

GEOMETRY OF THE MAPPING CLASS GROUPS I: BOUNDARY AMENABILITY

URSULA HAMENSTÄDT

ABSTRACT. We construct a geometric model for the mapping class group \mathcal{MCG} of a non-exceptional oriented surface S of genus g with k punctures and use it to show that the action of \mathcal{MCG} on the compact metrizable Hausdorff space of complete geodesic laminations for S is topologically amenable. As a consequence, the Novikov higher signature conjecture holds for every subgroup of \mathcal{MCG} .

CONTENTS

1. Introduction	2
2. Train tracks and geodesic laminations	6
2.1. Complete geodesic laminations	6
2.2. Complete train tracks	8
3. The complex of train tracks	13
4. Train tracks and the mapping class group	22
5. Cubical Euclidean cones	24
6. Quasi-isometrically embedded subspaces	30
7. Boundary amenability	37
8. A strong boundary for \mathcal{MCG}	40
9. Super-rigidity of cocycles	43
Appendix A. Splitting, Shifting and Carrying	50
References	60

Date: October 6, 2008.

AMS subject classification: 20M34

Research partially supported by Sonderforschungsbereich 611.

1. INTRODUCTION

A countable group Γ is called *boundary amenable* if it admits a *topologically amenable* action on a compact Hausdorff space X . This means that Γ acts on X as a group of homeomorphisms and that moreover the following holds. Let $\mathcal{P}(\Gamma)$ be the convex space of all probability measures on Γ . Note that $\mathcal{P}(\Gamma)$ can be viewed as a subset of the unit ball in the space $\ell^1(\Gamma)$ of summable functions on Γ and therefore it admits a natural norm $\|\cdot\|$. The group Γ acts on $(\mathcal{P}(\Gamma), \|\cdot\|)$ isometrically by left translation. We require that there is a sequence of weak*-continuous maps $\xi_n : X \rightarrow \mathcal{P}(\Gamma)$ with the property that $\|g\xi_n(x) - \xi_n(gx)\| \rightarrow 0$ ($n \rightarrow \infty$) uniformly on compact subsets of $\Gamma \times X$ (see [AR00] for more on amenable actions).

By a result of Higson [H00] which is based on earlier work of Yu [Y98, Y00], for any countable group Γ which is boundary amenable and for every separable $\Gamma - C^*$ -algebra A , the Baum-Connes assembly map

$$\mu : KK_*^\Gamma(\mathcal{E}\Gamma, A) \rightarrow KK(\mathbb{C}, C_r^*(\Gamma, A))$$

is split injective. As a consequence, the strong Novikov conjecture holds for Γ and hence the Novikov higher signature conjecture holds as well [BCH94, MV03].

Now let S be a non-exceptional oriented surface of finite type. This means that S is a closed surface of genus $g \geq 0$ from which $k \geq 0$ points, so-called *punctures*, have been deleted, and where $3g - 3 + k \geq 2$. The *mapping class group* \mathcal{MCG} of S is the group of all isotopy classes of orientation preserving self-homeomorphisms of S . The mapping class group is finitely presented. We refer to [I02] for a summary of the basic properties of the mapping class group and for references.

By assumption, the Euler characteristic of S is negative and hence S admits a complete hyperbolic Riemannian metric of finite volume. A *geodesic lamination* for such a hyperbolic structure is a *compact* subset of S which is foliated into simple geodesics. Call a geodesic lamination λ *complete* if its complementary components are all ideal triangles or once punctured monogons and if moreover λ can be approximated in the *Hausdorff topology* by simple closed geodesics. The space \mathcal{CL} of complete geodesic laminations on S equipped with the Hausdorff topology is compact and metrizable (see Section 2 of this paper). The mapping class group naturally acts on the space \mathcal{CL} as a group of homeomorphisms. In other words, \mathcal{CL} is a compact metrizable \mathcal{MCG} -space which we call the *Furstenberg boundary* of \mathcal{MCG} . We show.

Theorem 1. *The action of the mapping class group of a non-exceptional surface S of finite type on the space of complete geodesic laminations on S is topologically amenable.*

As a consequence, the mapping class groups are boundary amenable. Since boundary amenability is passed on to subgroups (see Chapter 5 of [AR00]), as explained above the following corollary is immediate from Theorem 1 and the work of Higson [H00].

Corollary 1. *The Novikov higher order signature conjecture holds for any subgroup of the mapping class group of a non-exceptional surface of finite type.*

An action of a countable group Γ on a compact Hausdorff space can only be topologically amenable if its point stabilizers are amenable subgroups of Γ . The point stabilizers of the action of \mathcal{MCG} on \mathcal{CL} are virtually abelian. For example, if $\lambda \in \mathcal{CL}$ is a complete geodesic lamination whose minimal components are simple closed curves, then the stabilizer of λ in \mathcal{MCG} contains the free abelian subgroup of \mathcal{MCG} generated by the Dehn twists about these components as a subgroup of finite index (this fact can easily be deduced from the results in [MP89]).

On the other hand, it is well known that an amenable subgroup Γ of \mathcal{MCG} is virtually abelian (see [I02] for references). Together with the structural results for subgroups of \mathcal{MCG} in [MP89], it is not very difficult to show that every amenable subgroup of \mathcal{MCG} has a subgroup of finite index which fixes a point in \mathcal{CL} (we do not need this fact for our purpose). This is analogous to the properties of the action of a simple Lie group of non-compact type and higher rank on its Furstenberg boundary and justifies the terminology Furstenberg boundary of \mathcal{MCG} for \mathcal{CL} . In contrast, the action of \mathcal{MCG} on the more familiar Thurston boundary of Teichmüller space of all projective measured geodesic laminations on S is not topologically amenable.

A simple closed curve on the surface S is called *essential* if it is neither contractible nor freely homotopic into a puncture. The *curve graph* $\mathcal{C}(S)$ of the surface S is a locally infinite metric graph whose vertices are the free homotopy classes of essential simple closed curves on S and where two such vertices are connected by an edge of length one if they can be realized disjointly. By an important result of Masur and Minsky [MM99], the curve graph is hyperbolic in the sense of Gromov. The mapping class group acts on $\mathcal{C}(S)$ as a group of isometries.

