Title	Mapping class factorization via fatgraph Nielsen reduction
Author(s)	Bene, Alex James
Citation	Osaka Journal of Mathematics. 2011, 48(4), p. 1047-1061
Version Type	VoR
URL	https://doi.org/10.18910/4939
rights	
Note	

MAPPING CLASS FACTORIZATION VIA FATGRAPH NIELSEN REDUCTION

Alex James Bene

Abstract

The mapping class group of a genus g surface $\Sigma_{g,1}$ with one boundary component is known to have a simple yet infinite presentation with generators given by elementary *Whitehead moves* on marked bordered fatgraphs. In this paper, we introduce an algorithm called *fatgraph Nielsen reduction* which, from the action of a mapping class $\varphi \in MC_{g,1}$ of $\Sigma_{g,1}$ on the fundamental group $\pi_1(\Sigma_{g,1})$ of $\Sigma_{g,1}$, determines a sequence of Whitehead moves representing φ beginning at any choice of marked bordered fatgraph. The algorithm utilizes a reduction of bordered fatgraphs to linear chord diagrams, where the desired sequence is given in terms of elementary *chord slide* moves which continuously decrease some *energy* function. As a consequence, this leads to an algorithm which factors any mapping class given by its action on $\pi(\Sigma_{g,1})$ in terms of any convenient generating set for $MC_{g,1}$.

1. Introduction

The combinatorial description of the mapping class group of a surface Σ in terms of ideal triangulations of Σ has a long history, going back to Whitehead who proved that any two ideal triangulations are related by a sequence of elementary "diagonal exchange" moves. Later, inspired by the geometric insights of Thurston and Mumford, this combinatorial description became more firmly established in the hyperbolic setting by Penner's decorated Teichmüller space [7] and in the conformal setting by Harer's utilization of Strebel's results on quadratic differentials [4, 10]. In these settings, the Poincaré dual viewpoint gained prevalence, where diagonal exchanges on ideal triangulations gave way to elementary moves called Whitehead moves on marked fatgraphs: vertex-oriented graphs embedded in Σ as a spine. In particular, any mapping class of Σ can be represented by a sequence of Whitehead moves, and this sequence is unique up to certain well-known relations.

The mapping class group $MC_{g,1}$ of a surface $\Sigma_{g,1}$ with one boundary component has particularly nice properties. Algebraically, it is classically known that $MC_{g,1}$ is a subgroup of the automorphism of a free group via its action on the fundamental

---

group of $\Sigma_{g,1}$. Combinatorially, the mapping class group $MC_{g,1}$ admits another related combinatorial description in terms of elementary moves called chord slides on a special type of *linear* fatgraph [1, 2] which have coincidentally been studied extensively in other contexts under the name of *linear chord diagrams*. We will choose to use this later terminology in this paper. Moreover, there exists a precise chordslide–Whitehead move correspondence which relates these two types of elementary moves. This "linear" variation of the theme appears to have some advantages. In particular, every marked linear chord diagram canonically determines a set of generators for the free group $\pi := \pi_1(\Sigma_{g,1})$ of Σ_g. We call a generating set arising in this fashion a combinatorial generating set, or simply CG set. Each CG set satisfies certain constraints dictated by the form of the linear chord diagram from which it arises, a special case of which is the well known condition that a standard set of symplectic generators $\{\alpha_i, \beta_i\}$ satisfies the relation $\prod_{i=1}^{g} [\alpha_i, \beta_i] = \partial \Sigma_{g,1}$.

While every mapping class $\varphi \in MC_{g,1}$ has a description in terms of elementary moves on fatgraphs, algorithms currently available for determining such a sequence have the disadvantage that they depend on resolving intersections of arcs or closed curves in the surface $\Sigma_{g,1}$ [6, 8]. For many reasons, it would be desirable to construct such a sequence from purely algebraic information, such as the action of φ on the fundamental group π of $\Sigma_{g,1}$. In this paper, we present just such an algorithm, which we call *fatgraph Nielsen reduction*.

To define this algorithm, we introduce an *energy function* $\| \cdot \|$ on CG sets which is derived from a lexicographical ordering on π extending the usual word length function with respect to some set of generators for π. In fact, each marked linear chord diagram $G \hookrightarrow \Sigma_{g,1}$, which can be considered as a choice of "basepoint in decorated Teichmüller space" (see Section 7 or [1]), determines its own energy function. By the correspondence between marked linear chord diagrams and CG sets, we can equivalently consider this energy $\| \cdot \|$ as a function on marked linear chord diagrams, and the main result of this paper can be stated as follows:

Theorem 1.1. *Let* $\| \cdot \|$ *denote the energy function with respect to a fixed "basepoint"* $G \hookrightarrow \Sigma_{g,1}$. *Given any marked linear chord diagram* G_0, *there exists a sequence of energy decreasing chord slides*

$$G_0 \to G_1 \to \cdots \to G_k, \quad \|G_i\| < \|G_{i-1}\|$$

with $G_k = G$.

The proof of this theorem relies on the combinatorics of linear chord diagrams and elementary cancellation theory.

