GEOMETRY OF THE MAPPING CLASS GROUPS IIA: CAT(0) CUBE COMPLEXES AND BOUNDARIES

URSULA HAMENSTÄDT

Abstract. We construct a uniformly finite CAT(0) cube complex which admits a coarsely vertex bijective Lipschitz map onto the mapping class group $\mathcal{MCG}$ of a surface S of genus g with m punctures $(3g-3+m \geq 2)$. We identify the regular Roller boundary of this complex with the space of complete geodesic laminations on S. Furthermore, we construct an explicit compactification of the mapping class group which is small at infinity. We define an electrification of the curve graph of S and use it to identify the Poisson boundary of a random walk on $\mathcal{MCG}$ with some mild moment condition as a stationary measure on the space of minimal complete geodesic laminations on S.

Contents

Date: May 10, 2021.
AMS subject classification: 20F65.

1. Introduction

The *mapping class group* $\mathcal{MCG}$ of a surface S of finite type, that is, S is a closed surface of genus $g \geq 0$ from which $m \geq 0$ points, so-called *punctures*, have been deleted, is the group of isotopy classes of diffeomorphisms of S. We assume that S is *non-exceptional*, that is, S is not a sphere with at most 4 punctures or a torus with at most 1 puncture.

Searching for geometric similarities between the mapping class group and non-positively curved groups has a long history, perhaps culminating in equivariant embedding results of the mapping class group into a finite product of quasi-trees [BBF19].

However, the mapping class group is *not* non-positively curved in spite of the fact that it admits an isometric action on a simply connected Kähler manifold of negative sectional curvature, namely, the Teichmüller space equipped with the Weil-Petersen metric. One manifestation of the failure of being nonpositively curved is the fact that for any action of $\mathcal{MCG}$ on a complete CAT(0)-space by semisimple isometries, Dehn twists act as elliptic elements [Br10].

The goal of this article is to add to the positive results concerning the relation of the mapping class group and non-positively curved spaces. To this end define a *cube complex* to be a metric space obtained by gluing standard cubes $[0,1]^n$ ($n \geq 0$) with isometries along proper faces. The cube complex is *uniformly finite* if the number of one-cubes incident on each vertex is uniformly bounded. A uniformly finite cube complex has finite dimension and is proper (closed balls of finite radius are compact).

The *Roller boundary* ∂X of a uniformly finite CAT(0)-cube complex X defines a compactification of X. It depends on the cubical structure, and it is intimately related to the fact that the one-skeleton of a CAT(0) cube complex is a *median space*. We refer to [FLM18] and Section 7 for an account on this construction.

The Roller boundary of a CAT(0) cube complex contains a closed (possibly empty) subset, called the *regular Roller boundary*. This set is invariant under the natural action of the automorphism group of the complex (Section 5.3 of [FLM18]).

We use these notions in our first main result. For its formulation, recall from [H09] that a *complete geodesic lamination* on S is a geodesic lamination which can be approximated in the Hausdorff topology by simple closed geodesics and which decomposes S into ideal triangles and once punctured monogons. The space $\mathcal{CL}$ of all complete geodesic laminations on S, equipped with the Hausdorff topology, is a compact totally disconnected $\mathcal{MCG}$-space. Furthermore, the action of $\mathcal{MCG}$ on $\mathcal{CL}$ is amenable [H09].

Theorem 1. *There exists a uniformly finite* CAT(0) *cube complex C with the following properties.*

(1) *There exists a proper coarsely surjective Lipschitz map $F : C \to \mathcal{MCG}$.*

(2) *There is a natural homeomorphism of the regular Roller boundary of C onto the space of complete geodesic laminations, equipped with the Hausdorff topology.*

The map F in the statement of the theorem is however not bilipschitz.

Among others, the interest in the second part of the theorem lies in the fact that three distinct point of the Roller boundary of a uniformly finite CAT(0) cube complex X have a unique *median* in X and hence there is a natural map from the space of pairwise distinct triples in ∂X into X. This map is proper, but it is not injective. By Theorem 1, the regular Roller boundary of the uniformly finite CAT(0) cube complex C is a compact $\mathcal{MCG}$-space. On the other hand, the action of a group with Kazhdan's property (T) by isometries on a median space has bounded orbits [CDH10]. This supportes the idea that the mapping class group does not have property (T).

A uniformly finite CAT(0) cube complex X can be compactified by adding its *geometric boundary* $\partial_\angle X$, which is the space of equivalence classes of geodesic rays starting at a fixed point. This boundary does not depend on the choice of the basepoint.

As the cube complex C is not bilipschitz equivalent to $\mathcal{MCG}$ and $\mathcal{MCG}$ does not act on C as a group of simplicial isometries, the geometric boundary of C is not a boundary for $\mathcal{MCG}$ in the sense of the following

Definition. A *boundary* of a finitely generated group Γ is a compact Γ-space Y with the following properties. There exists a topology on $\Gamma \cup Y$ which restricts to the discrete topology on Γ, to the given topology on Y and is such that $Y \cup \Gamma$ is compact. The left action of Γ on itself extends to the Γ-action on Y. The boundary is called *small* if the right action of Γ extends to the trivial action of Γ on Y.

The space $\mathcal{CL}$ of complete geodesic laminations is a compact $\mathcal{MCG}$-space, but it is *not* a boundary in the sense of the above definition as it does not define a compactification of $\mathcal{MCG}$.

The second goal of this article is to construct an explicit small boundary $\mathcal{X}$ for $\mathcal{MCG}$. As a set, this boundary $\mathcal{X}$ is given as follows.

The *curve complex* $\mathcal{CG}(S_0)$ of a subsurface S_0 of S different from a pair of pants or an annulus is the simplicial complex whose vertices are isotopy classes of simple closed curves and where k such curves span a $k-1$-simplex if they can be realized disjointly. If S_0 is a four-holed sphere or a one holed torus, then this definition has to be slightly modified. The curve complex is a hyperbolic geometric graph of infinite diameter [MM99]. Its *Gromov boundary* $\partial\mathcal{CG}(S_0)$ is the space of *minimal geodesic laminations* on S_0 which *fill* S_0, that is, which intersect every essential simple closed curve on S_0 transversely. The topology on $\partial\mathcal{CG}(S_0)$ is the *coarse Hausdorff topology.* With respect to this topology, a sequence λ_i of minimal filling laminations converges to the lamination λ if and only if the limit of any subsequence which converges in the Hausdorff topology on compact subsets of S_0 contains λ as a sublamination [H06, K99]. The space $\partial\mathcal{CG}(S_0)$ is separable and metrizable. Define the boundary of the curve complex of an essential annulus $A \subset S$ with core curve c to consist two points c^+, c^-.

If $S_1, \dots, S_k$ is a collection of isotopy classes of pairwise disjoint subsurfaces of S then we can form the join

$$\mathcal{X}(\cup_{i=1}^k S_i) = \partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k).$$

It can be viewed as the set of formal sums $\sum_i a_i\lambda_i$ where $a_i > 0$, $\sum_i a_i = 1$ and where $\lambda_i \in \partial\mathcal{CG}(S_i)$ for all i. This join is a separable metrizable topological space. Note that if $S_{i_1}, \dots, S_{i_s}$ is a subset of the set of surfaces $S_1, \dots, S_k$, then $\mathcal{X}(\cup_{j=1}^s S_{i_j})$ is naturally a closed subset of $\mathcal{X}(\cup_{i=1}^k S_i)$. Define

$$\mathcal{X} = \cup\mathcal{X}(\cup_{i=1}^k S_i)$$

where the union is over all collections of pairwise disjoint essential subsurfaces of S. Here we view an essential annulus A as an essential subsurface which is disjoint from any subsurface which can be moved off A by an isotopy. The union is not a disjoint union. Thus $\mathcal{X}$ is just the set of formal sums $\sum_i a_i\lambda_i$ where $a_i > 0, \sum_i a_i = 1$ and where $\lambda_1, \dots, \lambda_k$ are pairwise disjoint minimal geodesic laminations on S and where each simple closed curve component λ_i is equipped with an additional label $+, -$.

Define the *oriented curve complex* of an essential subsurface S_0 of S to be the complex whose vertices are isotopy classes of *oriented* simple closed curves and where $k+1$ such vertices span a k-simplex if they are pairwise distinct as unoriented simple closed curves and if they can be realized disjointly. With this terminoloy, the union of the oriented curve complex of S_0 with its Gromov boundary is an embedded subspace of $\mathcal{X}$ (a priory only as as set). Note that the mapping class group acts on $\mathcal{X}$ as a set.

Recall that an *embedding* of a topological space X into a topological space Y is an injective map $f : X \to Y$ which is a homeomorphism onto its image, equipped with the subspace topology. We show

Theorem 2. *The space $\mathcal{X}$ admits a $\mathcal{MCG}$-invariant topology $\mathcal{O}$ with the following properties.*

(1) *For any collection $S_1, \dots, S_k$ or pairwise disjoint subsurfaces of S, the inclusion $\mathcal{X}(\cup_{i=1}^k S_i) \to (\mathcal{X}, \mathcal{O})$ is an embedding.*
(2) *$(\mathcal{X}, \mathcal{O})$ is compact, and the $\mathcal{MCG}$-action is minimal and strongly proximal.*
(3) *$(\mathcal{X}, \mathcal{O})$ is a small boundary for $\mathcal{MCG}$.*

We call the space $(\mathcal{X}, \mathcal{O})$ the *geometric boundary* of $\mathcal{MCG}$. If follows immediately from the construction that the geometric boundary of any essential connected subsurface $S_0 \subset S$ is a closed embedded subspace of the geometric boundary of S. The description of the topology on $\mathcal{X}$ which gives $\mathcal{X}$ the structure of a small boundary of $\mathcal{MCG}$ is slightly involved and can be understood best by describing convergent sequences in $\mathcal{MCG}$. We give a detailed account in Section 8.

Another less explicit construction of a boundary for $\mathcal{MCG}$ is due to Durham, Hagen and Sisto [DHS17] taking advantage of a hierarchially hyperbolic structure of

$\mathcal{MCG}$. We do not know the relation between these constructions- in fact, hierarchial hyperbolicity for $\mathcal{MCG}$ does not play any role in this article. The advantage of our construction is that the space $\mathcal{X}$ and its topology as well as the action of the group $\mathcal{MCG}$ on $\mathcal{X}$ is completely explicit and can be used to study subgroups of $\mathcal{MCG}$, as for example in Koberda's work [Kb12] who constructed subgroups of $\mathcal{MCG}$ which are isomorphic to right angled Artin groups.

We conjecture that the geometric boundary $\mathcal{X}$ of $\mathcal{MCG}$ is in fact a *$\mathcal{Z}$-set* for any torsion free finite index subgroup of $\mathcal{MCG}$. We refer to [B96] for a nice account on $\mathcal{Z}$-sets for torsion free groups. Note that the work of Gabai [G14] relates the covering dimension of the boundary of the curve complex of S to the virtual cohomological dimension of $\mathcal{MCG}$, and Bestvina and Bromberg relate the asymptotic dimension of the curve complex to the virtual cohomological dimension of $\mathcal{MCG}$.

In the course of the proof of Theorem 1, we use an auxiliary tool which is defined as follows.

Definition. The *principal curve graph* of a closed surface S of genus $g \geq 2$ is the graph whose vertices are simple closed curves and where two such curves c, d are connected by an edge of length one if and only if $S - (c \cup d)$ has a complementary component which neither is a fourgon nor a sixgon nor a once punctured bigon.

We show in Section 6.

Theorem 3. *The principal curve graph is hyperbolic, and its Gromov boundary equals the space of minimal complete geodesic laminations equipped with the Hausdorff topology.*

As an application, we obtain the following strengthening of the main result of [GM18].

Theorem 4. *The Poisson boundary of a random walk on $\mathcal{MCG}$ generated by a probability measure μ whose support generates $\mathcal{MCG}$, with finite entropy and finite logarithmic moment for the action on the curve graph, can be realized by a stationary measure on the space of minimal complete geodesic laminations on S.*

The condition on the group generated by the support of the measure μ can be relaxed as in [GM18] to be non-elementary and to contain at least one pseudo-Anosov element with minimal complete attracting geodesic lamination. We do not address the question pursued in [GM18] here whether the attracting lamination of a typical element for the random walk is complete, however this can readily be derived from Theorem 4 and the results in [BGH20] which establishes precise information on the axis of a typical mapping class in Teichmüller space from information on the exit measure of the random walk.

We illustrate the usefulness of the tools developed for the proof of Theorem 1 by establishing three geometric results about subgroups of $\mathcal{MCG}(S)$. However, variants of these results are available in the literature.

Namely, a finite symmetric set $\mathcal{G}'$ of generators of a subgroup Γ' of a finitely generated group Γ can be extended to a finite symmetric set of generators of Γ.

Thus for any two word norms $|\,|$ on Γ and $|\,|'$ of Γ' there is a number $L > 1$ such that $|g| \leq L|g|'$ for every $g \in \Gamma'$. However, in general the word norm in Γ of an element $g \in \Gamma'$ can not be estimated from below by $L'|g|'$ for a universal constant $L' > 0$. Define the finitely generated subgroup Γ' of Γ to be *undistorted* in Γ if there is a constant $c > 1$ such that $|g|' \leq c|g|$ for all $g \in \Gamma'$. Thus $\Gamma' < \Gamma$ is undistorted if and only if the inclusion $\Gamma' \to \Gamma$ is a quasi-isometric embedding.

There are subgroups of $\mathcal{MCG}(S)$ with particularly nice geometric descriptions. To begin with, let $S_0 \subset S$ be an essential subsurface of S different from a pair of pants. A finite index subgroup $\mathcal{MCG}_0(S_0)$ of the mapping class group $\mathcal{MCG}(S_0)$ of S_0 can be identified with the subgroup of $\mathcal{MCG}(S)$ of all elements which can be represented by a homeomorphism of S fixing $S - S_0$ pointwise. We give an alternative proof of the following extension of a result of Masur and Minsky (the statement about distortion is implicitly but not explicitly contained in Theorem 6.12 of [MM00] and has been rediscovered in several other disguises ever since). For notational clarity, we write $\mathcal{MCG}(S)$ for the mapping class group of a surface S if this surface needs to be specified.

Theorem 5. *If $S_0 \subset S$ is an essential subsurface of a non-exceptional surface S of finite type, then $\mathcal{MCG}_0(S_0) < \mathcal{MCG}(S)$ is undistorted, and the union of $\mathcal{MCG}(S_0)$ with its geometric boundary embeds into the union of $\mathcal{MCG}(S)$ with its geometric boundary.*

Now let S be a closed surface of genus $g \geq 2$ and let $\Gamma < \mathcal{MCG}(S)$ be any *finite* subgroup. By the solution of the Nielsen realization problem [Ke83], Γ can be realized as a subgroup of the automorphism group of a marked complex structure h on S. Then the quotient $(S, h)/\Gamma$ is a compact Riemann surface, and the projection $S \to S/\Gamma$ is a branched covering ramified over a finite set Σ of points. Let $S_0 = S/\Gamma - \Sigma$ and let $N(\Gamma)$ be the normalizer of Γ in $\mathcal{MCG}(S)$. Then there is an exact sequence

$$0 \to \Gamma \to N(\Gamma) \to \mathcal{MCG}_0(S_0) \to 0$$

where $\mathcal{MCCG}_0(S_0)$ is the subgroup of the mapping class group of S_0 of all elements which can be represented by a homeomorphism which lifts to a homeomorphism of S [BH73]. We use this to observe (compare [RS07] for a similar statement)

Theorem 6. *Let S be a closed surface of genus $g \geq 2$ and let $\Gamma < \mathcal{MCG}(S)$ be a finite subgroup. Then the normalizer of Γ is undistorted in $\mathcal{MCG}(S)$.*

There are other relations between mapping class groups which can be described by exact sequences of groups. An example of such a relation is as follows. Let S_0 be any non-exceptional surface of finite type and let S be the surface obtained from S_0 by deleting a single point p. Then there is an exact sequence [B74]

$$0 \to \pi_1(S_0) \to \mathcal{MCG}(S) \stackrel{\Pi}{\to} \mathcal{MCG}(S_0) \to 0.$$

The image in $\mathcal{MCG}(S)$ of an element $\alpha \in \pi_1(S_0)$ is the mapping class obtained by dragging the puncture p of S along a simple closed curve in the homotopy class α. The projection $\Pi : \mathcal{MCG}(S) \to \mathcal{MCG}(S_0)$ is induced by the map $S \to S_0$ defined by closing the puncture p.

Define a *coarse section* for the projection Π to be a map $\Psi : \mathcal{MCG}(S_0) \to \mathcal{MCG}(S)$ with the property that there exists a number $\kappa > 0$ such that

$$d(\Pi\Psi(g), g) \leq \kappa$$

for all $g \in \mathcal{MCG}(S_0)$. The following result is a special case of a much more general result of Mosher [M96] (the quasi-isometric section lemma).

Theorem 7. *The projection $\mathcal{MCG}(S) \to \mathcal{MCG}(S_0)$ admits a coarse section which is a quasi-isometric embedding.*

Theorem 7 contrasts a result of Braddeus, Farb and Putman [BFP11] who showed that the normal subgroup $\pi_1(S_0)$ of $\mathcal{MCG}(S)$ is exponentially distorted.

The organization of this article is as follows.

The proof of Theorem 1 builds on the results of [H09]. In that paper we constructed a locally finite connected directed graph $\mathcal{TT}$ whose vertex set $\mathcal{V}(\mathcal{TT})$ is the set of all isotopy classes of complete train tracks on S. The mapping class group $\mathcal{MCG}(S)$ acts property and cocompactly on $\mathcal{TT}$ as a group of simplicial isometries.

In Section 5, we construct the CAT(0) cube complex C whose existence is stated in Theorem 1 from a subgraph of this complex. This construction depends on some rather technical results, established in Section 3, which shows that directed edge-paths in $\mathcal{TT}$ connect a coarsely dense set of pairs of points in $\mathcal{TT}$.

An inspection of hyperplanes in C leads to the introduction of the principal curve graph in Section 6, which is independent of the rest of the article. It contains the proof of Theorem 3 and Theorem 4. The principal curve graph is a key tool in Section 7 to control hyperplanes in the CAT(0) cube complex C and complete the proof of Theorem 1.

The construction of the geometric boundary of $\mathcal{MCG}$ is motivated by properties of the geometric boundary of the CAT(0) cube complex C, but it is independent of the rest of the article. It is contained in Section 8.

The first part of Theorems 5 as well as Theorem 6 and Theorem 7 are derived in Section 4. In Section 2 we summarize the properties of the train track complex $\mathcal{TT}$ of S which are needed for our purpose.

Acknowledgement: The work reported in this article was started in fall 2007 while the author was in residence at the Mathematical Sciences Research Institute in Berkeley, California. The Sections 2-4 of this article are essentially identical with Sections 2-4 of the preprint arXiv:0912.0137 whose second part is contained in part B of this work. The results in Sections 5-7 were accomplished while the author was in residence at the Mathematical Research Institute in Berkeley, California, during the fall 2016 semester. Both visits of the MSRI were supported by the National Science Foundation.

2. The complex of train tracks

In this section we summarize some results from [H09] which will be used throughout the paper.

Let S be an oriented surface of genus $g \geq 0$ with $m \geq 0$ punctures and where $3g - 3 + m \geq 2$. A *train track* on S is an embedded 1-complex $\tau \subset S$ whose edges (called *branches*) are smooth arcs with well-defined tangent vectors at the endpoints. At any vertex (called a *switch*) the incident edges are mutually tangent. Through each switch there is a path of class C^1 which is embedded in τ and contains the switch in its interior. In particular, the half-branches which are incident on a fixed switch are divided into two classes according to the orientation of an inward pointing tangent at the switch. Each closed curve component of τ has a unique bivalent switch, and all other switches are at least trivalent. The complementary regions of the train track have negative Euler characteristic, which means that they are different from discs with $0, 1$ or 2 cusps at the boundary and different from annuli and once-punctured discs with no cusps at the boundary. A train track is called *maximal* if each of its complementary components either is a trigon, i.e. a topological disc with three cusps at the boundary, or a once punctured monogon, i.e. a once punctured disc with one cusp at the boundary. We always identify train tracks which are isotopic. The book [PH92] contains a comprehensive treatment of train tracks which we refer to throughout the paper.

A train track is called *generic* if all switches are at most trivalent. The train track τ is called *transversely recurrent* if every branch b of τ is intersected by an embedded simple closed curve $c = c(b) \subset S$ of class C^1 which intersects τ transversely and is such that $S - \tau - c$ does not contain an embedded *bigon*, i.e. a disc with two corners at the boundary.

A *trainpath* on a train track τ is a C^1-immersion $\rho : [m, n] \to \tau \subset S$ which maps each interval $[k, k + 1]$ $(m \leq k \leq n - 1)$ onto a branch of τ. The integer $n - m$ is called the *length* of ρ. We sometimes identify a trainpath with its image in τ. Each complementary region of τ is bounded by a finite number of (not necessarily embedded) trainpaths which either are closed curves or terminate at the cusps of the region. A *subtrack* of a train track τ is a subset σ of τ which itself is a train track. Thus every switch of σ is also a switch of τ, and every branch of σ is a trainpath on τ. We write $\sigma < \tau$ if σ is a subtrack of τ.

A *transverse measure* on a train track τ is a nonnegative weight function μ on the branches of τ satisfying the *switch condition*: for every switch s of τ, the half-branches incident on s are divided into two classes, and the sums of the weights over all half-branches in each of the two classes coincide. The train track is called *recurrent* if it admits a transverse measure which is positive on every branch. We call such a transverse measure μ *positive*, and we write $\mu > 0$. If μ is any transverse measure on a train track τ then the subset of τ consisting of all branches with positive μ-weight is a recurrent subtrack of τ. A train track τ is called *birecurrent* if τ is recurrent and transversely recurrent. We call τ *complete* if τ is generic, maximal and birecurrent.

Remark: As in [H09], we require every train track to be generic. Unfortunately this leads to a slight inconsistency of our terminology with the terminology found in the literature.

There is a special collection of complete train tracks on S which were introduced by Penner and Harer [PH92]. Namely, a *pants decomposition* P for S is a collection of $3g-3+m$ simple closed curves which decompose S into $2g-2+m$ pairs of pants. Here a pair of pants is a planar orientable bordered surface of Euler characteristic -1 which may be non-compact. Define a *marking* of S (or complete clean marking in the terminology of [MM00]) to consist of a pants decomposition P for S and a system of *spanning curves* for P. For each pants curve $\gamma \in P$ there is a unique simple closed spanning curve which is contained in the connected component S_0 of $S - (P - \gamma)$ containing γ, which is not freely homotopic into the boundary or a puncture of this component and which intersects γ in the minimal number of points (one point if S_0 is a one-holed torus and two points if S_0 if a four-holed sphere). Note that any two choices of such a spanning curve differ by a *Dehn twist* about γ.

For each marking F of S we can construct a collection of finitely many maximal transversely recurrent train tracks as follows. Let P be the pants decomposition of the marking. Choose an open neighborhood A of P in S whose closure in S is homeomorphic to the disjoint union of $3g-3+m$ closed annuli. Then $S - A$ is the disjoint union of $2g-2+m$ pairs of pants. We require that each train track τ from our collection intersects a component of $S - A$ which does not contain a puncture of S in a train track with stops which is isotopic to one of the four *standard models* shown in Figure A (see Figure 2.6.2 of [PH92]).

Figure A

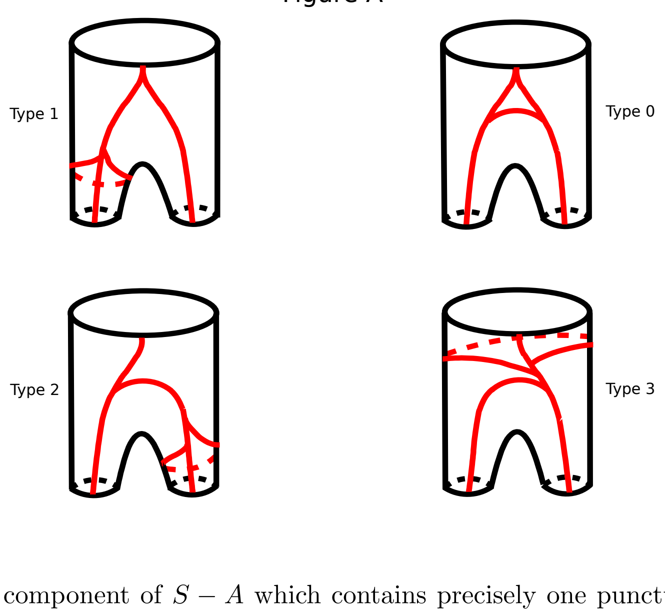

If S_0 is a component of $S - A$ which contains precisely one puncture, then we require that τ intersects S_0 in a train track with stops which we obtain up to diffeomorphism of S_0 from the standard model of type 2 or of type 3 by replacing the top boundary curve by a puncture and by deleting the branch which is incident on the stop of this boundary component. If S_0 is a component of $S - A$ which contains two punctures, then we require that τ intersects S_0 in a train track with stops which we obtain up to diffeomorphism of S_0 from the standard model of type 1 by replacing the two lower boundary components by a puncture and by deleting the two branches which are incident on the stops of these boundary components.

The intersection of τ with a component of the collection A of $3g - 3 + m$ annuli is one of the following four *standard connectors* which are shown in Figure B (see Figure 2.6.1 of [PH92]).

Figure B

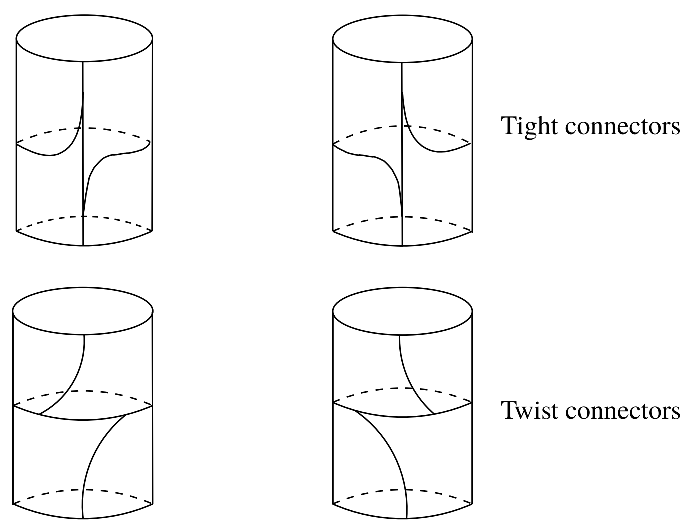

From the above standard pieces we can build a train track τ on S by choosing for each component of $S - A$ one of the standard models as described above and choosing for each component of A one of the four standard connectors. These train tracks with stops are then glued at their stops to a connected train track on S. Any two train tracks constructed in this way from the same pants decomposition P, the same choices of standard models for the components of $S - A$ and the same choices of connectors for the components of A differ by Dehn twists about the pants curves of P. The spanning curves of the marking F determine a specific choice of such a gluing [PH92]. We call each of the resulting train tracks *in standard form for* F provided that it is complete (see p.147 of [PH92] for examples of train tracks built in this way which are not recurrent and hence not complete).

A *geodesic lamination* for a complete hyperbolic structure on S of finite volume is a *compact* subset of S which is foliated into simple geodesics. A geodesic lamination λ is *minimal* if each of its half-leaves is dense in λ. A geodesic lamination is *maximal* if its complementary regions are all ideal triangles or once punctured monogons (note that a minimal geodesic lamination can also be maximal). The space of geodesic laminations on S equipped with the *Hausdorff topology* is a compact metrizable space.

A geodesic lamination λ is called *complete* if λ is maximal and can be approximated in the Hausdorff topology by simple closed geodesics. The space $\mathcal{CL}$ of all complete geodesic laminations equipped with the Hausdorff topology is compact. The mapping class group $\mathcal{MCG}(S)$ naturally acts on $\mathcal{CL}$ as a group of homeomorphisms. Every geodesic lamination λ which is a disjoint union of finitely many minimal components is a *sublamination* of a complete geodesic lamination, i.e. there is a complete geodesic lamination which contains λ as a closed subset (Lemma 2.2 of [H09]).

A train track or a geodesic lamination σ is *carried* by a transversely recurrent train track τ if there is a map $\varphi : S \to S$ of class C^1 which is homotopic to the identity and maps σ into τ in such a way that the restriction of the differential of φ to the tangent space of σ vanishes nowhere; note that this makes sense since a train track has a tangent line everywhere. We call the restriction of φ to σ a *carrying map* for σ. Write $\sigma \prec \tau$ if the train track σ is carried by the train track τ. Then every geodesic lamination λ which is carried by σ is also carried by τ. A train track τ is complete if and only if it is generic and transversely recurrent and if it carries a complete geodesic lamination. The space of complete geodesic laminations carried by a complete train track τ is open and closed in $\mathcal{CL}$ (Lemma 2.3 of [H09]). In particular, the space $\mathcal{CL}$ is totally disconnected.

For every pants decomposition P of S there is a finite set of complete geodesic laminations on S which contain (the geodesic representatives of) the components of P as their minimal components. We call such a geodesic lamination *in standard form* for P. If λ is a geodesic lamination in standard form for P then for each component S_0 of $S - P$ which does not contain a puncture of S, there are precisely three leaves of λ contained in S_0 which spiral about the three different boundary components of S_0. The leaves of λ spiraling from two different sides about a component γ of P define opposite orientations near γ (as shown in Figure A of [H09]). If S_0 contains exactly one puncture of S there there are two leaves of λ contained in S_0 which spiral about the two boundary components of S_0.

For every marking F of S with pants decomposition P and every train track τ in standard form for F with only twist connectors there is a unique complete geodesic lamination in standard form for P which is carried by τ. This implies that for every marking F of S with pants decomposition P, there is a bijection between the complete train tracks in standard form for F with only twist connectors and the complete geodesic laminations in standard form for P. The set of all complete geodesic laminations in standard form for some pants decomposition P is invariant under the action of the mapping class group, moreover there are only finitely many $\mathcal{MCG}(S)$-orbits of such complete geodesic laminations.

Define the *straightening* of a train track τ on S with respect to some complete finite volume hyperbolic structure g on S to be the edgewise immersed graph in S whose vertices are the switches of τ and whose edges are the unique geodesic arcs which are homotopic with fixed endpoints to the branches of τ. For a number $\epsilon > 0$ we say that the train track τ *ϵ-follows* a geodesic lamination λ if the tangent lines of the straightening of τ are contained in the ϵ-neighborhood of the projectivized tangent bundle $PT\lambda$ of λ (with respect to the distance function induced by the metric g) and if moreover the straightening of every trainpath on τ is a piecewise geodesic whose exterior angles at the breakpoints are not bigger than ϵ.

Lemma 2.1. *Let $\lambda \in \mathcal{CL}$ be a complete geodesic lamination on S in standard form for a pants decomposition P of S and let $\epsilon > 0$. Then there is a complete train track τ on S in standard form for a marking F of S with pants decomposition P which carries λ and ϵ-follows λ.*

Proof. The train track τ can be obtained from λ by collapsing a sufficiently small tubular neighborhood of λ. We refer to Theorem 1.6.5 of [PH92] and to Lemma 3.2 of [H09] and its proof for more details of this construction. □

Note that in Lemma 2.1, the marking F of S depends on the number ϵ as well as on choices made in the construction.

A *measured geodesic lamination* is a geodesic lamination equipped with a transverse translation invariant measure of full support. The space $\mathcal{ML}$ of measured geodesic laminations on S equipped with the weak*-topology is homeomorphic to the product of a sphere of dimension $6g - 7 + 2m$ with the real line. A measured geodesic lamination μ is carried by a train track τ if its support is carried by τ. Then μ defines a transverse measure on τ, and every transverse measure on τ arises in this way [PH92].

We use measured geodesic laminations to establish another relation between train tracks in standard form for a marking of S and complete geodesic laminations which is a variant of a result of Penner and Harer (Theorem 2.8.4 of [PH92]).

Lemma 2.2. *For any marking F of S, every complete geodesic lamination on S is carried by a unique train track in standard form for F.*

Proof. A complete geodesic lamination λ can be approximated in the Hausdorff topology by a sequence $\{c_i\}$ of simple closed geodesics. For a fixed marking F of S, each such geodesic is carried by a train track in standard form for F (this is contained in Theorem 2.8.4 of [PH92]). Since there are only finitely many train tracks in standard form for F, there is a fixed train track τ in standard form for F which carries infinitely many of the curves c_i. By Lemma 3.2 of [H09], the set of geodesic laminations carried by a fixed train track τ is closed in the Hausdorff topology and hence the geodesic lamination λ is carried by τ.

Now assume that there is a second train track η in standard form for F which carries λ. By Lemma 3.2 and Lemma 3.3 of [H09], there is a complete train track σ which is carried by both τ and η. The train track σ carries a measured geodesic lamination μ whose support is both minimal and maximal (see the top of p.556 of [H09] for a detailed discussion of this fact). Thus μ is carried by two distinct train tracks in standard form for F which violates Theorem 2.8.4 of [PH92]. □

A half-branch $\hat{b}$ in a generic train track τ incident on a switch v of τ is called *large* if every trainpath containing v in its interior passes through $\hat{b}$. A half-branch which is not large is called *small.* A branch b in a generic train track τ is called *large* if each of its two half-branches is large; in this case b is necessarily incident on two distinct switches. A branch is called *small* if each of its two half-branches is small. A branch is called *mixed* if one of its half-branches is large and the other half-branch is small (see p.118 of [PH92]).

