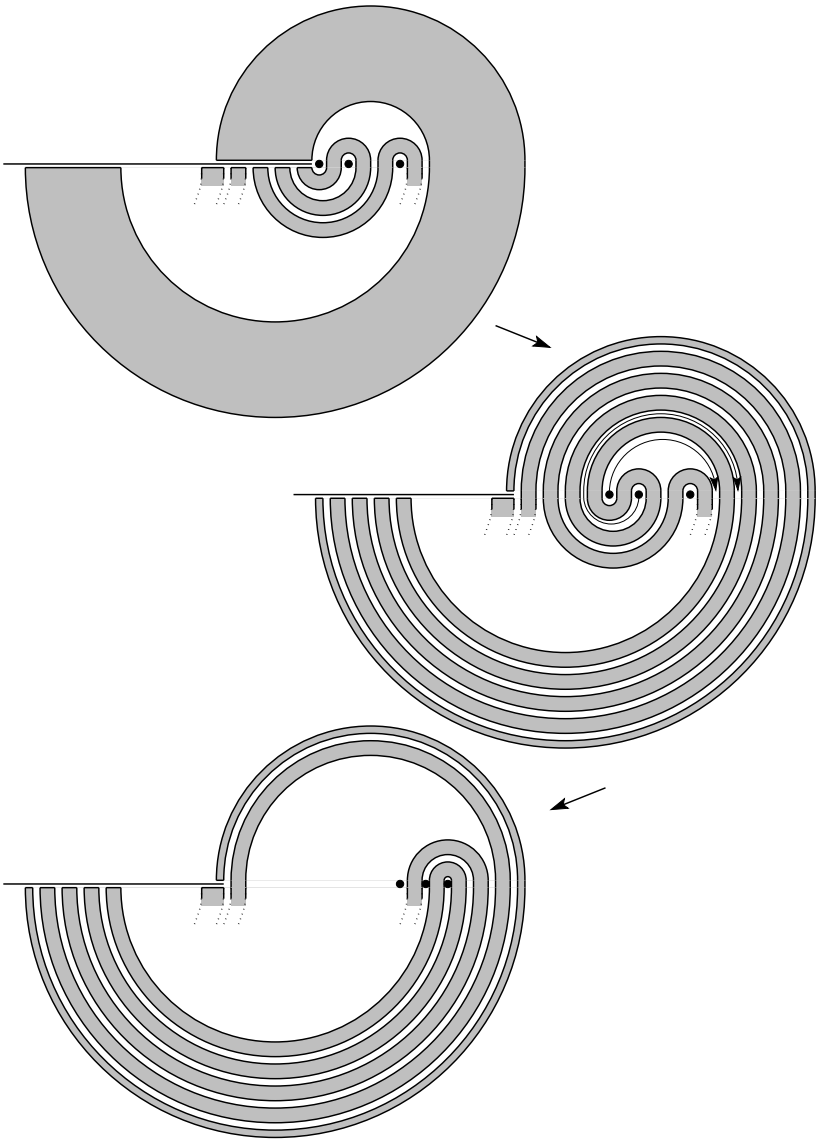

Fig. 13. Case where the numerator is of type BA, non-spiralling.

Case BA, non-spiralling. All the denominators are of BB type. Those of length one must have the other base further to the left. After the transmission, they give rise to BB-strips of length three, which do not need to be relaxed for the moment. The denominators of length three give rise to strips of length five or seven (see Fig. 13). The obstructing punctures are first slid along arcs parallel to essential ones, and then once along the numerator. As always, we have $\ell_\Delta(\beta'_k) \leq 3$.

Spiralling case. The difficulty with punctures that need to be picked up (see the previous section) now disappears, because there are no punctures in the interior of any strips. Thus we can simply apply $\Delta_{ij}^{2d'_k}$ in the BA case and $\Delta_{ij}^{-2d'_k}$ in the AB case, where the half-twist Δ_{ij} involves the punctures inside the spiral. We have $\ell_\Delta(\beta'_k) = \log_2(2d'_k) < 2\log_2(d'_k+1)$. Note that in the BA spiralling case all the denominators are of type BB and of length three. After the relaxation, the strips they give rise to are also of length three.

The transmission-relaxation process is terminated once the transmission "cutting edge" has arrived at $*$. Thus, during the process, whenever an obstructing puncture is slid along an arc in the lower half-plane, the arc will be above the outermost essential arc, and, therefore, on the right of $*$. This guarantees that condition (i) holds.

As a result of the process, some true punctures have moved to the left of $*$. At the very end of the relaxation process, we slide them back along arcs in the upper half-plane so as to get the simplest possible lamination. This yields a braid $\beta'_{p'+1}$ of Δ-length at most $2n$, because at most n punctures need to be slid (actually, it has Δ-length at most n, but even $8n$ would be good enough for our purposes).

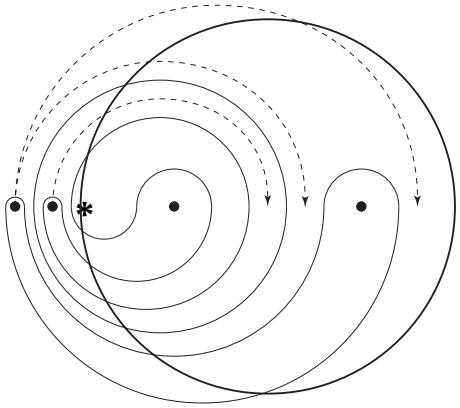

Fig. 14. The braid β'_i returns true punctures into the disk. In the situation shown here, there are two possibilities for the leftmost puncture, and just one possibility for the next one.

We remark that if the original lamination L contains even components, then there may be no canonical choice for $\beta'_{p'+1}$ (see Fig. 14).

This completes our description of the optimized untangling procedure, and thus the proof of Theorem 2.5.