The Gromov hyperbolic geodesic metric space $\mathcal{C}(S)$ admits a *Gromov boundary* $\partial\mathcal{C}(S)$ which is a metrizable topological space; however, it is not locally compact. The action of the mapping class group \mathcal{MCG} on $\mathcal{C}(S)$ extends to an action on $\partial\mathcal{C}(S)$ by homeomorphisms.

An action of a group on a topological space is called *universally amenable* if it is amenable with respect to every invariant Borel measure class. As another corollary of Theorem 1 we obtain the following result.

Corollary 2. *The action of \mathcal{MCG} on $\partial\mathcal{C}(S)$ is universally amenable.*

As in the case of a simple Lie group of non-compact type and higher rank, the Furstenberg boundary of \mathcal{MCG} has a geometric interpretation. Namely, an *L -quasi-isometric embedding* of a metric space (X, d) into a metric space (Y, d) is a map $F : X \rightarrow Y$ such that

$$d(x, y)/L - L \leq d(Fx, Fy) \leq Ld(x, y) + L \text{ for all } x, y \in X.$$

The map F is called an *L -quasi-isometry* if moreover its image $F(X)$ is *L -dense* in Y , i.e. if for every $y \in Y$ there is some $x \in X$ such that $d(Fx, y) \leq L$. An *L -quasi-geodesic* in (X, d) is an L -quasi-isometric embedding of a connected subset of the real line.

The mapping class group \mathcal{MCG} of S is finitely generated [I02] and hence a finite symmetric generating set \mathcal{G} defines a *word norm* $||$ on \mathcal{MCG} by assigning to an element $g \in \mathcal{MCG}$ the minimal length of a word in \mathcal{G} which represents g .

The word norm $||$ determines a distance d on \mathcal{MCG} which is invariant under the action of \mathcal{MCG} by left translation by defining $d(g, h) = |g^{-1}h|$. Any two such distance functions are quasi-isometric. In particular, quasi-isometrically embedded flat subspaces in \mathcal{MCG} are well defined.

To understand the geometry of the mapping class group we introduce a locally finite metric graph \mathcal{TT} whose set $\mathcal{V}(\mathcal{TT})$ of vertices are the isotopy classes of *complete train tracks* on S (see [PH92]) and where such a train track τ is connected to a train track σ by a directed edge of length one if σ can be obtained from τ by a single *split*. We define this graph in Section 3 and show that it is connected. In Section 4 we observe that the mapping class group of S acts properly and cocompactly as a group of simplicial isometries on \mathcal{TT} . Thus by a well known result of Švarc-Milnor (see [BH99]), for every $\tau \in \mathcal{TT}$ the orbit map $\mathcal{MCG} \rightarrow \mathcal{TT}$ which associates to $g \in \mathcal{MCG}$ the point $g\tau \in \mathcal{TT}$ is an \mathcal{MCG} -equivariant quasi-isometry.

The space \mathcal{CL} of all complete geodesic laminations on S is related to the metric graph \mathcal{TT} as follows. In Section 2 we observe that every complete train track τ carries a complete geodesic lamination. For a complete train track τ and a complete geodesic lamination λ carried by τ we then can consider the full subgraph $E(\tau, \lambda)$ of \mathcal{TT} whose set of vertices is the set of all complete train tracks which carry λ and which can be obtained from τ by a directed edge-path in \mathcal{TT} . Note that $E(\tau, \lambda)$ is determined by τ and λ . In Section 5 we show that there is a number $m > 0$ only depending on the topological type of S such that each of these full subgraphs $E(\tau, \lambda)$ of \mathcal{TT} equipped with the intrinsic path metric d_E is isometric to a full subgraph of the standard cubical cone in \mathbb{R}^m whose vertex set is the intersection of the integer lattice \mathbb{Z}^m with the cone $\{(x_1, \dots, x_m) \mid x_i \geq 0\}$. This justifies the name “cubical Euclidean cone” for these spaces.

Our main technical result relates the geometry of these cubical Euclidean cones to the geometry of \mathcal{TT} . We show in Section 6.

Theorem 2. *There is a number $L > 1$ such that for every $\tau \in \mathcal{V}(\mathcal{TT})$ and every complete geodesic lamination λ carried by τ the inclusion $(E(\tau, \lambda), d_E) \rightarrow \mathcal{TT}$ is an L -quasi-isometric embedding.*

As a consequence, a point in the Furstenberg boundary of \mathcal{MCG} can be thought of as a family of uniformly quasi-isometrically embedded subspaces of \mathcal{MCG} with an euclidean structure. Moreover, Theorem 2 gives a fairly precise picture on the large-scale geometry of the mapping class group which can be used to analyze for example the asymptotic cone of \mathcal{MCG} .

As an immediate corollary, we obtain information on uniform quasi-geodesics in \mathcal{MCG} . Namely, define a *splitting arc* in \mathcal{TT} to be a directed simplicial path $\gamma : [0, m] \rightarrow \mathcal{TT}$ for the standard simplicial structure on \mathbb{R} whose vertices are the integers. This means that for every $i < m$ the arc $\gamma[i, i + 1]$ is an edge in \mathcal{TT} connecting the train track $\gamma(i)$ to a train track $\gamma(i + 1)$ which can be obtained from $\gamma(i)$ by a single split. We obtain in Section 6.

Corollary 3. *There is a number $M > 0$ such that every splitting arc in \mathcal{TT} is an M -quasi-geodesic.*

Theorem 2 can also be used to investigate various geometric properties of the mapping class group. As an illustration, we include here a particularly easy corollary.