There are two simple ways to modify a complete train track τ to another complete train track. First, we can *shift* τ along a mixed branch b to a train track τ' as shown in Figure C. If τ is complete then the same is true for τ'. Moreover, a train track

Figure C

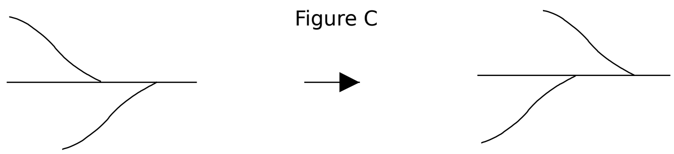

or a geodesic lamination is carried by τ if and only if it is carried by τ' (see [PH92] p.119). In particular, the shift τ' of τ is carried by τ. There is a natural bijection $\varphi(\tau,\tau')$ of the set of branches of τ onto the set of branches of τ' which is induced by the identity of the complement of a small neighborhood of b in S. The bijection $\varphi(\tau,\tau')$ also induces a bijection of the set of half-branches of τ onto the set of half-branches of τ' which we denote again by $\varphi(\tau,\tau')$.

Second, if e is a large branch of τ then we can perform a right or left *split* of τ at e as shown in Figure D. Note that a right split at e is uniquely determined by

Figure D

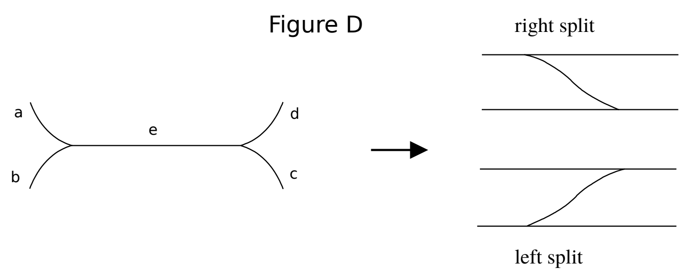

the orientation of S and does not depend on the orientation of e. Using the labels in the figure, in the case of a right split we call the branches a and c *winners* of the split, and the branches b,d are *losers* of the split. If we perform a left split, then the branches b,d are winners of the split, and the branches a,c are losers of the split. The split τ' of a train track τ is carried by τ, and there is a natural choice of a carrying map which maps the switches of τ' to the switches of τ. The image of a branch of τ' is then a trainpath on τ whose length either equals one or two. There is a natural bijection $\varphi(\tau,\tau')$ of the set of branches of τ onto the set of branches of τ' which maps the branch e to a small branch e' which we call the *diagonal* of the split. This bijection is induced by the identity on the complement of a small neighborhood of e in S. The map $\varphi(\tau,\tau')$ also induces a bijection of the set of half-branches of τ onto the set of half-branches of τ' again denoted by $\varphi(\tau,\tau')$.

Occasionally we also have to consider the *collision* of a train track η at a large branch e. This collision is obtained from η by a split at e and removal of the diagonal in the split track. Such a collision is shown in Figure 2.1.2 of [PH92].

A split of a maximal transversely recurrent generic train track is maximal, transversely recurrent and generic. If τ is a complete train track and if $\lambda \in \mathcal{CL}$ is carried by τ, then for every large branch e of τ there is a unique choice of a right or left split of τ at a large branch e of τ with the property that the split track τ' carries λ (see p. 557 of [H09] for a more complete discussion). We call such a split a λ-*split*. The train track τ' is complete. In particular, a complete train track τ can always be split at any large branch e to a complete train track τ'; however there may be a choice of a right or left split at e such that the resulting train track is not recurrent any more (compare p.120 in [PH92]).

For a number $L \geq 1$, an *L-quasi-isometric embedding* of a metric space (X, d) into a metric space (Y, d) is a map $\varphi : X \to Y$ such that

$$d(x, y)/L - L \leq d(\varphi(x), \varphi(y)) \leq Ld(x, y) + L$$

for all $x, y \in X$. The map φ is called an *L-quasi-isometry* if moreover the L-neighborhood of φX in Y is all of Y. An *L-quasi-geodesic* in a metric space (X, d) is an L-quasi-isometric embedding of a closed connected subset of $\mathbb{R}$ or of the intersection of such a closed connected subset of $\mathbb{R}$ with $\mathbb{Z}$.

Denote by $\mathcal{TT}$ the directed metric graph whose set $\mathcal{V}(\mathcal{TT})$ of vertices is the set of isotopy classes of complete train tracks on S and whose edges are determined as follows. The train track $\tau \in \mathcal{V}(\mathcal{TT})$ is connected to the train track τ' by a directed edge of length one if and only if τ' can be obtained from τ by a single split. The graph $\mathcal{TT}$ is connected (Corollary 2.7 of [H09]). The mapping class group $\mathcal{MCG}(S)$ of S acts properly and cocompactly on $\mathcal{TT}$ as a group of simplicial isometries. In particular, $\mathcal{TT}$ is $\mathcal{MCG}(S)$-equivariantly quasi-isometric to $\mathcal{MCG}(S)$ equipped with any word metric (Corollary 4.4 of [H09]).

Define a *splitting sequence* in $\mathcal{TT}$ to be a sequence $\{\alpha(i)\}_{0 \leq i \leq m} \subset \mathcal{V}(\mathcal{TT})$ with the property that for every $i \geq 0$ the train track $\alpha(i+1)$ can be obtained from $\alpha(i)$ by a single split. Thus splitting sequences in $\mathcal{V}(\mathcal{TT})$ correspond precisely to directed edge-paths in $\mathcal{TT}$. If τ can be connected to η by a splitting sequence then we say that τ is *splittable* to η. If $\{\alpha(i)\}_{0 \leq i \leq m}$ is a splitting sequence then the composition

$$\varphi(\alpha(0), \alpha(m)) = \varphi(\alpha(m-1), \alpha(m)) \circ \cdots \circ \varphi(\alpha(0), \alpha(1))$$

(read from right to left) is a bijection of the branches (or half-branches) of $\alpha(0)$ onto the branches (or half-branches) of $\alpha(m)$ which does not depend on the choice of the splitting sequence connecting $\alpha(0)$ to $\alpha(m)$ (Lemma 5.1 of [H09]).

3. Density of splitting sequences

The goal of this section is to show the following proposition which is the main technical tool for the proof of Theorem 1 and Theorems 5-7.

Proposition 3.1. *There is a number $d_0 > 0$ with the following property. For any train tracks $\tau, \sigma \in \mathcal{V}(\mathcal{TT})$ there is a train track τ' which is contained in the d_0-neighborhood of τ and which is splittable to a train track σ' contained in the d_0-neighborhood of σ.*

To simplify the argument we reduce Proposition 3.1 to the following

Proposition 3.2. *There is a number $d_1 > 0$ with the following property. Let F be any marking of S. Then for any train track $\sigma \in \mathcal{V}(\mathcal{TT})$ there is a train track $\tau \in \mathcal{V}(\mathcal{TT})$ in standard form for F which carries a train track σ' contained in the d_1-neighborhood of σ. If σ is in standard form for a marking G with pants decomposition Q, then σ' can be chosen to contain the pants decomposition Q as an embedded subtrack.*

We begin with explaining how Proposition 3.1 follows from Proposition 3.2. The mapping class group acts on the set of all markings of S, with finitely many orbits, and it acts properly and cocompactly on $\mathcal{TT}$ preserving the set of train tracks in standard form for some marking of S. Thus every complete train track is contained in a uniformly bounded neighborhood of a train track in standard form for some marking F of S. Furthermore, the diameter in $\mathcal{TT}$ of a set of train tracks in standard form for a fixed marking is bounded from above by a universal constant. Thus there is a number $d_2 > 0$, and for every complete train track τ on S there is a marking F of S such that $d(\tau, \eta) \leq d_2$ for any train track η in standard form for F (here d is the distance on $\mathcal{TT}$).

By Lemma 6.6 of [H09], there is a number $p > 0$, and for two complete train tracks $\sigma \prec \tau$ there is a train track ζ which can be obtained from τ by a splitting sequence and such that $d(\sigma, \zeta) \leq p$. As a consequence, Proposition 3.1 follows from Proposition 3.2.

The idea of proof for Proposition 3.2 is as follows. Define a *splitting and shifting sequence* to be a sequence $\{\alpha(i)\}_{0 \leq i \leq m}$ with the property that for every $i \geq 0$ the train track $\alpha(i+1)$ can be obtained from $\alpha(i)$ by a sequence of shifts followed by a single split. Theorem 2.4.1 of [PH92] relates splitting and shifting to carrying.

Proposition 3.3. *If $\sigma \in \mathcal{V}(\mathcal{TT})$ is carried by $\tau \in \mathcal{V}(\mathcal{TT})$ then τ can be connected to σ by a splitting and shifting sequence.*

Now let F, G be any two markings of S. We attempt to construct a splitting and shifting sequence connecting some train track in standard form for F to some train track in standard form for G.

A train track in standard form for G with only twist connectors carries a complete geodesic lamination λ in standard form for the pants decomposition Q of the marking G of S. By Lemma 2.2, every complete geodesic lamination λ on S is carried by a unique train track τ in standard form for F. We modify τ with a sequence of splits as efficiently as possible to a train track ξ which carries λ and contains the pants decomposition Q as a subtrack. This train track ξ carries a train track η in standard form for some marking of S with pants decomposition Q whose distance to ξ is uniformly bounded. There is a *multi-twist* φ (that is, a concatenation of mutually commuting Dehn twists) about the pants curves of Q which maps η to a train track $\varphi\eta$ in standard form for G. In general, $\varphi\eta$ is not carried by η, but we obtain enough control that we can find a perhaps different train track in standard form for F which can be connected with a splitting and shifting sequence to a train track in a uniformly bounded neighborhood of $\varphi\eta$ which is in standard form for a marking with pants decomposition Q.

To carry out this strategy we use the pants decomposition Q for the construction of splitting sequences. However, Q may not *fill* τ, that is, a carrying map $Q \to \tau$ may not be surjective. Therefore we are lead to investigate splitting sequences of complete train tracks which are determined by modifications of subtracks. This will occupy the major part of this section and will also be important in Section 7.

Fix a complete Riemannian metric on S of finite volume. With respect to this metric, a complementary region C of a train track σ on S is a hyperbolic surface whose metric completion $\overline{C}$ is a bordered surface with boundary ∂C. This boundary consists of a finite number of arcs of class C^1, called *sides* of C or of $\overline{C}$. Each side of C either is a closed curve of class C^1 (that is, the boundary component containing the side does not contain any cusp) or an arc with endpoints at two not necessarily distinct cusps of the component. We call a side of C which does not contain cusps a *smooth side* of C. The closure of C in S can be obtained from $\overline{C}$ by some identifications of subarcs of sides (the inclusion $C \to S$ extends to an immersion of each side of C, but the image arc may have tangential self-intersections or may meet another side tangentially). For simplicity we call the image in σ of a side of C a side of C as well. Most of the time we view a side of C as an immersed arc of class C^1 in σ. Using this abuse of notation, a side of C then is an immersed arc or an immersed closed curve of class C^1 in σ with only tangential self-intersections. However, we reserve the notation $\overline{C}$ for the metric completion of C.

If $T \subset \partial C$ is a smooth side of a complementary region C of σ then we mark a point on T. We view this point as a single point on the boundary of the completion $\overline{C}$ of C, even if the point corresponds to a point of tangential self-intersection of the image of ∂C in σ and hence its preimage in $\overline{C}$ under the natural map $\overline{C} \to S$ consists of more than one point.

If C is a complementary region of σ whose boundary contains precisely $k \geq 0$ cusps, then the *Euler characteristic* $\chi(C)$ is defined by $\chi(C) = \chi_0(C) - k/2$ where $\chi_0(C)$ is the usual Euler characteristic of the compact topological surface with boundary $\overline{C}$. Note that the sum of the Euler characteristics of the complementary regions of σ is just the Euler characteristic of S (see the discussion in Chapter 1.1 of [PH92]).

A *complete extension* of a train track σ is a complete train track τ containing σ as a subtrack. We require that the switches of τ are distinct from the images in σ of the marked points on smooth boundary components of complementary regions of σ. Such a complete extension τ intersects each complementary region C of σ in an embedded graph with smooth edges. The closure of $\tau \cap C$ in the completion $\overline{C}$ of C is a graph whose univalent vertices are contained in the complement of the cusps and marked points of the boundary ∂C of $\overline{C}$. At a univalent vertex, the graph is tangential to ∂C. We call two such graphs $\tau \cap C, \tau' \cap C$ *equivalent* if there is a smooth isotopy of $\overline{C}$ which fixes the cusps and the marked points in ∂C and which maps $\tau \cap C$ onto $\tau' \cap C$. The complete extensions τ, τ' of σ are called *σ-equivalent* if for each complementary region C of σ the graphs $\tau \cap C$ and $\tau' \cap C$ are equivalent in this sense. The purpose of marking a point on a smooth boundary component T of a complementary region of σ is to control the amount of relative twisting about T of two complete extensions τ, τ' of σ.

For two complete extensions τ, τ' of σ define the *intersection number* $i_\sigma(\tau, \tau')$ to be the minimal number of intersection points contained in $S - \sigma$ between any two complete extensions η, η' of σ which are σ-equivalent to τ, τ' and with the following additional properties.

a) A switch v of η (or η') is also a switch of η' (or η) if and only if v is a switch of σ.
b) A switch of η (or η') contained in the interior of a complementary region C of σ is not contained in η' (or η), that is, an intersection point of η with η' contained in C is an interior point of a branch of η and of a branch of η'.

Since the number of switches of a complete train track on S only depends on the topological type of S, for any complete extension τ of σ the intersection number $i_\sigma(\tau, \tau)$ is bounded from above by a constant only depending on S but neither on σ nor on τ. Moreover, for every number $m > 0$ there is a number $q(m) > 0$ not depending on σ so that for every complete extension τ of σ, the number of σ-equivalence classes of complete extensions τ' of σ with $i_\sigma(\tau, \tau') \leq m$ is bounded from above by $q(m)$.

To simplify the notation we do not distinguish between σ as a subgraph of τ (and hence containing switches of τ which are bivalent in σ) and σ viewed as a subtrack of τ, that is, the graph from which the bivalent switches not contained in simple closed curve components have been removed. A branch e of σ defines an embedded trainpath $\rho : [0, m] \to \tau$, unique up to orientation, whose image is precisely e. We call τ *tight* at e if e is a branch in τ, that is, if the length m of ρ equals one. If e is a large branch of σ, then ρ begins and ends with a large half-branch and hence $\rho[0, m]$ contains a large branch of τ (Lemma 2.7.2 of [PH92]).

A *proper subbranch* of a branch e of σ is a branch b of τ which is a proper subset of e. Then b is incident on at least one switch v of τ which is not a switch of σ. There is a half-branch c of τ which is incident on v and not contained in σ. We call c a *neighbor* of σ at v. We distinguish three different types of *large* proper subbranches b of a branch e of σ. These types are shown in Figure E. Note that a large branch of any train track on S is embedded in S.

Figure E

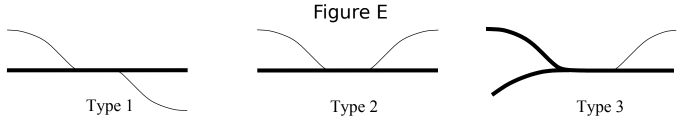

Type 1: b is contained in the interior of e and the two neighbors of σ at the two endpoints of b lie on different sides of e in a small tubular neighborhood of b in S.

Type 2: b is contained in the interior of e and both neighbors of σ at the two endpoints of b lie on the same side of e in a small tubular neighborhood of b in S.

Type 3: One endpoint of b is incident on a switch of σ.

A split of τ at a large proper subbranch b of σ (that is, of a large branch of τ which is a proper subbranch of a branch of σ) is called a *σ-split* if the split track contains σ as a subtrack. Note that such a split always exists. If b is a large proper subbranch of σ of type 2 then any split of τ at b is a σ-split.

Let q be the number of branches of a complete train track on S. The number of switches of a complete train track on S then equals $2q/3 < q$. For a subtrack σ of a complete train track τ let $\beta(\tau,\sigma)$ be the number of neighbors of σ in τ, that is, the number of half-branches of $\tau - \sigma$ which are incident on a switch contained in σ. If τ_1 is obtained from τ by a split at a large proper subbranch b of σ of type 2 then $\beta(\tau_1,\sigma) = \beta(\tau,\sigma) - 1$.

Now let σ be a recurrent train track on S. Then there is a measured geodesic lamination ν on S which is carried by σ and which defines a positive transverse measure on σ. We call such a measured geodesic lamination *filling* for σ. For every complete extension τ of σ there is a complete geodesic lamination λ which is carried by τ and contains the support of ν as a sublamination. Namely, the positive transverse measure on σ defined by ν can be approximated by positive transverse measures μ_i on τ which define a measured geodesic lamination whose support is a minimal and maximal geodesic lamination carried by τ (see p.556 of [H09] for a detailed proof of this fact). Since the space $\mathcal{CL}$ of all complete geodesic laminations on S is compact, as $\mu_i \to \nu$ in the space of transverse measures on τ, up to passing to a subsequence the supports of μ_i converge in the Hausdorff topology to a complete geodesic lamination λ which contains the support of ν as a sublamination. By Lemma 2.3 of [H09], λ is carried by τ. We call λ a *complete τ-extension* of ν.

The following observation will be used throughout several times in the sequel.

Lemma 3.4. *Let σ be a recurrent subtrack of a complete train track τ and let λ be a complete τ-extension of a σ-filling measured geodesic lamination. Then for every large branch e of σ there is a unique train track τ' with the following properties.*

1. *τ' can be obtained from τ by at most q^2 σ-splits at large proper subbranches of e. In particular, τ' contains σ as a subtrack.*
2. *τ' carries λ and is complete.*
3. *τ' is tight at e.*

Moreover, given any complete extension η of σ, if the branch e does not contain any marked point on a smooth side of a complementary region C of σ then

$$i_\sigma(\tau',\eta) \le i_\sigma(\tau,\eta) + q(\beta(\tau,\sigma) - \beta(\tau',\sigma)).$$

Otherwise we have $i_\sigma(\tau',\eta) \le i_\sigma(\tau,\eta) + q^3$.

Proof. If τ is tight at the large branch e of σ then $\tau = \tau'$ satisfies the requirements in the lemma.

Otherwise let a be a neighbor of σ at a switch of τ contained in e. There is a unique maximal trainpath $\rho : [-1, m] \to \tau$ with $\rho[-1/2, 0] = a$ and such that $\rho[0, m] \subset e$. Then $\rho(m)$ is a switch of σ on which e is incident. Let $c(a, e) = m \leq q$ be the length of the intersection of the trainpath ρ with e and let

$$c(\tau, e) = \sum_a c(a, e)$$

where the sum is taken over all neighbors of e in τ. Since e has at most q neighbors we have $c(\tau, e) \leq q^2$, and $c(\tau, e) = 0$ if and only if τ is tight at e.

Now note (see [PH92]) that any trainpath $\rho : [0, m] \to \tau$ which begins and ends with a large half-branch contains a large branch. This is clear is the length m of ρ equals one, so by induction, let us assume that it holds true whenever this length is at most $m - 1 \geq 1$. If $\rho : [0, m] \to \tau$ is trainpath beginning and ending with a large half-branch and if the branch $\rho[0, 1]$ is not large, then $\rho[0, 1]$ is mixed and $\rho[1, m]$ is a trainpath beginning and ending with a large half-branch. By induction hypothesis, this trainpath contains a large branch.

Thus let $b \subset \tau$ be a large proper subbranch of e. Let λ be a complete τ-extension of a σ-filling measured geodesic lamination ν and let τ_1 be the train track obtained from τ by a λ-split at b. We distinguish two cases according to the type of b.

If b is of type 1 or of type 3, then there is a unique choice of a right or left split of τ at b such that the split track τ_1' contains σ as a subtrack. Since ν is a σ-filling measured geodesic lamination, τ_1' is also the unique train track obtained from τ by a split at b which carries ν. Now λ is a complete extension of ν and therefore $\tau_1' = \tau_1$. The natural bijection $\varphi(\tau, \tau_1)$ of the half-branches of τ onto the half-branches of τ_1 maps any neighbor a of σ in τ to a neighbor $\varphi(\tau, \tau_1)(a)$ of σ in τ_1, and $c(\varphi(\tau, \tau_1)(a), e) \leq c(a, e)$. If the neighbor a of σ is incident at an endpoint of b, then we have $c(\varphi(\tau, \tau_1)(a), e) = c(a, e) - 1$ (see Figure E). Together this shows that $c(\tau_1, e) \leq c(\tau, e) - 1$.

If b is of type 2 then once again, the train track τ_1 contains σ as a subtrack. Moreover, we have $\beta(\tau_1, \sigma) = \beta(\tau, \sigma) - 1$ and $c(\tau_1, e) < c(\tau, e)$. As a consequence, a splitting sequence of length at most q^2 at large proper subbranches of e transforms τ to a train track τ' which contains σ as a subtrack, is tight at e and carries λ. By uniqueness of sequences of λ-splits up to order (Lemma 5.1 of [H09]), the train track τ' is uniquely determined by τ, σ, λ, e.

To estimate intersection numbers between τ, τ' and an arbitrary complete extension η of σ, let again b be a large proper subbranch of e of type 1 or type 3 and let τ_1 be the train track obtained from τ by a λ-split at b. If e does not contain the image of any marked point on a smooth boundary component of a complementary region of σ, then τ and τ_1 are σ-equivalent.

Otherwise there are one or two (not necessarily distinct) complementary regions C_1, C_2 of σ and smooth sides T_i of C_i whose images in σ contain e. Up to isotopy, a split of τ at b can be realized by moving one of the neighbors of σ incident on an endpoint of b, say the neighbor a, across b while leaving the second neighbor (or the branches of σ incident on an endpoint of e in case the branch b is of type 3) fixed. Assume that the half-branch a is contained in the complementary region C_1 of σ and that a terminates at a point in the smooth boundary component T_1 of C_1. There are at most q half-branches of η contained in $\eta \cap C_1$ which terminate at a point in T_1. Up to isotopy of $\overline{C_1} \cup \overline{C_2}$ preserving the cusps and the marked points, moving the half-branch a of τ across the marked point in T_1 increases the number of intersection points between τ and η by at most q. Namely, up to isotopy such a move creates at most one additional intersection point with any half-branch of η with endpoint T_1. As a consequence, we have

$$i_\sigma(\tau_1, \eta) \le i_\sigma(\tau, \eta) + q. \tag{1}$$

If τ_1 is obtained from τ by a λ-split at a large proper subbranch of e of type 2 then the split which modifies τ to τ_1 can be realized by moving one of the neighbors of σ incident on an endpoint of b across b to a half-branch which is incident on a point in the interior of the neighbor of σ at the second endpoint of b. As before, this implies that the inequality (1) holds true for every complete extension η of σ (independent of whether or not e contains the image of a marked point).

To summarize, if there is no marked point on the boundary of a complementary region of σ which is mapped into e and if η is any complete extension of σ, then the above discussion shows that only splits at large proper subbranches of e of type 2 change the intersection number between τ and η. A successive application of the estimate (1) yields that

$$i_\sigma(\tau', \eta) \le i_\sigma(\tau, \eta) + q(\beta(\tau, \sigma) - \beta(\tau', \sigma)).$$

The second estimate of intersection numbers stated in the lemma follows in the same way from the inequality (1). □

For a recurrent subtrack σ of a complete train track τ, for a large branch e of σ and a complete τ-extension λ of a σ-filling measured geodesic lamination, we call the complete train track τ' constructed in Lemma 3.4 the (e, λ)-*modification* of τ.

Remark 3.5. 1) Lemma 3.4 and its proof remain valid if the large branch e of a recurrent subtrack σ of τ is replaced by any embedded trainpath $\rho : [0, m] \to \tau$ which begins and ends with a large half-branch. A large branch e of a non-recurrent subtrack of τ is an example. In this case the complete geodesic lamination λ has to be replaced by a measured geodesic lamination whose support is minimal and complete and is carried by τ and which defines a transverse measure on τ giving positive weight to the set of arcs which are mapped homeomorphically *onto* $\rho[0, m]$ by a carrying map. We call such a complete geodesic lamination ρ-*filling*, and we call the train track obtained from τ, ρ, λ with the procedure from the proof of Lemma 3.4 the (ρ, λ)-*modification* of τ. Note that a ρ-filling complete geodesic lamination may not always exist.

2) Let e_1, e_2 be distinct large branches of a train track σ on S, let τ be a complete extension of σ and let λ be a complete τ-extension of a σ-filling measured geodesic lamination. Denote by τ_1, τ_2 the complete train tracks constructed from σ and λ as in Lemma 3.4 which are tight at the large branch e_1, e_2. Then up to isotopy,

for every neighborhood U_1, U_2 of e_1, e_2 in S the intersection $\tau_i \cap (S - U_i)$ coincides with the intersection $\tau \cap (S - U_i)$ $(i = 1, 2)$. As a consequence, the train track $\tau_{1,2}$ obtained from τ_1, σ, λ by the construction in Lemma 3.4 which is tight at the large branch e_2 coincides with the train track $\tau_{2,1}$ obtained from τ_2, σ, λ by the construction in Lemma 3.4 which is tight at e_1.

If σ is any train track on S and if τ, η are two complete extension of σ, then we defined an intersection number $i_\sigma(\tau, \eta)$ which depends on the choice of marked points, one on each smooth boundary component of a complementary region of σ. A different choice of a marked point only changes the intersection number up to a uniformly bounded amount (compare the proof of Lemma 3.4 for a more detailed explanation and recall that the choice of the marked point is needed to control twisting of η relative to τ along the smooth boundary components of σ). The last statement of the following proposition then means that there are choices of marked points on σ, σ_ℓ so that the stated inequality holds true for these choices.

For a precise formulation, for a train track τ which is splittable to a train track η (that is, such that τ can be connected to η by a splitting sequence) denote by

$$E(\tau, \eta) \subset \mathcal{TT} \tag{2}$$

the graph whose vertex set consists of all train tracks which can be obtained from τ by a splitting sequence and which are splittable to η and where such a vertex ξ is connected to a vertex ζ by a directed edge of length one if ζ can be obtained from ξ by a single split.

Call a splitting sequence $\{\sigma_i\}$ of train tracks on S *recurrent* if each of the train tracks σ_i is recurrent.

Proposition 3.6. *Given a recurrent splitting sequence $\{\sigma_i\}_{0 \le i \le \ell}$ of train tracks on S, there is an algorithm which associates to a complete extension τ of σ_0 and a complete τ-extension λ of a σ_ℓ-filling measured geodesic lamination ν a sequence $\{\tau_i\}_{0 \le i \le 2\ell} \subset \mathcal{V}(\mathcal{TT})$ with the following properties.*

1. *$\tau_0 = \tau$, and for each $i \le \ell$ the train tracks τ_{2i}, τ_{2i+1} contain σ_i as a subtrack and carry λ.*
2. *If σ_{i+1} is obtained from σ_i by a right (or left) split at a large branch e_i then τ_{2i+1} is the (e_i, λ)-modification of τ_{2i}, and τ_{2i+2} is obtained from τ_{2i+1} by a right (or left) split at e_i.*
3. *The train track $\tau_{2\ell}$ only depends on $\tau, \sigma, \sigma_\ell, \lambda$ but not on the choice of a splitting sequence connecting σ to σ_ℓ.*
4. *Every complete train track $\tau' \in E(\tau, \tau_{2\ell})$ contains a subtrack $\sigma' \in E(\sigma_0, \sigma_\ell)$.*
5. *If $\{\eta_i\}_{0 \le i \le 2\ell}$ is another such sequence beginning with a complete extension $\eta = \eta_0$ of σ then*
$$i_{\sigma_\ell}(\tau_{2\ell}, \eta_{2\ell}) \le i_\sigma(\tau, \eta) + 4q^5.$$

Proof. Let σ' be a train track which can be obtained from a train track σ by a single split at a large branch e. Let U be any neighborhood of e in S. Then up to modifying σ' with an isotopy we may assume that $\sigma' \cap (S - U) = \sigma \cap (S - U)$ and that there is a map $F : S \to S$ of class C^1 which equals the identity on $S - U$ and which restricts to a carrying map $\sigma' \to \sigma$. In particular, there is a natural bijection ψ between the complementary regions of σ and the complementary regions of σ' which preserves the topological type of the regions and which maps a complementary region C of σ to the complementary region $\psi(C)$ of σ' containing $C - U$ (here we assume that U is sufficiently small that $C - U \neq \emptyset$ for every complementary region C of σ).

If T is a smooth side of C then there is a smooth side T' of $\psi(C)$ whose image in σ' is mapped by the carrying map F onto the image of T in σ. Let $\rho : [0, n] \to \sigma$ be a trainpath which parameterizes the image of T in σ. Then ρ passes through any branch of σ at most twice, in opposite direction. In particular, the length n of ρ is at most $2q$ where as before, q is the number of branches of a complete train track on S (which is the maximal number of branches of any train track on S). If $\rho[0, n]$ contains the branch e, then the image in σ' of the side T' of $\psi(C)$ does not pass through the diagonal branch of the split. As a consequence, the length of a trainpath ρ' on σ' parameterizing the image of T' is strictly smaller than the length n of the trainpath on ρ parameterizing the image of T (see Figure E).

The number of distinct smooth boundary components of complementary regions of σ is bounded from above by $3g - 3 + m < q/2$. If $\{\sigma_i\}_{0\leq i\leq \ell}$ is any splitting sequence, then the discussion in the previous paragraph shows that there are at most q^2 numbers $i \in \{1, \dots, \ell\}$ such that σ_{i+1} is obtained from σ_i by a split at a large branch which is contained in the image of a smooth boundary component of a complementary region of σ_i.

Now let τ, η be complete extensions of a recurrent train track $\sigma = \sigma_0$. As in the beginning of this section, mark a point on each smooth boundary component of a complementary region of σ in such a way that no marked point of σ is a switch of either τ or η (this can always be achieved with a small isotopy of τ, η preserving σ as a set). Let $\{\sigma_i\}_{0\leq i\leq \ell}$ be a recurrent splitting sequence issuing from $\sigma = \sigma_0$ and let λ, μ be complete τ, η-extensions of a σ_ℓ-filling measured geodesic lamination ν. We construct sequences $\{\tau_i\}_{0\leq i\leq 2\ell}, \{\eta_i\}_{0\leq i\leq 2\ell} \subset \mathcal{V}(\mathcal{TT})$ with the properties stated in the proposition inductively as follows.

Let $\tau_0 = \tau, \eta_0 = \eta$ and assume that the train tracks τ_{2i}, η_{2i} have already been constructed for some $i \geq 0$. Assume that σ_{i+1} is obtained from σ_i by a right (or left) split at the large branch e_i. Define τ_{2i+1}, η_{2i+1} to be the (e_i, λ)-modification (or the (e_i, μ)-modification, respectively) of τ_{2i}, η_{2i}. By construction, these train tracks carry the geodesic laminations λ, μ, and they are tight at e_i.

Since ν is σ_ℓ-filling, the right (or left) split of σ_i at e_i is the unique split so that the split track carries ν. Namely, otherwise ν is carried by the train track obtained from σ_i by splitting at e_i and removing the diagonal of the split. But this then means that a carrying map $\nu \to \sigma_{i+1}$ is not surjective which violates the assumption that ν fills $\sigma_\ell \prec \sigma_{i+1}$. Define τ_{2i+2}, η_{2i+2} to be the train track obtained from τ_{2i+1}, η_{2i+1} by a right (or left) split at the large branch e_i. Then τ_{2i+2}, η_{2i+2} contains σ_{i+1} as a subtrack, and by the above reasoning, it is the unique train track obtained from τ_{2i+1}, η_{2i+1} by a split at e_i which carries ν. On the other hand, there is a unique choice of a split of τ_{2i+1}, η_{2i+1} at e_i so that the split track carries λ, μ and hence ν. But ν is a sublamination of λ, μ and therefore the train tracks τ_{2i+2}, η_{2i+2} carry λ, μ. In particular, these train tracks are complete.

As a consequence, the inductively defined sequences $\{\tau_i\}_{0\leq i\leq 2\ell}, \{\eta_i\}_{0\leq i\leq 2\ell}$ have properties 1)-2) stated in the proposition. The third property follows from the fact that a splitting sequence connecting σ to σ_ℓ is unique up to order (Lemma 5.1 of [H09] is also valid for splitting sequences of train tracks which are not complete since the assumption of completeness is nowhere used in the proof) and from the second remark after Lemma 3.4. Namely, by this remark, the train track obtained from a complete extension τ of σ and a complete τ-extension of a σ-filling measured geodesic lamination λ by two consecutive applications of Lemma 3.4 at distinct large branches e_1, e_2 of σ only depends on $\tau, \sigma, \lambda, e_1, e_2$ but not on the order in which these two applications of Lemma 3.4 are carried out.