Remark 2.5. In each braid $\beta_i \in B_{n+1}$ just described, there is one strand that corresponds to the false puncture $*$. By removing this strand, we obtain a braid $\beta''_i \in B_n$, and we have

$$\beta = \beta''_{p'+1} \dots \beta''_2 \beta''_1.$$

Consequently, applying braids $\beta''_1, \beta''_2, \dots$ to L may, in principle, eventually result in a more complicated lamination than the original one. This may occur when β''_i corresponds

to sliding a puncture to the left of $*$. So, though the braid β untangles the initial lamination L, the sequence

$$\|L\|,\ \|\beta_1'' \cdot L\|,\ \|\beta_2''\beta_1'' \cdot L\|,\ \dots$$

may not be decreasing.

3. Applications

3.1. Two equivalent metrics on B_n

The Δ-length of a braid defined in Section 1 possesses the following obvious properties:

$$\ell_\Delta(\beta) = \ell_\Delta(\beta^{-1}) \quad \text{and} \quad \ell_\Delta(\beta_1\beta_2) \le \ell_\Delta(\beta_1) + \ell_\Delta(\beta_2).$$

This means that the function

$$\rho_\Delta\colon\ B_n \times B_n \to \mathbb{R}, \qquad \rho_\Delta(\beta_1, \beta_2) = \ell_\Delta(\beta_1^{-1}\beta_2), \tag{14}$$

is a distance on B_n.

The analogous statement for the complexity function is not true: in general, $c(\beta) \ne c(\beta^{-1})$. However, it is true for the following modified complexity function. Put

$$\widetilde{c}(\beta) = \sup_{L\in\mathcal{L}_n} \left|\log_2 \|\beta \cdot L\| - \log_2 \|L\|\right|. \tag{15}$$

By definition, for all $\beta \in B_n$, we have $c(\beta) \le \widetilde{c}(\beta)$ and

$$\begin{aligned}\widetilde{c}(\beta) &= \sup_{L\in\mathcal{L}_n} \left|\log_2 \|\beta \cdot (\beta^{-1} \cdot L)\| - \log_2 \|\beta^{-1} \cdot L\|\right| \\ &= \sup_{L\in\mathcal{L}_n} \left|\log_2 \|L\| - \log_2 \|\beta^{-1} \cdot L\|\right| = \widetilde{c}(\beta^{-1}).\end{aligned}$$

For any $\beta_1, \beta_2 \in B_n$ we have

$$\begin{aligned}\widetilde{c}(\beta_1\beta_2) &= \sup_{L\in\mathcal{L}_n} \left|\log_2 \|\beta_1\beta_2 \cdot L\| - \log_2 \|L\|\right| \\ &= \sup_{L\in\mathcal{L}_n} \left|\log_2 \|\beta_1 \cdot (\beta_2 \cdot L)\| - \log_2 \|\beta_2 \cdot L\| + \log_2 \|\beta_2 \cdot L\| - \log_2 \|L\|\right| \\ &\le \sup_{L\in\mathcal{L}_n} \left|\log_2 \|\beta_1 \cdot L\| - \log_2 \|L\|\right| + \sup_{L\in\mathcal{L}_n} \left|\log_2 \|\beta_2 \cdot L\| - \log_2 \|L\|\right| \\ &= \widetilde{c}(\beta_1) + \widetilde{c}(\beta_2).\end{aligned}$$

Thus the formula

$$\rho_{\mathcal{L}}(\beta_1, \beta_2) = \widetilde{c}(\beta_1^{-1}\beta_2) = \sup_{L\in\mathcal{L}_n} \left|\log_2 \|\beta_1^{-1} \cdot L\| - \log_2 \|\beta_2^{-1} \cdot L\|\right|$$

defines another metric $\rho_{\mathcal{L}}$ on B_n. The proof of the following claim is easy.

Proposition 3.1. *For either of the metrics $\rho_{\mathcal{L}}$ and ρ_Δ, the standard injection $B_n \to B_{n+1}$ is distance-preserving.*

The main result of this section is

Theorem 3.2. *The metrics $\rho_{\mathcal{L}}$ and ρ_Δ on B_n are quasi-isometric, namely, for any $\beta_1, \beta_2 \in B_n$ with $\beta_1 \neq \beta_2$,*

$$\log_3 2 \leq \frac{\rho_\Delta(\beta_1, \beta_2)}{\rho_{\mathcal{L}}(\beta_1, \beta_2)} \leq 9n.$$

Proof. The first inequality is deduced from Lemma 2.4 by analogy with the proof of the easy part of Theorem 1.3. The second inequality follows from Theorem 2.5. □

One can immediately deduce an analogous result for the group $B_n/\langle\Delta^2\rangle$, i.e., the quotient of the n string braid group by its centre. This is the finite index subgroup of the mapping class group of the $n+1$ times punctured sphere S^2_{n+1} consisting of those elements which fix the $(n+1)$st puncture.

The metric ρ_Δ on B_n induces a metric on $B_n/\langle\Delta^2\rangle$, which we shall still denote ρ_Δ. So by definition the ρ_Δ-distance of an element β of $B_n/\langle\Delta^2\rangle$ from the trivial element is $\min_{k\in\mathbb{Z}} \rho_\Delta(\beta\Delta^{2k}, 1_{B_n})$. Similarly, we can define an analogue of the metric $\rho_{\mathcal{L}}$ for $B_n/\langle\Delta^2\rangle$. Geometrically, this means that two laminations on D_n are regarded as equivalent if one can be deformed into the other, where the deformation must preserve ∂D_n setwise, but not necessarily pointwise: all the $n-1$ endpoints of arcs can be slid simultaneously through an angle of $2\pi z$ ($z \in \mathbb{Z}$) along the boundary. This modified notion of equivalence yields a modified notion of complexity of a lamination, and thus an analogue of the metric $\rho_{\mathcal{L}}$ on $B_n/\langle\Delta^2\rangle$. Again, this metric shall still be denoted $\rho_{\mathcal{L}}$. As an immediate consequence of Theorem 3.2 we have