Since there is a canonical sequence of Whitehead moves which "linearizes" a given bordered fatgraph (given by the greedy algorithm of [1]) and every chord slide of a linear chord diagram can be described in terms of Whitehead moves, we obtain the following

Corollary 1.2. *There exists an explicit algorithm which from the action of a mapping class $\varphi \in MC_{g,1}$ on π determines a sequence of Whitehead moves representing φ, beginning at any "basepoint" marked bordered fatgraph $G \hookrightarrow \Sigma_{g,1}$.*

Moreover, since there is a map from Whitehead moves to mapping classes in $MC_{g,1}$ (see [1]) with finite image $\mathfrak{G}_g$, we have the following

Corollary 1.3. *Given any generating set $\mathfrak{S}$ for $MC_{g,1}$ described in terms of $\mathfrak{G}_g$, there is an algorithm for decomposing any mapping class $\varphi \in MC_{g,1}$ into a product of generators in $\mathfrak{S}$.*

2. Combinatorial generating sets

Let $\Sigma_{g,1}$ be a genus g surface with one boundary component, and let π denote its fundamental group $\pi := \pi_1(\Sigma_{g,1}, p)$ with respect to a basepoint $p \in \partial\Sigma_{g,1}$. It is well known that π is isomorphic to a free group on $2g$ generators, and an explicit isomorphism is equivalent to a choice of an ordered set of generators for π. For a given set of $2g$ generators $\mathcal{X} = \{x_i\}_{i=1}^{2g}$, we define the corresponding set of *letters* $\{x_i, \bar{x}_i\}_{i=1}^{2g}$ to be the set containing each element $x \in \mathcal{X}$ and its inverse $\bar{x}$.

We define a *standard symplectic set of generators* for π to be one of the form $\mathcal{S} = \{\alpha_i, \beta_i\}_{i=1}^{g}$ such that the element of π corresponding to the boundary $\partial\Sigma_{g,1}$ is represented by the word $\partial\Sigma_{g,1} = \prod_{i=1}^{g}[\alpha_i, \beta_i]$, where the bracket $[x, y]$ denotes the commutator $xy\bar{x}\bar{y}$. Again, we let $\bar{x}$ denote the inverse of an element x. Note that each letter of the generating set is used once in this product.

More generally, we make the

DEFINITION 2.1. An ordered set of generators $\mathcal{X}$ for π is a *combinatorial generating set*, or CG set, if the boundary element $\partial\Sigma_{g,1} \in \pi$ can be written as a reduced word using each letter of $\mathcal{X}$ exactly once. Two CG sets are equivalent if their corresponding sets of letters are the same.

We remark that not all generating sets are CG sets, and not all CG sets are (equivalent to) standard symplectic generating sets. Topologically, an equivalence class of CG sets corresponds to (an isotopy class of) a collection of $2g$ arcs based at p which decompose the surface $\Sigma_{g,1}$ into a $(4g + 1)$-gon, and a CG set itself corresponds to a labeling and choice of orientation for each such arc. However, we will find it more useful to consider the picture which is Poincaré dual to this, which will involve linear chord diagrams.

The mapping class group $MC_{g,1}$ of the surface $\Sigma_{g,1}$ is typically defined as $\pi_0(\text{Diff}(\Sigma_{g,1}, \partial\Sigma_{g,1}))$, the group of components of the group of self-diffeomorphisms of $\Sigma_{g,1}$ which fix the boundary pointwise. By a classical result usually attributed to Dehn and Nielsen (see [11]), $MC_{g,1}$ has an equivalent algebraic definition as the group

of automorphisms of π which fix the word representing the boundary $\partial \Sigma_{g,1}$. Equivalently, but more in line with our viewpoint, we can define the mapping class group $MC_{g,1}$ in the following way: mapping classes in $MC_{g,1}$ are exactly the automorphisms of π which take standard symplectic generating sets to standard symplectic generating sets, and there is a (non-canonical) one-to-one correspondence between (equivalence classes of) standard symplectic generating sets and elements of $MC_{g,1}$.

We now introduce the chord slide groupoid, which is a subgroupoid of the Ptolemy groupoid of $\Sigma_{g,1}$ (see Section 7 or [3]) and should be thought of as a groupoid "containing" the mapping class group $MC_{g,1}$. Recall that a groupoid can either be described as a set with a partial composition operation with inverses, or as a category all of whose morphisms are isomorphisms.

DEFINITION 2.2 (cf. Theorem 3.1). Fix a CG set $\mathcal{X}_i$ in each equivalence class, and let $\mathfrak{X} = \{\mathcal{X}_i\}$ denote the set of all representatives. The *chord slide groupoid* is defined as the category whose objects are copies of π, precisely one for every $\mathcal{X}_i \in \mathfrak{X}$, and whose morphisms are the homomorphisms $\psi_{i,j} \colon \pi \to \pi$ which take one CG set to another $\psi_{i,j} \colon \mathcal{X}_i \mapsto \mathcal{X}_j$.

Note that by the Hopfian property of π, the $\psi_{i,j}$ are necessarily isomorphisms. Also note that there is a unique morphism between each pair of objects, so that this groupoid is equivalent to a trivial groupoid (in the category-theoretic sense), and that up to isomorphism, this definition does not depend on the choice of representatives.

3. Linear chord diagrams

A linear chord diagram is a combinatorial object best described as a graph G immersed in the plane: the *core* of G consists of the (connected subgraph consisting of) edges of G embedded in the real line, and the remaining edges, the *chords* of G, are immersed arcs in the upper half plane with their endpoints attached at distinct points of the core. We call the attaching points the *chord ends* of G, and require that they correspond to integer points of the real line. (See [2] for a more thorough description.) For a given chord end v, we let $\bar{v}$ denote the opposite end of the chord attached at v.