Property 4) follows in the same way by induction on the length of a splitting sequence connecting τ to $\tau_{2\ell}$. If this length vanishes then there is nothing to show, so assume that the claim holds true whenever the length of such a sequence does not exceed $n-1$ for some $n \geq 1$. Under the hypotheses used throughout this proof, assume that the length of a splitting sequence connecting τ to $\tau_{2\ell}$ equals n.

Let $\tau' \in E(\tau, \tau_{2\ell})$. If $\tau' = \tau$ then τ' contains σ as a subtrack and there is nothing to show. Otherwise there is a train track $\tilde{\tau} \in E(\tau, \tau') \subset E(\tau, \tau_{2\ell})$ which can be obtained from τ by a single split at a large branch b. By uniqueness of splitting sequences (Lemma 5.1 of [H09]), we have $b \subset \sigma$.

If b is a large branch of σ (that is, if τ is tight at b), then it follows once again by uniqueness of splitting sequences that $\tilde{\tau}$ contains a subtrack $\tilde{\sigma} \in E(\sigma, \sigma_\ell)$ which can be obtained from σ by a single split at b. Property 4) now follow s from property 3) and the induction hypothesis, applied to $\tilde{\tau}, \tau_{2\ell}, \tilde{\sigma}, \sigma_\ell, \tau'$. Otherwise b is a large proper subbranch of σ. If b is of type 2 then any split of τ at b contains σ as a subtrack. If b is of type 1 or type 3 then there is a unique split of τ at b so that the split track contains σ as a subtrack, and by the previous discussion, the split track coincides with $\tilde{\tau}$. Once again, we can apply the induction hypothesis to $\tilde{\tau}, \tau_{2\ell}, \sigma, \sigma_\ell, \tau'$ to complete the induction step and hence the proof of property 4).

We are left with the verification of property 5). For this we control the increase of intersection numbers between the train tracks τ_{2i}, η_{2i} and τ_{2i+2}, η_{2i+2}. This is done by distinguishing two cases.

Case 1: No marked point of a smooth side of a complementary component of σ_i is mapped into the large branch e_i of σ_i.

By two applications of Lemma 3.4, in this case we have

$$\begin{aligned} i_{\sigma_i}(\tau_{2i+1}, \eta_{2i+1}) &\leq i_{\sigma_i}(\tau_{2i+1}, \eta_{2i}) + q(\beta(\eta_{2i}, \sigma_i) - \beta(\eta_{2i+1}, \sigma_i)) \\ &\leq i_{\sigma_i}(\tau_{2i}, \eta_{2i}) + q(\beta(\tau_{2i}, \sigma_i) - \beta(\tau_{2i+1}, \sigma_i)) + q(\beta(\eta_{2i}, \sigma_i) - \beta(\eta_{2i+1}, \sigma_i)). \end{aligned}$$

Let ψ be the natural bijection between the complementary regions of σ_i and the complementary regions of σ_{i+1} as introduced in the first paragraph of this proof. Up to isotopy, for an arbitrary given neighborhood U of e_i in S and for any complementary region C of σ_i, there is a diffeomorphism F of the completion $\overline{C}$ of C onto the completion $\overline{\psi(C)}$ of $\psi(C)$ respecting cusps and marked points and which equals the identity outside of U. Since τ_{2i+2} is obtained from τ_{2i+1} by a right or left split at the tight large branch b, and the corresponding right or left split of σ_i at b transforms σ_i to σ_{i+1}, the intersection $\tau_{2i+2}\cap\psi(C)$ is equivalent to $F(\tau_{2i+1}\cap C)$, and $\eta_{2i+2}\cap\psi(C)$ is equivalent to $F(\eta_{2i+1}\cap C)$. This shows that

$$\begin{aligned} &i_{\sigma_{i+1}}(\tau_{2i+2},\eta_{2i+2}) = i_{\sigma_i}(\tau_{2i+1},\eta_{2i+1}) \\ &\quad\le i_{\sigma_i}(\tau_{2i},\eta_{2i}) + q(\beta(\tau_{2i},\sigma_i)-\beta(\tau_{2i+1},\sigma_i)) + q(\beta(\eta_{2i},\sigma_i)-\beta(\eta_{2i+1},\sigma_i)). \end{aligned}$$

Case 2: There is a marked point on a smooth boundary component of a complementary region of σ_i which is mapped into the branch e_i.

In this case, e_i is contained in the image of one or two smooth boundary components T_1, T_2 of complementary regions of σ_i. By two applications of Lemma 3.4, we have

$$i_{\sigma_i}(\tau_{2i+1},\eta_{2i+1}) \le i_{\sigma_i}(\tau_{2i},\eta_{2i+1}) + q^3 \le i_{\sigma_i}(\tau_{2i},\eta_{2i}) + 2q^3.$$

The marked points on T_1, T_2 determine marked points on smooth sides T_1', T_2' of complementary regions of σ_{i+1} so that we have

$$i_{\sigma_{i+1}}(\tau_{2i+2},\eta_{2i+2}) = i_{\sigma_i}(\tau_{2i+1},\eta_{2i+1}) \le i_{\sigma_i}(\tau_{2i},\eta_{2i}) + 2q^3.$$

Now by the consideration in the beginning of this proof, Case 2 can occur at most q^2 times. Moreover, the number of neighbors of σ in τ,η is bounded from above by the upper bound q for the number of switches of τ,η and hence this number can not be decreased by more than q in this process. Together we conclude that there are at most $q^2+2q\le 2q^2$ among the numbers $0,\dots,\ell-1$ such that $i_{\sigma_{i+1}}(\tau_{2i+2},\eta_{2i+2}) \ne i_{\sigma_i}(\tau_{2i},\eta_{2i})$. Since

$$|i_{\sigma_{i+1}}(\tau_{2i+2},\eta_{2i+2}) - i_{\sigma_i}(\tau_{2i},\eta_{2i})| \le 2q^3$$

for all i, this completes the proof of the proposition. □

We call the sequence $\{\tau_j\}_{0\le j\le 2\ell}$ constructed in Proposition 3.6 from a recurrent splitting sequence $\{\sigma_i\}_{0\le i\le\ell}$, a complete extension τ of σ_0 and a complete τ-extension of a σ_ℓ-filling measured geodesic lamination a sequence *induced* by $\{\sigma_i\}$.

If σ_0 is an arbitrary (not necessarily recurrent) subtrack of a complete train track τ_0 and if $\{\sigma_i\}_{0\le i\le\ell}$ is a splitting sequence issuing from σ_0, then the definition of a sequence of complete train tracks $\{\tau_i\}_{0\le i\le 2\ell}$ induced by the splitting sequence $\{\sigma_i\}_{0\le i\le\ell}$ always makes sense. However, if the splitting sequence $\{\sigma_i\}$ is not recurrent then such an induced sequence of complete train tracks may not exist.

Given any train track σ and a complementary component C of σ which is a topological disc, the number of equivalence classes of graphs $\tau\cap C$ where C is a complete extension of σ is bounded from above by a universal constant. By invariance under the action of the mapping class group, this implies the following. There exists a number $k>0$ such that if σ is a recurrent train track whose complementary components are all topological discs or once punctured topological discs, and if τ,η are complete extensions of σ, then $d(\tau,\eta)\le k$.

Furthermore, splitting a complete train track τ along a subtrack σ does not change the intersection of τ with the complement of a neighborhood of σ in S. Moreover, it commutes with the action of the pure mapping class group of a bordered subsurface of S which is contained in a complementary component of σ. As twisting about smooth boundary components of complementary regions of σ is controlled via marked points along an induced splitting sequence, together we obtain the following consequence.

Corollary 3.7. *For every $R > 0$ there is a number $p(R) > 0$ with the following property. Let $\{\sigma_i\}_{0\leq i\leq \ell}$ be a recurrent splitting sequence of train tracks on S. Let τ, η be complete extensions of σ_0 with $d(\tau,\eta) \leq R$ and let τ', η' be the endpoints of a sequence induced by $\{\sigma_i\}$ and issuing from τ, η. Then*

$$d(\tau', \eta') \leq p(R).$$

For a simple geodesic multi-curve c and a train track τ which carries c we denote by $\tau(c) \subset \tau$ the subgraph of τ of all branches which are contained in the image of c under a carrying map. Note that $\tau(c)$ is a recurrent subtrack of τ and hence either it is a disjoint union of simple closed curves which define the multi-curve c, or it contains a large branch (Lemma 2.7.2 of [PH92]).

Now let more specifically Q be a pants decomposition of S. If C is any complementary component of $\tau(Q)$, then a simple closed curve c contained in C is disjoint from Q. Thus if c is neither contractible nor freely homotopic into a puncture of S then c is freely homotopic to a component of the pants decomposition Q. In particular, the Euler characteristic of the completion of a complementary component of $\tau(Q)$ is at least -1. Thus this completion is of one of the following seven types, where in our terminology, a pair of pants can be a twice punctured disc or a once punctured annulus or a planar compact bordered surface of Euler characteristic -1 with three boundary components.

(1) A *triangle*, that is, a disc with three cusps at the boundary.
(2) A *quadrangle*, that is, a disc with four cusps at the boundary.
(3) A punctured disc with one cusp at the boundary.
(4) A punctured disc with two cusps at the boundary.
(5) An annulus with one cusp at the boundary.
(6) An annulus with two cusps at the boundary.
(7) A pair of pants with no cusps at the boundary.

If C is a complementary component of $\tau(Q)$ which is an annulus then the core curve of this annulus is freely homotopic to a component of Q. Since $\tau(Q)$ carries Q, this implies that if the boundary of C contains two cusps then these cusps are contained in the same boundary component, that is, one of the boundary components of C is a smooth circle. In other words, every complementary component of type (6) is of the more restricted following type.

(6′) An annulus with one smooth boundary component and one boundary component containing two cusps.

For the proof of Proposition 3.2 we need some control on $\tau(Q)$ where Q is any pants decomposition of S and where τ is a train track in standard form for a marking F of S which carries Q. The next lemma provides such a control. It follows immediately from the work of Penner and Harer [PH92] and uses an argument due to Thurston (see [FLP91]). We present the lemma in the form needed in Section 5.

Lemma 3.8. *Let F be a marking for S with pants decomposition P. Let ν be a measured geodesic lamination; then there is a train track σ with the following properties:*

(1) *σ carries ν.*
(2) *ν fills σ.*
(3) *Every train track in standard form for F which carries ν contains σ as a subtrack.*

Proof. We begin with showing the lemma in the case that ν is supported in a simple geodesic multi-curve.

Thus let F be a marking of S with pants decomposition P and let c be a simple geodesic multi-curve. Let S_0 be a connected component of $S - P$ with boundary circles $\gamma_i \in P$ (the number of these circles is contained in $\{1, 2, 3\}$). Up to homotopy, the multi-curve c intersects S_0 in a (perhaps empty) collection of disjoint simple arcs with endpoints on the boundary of S_0 which are *essential*, that is, not homotopic with fixed endpoints into the boundary of S_0.

For each i let $n(\gamma_i)$ be the *intersection number* between c and γ_i. Note that if γ_i is a component of c then $n(\gamma_i) = 0$. Since any two essential simple arcs in S_0 with endpoints on the same boundary component of S_0 are isotopic in S_0 relative to the boundary, there is up to isotopy a unique configuration of mutually disjoint simple arcs in S_0 with endpoints on the boundary of S_0 which realizes the intersection numbers $n(\gamma_i)$ (see [FLP91] for details). For this configuration there is a unique isotopy class of a train track (with stops) in S_0 which carries the configuration with a surjective carrying map and which can be obtained from a standard model as shown in Figure A by removing some (perhaps all) of the branches (Figure 2.6.2 of [PH92] shows in detail how to remove some of the branches of a standard model). These train tracks with stops can be glued to connectors obtained from the standard models shown in Figure B by removing some of the branches (see Figure 2.6.1 of [PH92]) in such a way that the resulting train track σ has the following properties.

(1) σ carries c.
(2) c fills σ.
(3) There is a train track in standard form for F which contains σ as a subtrack.

Note that the direction of the winding of a component of c relative to a curve from the marking F which intersects the pants curve γ_i determines the connector about γ_i in σ. It is immediate from the construction that a complete train track in standard form for F carrying c is an extension of σ (compare the discussion in [PH92]).

Now if ν is an arbitrary measured geodesic lamination then the support of ν can be approximated in the Hausdorff topology by a sequence $\{c_i\}$ of simple geodesic multi-curves [CEG87]. There are only finitely many train tracks which are subtracks of a train track in standard form for F. Thus if $\{\sigma_i\}$ is a sequence of train tracks as above for the multi-curves c_i then there is some train track σ so that $\sigma = \sigma_i$ for infinitely many i. Since the set of all geodesic laminations carried by the fixed train track σ is closed in the Hausdorff topology, σ carries ν and satisfies the requirements in the lemma. ☐

We use this to show

Lemma 3.9. *There is a number $\kappa > 0$ with the following properties. Let F be a marking for S and let Q by any pants decomposition of S. Then there is a set $\mathcal{D}$ of complete train tracks in standard form for some marking with pants decomposition Q (the marking may depend on the train track from $\mathcal{D}$) with the following properties.*

1. *Every geodesic lamination in standard form for Q is carried by some train track in the set $\mathcal{D}$.*
2. *The diameter of $\mathcal{D}$ in $\mathcal{TT}$ is at most κ.*
3. *For every $\eta \in \mathcal{D}$ there is a train track in standard form for F which carries η.*

Proof. By the discussion in the beginning of this section, it suffices to show the existence of a number $\chi > 0$ with the following properties. Let F be a marking of S and let Q by any pants decomposition of S. Then for every geodesic lamination λ in standard form for Q there is a train track $\tau(\lambda)$ with the following properties.

1. $\tau(\lambda)$ carries λ.
2. $\tau(\lambda)$ is in standard form for a marking with pants decomposition Q.
3. $\tau(\lambda)$ is carried by a train track $\tau_0(\lambda)$ in standard form for F.
4. If λ' is any other geodesic lamination in standard form for Q then we have $i_Q(\tau(\lambda), \tau(\lambda')) \leq \chi$.

Note that (4) above makes sense since by the definition of a train track in standard form for a marking with pants decomposition Q and by the fact that a geodesic lamination in standard form for Q contains Q as a sublamination, the pants decomposition Q is a subtrack of $\tau(\lambda)$.

Thus let F be a marking for S with pants decomposition P, let Q be a second pants decomposition and let λ, λ' be two geodesic laminations in standard form for Q. By Lemma 2.2, there are unique train tracks τ, τ' in standard form for F which carry λ, λ'; in particular, τ, τ' carry Q. Let as before $\tau(Q), \tau'(Q)$ be the subtrack of τ, τ' of all branches of positive Q-weight. By Lemma 3.8, the train tracks $\tau(Q), \tau'(Q)$ are isotopic. This means that τ, τ' are complete extensions of $\tau(Q)$. We equip the smooth boundary components of complementary regions of $\tau(Q)$ with marked points and use these marked points to define the intersection number between the complete extensions τ, τ' of $\tau(Q)$. Then $i_{\tau(Q)}(\tau, \tau')$ is bounded from above by a universal constant.

Assume first that $\tau(Q) = Q$, that is, that Q is a subtrack of τ. By invariance under the action of the mapping class group and cocompactness, Lemma 2.1 together with Lemma 3.3 of [H09] shows that there is a number $\beta > 0$, and there is a train track η in standard form for some marking of S with pants decomposition Q which carries λ, which is carried by τ and such that $d(\tau, \eta) \leq \beta$. Thus if we put $\tau(\lambda) = \eta$, then the distance in $\mathcal{TT}$ between the train track τ in standard form for F and the train track $\tau(\lambda)$ is uniformly bounded. As the same holds true for $\tau(\lambda')$ where λ' is another complete geodesic lamination on S in standard form for Q, the intersection number $i_Q(\tau(\lambda), \tau(\lambda'))$ is uniformly bounded as this number measures the twisting of $\tau(\lambda')$ about the components of Q relative to $\tau(\lambda)$. Therefore we may assume that $\tau(Q)$ contains a large branch.

Let $\{\sigma_i\}_{0 \leq i \leq s}$ be a splitting sequence issuing from $\sigma_0 = \tau(Q)$ so that for each $i \leq s$ the pants decomposition Q is carried by σ_i and fills σ_i. Then for each i the pants decomposition Q defines an integral transverse counting measure on σ_i by assigning to a branch b the number of connected components of the preimage of b under a carrying map $Q \to \sigma_i$. For $i < s$ the total Q-weight of σ_{i+1}, i.e. the sum of the weights of this counting measure over all branches of σ_{i+1}, is bounded from above by the total Q-weight of σ_i minus two. Namely, if σ_{i+1} is obtained from σ_i by a split at the large branch e and if e' is the diagonal branch of the split in σ_{i+1}, then the Q-weight of e equals the sum of the Q-weights of e' and the Q-weights of the two losing branches of the split. As by assumption that Q fills σ_i these weights are all positive and integral, the weight of e' does not exceed the weight of e minus two. The weights of the branches of σ_i which are distinct from e coincide with the weights of their images in σ_{i+1} under the natural bijection $\varphi(\sigma_i, \sigma_{i+1})$ of the branches of σ_i onto the branches of σ_{i+1} (which maps e to e'). Therefore the length of the splitting sequence $\{\sigma_i\}$ is bounded from above by the total Q-weight of $\tau(Q)$.

Assume that the sequence $\{\sigma_i\}_{0 \leq i \leq s}$ is of maximal length. This means that for every large branch e of σ_s, the pants decomposition Q is carried by a collision of σ_s at e (that is, a split followed by the removal of the diagonal).

By Proposition 3.6, there are complete extensions τ_1, τ_1' of σ_s with $\tau_1(Q) = \tau_1'(Q) = \sigma_s$ so that τ_1 carries λ, τ_1' carries λ' and that moreover

$$i_{\tau_1(Q)}(\tau_1, \tau_1') \leq i_{\tau(Q)}(\tau, \tau') + 4q^5. \tag{3}$$

Let e be any large branch of $\tau_1(Q) = \sigma_s$. The pants decomposition Q is carried by the train track ξ obtained from σ_s by the collision at e. Let τ_2 (or τ_2') be the $(\tau_1(Q), \lambda)$-modification (or the $(\tau_1'(Q), \lambda')$-modification) of τ_1 (or τ_1') at e. Two applications of Lemma 3.4 show that

$$i_{\tau_2(Q)}(\tau_2, \tau_2') \leq i_{\tau_1(Q)}(\tau_1, \tau_1') + 2q^3. \tag{4}$$

The train tracks τ_2, τ_2' are tight at e. Let τ_3, τ_3' be the train tracks obtained from τ_2, τ_2' by a split at e with the property that the split tracks carry λ, λ'. The number of branches of $\tau_3(Q) = \tau_3'(Q) = \xi$ is strictly smaller than the number of branches of $\tau_1(Q)$. The diagonal branch $d = \varphi(\tau_2, \tau_3)(e)$ of the split of τ_2 at e is a small branch of τ_3 which is contained in a complementary region C of $\xi = \tau_3(Q)$ and which is attached at both endpoints to a side of C. Let $d' = \varphi(\tau_2', \tau_3')(e)$ be the diagonal of the split of τ_2' at e.

Let U be a neighborhood of e in S which is sufficiently small that σ_s intersects U in the union of e with the four half-branches of σ_s which are incident on the endpoints of e. Up to isotopy, the intersections of σ_s, ξ with $S - U$ coincide. Thus if there is a complementary region C of ξ containing a smooth boundary component T which newly arises in the process, then this boundary component intersects U. Mark a point on T which is mapped into U. Note that any marked point on a smooth boundary component of a complementary region of σ_s induces a marked point on a smooth boundary component of a complementary region of ξ.

Let ζ, ζ' be complete extensions of σ_s which are σ_s-equivalent to τ_2, τ_2' and which have the minimal number of intersection points in $S - \sigma_s$, that is, which realize the intersection number $i_{\sigma_s}(\tau_2, \tau_2')$. By the definition of equivalence, after perhaps replacing ζ, ζ' by equivalent train tracks and after perhaps a modification with an isotopy, we may assume that ζ, ζ' are tight at e.

Using once more the definition of equivalence, the train tracks ζ, ζ' can be modified with a single split at e to train tracks ζ_0, ζ_0' which are complete extensions of ξ and which are equivalent to the complete extensions τ_3, τ_3' of ξ. If τ_3, τ_3' is obtained from τ_2, τ_2' by a right (or left) split at e then ζ_0, ζ_0' is obtained from ζ, ζ' by a right (or left) split at e.

Let d_0, d_0' be the diagonal of the split in ζ_0, ζ_0'. The train track ξ intersects U in two disjoint embedded arcs which are joined by the two branches d_0, d_0' of ζ_0, ζ_0'. We may assume that the branches d_0, d_0' either are disjoint (if the train tracks τ_3, τ_3' are both obtained from τ_2, τ_2' by the same type of split, left or right) or that they intersect transversely in a single point. By the definition of intersection numbers, this shows that

$$i_{\tau_3(Q)}(\tau_3, \tau_3') = i_\xi(\zeta_0, \zeta_0') \leq i_{\tau_2(Q)}(\tau_2, \tau_2') + 1. \tag{5}$$

Inequalities (5) and (4) now yield that

$$i_{\tau_3(Q)}(\tau_3, \tau_3') \leq i_{\tau_2(Q)}(\tau_2, \tau_2') + 1 \leq i_{\tau_1(Q)}(\tau_1, \tau_1') + 2q^3 + 1$$

and hence from the estimate (3) we obtain (since $q \geq 2$)

$$i_{\tau_3(Q)}(\tau_3, \tau_3') \leq i_{\tau(Q)}(\tau, \tau') + 6q^5.$$

Repeat this procedure with the train track $\xi = \tau_3(Q)$. After at most $k \leq q$ such steps where k is the number of branches of $\tau(Q)$, we arrive at train tracks η, η' which contain Q as a disjoint union of simple closed curves and carry λ, λ'.

To summarize, we obtain in at most q steps two splitting sequences connecting τ, τ' to train tracks η, η' so that $\eta(Q) = \eta'(Q) = Q$ and that η, η' carry λ, λ'. Each of these steps increases intersection numbers by at most $6q^5$. In particular, the intersection number $i_Q(\eta, \eta')$ is uniformly bounded and hence the distance in $\mathcal{TT}$ between η, η' is uniformly bounded as well.

Now λ, λ' is carried by η, η' and is in standard form for Q. Hence by the reasoning in the third paragraph of this proof, there are train tracks β, β' in standard form for some marking with pants decomposition Q which carry λ, λ', which are carried by η, η' and such that $d(\eta, \beta) \leq \kappa, d(\eta', \beta') \leq \kappa$. As a consequence, the distance between β, β' is uniformly bounded. Since λ, λ' were arbitrarily chosen geodesic laminations in standard form for Q the lemma follows. □

In the situation of Lemma 3.9, it may happen that the pants decomposition Q fills a train track τ in standard form for F, that is, the subtrack $\tau(Q)$ of Q coincides with τ. In this case each of the train tracks in standard form for Q is carried by τ. As in this case the proof of Lemma 3.9 does not use any further information on the marking F, we also obtain the following statement, where the constant $\kappa > 0$ is the constant from Lemma 3.9.

Corollary 3.10. *There is a number $\kappa > 0$ with the following properties. Let Q be a pants decomposition of S and let τ be a complete train track which carries all complete geodesic laminations λ in standard form for Q. Then there is a set $\mathcal{D}$ of complete train tracks in standard form for some marking with pants decomposition Q (where the marking may depend on the train track from $\mathcal{D}$) with the following properties.*

(1) *Every geodesic lamination in standard form for Q is carried by some train track in the set $\mathcal{D}$.*
(2) *The diameter of $\mathcal{D}$ in $\mathcal{TT}$ is at most κ.*
(3) *Every $\eta \in \mathcal{D}$ is carried by τ.*

Now we are ready for the proof of Proposition 3.2. For later use, we formulate it as a consequence of Lemma 3.9 and Corollary 3.10.

Lemma 3.11. *There exists a number $d_2 > 0$ with the following properties. Let G be a marking with pants decomposition Q. For a complete geodesic lamination λ in standard form for Q let $\tau(\lambda)$ be the train track constructed in Lemma 3.9 or in Corollary 3.10. Then there exists λ such that $\tau(\lambda)$ carries a train track ξ in the d_2-neighborhood of a train track in standard form for G. Furthermore, ξ is in standard form for a marking with pants decomposition Q.*

Proof. Let λ be a complete geodesic lamination in standard form for Q and let $\tau(\lambda) \in \mathcal{V}(\mathcal{TT})$ be as in Lemma 3.9 or in Corollary 3.10. Then $\tau(\lambda)$ is in standard form for a marking G' with pants decomposition Q. Any two markings with pants decomposition Q differ from each other by a multi-twist about the pants curves of Q. Thus if we write $k = 3g-3+m$ for simplicity of notation and if we let $\theta_1, \dots, \theta_k$ be the positive Dehn twists about the components $\gamma_1, \dots, \gamma_k$ of Q then there is an integral vector $(n_1, \dots, n_k) \in \mathbb{Z}^k$ such that

$$G = \theta_1^{n_1} \cdots \theta_k^{n_k} G'.$$

Every pants curve γ_i of Q is the core curve of a twist connector for $\tau(\lambda)$. Splitting a standard twist connector at the large branch, with the small branch of the connector as a winner, results in deforming a train track by a (positive or negative) Dehn twist about the core curve of the connector. The sign of the twist is determined by the type of the twist connector which in turn is determined by the spiraling direction of the geodesic lamination λ in standard form for Q about the pants curve γ_i.

Assume after reordering that for some $p \leq k$ and for every $i \leq p$, either $n_i = 0$ or the sign of n_i coincides with the sign determined by the spiraling direction of λ about γ_i, and that for $i > p$ we have $n_i \neq 0$ and the sign of n_i differs from this direction. Let σ be a train track in standard form for the marking G' which is obtained from $\tau(\lambda)$ by reversing the directions in the twist connectors about the curves $\gamma_{p+1}, \dots, \gamma_k$. By Lemma 2.6.1 of [PH92], σ is complete, and σ is splittable to the train track $\theta_1^{n_1} \cdots \theta_k^{n_k}\sigma$ in standard form for G. The train track σ carries a complete geodesic lamination λ' in standard form for Q. By equivariance under the action of the mapping class group and cocompactness, there is a universal constant $\chi > 0$ such that

$$d(\sigma, \tau(\lambda)) \leq \chi.$$

By Lemma 3.9 or Corollary 3.10 there is a train track $\tau(\lambda')$ which is in standard form for a marking with pants decomposition Q, which carries λ' and such that

$$d(\tau(\lambda), \tau(\lambda')) \leq \kappa.$$

Since $\mathcal{MCG}(S)$ acts isometrically on $\mathcal{TT}$, we have

$$d(\theta_1^{n_1} \cdots \theta_k^{n_k}\sigma,\, \theta_1^{n_1} \cdots \theta_k^{n_k}\tau(\lambda')) \leq \kappa + \chi.$$

But $\theta_1^{n_1} \cdots \theta_k^{n_k}\sigma$ is in standard form for G and therefore the train track $\eta = \theta_1^{n_1} \cdots \theta_k^{n_k}\tau(\lambda')$ is at distance at most $\kappa + \chi$ from a train track in standard form for G. It contains the pants decomposition Q as an embedded subtrack. This is what we wanted to show. □

Proposition 3.2 is an immediate consequence of Lemma 3.11. Namely, given any two markings F, G, by Lemma 3.9 the train track $\tau(\lambda)$ found in Lemma 3.11 is carried by some train track η in standard form for F and hence η carries a train track σ in the d_2-neighborhood of a train track in standard form for G. This is precisely what is claimed in Proposition 3.2. The same reasoning using Corollary 3.10 then yields the following

Corollary 3.12. *There exists a number $d_3 > 0$ with the following property. Let Q be a pants decomposition of S and let τ be a complete train track which carries Q and such that Q fills τ. Then for any marking G of S with pants decomposition Q, the train track τ is splittable to a train track in the d_3-neighborhood of a train track in standard form for G.*

Remark 3.13. The results in this section are also valid if the surface S is a once punctured torus or a four punctured sphere.

4. Quasi-isometric embeddings

In this section we use the results from Section 3 to give a unified proof of Theorems 5-7 from the introduction.

We begin with an investigation of the mapping class group of an essential subsurface S_0 of S. This means that S_0 is a bordered subsurface of S with the property that the inclusion $S_0 \to S$ induces an injection $\pi_1(S_0) \to \pi_1(S)$ of fundamental groups and that moreover every boundary component of S_0 is an essential simple closed curve in S. Let $\mathcal{PMCG}(S_0)$ be the *pure mapping class group* of S_0, i.e. the subgroup of the mapping class group $\mathcal{MCG}(S_0)$ of S_0 of all mapping classes which fix each of the boundary components and each of the punctures. Then $\mathcal{PMCG}(S_0)$ is a subgroup of $\mathcal{MCG}(S_0)$ of finite index. It can be identified with the subgroup of the mapping class group of S of all elements which can be realized by a homeomorphism preserving $S - S_0$ pointwise and fixing each of the punctures of S.

As in the introduction, call a finitely generated subgroup Γ of $\mathcal{MCG}(S)$ *undistorted* if the inclusion $\Gamma \to \mathcal{MCG}(S)$ is a quasi-isometric embedding. For example, every subgroup of $\mathcal{MCG}(S)$ of finite index is undistorted. The following result is implicitly but not explicitly contained in Theorem 6.12 of [MM00].

Proposition 4.1. *For an essential subsurface $S_0 \subset S$ the subgroup $\mathcal{PMCG}(S_0)$ of $\mathcal{MCG}(S)$ is undistorted.*

Proof. If $S_0 = S_1 \cup S_2$ for two disjoint essential subsurfaces S_1, S_2 of S whose fundamental groups as subgroups of $\pi_1(S)$ have trivial intersection then

$$\mathcal{PMCG}(S_0) = \mathcal{PMCG}(S_1) \times \mathcal{PMCG}(S_2).$$

Now a subgroup of a finitely generated group which is a direct product of two undistorted subgroups is undistorted and hence it suffices to show the proposition for connected essential subsurfaces of S. The case that S_0 is an essential annulus is treated in detail in [FLM01, H09], so we assume that the Euler characteristic of S_0 is negative. If S_0 is a thrice punctured sphere then $\mathcal{PMCG}(S_0)$ equals the free abelian group of Dehn twists about the boundary components of S_0. Thus we also may assume that S_0 is different from a thrice punctured sphere.

Our goal is to show that any two elements of $\mathcal{PMCG}(S_0)$ can be connected by a uniform quasi-geodesic in $\mathcal{MCG}(S)$ which is entirely contained in $\mathcal{PMCG}(S_0)$. For this let $\hat{S}_0$ be the surface which we obtain from S_0 by replacing each boundary component by a puncture. There is an exact sequence

$$0 \to \mathbb{Z}^p \to \mathcal{PMCG}(S_0) \stackrel{\Pi}{\to} \mathcal{PMCG}(\hat{S}_0) \to 0$$

where $\mathbb{Z}^p$ is identified with the free abelian group of Dehn twists about the boundary components of S_0.

Choose a pants decomposition P for S which contains the boundary of S_0 as a subset. Let τ be a complete train track in standard form for a marking F with pants decomposition P and only twist connectors. Let τ_1 be the subtrack of τ which we obtain from τ by removing all branches contained in the interior of S_0. We can choose τ in such a way that any two points in the same connected component of τ_1 can be connected by a trainpath in τ_1 (however, in general τ_1 is neitiher connected not recurrent).

Let $c_1, \dots, c_p$ be the boundary circles of S_0. Every complete train track σ on $\hat{S}_0$ is a subtrack of a complete train track η on S which contains τ_1 as a subtrack. Namely, up to isotopy, each boundary component c_i of S_0 is contained in a complementary once punctured monogon region C_i of σ. It cuts C_i into an annulus $A_i \subset S_0$ and a once punctured disc. Add a small branch b_i to $\sigma \cup \tau_1$ which is contained in the closure $\overline{A_i}$ of the annulus A_i and connects the boundary of C_i to the boundary circle c_i of S_0. Since σ is complete and hence non-orientable, if for each $i \leq p$ we connect the branch b_i to the circle c_i in such a way that the resulting train track η intersects an annulus neighborhood of c_i in a twist connector as shown in Figure B, then η is recurrent [PH92]. The train track η is also very easily seen to be transversely recurrent and hence it is complete.

The train track η is not uniquely determined by τ_1 and σ. The choices made are the positions of the additional switches on the boundaries of the complementary regions C_i, the inward pointing tangents of the added branches b_i at these switches and the homotopy class with fixed endpoints of the branch $b_i \subset \overline{A_i}$. By invariance under the action of the group $\mathcal{PMCG}(\hat{S}_0)$ and cocompactness, for any two such choices η, η' there is a multi-twist φ about the multi-curve $c = \cup_i c_i$ such that the distance in $\mathcal{TT}$ between $\eta, \varphi\eta'$ is uniformly bounded. The set $\mathcal{E}$ of all such extensions of all complete train tracks on $\hat{S}_0$ is invariant under the action of the group $\mathcal{PMCG}(S_0)$, with finitely many orbits and finite point stabilizers.