Corollary 3.3. *The metrics $\rho_{\mathcal{L}}$ and ρ_Δ on $B_n/\langle\Delta^2\rangle$ are quasi-isometric.*

3.2. Teichmüller spaces

Among the metrics which are habitually imposed upon the Teichmüller space $\mathcal{T}(S)$ of a surface S, there are notably the *Teichmüller metric* (which can be interpreted in terms of stretch factors of extremal lengths of curves on the surface, see [17]), and W. Thurston's *Lipschitz metric* [25]. The latter metric can be interpreted in terms of stretch factors of hyperbolic lengths of curves on the surface. This interpretation, which is due to Thurston, will be recalled below. A theorem of Choi and Rafi [7] states that the distances $d(\sigma, \tau)$ between two points σ, τ in Teichmüller space according to the two metrics are the same up to an additive constant, provided the two points lie in the *thick part* of Teichmüller space, meaning that they possess no hyperbolic geodesics shorter than the Margulis constant. In particular, the thick parts of Teichmüller space, equipped with these two metrics, are quasi-isometric.

It should be mentioned that the Lipschitz "metric" is not actually symmetric, but could easily be turned into a metric by symmetrizing. Moreover, the Lipschitz metric and its symmetrized version are quasi-isometric on the thick part of Teichmüller space.

We recall that there is a natural action of the mapping class group of S^2_{n+1}, and hence of $B_n/\langle\Delta^2\rangle$, on the Teichmüller space $\mathcal{T}(S^2_{n+1})$, and this restricts to a cocompact action on the thick part $\mathcal{T}_{\text{thick}}(S^2_{n+1})$ of Teichmüller space. Thus for any fixed point σ_* in the thick part, the orbit of σ_* under the action provides an embedding $\Phi\colon B_n/\langle\Delta^2\rangle \to \mathcal{T}_{\text{thick}}(S^2_{n+1})$, $\beta \mapsto \beta\cdot\sigma_*$. Let us now equip $\mathcal{T}_{\text{thick}}(S^2_{n+1})$ with the restriction of the Teichmüller, or equivalently, the Lipschitz metric on the full Teichmüller space (see [7]). The aim of this section is to prove that the pullback metric on the braid group is, up to quasi-isometry, either one of the metrics defined in the previous section. Thus the metric spaces constructed in Section 3.1 turn out to be combinatorial models for the thick part of Teichmüller space.

Proposition 3.4. *The embedding* $\Phi\colon (B_n/\langle\Delta^2\rangle, \rho_{\mathcal{L}}) \to (\mathcal{T}(S^2_{n+1}), d_{\text{Lipschitz}})$ *is quasi-isometric.*

It should be stressed that this result is quite easy to prove, and certainly not original. What is more surprising is that, using Corollary 3.3 and Choi and Rafi's comparison between Teichmüller and Lipschitz metric on Teichmüller space [7], we obtain

Corollary 3.5. *The following four spaces are mutually quasi-isometric:*

(1) $(B_n/\langle\Delta^2\rangle, \rho_\Delta)$, (2) $(B_n/\langle\Delta^2\rangle, \rho_{\mathcal{L}})$,
(3) $(\mathcal{T}_{\text{thick}}(S^2_{n+1}), d_{\text{Lipschitz}})$, (4) $(\mathcal{T}_{\text{thick}}(S^2_{n+1}), d_{\text{Teichm.}})$.

We understand that very similar results were independently obtained by K. Rafi [24] for more general surfaces, using deep theorems like those of Masur and Minsky [19, 20] and Minsky [21].

Proof of Proposition 3.4. We shall use the following notation. If $f, g\colon X \to \mathbb{R}$ are two functions, where X is any set, then we say f and g are *comparable*, and write $f \asymp g$, if there exist constants $C \geq 1$ and $d \geq 0$ such that $(1/C)\cdot g(x) - d \leq f(x) \leq C\cdot g(x) + d$.

Now, for α an isotopy class of simple closed curves in D_n, and σ a hyperbolic structure on D_n (i.e., a point in $\mathcal{T}$) we shall denote by $l_\sigma(\alpha)$ the hyperbolic length of the shortest representative of α, measured in the metric σ. According to Thurston [25], there are two equivalent definitions of the Lipschitz metric, among them the following: if σ, τ are two hyperbolic structures, then

$$d_{\text{Lipschitz}}(\sigma,\tau) = \sup_\alpha(\log(l_\sigma(\alpha)) - \log(l_\tau(\alpha)))$$

where the supremum is taken over all simple closed curves in D_n. In particular, if we try to measure the distance between σ and its translate under the action of a braid β, we obtain

$$d_{\text{Lipschitz}}(\sigma,\beta\cdot\sigma) = \sup_\alpha(\log(l_\sigma(\alpha)) - \log(l_\sigma(\beta\cdot\alpha)))$$

Now we recall the well-known fact that for any fixed point $\sigma_* \in \mathcal{T}$, there exist constants $c, C > 0$ such that

$$c\cdot l_{\sigma_*}(\alpha) \leq \|\alpha\| \leq C\cdot l_{\sigma_*}(\alpha)$$

for all α. That is, $l_{\sigma_*}(\alpha)$ and $\|\alpha\|$ are in bilipschitz correspondence, and in particular, we have $l_{\sigma_*}(\alpha) \asymp \|\alpha\|$. (The reason why this is true is that for *simple* closed geodesic curves α in S^2_{n+1}, equipped with the metric σ_*, there are lower and upper bounds for the lengths of the components of intersection of α with the lower and upper half-planes.)