Namely, a finite symmetric generating set \mathcal{G}_Γ of a finitely generated subgroup $\Gamma < \mathcal{MCG}$ defines a distance function d_Γ on Γ . Since we can always extend \mathcal{G}_Γ to a finite symmetric generating set of \mathcal{MCG} , for every distance function d on \mathcal{MCG} defined by a word norm there is a number $q > 0$ such that the natural inclusion $(\Gamma, d_\Gamma) \rightarrow (\mathcal{MCG}, d)$ is q -Lipschitz. However, in general the word norm in Γ of an element $g \in \Gamma$ can not be estimated from above by a constant multiple of its word norm in \mathcal{MCG} . The group $\Gamma < \mathcal{MCG}$ is called *undistorted* if there is a constant $c > 1$ such that $d_\Gamma(g, h) \leq cd(g, h)$ for all $g, h \in \Gamma$. This is equivalent to stating that the natural inclusion $(\Gamma, d_\Gamma) \rightarrow (\mathcal{MCG}, d)$ is a quasi-isometric embedding.

A *pants decomposition* for S is a collection of $3g - 3 + k$ simple closed mutually disjoint *pants curves* which decompose S into $2g - 2 + k$ *pairs of pants*, i.e. three-holed spheres. Such a pants decomposition determines a free abelian subgroup of \mathcal{MCG} of rank $3g - 3 + k$ which is generated by the *Dehn twists* about the pants curves. As an immediate application of Theorem 3 we obtain a new proof of the following result of Farb, Lubotzky and Minsky [FLM01].

Corollary 4. *For every pants decomposition P for S , the free abelian subgroup of \mathcal{MCG} of rank $3g - 3 + k$ which is generated by the Dehn twists about the pants curves of P is an undistorted subgroup of \mathcal{MCG} .*

Now let $n \geq 2$ and for $i \leq n$ let G_i be a locally compact second countable topological group. A subgroup Γ of the group $G = G_1 \times \cdots \times G_n$ is an *irreducible lattice* in G if the volume of G/Γ with respect to a Haar measure is finite and if moreover the projection of Γ to each factor is dense. We allow the groups G_i to be discrete. Let X be a standard Borel Γ -space and let μ be a Γ -invariant ergodic probability measure on X . The action of Γ on (X, μ) is called *mildly mixing* if there are no non-trivial recurrent sets, i.e. if for any measurable set $A \subset X$ and any sequence $\varphi_i \rightarrow \infty$ in Γ , one has $\mu(\varphi_i A \triangle A) \rightarrow 0$ only when $\mu(A) = 0$ or $\mu(X - A) = 0$. An \mathcal{MCG} -valued *cocycle* for this action is a measurable map $\alpha : \Gamma \times X \rightarrow \mathcal{MCG}$ such that

$$\alpha(gh, x) = \alpha(g, hx)\alpha(h, x)$$

for all $g, h \in \Gamma$ and μ -almost every $x \in X$. The cocycle α is *cohomologous* to a cocycle $\beta : \Gamma \times X \rightarrow \mathcal{MCG}$ if there is a measurable map $\varphi : X \rightarrow \mathcal{MCG}$ such that $\varphi(gx)\alpha(g, x) = \beta(g, x)\varphi(x)$ for all $g \in \Gamma$, μ -almost every $x \in X$.

For every continuous homomorphism ρ of a topological group H into a topological group L , the composition with ρ of an H -valued cocycle is a cocycle with values in L . We use the space of complete geodesic laminations on S to show the following super-rigidity result for \mathcal{MCG} -valued cocycles.

Theorem 3. *Let $n \geq 2$, let $\Gamma < G_1 \times \cdots \times G_n$ be an irreducible lattice, let X be a mildly mixing Γ -space and let $\alpha : \Gamma \times X \rightarrow \mathcal{MCG}$ be any cocycle. Then α is cohomologous to a cocycle α' with values in a subgroup $H = H_0 \times H_1$ of \mathcal{MCG} where H_0 is virtually abelian and where H_1 contains a finite normal subgroup K such that the projection of α' into H_1/K defines a continuous homomorphism $G \rightarrow H_1/K$.*

The organization of this paper is as follows. In Section 2 we introduce in more detail the space of complete geodesic laminations. We also summarize some results on geodesic laminations and train tracks from the literature which are used throughout the paper. As mentioned above, in Section 3 we define the train track complex \mathcal{TT} , and we show that \mathcal{TT} is a connected metric graph. In Section 4 we show that the mapping class groups acts properly and cocompactly as a group of isometries on \mathcal{TT} .

In Section 5 we define the family of cubical Euclidean cones $E(\tau, \lambda) \subset \mathcal{TT}$ which appear in Theorem 2 and investigate their intrinsic geometry. Section 6 contains the proof of Theorem 2 and is the central part of this paper. This immediately implies Corollary 3 and Corollary 4. In Section 7 we use the results from the previous sections to show Theorem 1. In Section 8 we construct an explicit *strong boundary* for the mapping class group and derive Corollary 2. The proof of Theorem 3 is contained in Section 9.

The results in Section 6 depend in an essential way on a technical property of train tracks which is established in the appendix. The appendix only uses some results from the literature which are collected in Section 2, and it is independent from the rest of the paper.

After we completed this work we obtained the preprint of Kida [Ki05] who shows with an inductive argument that the action of the mapping class group on its Stone Čech compactification is topologically amenable. By an observation of Higson and Roe [HR00], this is equivalent to boundary amenability for the mapping class groups. Kida also obtains Corollary 2.

2. TRAIN TRACKS AND GEODESIC LAMINATIONS

This introductory section is divided into two subsections. In the first subsection, we introduce the space of complete geodesic laminations for an oriented surface S of genus $g \geq 0$ with $k \geq 0$ punctures and where $3g - 3 + k \geq 2$. The second part summarizes some properties of train tracks on S which are used throughout the paper.

2.1. Complete geodesic laminations. Fix a complete hyperbolic metric h of finite volume on the surface S . A *geodesic lamination* for the metric h is a *compact* subset of S which is foliated into simple geodesics. Particular geodesic laminations are simple closed geodesics, i.e. laminations which consist of a single leaf.