Let F be a marking for $\hat{S}_0$. As by Corollary 3 of [H09] splitting sequences are uniform quasi-geodesics, we deduce from Proposition 3.2 and the following remark that any complete train track η on $\hat{S}_0$ can be obtained from a train track σ in standard form for F by a uniform quasi-geodesic in the train track complex $\mathcal{TT}(\hat{S}_0)$ of $\hat{S}_0$ which is a concatenation of a splitting sequence with an edge-path of uniformly bounded length.

For every train track σ on $\hat{S}_0$ in standard form for F choose an extension $\Psi(\sigma) \in \mathcal{V}(\mathcal{TT})$ as above. By Proposition 3.6, there is a universal number $p > 0$ (depending on the topological type of S) and for every splitting sequence $\{\sigma_i\}_{0\leq i\leq m}$ of complete train tracks on $\hat{S}_0$ issuing from a train track $\sigma_0 = \sigma$ in standard form for F and for every complete $\Psi(\sigma)$-extension of a σ_m-filling measured geodesic lamination there is an induced sequence $\{\tau_j\}_{0\leq j\leq 2m} \subset \mathcal{TT}$ connecting $\tau_0 = \Psi(\sigma)$ to a train track τ_{2m} which contains σ_m as well as τ_1 as a subtrack. There is a universal number $s > 0$ such that for every $i < m$ the train track track τ_{2i+2} can be obtained from τ_{2i} by a splitting sequence whose length is contained $[1, s]$.

By Corollary 3 of [H09], splitting sequences are uniform quasi-geodesics in both $\mathcal{TT}$ and the train track complex $\mathcal{TT}(\hat{S}_0)$ of $\hat{S}_0$. As a consequence, there is a number $c > 0$ with the following property. For every train track $\xi \in \mathcal{V}(\mathcal{TT}(\hat{S}_0))$ there is a train track $\Psi(\xi) \in \mathcal{E}$ which contains both ξ and τ_1 as a subtrack and is such that the distance in $\mathcal{TT}(\hat{S}_0)$ between ξ and a train track σ in standard form for F is not bigger than $cd(\Psi(\sigma), \Psi(\xi)) + c$.

The resulting map $\Psi : \mathcal{V}(\mathcal{TT}(\hat{S}_0)) \to \mathcal{E}$ is used to define a map $\rho : \mathcal{PMCG}(\hat{S}_0) \to \mathcal{PMCG}(S_0)$ as follows. Let σ be a fixed train track in standard form for F. For $g \in \mathcal{PMCG}(\hat{S}_0)$ define $\rho(g) \in \mathcal{PMCG}(S_0)$ in such a way that the distance between $\Psi(g\sigma)$ and $\rho(g)(\Psi(\sigma))$ is uniformly bounded. Since $\mathcal{PMCG}(S_0)$ acts on $\mathcal{E}$ with finitely many orbits and finite point stabilizers and since $\mathcal{TT}(\hat{S}_0)$ is equivariantly quasi-isometric to $\mathcal{PMCG}(\hat{S}_0)$, the map ρ is a coarse section of the projection Π.

By this we mean that there is a universal constant $\kappa > 0$ such that $d(\Pi\rho(g), g) \le \kappa$ for all g.

The image of $\mathcal{V}(\mathcal{TT}(\hat{S}_0))$ under the map Ψ consists of train tracks which contain each boundary component c_i of S_0 as the core curve of a twist connector. Splitting such a train track τ at the large branch in this twist connector, with the small branch as the winner, results in replacing τ by $\theta_c(\tau)$ where θ_c is a Dehn twist about c whose direction (positive or negative) depends on the twist connector. Thus if Γ denotes the semi-group of Dehn twists about the boundary components of S_0 determined by the train track τ_1 then for every $g \in \rho(\mathcal{PMCG}(\hat{S}_0))$ and every $\varphi \in \Gamma$ there is a uniform quasi-geodesic in $\mathcal{MCG}(S)$ connecting the identity to $\varphi\rho(g)$ and which is entirely contained in $\mathcal{PMCG}(S_0)$. However, the choice of the twist connector in the train track τ_1 was arbitrary and consequently the unit element in $\mathcal{MCG}(S)$ can be connected to any mapping class $g \in \mathcal{PMCG}(S_0)$ by a uniform quasi-geodesic in $\mathcal{MCG}(S)$ which is entirely contained in $\mathcal{PMCG}(S_0)$. By invariance of the word metrics under left translation, this just means that $\mathcal{PMCG}(S_0) < \mathcal{MCG}(S)$ is undistorted. □

Now let S_0 be any non-exceptional surface of genus $g \ge 0$ with $m \ge 0$ punctures and let S be the surface S_0 punctured at one additional point p. There is an exact sequence [B74]

$$0 \to \pi_1(S_0) \to \mathcal{MCG}(S) \xrightarrow{\Pi} \mathcal{MCG}(S_0) \to 0.$$

The projection Π is induced by the map $S \to S_0$ which consists in closing the puncture p. An element α of the fundamental group $\pi_1(S_0)$ of S_0 is mapped to the element of $\mathcal{MCG}(S)$ obtained by dragging the point p along a loop in S_0 in the homotopy class α. Braddeus, Farb and Putman [BFP11] showed that $\pi_1(S_0)$ is an exponentially distorted subgroup of $\mathcal{MCG}(S)$. We next observe that in contrast, the projection $\Pi : \mathcal{MCG}(S) \to \mathcal{MCG}(S_0)$ has a coarse section which is a quasi-isometric embedding. Here by a coarse section we mean a map $\Psi : \mathcal{MCG}(S_0) \to \mathcal{MCG}(S)$ such that

$$d(\Pi\Psi(g), g) \le \kappa$$

for all $g \in \mathcal{MCG}(S)$ where $\kappa \ge 0$ is a universal constant. The following proposition is a special case of a general result of Mosher [M96].

Proposition 4.2. *Let S_0 be a non-exceptional surface of genus $g \ge 0$ with $m \ge 0$ punctures and let S be the surface of genus g with $m+1$ punctures. Then there is a coarse section for the projection $\Pi : \mathcal{MCG}(S) \to \mathcal{MCG}(S_0)$ which is a quasi-isometric embedding.*

Proof. Let $\mathcal{TT}(S)$ and $\mathcal{TT}(S_0)$ be the train track complex of S and of S_0. We first define a map $\Psi : \mathcal{V}(\mathcal{TT}(S_0)) \to \mathcal{V}(\mathcal{TT}(S))$ as follows.

For a complete train track τ on S_0 choose any complementary trigon C of τ. Mark a point p in the interior of C and add two switches v_1, v_2 and two branches $b_1, b_2 \subset C - \{p\}$ to τ in the following way. The switch v_1 is an interior point of a branch of τ contained in a side of C, $v_2 \neq p$ is a point in the interior of C, b_1 connects v_1 to v_2 and b_2 is a small branch contained in the interior of C whose endpoints are both incident on v_2 and which is the boundary of a subdisc of C containing p in its interior. Since τ is complete, Proposition 1.3.7 of [PH92] shows that the resulting train track η_0 on $S = S_0 - \{p\}$ is recurrent. It is also easily seen to be transversely recurrent. The train track η_0 decomposes S into trigons, once punctured monogons and one fourgon. The fourgon can be subdivided into two trigons by adding a single small branch. The resulting train track η on S is complete, and it contains τ as a subtrack. This construction defines a map $\Psi : \mathcal{V}(\mathcal{TT}(S_0)) \to \mathcal{V}(\mathcal{TT}(S))$.

The map Ψ depends on some choices among a finite set of possibilities: The choice of the complementary trigon C, the choice of the position of the switch v_1 on a side of C, the orientation of the inward pointing tangent of the branch b_1 at the switch v_1 and the choice of the small branch subdividing the fourgon. Any train track constructed in this way contains τ as a subtrack. Moreover, for $g \in \mathcal{MCG}(S_0)$ there is some $h \in \mathcal{MCG}(S)$ such that $h(\Psi(\tau))$ is one of the possibilities for $\Psi(g\tau)$. Since there are only finitely many orbits of complete train tracks on S_0 under the action of the mapping class group, by coarse equivariance of the construction we conclude that there is a universal number $\kappa_0 > 0$ such that for any other choice Ψ' of such a map we have $d(\Psi(\tau), \Psi'(\tau)) \leq \kappa_0$ for all $\tau \in \mathcal{V}(\mathcal{TT}(S_0))$.

We use the map Ψ to define a map $\Phi : \mathcal{MCG}(S_0) \to \mathcal{MCG}(S)$ as follows. The mapping class groups of S_0, S act properly and cocompactly on $\mathcal{TT}(S_0)$, $\mathcal{TT}(S)$. Choose $\tau \in \mathcal{V}(\mathcal{TT}(S_0))$ and a fundamental domain D for the action of $\mathcal{MCG}(S)$ on $\mathcal{TT}(S)$ containing $\Psi(\tau)$. For $g \in \mathcal{MCG}(S_0)$ choose $\Phi(g) \in \mathcal{MCG}(S)$ in such a way that $\Psi(g\tau) \in \Phi(g)D$. If Φ' is any other such map then $d(\Phi(g), \Phi'(g)) \leq \kappa_1$ where $\kappa_1 > 0$ is a universal constant (and d is any distance on $\mathcal{MCG}(S)$ defined by a word norm of a finite symmetric generating set).

By construction, the map $\Phi : \mathcal{MCG}(S_0) \to \mathcal{MCG}(S)$ is a coarse section for the projection $\mathcal{MCG}(S) \to \mathcal{MCG}(S_0)$. Thus we are left with showing that Φ is a quasi-isometric embedding, and this holds true if this is the case for the map Ψ. To this end, note that τ is a subtrack of $\Psi(\tau)$. By Proposition 3.6, a splitting sequence $\{\tau_i\}_{0 \leq i \leq \ell} \subset \mathcal{TT}(S_0)$ issuing from $\tau_0 = \tau$ induces a splitting sequence in $\mathcal{TT}(S)$ issuing from $\Psi(\tau)$. The length of this sequence is not smaller than the length ℓ of the splitting sequence of τ, and it is not bigger than $q\ell$ for a universal constant $q > 0$. On the other hand, a point on the induced sequence which contains τ_i as a subtrack is a possible choice for $\Psi(\tau_i)$ and hence it is at uniformly bounded distance to $\Psi(\tau_i)$.

By Corollary 3 of [H09], splitting sequences in $\mathcal{TT}(S)$ and $\mathcal{TT}(S_0)$ are uniform quasi-geodesics. As a consequence, there is a number $c > 1$ such that

$$d(\tau_0, \tau_\ell)/c - c \leq d(\Psi(\tau_0), \Psi(\tau_\ell)) \leq cd(\tau_0, \tau_\ell) + c$$

whenever $\tau_0 \in \mathcal{V}(\mathcal{TT}(S_0))$ is splittable to $\tau_\ell \in \mathcal{V}(\mathcal{TT}(S_0))$. By Proposition 3.1, splitting sequences connect a coarsely dense set of pairs of points in the train track complex $\mathcal{TT}(S_0)$. This implies that the map $\Psi : \mathcal{V}(\mathcal{TT}(S_0)) \to \mathcal{V}(\mathcal{TT}(S))$ defines a quasi-isometric embedding and hence the same holds true for $\Phi : \mathcal{MCG}(S_0) \to \mathcal{MCG}(S)$. □

Finally, for a closed surface of genus $g \geq 2$ we investigate the normalizer of a finite subgroup Γ of $\mathcal{MCG}(S)$ (see [RS07] for an earlier proof of this result, stated a bit differently).

Proposition 4.3. *For a closed surface S of genus $g \geq 2$, the normalizer of a finite subgroup of $\mathcal{MCG}(S)$ is undistorted.*

Proof. Let $\mathcal{T}(S)$ be the Teichmüller space of S. By the Nielsen realization problem, a finite subgroup Γ of $\mathcal{MCG}(S)$ fixes a point $x \in \mathcal{T}(S)$ [Ke83]. This means that Γ can be realized as a finite group of biholomorphisms of (S, x). The quotient $(S, x)/\Gamma$ is a Riemann surface, and the projection $\pi : (S, x) \to (S, x)/\Gamma$ is a branched covering ramified over a finite number of points $p_1, \dots, p_\ell \in (S, x)/\Gamma$. The marked complex structure x on S projects to a marked complex structure on $(S, x)/\Gamma$.

Let $(S_1, x_1), (S_0, x_0)$ be the punctured Riemann surfaces which are obtained from $(S, x), (S, x)/\Gamma$ by removing the branch points of the covering $(S, x) \to (S, x)/\Gamma$. The projection π restricts to an unbranched covering $S_1 \to S_0$. The Teichmüller spaces $\mathcal{T}(S_0)$ of S_0, $\mathcal{T}(S_1)$ of S_1 are contractible. For every point $y \in \mathcal{T}(S_0)$ which is sufficiently close to x_0 there is a covering $\Psi(y) \in \mathcal{T}(S_1)$ of the Riemann surface (S_0, y) which is of the same topological type as the covering $(S_1, x_1) \to (S_0, x_0)$. The marking of $\Psi(y)$ is determined in such a way that the map Ψ is continuous near x_0. This construction defines a developing map $\Psi : \mathcal{T}(S_0) \to \mathcal{T}(S_1)$. Since $\mathcal{T}(S_0)$ is simply connected, the developing map is in fact single-valued. Moreover, it is clearly injective and hence an embedding. (In fact, its is not hard to see that this construction defines an isometric embedding of $\mathcal{T}(S_0)$ into $\mathcal{T}(S_1)$ for the Teichmüller metrics). There is a natural projection $\Pi : \mathcal{T}(S_1) \to \mathcal{T}(S)$ defined by filling in the punctures.

Let $\mathcal{MCG}_0(S_0)$ be the subgroup of the mapping class group $\mathcal{MCG}(S_0)$ of S_0 of all mapping classes realizable by a homeomorphism of S_0 which lifts to a homeomorphism of S. Let $N(\Gamma)$ be the normalizer of Γ in $\mathcal{MCG}(S)$. Then there is an exact sequence [BH73]

$$0 \to \Gamma \to N(\Gamma) \to \mathcal{MCG}_0(S_0) \to 0.$$

(Theorem 3 in [BH73] states this only in the case that the group Γ is cyclic. However, as pointed out explicitly in [BH73], the result for all finite groups is immediate from the argument given there and the Nielsen realization problem).

Since the group Γ is finite, the groups $N(\Gamma)$ and $\mathcal{MCG}_0(S_0)$ are quasi-isometric. Thus to show the proposition it is enough to show that there is quasi-isometric embedding of $\mathcal{MCG}_0(S_0)$ into $\mathcal{MCG}(S)$ whose image is contained in a uniformly boundedd neighborhood of $N(\Gamma)$. Following Proposition 4.2, it suffices in fact to show that there is a quasi-isometric embedding of $\mathcal{MCG}(S_0)$ into $\mathcal{MCG}(S_1)$ whose image is contained in a uniformly bounded neighborhood of the image of $N(\Gamma)$ under a a coarse section for the projection $\mathcal{MCG}(S_1) \to \mathcal{MCG}(S)$.

For this note that the preimage of a complete train track τ on S_0 under the covering $S_1 \to S_0$ is a Γ-invariant graph ξ in S which decomposes S into polygons and once punctured polygons. The preimage of each trigon component of τ is a union of n trigon components of ξ where $n = |\Gamma|$ is the number of sheets of the covering. Each once punctured monogon in τ encloses one of the points p_i and lifts to a punctured m_i-gon in $S_1 - \xi$ where $2 \leq m_i \leq n$ is the ramification index of p_i.

The branch points of the covering define a set of marked points contained in complementary regions of ξ. Each complementary region of ξ contains at most one such point. Thus ξ defines a (non-complete) train track on the punctured surface S_1, again denoted by ξ. Now a positive transverse measure on τ lifts to a positive transverse measure on ξ and therefore ξ is recurrent. The same argument also shows that ξ is transversely recurrent. Then ξ is a subtrack of a complete train track on S_1 obtained by subdividing some of the complementary regions as in the proof of Proposition 4.2. As before, the resulting complete train track η depends on choices among a uniformly bounded number of possibilities (compare the proof of Proposition 4.2).

Now if the complete train track τ_1 on S_0 is obtained from the complete train track τ by a single split at a large branch e then the preimage ξ_1 of τ_1 can be obtained from the preimage ξ of τ by a splitting sequence of length n. Namely, the preimage of any large branch of τ is the union of n large branches of ξ. Such a splitting sequence then induces a splitting sequence of length at most qn of the complete train track η on S_1 constructed in the previous paragraph where $q > 0$ is a universal constant.

By Proposition 3.1, splitting sequences in the train track complex $\mathcal{TT}(S_0)$ of S_0 connnect a coarsely dense set of pairs of points. By Corollary 3 of [H09], each such splitting sequence defines a uniform quasi-geodesic in the subgroup $\mathcal{MCG}_0(S_0)$ of the mapping class group of S_0. This quasi-geodesic lifts to a uniform quasi-geodesic in $\mathcal{MCG}(S_1)$ contained in a uniformly bounded neighborhood of the image of $N(\Gamma)$ under the coarse section for the projection $\mathcal{MCG}(S_1) \to \mathcal{MCG}(S)$ constructed in Proposition 4.2. As a consequence, the normalizer $N(\Gamma)$ of Γ is undistorted. □

Remark: 1) Since splitting sequences define quasi-geodesics in the *curve graph* of a surface of finite type [H06], the above argument immediately implies the following. Let S be a closed surface and let Γ be a finite subgroup of $\mathcal{MCG}(S)$. Then there is a quasi-isometric embedding of the curve graph of S/Γ into the curve graph of S. This was shown in [RS07].

2) In [ALS09], Aramayona, Leininger and Souto constructed for infinitely many $g_i > 0$ injective homomorphisms of the mapping class group of a closed surface of genus g_i into the mapping class group of a closed surface of strictly bigger genus using unbranched coverings. The reasoning in the proof of Proposition 4.3 can be used to show that these homomorphisms are quasi-isometric embeddings.

5. A CAT(0) cube complex related to the mapping class group

This section is devoted to the proof of Theorem 1 from the introduction. Our goal is to construct a CAT(0) cube complex from a subgraph of the train track complex $\mathcal{TT}$.

A subgraph G_0 of a metric graph G is called *coarsely dense* if there exists a number $\kappa > 0$ such that any point in G is of distance at most κ from a point in G_0. We show

Theorem 5.1. *There exists a $\mathcal{MCG}$-invariant family $\mathcal{G}_i$ of complete connected subgraphs of $\mathcal{TT}$ with the following properties.*

(1) *The subgraphs $\mathcal{G}_i$ are mutually isometric, and they are isometric to the one-skeleton of a* CAT(0) *cube complex.*
(2) *The image of each of the subgraphs under the inclusion $\mathcal{G}_i \to \mathcal{TT}$ is coarsely dense in $\mathcal{TT}$.*

Theorem 1 follows from Theorem 5.1 and the fact that the train track complex $\mathcal{TT}$ is connected, and the mapping class group $\mathcal{MCG}$ of S acts on $\mathcal{TT}$ properly and cocompactly [H09].

We begin the discussion with looking at subcomplexes of the cube complex obtained by subdividing $\mathbb{R}^n$ into cubes with integral vertices. Namely, let $e_1, \dots, e_n$ be the standard basis of $\mathbb{R}^n$. We call the set $\{\pm e_i \mid i \leq n\}$ the set of *standard basis vectors*. If we talk about a standard basis vector in the sequel, then we mean a point in this set.

Define a *cubical graph* in $\mathbb{R}^n$ to be an embedded connected directed graph G whose vertices are points with integer coordinates (that is, points contained in $\mathbb{Z}^n$) and whose edges are oriented line segments of length one connecting a vertex v_1 to a vertex v_2 with $v_2 - v_1 = e_i$ for some $i \leq n$. The *full grid* $\mathcal{G}$ is the maximal cubical graph in $\mathbb{R}^n$ with vertex set $\mathbb{Z}^n$. It is the one-skeleton of a CAT(0) cube complex which is isometric to $\mathbb{R}^n$.

Call an embedded cube $C \subset \mathbb{R}^n$ of dimension $0 \leq k \leq n$ *standard* if it is isometric to $[0,1]^k$ and if its vertices are contained in $\mathbb{Z}^n$; then its one-skeleton consists of arcs of length one parallel to the axes. A cubical graph G in $\mathbb{R}^n$ is naturally the one-skeleton of a cube complex $\mathcal{C}(G)$, called the *span* of G, as follows. A standard cube C of dimension $k \geq 2$ is contained in $\mathcal{C}(G)$ if the one-skeleton of C is entirely contained in G.

The next lemma establishes a criterion for a special class of cubical graphs G which is necessary and sufficient for their span $\mathcal{C}(G)$ to be nonpositively curved. To define this class we say that a cubical graph $G \subset \mathbb{R}^n$ is *locally determined by its* 0*-skeleton* if whenever x is a vertex of G and a is a standard basis vector of $\mathbb{R}^n$ such that $x + a \in G$, then the edge connecting x to $x + a$ is contained in G. The cubical graph G *satisfies the 3-cube condition* if the following holds true. Let x be a vertex of G, let $a_{i_1}, a_{i_2}, a_{i_3}$ be linearly independent standard basis vectors and assume that $x, x + a_{i_j}, x + a_{i_j} + a_{i_k} \in G$ for $k \neq j \in \{1,2,3\}$; then $x + \sum_j a_{i_j} \in G$. We have

Lemma 5.2. *Let $G \subset \mathbb{R}^n$ be the cubical graph which is locally determined by its* 0*-skeleton. Then the span $\mathcal{C}(G)$ of G is nonpositively curved if and only if G satisfies the 3-cube condition.*

Proof. By a result of Gromov (Theorem 5.18 of [BH99]), $\mathcal{C}(G)$ is nonpositively curved if and only if for every vertex $x \in G \subset \mathcal{C}(G)$, the link complex of x is a flag complex. This means that if there are k vertices in the link complex which are pairwise connected by an edge, then these vertices span a simplex in the link complex of dimension $k-1$.

By the definition of the span of a cubical graph $G \subset \mathbb{R}^n$ which is locally determined by its 0-skeleton, for each vertex $x \in G$, an edge in G incident on x is determined by a standard basis vector a of $\mathbb{R}^n$ with the property that $x + a \in G$. The vertices in the link complex of $\mathcal{C}(G)$ at the vertex x are precisely the edges in G defined in this way.

Assume that the link complex of $\mathcal{C}(G)$ is a flag complex. By the definition of the span $\mathcal{C}(G)$ of G, edges in the link complex of x are defined by pairs $(x, a_i), (x, a_j)$ where a_i, a_j are linearly independent standard basis vectors and such that $x, x + a_i, x + a_j, x + a_i + a_j \in G$. Now let $x, x + a_{i_1}, x + a_{i_2}, x + a_{i_3} \in G$ for linearly independent standard basis vectors a_{i_j}. Assume furthermore that $x + a_{i_u} + a_{i_s} \in G$ for any pair of distinct points $u \neq s \in \{1, 2, 3\}$. This implies that the one-skeletons of the three faces of a 3-cube $C \subset \mathbb{R}^n$ which has x as one of its vertices are contained in G. Since the link complex of $\mathcal{C}(G)$ is a flag complex by assumption, the 3-cube C containing these three faces in its boundary is contained in $\mathcal{C}(G)$. But the vertex of C opposite to x is the vertex $x + a_{i_1} + a_{i_2} + a_{i_3}$ and hence this vertex is contained in G. This shows that G satisfies the 3-cube condition.

Now let us assume that the cubical graph G which is locally determined by its 0-skeleton satisfies the 3-cube condition. Let $x \in G$ and assume that for a collection $\mathcal{A} = \{a_{i_1}, \dots, a_{i_k}\}$ of $k \geq 3$ linearly independent standard basis vectors and any pair i_j, i_ℓ the vertices $x + a_{i_j}, x + a_{i_\ell}, x + a_{i_j} + a_{i_\ell}$ are all contained in G. Since G is locally determined by its 0-skeleton, this implies that $\mathcal{C}(G)$ contains the square spanned by these vertices. Thus the link complex of x contains the one-skeleton a $(k-1)$-simplex which is the span of the vertices defined by the edges in G connecting x to $x + a_{i_j}$.

As G satisfies the 3-cube condition, for any three distinct of the vectors a_{i_j}, say the vectors $a_{i_j}, a_{i_\ell}, a_{i_s}$, the point $x + a_{i_j} + a_{i_\ell} + a_{i_s}$ is a vertex of G. As G is locally determined by its 0-skeleton, G contains the 1-skeleton of the cube with these vertices and hence $\mathcal{C}(G)$ contains the 3-cube spanned by these vertices.

Now let a_{i_u} be an element of $\mathcal{A} - \{a_{i_j}, a_{i_\ell}, a_{i_s}\}$. Apply the 3-cube condition to the vertex $y = x + a_{i_j}$ and the vectors $a_{i_\ell}, a_{i_s}, a_{i_u}$ to conclude that $x + a_{i_j} + a_{i_\ell} + a_{i_s} + a_{i_u} \in G$. By induction on $3 \leq m \leq k$, we obtain in this way that for any subset $a_{i_{j(1)}}, \dots, a_{i_{j(m)}}$ of $\mathcal{A}$ of cardinality m, the vertex $x + \sum_s a_{i_{j(s)}}$ is contained in G. Since G is locally determined by its 0-skeleton, by the definition of $\mathcal{C}(G)$, the cube C in $\mathbb{R}^n$ which is spanned by the edges connecting x to $x + a_{i_j}$ is contained in $\mathcal{C}(G)$. Since the link complex of the cube C at the vertex x is the simplex spanned by these edges, this yields that the link complex of x in $\mathcal{C}(G)$ satisfies the flag condition. Thus $\mathcal{C}(G)$ is nonpositively curved which is what we wanted to show. □

For a complete train track τ on S let

$$E(\tau) \subset \mathcal{TT}$$

be the full subgraph of $\mathcal{TT}$ whose vertex set is the set of all train tracks which can be obtained from τ by a splitting sequence. By Lemma 5.1 of [H09], such a splitting sequence is unique up to the order of commuting splits, and the *splitting distance* to τ of a train track $\eta \in E(\tau)$, defined to be the shortest length of an edge path in $E(\tau)$ connecting τ to η, is the length of a splitting sequence connecting τ to η. Moreover, if η' is obtained from η by a single split, then the splitting distance of η' to τ equals the splitting distance of η to τ plus one. In particular, edges in $E(\tau)$ either strictly increase or strictly decrease the splitting distance to τ and consequently any closed edge path in $E(\tau)$ has even length.

For a complete geodesic lamination ξ carried by τ define $E(\tau, \xi)$ to be the complete subgraph of $E(\tau)$ consisting of all train tracks which carry ξ. If $\sigma \in E(\tau)$, then let $E(\tau, \sigma) \subset E(\tau)$ be the subgraph of all train tracks which are splittable to σ.

The following is an extension of Lemma 5.4 of [H09].

Lemma 5.3. *Let $\sigma, \eta \in E(\tau)$ be any two vertices. Then there is a unique train track $\Theta_-(\sigma, \eta) \in E(\tau)$ such that $\sigma, \eta \in E(\Theta_-(\sigma, \eta))$ and that the splitting distance between τ and $\Theta_-(\sigma, \eta)$ is maximal with this property.*

Proof. As τ is splittable to both σ, η, and σ, η have finite splitting distance to τ, there exists a train track $\Theta_-(\sigma, \eta) \in E(\tau)$ so that $\sigma, \eta \in E(\Theta_-(\sigma, \eta))$ and such that the splitting distance between τ and $\Theta_-(\sigma, \eta)$ is maximal with this property. We have to show that $\Theta_-(\sigma, \eta)$ is unique.

To this end assume that there exists a second such train track ξ, of the same splitting distance to τ. By Lemma 5.4 of [H09], as $\Theta_-(\sigma, \eta), \xi \in E(\tau, \sigma)$, there exists a unique train track $\beta \in E(\tau, \sigma)$ which is splittable to both $\Theta_-(\sigma, \eta), \xi$ and which has maximal splitting distance to τ with this property. Furthermore, there is a partition of the branches of β into two disjoint subsets B_1, B_2 so that $\Theta_-(\sigma, \eta)$ is obtained from β by a splitting sequence only involving splits at the branches in B_1, and ξ is obtained from β by a splitting sequence only involving splits at the branches in B_2. As $\xi \neq \Theta_-(\sigma, \eta)$ and as both $\Theta_-(\sigma, \eta), \xi$ have the same splitting distance to τ and hence the same splitting distance to β, both sets B_1, B_2 contain a large branch, and a splitting sequence connecting β to $\Theta_-(\sigma, \eta), \xi$ is of positive length.

But if $e_1 \in B_1$ is a large branch so that a (right or left) split β' of β at e_1 is splittable to $\Theta_-(\sigma, \eta)$, then since $\Theta_-(\sigma, \eta)$ is splittable to both σ, η, β' is splittable to both σ, η. Furthermore, since the splitting sequence connecting β to ξ does not involve a split at e_1, under the natural identification of branches of β with branches of ξ, the branch e_1 in ξ is large. By uniqueness of splitting sequences, this implies that a (right or left) split ξ' of ξ at e_1 is splittable to both σ, η. But the splitting distance between τ and ξ' is strictly larger than the splitting distance between τ and ξ which contradicts the assumption on ξ. This contradiction shows the lemma. □

We use these observations to establish the following strengthening of Lemma 5.1 and Lemma 5.4 of [H09]. For its formulation, let $m \geq 2$ be the number of branches of a complete train track τ on S.

Proposition 5.4. *If λ is a complete geodesic lamination carried by τ, then there is an isometry Φ of $E(\tau,\lambda)$, equipped with the intrinsic path metric d_E, onto a cubical graph in $\mathbb{R}^m$ which maps any splitting arc in $E(\tau,\lambda)$ to a directed edge path in $\Phi(E(\tau,\lambda))$. The cubical graph $\Phi(E(\tau,\lambda))$ is locally determined by its 0-skeleton, and its span $\mathcal{C}(\tau,\lambda)$ is nonpositively curved.*

Proof. The first part of the proposition is just the statement of Lemma 5.1 of [H09]. We have to show that the image cubical graph $\Phi(E(\tau,\lambda))$ is locally determined by its 0-skeleton, and is nonpositively curved.

To see that this is the case, we recall from the proof of Lemma 5.1 of [H09] the construction of the map Φ. Namely, number the vertices of τ in an arbitrary way with numbers $1,\dots,m$. If η is obtained from τ by a single split at a large branch e, then there is a natural bijection $\varphi(\tau,\eta)$ of the branches of τ onto the branches of η which induces a numbering of the branches of η. For this numbering, a branch not incident on an endpoint of e is mapped to the corresponding branch in η, where we view this branch as being contained in the complement of an open neighborhood of e in τ which is unchanged in the splitting process. The branch e is mapped to the diagonal of the split. Thus we can talk about numbered splitting sequences.

Choose a point $\Phi(\tau)\in\mathbb{Z}^m$ in an arbitrary way, say $\Phi(\tau)=0$. For $\eta\in E(\tau,\lambda)$ define $\Phi(\eta)$ as follows. Connect τ to η by a splitting sequence, say the sequence $(\tau_i)_{0\leq i\leq k}$, with $\tau_0=\tau$ and $\tau_k=\eta$, and let $\Phi(\eta)=\Phi(\tau_{k-1})+e_j$ where j is the number of the large branch of τ_{k-1} defining the split and where e_j is the j-th standard basis vector of $\mathbb{R}^n$. Lemma 5.1 of [H09] shows that this is well defined, that is, it does not depend on the choice of the splitting sequence. Connect $\Phi(\tau_{k-1})$ with $\Phi(\tau_k)$ by an edge.

We claim that $\Phi(E(\tau,\lambda))=G$ is locally determined by its 0-skeleton. To this end let $x=\Phi(\xi)\in\Phi(E(\tau,\lambda))$ and assume that $x+e_j=\Phi(\eta)\in\Phi(E(\tau,\lambda))$ for some j. Let k be the splitting distance between τ and ξ; then the splitting distance between τ and η equals $k+1$.

By Lemma 5.3, there exists a unique train track $\Theta_-(\xi,\eta)\in E(\tau,\lambda)$ so that $\xi,\eta\in E(\Theta_-(\xi,\eta))$ and such that $\Theta_-(\xi,\eta)$ has maximal splitting distance to τ with this property. Furthermore, there is a partition of the branches of $\Theta_-(\xi,\eta)$ into two disjoint subsets B_1,B_2 so that ξ can be obtained from $\Theta_-(\xi,\eta)$ by a splitting sequence at branches in B_1, and η can be obtained from $\Theta_-(\xi,\eta)$ by a splitting sequence at branches in B_2. Thus $\Phi(\xi)-\Phi(\Theta_-(\xi,\eta))$ is contained in the linear span of the basis vectors e_ℓ for $\ell\in B_1$, and $\Phi(\eta)-\Phi(\Theta_-(\xi,\eta))$ is contained in the linear span of the basis vectors e_u for $u\in B_2$. Since $\Phi(\eta)=\Phi(\xi)+e_j$, this implies that $\xi=\Theta_-(\xi,\eta)$ and indeed, η is obtained from ξ by a split at the branch j and x is connected to $x+e_j$ by an edge. As a consequence, the cubical graph G is locally determined by its 0-skeleton.