Moreover, there are global lower bounds on $l_{\sigma_*}(\alpha)$ (namely the Margulis constant) and on $\|\alpha\|$ (namely 2). Thus we can deduce that the logarithms of these quantities are also comparable:

$$\log(l_{\sigma_*}(\alpha)) \asymp \log(\|\alpha\|).$$

Now let $\alpha_1, \dots, \alpha_k$ denote any finite family of simple closed curves with the property that every simple closed curve in S^2_{n-1}, except those enclosing a single puncture, can be obtained from one of the α_is by the action of some braid. We calculate

$$\begin{aligned} d_{\text{Lipschitz}}(\sigma_*, \beta\cdot\sigma_*) &\asymp \sup_{\alpha} (\log(\|\alpha\|) - \log(\|\beta\cdot\alpha\|)) \\ &= \sup_{\zeta\in B_n} \sup_{i=1,\dots,k} (\log(\|\zeta\cdot\alpha_i\|) - \log(\|\beta\zeta\cdot\alpha_i\|)) \\ &\asymp \sup_{\zeta\in B_n} \Big(\sum_{i=1}^{k} (\log(\|\zeta\cdot\alpha_i\|) - \log(\|\beta\zeta\cdot\alpha_i\|))\Big) \\ &\asymp \sup_{\zeta\in B_n} \Big(\log\Big(\sum_{i=1}^{k} \|\zeta\cdot\alpha_i\|\Big) - \log\Big(\sum_{i=1}^{k} \|\beta\zeta\cdot\alpha_i\|\Big)\Big). \end{aligned}$$

We shall fix one very particular choice for the family $\alpha_1, \dots, \alpha_k$, namely the one indicated in Figure 15—in particular, in our choice we have $k = n-1$. With this particular

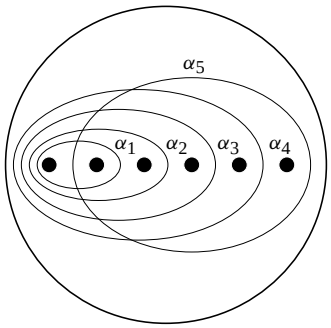

Fig. 15. The curves $\alpha_1, \dots, \alpha_{n-1}$.

choice for the family $\alpha_1, \dots, \alpha_k$ we have the following comparison with the complexity of the curve diagram $\zeta \cdot E$:

$$\|\zeta\cdot E\| - n + 1 \le \sum_{i=1}^{n-1} \|\zeta\cdot\alpha_i\| \le \|\zeta\cdot E\| + n - 1.$$

In particular, we obtain

$$d_{\text{Lipschitz}}(\sigma_*, \beta\cdot\sigma_*) \asymp \sup_{\zeta\in B_n}\Big(\log\Big(\sum_{i=1}^{n-1}\|\zeta\cdot\alpha_i\|\Big) - \log\Big(\sum_{i=1}^{n-1}\|\beta\zeta\cdot\alpha_i\|\Big)\Big)$$
$$\asymp \sup_{\zeta\in B_n}(\log(\|\zeta\cdot E\|) - \log(\|\beta\zeta\cdot E\|)) \asymp \rho_{\mathcal{L}}(1_{B_n}, \beta),$$

which is what we wanted to prove. □

3.3. Dehornoy braid ordering

In 1991 Patrick Dehornoy discovered that the braid group B_n admits a left-invariant ordering [8]. His purely algebraic construction was based on the following notion of σ-positive braid word.

Definition 3.6. *A braid word w is said to be* σ_k-positive *(respectively,* σ_k-negative*) if it contains σ_k, but does not contain σ_k^{-1} and $\sigma_i^{\pm1}$ with $i < k$ (respectively, contains σ_k^{-1}, but not σ_k and $\sigma_i^{\pm1}$ with $i < k$). If w contains no $\sigma_i^{\pm1}$ with $i \le k$, it is called* σ_k-neutral.

A braid word w is said to be σ-positive *(respectively,* σ-negative*) if it is σ_k-positive (respectively, σ_k-negative) for some $k \le n-1$. A braid word w is said to be* σ-consistent *if it is either trivial, σ-positive, or σ-negative.*

Theorem 3.7 (Dehornoy [8])**.** *For any braid $\beta \in B_n$, exactly one of the following is true:*

1) *β is trivial;*
2) *β can be represented by a σ_k-positive braid word for some k;*
3) *β can be represented by a σ_k-negative braid word for some k.*

In the latter two cases k is unique.

Thus, it makes sense to speak about σ-positive and σ_k-positive (or σ-, σ_k-negative) braids. It is then an immediate consequence that the relation $<$ on B_n defined by the rule: $\beta_1 < \beta_2$ if and only if $\beta_1^{-1}\beta_2$ is σ-positive, is a left-invariant linear ordering on B_n (see [8]).

It was noticed in [14] that this notion of σ-positivity has a nice geometric interpretation in terms of curve diagrams. We refer the reader to the monograph [10] for a survey of this and many other explanations of the phenomenon and different proofs of Dehornoy's theorem.

Although many approaches to σ-ordering have been developed since Dehornoy's discovery, the following question remains unsettled: is there a polynomial in ℓ which is an upper bound on the length of the shortest σ-consistent braid word representing a braid of length ℓ? Dehornoy's original algorithms (in [8], and handle reduction [9]) and the algorithm from [14] give only an exponential bound on the length of the shortest σ-consistent representative.

At the end of the paper we shall present some further reasons for believing that a linear bound exists. The aim of the current section is to solve a closely related problem, namely, to give a positive answer to the question above with the ordinary braid length replaced

by the Δ-length. This assumes the following extension of the notion of σ-positive braid word: a word in the alphabet $\{\Delta_{ij}\}_{0<i<j<n}$ is said to be σ-positive if, for some $k < l$, it contains Δ_{kl}, and contains neither Δ_{kj}^{-1} nor $\Delta_{ij}^{\pm 1}$ with $i < k$ and any j. In other words, a word w in letters Δ_{ij} is σ-positive (negative, neutral) if the word in standard generators σ_i obtained from w by expansion (1) is.