A geodesic lamination λ is called *minimal* if each of its half-leaves is dense in λ . Thus a geodesic lamination is minimal if it does not contain any proper *sublamination*, i.e. a proper closed subset which is a geodesic lamination. A simple closed geodesic is a minimal geodesic lamination. A minimal geodesic lamination with more than one leaf has uncountably many leaves and is called *minimal arational*. Every geodesic lamination λ is a disjoint union of finitely many minimal sublamination and a finite number of isolated leaves. Each of the isolated leaves of λ either is an isolated closed geodesic and hence a minimal component, or it *spirals*

about one or two minimal components. We refer to [CB88, CEG87] for a detailed discussion of the structure of a geodesic lamination.

A geodesic lamination is *finite* if it contains only finitely many leaves, and this is the case if and only if each minimal component is a closed geodesic. A geodesic lamination is *maximal* if it is not a proper sublamination of another geodesic lamination, and this is the case if and only if all complementary regions are ideal triangles or once punctured monogons [CEG87]. Note that a geodesic lamination can be both minimal and maximal.

Since each geodesic lamination is a compact subset of the surface S , the space of all geodesic laminations can be equipped with the restriction of the *Hausdorff topology* for compact subsets of S . Moreover, the tangent lines of a geodesic lamination define a compact subset of the *projectivized tangent bundle* PTS of S . We therefore can equip the space of geodesic laminations on S with the Hausdorff topology for compact subsets of PTS . However, these two topologies coincide [CB88], and in the sequel we shall freely use these two descriptions interchangeably. With this topology, the space of all geodesic laminations is compact, and it contains the space of all maximal geodesic laminations as a compact subset [CB88].

Definition 2.1. A *complete geodesic lamination* is a maximal geodesic lamination which can be approximated in the Hausdorff topology by simple closed geodesics.


Since every minimal geodesic lamination can be approximated in the Hausdorff topology by simple closed geodesics [CEG87], laminations which are both maximal and minimal are complete. There are also complete finite geodesic laminations. Namely, let P be a *pants decomposition* of S , i.e. P is the union of $3g - 3 + m$ pairwise disjoint simple closed geodesics which decompose S into $2g - 2 + m$ planar surfaces of Euler characteristic -1 . A geodesic lamination λ with the following properties is complete.

- (1) The pants decomposition P is the union of the minimal components of λ .
- (2) For each component Q of $S - P$ and every pair $\gamma_1 \neq \gamma_2$ of boundary geodesics of Q there is a leaf of λ contained in Q which spirals in one direction about γ_1 , in the other direction about γ_2 .
- (3) Each component Q of $S - P$ containing a puncture of S contains a leaf of λ which goes around a puncture and spirals in both directions about the same boundary component of Q .
- (4) For every component γ of P , the leaves of λ which spiral about γ from each side of γ define opposite orientations near γ as in Figure A.

The fourth condition guarantees that we can approximate λ by smooth simple closed curves which pass from a leaf spiraling about a component γ of P from one side to a leaf spiraling about γ from the other side and whose tangents are close to the tangent lines of λ .

Since maximal geodesic laminations form a closed subset of the space of all geodesic laminations, the set \mathcal{CL} of all complete geodesic laminations is a closed subset of the space of all geodesic laminations with the Hausdorff topology. Thus

Figure A



\mathcal{CL} is a compact topological space. A *simple geodesic multi-curve* is a disjoint union of simple closed geodesics. We have.

- Lemma 2.2.** (1) *Every geodesic lamination on S which can be approximated in the Hausdorff topology by simple geodesic multi-curves is a sublamination of a complete geodesic lamination.*
- (2) *Finite complete geodesic laminations are dense in the space of all complete geodesic laminations.*

Proof. Let λ be any geodesic lamination on S which can be approximated in the Hausdorff topology by a sequence $\{c_i\}$ of simple geodesic multi-curves. For each i let μ_i be a complete finite geodesic lamination which contains c_i as a union of minimal components. Such a geodesic lamination exists by the above example since each simple geodesic multi-curve is a subset of a pants decomposition for S . By passing to a subsequence we may assume that the geodesic laminations μ_i converge as $i \rightarrow \infty$ in the Hausdorff topology to a complete geodesic lamination μ . Since $c_i \rightarrow \lambda$ ($i \rightarrow \infty$) in the Hausdorff topology, the lamination μ contains λ as a sublamination. This shows the first part of the lemma.

Since by definition a complete geodesic lamination can be approximated in the Hausdorff topology by simple closed geodesics, the second part follows from the same argument together with the observation that a geodesic lamination which contains a complete geodesic lamination λ as a sublamination coincides with λ . \square

2.2. Complete train tracks. 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 of τ are divided into two classes according to the orientation of the 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.

Train tracks were invented by Thurston [T79] and provide a powerful tool for the investigation of surfaces and hyperbolic 3-manifolds. A detailed account on train tracks can be found in the book [PH92] of Penner with Harer which we use as our main reference. The more recent unpublished manuscript [M03] of Mosher contains a discussion of train tracks from a somewhat different viewpoint, however it will not be used in this paper.

A *trainpath* on a train track τ is a C^1 -immersion $\rho : [m, n] \rightarrow \tau \subset S$ which maps each interval $[k, k + 1]$ ($m \leq k \leq n - 1$) onto a branch of τ . The integer $n - m$ is then called the *length* of ρ . We sometimes identify a trainpath on S with its image in τ . Each complementary region of τ is bounded by a finite number of trainpaths which either are simple closed curves or terminate at the cusps of the region.

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 $c = c(b) \subset S$ 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.

Remark: 1) We chose to use transversely recurrent train tracks for our purpose even though this property is nowhere needed. The main reason for using transversely recurrent train tracks is convenience of reference to the existing literature.

2) Throughout the paper, we require every train track to be generic, and this is indeed necessary for many of our constructions. Unfortunately this leads to a slight inconsistency of our terminology with the terminology found in the literature.