The above discussion implies the following.

(1) If $x\in G$ and if $i_1\neq i_2$ are such that $x+e_{i_1},x+e_{i_2}\in G$, then $x+e_{i_1}+e_{i_2}\in G$.
(2) If $x\in G$ and if $i_i\neq i_2$ are such that $x-e_{i_1},x-e_{i_2}\in G$ then $x-e_{i_1}-e_{i_2}\in G$.

Write $\xi < \eta$ if ξ is splittable to η. This defines a partial order on $E(\tau, \lambda)$. If we denote by $\rho_1, \dots, \rho_n$ the basis of $(\mathbb{R}^n)^*$ dual to the basis $e_1, \dots, e_n$, then we have $\xi < \eta$ if and only if $\rho_j(\Phi(\xi)) \leq \rho_j(\Phi(\eta))$ for all $j \leq n$. Since the cubical graph G is locally determined by its 0-skeleton, Lemma 5.2 shows that its span is nonpositively curved if G fulfills the 3-cube condition.

To verify that this is indeed the case let $x \in G$ and assume that for some $j = 1, 2, 3$, the vertices $x + \iota_{i_j} e_{i_j}$ are contained in G where $\iota_{i_j} = \pm 1$, and that the same holds true for $x + \iota_{i_j} e_{i_j} + \iota_{i_s} e_{i_s}$ for $j \neq s$. We distinguish three cases.

Case 1: $\iota_{i_j} = -1$ for all $j \in \{1, 2, 3\}$.

By property (2) above, applied to each of the points $x - e_{i_j}$, we know that $x - e_{i_j} - e_{i_s} \in G$ for all $j \neq s$. We can now apply property (2) to the vertices $x - e_{i_1}, x - e_{i_1} - e_{i_2}, x - e_{i_1} - e_{i_3}$ to conclude that $x - e_{i_1} - e_{i_2} - e_{i_3} \in G$. This verifies the 3-cube condition in this case.

Case 2: $\iota_{i_j} = 1$ for all $j \in \{1, 2, 3\}$.

If we use property (1) instead of property (2) above, and change signs, then the argument in Case 1 is valid to show the 3-cube condition.

Case 3: Up to permuting labels, $\iota_{i_1} = 1$ and $\iota_{i_2} = \iota_{i_3} = -1$.

By assumption, for $j = 2, 3$ the vertices $x, x + e_{i_1}, x - e_{i_j}$ are contained in a square whose fourth vertex is $x + e_{i_1} - e_{i_j}$. Writing $y = x + e_{i_1}$, we observe that $y, y - e_{i_1} = x, y - e_{i_2} = x + e_{i_1} - e_{i_2}, y - e_{i_3} = x + e_{i_1} - e_{i_3} \in G$. Using property (2) for this quadruple of points as in Case 1 above yields the the 3-cube condition.

Case 4: Up to permuting labels, $\iota_{i_1} = -1$ and $\iota(i_2) = \iota(i_3) = 1$.

Arguing as in Case 3 above reduces the claim to Case 2.

As these cases exhaust all possibilities, we conclude that the 3-cube condition is fulfilled for the graph G. Since G is locally determined by its 0-skeleton, the span $\mathcal{C}(G)$ of G is a nonpositively curved cube complex as claimed. $\square$

To summarize, for any complete geodesic lamination λ carried by τ, the cubical graph $E(\tau, \lambda)$ is the one-skeleton of a non-positively curved cube complex $\mathcal{C}(\tau, \lambda)$. Our next goal is to show that the same holds true for $E(\tau)$. We begin with having a closer look at cycles in $E(\tau)$ of length four.

Lemma 5.5. *Let $c \subset E(\tau)$ by an embedded closed edge-path of length four. Then there is a train track $\eta \in E(\tau)$, and there are two large branches $a \neq b$ of η such that the vertices of the edge path c are η, two train tracks which are obtained from η by a split at a, b, respectively, and a train track σ obtained from η by a split at both a, b. In particular, $c \subset E(\eta, \sigma) \subset E(\tau, \sigma)$.*

Proof. Let c be a closed embedded edge-path in $E(\tau)$ of length four. Let η be a vertex in this edge path of minimal splitting distance to τ. Let ζ, ξ be the two vertices on the edge path c connected to η by an edge. As edges in $E(\tau)$ either strictly increase or strictly decrease the splitting distance to τ (compare Section 5 of [H09]) and η is of minimal splitting distance to τ, the train tracks ζ, ξ are obtained from η by a single split. Thus there are two (a priori not distinct) large branches a, b of η such that ζ, ξ can be obtained from η by a split at a, b.

The vertex σ on c opposite to η is connected to both ζ, ξ by an edge. Thus its splitting distance to τ either coincides with the splitting distance of η, or it equals this splitting distance plus two.

If the splitting distance of σ to τ equals the splitting distance of η plus two, then both ζ, ξ are splittable to σ. Together with uniqueness of splitting sequences, this implies that the large branches a, b of η are distinct, that is, ζ, ξ are not obtained from η by a right and left split, respectively, at the same large branch. As ζ is obtained from η by a split at the large branch a, it contains the large branch b (using the natural bijection between the large branches of η and ζ as before), and σ is obtained from ζ by a split at b. By symmetry, a is a large branch in ξ, and σ is obtained from ξ by a split at a. Thus the statement of the lemma holds true in this case.

To complete the proof of the lemma it now suffices to show that the splitting distance between τ and σ can not coincide with the splitting distance between τ and η. We argue by contradiction and assume that this is the case. Then $\sigma \neq \eta \in E(\tau)$ are of the same splitting distance to τ, and they are both splittable with a single split to ζ, ξ. This violates uniqueness of the minimizer $\Theta_-(\zeta, \xi)$ of maximal splitting distance to τ in Lemma 5.3. The lemma follows. □

We are now ready to construct a cube complex $\mathcal{C}(\tau)$ from the graph $E(\tau)$ as follows.

If $c \subset E(\tau)$ is an embedded cycle of length 4, then we attach to c an euclidean square of side length one whose vertices are the vertices of $E(\tau)$ contained in c. If a collection of such squares is isometric to the boundary of a standard cube of dimension 3, then we attach a 3-cube to these squares, with an isometry of the boundary, and we continue inductively. Let $\mathcal{C}(\tau)$ be the resulting cube complex. Observe that for any complete geodesic lamination λ carried by τ, the nonpositively curved cube complex $\mathcal{C}(\tau, \lambda)$ is isometrically embedded in $\mathcal{C}(\tau)$. Furthermore, the number of edges of $\mathcal{C}(\tau)$ incident on a fixed vertex is bounded from above by twice the number of branches of a complete train track on S and hence $\mathcal{C}(\tau)$ is *uniformly finite* (which is defined to be a cube complex such that the number of edges incident on any vertex is uniformly bounded).

A *combinatorial geodesic* in a CAT(0) cube complex X is an edge path in X which is a geodesic in the one-skeleton of the complex. A *splitting arc* in $\mathcal{TT}$ to be a simplicial edge path with consecutive vertices τ_i such that for each i, τ_{i+1} can be obtained from τ_i by a single split. Let as before $m > 0$ be the number of branches of a complete train track τ on S. We are now ready to show the main result of this section.

Proposition 5.6. *Let τ be a complete train track on S.*

1. *The cube complex $\mathcal{C}(\tau)$ is* CAT(0)*.*
2. *Splitting arcs are combinatorial geodesics in $\mathcal{C}(\tau)$.*
3. *There exists a natural locally injective cubical map $\Phi : \mathcal{C}(\tau) \to \mathbb{R}^{2m}$ where $\mathbb{R}^{2m}$ is equipped with the standard grid cubulation.*

Proof. We first construct a map $\Phi : C(\tau) \to \mathbb{R}^{2m}$ as stated in part (3) of the proposition.

Number the m branches of τ in an arbitrary way with numbers $1, \dots, m$. If (τ_i) is any splitting sequence beginning with $\tau_0 = \tau$, then the numbering of the branches of τ induces a numbering of the branches of τ_i for all i, and this numbering only depends on the numbering of the branches of τ but not on the choice of a splitting sequence connecting τ to τ_i (see [H09] for details). If we keep track of these numberings then we talk about numbered splitting sequences.

Extending the construction in Section 5 of [H09], define a map $\Phi : E(\tau) \to \mathbb{R}^{2m}$ as follows. Let $e_1, \dots, e_{2m}$ be the standard basis of $\mathbb{R}^{2m}$, and let $\rho_1, \dots, \rho_{2m}$ be the dual basis of $(\mathbb{R}^{2m})^*$. Define $\Phi(\tau) = 0$. Given a splitting sequence (τ_i) starting from $\tau_0 = \tau$, we extend Φ to (τ_i) inductively as follows.

Let us assume that we defined already $\Phi(\tau_i)$. Assume that τ_{i+1} is obtained from τ_i by a single split at the large branch j. Define $\Phi(\tau_{i+1}) = \Phi(\tau_i) + e_j$ if the split connecting τ_i to τ_{i+1} is a right split, and define $\Phi(\tau_{i+1}) = \Phi(\tau_i) + e_{m+j}$ if this split is a left split. Connect $\Phi(\tau_i)$ and $\Phi(\tau_{i+1})$ by an edge of the standard grid.

It follows from Lemma 5.1 of [H09] that with this definition, the image $\Phi(\eta)$ of $\eta \in E(\tau)$ only depends on the choice of $\Phi(\tau)$ and a choice of a numbering of the branches of τ but not on the choice of a splitting sequence connecting τ to η. The thus defined map $\Phi : E(\tau) \to \mathbb{R}^{2m}$ maps vertices of $E(\tau)$ to points in $\mathbb{Z}^{2m}$, and it maps edges to edges in the standard grid of $\mathbb{R}^{2m}$. Every vertex in the image of Φ can be obtained from 0 by a directed edge path, and splitting arcs are mapped by Φ to directed edge paths in the standard grid. Here an edge path in the standard grid of $\mathbb{R}^{2m}$ is called directed if the restriction of every standard coordinate function ρ_i to the path is non-decreasing. Furthermore, for $\eta \in E(\tau)$ the splitting distance between τ and η coincides with the shortest length of a path in the standard grid of $\mathbb{R}^{2m}$ between $0 = \Phi(\tau)$ and $\Phi(\eta)$, and this is just the sum $\sum_i (\rho_i(\Phi(\eta))$.

By Lemma 5.5 and the definition of the map Φ, closed edge paths in $E(\tau)$ of length 4 are mapped by Φ to the boundary of a square, that is, a cube of dimension 2. Equivalently, Φ maps a square in $\mathcal{C}(\tau)$ to a square in $\mathbb{R}^{2m}$. As $\mathcal{C}(\tau)$ is determined by its two-skeleton, we conclude that the map Φ extends to a locally injective cubical map of $\mathcal{C}(\tau)$ onto the span $\mathcal{C}$ of the cubical graph $\Phi(E(\tau))$. Furthermore, for every complete geodesic lamination λ carried by τ, it maps $\mathcal{C}(\tau, \lambda)$ onto a nonpositively curved subcomplex of $\mathcal{C}$.

Since nonpositive curvature of a cube complex is equivalent to the statement that the link complex of every vertex is a flag complex, to show that $\mathcal{C}(\tau)$ is nonpositively curved it now suffices to show the following claim: Let $\eta \in \mathcal{C}(\tau)$ and assume that for some $k \geq 3$ there are train tracks $\xi_1, \dots, \xi_k$ connected to η by an edge such that for any $i \neq j$, the vertices η, ξ_i, ξ_j are contained in a closed edge path in $E(\tau)$ of length 4, that is, their images under Φ define the one-skeleton of a $k-1$-simplex in the link complex of $\Phi(\eta)$. Then there exists a complete geodesic lamination λ which is carried by each of the train tracks η, ξ_i. Namely, if this is the case then these edges are contained in the link complex of the nonpositively curved cube complex $\mathcal{C}(\tau, \lambda)$ (see Proposition 5.4), and as the link complex of a nonpositively curved cube complex is a flag complex, this simplex is contained in the link complex of a cube $C \in \mathcal{C}(\tau, \lambda) \subset \mathcal{C}(\tau)$.

To see that the claim holds indeed true, assume by reordering that $\xi_1, \dots, \xi_s$ are obtained from η by a collapse (the inverse of a split), and that $\xi_{s+1}, \dots, \xi_k$ are obtained from η by a split. By Lemma 5.5, as for $i \neq j$ the train tracks η, ξ_i, ξ_j are contained in a square, no two distinct of the train tracks ξ_j for $s+1 \leq j \leq k$ are obtained from η by a split at the same large branch. But this means that there are pairwise distinct large branches $a_{s+1}, \dots, a_k$ of η so that ξ_j is obtained from η by a (right or left) split at a_j. The train track ζ obtained from η by a (right or left according to the split which results in ξ_i) split at each of the branches a_j has the property that η, ξ_i are all splittable to ζ. Furthermore, as each of the train tracks ξ_i is recurrent, the same holds true for ζ and hence ζ is complete. Thus if λ is any complete geodesic lamination carried by ζ, then $\eta, \xi_i \in E(\tau, \lambda)$ for all $i \leq k$ as required. This completes the proof that $\mathcal{C}(\tau)$ is nonpositively curved. Furthermore, splitting arcs are combinatorial geodesics in $\mathcal{C}(\tau)$ since the combinatorial distance in $\mathcal{C}(\tau)$ between τ and $\zeta \in \mathcal{C}(\tau)$ equals precisely the splitting distance between τ and ζ.

To conclude that $\mathcal{C}(\tau)$ is CAT(0) it now suffices to show that $\mathcal{C}(\tau)$ is simply connected, and this is the case if any closed edge-path γ in $\mathcal{C}(\tau)$ with basepoint τ is homotopic to the constant path.

We proceed by induction on the maximal splitting distance to τ of a point on the path γ. If this distance equals zero then γ is constant and there is nothing to show. Thus assume by induction that contractibility holds true for all edge loops in $\mathcal{C}(\tau)$ based at τ which only pass through vertices of splitting distance at most $k \geq 0$ to τ. Let γ be an edge loop based at τ which only meets vertices of splitting distance at most $k+1$ to τ. Assume that there are $n \geq 0$ vertices of splitting distance $k+1$ on γ. We successively remove these points with a homotopy as follows.

Assume without loss of generality that $n \geq 1$. Let $\ell > 0$ be such that the splitting distance between $\gamma(\ell)$ and τ equals $k+1$. As an edge in $E(\tau)$ either increases or decreases the splitting distance to τ by one, by Lemma 5.5 and its proof, the splitting distance of $\gamma(\ell-1)$, $\gamma(\ell+1)$ to τ equals k and therefore $\gamma(\ell-1), \gamma(\ell+1)$ are both splittable to $\gamma(\ell)$. If $\gamma(\ell-1) = \gamma(\ell+1)$ then $\gamma[\ell-1, \ell+1]$ passes through the same edge twice, in opposite direction, and we can homotope γ to an edge-path with $n-1$ points of splitting distance $k+1$ to τ.

Otherwise by Lemma 5.5 and Lemma 5.3 and their proofs, the arc $\gamma[\ell-1, \ell+1]$ is contained in the boundary of a two-dimensional cube C in $\mathcal{C}(\tau)$ whose fourth vertex $x = \Theta_-(\gamma(\ell-1), \gamma(\ell+1))$ opposite to $\gamma(\ell)$ has splitting distance $k-1$ to τ. Then $\gamma[\ell-1, \ell+1]$ can be homotoped in C with fixed endpoints to an edge path of length two passing through the vertices $\gamma(\ell-1), x, \gamma(\ell+1)$. This homotopy reduces the number of vertices on γ with splitting distance $k+1$ to τ by one. Removing in n steps each of vertices on γ of splitting distance $k+1$ to τ in this way yields a closed edge path which is homotopic to γ and which does not pass through a vertex of splitting distance $k+1$ to τ. An application of the induction hypothesis then yields the induction step and completes the proof that $\mathcal{C}(\tau)$ is CAT(0). □

The *interval* $\mathcal{I}(x,y)$ defined by two vertices x,y in the CAT(0) cube complex $\mathcal{C}(\tau)$ is the set of all vertices z such that $d(x,y)=d(x,z)+d(z,y)$ where d denotes the *combinatorial distance*, that is, the shortest length of a path in the one-skeleton of $\mathcal{C}(\tau)$ connecting x to y. The *median* of three distinct vertices x,y,z is the intersection $\mathcal{I}(x,y)\cap\mathcal{I}(y,z)\cap\mathcal{I}(z,x)$, and this median is unique. We have

Lemma 5.7. *Let $\sigma,\eta\in E(\tau)$; then the vertex $\Theta_-(\sigma,\eta)$ is the median of the points σ,η,τ in $\mathcal{C}(\tau)$.*

Proof. Let $\sigma,\eta\in E(\tau)$. Since τ is splittable to $\Theta_-(\sigma,\eta)$ and since splitting arcs are combinatorial geodesic in $\mathcal{C}(\tau)$, it suffices to show that there is a combinatorial geodesic connecting σ to η which passes through $\Theta_-(\sigma,\eta)$.

This statement is a consequence of Lemma 5.4 of [H09] if there exists some complete geodesic lamination λ such that $\sigma,\eta\in E(\tau,\lambda)$. To extend the claim to the more general situation at hand, we proceed by induction on the sum of the lengths of the two splitting sequences connecting $\Theta_-(\sigma,\eta)$ to σ,η.

If this length is zero there is nothing to show, so assume that the claim holds true if the sum of these lengths is at most $k-1$ for some $k-1\geq 0$. Let now σ,η be such that the sum of these lengths equals k. Let σ_i,η_j be a splitting sequence connecting $\Theta_-(\sigma,\eta)$ to σ,η. We may assume that the length of these sequences is positive. Connect σ to η by a combinatorial geodesic ζ_i.

We consider two cases. In the first case, ζ_1 is obtained from $\zeta_0=\sigma$ by a single split. We claim that in this case $\Theta_-(\zeta_1,\eta)$ is obtained from $\Theta_-(\sigma,\eta)$ by a single split at a large branch e with the following properties. Via the natural identification of the branches of $\Theta_-(\sigma,\eta)$ with the branches of σ defined by some (and hence any) splitting sequence connecting $\Theta_-(\sigma,\eta)$ to σ, the branch e is a large branch in σ, and the splitting sequence connecting $\Theta_-(\sigma,\eta)$ to η contains a split at e.

To see that this holds true note that by induction assumption, the length of a combinatorial geodesic connecting ζ_1 to η is the sum of the lengths of a splitting sequence connecting $\Theta_-(\zeta_1,\eta)$ to ζ_1,η, and hence the sum of these lengths has to be strictly smaller than the sum of the lengths of the splitting sequences σ_i,η_j. Since $\Theta_-(\sigma,\eta)$ is splittable to both ζ_1,η, we conclude that $\Theta_-(\zeta_1,\eta)$ is obtained from $\Theta_-(\sigma,\eta)$ by a nontrivial splitting sequence. By Lemma 5.3 and its proof, this splitting sequence is of length one and of the above form. But then the sum of the combinatorial distances between σ and $\Theta_-(\sigma,\eta)$ and η and $\Theta_-(\sigma,\eta)$ equals the combinatorial distance between σ,η and we are done.

In the second case, ζ_1 is obtained from σ by a single collapse. Using once more Lemma 5.3 and its proof, we conclude that we have $\Theta_-(\zeta_1,\eta)=\Theta_-(\sigma,\eta)$ in this case and the claim follows as before. This completes the proof of the lemma. □

Let φ be a pseudo Anosov mapping class whose horizontal and vertical measured geodesic laminations are complete (that is, minimal and maximal). Let τ be a train track defining a *splittable expansion* of φ. This means that τ can be connected to $\varphi(\tau)$ by a splitting sequence (τ_i), with $\tau_0=\tau$ and $\tau_n=\varphi(\tau)$. We require furthermore that a carrying map $\tau_{i+n}\to\tau_i$ maps every edge of τ_{i+n} *onto* τ_i. As every pseudo-Anosov mapping class admits a train track expansion, a pair consisting of a pseudo-Anosov mapping class φ with minimal and maximal horizontal and vertical measured geodesic lamination and a complete train track τ which is splittable to $\varphi(\tau)$ and has the above properties can be found by perhaps replacing a given pseudo-Anosov mapping class by a positive power.

Denote as before by $E(\tau_i)\subset\mathcal{TT}$ the complete subgraph of all train tracks which can be obtained from τ_i by a splitting sequence. Then $\varphi^{-1}(E(\tau_{i+n}))=E(\tau_i)\supset E(\tau_{i+n})$ for all i. Hence

$$E=\cup_j\varphi^{-j}E(\tau_0)$$

is the one-skeleton of a complete uniformly finite CAT(0)-cube complex

$$\mathcal{C}(E)$$

which contains each of the cube complexes $\mathcal{C}(E(\tau_i))$. Note that E is a complete subgraph of $\mathcal{TT}$.

For $i\in\mathbb{Z}$ and $j=0,\dots,n-1$ write $\tau_{in+j}=\varphi^i\tau_j$. Then for each i, the train track τ_i is splittable to τ_{i+1}.

Let ν be the support of the repelling measured geodesic lamination of the pseudo-Anosov mapping class φ. Then ν is a minimal complete uniquely ergodic geodesic lamination. Let $\mathcal{CL}$ be the space of complete geodesic laminations on S equipped with the Hausdorff topology. The following is taken from [H09], we include the short argument for completeness.

Lemma 5.8. *For every complete geodesic lamination $\mu\in\mathcal{CL}-\nu$ there exists some i such that μ is carried by τ_i.*

Proof. By Lemma 2.3 of [H09], the set U of all complete geodesic laminations carried by τ_0 is open and closed in $\mathcal{CL}$, and it contains the support of the attracting measured geodesic lamination for φ. As φ acts on $\mathcal{CL}$ with north-south dynamics, we conclude that $\cup_i\varphi^{-i}U=\mathcal{CL}-\{\nu\}$. Now a lamination in the set $\varphi^{-i}U$ is carried by $\varphi^{-i}\tau_0=\tau_i$ which shows the lemma. □

A map between metric spaces $F:X\to Y$ is called *coarsely surjective* if there exists a number $R>0$ such that every point in Y is at distance at most R to $F(X)$. We have

Proposition 5.9. *The map $\mathcal{C}(E)\to\mathcal{TT}$ which is defined by the vertex inclusion is coarsely surjective.*

Proof. Let F be any marking of S with pants decomposition Q. By Lemma 5.8, there exists i such that the train track τ_i carries each complete geodesic lamination ν in standard form for Q and that furthermore τ_i is filled by Q. The proposition now follows from Corollary 3.12. □

6. The principal curve graph

In this section we consider the *curve graph* $\mathcal{CG}(S)$ of a non-exceptional surface S of finite type. Its vertices are isotopy classes of essential simple closed curves on S, that is, simple closed curves which are not contractible or homotopic into a puncture, and where two such vertices are connected by an edge of length one if they can be realized disjointly. We introduce an electrification of the curve graph, called the *principal curve graph*, and show that it is hyperbolic, of infinite diameter. We study the Gromov boundary of this graph and prove Theorem 4 from the introduction. In Section 7, we use the principal curve graph to identify the *regular Roller boundary* of the CAT(0) cube complex $\mathcal{C}(E)$. This section can be read independently of the rest of the article.

Let c, d be two simple closed curves on S in minimal position, that is, c, d intersect in the minimal number of points. In the sequel we always assume that this is the case. Then $S - (c \cup d)$ is a union of complementary regions whose boundaries consist of subarcs of c and d in alternating order. In particular, a complementary component which is simply connected is a polygon with an even number of sides, and a complementary component which is a punctured disk is a punctured polygon with an even number of sides.

The curves c, d *bind* S if each component of $S-(c\cup d)$ is a disk or a once punctured disk. This is equivalent to stating that there is no component of $S - (c \cup d)$ which contains an essential simple closed curve.

Let us assume in the sequel that c, d bind S. By reasons of Euler characteristic, there is at least one complementary component which is a polygon with at least 6 sides or once punctured polygons with at least 4 sides.

Definition 6.1. The *principal curve graph* $\mathcal{PC}(S)$ of S is the graph whose vertices are isotopy classes of essential simple closed curves on S and where two such curves c, d are connected by an edge of length one if and only if there is a component of $S - (c \cup d)$ which either contains an essential simple closed curve of S, or is a polygon with at least 8 sides or a once punctured polygon with at least four sides.

By construction, the mapping class group $\mathcal{MCG}$ acts on $\mathcal{PC}(S)$ as a group of simplicial automorphisms. Furthermore, $\mathcal{PC}(S)$ contains a $\mathcal{MCG}$-invariant subgraph $\mathcal{CG}_0(S)$, with the same set of vertices, that is, vertices are simple closed curves, and where two such vertices c, d are connected by an edge of length one if and only if there is a component of $S - (c \cup d)$ which contains an essential simple closed curve of S. Then this curve is disjoint from both c, d and therefore the distance in the curve graph $\mathcal{CG}(S)$ of S between c, d is at most two. Vice versa, if c, d are connected in the curve graph by an edge, then they are disjoint, and $S - (c \cup d)$ contains a component which is not simply connected. Thus c, d are connected in $\mathcal{CG}_0(S)$ by an edge. This shows

Lemma 6.2. *$\mathcal{CG}_0(S)$ is two-quasi isometric to the curve graph $\mathcal{CG}(S)$ of S.*

As $\mathcal{PC}(S)$ is obtained from $\mathcal{CG}_0(S)$ by adding some edges, the graph $\mathcal{PC}(S)$ can be thought of as an electrification of $\mathcal{CG}(S)$. If $d_{\mathcal{PC}}(c, d) \geq 2$ then we say that c, d *completely fill* S.

The following lemma is based on a construction due to Masur and Minsky (Section 4 of [MM04]).

Lemma 6.3. *Let c, d be any two simple closed curves on S which bind S. Then there exist a recurrent one-switch train track $\eta(c, d)$ which carries d and intersects c in a single point. This train track is maximal only if c, d completely fill S.*

Proof. Using the notations from the lemma, choose a component I of $c - d$ contained in the boundary of a polygonal component of $S - (c \cup d)$ with at least 6 sides or a once punctured polygonal component of $S - (c \cup d)$ with at least four sides. Such a component C always exists by reasons of Euler characteristic. If the distance between c, d in $\mathcal{PC}(S)$ equals one, then we may assume that C is a polygonal component with at least 8 sides or a once punctured polygonal component with at least 4 sides. Contract $c - I$ to a point. The graph G obtained from $c \cup d$ in this way has a single vertex p which is the image of $c - I$. Furthermore, it intersects c in a single point. Impose a switch structure at the vertex p as follows.

Declare all half-edges of G which are subarcs of d and leave c to a fixed side to be incoming, and declare the half-edges of G which are subarcs of d and leave c to the opposite side to be outgoing. The result of this construction is a *bigon track* $\hat{\sigma}$, that is, a graph which has all properties of a train track except that it may contain bigons. These bigons can be collapsed to yield a train track σ which carries d and intersects c in a single point. As σ carries the simple closed curve d and is filled by d, the train track σ is recurrent.

Let us inspect the complementary regions of $\hat{\sigma}$. If E is a component of $S - (c \cup d)$ with 2ℓ sides not containing I, then each side e of E contained in c is contracted to a point in the construction of $\hat{\sigma}$, and the two sides of E adjacent to e meet at a cusp of the complementary region of $\hat{\sigma}$ which is the collapse of E. Thus any complementary polygon (or once punctured complementary polygon) of $S - (c \cup d)$ with 2ℓ sides not containing I gives rise to a complementary component of $\hat{\sigma}$ which is a topological disk (or once punctured disk) with ℓ cusps in the boundary. Complementary quadrangles collapse to complementary bigons.

The side I of the component C is removed, and the component C merges with the second component C' of $S - (c \cup d)$ which contains I in its boundary to a complementary component D of $\hat{\sigma}$. To analyze this component, we distinguish three cases.

If $C' = C$ then D contains an essential simple closed curve. Namely, in this case there is a simple closed curve in $C \cup I$ which intersects c in a single point contained in I, and this simple closed curve is contained in the complementary region D of $\hat{\sigma}$.

Now let us assume that C is a polygon without puncture and $2\ell \geq 6$ sides and that $C' \neq C$. If C' is a polygon with $2k \geq 4$ sides, then D is disk with $\ell + k - 2 \geq \ell$ cusps in the boundary. Similarly, if C' is a once punctured polygon with $2k \geq 2$ sides, then D is a once punctured disk with $\ell + k - 2 \geq \ell - 1$ cusps in the boundary.

By symmetry, we are left with the case that both C, C' are once punctured polygons. But then the component D contains two punctures and hence it contains an essential simple closed curve on S surrounding the punctures.

As a consequence, if C is a polygon with at least 8 sides or a once punctured polygon with at least 4 sides, then the bigon track $\hat{\sigma}$ and hence the train track σ contains a complementary component which either contains an essential simple closed curve, or is a polygon with at least 8 sides, or is a once punctured polygon with at least 4 sides. In particular, σ is not maximal. This completes the proof of the lemma. □

Denote by

$$\Psi : \mathcal{CG}(S) \to \mathcal{PC}(S)$$

the map induced by the vertex inclusion. This map is one-Lipschitz. An *unparameterized L-quasi-geodesic* in a geodesic metric space X is a map $\psi : [a, b] \to X$ with the property that there exists an increasing homeomorphism $\rho : [0, c] \to [a, b]$ such that $\psi \circ \rho : [0, c] \to X$ is an L-quasi-geodesic.

A *vertex cycle* of a recurrent train track τ is an immersed simple closed curve in τ which is an extreme point for the cone of all transverse measures for τ. Such a vertex cycle either is an embedded simple closed curve in τ or a dumbbell. In particular, a closed trainpath defined by a vertex cycle passes through any branch at most twice. As a consequence, the geometric intersection number between any two vertex cycles on τ is uniformly bounded, and there is a coarsely well defined map

$$\Upsilon : \mathcal{TT} \to \mathcal{CG}(S) \tag{6}$$

which associates to a train track one of its vertex cycles. Here coarsely well defined means that Υ depends on choices, but any two choices give rise to maps which map a given point to images of uniformly bounded distance. We refer to [H06] for more information on this construction.

Proposition 6.4. *The principal curve graph $\mathcal{PC}(S)$ is hyperbolic. Geodesics in $\mathcal{CG}(S)$ map by Ψ to uniform unparameterized quasi-geodesics in $\mathcal{PC}(S)$.*

Proof. The proposition follows from the work of Kapovich and Rafi [KR14] if we can verify that the conditions in Corollary 2.4 of [KR14] are fulfilled. For this it suffices to show the existence of a number $L > 1$ with the following property. Whenever c, d are curves whose distance in the principal curve graph is one, then there exists an L-quasi-geodesic $\rho : [0, m] \to \mathcal{CG}(S)$ connecting $\rho(0) = c$ to $\rho(m) = d$ such that the diameter of $\Psi(\rho[0, m]) \subset \mathcal{PC}(S)$ is at most L.

This is obvious if c, d do not bind S as in this case, their distance in the curve graph is at most two. If c, d bind S then we construct such a quasi-geodesic using the fact that splitting sequences of (not necessarily complete) train tracks on S are mapped by the map Υ to uniform unparameterized quasi-geodesics in $\mathcal{CG}(S)$ [H06].

Let σ be a one-switch train track as in Lemma 6.3. Then σ has a complementary component different from a trigon or a once punctured monogon. In particular, any two simple closed curves carried by σ have distance one in $\mathcal{PC}(S)$. Furthermore, the distance between a vertex cycle of σ and the curve c is uniformly bounded since this holds true for their distance in the curve graph.

There is a splitting and collision sequence σ_i issuing from σ which consists of train tracks which carry d and which connects σ to the train track which consists of the single curve d [PH92]. Associating to each train track σ_i one of its vertex cycles c_i defines a uniform unparameterized quasi-geodesic in $\mathcal{CG}(S)$ [H06] which connects a vertex cycle c_0 of $\sigma = \sigma_0$, that is, a simple closed curve in a uniformly bounded neighborhood of c in $\mathcal{CG}(S)$, to the curve d. Since each of the curves c_i is carried by σ, this quasi-geodesic consists of curves whose distance to d in the principal curve graph equals at most one.

Hyperbolicity of the principal curve graph now follows from Corollary 2.4 of [KR14]. This result also shows that geodesics in $\mathcal{CG}(S)$ map to uniform unparameterized quasi-geodesic in $\mathcal{PC}(S)$. This is what we wanted to show. □

Proposition 6.4 does not imply that the principal curve graph has infinite diameter. We next show that this is indeed the case. To this end recall that by a construction due to Thurston and Veech, simple closed curves c, d which bind S determine a line of *area one holomorphic quadratic differentials* on S with horizontal and vertical measured geodesic lamination supported in c, d, respectively. These differentials define the cotangent line of a *Teichmüller geodesic*. If c, d are connected by an edge in the principal curve graph, then these quadratic differentials either have a zero of order at least two, or one of the marked points (punctures) is a regular point or a zero of the differential and not a simple pole.