Theorem 3.8. *Any braid $\beta \in B_n$ can be represented by a σ-consistent word w in the alphabet $\{\Delta_{ij}\}$ such that*

$$\ell_\Delta(w) \leq 30n \cdot \ell_\Delta(\beta). \tag{16}$$

The following lemma plays a key rôle in the proof. Denote by E_2 the lamination in D_n whose closure is a circle in the complex plane surrounding the leftmost puncture and the leftmost point $*$ of D_n. (The notation is motivated by the fact that this circle coincides with the trivial curve diagram E in the case $n = 2$.)

Lemma 3.9. *Let $\beta \in B_n$ be a σ_1-positive braid, and let u be the braid word spelt out by the untangling procedure of Section* 2.4 *applied to the lamination $\beta \cdot E_2$. Then the braid word u is σ_1-negative, the braid $u \cdot \beta$ is σ_1-neutral, and*

$$\ell_\Delta(u) \leq \log_2 3 \cdot 9n \cdot \ell_\Delta(\beta). \tag{17}$$

Proof. First of all, we remark that the algorithm of Section 2.4 constructs a sequence of puncture slidings rather than a concrete braid word. In some cases such a sliding can be written in two different ways; an example is shown in Fig. 6. More precisely then, what we are going to prove is that, *under an appropriate choice* of the decomposition of the slidings into a product of Δs at each step of the untangling procedure, we get a σ_1-negative word u.

A sliding of a group of punctures, like the one shown in Fig. 6, is not σ_1-neutral if and only if one of the following occurs: the leftmost puncture is slid; or punctures are slid over or under the leftmost (true) one. If none of these takes place, it does not matter which decomposition, $\Delta_{ij}^{\epsilon}\Delta_{ij'}^{-\epsilon}$ or $\Delta_{ij}^{\epsilon}\Delta_{i'j}^{-\epsilon}$, we choose—both are σ_1-neutral.

If a sliding of a group of punctures is not σ_1-neutral, then its decomposition of the form $\Delta_{ij}^{\epsilon}\Delta_{i'j}^{-\epsilon}$ is σ_1-definite, and it is σ_1-positive (respectively, negative) if and only if the punctures are slid clockwise (respectively, counterclockwise).

Thus, in order to prove that u is σ_1-negative it suffices to show that all the clockwise slidings that occur during untangling the lamination $\beta \cdot E_2$ are σ_1-neutral, but the whole word u is not σ_1-neutral.

By construction, the braid $u \cdot \beta$ preserves E_2, which implies that this braid is σ_1-neutral. Since β is not σ_1-neutral, the braid represented by u cannot be σ_1-neutral, either. Let us check that σ_1-positive slidings do not appear in u.

To this end, we must examine all the clockwise slidings and make sure that they do not involve the leftmost puncture, i.e. the arcs of all clockwise sliding are always on the right of the leftmost puncture.

Similarly to [14] one can show that σ_1-positivity of the braid β is equivalent to the following property of the lamination $L = \beta \cdot E_2$, which is assumed to be tight with respect to $\mathbb{R}$:

there is an arc $\gamma \subset \overline{L}$ lying in the lower half-plane with endpoints at $\mathbb{R}$ such that the segment of $\mathbb{R}$ bounded by the endpoints of γ contains $*$ and does not contain a true puncture. $(\star)$

See also [10] for an explanation of the relation between σ-positivity and a $(\star)$-like property of curve diagrams.

The lamination $\overline{L}$ intersects the ray $(-\infty, *) \subset \mathbb{R}$ exactly once, at the left endpoint of γ. (This means, in particular, that there is only one arc γ satisfying $(\star)$.) This implies that no puncture is slid out of the disk D_n during untangling L. Indeed, in the general case, punctures can be slid out of the disk at the relaxation step that follows an AA- or AB-transmission. One can see that, in both situations, any arc in the upper half-plane along which a puncture is slid counterclockwise must lie in between two parallel arcs of the lamination. Thus, sliding a puncture out of the disk may be forced only if the closure of the lamination L intersects the ray $(-\infty, *)$ at least twice, which is not the case for $L = \beta \cdot E_2$.

Thus, the braid $\beta'_{p'+1}$ that is composed of clockwise slidings and is a potential source of a σ_1-positive contribution to u is trivial.

Further, we claim that property $(\star)$ holds during the whole untangling process. Indeed, the arc γ cannot be essential, so it always remains untouched. Thus property $(\star)$ is violated only when a true puncture is slid to a point in between the endpoints of γ, which is the moment when the lamination gets untangled completely.

Now we can show that property $(\star)$ guarantees that no clockwise sliding involves the leftmost puncture. That is, we have to exhibit a true puncture to the left of all arcs along which clockwise slidings occur. This is done by revising, case by case, the relaxations following different types of transmissions. The AB spiralling and non-spiralling cases are trivial because the corresponding slidings are counterclockwise. In all the other cases, just before the transmission, there must be an AA-strip whose left base is further to the left than the right base of the numerator. Let R be the innermost such strip (R may be the numerator in the AA case).

According to our transmission-relaxation rules, there must be a puncture (true or false) P, say, between the bases of R. In each individual case it is easy to show that P is not obstructing and that it lies to the left of the arcs along which clockwise slidings occur. Thus if P is a true puncture, then the clockwise slidings are σ_1-neutral. We conclude by noticing that P cannot be the false puncture because otherwise property $(\star)$ would imply the presence of a circle surrounding just $*$, which is absurd.