There is a natural linear chord diagram associated to every CG set, but before we describe this association, we first develop some notation concerning CG sets. Given a CG set $\mathcal{X} = \{x_i\}_{i=1}^{2g}$ for π, let

$$\mathcal{C}_{\mathcal{X}} = \{c_j\}_{j=1}^{4g}, \quad c_j = x_{i_j}^{\varepsilon_j}$$

denote the (unique) ordered set of letters corresponding to $\mathcal{X}$, ordered such that

$$\prod_{j=1}^{4g} c_j = \prod_{j=1}^{4g} x_{i_j}^{\varepsilon_j} = \overline{\partial \Sigma_{g,1}}$$

$$\beta_g \quad \alpha_g \quad \overline{\beta_g} \quad \overline{\alpha_g} \quad \beta_{g-1} \quad \alpha_{g-1} \quad \overline{\beta_{g-1}} \quad \overline{\alpha_{g-1}} \qquad \beta_1 \quad \alpha_1 \quad \overline{\beta_1} \quad \overline{\alpha_1}$$

Fig. 3.1. A symplectic CG set and linear chord diagram.

with $\varepsilon_i = \pm 1$. (The reason for the appearance of $\overline{\partial \Sigma_{g,1}}$ rather than $\partial \Sigma_{g,1}$ in the above formula will be discussed in Section 7.) For example, for a standard symplectic CG set $\{\alpha_i, \beta_i\}_{i=1}^g$, the corresponding letters are $c_1 = \beta_g$, $c_2 = \alpha_g$, $c_3 = \bar{\beta}_g$, etc. Finally, we define integers l_i and r_i for $i = 1, \dots, 2g$ by

$$c_{l_i} = x_i^{\pm 1}, \quad c_{r_i} = x_i^{\pm 1}, \quad l_i < r_i.$$

Note that the set $\{c_{l_i}\}_{i=1}^{2g}$ of "left ends" gives a canonical representative for the equivalence class containing $\mathcal{X}$, as does the set $\{c_{r_i}\}_{i=1}^{2g}$ of "right ends".

Given a CG set $\mathcal{X}$ for π, we construct a linear chord diagram, denoted $G_\mathcal{X}$, as follows. The core of $G_\mathcal{X}$ is taken to be the interval $[1, 4g]$ of the real line, the chords of $G_\mathcal{X}$ correspond to the generators of $\mathcal{X}$, and the ends of a chord corresponding to $x_i \in \mathcal{X}$ are prescribed by l_i and r_i. Using this correspondence, we will often abuse notation slightly and refer to the generators x_i as *chords* and the letters c_j as *chord ends*. See Fig. 3.1 for the linear chord diagram associated to a standard symplectic set of generators. By the uniqueness of reduced words in π (see [1]), the linear chord diagram $G_\mathcal{X}$ is uniquely determined by the CG set $\mathcal{X}$. We call a linear chord diagram arising from a CG set for π a *genus g chord diagram*. By a *marked linear chord diagram*, we shall mean a linear chord diagram $G_\mathcal{X}$ together with a labeling of its chord ends by the letters $\mathcal{C}_\mathcal{X}$ of the CG set $\mathcal{X}$. By abuse of notation, we shall denote this marked linear chord diagram also by $G_\mathcal{X}$.

We now describe the elementary move on linear chord diagrams, the *chord slide*. Let v be a chord end which is not farthest to the right of a linear chord diagram G. The chord slide of v to the right results in a new linear chord diagram G' obtained by "sliding v over" the chord to its right, thus relocating v to a new position along the core (and isotoping chord ends so that they lie at distinct integer points). The chord slide to the left is completely analogous. See Fig. 3.2.

We can extend chord slides so that they act on the set $\mathfrak{X}$ of (representatives of) CG sets in a natural way. This action is most easily described pictorially by the action of a chord slide on the chord ends associated to a CG set. There are four possible types of chord slides, and we illustrate the action of each in Fig. 3.2, where all chord ends not depicted in the figure are understood to be fixed by the chord slide. While not explicitly needed, we note that the resulting effect on a CG set $\mathcal{X}$ takes the following form up to

Fig. 3.2. Two slides to the right and two slides to the left.

permutation, which is a (composition of) Nielsen transformation(s) for the free group π:

$$x_j \mapsto x_j^{\pm 1}, \quad j \neq i$$
$$x_i \mapsto (x_i^{\pm 1} x_k^{\pm 1})^{\pm 1}, \quad \text{for some} \quad k \neq i,$$

where $x_i \in \mathcal{X}$ is the generator corresponding to the chord which a chord end is slid over, and the exact form depends on the choice of representative CG sets.

Note that under obvious conditions, chord slides on CG sets can be composed, thus they generate a groupoid. In fact, by their actions on (equivalence classes of) CG sets, chord slides can be considered as elements (i.e., morphisms) of the chord slide groupoid, and we now recall the main result of [2]:

Theorem 3.1 ([2]). *The chord slide groupoid is generated by chord slides on CG sets, and all relations are compositions of five explicit ones* (*which we do not list here*).

The importance of this theorem is that it yields a combinatorial presentation of the mapping class group $MC_{g,1}$. More precisely, the action of the mapping class group $MC_{g,1}$ on π extends to an obvious action on the set $\mathfrak{X}$ of (equivalence classes) of CG sets, and letting $\varphi(\mathcal{X})$ denote the image of $\mathcal{X}$ under $\varphi \in MC_{g,1}$, the theorem states that φ can be represented by a sequence of chord slides

$$G_{\mathcal{X}} \to G_{\mathcal{X}_1} \to \cdots \to G_{\varphi(\mathcal{X})}$$

on marked linear chord diagrams, unique up to some known relations. The main result of this paper is an algorithm which will determine such a sequence. Note that as a consequence, we arrive at another proof that the chord slide groupoid is generated by chord slides.