A simple closed curve admits a unique transverse measure up to scale and hence can be viewed as a projective measured geodesic lamination. Using compactness of the space $\mathcal{PML}$ of projective measured geodesic laminations we observe

Lemma 6.5. *Let (c_i, d_i) be a sequence of pairs of simple closed curves such that $d_{\mathcal{PC}}(c_i, d_i) \leq 1$ for all i. Assume that the sequence c_i converges as $i \to \infty$ in $\mathcal{PML}$ to a projective measured geodesic lamination whose support μ is both minimal and maximal. Then after perhaps passing to a subsequence, the sequence d_i converges to a projective measured geodesic lamination with support μ.*

Proof. Let $\rho \in \mathcal{PML}$ be a limit of the sequence c_i. The support of ρ equals μ. Using the notations from the lemma, by compactness and passing to a subsequence of the sequence d_i, we may assume that d_i converges in $\mathcal{PML}$ to a projective measured geodesic lamination ξ. We argue by contradiction and assume that the support of ξ is distinct from μ. Since μ is minimal and maximal, this implies that ξ together with the projective measured geodesic lamination ρ determines a Teichmüller geodesic whose cotangent line γ consists of quadratic differentials with vertical and horizontal projective measured geodesic laminations ρ, ξ, respectively.

Now two projective measured geodesic laminations ρ, ξ determine a Teichmüller geodesic in this way if and only if for every measured geodesic lamination ζ, the geometric intersection number between ζ and at least one of the two laminations ρ, ξ is positive. Note that this property is invariant under scaling and hence makes sense for projective measured geodesic laminations. As this is an open condition, we conclude that for sufficiently large i the pair (c_i, d_i) defines the cotangent line γ_i of a Teichmüller geodesic, that is, an orbit of the Teichmüller flow on bundle of area one quadratic differentials. Furthermore, by continuity, the cotangent lines γ_i of these Teichmüller geodesics converge locally uniformly to γ in the bundle over Teichmüller space whose fiber over a Riemann surface X is the sphere of area one quadratic differentials on X.

On the other hand, as the distance in the principal curve graph between c_i, d_i equals one, the cotangent line γ_i is *not* contained in the principal stratum of quadratic differentials with only simple zeros and simple poles at the marked points (here we view an abelian differential as a quadratic differential with all zeros of even order). As the complement of the principal stratum in the Teichmüller space of quadratic differentials is closed and invariant under the Teichmüller flow, the cotangent line of the limiting Teichmüller geodesic is contained in the complement of the principal stratum as well. But any quadratic differential whose horizontal measured lamination is supported in a minimal complete geodesic lamination is contained in the principal stratum. This is a contradiction which shows the lemma. □

We are now ready to complete the proof that the diameter of $\mathcal{PC}(S)$ is infinite. The following proposition gives a more precise information used for the investigation of hyperplanes in $\mathcal{C}(E)$. For its purpose and later use, define a *full split* of a train track η to be a train track obtained from η by a single split at each large branch. A *full splitting sequence* is a sequence (ξ_i) of complete train tracks such that for each i, the train track ξ_{i+1} can be obtained from ξ_i by a full split. If λ is a complete geodesic lamination carried by η, then a *full λ-split* of η is a full split with the property that the split track carries λ.

For the purpose of the proof of the following propopsition and later use, recall [H06, K99] that the *Gromov boundary* of the curve graph is the space of minimal geodesic laminations which *fill* S, that is, which intersect each simple closed curve on S transversely, equipped with the *coarse Hausdorff topology.* In this topology, a sequence ν_i of minimal geodesic laminations which fill converges to a minimal filling lamination ν if any limit of ν_i in the Hausdorff topology contains ν as a sublamination. Note that it also makes sense for a sequence of simple closed curves to converge to a minimal filling lamination in the coarse Hausdorff topology.

Proposition 6.6. *Let η be a complete train track and let λ be a minimal complete geodesic lamination carried by η. Let (η_i) be a sequence of full λ-splits issuing from $\eta_0 = \eta$. If c_i is a vertex cycle of η_i then $d_{\mathcal{PS}}(c_0, c_i) \to \infty$. In particular, the diameter of $\mathcal{PC}(S)$ is infinite.*

Proof. The proof of this proposition is a variation of an argument of Luo as used in [MM99].

Let (η_i) be any full splitting sequence consisting of λ-splits for a minimal complete geodesic lamination λ. The map which associates to i the simple closed curve $\Upsilon(\eta_i) = c_i$ is a uniform unparameterized quasi-geodesic in the curve graph [H06] converging to λ, viewed as a point in the Gromov boundary of $\mathcal{CG}(S)$. In particular, c_i converges in the Hausdorff topology to λ as $i \to \infty$ [H06] (more precisely, we would have to consider convergence in the coarse Hausdorff topology, but as λ is minimal and complete, this is equivalent to convergence in the Hausdorff topology).

By Proposition 6.4, the assignment $i \to c_i = \Upsilon(\eta_i)$ also defines a uniform unparameterized quasi-geodesic in the principal curve graph. Let us assume to the contrary that this quasi-geodesic is of finite diameter, that is, that $d_{\mathcal{PC}}(c_0, c_i)$ is uniformly bounded. Then by passing to a subsequence, we may assume that there is a number $k > 0$ so that $d_{\mathcal{PC}}(c_0, c_i) = k$ for all i. Let $c_i^1 \in \mathcal{PC}(S)$ be a vertex so that $d_{\mathcal{PC}}(c_0, c_i^1) = k - 1$ and $d_{\mathcal{PC}}(c_i^1, c_i) = 1$ for all i.

By Lemma 6.5, we know that $c_i^1 \to \lambda$ in the Hausdorff topology. Repeat this argument with the sequence c_i^1. After k such steps we conclude that $c_0 \to \lambda$ in the Hausdorff topology, which is a contradiction. □

We use this to show

Proposition 6.7. *The Gromov boundary of the principal curve graph is the space of minimal complete geodesic laminations, equipped with the Hausdorff topology.*

Proof. A point in the Gromov boundary $\partial\mathcal{CG}(S)$ of $\mathcal{CG}(S)$ can be viewed as an equivalence class of uniform quasi-geodesic rays in $\mathcal{CG}(S)$ where two such quasi-geodesic rays are equivalent if their Hausdorff distance in $\mathcal{CG}(S)$ in finite (this is true in the situation at hand in spite of the fact that the curve graph is not locally finite). The same holds true for the Gromov boundary of the principal curve graph.

Since $\mathcal{CG}(S)$ and $\mathcal{PC}(S)$ have the same vertices and, by Proposition 6.4, the same uniform quasi-geodesics up to parameterization, we conclude that the Gromov boundary of the principal curve graph is a quotient of the subspace of the Gromov boundary of $\mathcal{CG}(S)$ consisting of equivalence classes of those quasi-geodesic rays in $\mathcal{CG}(S)$ whose diameter in $\mathcal{PC}(S)$ are infinite.

By [H06], the image under the map Υ of a full splitting sequence consisting of λ-splits for a minimal complete geodesic lamination λ is a uniform unparameterized quasi-geodesic in the curve graph converging to $\lambda \in \partial\mathcal{CG}(S)$, and by Proposition 6.6, this unparameterized quasi-geodesic has infinite diameter in $\mathcal{PC}(S)$. Furthermore, two distinct such minimal complete geodesic laminations define non-equivalent quasi-geodesic rays in $\mathcal{PC}(S)$ and hence distinct points in the Gromov boundary of $\mathcal{PC}(S)$. Namely, such a pair of points determines a biinfinite uniform quasi-geodesic in the curve graph and hence by Proposition 6.4 a biinfinite unparameterized quasi-geodesic in $\mathcal{PC}(S)$ whose two half-rays have infinite diameter by Proposition 6.6. As a consequence, the subspace of $\partial\mathcal{CG}(S)$ of all minimal complete geodesic laminations embeds (as a set) into the Gromov boundary of $\mathcal{PC}(S)$.

We show next that a minimal filling geodesic lamination which is not complete defines an equivalence class of uniform quasi-geodesic rays in $\mathcal{CG}(S)$ which have finite diameter in $\mathcal{PC}(S)$.

Thus let λ be a minimal geodesic lamination which fills S but which is not complete. Let $\hat\lambda$ be a complete geodesic lamination which contains λ as its minimal component. By Lemma 3.2 of [H09], there exists a complete train track η containing a subtrack ξ so that λ is carried by ξ and such that there is a bijection between the complementary components of λ and the complementary components of ξ. In particular, at least one of the complementary components of ξ is a polygon with more than three sides or a once punctured polygon with at least two sides.

Let (η_i) be a splitting sequence beginning at $\eta_0 = \eta$ which is induced from a full sequence of λ-splits ξ_i of ξ. For each i let c_i be a vertex cycle of ξ_i. As both c_0, c_i are carried by ξ, there is at least one component of $S - (c_0 \cup c_i)$ which is a polygon with at least 8 sides or a once punctured polygon with at least four sides. This yields that the distance in $\mathcal{PC}(S)$ between c_0 and c_i is at most one for all i. Since $i \to c_i$ is an unparameterized quasi-geodesic in the curve graph of S which converges in $\mathcal{CG}(S) \cup \partial\mathcal{CG}(S)$ to $\lambda \in \partial\mathcal{CG}(S)$ [H06], this implies that λ does not define a point in the Gromov boundary of $\mathcal{PC}(S)$.

To summarize, as a set, the Gromov boundary of $\mathcal{PC}(S)$ can be identified with the subspace $\partial\mathcal{PC}(S)$ of $\partial\mathcal{CG}(S)$ of all minimal complete geodesic laminations. That this inclusion is a homeomorphism onto its image can be seen with the same arguments used before. Namely, as the inclusion $\mathcal{CG}(S) \to \mathcal{PC}(S)$ is one-Lipschitz, the topology on $\partial\mathcal{PC}(S)$ as a subspace of $\partial\mathcal{CG}(S)$ is finer than the topology on $\partial\mathcal{PC}(S)$ as the Gromov boundary of $\mathcal{PC}(S)$. Thus it suffices to show that the inclusion map $\partial\mathcal{PC}(S) \to \partial\mathcal{CG}(S)$ is continuous where $\partial\mathcal{PC}(S)$ is equipped with the topology as the Gromov boundary of $\mathcal{PC}(S)$, and this is equivalent to stating that if $\lambda_i \to \lambda$ in $\partial\mathcal{PC}(S)$, then λ_i converges to λ in the Hausdorff topology.

To establish this claim we argue by contradiction and we assume that there exists a sequence $\lambda_i \subset \partial\mathcal{PC}(S)$ which converges to $\lambda \in \partial\mathcal{PC}(S)$ but does not converge to λ in $\partial\mathcal{CG}(S)$. By compactness of the space of geodesic laminations equipped with the Hausdorff topology, up to passing to a subsequence, we may assume that $\lambda_i \to \mu$ in the Hausdorff topology where $\mu \neq \lambda$. Assume without loss of generality that $\lambda_i \neq \lambda$ for all i. Equip λ_i, λ with projective transverse measures α_i, β. Then for each i, the pair (α_i, β) of projective measured geodesic laminations determines a Teichmüller geodesic γ_i.

By passing to another subsequence, we may assume that the Teichmüller geodesics γ_i converge locally uniformly to a Teichmüller geodesic γ defined by a pair (α, β) of projective measured geodesic laminations, where α is supported in $\mu \neq \lambda$. Now by [MM99], associating to a point X on the Teichmüller geodesic γ a simple closed curve in $\mathcal{CG}(S)$ whose length for the hyperbolic metric X is uniformly bounded defines a uniform unparameterized quasi-geodesic in $\mathcal{CG}(S)$. As the Teichmüller geodesics γ_i, γ all pass through a fixed compact subset of Teichmüller space, together with Proposition 6.4, this implies that there are uniform quasi-geodesics in $\mathcal{PC}(S)$ connecting λ to λ_i which pass through a fixed subset of $\mathcal{PC}(S)$

of uniformly bounded diameter, violating the assumption that $\lambda_i \to \lambda$ in $\partial\mathcal{PC}(S)$. This completes the proof of the proposition. □

Recall that any pseudo-Anosov mapping class $\varphi \in \mathcal{MCG}$ has a unique *axis* in Teichmüller space $\mathcal{T}(S)$ of S, that is, an invariant Teichüller geodesic on which φ acts as a nontrivial translation.

Corollary 6.8. *Let $\varphi \in \mathcal{MCG}$ be a pseudo-Anosov mapping class. The φ acts as a hyperbolic isometry on the principal curve graph if and only if the cotangent line of its axis is contained in the principal stratum of quadratic differentials with only simple zeros.*

Proof. For a point $X \in \mathcal{T}(S)$ define $\Upsilon_0(X) \in \mathcal{CG}(S)$ to be a simple closed curve of uniformly bounded length. If γ is any Teichmüller geodesic, then the assignment $t \to \Upsilon_0(\gamma(t))$ is a uniform unparameterized quasi-geodesic in $\mathcal{CG}(S)$ [MM99].

Let $\gamma \subset \mathcal{T}(S)$ be the axis of a pseudo-Anosov element φ and assume that the cotangent line of γ is defined by quadratic differentials in the principal stratum. Then the horizontal and vertical measured geodesic laminations are minimal and complete as there can not be any horizontal or vertical saddle connections. By Proposition 6.4 and Proposition 6.6, the image under Υ_0 of the line γ is an unparameterized quasi-geodesic in $\mathcal{PC}(S)$ of infinite diameter. As φ acts on this quasi-geodesic as a nontrivial translation, this quasi-geodesic is a quasi-axis for φ acting on $\mathcal{PC}(S)$ and hence φ acts as a hyperbolic isometric on $\mathcal{PC}(S)$.

Vice versa, if the axis of φ is not contained in the principal stratum, then the support of the horizontal and vertical measured geodesic laminations defined by this axis is not complete. By Proposition 6.7, this implies that the image of the axis under the map Υ_0 has finite diameter in $\mathcal{PC}(S)$. The corollary follows. □

We use Proposition 6.7 and Corollary 6.8 to show Theorem 4 from the introduction. For its formulation, call a subgroup Γ of $\mathcal{MCG}$ *non-elementary* if it contains at least two independent pseudo-Anosov elements.

The *entropy* of a probability measure μ on the group $\mathcal{MCG}$ is defined as

$$H(\mu) = -\sum_{g\in\mathcal{MCG}} \mu(g)\log\mu(g).$$

The measure μ is said to have *finite logarithmic moment* for the action of $\mathcal{MCG}$ on the curve graph $\mathcal{CG}(S)$ if

$$\sum_{g\in\mathcal{MCG}} \mu(g)|\log d_{\mathcal{CG}(S)}(c, gc)| < \infty.$$

Theorem 6.9. *Let Ω be a random walk on $\mathcal{MCG}$ generated by a probability measure μ on $\mathcal{MCG}$ with the following properties.*

1. *The support of μ generates a non-elementary subgroup Γ of $\mathcal{MCG}$ as a semi-group which contains at least one pseudo-Anosov element whose attracting fixed point is a minimal complete geodesic lamination.*

(2) *μ has finite entropy and finite logarithmic moment for the action of $\mathcal{MCG}$ on the curve graph.*

Then the Poisson boundary of the walk can be realized as a $\mathcal{MCG}$-invariant measure class on the space $\partial\mathcal{PC}(S)$ of minimal complete geodesic laminations.

Proof. Since by assumption and Corollary 6.8, the semigroup generated by the support of μ is non-elementary and contains at least one pseudo-Anosov element φ which acts as a hyperbolic isometry on $\mathcal{PC}(S)$, this group is non-elementary as a group of isometries acting on $\mathcal{PC}(S)$ (apply Proposition 6.7 and Corollary 6.8 to conjugates of φ).

Theorem 1 of [MT18] now shows that for any vertex $x \in \mathcal{PC}(S)$, almost every sample path $(\omega_n x) \subset \mathcal{PC}(S)$ converges to a point $\omega_+ \in \partial\mathcal{PC}(S)$. The resulting hitting measure ν is non-atomic, and it is the unique μ-stationary measure on $\partial\mathcal{PC}(S)$.

The same construction also applies for the action of the random walk on $\mathcal{CG}(S)$. As $\partial\mathcal{PC}(S) \subset \partial\mathcal{GG}(S)$ by Proposition 6.7, the μ-stationary measure ν on $\partial\mathcal{PC}(S)$ also can be viewed as a μ-stationary measure on $\partial\mathcal{CG}(S)$. By uniqueness, ν equals the hitting measure of the random walk on $\mathcal{CG}(S)$.

Now the action of $\mathcal{MCG}$ on $\mathcal{CG}(S)$ is acylindrical [Bw08] and therefore by Theorem 1.5 of [MT18], the Poisson boundary of $(\mathcal{MCG}, \mu)$ equals the Gromov boundary $\partial\mathcal{CG}(S)$ of $\mathcal{CG}(S)$, equipped with the hitting measure ν. But this hitting measure is just the unique stationary measure for the action of $\mathcal{MCG}$ on the boundary of the principal curve graph, which completes the proof of the theorem. □

7. The regular Roller boundary and geodesic laminations

In this section we resume the discussion of the large scale geometry of the CAT(0) cube complex $\mathcal{C}(E)$ constructed in Section 5 and prove the second part of Theorem 1 from the introduction. When discussing general properties of CAT(0) cube complexes, we follow [FLM18] which contains an excellent summary of the statements we need.

We begin with defining the Roller boundary of a CAT(0) cube complex X. To this end define a *midcube* of a cube $[0, 1]^n$ to be the preimage of $1/2$ under one of the i coordinate projections. A *hyperplane* in X is a CAT(0)-convex subset whose intersection with each cube is either a midcube or empty. The complement of a hyperplane in X has two connected components. The intersection of one of these components with the vertex set of X is called a *halfspace* h. The intersection of the second component with the vertex set of X is the *complementary half-space* h^*. Let $\mathfrak{H}$ be the collection of all halfspaces in X.

For a vertex $x \in X$ consider the collection $U_x = \{h \in \mathfrak{H} \mid x \in h\}$. The vertex x is uniquely determined by the set U_x. This yields an embedding $X \to 2^{\mathfrak{H}}$ obtained by $x \to U_x$.

Definition 7.1. The *Roller compactification* $\overline{X}$ of X is the closure of X in $2^{\mathfrak{H}}$. The *Roller boundary* is then $\partial X = \overline{X} - X$.

Note that for any $\xi \in \partial X$ we can consider the set U_ξ of half-spaces defined by ξ, an analog of the set of half-spaces containing a vertex of X. The set U_ξ fulfills the axioms of an ultrafilter on $\mathcal{H}$:

(1) For all $h \in \mathcal{H}$, either $h \in U_\xi$ or $h^* \in U_\xi$ but not both.
(2) If $h \in U_\xi$ and $h' \supset h$ then $h' \in U_\xi$.

Two half-spaces h, k are called *transverse* if the four intersections $h \cap k$, $h \cap k^*$, $h^* \cap k, h^* \cap k^*$ are all nonempty. We use this terminology to define the *regular Roller boundary* (Section 5.3 of [FLM18]).

Definition 7.2. Two half-spaces h and k are *strongly separated* if there is no half-space which is transverse to both h, k. Two hyperplanes are *strongly separated* if this holds true for the half-spaces they bound.

Definition 7.3. A point $\xi \in \partial X$ is called *regular* if for every $h_1, h_2 \in U_\xi$ there is $k \in U_\xi$ such that $k \subset h_1 \cap h_2$ and that k is strongly separated from both h_1, and h_2. The *regular Roller boundary* is the closure in ∂X of the set $\partial_r X$ of regular points.

By definition, the regular Roller boundary of X is a compact totally disconnected topological space. The goal of this section is to show

Theorem 7.4. *There is a natural homeomorphism from the space $\mathcal{CL}$ of complete geodesic laminations on S, equipped with the Hausdorff topology, onto the regular Roller boundary of $\mathcal{C}(E)$.*

The idea of the proof of Theorem 7.4 is to relate hyperplanes in $\mathcal{C}(E)$ to the principal curve graph $\mathcal{PC}(S)$ of S. To this end note that a hyperplane H in $\mathcal{C}(E)$ is determined by an edge in E and hence by a large branch b in a train track $\eta \in \mathcal{C}(E)$ and the choice of a right or left split of b, viewed as an edge in E.

Any hyperplane H in a CAT(0) cube complex is parallel to two *combinatorial hyperplanes* which are subcomplexes of the cube complex. As the graph E is directed, for a hyperplane $H \subset \mathcal{C}(E)$ determined by a train track η and a large branch b in η we can distinguish the *negative* combinatorial hyperplane H_- consisting of train tracks which contain the large branch b. The parallel combinatorial hyperplane consists of train tracks which can be obtained from a train track in H_- by a single right (or left) split at the distinguished large branch and is called the *positive* combinatorial hyperplane H_+ of H. Note that if both the right and the left split of η are complete, then H_- is the negative combinatorial hyperplane of two distinct (and in fact disjoint) positive combinatorial hyperplanes in $\mathcal{C}(E)$.

Denote by $d_{\mathcal{PC}}$ the distance in the principal curve graph. We show

Lemma 7.5. *There is a number $\chi > 0$ with the following property. Let $H_- \subset \mathcal{C}(E)$ be a negative combinatorial hyperplane and let $\eta, \xi \in H_-$ be such that η is splittable to ξ. Then for any vertex cycles c of η, d of ξ we have $d_{\mathcal{PC}}(c, d) \le \chi$.*

Proof. Let b be a large branch of a train track η. The large branch b determines a negative hyperplane H_- in $\mathcal{C}(E)$. If $\xi \in H_-$ can be obtained from η by a splitting sequence, then such a splitting sequence does not contain a split at b. In particular, b is a large branch in ξ.

Let ξ' be obtained from ξ by a right (or left) split at b, chosen in such a way that ξ' is complete. Let η' be obtained from η by a right (or left) split at b. Then η' is transversely recurrent and splittable to the complete train track ξ' and hence it is complete. The train track ξ'' obtained from ξ' by removal of the small branch of the split is carried by the train track η'' obtained from η' by removal of the small branch of the split.

Now the complementary component of η'' containing the diagonal of the split which connects η to η' is neither a trigon nor a once punctured monogon. Since ξ'' is carried by η'', the distance in the principal curve graph between a vertex cycle of η'' and ξ'' is at most one. But the distance in the curve graph between a vertex cycle of η and a vertex cycle of η'' is uniformly bounded, and the same holds true for the distance in the curve graph between a vertex cycle of ξ and a vertex cycle of ξ''. Hence by the triangle inequality, the distance in the principal curve graph between a vertex cycle of η and a vertex cycle of ξ is uniformly bounded. This shows the lemma. □

The next example shows that hyperplanes give a finer information than a decomposition of S into subsurfaces.

Example 7.6. Let η be a non-complete recurrent transversely recurrent train track which decomposes S into trigons and one sixgon. This train track can be extended to a complete train track ξ by inserting in the interior of the sixgon the union of a large branch e and 4 small branches, each with one endpoint on a different side of the sixgon. Assuming that ξ is a vertex of the cube complex $\mathcal{C}(E)$, any splitting sequence ξ_i starting at $\xi_0 = \xi$ which is induced by a splitting sequence of η is contained in the negative combinatorial hyperplane defined by the large branch e. This also extends to train tracks obtained from η by a collapse at a small branch not contained in the interior of the sixgon as long as this train track is still contained in $\mathcal{C}(E)$.

We use Lemma 7.5 to show.

Corollary 7.7. *Let c, d be any two vertex cycles of train tracks η, ξ contained in the same negative combinatorial hyperplane of $\mathcal{C}(E)$. Then $d_{\mathcal{PC}}(c, d) \leq 2\chi$ where $\chi > 0$ is as in Lemma 7.5.*

Proof. Let H_- be a negative combinatorial hyperplane in $\mathcal{C}(E)$ and let $\xi, \eta \in H_-$. Consider the train track $\Theta_-(\xi, \eta) \in E$ constructed in Lemma 5.3. By construction, $\Theta_-(\xi, \eta)$ is splittable to ξ, η. Furthermore, if b is any large branch of $\Theta_-(\xi, \eta)$ then there are three possibilities.

(1) A splitting sequence connecting $\Theta_-(\xi, \eta)$ to both ξ, η does not contain a split at b. We call b *neutral* in this case.

(2) Up to exchanging ξ, η, a splitting sequence connecting $\Theta_-(\xi, \eta)$ to ξ contains a split at b, but this is not true for a splitting sequence connecting $\Theta_-(\xi, \eta)$ to η.
(3) A splitting sequence connecting $\Theta_-(\xi, \eta)$ to ξ contains a right (or left) split at b, and a splitting sequence connecting $\Theta_-(\xi, \eta)$ to η contains a left (or right) split at b.

Since ξ, η are contained in the same negative combinatorial hyperplane H_-, there exists a neutral large branch b in $\Theta_-(\xi, \eta)$ defining the hyperplane, and $\Theta_-(\xi, \eta) \in H_-$. Since $\Theta_-(\xi, \eta)$ is splittable to both ξ, η, Proposition 7.5 shows that the distance in the principal curve graph between a vertex cycle of $\Theta_-(\xi, \eta)$ and a vertex cycle of η is at most χ, and the same holds true for the distance in the principal curve graph between a vertex cycle of $\Theta_-(\xi, \eta)$ and a vertex cycle of ξ. The corollary now follows from the triangle inequality. □

Fix a vertex $\tau \in E$. Associate to a vertex $\eta \in E$ the one-Lipschitz function

$$h_\eta : E \to \mathbb{R}, \quad h_\eta(x) = d(\eta, x) - d(\eta, \tau)$$

where d is the *combinatorial distance* on E, that is, the distance in the graph E which is the one-skeleton of $\mathcal{C}(E)$. The *combinatorial horoboundary* of E is the boundary of the closure of the set of these functions in the space of all one-Lipschitz functions on E. It does not depend on the choice of the basepoint τ.

The following is Proposition 6.20 of [FLM18].

Proposition 7.8. *There exists a natural homeomorphism $h : \xi \to h_\xi$ of the Roller boundary $\partial\mathcal{C}(E)$ of $\mathcal{C}(E)$ onto the combinatorial horoboundary of $\mathcal{C}(E)$. If for $\xi \in \partial\mathcal{C}(E)$ and $x \in E$ we denote by $m(x, \xi)$ the median between x, τ, ξ, then*

$$h_\xi(x) = d(m(x, \xi), x) - d(m(x, \xi), \tau).$$

Recall from Section 5 that for any two points $\eta, \zeta \in E$, there exists a combinatorial geodesic in $\mathcal{C}(E)$ connecting η to ζ which is the composition of the inverse of a splitting sequence connecting $\Theta_-(\eta, \zeta)$ to η and a splitting sequence connecting $\Theta_-(\eta, \zeta)$ to ζ. We call such a combinatorial geodesic a *tail splitting sequence*. We next compute in more detail the interval $\mathcal{I}(\eta, \zeta)$ between η, ζ. To this end define a *full rearrangement* of a full splitting sequence $(\eta_i) \subset E(\eta, \lambda)$ to be a splitting sequence $\hat{\eta}_i$ beginning at $\eta_0 = \hat{\eta}_0 = \eta$ so that for each j there is i such that η_j is splittable to $\hat{\eta}_i$.

Let as before $\mathcal{CL}$ be the space of complete geodesic laminations, equipped with the Hausdorff topology. This is a compact totally disconnected topological space [H09]. Proposition 7.8 is used to establish the following

Lemma 7.9. *Let $\lambda \in \mathcal{CL}$ be arbitrary and let $(\eta_i) \subset E(\eta, \lambda)$ be a full splitting sequence starting at $\eta_0 = \eta$. Then (η_i) determines a point $\Psi(\lambda) \in \partial\mathcal{C}(E)$, and $\mathcal{I}(\eta, \Psi(\lambda)) = E(\eta, \lambda)$. Any full rearrangement of η_i defines the same point $\Psi(\lambda)$. If $\zeta_j \subset E(\eta, \lambda)$ is a splitting sequence which is not a full rearrangement of η_j, then there exists a negative combinatorial hyperplane H_- containing ζ_j for all large enough j.*

Proof. If $\zeta \in E$ is obtained from η by a splitting sequence, then we have $\mathcal{I}(\eta,\zeta) = E(\eta,\zeta)$. This implies the following. Define the interval of a sequence η_i of λ-splits to be $\cup_i \mathcal{I}(\eta_0,\eta_i)$. If the splitting sequence η_i is a full rearrangement of a full sequence of λ-splits, then this interval equals the convex subspace $E(\eta,\lambda)$ of E.

Any one sided infinite splitting sequence in E is a combinatorial geodesic in $\mathcal{C}(E)$ and hence it defines a point in the combinatorial horoboundary of $\mathcal{C}(E)$. If $\zeta_i \subset E(\eta,\lambda)$ is a splitting sequence which is not a full rearrangement of the full splitting sequence (η_i), then there exists some j so that η_j is not splittable to any of the train tracks ζ_i. As a consequence, there is some $\xi \in E(\eta,\lambda)$ which is splittable to ζ_i for some i, and there is a large branch b of ξ such that a splitting sequence connecting ξ to ζ_j for any large enough j does not contain a split at b. But then for all large enough j the train track ζ_j is contained in the negative combinatorial hyperplane of the CAT(0) cube complex $E(\eta,\lambda)$ defined by the large branch b, and the train track obtained by splitting ξ at b once is not contained in the interval defined by the sequence ζ_i. □

Before we complete the proof of Theorem 7.4 we introduce some more terminology.

Definition 7.10. Let λ be a minimal geodesic lamination on S. The *subsurface of S filled by* λ is the intersection S_λ of all subsurfaces of S which contain λ.

Note that if λ is a simple closed curve, then the subsurface S_λ of S filled by λ is an annulus with core curve λ. More generally, S_λ is characterized by the property that S_λ contains λ and that any simple closed non-peripheral curve in S_λ intersects λ.

The following statement is main step towards the proof of Theorem 7.4.

Proposition 7.11. *There exists a continuous injective map from the space $\mathcal{CL}$ of complete geodesic laminations into the Roller boundary of $\mathcal{C}(E)$.*

Proof. By Lemma 7.9, a point in $\mathcal{CL}$ determines a point in the combinatorial horoboundary of $\mathcal{C}(E)$. Thus by Proposition 7.8, there is a map $\Psi : \mathcal{CL} \to \partial\mathcal{C}(E)$. We have to show that Ψ is injective and continuous, and that its image is the regular Roller boundary of $\mathcal{C}(E)$.

To show that Ψ is injective, let $\lambda \neq \mu \in \mathcal{CL}$; we have to show that $\Psi(\lambda) \neq \Psi(\mu)$. This is equivalent to stating that for some $\eta \in E$, the intervals $E(\eta,\lambda)$ and $E(\eta,\mu)$ are distinct.

It follows easily from the above discussion that this is the case if one of λ,μ is the repelling fixed point of the pseudo-Anosov mapping class used in the construction of E, so assume without loss of generality that λ,μ are both carried by the same train track $\eta \in E$. Let $\lambda^1,\dots,\lambda^k$ be the minimal components of λ and write $\hat{\lambda} = \cup_i \lambda^i$.

By Corollary 2.4.3 of [PH92], there exists a splitting and collision sequence which connects τ to a train track $\xi = \cup_i \xi^i$ (disjoint union) where ξ_i is contained in the subsurface S^i filled by λ^i and carries λ^i. As a consequence of the results in Section

3, τ is splittable to a complete train track ζ which carries λ and such that ξ^i is a subtrack of ζ for each i. Furthermore, we conclude the following.

For each i, the train track ξ^i admits an infinite sequence of λ^i-splits, and a sequence of vertex cycles on these train tracks converge in the coarse Hausdorff topology to λ^i. By Section 3, these splitting sequences induce splitting sequences of ζ contained in $E(\tau,\lambda)$. As a consequence, if there is a minimal component λ^i of λ which is not a minimal component of μ, then $E(\tau,\lambda) \neq E(\tau,\mu)$. Thus by symmetry, if $E(\tau,\lambda) = E(\tau,\mu)$ then the minimal components of λ,μ coincide.

To summarize, if $E(\tau,\lambda) = E(\tau,\mu)$ and if $\lambda \neq \mu$, then λ and μ differ by finitely many isolated leaves. Now consider as before a train track $\zeta \in E(\tau,\lambda)$ which carries λ and has the additional property that ζ contains a subtrack ξ which is a disjoint union of train tracks $\xi^1,\dots,\xi^k$ contained in the surfaces S^i filled by the components ξ^i and which carry the laminations λ^i. Since λ is complete by assumption, every branch of $\zeta - \xi$ is passed through by an isolated leaf of λ. Since an isolated leaf of λ spirals about some minimal components of λ and hence passes through any branch of $\zeta - \xi$ only finitely many times, the sum over all branches a of $\zeta - \xi$ of the number of times the branch a is passed through by a leaf of λ is finite, and it strictly decreases if ζ is split at a large branch in $\zeta - \xi$ to a train track ζ' which carries λ.