Every generic train track τ on S is contained in a closed subset A of S with dense interior and piecewise smooth boundary which is foliated by smooth arcs transverse to the branches of τ . These arcs are called *ties* for τ , and the set A is called a *foliated neighborhood* of τ (even though A is not a neighborhood of τ in the usual sense). Each of the ties intersects τ in a single point, and the switches of τ are the intersection points of τ with the singular ties, i.e. the ties which are not contained in the interior of a foliated rectangle $R \subset A$. Collapsing each tie to its intersection point with τ defines a map $F : A \rightarrow \tau$ of class C^1 which we call a *collapsing map*. The map F is the restriction to A of a map $\tilde{F} : S \rightarrow S$ of class C^1 which is homotopic to the identity.

A train track or a geodesic lamination λ is *carried* by a transversely recurrent train track τ if there is an isotopy φ of S such that λ is contained in a foliated neighborhood A of $\varphi(\tau)$ and is transverse to the ties. This is equivalent to stating that there is a map $G : S \rightarrow S$ of class C^1 which is homotopic to the identity and which maps λ to τ in such a way that the restriction of the differential of G to every tangent space of λ is non-singular. Note that this makes sense since a train track has a tangent line everywhere (compare Theorem 1.6.6 of [PH92]). We then call the restriction of G to λ a *carrying map*. Every train track τ which carries a maximal

geodesic lamination is necessarily maximal [PH92]. The set of geodesic laminations which are carried by a transversely recurrent train track τ is a closed subset of the space of all geodesic laminations on S with respect to the Hausdorff topology (Theorem 1.5.4 of [PH92], or see [CB88] for the same result for train tracks which are not necessarily transversely recurrent).

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 at s are divided into two classes according to the orientation of their inward pointing tangent at s . We require that the sums of the weights over all 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 $\mu > 0$. The set $V(\tau)$ of all transverse measures on τ is a closed convex cone in a linear space and hence topologically it is a closed cell. For every recurrent train track τ , positive measures define the interior of the convex cone $V(\tau)$. A train track τ is called *birecurrent* if τ is recurrent and transversely recurrent.

A *measured geodesic lamination* is a geodesic lamination equipped with a transverse Borel measure of full support which is invariant under holonomy [PH92]. A measured geodesic lamination can be viewed as a locally finite Borel measure on the space S_∞ of unoriented geodesics in the universal covering \mathbf{H}^2 of S which is invariant under the action of $\pi_1(S)$ and supported in the closed set of geodesics whose endpoints are not separated under the action of $\pi_1(S)$. The weak*-topology on the space of $\pi_1(S)$ -invariant locally finite Borel measures on S_∞ then restricts to a natural topology on the space \mathcal{ML} of all measured geodesic laminations.

If λ is a geodesic lamination with transverse measure μ and if λ is carried by a train track τ , then the transverse measure μ induces via a carrying map $\lambda \rightarrow \tau$ a transverse measure on τ , and every transverse measure on τ arises in this way (Theorem 1.7.12 in [PH92]). Moreover, if τ is maximal and birecurrent, then the set of measured geodesic laminations which correspond to positive transverse measures on τ in this way is an open subset U of \mathcal{ML} (Lemma 3.1.2 in [PH92]).

We use measured geodesic laminations to relate maximal birecurrent train tracks to complete geodesic laminations. Note that the first part of the following lemma is well known and reflects the fact that the space \mathcal{CL} of complete geodesic laminations is totally disconnected. We refer to [Bo97] and [ZB04] for more details and references about the Hausdorff topology on the space of all geodesic laminations.

Lemma 2.3. *Let τ be a maximal transversely recurrent train track. Then the set of all complete geodesic laminations on S which are carried by τ is open and closed in \mathcal{CL} . This set is non-empty if and only if τ is recurrent.*

Proof. Let τ be a maximal transversely recurrent train track on S . Then the subset of \mathcal{CL} of all complete geodesic laminations which are carried by τ is closed (Theorem 1.5.4 of [PH92]).

On the other hand, if $\lambda \in \mathcal{CL}$ is carried by τ then after possibly modifying τ with an isotopy we may assume that there is a foliated neighborhood A of τ with

collapsing map $F : A \rightarrow \tau$ which contains λ in its interior and such that λ is transverse to the ties of A . Since the space of tangent lines $PT\lambda$ of λ is a compact subset of the projectivized tangent bundle PTS of S , there is a neighborhood U of $PT\lambda$ in PTS which is mapped by the canonical projection $PTS \rightarrow S$ into A and such that the differential of F is non-singular on each line in U . As a consequence, a geodesic lamination whose space of tangent lines is contained in U is carried by τ . Since the Hausdorff topology on the space of geodesic laminations coincides with the Hausdorff topology for their projectivized tangent bundles, the set of all such geodesic laminations which are moreover complete is a neighborhood of λ in \mathcal{CL} . This shows that the set of all complete geodesic laminations which are carried by τ is open and closed in \mathcal{CL} .

To show the second part of the lemma, let τ be a maximal transversely recurrent train track on S which carries a complete geodesic lamination $\lambda \in \mathcal{CL}$. To see that τ is recurrent, let $\{c_i\}$ be a sequence of simple closed geodesics which approximate λ in the Hausdorff topology. By the above consideration, the curves c_i are carried by τ for all sufficiently large i . Let $F_i : c_i \rightarrow \tau$ be a carrying map. Then for each sufficiently large i , the curve c_i defines a *counting measure* on τ by associating to each branch b of τ the number of components of $F_i^{-1}(b)$. This counting measure satisfies the switch condition. Now a carrying map $F : \lambda \rightarrow \tau$ maps λ *onto* τ since λ is maximal, and hence the same is true for the carrying map $F_i : c_i \rightarrow \tau$ provided that i is sufficiently large. Thus for sufficiently large i , the counting measure defined by c_i is positive and therefore τ is recurrent.

To show that a maximal birecurrent train track carries a complete geodesic lamination, note that a geodesic lamination which is both minimal and maximal is complete and supports a transverse measure [CEG87]. Since the set of all measured geodesic laminations carried by a maximal birecurrent train track τ has non-empty interior, it is enough to show that the set of measured geodesic laminations whose support is minimal and maximal is dense in the space \mathcal{ML} of all measured geodesic laminations.