It remains to prove estimate (17). This is done by applying the first inequality in Theorem 1.3 and Theorem 2.5, and using the fact that u is precisely the braid word for which the estimate from Theorem 2.5 has been obtained:

$$\ell_\Delta(u) \le 9n \cdot \log_2 \|L\| \le 9n \cdot (\log_2 3 \cdot \ell_\Delta(\beta) + \log_2 \|E_2\|) = 9n \cdot \log_2 3 \cdot \ell_\Delta(\beta),$$

since $\|E_2\| = 1$. □

Lemma 3.10. *If a braid β is σ_1-neutral then the shortest braid word representing β (where either the ordinary braid length or the Δ-length is used) is σ_1-neutral.*

Proof. If β is σ_1-neutral, then the string that starts in the leftmost position also ends in the leftmost position. Now if w is any braid word representing β, then we can obtain another braid word $\widetilde{w}$ representing β which is σ_1-neutral and at most as long as w in the following way: we delete the string starting and ending in the leftmost position, and replace it by a string that lies entirely to the left of all the other (true) strings. The word $\widetilde{w}$ is then obtained from w by removing some letters $\sigma_i^{\pm 1}$, shifting indices of others: $\sigma_i^{\pm 1} \mapsto \sigma_{i-1}^{\pm 1}$, $\Delta_{ij}^p \mapsto \Delta_{i-1,j-1}^p$, and replacing some Δ_{ij}^p with $\Delta_{i,j-1}^p$. If w was not σ_1-neutral, then $\widetilde{w}$ is strictly shorter than w. □

Proof of Theorem 3.8. Without loss of generality we may assume that β is σ_k-positive with some $k < n$. The proof is by induction on k. We skip the details of the induction step, which is an easy consequence of Lemma 3.10, and show that the assertion of the theorem holds for a σ_1-positive braid.

Indeed, take u from Lemma 3.9. Since the braid $\beta' = u \cdot \beta$ is σ_1-neutral, Lemma 3.10 implies that β' can be represented by a σ_1-neutral braid word v of Δ-length

$$\ell_\Delta(v) \le \ell_\Delta(u \cdot \beta) \le \ell_\Delta(u) + \ell_\Delta(\beta) \le (\log_2 3 \cdot 9n + 1) \cdot \ell_\Delta(\beta).$$

The braid word $w = u^{-1}v$ is then σ_1-positive, it represents β and has Δ-length

$$\ell_\Delta(w) \le \ell_\Delta(u) + \ell_\Delta(v) \le (2 \log_2 3 \cdot 9n + 1) \cdot \ell_\Delta(\beta) \le 30n \cdot \ell_\Delta(\beta). \qquad \square$$

3.4. Algorithmical issues

The proof of Theorem 2.5 presented above consists in an explicit description of an algorithm that, given an integral lamination L, finds a word-representative of a braid that untangles L. In this section we discuss the efficiency of the algorithm and, more generally, of the algorithmical treatment of the braid groups based on it.

First of all, we remark that our estimations for the running time of algorithms will be made for the computational model called *Random Access Memory Machine*. Roughly speaking, this means that we assume the input to be in a reasonable range, and estimate the number of elementary operations of a realistic computer needed to implement the algorithm. More precisely, we assume that the number of strands n is "small enough", so that its record fits one standard unit of memory, and that any arithmetic operation on integers between 0 and n takes constant time. This is a reasonable assumption because actually it allows n to be very large: if, say, four bytes are used to store an integer (which is quite usual), then n can be as large as 2^{31}. For other integers (which are not indices) we will assume that their logarithm is "small" (i.e. smaller than 2^{31}).

Since even for reasonably long braid words the implementation of our algorithms may need to operate with "large" integers, we will pay attention to the number of elementary machine operations that are needed to perform an arithmetic operation on those integers. The most frequently used operations will be addition, subtraction, and comparison. They consume logarithmic time in the value of the larger operand, and we call them *simple*. Sometimes we will also need to perform divisions and multiplications. We use the fact

that dividing k by l takes $O(\log l \cdot \log(k/l))$ elementary operations, whereas multiplying k by l consumes $O(\log k \cdot \log l)$ operations.

The next important question is how the input/output data and the objects used in the algorithm are presented. Our main objects are: braid words, laminations, and strip systems, so we briefly discuss their numerical presentations.

We will assume that words of the form (2) are presented by the corresponding sequences of integers:

$$s; i_1, j_1, k_1, i_2, j_2, k_2, \ldots, i_s, j_s, k_s.$$

Here s and i_t, j_t with $t = 1, \ldots, s$ are "small" integers, whereas k_t, $t = 1, \ldots, s$, can be "large". This implies that the length of such a presentation of w is of order $O(\ell_\Delta(w))$.

The set $\mathcal{L}_n$ of laminations in D_n can be naturally identified with $\mathbb{Z}^{2n-2} \setminus \{0\}$ so that the norm $\|L\|$ will become a norm in $\mathbb{Z}^{2n-2}$, and the action of each generator $\sigma_i^{\pm 1} \in B_n$ will be given by a finite number of simple arithmetic operations on the coordinates of the lamination. See [10, Chapter 8] and [11] for details, where a slightly different definition of laminations is used, which results in two additional coordinates appearing in the "code" of a lamination. (In order for the formulas in [11, 10] to work in our current settings, one should set the two additional coordinates to $a_n = 0$, $b_n = +\infty$.) This implies the following

Proposition 3.11. *There exists an algorithm A_1 that, given a word w in the generators Δ_{ij} representing a braid β, computes the curve diagram $\beta \cdot E \in \mathcal{L}_n = \mathbb{Z}^{2n-2} \setminus \{0\}$ in time $O(\ell(w) \cdot \ell_\Delta(w) + n)$.*

The algorithm A_1 expands the given word w by using (1), thus obtaining a word w' in $\sigma_i^{\pm 1}$-generators recorded in the usual way. Then it generates the initial lamination E and applies, one after another, the letters of w' (from right to left) to the lamination.

This may be very inefficient if w contains a subword Δ_{ij}^N with a very large N. However, the action of the braid Δ_{ij}^N on a lamination L can be computed without expanding the braid into a product of σ_is.