As any train track carrying a complete geodesic lamination μ can be split at any large branch to another complete train track which carries λ, this means that ζ can be split to a train track $\sigma \in E(\tau,\lambda)$ which contains ξ as a subtrack and such that $\sigma - \xi$ does not contain a large branch. Thus it now suffices to show that any complete extension of $\cup_i \lambda^i$ which is carried by σ coincides with λ.

However, this follows from the fact that as $\sigma - \xi$ does not contain a large branch, any two simple arcs α_1,α_2 carried by $\sigma - \xi$ are disjoint up to homotopy. Namely, let α_1,α_2 be two oriented simple arcs carried by $\sigma - \xi$ which pass through the same branch of $\sigma - \xi$ in the same direction. Let us assume that for the given orientation, their first intersection point on $\sigma - \xi$ is a switch $v \in \sigma - \xi$. Then the large half-branch b of $\sigma - \xi$ which is incident on v is traveled through by both α_1,α_2 as oriented arcs starting at v. Since $\sigma - \xi$ does not contain a large branch, it does not contain a large embedded trainpath starting at v [PH92], see also Section 3. Thus the trainpaths beginning at v which are defined by α_1,α_2 and a carrying map coincide up to their first intersection with ξ [PH92]. In particular, the arcs α_1,α_2 can be homotoped to be disjoint. Now if $\mu \neq \lambda$ and if the minimal components of μ coincide with the minimal components of λ, then there are isolated leaves of μ,λ which have essential intersections. Then μ can not be carried by σ. This completes the proof that the map Ψ is injective.

Since both $\mathcal{CL}$ and $\partial\mathcal{C}(E)$ are compact separable Hausdorff spaces, to show that the map Ψ is continuous it suffices to show the following. Let $\lambda_i \subset \mathcal{CL}$ be a sequence which converges in the Hausdorff topology to a lamination λ; then $\Psi(\lambda_i) \to \Psi(\lambda)$.

To this end let $(\tau_i) \subset E$ be a full splitting sequence converging to λ. Then for each i, the set $\mathcal{CL}(\tau_i)$ of all complete geodesic laminations carried by τ_i is a neighborhood of λ in $\mathcal{CL}$ [H09]. Thus for any $m > 0$ there exists a number $i(m) > 0$ such that for $i \geq i(m)$, the train track τ_m carries λ_i and hence it is contained in the interval $E(\tau, \lambda_i)$ defined by λ_i. But this means that the points $\Psi(\lambda_i) \in \partial\mathcal{C}(E)$ converge to $\Psi(\lambda)$ showing continuity.

Since both $\mathcal{CL}$ and $\partial\mathcal{C}(E)$ are compact, the map Ψ is closed. As Ψ is also injective, it is open as a map onto its image and hence it is a homeomorphism onto its image. This completes the proof that $\Psi : \mathcal{CL} \to \partial\mathcal{C}(E)$ is an embedding and shows the proposition. □

Proof of Theorem 7.4. By Proposition 7.11, we are left with showing that the image of the continuous injective map $\Psi : \mathcal{CL} \to \partial\mathcal{C}(E)$ equals the regular Roller boundary of $\mathcal{C}(E)$.

We begin with showing the following

Claim: The image under Ψ of a minimal complete geodesic lamination is a regular point in $\partial\mathcal{C}(E)$.

Since the subset of minimal complete geodesic laminations is dense in $\mathcal{CL}$, by continuity this then implies that the subspace $\Psi(\mathcal{CL}) \subset \partial\mathcal{C}(E)$ is contained in the regular Roller boundary of $\mathcal{C}(E)$.

To see that the claim holds indeed true recall that splitting sequences contained in E are combinatorial geodesics in $\mathcal{C}(E)$. Furthermore, if φ is the pseudo-Anosov mapping class used in the construction of the graph E, then the infinite cyclic subgroup of $\mathcal{MCG}$ generated by φ acts on $\mathcal{C}(E)$ as an infinite cyclic group of combinatorial isometries. This group preserves a φ-invariant biinfinite splitting sequence. Therefore if the forward ray of this splitting sequence defines a point in the regular Roller boundary of $\mathcal{C}(E)$, then by symmetry and the definition of the regular Roller boundary via strongly separated hyperplanes, the same holds true for the backward ray. Thus to show the claim it suffices to show the following.

Let λ be a minimal complete geodesic lamination different from the minimal complete geodesic lamination which is the repelling fixed point of the action of φ on $\mathcal{CL}$. Let $(\eta_i)_{i\geq 0} \subset E$ be a full splitting sequence consisting of train tracks which carry λ. For each i let H_i be a negative combinatorial hyperplane containing η_i. Then there is a sequence $i_j \to \infty$ such that for all j, the hyperplanes H_{i_j} and $H_{i_{j+1}}$ are strongly separated.

To this end let $\chi \geq 2$ be as in Corollary 7.7. Let i_j be a subsequence of the positive integers such that $d_{\mathcal{PC}}(c_{i_j}, c_{i_{j+1}}) \geq 8\chi$ for any vertex cycle c_{i_j} of η_{i_j} and any vertex cycle $c_{i_{j+1}}$ of $\eta_{i_{j+1}}$. Such a sequence exists by Proposition 6.6. We claim that there is no negative combinatorial hyperplane Q which intersects both $H_{i_j}, H_{i_{j+1}}$.

Namely, assume to the contrary that there is a negative combinatorial hyperplane Q which intersects H_{i_j} and $H_{i_{j+1}}$ in points ζ_1, ζ_2. Let $c(\zeta_i)$ be a vertex cycle of ζ_i. By Corollary 7.7, applied to the hyperplane Q, we have $d_{\mathcal{PC}}(c(\zeta_1), c(\zeta_2)) \leq 2\chi$, and similarly we obtain $d_{\mathcal{PC}}(c(\zeta_1), c_{i_1}) \leq 2\chi$, $d_{\mathcal{PC}}(c(\zeta_2), c_{i_{j+1}}) \leq 2\chi$. This contradicts the choice of the train tracks $\eta_{i_j}, \eta_{i_{j+1}}$.

As a consequence, the combinatorial geodesic $(\eta_i) \subset \mathcal{C}(E)$ determines an infinite descending chain of pairwise strongly separated half-spaces in $\mathcal{C}(E)$ and hence by Proposition 5.10 of [FLM18], it determines a regular point in the Roller boundary of $\mathcal{C}(E)$. On the other hand, by Lemma 6.5, the quasi-geodesic in $\mathcal{PC}(S)$ which is the image of the splitting sequence under the map Υ is of infinite diameter and hence converges to a point in the Gromov boundary of $\mathcal{PC}(S)$. This point is just the minimal complete geodesic lamination λ. As a consequence, the image under Ψ of a minimal complete geodesic lamination is a regular point of the Roller boundary of $\mathcal{C}(E)$.

We are left with showing that $\Psi(\mathcal{CL})$ coincides with the regular Roller boundary. Since $\Psi(\mathcal{CL})$ is a closed subset of $\partial\mathcal{C}(E)$, to this end it suffices to show that a regular point $\xi \in \partial_r\mathcal{C}(E)$ is contained in $\Psi(\mathcal{CL})$. Now by definition of the regular Roller boundary, if $\xi \in \partial_r\mathcal{C}(E)$ then ξ it can be approximated by a nested strongly separated sequence of hyperplanes. On the other hand, any combinatorial geodesic in $\mathcal{C}(E)$ different from a geodesic parallel to the axis of the isometry φ is a tail splitting sequence and therefore all we need to show is that if $\xi \in \partial_r\mathcal{C}(E)$, then a splitting sequence defining ξ is a full rearrangement of a full splitting sequence. But this is an immediate consequence of Lemma 7.9. □

Remark 7.12. Theorem 7.4 does *not* state that the set of regular points in the Roller boundary of $\mathcal{C}(E)$ corresponds to the minimal complete geodesic laminations, and we expect that there are additional regular points.

8. A small boundary for $\mathcal{MCG}$

A proper Cat(0) cube complex X not only has a Roller boundary, but also has a CAT(0) boundary $\partial_\angle X$, called *geometric boundary* in the sequel. Given a basepoint $x \in X$, the geometric boundary is the space of geodesic rays for the CAT(0) metric starting at x, equipped with the topology of uniform convergence on compact sets. It does not depend on the choice of the basepoint x. The boundary $\partial_\angle X$ is compact, and the same holds true for $\overline{X} = X \cup \partial_\angle X$. Furthermore, the geometric boundary $\partial_\angle X$ of X is a $\mathcal{Z}$*-set* for the compactification $\overline{X}$ of X. Thus if Γ is any torsion free group acting properly and cocompactly on X, then the cohomology of ∂X can be used to obtain information on the cohomology of Γ with coefficients in the group ring of Γ [B96].

For the CAT(0) cube complex $\mathcal{C}(E)$, we can try to analyze the geometric boundary and use it to obtain information on the large scale geometry of $\mathcal{MCG}$. One difficulty in this endeavor is that there are distinct geodesic rays in $\mathcal{C}(E)$ whose images under the natural embedding $\mathcal{C}(E) \to \mathcal{TT}$ are of uniformly bounded distance in $\mathcal{TT}$.

Nevertheless, we shall use the geometric information collected so far to make a guess for a boundary of $\mathcal{MCG}$ which can be thought of as a quotient of the geometric boundary of $\mathcal{C}(E)$, and show that it is indeed a small boundary for $\mathcal{MCG}$. This leads to the proof of Theorem 2 from the introduction.

Recall that the *join* $X_1 * X_2$ of two topological spaces X_1, X_2 is the quotient $X_1 \times X_2 \times [0,1]/\sim$ where the equivalence relation $\sim$ collapses $X_1 \times X_2 \times \{0\}$ to X_1 and collapses $X_1 \times X_2 \times \{1\}$ to X_2. For example, the join $S^0_1 * S^0_2$ of two 0-spheres is the circle S^1, thought of as a union of four intervals glued at the endpoints, where each interval has one endpoint in S^0_1 and the second endpoint in S^0_2. The join of two spaces X, Y contains an embedded copy of X, Y.

Let $\mathcal{OG}(S)$ be the *oriented curve complex* of S. This is the complex whose vertices are isotopy classes of oriented simple closed curves in S and where two such vertices are connected by an edge (of length 1) if they can be realized disjointly and are not homotopic up to orientation. Thus any simple closed curve in S defines two distinct vertices in $\mathcal{OG}(S)$, and these vertices are not connected by an edge. Furthermore, we require that any $k \geq 2$ oriented disjoint simple closed curves span a simplex. The union of these simplices corresponding to this collection of curves equipped with all combinations of orientations is a sphere of dimension $k-1$. Note that a point in $\mathcal{OG}(S)$ can be viewed as a formal linear combination $\sum_{i=1}^k a_i \lambda_i$ where for some $k \geq 1$, $\lambda_1, \dots, \lambda_k$ are pairwise disjoint oriented simple closed curves, where $a_i > 0$ for all i and $\sum_i a_i = 1$.

Remark 8.1. If we choose the length of the edges of the oriented curve complex to be $\pi/2$, then this is consistent with the idea that the oriented curve complex can be thought of as being contained in the Tits boundary of $\mathcal{MCG}$, equipped with the angular length metric which identifies each sphere with a sphere of constant curvature one.

A simple closed curve c is the core curve of an embedded annulus $A(c) \subset S$. The "curve graph" $\mathcal{CG}(A(c))$ of the annulus $A(c)$ is a graph of isotopy classes of arcs connecting the two boundary components and whose endpoints are allowed to move freely in the complement of a fixed point on each of the two boundary circles. The curve graph of $A(c)$ is a simplicial line. If α is a fixed vertex of $\mathcal{CG}(A(c))$, then any other isotopy class of arcs can be represented by an arc which is the image of α by a Dehn multi-twist about c. With this viewpoint, the choice of an orientation of c can be thought of as a spiraling direction about c. However, the distinction between a positive and a negative Dehn twists about c only depends on the orientation of S but not on the orientation of c. In the sequel we denote by c^+ the point in the Gromov boundary of $\mathcal{CG}(A(c))$ (which consists of two points) defined by iteration of positive Dehn twists about c, and we denote by c^- the point in the Gromov boundary of $\mathcal{CG}(A(c))$ defined by iteration of negative Dehn twists about c. Write $\mathcal{X}(c) = \{c^+, c^-\}$. It will be convenient in the sequel to think about $\mathcal{X}(c)$ as two points in the oriented curve complex, with the same underlying curve.

If S_0 is a subsurface of S different from a pair of pants or an annulus, then we denote its (non-oriented) curve complex by $\mathcal{CG}(S_0)$. Note that $\mathcal{CG}(S_0)$ is a subcomplex of the curve complex of S unless S_0 is a one-holed torus or a four punctured sphere; in this case, the edges of $\mathcal{CG}(S_0)$ are not edges in $\mathcal{CG}(S)$. The curve complex of S_0 is hyperbolic and hence it has a *Gromov boundary* $\partial\mathcal{CG}(S_0)$. As a set, the Gromov boundary $\partial\mathcal{CG}(S_0)$ is the set of all minimal filling geodesic laminations on S_0.

There is a natural topology on the union $\overline{\mathcal{CG}}(S_0)$ of $\mathcal{CG}(S_0)$ with its Gromov boundary, called the *coarse Hausdorff topology.* With respect to this topology, the subspace $\mathcal{CG}(S_0)$, equipped with its natural topology, is an open dense subset. The topology is metrizable, and a sequence $\lambda_i \subset \mathcal{CG}(S_0) \cup \partial\mathcal{CG}(S_0)$ converges in this topology to $\lambda \in \partial\mathcal{CG}(S_0)$ if and only if the limit of any converging subsequence of λ_i in the Hausdorff topology on compact subsets of S contains λ as a sublamination [H06]. Define $\mathcal{X}(S_0) = \partial\mathcal{CG}(S_0)$, equipped with the topology as a subset of $\overline{\mathcal{CG}}(S_0)$. If S_0 is a pair of pants, then we define $\mathcal{X}(S_0) = \emptyset$.

If $S_1, \dots, S_k$ are *disjoint* connected subsurfaces of S (we allow that they share boundary components, and annuli about such boundary components may be included in the list), then we define

$$\mathcal{X}(\cup_i S_i) = \partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k)$$

to be the join of the spaces $\mathcal{X}(S_i) = \partial\mathcal{CG}(S_i)$. For example, if $S_1 \subset S$ is a subsurface which is the complement of a non-separating simple closed curve c, then $\mathcal{X}(S_1 \cup A(c)) = \partial\mathcal{CG}(S_1) * \{c^+, c^-\}$. A point in $\mathcal{X}(S_1 \cup \cdots \cup S_k)$ can be viewed as a formal linear combination $\xi = \sum_i a_i\xi_i$ where $\xi_i \in \partial\mathcal{CG}(S_i)$, $a_i \geq 0$ for all i and, furthermore, $\sum_i a_i = 1$. The union $\hat{\xi} = \cup_{a_i>0}\xi_i$ is a geodesic lamination with minimal components ξ_i, and ξ can be viewed as a weighted (and partially labeled if there are simple closed curve components) geodesic lamination. For all $u \leq k$ there is an inclusion $\mathcal{X}(S_1 \cup \cdots \cup S_u) \subset \mathcal{X}(S_1 \cup \cdots \cup S_k)$ which is a topological embedding.

A collection $S_1, \dots, S_k$ of disjoint connected subsurfaces of S is called *maximal* if $S - \cup_i S_i = \emptyset$. By convention, this means that if c is a boundary component of one of the surfaces S_i, then $A(c)$ is contained in the collection. Any collection of disjoint connected subsurfaces of S is contained in a maximal collection of such subsurfaces, however this maximal collection is in general not unique. Note that there is a canonical maximal collection which is comprised of the surfaces S_i, the annuli $A(c)$ where c runs through all boundary components of $\cup_i S_i$ which are not already contained in that list and all connected components of $S - \cup_i S_i$.

Define

$$\mathcal{X} = \cup\mathcal{X}(S_1 \cup \cdots \cup S_k)/\sim$$

where the union is over all collections of disjoint subsurfaces $S_1, \dots, S_k$ of S. The equivalence relation $\sim$ identifies two points $\sum_i a_i\xi_i$ and $\sum_j b_j\zeta_j$ if they coincide as weighted labeled geodesic laminations. Thus a point in $\mathcal{X}$ is nothing else but a formal sum $\sum_{i=1}^k a_i\xi_i$ where $a_i > 0, \sum_i a_i = 1$, where $\xi_1, \dots, \xi_k$ are pairwise disjoint minimal geodesic laminations on S and where every simple closed curve component of this collection is in addition equipped with a label $\pm$. Note that the oriented curve complex $\mathcal{OG}(S)$ of S naturally is a subset of $\mathcal{X}$, and the same holds true for its Gromov boundary (which is just the Gromov boundary $\partial\mathcal{CG}(S)$ of the non-oriented curve graph).

The oriented curve graph of S is connected, and any non-filling geodesic lamination, that is, a geodesic lamination which is disjoint from some simple closed curve, is disjoint from some vertex of $\mathcal{OG}(S)$. Thus if we equip $\mathcal{X} - \partial\mathcal{CG}(S)$ with the topology of a simplicial complex whose edges are the joins of two disjoint labeled geodesic laminations, then this complex is connected. As a consequence, the *set* $\mathcal{X}$ can be equipped with a topology which coincides with the topology of a (non-locally finite) simplicial complex on $\mathcal{X} - \partial\mathcal{CG}(S)$ and is such that each point in $\partial\mathcal{CG}(S)$ is isolated. We write $\mathcal{X}_T$ for $\mathcal{X}$ equipped with this topology (having the Tits boundary of a CAT(0) space as guidance). From this description, we obtain

Lemma 8.2. *The mapping class group $\mathcal{MCG}$ of S acts on $\mathcal{X}_T$ as a group of automorphisms.*

Proof. The mapping class group acts on the oriented curve graph of S as a group of simplicial automorphisms, and this action extends to an action on the space of formal sums of weighted disjoint minimal geodesic laminations preserving weight and disjointness. Furthermore, it acts on $\partial\mathcal{CG}(S)$ as a group of transformations. Since the topology on $\mathcal{X}_T$ is the topology of a disconnected simplicial complex, constructed from the curve graphs of subsurfaces, the lemma follows. □

The Tits boundary of a CAT(0) space X can be viewed as the geometric boundary (that is, the CAT(0) boundary) of X, equipped with a topology which in general is finer than the geometric topology. Our goal is to equip the set $\mathcal{X}$ with a topology which is coarser than the Tits topology so that for this topology, $\mathcal{X}$ becomes a compact space which indeed defines a small compactification of $\mathcal{MCG}$.

To achieve this goal we use markings of subsurfaces of the surface S (see Section 2). There is a natural way to equip the set of all markings on S with the structure of a locally finite connected graph on which $\mathcal{MCG}$ acts properly and cocompactly.

Choose a marking μ on S as a basepoint for $\mathcal{MCG}$. For every subsurface S_0 of S which is distinct from a pair of pants, this marking determines a coarsely well defined marking $\mu(S_0)$ of S_0. Namely, the intersection of each marking curve c with S_0 either is a simple closed curve contained in S_0, a collection of pairwise disjoint arcs with endpoints on the boundary of S_0 or empty. The union of these intersections over all curves from the marking decompose S_0 into simply connected regions and hence coarsely defines a marking $\mu(S_0)$ of S_0, called the *subsurface projection* of μ [MM00]. Here coarse definition means that the construction depends on choices, but any two choices give rise to markings which are uniformly close in the marking graph of S_0, independent of the subsurface S_0.

Now let $S = \cup_{i=1}^k S_i$ be a collection of pairwise disjoint subsurfaces of S. By the above, each of the surfaces S_i is equipped with a coarsely well defined marking $\mu(S_i)$. Let x_i be one of the marking curves of $\mu(S_i)$. We then obtain a based product space

$$(\mathcal{CG}(\cup_i S_i), x) = (\mathcal{CG}(S_1) \times \cdots \times \mathcal{CG}(S_k), x)$$

where the basepoint $x = (x_1, \ldots, x_k)$ is the product of the coarsely well defined basepoints in $\mathcal{CG}(S_i)$.

For any subsurface S_0 of S, denote by $\mathrm{pr}_{S_0} : \mathcal{CG}(S) \to \mathcal{CG}(S_0)$ the (coarsely well defined) subsurface projection. Define a topology on the union

$$\mathcal{Y}(\cup_i S_i) = \mathcal{CG}(\cup_i S_i) \cup \partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k) = \mathcal{CG}(\cup_i S_i) \cup \mathcal{X}(\cup_i S_i),$$

by the following requirement. The product space $\mathcal{CG}(\cup_i S_i)$ is equipped with the product topology, and the topology on the subspace $\mathcal{X}(\cup_i S_i)$ is the topology described above as a join of the Gromov boundaries of the curve graphs of S_i. Furthermore, a sequence of points $(y_1^j, \dots, y_k^j) \in \mathcal{CG}(\cup_i S_i)$ converges to $\sum_i a_i \xi_i \in \partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k)$ if the following two conditions are fulfilled.

(1) For each i with $a_i > 0$, the components y_i^j converge as $j \to \infty$ to ξ_i in the coarse Hausdorff topology (and hence they converge in $\overline{\mathcal{CG}}(S_i)$ to ξ_i). In particular, we have $d_{\mathcal{CG}(S_i)}(y_i^j, x_i) \to \infty$.
(2) Assume without loss of generality that $a_1 > 0$. Then for all $i \geq 2$ we have

$$\frac{d_{\mathcal{CG}(S_i)}(y_i^j, x_i)}{d_{\mathcal{CG}(S_1)}(y_1^j, x_1)} \to \frac{a_i}{a_1} \quad (j \to \infty).$$

Lemma 8.3. *This notion of convergence defines a topology on $\mathcal{Y}(\cup_i S_i)$ which restricts to the given topology on $\mathcal{CG}(\cup S_i)$ and on $\partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k)$. The subspace $\partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k)$ is closed in $\mathcal{Y}(\cup_i S_i)$.*

Proof. Define a subset A of $\mathcal{Y}(\cup S_i)$ to be *closed* if $A_1 = A \cap \mathcal{CG}(\cup_i S_i)$ is closed, $A_2 = A \cap \partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k)$ is closed and if furthermore the following holds true. If $y_i \subset A_1$ is a sequence which converges in the sense described above to a point $y \in \partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k)$, then $y \in A_2$. Note that by definition, the empty set is closed, and the same holds true for the total space.

We have to show that complements of closed sets defined in this way fulfill the axioms of a topology, that is, they are stable under arbitrary unions and finite intersections. Equivalently, the family of closed sets is stable under arbitrary intersections and finite unions. As this holds true for the closed subsets of $\mathcal{CG}(\cup_i S_i)$ and for the closed subsets of $\partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k)$, all we need to observe is that taking arbitrary intersections and finite unions is compatible with the notion of convergence of points in $\mathcal{CG}(\cup_i S_i)$ to points in the join $\partial\mathcal{CG}(S_1) * \cdots * \partial\mathcal{CG}(S_k)$ in the sense specified above.

Consistency with arbitrary intersections is straightforward. To show consistency with finite unions let $B_1, \dots, B_\ell \subset \mathcal{Y}(\cup_i S_i)$ be closed in the above sense. Let $y_j \subset \cup_k (B_k \cap \mathcal{CG}(\cup_i S_i))$ be any sequence which converges to a point in $\mathcal{X}(\cup_i S_i)$ according to the definition of convergence. By passing to a subsequence, we may assume that $y_j \in B_m$ for a fixed m and all j. As B_m is closed and the subsequence also fulfills the requirements for convergence, its limit is contained in $B_m \subset \cup_k B_k$. Hence indeed, the notion of a closed set is consistent with taking finite unions. □

So far we have constructed a topology on the spaces $\mathcal{Y}(\cup_i S_i)$ where $S_1, \dots, S_k$ is a collection of disjoint subsurfaces of S. We now use these spaces to define convergent sequences in $\mathcal{X}$ and use this notion of convergent sequence to construct a topology on $\mathcal{X}$ which gives $\mathcal{X}$ the structure of a compact Hausdorff space.

Thus let $\xi^j = \sum_i a_i^j \xi_i^j$ be a sequence in $\mathcal{X}$. We shall impose 3 requirements for the sequence to converge to a point $\sum_{i=1}^k b_i \zeta_i \in \mathcal{X}$ (here as before, we require that

$a_i^j > 0, b_i > 0, \sum_i b_i = 1 = \sum_i a_i^j$ for all j and that furthermore, $\zeta_1, \dots, \zeta_k$ are disjoint minimal geodesic laminations).

By definition of the space $\mathcal{X}$, for each j the union

$$\hat{\xi}^j = \cup_i \xi_i^j$$

is a geodesic lamination on S, and the same holds true for $\hat{\zeta} = \cup_i \zeta_i$. Recall that the space of geodesic laminations on S is compact with respect to the Hausdorff topology.

Requirement 1: *Convergence in the coarse Hausdorff topology*
Let ξ^{ℓ_n} be any subsequence of the sequence ξ^j such that the geodesic laminations $\hat{\xi}^{\ell_n}$ converge in the Hausdorff topology to a geodesic lamination β. Then β contains $\hat{\zeta}$ as a sublamination.

Let Min$\mathcal{L}$ be the set of all minimal geodesic laminations on S where as before, a simple closed curve carries in addition a label $\pm$. We next define for each collection $S_1, \dots, S_k$ of pairwise disjoint subsurfaces of S a coarsely well defined projection

$$\mathrm{pr}_{\mathcal{Y}(\cup_i S_i)} : \mathrm{Min}\mathcal{L} \to \mathcal{Y}(\cup_i S_i) \tag{7}$$

as follows.

Let $\nu \in \mathrm{Min}\mathcal{L}$. The either its subsurface projection into a given subsurface S_i of S is empty, or it equals ν if ν is supported in S_i, or it consists of a collection of simple arcs with endpoints on the boundary which coarsely define a point in $\mathcal{CG}(S_i)$.

(1) If $\nu \in \partial\mathcal{CG}(S_i)$ for some i, then $\mathrm{pr}_{\mathcal{Y}(\cup_i S_i)}(\nu) = \nu \in \partial\mathcal{CG}(S_1) * \dots * \partial\mathcal{CG}(S_k)$.
(2) If $\nu \subset S_i$ for some i but if ν is disjoint from an essential simple closed curve $c \subset S_i$ then define $\mathrm{pr}_{\mathcal{Y}(\cup_i S_i)}(\nu) = (x_1, \dots, c, \dots, x_k)$ where x_j is the basepoint in $\mathcal{CG}(S_j)$.
(3) If $\nu \not\subset S_i$ for any i, then the subsurface projections of ν into S_i are either coarsely well defined simple closed curves or empty. Let $\mathrm{pr}_{\mathcal{Y}(\cup_i S_i)}(\nu) = (\nu_1, \dots, \nu_k)$ where for each i, the component ν_i either is a subsurface projection of ν into S_i if ν intersects S_i, or $\nu_i = x_i$ for the basepoint $x_i \in \mathcal{CG}(S_i)$.

We are now ready to define convergence of a sequence $\xi^j \subset \mathcal{X}$ of minimal geodesic laminations to a limit point $\zeta = \sum_i b_i \zeta_i$.

Requirement 2: Assume that ξ^j is a minimal geodesic lamination for all j. Let $\xi^{j_\ell} \subset \xi^j$ be any subsequence which converges in the Hausdorff topology to a lamination β. By the first requirement, we have $\beta \supset \hat{\zeta}$. Let $\beta_1, \dots, \beta_s$ be the minimal components of β, ordered in such a way that $\beta_i = \zeta_i$ for $i \le k$. For each i let S_i be the subsurface of S filled by β_i, that is, the intersection of all surfaces of S containing β; then $\mathrm{pr}_{\mathcal{Y}(\cup_i S_i)} \xi^{j_\ell} \to \zeta$ in $\mathcal{Y}(\cup_i S_i)$.

Example 8.4. i) Let $\varphi \in \mathcal{MCG}$ be a pseudo-Anosov element, with attracting geodesic lamination $\nu \in \partial\mathcal{CG}(S)$. Let $\mu \in \mathcal{X}$ be any minimal geodesic lamination which is different from the repelling fixed point ζ for the action of φ on $\partial\mathcal{CG}(S)$. Then $\varphi^j \mu \to \nu$ $(j \to \infty)$ in the coarse Hausdorff topology and therefore $\varphi^j \mu \to \nu$ in $\mathcal{X}$.

ii) Now let us assume that $S_0 \subset S$ is a proper connected subsurface different from an annulus and a pair of pants and that $\varphi \in \mathcal{MCG}$ restricts to a pseudo-Anosov mapping class on S_0 and to the trivial mapping class on $S - S_0$. Let $\nu \in \partial\mathcal{CG}(S_0)$ be the attracting geodesic lamination for the action of φ on S_0. Let furthermore $\mu \in \mathcal{X}$ be any *minimal* geodesic lamination on S which is different from the repelling fixed point ζ for the action of φ on $\mathcal{CG}(S_0)$. Then there are two possibilities. In the first case, μ is disjoint from ν, that is, μ is supported in $S - S_0$. Then we have $\varphi^j(\mu) = \mu$ for all j. However, if μ intersects S_0, then μ intersects ν and we have $\varphi^j(\mu) \to \nu$ $(j \to \infty)$ in $\mathcal{X}$.

The definition of convergence of a general sequence $\xi^j = \sum_i a_i^j \xi_i^j \subset \mathcal{X}$ to a limit point $\sum_{i=1}^k b_i\zeta_i$ is a bit more involved. To obtain a better understanding of what it captures, for a collection $\cup_{i=1}^k S_i$ of disjoint subsurfaces of S we define

$$\mathcal{Z}(\cup_i S_i) = \{a\zeta_1 + (1-a)\zeta_2 \mid a \in [0,1], \zeta_1 \in \mathcal{X}(\cup_{j\leq s} S_{i_j}), \zeta_2 \in \mathcal{Y}(\cup_{u\notin\{i_1,\dots,i_s\}} S_u)\}$$

to be the union of the joins of the spaces $\mathcal{X}(\cup_j S_{i_j})$ and $\mathcal{Y}(\cup_{u\notin\{i_1,\dots,i_s\}} S_u)$ where $\{i_1,\dots,i_s\}$ runs through all (possibly empty) subsets of the set $\{1,\dots,k\}$. Note that $\mathcal{Z}(\cup_i S_i)$ contains $\mathcal{Y}(\cup_i S_i)$ as a subset. However, the union is not meant to be a disjoint union as we identify points if they correspond to the same weighted geodesic laminations.

Construct next a projection $\mathrm{pr}_{\mathcal{Z}(\cup S_i)} : \mathcal{X} \to \mathcal{Z}(\cup_i S_i)$ as follows. Let $\xi = \sum_{i=1}^m a_i\xi_i \in \mathcal{X}$ with $a_i > 0$ and $\sum_i a_i = 1$ and write as before $\hat{\xi} = \cup_i \xi_i$. Assume that after perhaps a reordering that for some $u \leq \min\{k,m\}$ the components $\xi_1,\dots,\xi_u$ fill the subsurfaces $S_1,\dots,S_u$, that is, they define points in $\partial\mathcal{CG}(S_i)$, and that for no $i > u$, the component ξ_i fills any of the surfaces S_j. As the components of $\hat{\xi}$ are disjoint, this implies that if $s,t > u$, if $j \in \{u+1,\dots,k\}$ and if the subsurface projections of ξ_s, ξ_t into S_j are not empty, then they are points of uniformly bounded distance in $\mathcal{CG}(S_j)$ (where we adopt the convention to associate to any non-filling geodesic lamination in S_j a disjoint essential simple closed curve).

Define

$$\mathrm{pr}_{\mathcal{Z}(\cup S_i)} = \sum_{i=1}^u a_i\xi_i + (1 - \sum_{i=1}^u a_i)(\mathrm{pr}_{\mathcal{Y}(\cup_{i\geq u+1} S_i)} \cup_{i\geq u+1} \xi_i).$$

Here the term on the right hand side is understood in the following sense. Let us consider a subsurface S_j for some $j > u$. If there exists some $s > u$ such that ξ_s intersects S_j, then the component in S_j of the projection $\mathrm{pr}_{\mathcal{Y}(\cup_{i\geq u+1} S_i)}(\cup_{i\geq u+1}\xi_i)$ is a point in $\mathcal{CG}(S_j)$ which is coarsely determined by this projection. The above remark shows that this projection coarsely does not depend on choices, nor on the component ξ_s of ξ intersecting S_j. If the lamination $\hat{\xi} = \cup_i \xi^i$ is disjoint from the subsurface S_j, then the projection component is defined to be the basepoint of $\mathcal{CG}(S_j)$ constructed from the base marking.

Requirement 3: Let ξ^{j_s} be any subsequence of the sequence ξ^j so that the laminations $\hat{\xi}^{j_s}$ converge as $s \to \infty$ in the Hausdorff topology to a lamination β with minimal components $\beta_1,\dots,\beta_n$ for some $n \geq k$. By the first requirement, we have $\beta = \cup_i \beta_i \supset \hat{\zeta}$. Assume by reordering that $\beta_i = \zeta_i$ for $i \leq k$. For each i let S_i be the subsurface filled by β_i; then $\mathrm{pr}_{\mathcal{Z}(\cup_i S_i)}(\xi^{j_s}) \to \zeta$ in $\mathcal{Z}(\cup_i S_i) \supset \mathcal{Z}(\cup_{i\leq k} S_i) \supset \mathcal{X}(\cup_i S_i)$.