However, as was pointed out to me by McMullen, this follows for example from the work of Kerckhoff, Masur and Smillie [KMS86]. Namely, the *vertical foliation* of a *quadratic differential* defines a measured geodesic lamination, and every measured geodesic lamination is of this form [T79]. The set of measured geodesic laminations which arise from quadratic differentials with only simple zeros is dense in the space of all measured geodesic laminations (compare the discussion in [KMS86]). Given such a quadratic differential q , there are only countably many $\theta \in [0, 2\pi]$ such that there is a vertical *saddle connection* for $e^{i\theta}q$, i.e. compact vertical arc connecting two zeros of q . On the other hand, for each quadratic differential q the set of all $\theta \in [0, 2\pi]$ such that the vertical foliation of $e^{i\theta}q$ is *uniquely ergodic* and hence minimal has full Lebesgue measure. Now if the vertical foliation of $e^{i\theta}q$ is uniquely ergodic and without saddle connection, then this foliation is both minimal and maximal. Thus there is a dense set of points $\theta \in [0, 2\pi]$ with the property that the vertical foliation of $e^{i\theta}q$ is both minimal and maximal.

As a consequence, a maximal birecurrent train track carries a geodesic lamination which is both maximal and minimal and hence complete. \square


Definition 2.4. A train track τ on S which is generic, maximal and birecurrent is called *complete*.

Note that this definition of a complete train track slightly differs from the one in [PH92] since we require a complete train track to be generic. By Lemma 2.3, a generic transversely recurrent train track is complete if and only if it carries a complete geodesic lamination.

There are two basic ways to modify a complete train track to another complete train track. Namely, let \tilde{b} be a half-branch of a generic train track τ and let v be a trivalent switch of τ on which \tilde{b} is incident. Then \tilde{b} is called *large* if any immersed arc of class C^1 in τ passing through v intersects \tilde{b} . The branch b of τ containing \tilde{b} is called *large at v* . 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. A large branch b is necessarily incident on two distinct switches, and it is large at both of them. 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 (for all this, see [PH92] p.118).

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


Figure B



train track 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 . Note that there is a natural bijection of the set of branches of τ onto the set of branches of τ' .

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 C. Note that a right split at e is uniquely determined by

Figure C



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 (see [MM99]). The split τ' of a train track τ is carried by τ . There is a natural bijection of the set of branches of τ onto the set of branches of τ' which maps the branch e to the diagonal e' of the split. The split of a generic maximal transversely recurrent train track is generic, maximal and transversely recurrent. If τ is complete and if $\lambda \in \mathcal{CL}$ is carried by τ , then there is a unique choice of a right or left split of τ at e with the property that the split track τ' carries λ . We call τ' the λ -*split* of τ . By Lemma 2.3, the train track τ' is recurrent and hence 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 track is not recurrent any more (compare p.120 in [PH92]). The reverse of a split is called a *collapse*.

3. THE COMPLEX OF TRAIN TRACKS

Let as before S be an oriented surface of genus $g \geq 0$ with $k \geq 0$ punctures and where $3g - 3 + k \geq 2$. Using the notations from Section 2, define \mathcal{TT} to be the directed graph whose set of vertices is the set $\mathcal{V}(\mathcal{TT})$ of isotopy classes of complete train tracks on S and whose edges are determined as follows. The train track τ is connected in \mathcal{TT} to the train track τ' by a directed edge if and only if τ' can be obtained from τ by a single right or left split. The goal of this section is to show that \mathcal{TT} is connected.

The main ingredient for the proof of this result is Corollary 2.4.3 of [PH92] which can in our context be stated as follows:

Proposition 3.1. *Let $\tau, \sigma \in \mathcal{V}(\mathcal{TT})$ and assume that τ carries σ and that σ carries a geodesic lamination λ which is both minimal and maximal. Then there is a complete train track η which carries λ , and there are directed edge-paths in \mathcal{TT} connecting τ, σ to η .*

Namely, a geodesic lamination λ which is both minimal and maximal is the support of a transverse measure μ . Let $\tau, \sigma \in \mathcal{V}(\mathcal{TT})$, and assume that τ carries σ and that σ carries λ . By Corollary 2.4.3 of [PH92], there is a birecurrent train track η which carries the measured geodesic lamination defined by (λ, μ) and such that both τ and σ can be connected to η by a sequence of splits (where a splitting move in the sense of Penner and Harer may be a *collision* which is defined to be a split followed by the removal of the diagonal of the split). Since the support λ of the measured geodesic lamination μ is a maximal geodesic lamination, the train track η is maximal and hence complete (and no collision can have occurred in the process). As a consequence, τ and σ can be connected by a path in \mathcal{TT} which is the concatenation of a directed edge path connecting τ to η and the inverse of a directed edge path connecting σ to η .

To exploit this result, we verify three properties of complete train tracks from which connectedness of \mathcal{TT} easily follows.

- (1) Any complete train track carries a geodesic lamination which is both minimal and maximal.
- (2) Two complete train tracks which carry a common complete geodesic lamination λ also carry a common complete train track which “approximates” λ in the C^1 -topology.
- (3) For any two minimal and maximal geodesic laminations μ, λ there is a complete train track which carries both λ and μ .

The first property is well known, but hard to trace in the literature. The second property is more subtle. The notion of approximation in the C^1 -topology of a geodesic lamination λ by train tracks we use is the approximation of the projectivized tangent bundle of λ in the Hausdorff topology for compact subsets of the projectivized tangent bundle PTS of S .

Given a complete hyperbolic metric of finite volume on S , such an approximation can be made quantitative by using the metric to determine a geometric representative for a given train track τ and compare its projectivized tangent bundle to the projectivized tangent bundle of λ . The representative of τ we use is its *straightening* with respect to the hyperbolic metric on S which is defined to be the 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 . The straightening of a train track τ is in general not a train track. Since carrying is a relation determined by maps of class C^1 , with restrictions on the tangent map, we observe later in this section that if τ is a complete train track on S which carries a complete geodesic lamination λ and if the projectivized tangent bundle of the straightening of a train track η is sufficiently close to the projectivized tangent bundle of λ and does not have a short branch which makes a sharp turn in the wrong direction then η is carried by τ and carries λ . Note that this is not true for train tracks η which are merely close to λ in the Hausdorff topology for compact subsets of S .