4. The length and energy functions

We now introduce the length and energy functions that we will use to develop our algorithm. For the rest of the paper, we assume that the genus g is fixed, and we fix

a (not necessarily standard symplectic) CG set $\mathcal{S} = \{s_i\}_{i=1}^{2g}$ for π with corresponding letters $\mathcal{C}_{\mathcal{S}} = \{\sigma_j\}_{j=1}^{4g}$ so that $\prod_{j=1}^{4g} \sigma_j = \overline{\partial \Sigma_{g,1}}$.

Firstly, we recall the standard word length function with respect to $\mathcal{S}$, which we denote by $|\cdot|$. For an element $w \in \pi$, we define $|w| = k$ if w is written as a reduced word $w = \prod_{i=1}^{k} \sigma_{j_i}$ using k (possibly repeating) letters of $\mathcal{C}_{\mathcal{S}}$. We also set $|id| = 0$. For any CG set $\mathcal{X} = \{x_i\}$ with corresponding chord ends $\mathcal{C}_{\mathcal{X}} = \{c_j\}$, we define the length of $\mathcal{X}$ by

$$|\mathcal{X}| := 2 \sum_{i=1}^{2g} |x_i| = \sum_{j=1}^{4g} |c_j|.$$

Note that $|\mathcal{X}| = 4g$ if and only if $\mathcal{X}$ is equivalent to $\mathcal{S}$ (i.e., $\mathcal{C}_{\mathcal{X}} = \mathcal{C}_{\mathcal{S}}$) by the uniqueness of reduced words. We also define the length function of a marked linear chord diagram $G_{\mathcal{X}}$ by $|G_{\mathcal{X}}| := |\mathcal{X}|$.

We shall also require a refinement of the above word length function, which we call the *energy function*. This energy function is derived from a lexicographical ordering for π, and we begin by defining the value of the energy function on the letters in $\mathcal{C}_{\mathcal{S}}$. We set

$$\|\sigma_j\| = j.$$

We then define the energy of a reduced word $w = \prod_{i=1}^{k} \sigma_{j_i} \in \pi$ by

$$\|w\| = \sum_{i=1}^{k} (4g + 1)^{k-i} \|\sigma_{j_i}\|.$$

For example, the word $w = \alpha_2 \bar{\beta}_1 = \sigma_2 \sigma_7$ has energy $\|w\| = 25$ with respect to the standard symplectic generating set $\{\alpha_i, \beta_i\}_{i=1}^{2}$ for the genus two surface. Note that this function defines a well-ordering on elements of π. Moreover, this extends the word length function since by the definition, we have

$$(4g + 1)^{|w|-1} \leq \|w\| < (4g + 1)^{|w|}$$

so that

$$|w| < |v|, \quad \text{implies} \quad \|w\| < \|v\|.$$

Finally, for any CG set $\mathcal{X} = \{x_i\}$ with corresponding chord ends $\mathcal{C}_{\mathcal{X}} = \{c_j\}$, we define the energy of $\mathcal{X}$ to be

$$\|\mathcal{X}\| = \sum_{j=1}^{4g} \|c_j\|.$$

Similarly, we define the energy of a marked linear chord diagram $G_{\mathcal{X}}$ by $\|G_{\mathcal{X}}\| := \|\mathcal{X}\|$.

As a final remark, we state the easy

Lemma 4.1. *If $p, q \in \pi$ with $\|p\| < \|q\|$, and $w \in \pi$ is written as a reduced word $w = qx$ with $|q| \leq |x|$, then*

$$\|px\| + \|\bar{x}\bar{p}\| < \|qx\| + \|\bar{x}\bar{q}\|.$$

Proof.   Letting $k = |x|$, we have

$$\|\bar{x}\bar{p}\| - \|\bar{x}\bar{q}\| = \|\bar{p}\| - \|\bar{q}\| < (4g + 1)^k,$$

while

$$\|qx\| - \|px\| = (\|q\| - \|p\|)(4g + 1)^k \geq (4g + 1)^k,$$

and the result follows. $\square$

5.   Some cancellation theory

Recall that we have fixed a "basepoint" CG set $\mathcal{S} = \{s_i\}_{i=1}^{2g}$ with corresponding letters $\mathcal{C}_\mathcal{S} = \{\sigma_j\}_{j=1}^{4g}$ and length and energy functions $|\cdot|$ and $\|\cdot\|$.

Now, consider an arbitrary CG set $\mathcal{X} = \{x_i\}_{i=1}^{2g}$ with corresponding chord ends $\mathcal{C}_\mathcal{X} = \{c_j\}_{j=1}^{4g}$ and linear chord diagram $G_\mathcal{X}$. For a given chord end $c_j \in \mathcal{C}_\mathcal{X}$ with $j > 1$, let $l(c_j)$ denote the amount of cancellation of c_j on the left, meaning the number of letters that get cancelled when reducing the word $c_{j-1}c_j$. More precisely,

$$(1) \qquad\qquad l(c_j) = \frac{1}{2}(|c_{j-1}| + |c_j| - |c_{j-1}c_j|).$$

Similarly, define $r(c_j)$ to be the number of cancellations on the right, so that $r(c_j) = l(c_{j+1})$ with $1 \leq j < 4g$. We also formally set $l(c_1) = r(c_{4g}) = 0$.