Lemma 3.12. *The action of Δ_{ij}^N on L can be computed in $O((n + \log N) \cdot \log \|L\|)$ operations.*

Proof. Let us look at the sequence of laminations $\Delta_{ij}^k \cdot L$, where $k = \ldots, -2, -1, 0, 1, 2, \ldots$. For large $|k|$ the laminations $\Delta_{ij}^k \cdot L$ have a big "spiral" surrounding punctures i through j, and the "thickness" of this spiral grows linearly with k. More formally, this means the following.

Let k_0 be an integer for which $\|\Delta_{ij}^{k_0} \cdot L\|$ is as small as possible. Let us cut L along the real axis and count the number of the resulting arcs having one endpoint between the ith and jth punctures, and the other endpoint outside this segment. In a sense, this is twice the number of strings involved in the spirals of $\Delta_{ij}^N \cdot L$. Let this number be m and let L_0 be the lamination consisting of m circles surrounding punctures i through j. Let us think of laminations as points in $\mathbb{Z}^{2n-2}$. Then for any $p \geq 1$,

$$\Delta_{ij}^{k_0+2p} \cdot L = \Delta_{ij}^{k_0+2} \cdot L + (p-1) \cdot L_0, \qquad \Delta_{ij}^{k_0-2p} \cdot L = \Delta_{ij}^{k_0-2} \cdot L - (p-1) \cdot L_0.$$

So, we start by establishing the structure of a spiral in L surrounding punctures i through j, if there is one. Even if there is no spiral, we compute m, i.e. the lamination L_0. If there is a spiral, we also need to find its "thickness" θ and its direction (clockwise or counterclockwise).

By using flips of triangulations in a similar way as described in [10], we can do all this job in $O(j-i) \le O(n)$ simple operations on integers of order $\|L\|$. So, the structure of the spiral can be discovered for $O(n \cdot \log \|L\|)$ elementary operations.

Depending on the direction of the spiral and the sign of N different cases are possible. It may happen that Δ_{ij}^N twists the spiral further, in which case we are lucky, because we have $\Delta_{ij}^N \cdot L = L + N/2 \cdot L_0$, provided that N is even. If N is odd, we shall also need to apply one Δ_{ij} "explicitly", which, by the same "flip argument", takes $O(n)$ simple operations on integers of order $\|L\|$. So, the total work in this case is $O(n \cdot \log \|L\|) + O(\log N \cdot \log \|L\|)$, where the second summand appears because we need to multiply L_0 by $N/2$. (Note that L_0 has only two non-zero coordinates.)

The same estimate works if Δ_{ij}^N untwists the spiral partially. The most involved case is when Δ_{ij}^N untwists the spiral completely and then twists in the opposite direction. In addition to the previous cases, we shall need to apply a few more Δ_{ij}s explicitly, and compute the number of twists in the original spiral. The latter is done by computing $[\theta/m]$, which consumes $O(\log N \cdot \log \|L\|)$ elementary operations (because $\theta/m < N$). □

By using induction we deduce the following from Lemma 3.12.

Proposition 3.13. *There exists an algorithm A_2 that computes the curve diagram $\beta \cdot E$ of a braid β given by a braid word w in time $O(n \cdot \ell_\Delta(w)^2)$.*

In order to implement the algorithm of Section 2.4 one needs to choose a presentation method for strip systems. The most straightforward way to present a strip system (L, S) is to provide coordinates of L, list all interval identifications of S, and specify the positions of the punctures. However, in order to make the algorithm more efficient it is useful to include even more information in the object. For example, one may keep a bi-directed list of "significant" points of the axis, which are positions of punctures and the endpoints of the bases of strips, and a collection of cross-references between those points and the related objects (punctures, bases of strips), so as to be able, say, for any base of a strip to find the "next" one in a bounded number of simple operations. We skip the boring details.

One can show that for an appropriate encoding of strip systems, each non-spiralling transmission and the subsequent relaxation described in Sections 2.2–2.4 can be performed in $O(n)$ simple arithmetic operations whose operands are of order $O(\|L\|)$. For performing a d-times spiralling transmission on a strip system (L, S) we additionally need to implement one division (the width of the numerator is divided by the sum of the widths of the bases of the denominators participating in the transmission), which consumes $O(\log d \cdot \log \|L\|)$ operations. Together with Theorem 2.5 this implies the following.

Proposition 3.14. *There exist algorithms A_3, A_4 such that*

1) *given a lamination L, A_3 computes a braid untangling L in time $O(n^2 \cdot (\log \|L\|)^2)$, thus detecting whether L is the curve diagram of some braid;*
2) *given the curve diagrams L_1, L_2 of braids β_1, β_2, A_4 computes the curve diagram of $\beta_1 \circ \beta_2$ in time $O(n^2 \cdot (\ell_\Delta(\beta_1) + \ell_\Delta(\beta_2)) \cdot \ell_\Delta(\beta_1))$.*

We shall finish this paper with some remarks concerning open problems and possible further developments of our results.

First, all our results concern punctured disks and spheres and their mapping class groups. Rafi's work notwithstanding, it would be useful to find generalizations of our methods to all surfaces.

Secondly, we conjecture that the untangling procedures defined in Sections 2.3 and 3.3 describe paths in the Cayley graph of B_n which are uniform quasi-geodesics with respect to the standard metric (not our Δ-metric) on B_n. Indeed, these paths look very much like train track splitting sequences, which are known to be quasigeodesics by a theorem of Hamenstädt ([16, Proposition 3.1]). However, the exact technical conditions of Hamenstädt's theorem, and in particular the genericity condition, are not easy to satisfy. Our conjecture would in particular imply that every braid has a σ_1-consistent representative whose length is bounded linearly by the length of the braid—the existence of such a representative is still an open problem [9, 10]. If the conjecture were true, then our untangling paths would have the interesting property that they are short with respect to both the usual and the Δ-metric on B_n.