Fix again a complete hyperbolic metric of finite volume of S . For a given train track τ , the tangent lines of the straightening of τ define a closed subset of the projectivized tangent bundle PTS of S . The hyperbolic metric on S naturally induces a Riemannian metric and hence a distance function on PTS . For a number $\epsilon > 0$ we say that the train track τ *ϵ -follows* the 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 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 (here a vanishing exterior angle means that the arc is smooth—this condition prevents sharp turns in the wrong direction). The train track τ is *a -long* for a number $a > 0$ if the length of every edge of the straightening of τ is at least a .

The following technical lemma is a strengthening of Theorem 1.6.5 of [PH92] which follows from the same line of arguments. It shows that a geodesic lamination λ can be approximated in the sense described above by train tracks which carry λ .

Lemma 3.2. *There is a number $a > 0$ with the following property. Let λ be any geodesic lamination on S . Then for every $\epsilon > 0$ there is an a -long generic transversely recurrent train track τ which carries λ and which ϵ -follows λ .*

Proof. By the collar theorem for hyperbolic surfaces (see [B92]), a complete simple geodesic on S which is contained in a compact subset of S does not enter deeply into a cusp. This means that there is a compact bordered subsurface S_0 of S which contains the 1-neighborhood of every geodesic lamination on S . Let $k > 0$ be the maximal number of branches of any train track on S (this only depends on the topological type of S) and let $a \in (0, 1)$ be sufficiently small that $(5k + 2)a$ is smaller than the smallest length of any non-contractible closed curve in S_0 .

For $\epsilon > 0$ there is a number $\delta_0 = \delta_0(\epsilon) \in (0, a/2)$ with the following property. Let γ be a geodesic line in the hyperbolic plane \mathbf{H}^2 and let ζ be an arc in \mathbf{H}^2 which consists of two geodesic segments of length at least a . Assume that ζ is contained in the δ_0 -neighborhood of γ and that the length of ζ does not exceed the sum of the distance of its endpoints and $4\delta_0$. Then the tangent lines of ζ are contained in the ϵ -neighborhood of the tangent lines of γ as subsets of the projectivized tangent bundle of \mathbf{H}^2 , and the exterior angle at the breakpoint of ζ is at most ϵ .

We now use the arguments of Casson and Bleiler as described in the proof of Theorem 1.6.5 of [PH92]. Let λ be a geodesic lamination on S and for $\delta > 0$ let N_δ be the closed δ -neighborhood of λ in S . The preimage of this neighborhood in the hyperbolic plane \mathbf{H}^2 is the δ -neighborhood \tilde{N}_δ of the preimage $\tilde{\lambda}$ of λ . For sufficiently small δ , say for all $\delta < \delta_1$, each component of $\mathbf{H}^2 - \tilde{\lambda}$ contains precisely one component of $\mathbf{H}^2 - \tilde{N}_\delta$. These components are polygons whose sides are arcs of constant geodesic curvature, and this curvature tends to 0 with δ .

Let $\delta < \min\{\delta_0, \delta_1\}$ be sufficiently small that N_δ can be foliated by smooth arcs which are transverse to the leaves of λ and of length smaller than δ_0 . We may assume that the arcs through the finitely many corners of the complementary components of N_δ are geodesics (p.74-75 in [PH92]). We call this foliation of N_δ the *vertical foliation*, and we denote it by \mathcal{F}^\perp . The singular leaves of \mathcal{F}^\perp , i.e. the leaves through the corners of $S - N_\delta$, decompose N_δ into closed foliated rectangles with embedded interior. Two opposite sides of these rectangles are subarcs of singular leaves of the vertical foliation and hence they are geodesics. The other two sides are arcs of constant curvature. The intersection of any two distinct such rectangles is contained in the singular leaves of \mathcal{F}^\perp . Moreover, the rectangles are projections to S of *convex* subsets of the hyperbolic plane. After a small adjustment near the corners of the complementary components of N_δ , collapsing each leaf of \mathcal{F}^\perp to a suitably chosen point in its interior defines a train track τ on S and a map $F : N_\delta \rightarrow \tau$ of class C^1 which is homotopic to the identity and whose restriction to λ is a carrying map $\lambda \rightarrow \tau$. The switches of τ are precisely the collapses of the singular leaves of \mathcal{F}^\perp , and each branch of τ is a collapse of one of the foliated rectangles. The set N_δ is a foliated neighborhood of τ . Every edge of the straightening of τ is a geodesic arc which is contained in one of the rectangles and connects the two sides of the rectangle contained in leaves of \mathcal{F}^\perp . Via slightly changing δ we may assume that τ is generic; this is equivalent to saying that each singular leaf of \mathcal{F}^\perp contains

precisely one corner of a component of $S - N_\delta$. The train track τ is transversely recurrent [PH92].

Define the *length* $\ell(R)$ of a foliated rectangle $R \subset N_\delta$ as above to be the *intrinsic* distance in R between its two sides which are contained in \mathcal{F}^\perp . Since a rectangle R lifts to a convex subset \tilde{R} of the hyperbolic plane, this length is just the distance between the two geodesic sides of \tilde{R} . Let $R_1 \neq R_2$ be two rectangles which intersect along a nontrivial subarc c of a singular leaf of \mathcal{F}^\perp . Let a_i be the side of R_i opposite to the side containing c . Then $R_1 \cup R_2$ contains a geodesic segment γ which connects a_1 to a_2 and intersects c . The segment γ is a subarc of a leaf of λ , and its length is not smaller than $\ell(R_1) + \ell(R_2)$. Hence the length of any curve in S with one endpoint in a_1 and the second endpoint in a_2 which is homotopic to γ relative to $a_1 \cup a_2$ is not smaller than $\ell(R_1) + \ell(R_2) - 2\delta_0$. On the other hand, the length of a geodesic arc contained in R_i and connecting a_i to c which is homotopic to $\gamma \cap R_i$ relative to $a_i \cup c$ is not bigger than $\ell(R_i) + 2\delta_0$.