We begin by stating the obvious

Lemma 5.1. *If $l(c_j) > |c_j|/2$ for some $j > 1$, then the chord slide of c_j to the left results in a new CG set $\mathcal{X}'$ with $|\mathcal{X}'| < |\mathcal{X}|$. Similarly, if $r(c_j) > |c_j|/2$ for some $j < 4g$, the slide of c_j to the right reduces the length of the CG set.*

Proof.   The slide of c_j to the left replaces the chord ends c_{j-1} and $\bar{c}_{j-1}$ with $c_{j-1}c_j$ and $\bar{c}_j\bar{c}_{j-1}$ respectively and leaves the others unchanged (up to permutation). Thus, if $l(c_j) > |c_j|/2$, then $|c_{j-1}c_j| < |c_{j-1}|$ by (1), so the result follows. The argument for the slide to the right is analogous. $\square$

If we were able to show that every CG set always had an associated chord end c_j with $l(c_j) > |c_j|/2$ or $r(c_j) > |c_j|/2$, then there would always be a chord slide which reduced the length, and we would obtain our main result Theorem 1.1. However, this is not always the case, and we will have to make use of a more subtle argument using the energy function $\|\cdot\|$. Before stating the next Lemma, we develop some more notation.

$$x\bar{p} \;\; p\bar{q}\,qw \;\; \cdots \;\; \bar{w}\bar{q} \qquad\qquad x\bar{p} \;\; pw \;\; \cdots \;\; \bar{w}\bar{p}\,p\bar{q}$$

Fig. 5.1. A balanced chord end slid to the right.

DEFINITION 5.2. We say that a chord end c_j is *balanced* if $l(c_j) = r(c_j) = |c_j|/2$.

Note that if c_j is a balanced chord end corresponding to a CG set $\mathcal{X}$, then sliding c_j in either direction does not change the length $|\mathcal{X}|$ of the CG set. However, we have the following

Lemma 5.3. *Assume that no chord slide reduces the length of a CG set $\mathcal{X}$. Then, if c_j is a balanced chord end of $\mathcal{C}_\mathcal{X}$, either the chord slide of c_j to the left decreases the energy of $\mathcal{X}$ or else the slide of c_j to the right does.*

Proof. Since c_j is balanced, it can be written (as a reduced word)

$$c_j = p\bar{q}$$

with $|p| = |q|$. Note that necessarily, $p \neq q$. Assume, without loss of generality, that $\|p\| < \|q\|$. Then we claim that the slide of c_j to the right reduces the energy of the CG set. (In a completely analogous way, the slide to the left would decrease the energy if $\|q\| < \|p\|$.) To see this, write

$$c_{j+1} = qw.$$

We know that $|q| \leq |w|$, or else the chord slide of c_{j+1} to the left would decrease the length of $\mathcal{X}$.

After the slide of c_j to the right, the letters of the CG set get altered (up to permutation) by $c_{j+1} = qw$ being replaced by pw. See Fig. 5.1. Similarly, the opposite end $\bar{c}_{j+1} = \bar{w}\bar{q}$ is replaced by $\bar{w}\bar{p}$. The result then follows from Lemma 4.1. $\square$

Given a CG set $\mathcal{X}$ with chord ends $\mathcal{C}_\mathcal{X} = \{c_j\}$, consider the unreduced word $\prod_{j=1}^{4g} c_j$. We know that the word reduces to $\overline{\partial \Sigma_{g,1}}$, so there must exist

$$\frac{1}{2}\left(\sum_j |c_j| - 4g \right)$$

many cancellations. Note that a priori there may be different choices of cancellation "schemes" available. So let us fix one choice of cancellation scheme. We write W_j

for the subword of c_j which is not cancelled under this process, L_j for the subword cancelled on the left, and R_j for the subword cancelled on the right. Thus,

$$c_j = L_j W_j R_j, \quad \text{and} \quad \prod_{j=1}^{4g} W_j = \overline{\partial \Sigma_{g,1}}$$

with both L_1 and R_{4g} the empty word. Note that in general, there is no correlation between $l(c_j)$ and $|L_j|$.

Lemma 5.4. *For a given CG set $\mathcal{X}$, assume that no chord slide reduces the length $|\mathcal{X}|$ of $\mathcal{X}$. Then if there is a cancellation scheme $\{c_j = L_j W_j R_j\}$ for the word $\prod_{j=1}^{4g} c_j$ with some W_j the empty word, then some chord end c_i is balanced.*

Proof.   Note that it suffices to prove that there is some i with $l(c_i) + r(c_i) \geq |c_i|$, as then either $l(c_i) = r(c_i) = |c_i|/2$, $l(c_i) > |c_i|/2$, or $r(c_i) > |c_i|/2$, and the latter two cases cannot occur due to our assumption that no chord slide reduces the length of $\mathcal{X}$.

Choose a cancellation scheme $\{c_j = L_j W_j R_j\}$ and assume that W_j is the empty word for some $1 \leq j \leq 4g$. If $j = 1$ or $j = 4g$, then the chord slides of c_1 to the right or c_{4g} to the left respectively decrease the length of $\mathcal{X}$, so assume that $1 < j < 4g$. Let $I = [m, n] \cap \mathbb{Z} = \{m, m + 1, \ldots, n\}$ be the maximal index interval containing j such that for all $i \in I$, W_i is the empty word. In other words, the product

$$(2) \qquad W = R_{m-1} \left(\prod_{i=m}^{n} c_i \right) L_{n+1} = \prod_{i=m}^{n+1} R_{i-1} L_i$$

reduces to the identity element in π.