Thirdly, it might be useful to give substance to the intuition that every "spiral" that appears during our untangling algorithm is somehow "visible" in every reasonably short representative of the braid, and in particular in the Garside normal form. The idea here is that spirals correspond to passages of the Teichmüller geodesic through the thin part of Teichmüller space. The work of Masur–Minsky and Rafi is probably relevant here.

Finally, there might be applications of our results to the conjugacy problem in B_n, and more precisely to the study of sets of "short" elements in the conjugacy class of a braid (e.g., the super summit set). The reason for this hope is that conjugacy classes of braids correspond to free homotopy classes of closed curves in moduli space.

Acknowledgments. We thank Ian Agol, who, in a conversation with I. Dynnikov, first suggested applying the techniques from [3] to our problem. Jason Behrstock suggested the connection with Teichmüller spaces, and made some very helpful remarks on that subject. We thank Lee Mosher and Jérôme Los for very helpful discussions, and the referee for a very careful and critical reading of the manuscript which gave rise to a major revision of the paper. The work of I. Dynnikov was supported in part by the Russian Foundation for Basic Research (grant no. 02-01-00659). Finally, we thank the CNRS and the Russian Academy of Sciences for their financial support: their joint Franco-Russian exchange program paid for a two-week stay of B. Wiest at the Steklov Institute, Moscow, and for a one-week stay of I. Dynnikov at Rennes University.

References

[1] Artin, E.: Theorie der Zöpfe. Abh. Math. Sem. Hamburg Univ. **4**, 47–72 (1925) JFM 51.0450.01

[2] Artin, E.: Theory of braids. Ann. of Math. **48**, 101–126 (1947) Zbl 0030.17703 MR 0019087

[3] Agol, I., Hass, J., Thurston, W.: The computational complexity of knot genus and spanning area. Trans. Amer. Math. Soc. **358**, 3821–3850 (2006) Zbl 1098.57003 MR 2219001

[4] Birman, J.: Braids, Links, and Mapping Class Groups. Ann. of Math. Stud. 82, Princeton Univ. Press, Princeton (1975) Zbl 0305.57013 MR 0375281

[5] Birman, J., Ko, K. H., Lee, S. J.: A new approach to the word and conjugacy problems in the braid groups. Adv. Math. **139**, 322–353 (1998) Zbl 0937.20016 MR 1654165

[6] Bonnefont, M., Hillion, E.: computer program (maple worksheet), http://perso.univ-rennes1.fr/bertold.wiest

[7] Choi, Y.-E., Rafi, K.: Comparison between Teichmüller and Lipschitz metrics. arXiv: math.GT/0510136

[8] Dehornoy, P.: Braid groups and left distributive operations. Trans. Amer. Math. Soc. **345**, 115–151 (1994) Zbl 0837.20048 MR 1214782

[9] Dehornoy, P.: A fast method for comparing braids. Adv. Math. **125**, 200–235 (1997) Zbl 0882.20021 MR 1434111

[10] Dehornoy, P., Dynnikov, I., Rolfsen, D., Wiest, B.: Why are Braids Orderable?. Panoramas et Synthèses 14, Soc. Math. France (2002) Zbl 1048.20021 MR 1988550

[11] Dynnikov, I. A.: On a Yang–Baxter mapping and the Dehornoy ordering. Russian Math. Surveys **57**, 592–594 (2002) Zbl 1067.20501 MR 1918864

[12] Epstein, D. B. A., Cannon, J. W., Holt, D. F., Levy, S. V. F., Patterson, M. S., Thurston, W.: Word Processing in Groups. Jones and Bartlett, Boston and London (1992) Zbl 0764.20017 MR 1161694

[13] Fathi, A., Laudenbach, F., Poenaru, V.: Travaux de Thurston sur les surfaces. Séminaire Orsay, Astérisque **66–67** (1979), reprinted 1991

[14] Fenn, R., Greene, M. T., Rolfsen, D., Rourke, C., Wiest, B.: Ordering the braid groups. Pacific J. Math. **191**, 49–74 (1999) Zbl 1009.20042 MR 1725462

[15] Garside, F. A.: The braid group and other groups. Quart. J. Math. Oxford (2) **20**, 235–254 (1969) Zbl 0194.03303 MR 0248801

[16] Hamenstädt, U.: Geometry of the mapping class group II: subsurfaces. arXiv:math.GR/0511349v3

[17] Kerckhoff, S. P.: The asymptotic geometry of Teichmüller space. Topology **19**, 23–41 (1980) Zbl 0439.30012 MR 0559474

[18] Malyutin, A. V.: Fast algorithms for the recognition and comparison of braids. Zap. Nauchn. Sem. POMI **279**, 197–217 (2001) Zbl 1070.20044 MR 1846081

[19] Masur, H. A., Minsky, Y. N.: Geometry of the complex of curves I: hyperbolicity. Invent. Math. **138**, 103–149 (1999) Zbl 0941.32012 MR 1714338

[20] Masur, H. A., Minsky, Y. N.: Geometry of the complex of curves II: hierarchical structure. Geom. Funct. Anal. **10**, 902–974 (2000) Zbl 0972.32011 MR 1791145

[21] Minsky, Y. N.: Extremal length estimates and product regions in Teichmüller space. Duke Math. J. **83**, 249–286 (1996) Zbl 0861.32015 MR 1390649

[22] Mosher, L.: Mapping class groups are automatic. Ann. of Math. **142**, 303–384 (1995) Zbl 0867.57004 MR 1343324

[23] Penner, R. C., Harer, J. L.: Combinatorics of Train Tracks. Princeton Univ. Press (1992) Zbl 0765.57001 MR 1144770

[24] Rafi, K.: A combinatorial model for the Teichmüller metric. arXiv:math.GT/0509584

[25] Thurston, W. P.: Minimal stretch maps between hyperbolic surfaces. arXiv:math.GT/9801039

[26] Wiest, B.: An algorithm for the word problem in braid groups. arXiv:math.GT/0211169