Remark 8.5. It follows from the above description that for this notion of convergence, the following holds true. Let ξ^j be a sequence in $\mathcal{X}$ consisting of minimal geodesic laminations which converges to a point $\zeta = \sum_u b_u \zeta_u$.

(1) The lamination $\cup_u \zeta_u$ is a sublamination of the limit in the coarse Hausdorff topology of any convergent subsequence of the sequence $\hat{\xi}^j = \cup_i \xi_i^j$.
(2) For each j let η^j be a minimal geodesic lamination disjoint from ξ^j (we allow $\eta^j = \xi^j$) and let $s_i \in [0,1]$. Then any limit of a convergent subsequence of the sequence $\nu^j = s_i \xi^j + (1-s_i)\eta^j$ is of the form $s\zeta + (1-s)\eta$ where η is a limit of a subsequence of the sequence η^j and where $s \in [0,1]$.

Definition 8.6. A subset $A \subset \mathcal{X}$ is called *closed for the geometric topology of* $\mathcal{X}$ if the following holds true. Let $\xi_i \subset A$ be any sequence which converges to a point $\xi \in \mathcal{X}$ in the sense described by the requirements (1),(2),(3); then $\xi \in A$.

An *embedding* of a topological space X into a topological space Y is an injective map $f : X \to Y$ which is a homeomorphism onto its image, equipped with the subspace topology. Recall that for any collection $S_1, \dots, S_k$ of pairwise disjoint subsurfaces of S, the space $\mathcal{X}(\cup_{i=1}^k S_i)$ is equipped with a natural topology as a join of the Gromov boundaries of the curve graphs of the subsurfaces S_i. The following statement is the key step towards the proof of Theorem 2.

Proposition 8.7. (1) *Closed subsets of* $\mathcal{X}$ *in the sense of Definition 8.6 define a topology* $\mathcal{O}$ *on* $\mathcal{X}$.
(2) *With respect to this topology,* $\mathcal{X}$ *is a compact separable Hausdorff space.*
(3) *For any collection* $S_1, \dots, S_k$ *of pairwise disjoint subsurfaces, the natural inclusion* $\mathcal{X}(\cup_{i=1}^k S_i) \to (\mathcal{X}, \mathcal{O})$ *is an embedding.*
(4) *The group* $\mathcal{MCG}$ *acts on* $\mathcal{X}$ *as a group of transformations.*

Proof. Let $\mathcal{O} \subset \mathcal{X}$ be the family of all subsets of $\mathcal{X}$ whose complement is closed in the above sense. As the empty set and the entire space are open, to show that $\mathcal{O}$ is indeed a topology on $\mathcal{X}$ it suffices to show that arbitrary unions of open sets are open, and that finite intersections of open sets are open as well. Or, equivalently, arbitrary intersections of closed sets are closed, and finite unions of closed sets are closed. However, this can be established using exactly the same reasoning as in the proof of Lemma 8.3.

We show next the third property claimed in the proposition. Thus let $S_1, \dots, S_k$ be a collection of pairwise disjoint subsurfaces of S. Our goal is to show that the inclusion $\mathcal{X}(\cup_{i=1}^k S_i) \to (\mathcal{X}, \mathcal{O})$ is an embedding. Since $\mathcal{X}(\cup_{i=1}^k S_i)$ is a separable Hausdorff space, for this it suffices to verify the following

Claim: Let $\xi^j = \sum_i a_i^j \xi_i^j \subset \mathcal{X}(\cup_{i=1}^k S_i)$ be any sequence. Then ξ^j converges in $\mathcal{X}(\cup_{i=1}^k S_i)$ to a point $\zeta = \sum_i b_i \zeta_i \in \mathcal{X}(\cup_{i=1}^k S_i)$ if and only if $\xi^j \to \zeta$ in $\mathcal{X}$.

Consider first a sequence $\xi^j = \sum_i a_i^j \xi_i^j \subset \mathcal{X}(\cup_{i=1}^k S_i)$ which converges in the space $\mathcal{X}(\cup_{i=1}^k S_i)$ to $\zeta = \sum_i b_i \zeta_i$. By reordering, assume that $1 \leq m \leq k$ is such that $b_i > 0$ if and only if $i \leq m$. Let β be a limit in the coarse Hausdorff topology of a subsequence of the sequence of laminations $\hat{\xi}^j = \cup_{a_i^j > 0} \xi_i^j$. By the definition of convergence in $\mathcal{X}(\cup_{i=1}^k S_i)$, up to reordering, we may assume that for some $m \leq n \leq k$, we have $\beta = \cup_{i\leq n}\beta_i$ where β_i is a (not necessarily minimal and not necessarily filling) geodesic lamination on the surface S_i, and $\beta_i = \zeta_i$ for $1 \leq i \leq m$. The fact that n may be strictly smaller than k arises from the possibility that the formal sum describing ξ^j may not have a positive coefficient corresponding to a surface S_ℓ for $\ell > m$.

Since $\xi^j \in \mathcal{X}(\cup_{i=1}^k S_i)$ for all j, we know that the projection $\mathrm{pr}_{\mathcal{Z}(\cup_{i\leq m} S_i)}\xi^j$ of ξ^j to $\mathcal{Z}(\cup_{i\leq m} S_i)$ is contained in $\mathcal{X}(\cup_{i=1}^m S_i)$. Now by definition of the topology on $\mathcal{X}(\cup_{i=1}^k S_i)$, the subset $\mathcal{X}(\cup_{i=1}^m S_i)$ is an embedded subspace of $\mathcal{X}(\cup_{i=1}^k S_i)$, and the surfaces S_i for $i > m$ are precisely those surfaces with the property that the coefficients a_i^j of the components ξ_i^j of ξ^j in S_i tend to zero as $j \to \infty$. Furthermore, for $i \leq m$ the coefficients a_i^j converge to b_i. Thus an application of the first and third requirement in the definition of convergent sequences for $\mathcal{O}$ shows that indeed, $\xi^j \to \zeta \in \mathcal{X}$.

To summarize, we showed that a sequence $\xi^j \subset \mathcal{X}(\cup_{i=1}^k S_i)$ which converges in $\mathcal{X}(\cup_{i=1}^k S_i)$ to a limit point $\zeta = \sum_i b_i\zeta_i$ also converges in $(\mathcal{X}, \mathcal{O})$ to the same limit point. To complete the proof of the claim, we have to show that a sequence in $\mathcal{X}(\cup_{i=1}^k S_i)$ which converges in $(\mathcal{X}, \mathcal{O})$ to a limit point $\zeta = \sum_i b_i\zeta_i \in \mathcal{X}(\cup_{i=1}^k S_i)$ also converges in $\mathcal{X}(\cup_{i=1}^k S_i)$ to the same limit point. However, this can be established with essentially the same argument and will be omitted. The third part of the proposition follows.

Establishing the second property in the proposition is the most involved part of the proof. Note first that each of the countably many embedded subspaces $\mathcal{X}(\cup_{i=1}^k S_i)$ of $\mathcal{X}$ is a join of finitely many separable spaces and hence separable, and their union is all of $\mathcal{X}$. Thus it follows from the third property in the proposition that the space $(\mathcal{X}, \mathcal{O})$ is separable as well.

We show next that the topology $\mathcal{O}$ is Hausdorff. Thus let $\xi = \sum_i a_i\xi_i \neq \zeta = \sum_j b_j\zeta_j \in \mathcal{X}$. We have to show that ξ, ζ have disjoint neighborhoods. If this is not the case, then any neighborhoods U_ξ of ξ and U_ζ of ζ intersect nontrivially. Since $\mathcal{X}$ is separable, and since points in $\mathcal{X}$ are closed by construction, we conclude that there is a sequence $\xi^j \subset \mathcal{X}$ which converges both to ξ, ζ. But for the notion of convergence used to define the topology $\mathcal{O}$, the limit of a converging sequence is unique. Thus $\mathcal{O}$ is Hausdorff as stated.

For the completion of the proof of the third property in the proposition, we are left with showing that the topological space $(\mathcal{X}, \mathcal{O})$ is compact. As $\mathcal{X}$ is a separable Hausdorff space, this is equivalent to being sequentially compact.

Thus let $\xi^j = \sum_i a_i^j\xi_i^j \subset \mathcal{X}$ be any sequence. We have to construct a convergent subsequence. Since the space of geodesic laminations equipped with the Hausdorff topology is compact, by passing to a subsequence we may assume that the geodesic laminations $\hat{\xi}^j = \cup_i\xi_i^j$ converge in the Hausdorff topology to a geodesic lamination $\hat{\zeta}$ with minimal components $\zeta_1, \dots, \zeta_k$.

Let $S_i\subset S$ be the subsurface of S filled by ζ_i. Assume by passing to another subsequence that for each component ζ_i of $\hat\zeta$, either this component also is a component of $\hat\xi^j$ for all j, or it is not a component of $\hat\xi^j$ for all j.

For some $u\leq k$ let $\zeta_1,\dots,\zeta_u$ be those components of $\hat\zeta$ which are also components of $\hat\xi^j$ for all j. By reordering, we then can write

$$\xi^j=(\sum_{i=1}^u a_i^j\zeta_i^\pm)+\sum_{\ell>u}a_\ell^j\xi_\ell^j.$$

By convention, the label $\pm$ is only relevant if ζ_i is a simple closed curve component.

By passing to another subsequence, we may assume that for $i\leq u$, the labels of the components ξ_i^j are constant along the sequence, and that the weights $a_i^j\in(0,1]$ of the components ξ_i^j converge to weights $b_i\geq 0$. In particular, the sums $1-\sum_{i\leq u}a_i^j$ converge to $1-\sum_{i\leq u}b_i=\kappa$. If $\kappa=0$, then by the definition of convergent sequences in $\mathcal{X}$, the sequence ξ^j converges to $\sum_i b_i\zeta_i^\pm$ and we are done.

Now assume that $\kappa\neq 0$ and hence $\sum_{i>u}a_i^j>\kappa/2>0$ for all sufficiently large j. By passing to a subsequence, we may assume that this holds true for all j. For all $i>u$ and for all j, the subsurface of S filled by ξ_i^j is disjoint from the subsurfaces $S_1,\dots,S_u$ filled by the laminations $\zeta_1,\dots,\zeta_u$. In other words, if we denote by $\Sigma_{u+1},\dots,\Sigma_n$ the components of $S-\cup_{i\leq u}S_i$, then for $i\geq u+1$, each of the laminations ξ_i^j,ζ_i is supported in $\cup_{i\geq u+1}\Sigma_i$. Thus by the definition of the topology on $\mathcal{X}$ and writing $\xi^j=(\sum_{i\leq u}a_i^j\xi^j)+(\sum_{i\geq u+1}a_i^j\xi^j)$, viewed as points in the join of two subspaces of $\mathcal{X}$ and similarly for ζ, we conclude that it suffices to construct a convergent subsequence of a sequence ξ^j under the additional assumption that for all j, no component ξ_i^j of $\hat\xi^j$ coincides with a component of the limit $\hat\zeta=\cup_{i\leq k}\zeta_i$ in the coarse Hausdorff topology.

From now on we assume that the latter assumption holds true. Let as before S_i be the subsurface of S filled by ζ_i. Up to passing to a subsequence, we may assume that there is a number $u\leq k$ such that for each $i\leq u$ and each j, the geodesic lamination $\hat\xi^j$ has a component ξ_i^j which is supported in S_i and fills S_i. Since $\hat\xi^j$ converges as $j\to\infty$ in the Hausdorff topology to a geodesic lamination with minimal components $\zeta_1,\dots,\zeta_k$, we conclude that for $i\leq u$, the laminations ξ_i^j converge as $j\to\infty$ in the coarse Hausdorff topology to ζ_i. By passing to another subsequence, we may assume that for each $i\leq u$, the coefficients a_i^j converge as $j\to\infty$ to a coefficient b_j. As above, if $\sum_{i\leq u}b_i=1$, then by the definition of a convergent sequence in $\mathcal{X}$, we know that $\xi^j\to\sum_i b_i\zeta_i$ and hence once again, we are done.

According to what we established so far, it now suffices to assume that for no j there exists a component of $\hat\xi^j$ which fills any of the subsurfaces S_i. Then for each i, we can consider the subsurface projection $\mathrm{pr}_{S_i}(\hat\xi^j)$ of $\hat\xi^j$ into the surface S_i. Furthermore, by passing to another subsequence, we may assume that for all j and all $i\leq k$, this subsurface projection it is non-empty since the geodesic lamination ζ_i which fills S_i is contained in the limit with respect to the Hausdorff topology of the sequence of laminations $\hat\xi^j$. Put differently, we may assume that for each i and all

j, the subsurface projection $\mathrm{pr}_{S_i}(\hat{\xi}^j)$ of the lamination $\hat{\xi}^j$ into the subsurface S_i is a coarsely well defined point in $\mathcal{CG}(S_i)$. Furthermore, using once more that ζ_i fills S_i and that ζ_i is contained in the Hausdorff limit of the sequence $\hat{\xi}^j$, if we denote by x_i the fixed basepoint in $\mathcal{CG}(S_i)$, then we know that $d_{\mathcal{CG}(S_i)}(\mathrm{pr}_{S_i}(\hat{\xi}^j), x_i) \to \infty$ $(j \to \infty)$.

By passing to another subsequence and reordering indices, we may assume that

$$a_1^j = d_{\mathcal{CG}(S_1)}(\mathrm{pr}_{S_1}(\xi^j), x_1) \geq a_i^j = d_{\mathcal{CG}(S_i)}(\mathrm{pr}_{S_i}(\xi^j), x_i)$$

for all $i \geq 2$ and all j. Passing to another subsequence, we may assume furthermore that $a_i^j/a_1^j \to a_i \in [0,1]$ for all $i \geq 2$. Put $a_1 = 1$; then we have $\sum_u a_u \geq 1$ and hence defining $b_i = a_i/\sum_u a_u > 0$, we conclude that $\sum_u b_u = 1$. It now follows from the definition of the topology on $\mathcal{X}$ that $\xi^j \to \sum_i b_i\zeta_i$. This completes the proof that $\mathcal{X}$ is sequentially compact.

We are left with showing that the mapping class group $\mathcal{MCG}$ acts on $\mathcal{X}$ as a group of transformations. To this end observe first that by construction, $\mathcal{MCG}$ acts on $\mathcal{X}$ as a group of bijections. Thus it suffices to show that this action is continuous for the topology $\mathcal{O}$.

By the definition of $\mathcal{O}$, for this it suffices to show the following. Let ξ^j be a sequence converging for the topology $\mathcal{O}$ to a point ξ. Then for every $\varphi \in \mathcal{MCG}$, the sequence $\varphi(\xi^j)$ converges to $\varphi(\xi)$.

That the first defining requirement for convergence is passed on to the image sequence follows from continuity of the action of φ on the space of geodesic laminations, equipped with the Hausdorff topology.

For the second requirement, if $S_1, \dots, S_k$ is a partition of S into disjoint subsurfaces, then the same holds true for $\varphi(S_1), \dots, \varphi(S_k)$, and for any geodesic lamination ν, we have $\mathrm{pr}_{\mathcal{Y}(\cup_i \varphi(S_i))}(\varphi(\nu)) = \varphi(\mathrm{pr}_{\mathcal{Y}(\cup_i S_i)}(\nu))$ up to replacing the basepoints y_i of $\mathcal{CG}(\varphi(S_i))$ by $\varphi(x_i)$. As for all i, we have $d_{\mathcal{CG}(\varphi(S_i))}(\mathrm{pr}_{\varphi(S_i)}(\xi^j), \varphi(x_i)) = d_{\mathcal{CG}(S_i)}(\xi^j, x_i) \to \infty$ $(j \to \infty)$ and the determination of the weights of the limit points are computed using ratios of distances to the basepoint defined by subsurface projections, with the distances tending to infinity along the sequence, we conclude that the second requirement in the definition of convergence is fulfilled for $\varphi(\xi^i)$ if it is fulfilled for ξ^i. The same reasoning also applies to the third requirement. Thus indeed, $\mathcal{MCG}$ acts on $\mathcal{X}$ as a group of transformations. This completes the proof of the proposition. □

Let us note another naturality property of the geometric boundary of $\mathcal{MCG}$. Namely, if $S_0 \subset S$ is any essential subsuface, then we can construct a geometric boundary $\mathcal{X}_0$ for $\mathcal{MCG}(S_0)$. As a set, this is a subset of the geometric boundary of S. The above construction immediately yields

Corollary 8.8. *If $S_0 \subset S$ is any subsurface of S, then the geometric boundary of $\mathcal{MCG}(S_0)$ is a closed subspace of the geometric boundary of $\mathcal{MCG}(S)$.*

In the remainder of this article, we show that the geometric boundary $\mathcal{X}$ is indeed a small boundary for $\mathcal{MCG}$. To this end recall that the mapping class group acts properly and cocompactly on $\mathcal{TT}$. Thus to show that $\mathcal{X}$ is a boundary for $\mathcal{MCG}$ it suffices to construct a topology $\mathcal{O}_0$ on $\overline{\mathcal{TT}} = \mathcal{TT} \cup \mathcal{X}$ with the following property.

(1) $\mathcal{O}_0$ restricts to the simplicial topology on $\mathcal{TT}$ and to the topology $\mathcal{O}$ on $\mathcal{X}$.
(2) $\overline{\mathcal{TT}}$ is compact.
(3) The group $\mathcal{MCG}$ acts on $\overline{\mathcal{TT}}$ as a group of transformations.

As before, we define a topology $\mathcal{O}_0$ on $\mathcal{TT} \cup \mathcal{X}$ by defining what it means for a sequence $(\tau_i) \subset \mathcal{TT}$ to converge to a point $\zeta = \sum_i b_i \zeta_i \in \mathcal{X}$. Namely, let $\mathcal{VC}(\tau)$ be the set of all vertex cycles of τ, viewed as simple closed curves. Since τ is complete, for every vertex cycle $v \in \mathcal{VC}(\tau)$ there exists a second vertex cycle $w \in \mathcal{VC}(\tau)$ which intersects v transversely. The number of intersection points between two vertex cycles of a train track τ is bounded from above independent of τ and hence the distance of their subsurface projections into any subsurface S_0 of S is uniformly bounded as well provided these projections are non-empty. Note also that for any subsurface S_0 of S, there exists a vertex cycle $v \in \mathcal{VC}(\tau)$ which intersects S_0. Thus we can talk about the (coarsely well defined) subsurface projections of $\mathcal{VC}(\tau)$ into S_0. We denote this projection by $\mathrm{pr}_{S_0}(\tau)$.

As before, for a subsurface S_0 of S let $\mu(S_0)$ be a marking of S_0 obtained by projecting a fixed marking of S. This choice coarsely defines a basepoint in $\mathcal{CG}(S_0)$.

Definition 8.9. A sequence $\tau_j \subset \mathcal{TT}$ *converges* to a point $\xi = \sum_i a_i \xi_i \in \mathcal{X}$ if and only if the following holds true.

(1) Let $v_i \in \mathcal{VC}(\tau_i)$ be an arbitrary vertex cycle. Then any limit of v_i in the Hausdorff topology is disjoint from $\hat{\xi} = \cup_i \xi_i$.
(2) Let S_i be the subsurface of S filled by ξ_i and let U be any subsurface of S disjoint from the surfaces S_i. Then for each i, the projections $\mathrm{pr}_{S_i}(\tau_j)$ converge in $\mathcal{CG}(S_i)$ to ξ_i. Furthermore, if x_i is the basepoint in $\mathcal{CG}(S_i)$ and x_U is the basepoint of $\mathcal{CG}(U)$ then we have
$$d_{\mathcal{CG}(S_i)}(\mathrm{pr}_{S_i}(\tau_j), x_i)/d_{\mathcal{CG}(S_1)}(\mathrm{pr}_{S_1}(\tau_j), x_1) \to a_i/a_1 \text{ for all } i,$$
and $d_{\mathcal{CG}(U)}(\mathrm{pr}_U(\tau_j), x_U)/d_{\mathcal{CG}(S_1)}(\mathrm{pr}_{S_1}(\tau_j), x_1) \to 0$.

We first observe that this notion of converge gives indeed rise to a topology on $\mathcal{TT} \cup \mathcal{X}$.

Lemma 8.10. *There exists a topology $\mathcal{O}_0$ on $\overline{\mathcal{TT}} = \mathcal{TT} \cup \mathcal{X}$ with the property that a set $A \subset \overline{\mathcal{TT}}$ is closed for $\mathcal{O}_0$ if and only if the following holds true.*

(1) *$A \cap \mathcal{TT}$ is closed in $\mathcal{TT}$, and $A \cap \mathcal{X}$ is closed in $\mathcal{X}$.*
(2) *If $\tau_i \subset A \cap \mathcal{TT}$ converges in the above sense to a point $\xi \in \mathcal{X}$, then $\xi \in A$.*

Proof. The proof of completely analogous to the proof of Lemma 8.3 and will be omitted. ☐

The following is the main remaining step towards a proof of Theorem 2.

Proposition 8.11. *The topological space* $(\overline{\mathcal{TT}}, \mathcal{O}_0)$ *has the following properties.*

(1) $\overline{\mathcal{TT}}$ *is a compact separable Hausdorff space.*
(2) *The mapping class group acts on* $\overline{\mathcal{TT}}$ *as a group of transformations.*

Proof. $\overline{\mathcal{TT}}$ is clearly separable since this holds true for $\mathcal{X}$ and $\mathcal{TT}$. We show next that $\overline{\mathcal{TT}}$ is a Hausdorff space.

Since $\mathcal{TT}$ is open in $\overline{\mathcal{TT}}$ and a Hausdorff space, all we need to show is that two points $\xi \neq \eta \in \mathcal{X}$ have disjoint neighborhoods. Now ξ, η have disjoint neighborhoods in $\mathcal{X}$ and hence it suffices to show that the limit of any sequence $\tau_i \subset \mathcal{TT}$ converging to a point in $\mathcal{X}$ is unique. But this is clear from the definitions.

We show next that $\overline{\mathcal{TT}}$ is compact. Since $\mathcal{X}$ is compact, for this it suffices to show that any sequence $\tau_i \subset \mathcal{TT}$ has a convergent subsequence.

If the sequence has a bounded subsequence in $\mathcal{TT}$, then as $\mathcal{TT}$ is proper, we can extract a converging subsequence. Thus we may assume that the sequence is unbounded.

Since the space of geodesic laminations on S equipped with the Hausdorff topology is compact, by extracting a subsequence we may assume that the sets $\mathcal{VC}(\tau_i)$ converge in the Hausdorff topology to a finite union of geodesic laminations. Note that the number of such laminations is bounded from above by the maximal number of vertex cycles of a train track τ and hence this number is uniformly bounded.

Let $\zeta_1, \dots, \zeta_s$ be those of the components of the limit laminations which are distinct from simple closed curves. The number of such components is finite. Each of the laminations ζ_i fills a subsurface S_i of S which is different from an annulus or a pair of pants. Furthermore, the subsurface projections to $\mathcal{CG}(S_i)$ of the sets $\mathcal{VC}(\tau_j)$ converge as $j \to \infty$ in $\overline{\mathcal{CG}}(S_i)$ to the lamination $\zeta_i \in \partial\mathcal{CG}(S_i)$. Since this holds true for any sequence of subsurface projections of some choice of vertex cycle of τ_j provided that these projections are not empty, this implies that none of the limits in the Hausdorff topology of any sequence of vertex cycles of τ_j can intersect ζ_i.

By a similar argument, if ζ_i is a closed curve component, then we can consider the subsurface projections of a vertex cycle of τ_j to an annulus $A(\zeta_i)$ with core curve ζ_i. Up to passing to a further subsequence, we may assume that these projections are either bounded along the sequence, or converge to one of the two boundary components of the curve graph of $A(\zeta_i)$. In the first case call ζ_i *unlabeled.* In the second case, label ζ_i with the corresponding point in the Gromov boundary of the curve graph of $A(\zeta_i)$ and note by the reasoning used in the previous paragraph, no labeled simple closed curve component ζ_i can be intersected by another component ζ_j.

By reordering, let $\zeta_1, \dots, \zeta_k$ be the components of the limit laminations which either are distinct from simple closed curves or which are labeled simple closed curves. By the above discussion, we know that $\hat{\zeta} = \cup_{i=1}^k \zeta_i$ is a geodesic lamination. Furthermore, we know that if S_i is the subsurface of S filled by S_i then

$d_{\mathcal{CG}}(\mathrm{pr}_{S_i}(\mathcal{VC}(\tau_j)), x_i) \to \infty$ where as before, $x_i \in \mathcal{CG}(S_i)$ is a fixed basepoint for $\mathcal{CG}(S_i)$.

By passing to a subsequence and reordering, we may assume that

$$d_{\mathcal{CG}(S_i)}(\mathrm{pr}_{S_1}(\mathcal{VG}(\tau_j)), x_1) \geq d_{\mathcal{CG}(S_i)}(\mathrm{pr}_{S_i}(\mathcal{VG}(\tau_j)), x_i) + b$$

for all i, j where b is twice the maximal diameter of the subsurface projection of any of the sets $\mathcal{VC}(\tau_j)$. Then by passing to another subsequence, we may assume that

$$d_{\mathcal{CG}(S_i)}(\mathrm{pr}_{S_i}(\mathcal{VG}(\tau_j)), x_i)/d_{\mathcal{CG}(S_i)}(\mathrm{pr}_{S_1}(\mathcal{VG}(\tau_j)), x_1) \to b_i \leq 1.$$

Define $a_i = b_i / \sum_j b_j$ and let $\xi = \sum_{i=1}^k a_i \zeta_i$.

We claim that $\tau_i \to \xi \in (\overline{\mathcal{TT}}, \mathcal{O}_0)$. To this end note that the first property in the definition of a convergent sequence is fulfilled by the above discussion, and the second holds true by the observation that if there exists a subsurface U different from the surfaces S_i $(i \leq k)$ such that for some subsequence, the subsurface projections of $\mathcal{VC}(\tau_j)$ to U are unbounded, then the subsurface projections $\mathrm{pr}_U(\mathcal{VC}(\tau_j))$ converge up to passing to a subsequence in the Hausdorff topology to a lamination which fills U, violating the choice of the laminations ζ_i.

To summarize, we showed so far that $\overline{\mathcal{TT}}$ is a compact Hausdorff space. We are left with showing that $\mathcal{MCG}$ acts on $\overline{\mathcal{TT}}$ as a group of transformations. However, as $\mathcal{MCG}$ acts on $\mathcal{TT}$ and on $\mathcal{X}$ as a group of transformations, and as the definition of convergence which determines the topology $\mathcal{O}_0$ is natural with respect to the action of $\mathcal{MCG}$ on subsurfaces and subsurface projections, this is indeed the case. The proposition is proven. □

Theorem 8.12. *$\mathcal{X}$ is a small boundary for $\mathcal{MCG}$. A pseudo-Anosov mapping class acts on $\mathcal{X}$ with north-south dynamics. In particular, the action of $\mathcal{MCG}$ on $\mathcal{X}$ is strongly proximal.*

Proof. We showed so far that $\mathcal{X}$ defines a boundary of $\mathcal{TT}$ and hence of $\mathcal{MCG}$. Furthermore, a pseudo-Anosov element acts on $\mathcal{X}$ with north-south dynamics and hence the action of $\mathcal{MCG}$ on $\mathcal{X}$ is strongly proximal.

We are left with showing that the right action of $\mathcal{MCG}$ induces the identity. However, this action just consists of a change of basepoint. As a sequence of points of uniformly bounded distance from a convergent sequence converges to the same point, this yields the statement of the theorem. □

References

[ALS09] J. Aramayona, C. Leininger, J. Souto, *Injections of mapping class groups*, Geom. Topol. 13 (2009), 2523–2541.

[BGH20] H Baik, I. Gekhtman, U. Hamenstädt, *The smallest positive eigenvalue of fibered hyperbolic 3-manifolds*, Proc. Lond. Math. Soc. 120 (2020), 704–741.

[B96] M. Bestvina, *Local homology properties of boundaries of groups*, Michigan Math. J. 1996, 123–141.

[BB19] M. Bestvina, K. Bromberg, *On the asymptotic dimension of the curve complex*, Geometry & Topology 23 (2019), 2227–2276.

[BBF19] M. Bestvina, K. Bromberg, K. Fujiwara, *Proper actions on finite products of quasi-trees*, arXiv:1905.10813.

[BH73] J. Birman, H. Hilden, *On isotopies of homeomorphisms of Riemann surfaces*, Ann. Math. 97 (1973), 424-439.

[B74] J. Birman, *Braids, links and mapping class groups*, Ann. Math. Studies, Princeton Univ. Press, Princeton 1974.

[Br10] M. Bridson, *Semisimple actions of mapping class groups on* CAT(0) *spaces*, in "Geometry of Riemann surfaces", 1-14, London Math. Soc. Lecture Note Ser., 368, Cambridge Univ. Press, Cambridge, 2010.

[BH99] M. Bridson, A. Haefliger, *Metric spaces of non-positive curvature*, Springer Grundlehren 319, Springer, Berlin 1999.

[Bw08] B. Bowditch, *Tight geodesics in the curve complex*, Invent. Math. 171 (2008), 281–300.

[BFP11] N. Braddeus, B. Farb, A. Putman, *Irreducible Sp-representations and subgroup distortion in the mapping class group*, Comment. Math. Helv. 86 (2011), 537–556.

[CEG87] R. Canary, D. Epstein, P. Green, *Notes on notes of Thurston*, in "Analytical and geometric aspects of hyperbolic space", edited by D. Epstein, London Math. Soc. Lecture Notes 111, Cambridge University Press, Cambridge 1987.

[CDH10] I. Chatterji, C. Drutu, F. Haglund, *Kazhdan and Haagerup properties from the median viewpoint*, Adv. Math. 225 (2010), 882–921.

[DHS17] M. Durham, M. Hagen and A. Sisto, *Boundaries and automorphisms of hierarchically hyperbolic spaces*, Geometry & Topology 21 (2017), 3659–3758.

[FLM01] B. Farb, A. Lubotzky, Y. Minsky, *Rank one phenomena for mapping class groups*, Duke Math. J. 106 (2001), 581-597.

[FLP91] A. Fathi, F. Laudenbach, V. Poénaru, *Travaux de Thurston sur les surfaces,* Astérisque 1991.

[FLM18] T. Fernos, J. Lecureux and F. Mathéus, *Random walks and boundaries of* CAT(0) *cubical complexes*, Comm. Math. Helv. 93 (2018), 291–333.

[G14] D. Gabai, *On the topology of ending lamination space*, Geom. & Top. 18 (2014), 2683–2745.

[GM18] V. Gadre, J. Maher, *The stratum of random mapping classes*, Erg. Th. & Dyn. Sys. 38 (2018), 2666–2682.

[H06] U. Hamenstädt, *Train tracks and the Gromov boundary of the complex of curves*, in "Spaces of Kleinian groups" (Y. Minsky, M. Sakuma, C. Series, eds.), London Math. Soc. Lec. Notes 329 (2006), 187–207.

[H09] U. Hamenstädt, *Geometry of the mapping class groups I: Boundary amenability*, Invent. Math. 175 (2009), 545–609.

[KR14] I. Kapovich, K. Rafi, *On hyperbolicity of free splitting and free factor complexes*, Groups, Geom. Dyn. 8 (2014), 391–414.

[Ke83] S. Kerckhoff, *The Nielsen realization problem*, Ann. Math. 117 (1983), 235–265.

[K99] E. Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space*, unpublished manuscript. Ann Arbor (1999).

[Kb12] T. Koberda, *Right angled Artin groups and a generalized isomorphism problem for finitely generated subgroups of mapping class groups*, Geom. Funct. Anal. 22 (2012), 1541–1590.

[MT18] J. Maher, G. Tiozzo, *Random walks on weakly hyperbolic groups*, J. Reine Angew. Math. 742 (2018), 187–239.

[MM99] H. Masur, Y. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, Invent. Math. 138 (1999), 103-149.

[MM00] H. Masur, Y. Minsky, *Geometry of the complex of curves II: Hierarchical structure*, GAFA 10 (2000), 902-974.

[MM04] H. Masur, Y. Minsky, *Quasi-convexity in the curve complex*, Cont. Math. 355, Amer. Math. Soc., Providence, (2004), 309–320.

[M96] L. Mosher, *Hyperbolic extensions of groups*, J. pure appl. Algebra 110 (1996), 305–314.

[RS07] K. Rafi, S. Schleimer, *Covers and the curve complex*, Geom. Topol. 13 (2009), 2141–2162.

[PH92] R. Penner with J. Harer, *Combinatorics of train tracks*, Ann. Math. Studies 125, Princeton University Press, Princeton 1992.

[S95] M. Sageev, *Ends of group pairs and non-positively curved cube complexes*, Proc. London Math. Soc. 71 (1995), 585–617.

MATHEMATISCHES INSTITUT DER UNIVERSITÄT BONN
ENDENICHER ALLEE 60, D-53115 BONN, GERMANY