We now assume that $l(c_i) + r(c_i) < |c_i|$ for all $1 \leq i \leq 4g$ and derive a contradiction. Define a new factorization of each c_i with $m \leq i \leq n$ by

$$c_i = \mathcal{L}_i \mathcal{W}_i \mathcal{R}_i \quad \text{with} \quad |\mathcal{L}_i| = l(c_i), \quad |\mathcal{R}_i| = r(c_i), \quad |\mathcal{W}_i| > 0.$$

Similarly, write $R_{m-1} = \mathcal{W}_{m-1} \mathcal{R}_{m-1}$ and $L_{n+1} = \mathcal{L}_{n+1} \mathcal{W}_{n+1}$ with $|\mathcal{R}_{m-1}| = r(c_{m-1})$ and $|\mathcal{L}_{n+1}| = l(c_{n+1})$. Note that without our assumption $l(c_i) + r(c_i) < |c_i|$, the existence of such a factorization is not guaranteed.

With this new factorization, the product W of (2) can be written as

$$W = \mathcal{W}_{m-1} \mathcal{R}_{m-1} \left(\prod_{i=m}^{n} \mathcal{L}_i \mathcal{W}_i \mathcal{R}_i \right) \mathcal{L}_{n+1} \mathcal{W}_{n+1},$$

which reduces to

$$W' = \prod_{i=m-1}^{n+1} \mathcal{W}_i, \quad |W'| > 0$$

by the definitions of the numbers $l(c_i)$ and $r(c_i)$. However, the definitions of $l(c_i)$ and $r(c_i)$ also ensure that W' cannot be reduced any further, while on the other hand, we know that W' reduces to the empty word. Thus, we have our contradiction and the lemma is proved. $\square$

In the opposite direction, we finish this section with our final lemma. Before we state it, we introduce some helpful terminology. Given a word $w = \prod_{i=1}^{k} x_i$, we call the subwords

$$\prod_{i \leq k/2} x_i \quad \text{and} \quad \prod_{i \geq k/2} x_i$$

the *left half* and *right half* of w respectively.

Lemma 5.5. *Assume that there is no length reducing chord slide for a given CG set $\mathcal{X}$. If for some cancellation scheme $\{c_j = L_j W_j R_j\}$ there is no i with W_i the empty word, then $W_j = \sigma_j$ for all $1 \leq j \leq 4g$. Moreover, in this case, if $c_j = \sigma_j$ for some j, then no chord end other than c_j contains the letter σ_j.*

Proof.   We begin with some observations. First of all, since $\sum_{j=1}^{4g} |W_j| = |\overline{\partial \Sigma_{g,1}}| = 4g$, it is immediate that if no W_j is the empty word, then $|W_j| = 1$ for all j. In fact, it is clear that we must have $W_j = \sigma_j$ for all j and that $R_j = \bar{L}_{j+1}$ for $1 \leq j < 4g$. Second, using Lemma 5.1, the condition that no length reducing chord slide exists forces $\big||L_j| - |R_j|\big| \leq |W_j|$ for all j, which in our case gives $\big||L_j| - |R_j|\big| \leq 1$, meaning each W_j lies "near the middle" of c_j. As a consequence, any letter $\sigma_i \neq \sigma_j$ appears in the left half of c_j if and only if it appears in L_j, and it appears in the right half if and only if it appears in R_j.

Now assume that for some j we have $c_j = \sigma_j$ and that the letter σ_j appears in the chord end c_k with $k \neq j$. Since $\sigma_k \neq \sigma_j$, σ_j appears in either L_k or R_k, and we assume without loss of generality that it appears in L_k, in which case, k must be greater than 1. Then $\bar{\sigma}_j$ must appear in R_{k-1}. Let c_{k_1} be the opposite end of c_{k-1} so that we have σ_j appearing in c_{k_1}. Note that we cannot have $k_1 = j$ since $|c_{k_1}| > 1 = |c_j|$. Since $\bar{\sigma}_j$ lies in the right half of c_{k-1}, σ_j must lie in the left half of c_{k_1}. Since $W_{k_1} \neq \sigma_j$, σ_j must appear in L_{k_1}. Similarly, we see that $\bar{\sigma}_j$ appears in R_{k_1-1}, thus σ_j appears in the left half of the opposite end of c_{k_1-1}. Continuing in this way, we eventually see that either σ_j lies in L_1 or L_j, or that $\bar{\sigma}_j$ lies in R_j or R_{4g} (see the last section where we discuss the boundary cycle of a fatgraph), which gives us a contradiction since each L_1, R_{4g}, L_j, and R_j is the empty word. $\square$

6.   Main result

Again, we fix a basepoint $\mathcal{S}$ and corresponding length functions. We now present our main theorem:

Fig. 6.1. The case $|R_1| = 1$.

Theorem 6.1. *Given a CG set $\mathcal{X}$ not equivalent to $\mathcal{S}$, there is some chord slide which reduces the energy $\|\mathcal{X}\|$ of $\mathcal{X}$.*

Proof. If any chord slide reduces the word length of $\mathcal{X}$, then we are done, so assume otherwise. Also, if under any cancellation scheme $\{c_j = L_j W_j R_j\}$ there is some chord end c_j with W_j the empty word, we are done by Lemmas 5.3 and 5.4, so assume that this is not the case.

With these assumptions, the conditions of Lemma 5.5 are satisfied so that each W_j is the single letter word $W_j = \sigma_j$. In this case, we know that $c_1 = \sigma_1 R_1$ and $c_2 = L_2 \sigma_2 R_2$ with $L_2 = \bar{R}_1$, etc. Furthermore, we must have $|R_1| \leq 1$ otherwise we could slide c_1 to the right and decrease the length of c_2.

Now, if $|R_1| = 1$, then sliding c_1 to the right would have the effect of changing the first letter of c_2 from L_2 to σ_1, changing the last letter of $\bar{c}_2$ from $\bar{L}_2$ to $\bar{\sigma}_1$, and keeping all other chord ends unchanged. See Fig. 6.1. Since $\sigma_1 \neq \bar{R}_1 = L_2$, we would then have $1 = \|\sigma_1\| < \|L_2\|$. Thus,

$$\|\sigma_1 \sigma_2 R_2\| + \|\bar{R}_2 \bar{\sigma}_2 \bar{\sigma}_1\| < \|L_2 \sigma_2 R_2\| + \|\bar{R}_2 \bar{\sigma}_2 \bar{L}_2\|$$

by Lemma 4.1. In other words, this would reduce the energy $\|\mathcal{X}\|$ of $\mathcal{X}$, and we would be done.

So we are left with considering the case where $|R_1| = 0$, or in other words, the case where $c_1 = \sigma_1$. In this case, we must have $c_2 = \sigma_2 R_2$. Lemma 5.5 says that R_2 cannot contain the letter σ_1. Thus, if $|R_2| \neq 0$, the same argument as above shows that the chord slide of c_2 to the right reduces the energy. If $|R_2| = 0$, then we look at c_3, etc., until we find the minimum j such that $|R_j| \neq 0$. Note that a minimum j must exist since $\mathcal{X}$ is not equivalent to $\mathcal{S}$. In this case, we have $c_j = \sigma_j R_j$ with $|R_j| = 1$. Again, Lemma 5.5 says R_j cannot contain any of the letters σ_i for $i < j$ so that the chord slide of c_j to the right reduces the energy. $\square$

We note that if we assume our basepoint CG set $\mathcal{S}$ to be symplectic in the above theorem, we can prove slightly more: if no chord slide reduces the word length of $\mathcal{X}$, then some chord end is balanced. Moreover, it is tempting to conjecture that this holds regardless of our choice of basepoint $\mathcal{S}$.

As an immediate consequence of the above theorem, we obtain the

Theorem 6.2. *There is an algorithm which from a given marked linear chord diagram G_0 provides a sequence of energy decreasing chord slides*

$$G_0 \to G_1 \to \cdots \to G_k, \quad \|G_i\| < \|G_{i-1}\|$$

with $G_k = G_{\mathcal{S}}$. In particular, by applying this to $G_0 = G_{\varphi(\mathcal{S})}$ for some $\varphi \in MC_{g,1}$, we obtain an algorithm which factors mapping classes into sequences of chord slides.

Proof. At every step, the existence of an energy reducing chord slide is provided by Theorem 6.2. That this determines an algorithm is guaranteed by the fact that for any linear chord diagram, only finitely many chord slides are possible. $\square$

7. Relation to fatgraphs

In this final section, we re-articulate the main result of the previous section into statements about factorizations of mapping classes. As the name suggests, the original motivation for the fatgraph Nielsen reduction algorithm comes from a desire to obtain a factorization into sequences of Whitehead moves on fatgraphs, and we begin this section by briefly outlining how our results on linear chord diagrams transfer to this fatgraph context. We end the section with a purely mapping class group factorization algorithm where a mapping class is factored into a product of certain generators for $MC_{g,1}$.

A *fatgraph* is a finite graph with a cyclic ordering given to the half-edges incident to each vertex. Given a fatgraph G, we can define the boundary cycles of G to be the cyclically ordered sequences of oriented edges where an incoming edge at a vertex v is followed by the outgoing edge which is next in the cyclic ordering at v. By gluing a disc onto each boundary component of a fatgraph G, one obtains a closed orientable surface Σ_G and we define the *genus* of the fatgraph to be the genus of the surface Σ_G. By a genus g bordered fatgraph, we shall mean a genus g fatgraph with one boundary component with all vertices trivalent except for one which is univalent. We call the edge incident to the univalent vertex the *tail t* and give it the orientation that it points away from the univalent vertex.

Fix a genus g surface $\Sigma_{g,1}$ with one boundary component $\partial\Sigma_{g,1}$, and let $q \in \partial\Sigma_{g,1}$ be a point on the boundary distinct from the basepoint $p \in \partial\Sigma_{g,1}$. By a marking of a genus g bordered fatgraph, we shall mean an isotopy class of embedding of G into $\Sigma_{g,1}$ as a spine such that the image of the end of the tail is the point q. Poincaré dual to a marking of a bordered fatgraph $G \hookrightarrow \Sigma_{g,1}$ is a collection of arcs, which can be chosen to be based at p. In this way, we get a map from the oriented edges of G to $\pi = \pi_1(\Sigma_{g,1}, p)$, which we call a π-marking. A π-marking satisfies certain conditions referred to as orientation and vertex compatibility, surjectivity, and geometricity, the last of which states that the π-marking of the tail is the inverse of the element representing the boundary of $\Sigma_{g,1}$, $\pi(t) = \overline{\partial\Sigma_{g,1}}$.

We can then define the Ptolemy groupoid as the groupoid with one object for each (isomorphism class of) marked bordered fatgraph and a unique morphism between each pair of objects. This can be considered as a discrete combinatorial subgroupoid of the fundamental path groupoid of decorated Teichmüller space given by the 2-skeleton of the dual fatgraph complex for $\Sigma_{g,1}$, but we will not need this perspective here (for more details, see for example [7, 1]).

By a Whitehead move $W\colon G \to G'$ on a bordered fatgraph G, we shall mean the process of collapsing a non-tail edge e of G and expanding the resulting four-valent vertex in the unique opposite way. It is clear that markings, thus also π-markings, evolve unambiguously under Whitehead moves. Note that under obvious conditions, Whitehead moves can be composed. There are certain sequences of Whitehead moves called the involutivity, commutativity, and pentagon relations which always take a marked fatgraph to itself. As a result of decorated Teichmüller theory or the theory of Strebel differentials that the Ptolemy groupoid of $\Sigma_{g,1}$ can be equivalently defined as the groupoid generated by Whitehead moves and with relations given by the involutivity, commutativity, and pentagon relations.

7.1. Linear chord diagrams as bordered fatgraphs.

Note that the immersion of a linear chord diagram G in the plane endows G with the structure of a fatgraph. Moreover, if G is the linear chord diagram associated to some CG set for π, then G is a genus g fatgraph with one boundary component. In fact, we can consider G as a bordered fatgraph by ignoring the bivalent vertex corresponding to the rightmost chord end and "growing a tail" on the left: for example, by extending the core $[1, 4g] \in \mathbb{R}$ to the larger interval $[0, 4g]$. In this way, one can easily show that a marking of a linear chord diagram G (as described in Section 3) is equivalent to a π-marking of G as a bordered fatgraph.

As observed in [1], a chord slide on a marked linear chord diagram can be decomposed as a sequence of two Whitehead moves on the corresponding bordered fatgraph. In fact, Theorem 6.2 gives an alternate proof of the fact, first proven in [2], which states that the chordslide groupoid is a subgroupoid of the Ptolemy groupoid. Moreover, there is an algorithm which takes any bordered fatgraph to its "nearest linear chord diagram" via the so-called greedy algorithm [1]. Thus, Theorem 6.2 can be restated as follows:

Theorem 7.1. *Given any "basepoint" marked bordered fatgraph $G_0 \hookrightarrow \Sigma_{g,1}$, there is an algorithm which determines a sequence of Whitehead moves*

$$G_0 \to G_1 \to \cdots \to G_k = \varphi(G_0)$$

representing a mapping class $\varphi \in MC_{g,1}$ purely from the action of φ on π.

We note that other algorithms have been previously presented [6, 8] which rely on representing mapping classes by their action on ideal arcs in $\Sigma_{g,1}$ and resolving intersections of such arcs.

Finally, in [1], the so-called chordslide algorithm was discussed and was shown to give an $MC_{g,1}$-equivariant map from morphisms of the Ptolemy groupoid of $\Sigma_{g,1}$ (i.e., sequences of Whitehead moves) to $MC_{g,1}$. For any genus g, there are only finitely many combinatorial types of Whitehead moves (note that this shows that $MC_{g,1}$ itself is finitely generated) and we denote by $\mathfrak{G}_g$ the finite image of all Whitehead moves (i.e., generators of the Ptolemy groupoid) under this map.

Theorem 7.2.   *If the elements of a generating set $\mathfrak{S}$ for $MC_{g,1}$ can be explicitly written as products of elements of $\mathfrak{G}_g$, then* Theorem 7.1 *provides an algorithm to decompose any mapping class $\varphi \in MC_{g,1}$ as a product of elements of $\mathfrak{S}$, purely from the action of φ on π.*

References

[1]   J.E. Andersen, A.J. Bene and R.C. Penner: *Groupoid extensions of mapping class representations for bordered surfaces*, Topology Appl. **156** (2009), 2713–2725.

[2]   A.J. Bene: *A chord diagrammatic presentation of the mapping class group of a once bordered surface*, Geom. Dedicata **144** (2010), 171–190.

[3]   A.J. Bene, N. Kawazumi and R.C. Penner: *Canonical extensions of the Johnson homomorphisms to the Torelli groupoid*, Adv. Math. **221** (2009), 627–659.

[4]   J.L. Harer: *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. **84** (1986), 157–176.

[5]   S. Morita and R.C. Penner: *Torelli groups*, *extended Johnson homomorphisms*, *and new cycles on the moduli space of curves*, Math. Proc. Cambridge Philos. Soc. **144** (2008), 651–671.

[6]   L. Mosher: *Mapping class groups are automatic*, Ann. of Math. (2) **142** (1995), 303–384.

[7]   R.C. Penner: *The decorated Teichmüller space of punctured surfaces*, Comm. Math. Phys. **113** (1987), 299–339.

[8]   R.C. Penner: *Universal constructions in Teichmüller theory*, Adv. Math. **98** (1993), 143–215.

[9]   R.C. Penner: *Decorated Teichmüller theory of bordered surfaces*, Comm. Anal. Geom. **12** (2004), 793–820.

[10]   K. Strebel: Quadratic Differentials, Springer, Berlin, 1984.

[11]   H. Zieschang, E. Vogt and H.-D. Coldewey: Surfaces and Planar Discontinuous Groups, Lecture Notes in Mathematics **835**, Springer, Berlin, 1980.

Institute for the Physics and Mathematics of the Universe
University of Tokyo
5-l-5 Kashiwanoha
Kashiwa, 277-8568
Japan
e-mail: alex.bene@ipmu.jp