Let ζ be the subarc of the straightening of τ which consists of the straightening of the collapses of the rectangles R_1 and R_2 . By convexity, $\zeta \cap R_i$ is contained in the δ_0 -neighborhood of the geodesic arc $\gamma \cap R_i$, and the length of ζ does not exceed the sum of the length of the geodesic homotopic to ζ with fixed endpoints and $4\delta_0$. Thus by the choice of δ_0 , the train track τ ϵ -follows λ provided that the length of each edge from the straightening of τ is at least a , and this is the case if the length of each of the rectangles is at least a .

We now successively modify the train track τ to a train track τ' which is embedded in N_δ and transverse to the leaves of the vertical foliation, which carries λ and such that the length of each edge from the straightening of τ' is at least a as follows. Call a rectangle R *short* if its length is at most a . Let $n \geq 0$ be the number of short rectangles in N_δ . If $n = 0$ then τ is a -long, so assume that $n > 0$. We remove from N_δ a suitably chosen geodesic arc β_0 so that $N_\delta - \beta_0$ is partitioned into rectangles which are foliated by the restriction of the vertical foliation \mathcal{F}^\perp and for which the number of short rectangles is at most $n - 1$.


This *unzipping* of the train track τ ([PH92] p.74-75) is done as follows. Let $R \subset N_\delta$ be a short rectangle. The boundary ∂R of R contains a corner x of a complementary component T of N_δ . This corner projects to a switch $F(x)$ of τ , and the collapse of R is incident on $F(x)$. Let \tilde{T} be a lift of the complementary component T to \mathbf{H}^2 and let \tilde{x} be the lift of x to \tilde{T} . Let $\tilde{\lambda}$ be the lift of λ to \mathbf{H}^2 ; then \tilde{T} is contained in a unique component C of $\mathbf{H}^2 - \tilde{\lambda}$. The point \tilde{x} is at distance δ to two frontier leaves of C . Since every complementary component of $\tilde{\lambda}$ contains precisely one complementary component of $\mathbf{H}^2 - \tilde{N}_\delta$, these leaves have a common endpoint ξ in the ideal boundary of \mathbf{H}^2 . There is a unique geodesic ray $\tilde{\beta}$ in \mathbf{H}^2 which connects \tilde{x} to ξ , and this ray is entirely contained in the intersection of \tilde{N}_δ with the component C of $\mathbf{H}^2 - \tilde{\lambda}$. The projection of $\tilde{\beta}$ to S is a one-sided infinite simple geodesic β beginning at x which is contained in N_δ and is disjoint from λ . We may assume that β is parametrized by arc length and is everywhere transverse to the leaves of the foliation \mathcal{F}^\perp . In particular, the collapsing map F maps β up to parametrization to a one-sided infinite trainpath on τ .

We claim that the finite subarc β_0 of β of length $5ka$ which begins at x is mapped by F injectively into τ . For this note that by the choice of the constant $a > 0$, the length of the union of β_0 with any leaf of the vertical foliation \mathcal{F}^\perp is smaller than the minimal length of a non-contractible closed curve in the compact surface $S_0 \supset N_\delta$. If β_0 intersects the same leaf of the vertical foliation twice then there is a subarc β_1 of β_0 whose concatenation with a subarc of a leaf of \mathcal{F}^\perp is a closed curve c of length smaller than $5ka + a$ which is contained in $N_\delta \subset S_0$. By the choice of the constant a , the curve c is contractible in S . On the other hand, c is freely homotopic to the image of β_1 under the collapsing map F and hence to a subarc of a trainpath on τ which begins and ends at the same point (the induced orientations of its tangent line at that point may be opposite). By the definition of a train track, such a curve is homotopically nontrivial in S which is a contradiction.

Therefore β_0 intersects each leaf of the vertical foliation \mathcal{F}^\perp at most once and hence it is mapped by F injectively into τ . The length of the intersection of β_0 with a rectangle R' from our system of rectangles is at most $\ell(R') + 2\delta_0 \leq \ell(R') + a$. On the other hand, there are at most k distinct rectangles in our system of rectangles and hence β_0 intersects the interior of a rectangle of length at least $4a$.

Let $t \in [0, 5ka]$ be the infimum of all numbers $s > 0$ such that $\beta(s)$ is contained in the interior of a rectangle \hat{R} of length at least $4a$. Then $F(\beta(t))$ is a switch in τ , and $\beta(t+2a)$ is an interior point of \hat{R} . We may assume that the leaf of the vertical foliation through $\beta(t+2a)$ is a geodesic. Then this leaf subdivides the rectangle \hat{R} into two foliated rectangles of length at least a . Cut N_δ open along $\beta[0, t+2a)$ as shown in Figure D. Since β is transverse to the vertical foliation \mathcal{F}^\perp , the foliation

Figure D



\mathcal{F}^\perp restricts to a foliation of $N_\delta - \beta[0, t+2a)$. Now β is disjoint from λ , and hence collapsing each leaf of $\mathcal{F}^\perp|_{N_\delta - \beta[0, t+2a)}$ to a suitably chosen point in its interior yields a train track σ which carries λ . The switches of σ are the collapses of the singular leaves of $\mathcal{F}^\perp|_{N_\delta - \beta[0, t+2a)}$. Thus $N_\delta - \beta[0, t+2a)$ is partitioned into a finite number of rectangles as before. Each of these rectangles lifts to a convex subset of the hyperbolic plane. The restriction to σ of the collapsing map F defines a map $\sigma \rightarrow \tau$ with the property that the preimage of every switch $w \neq F(x)$ of τ consists of exactly one switch of σ . Note that σ is transversely recurrent [PH92].

We claim that the number of short rectangles of $N_\delta - \beta[0, t+2a)$ is at most $n-1$. For this let e be the branch of τ which is incident and large at the switch $F(x)$ and write $R_0 = F^{-1}(e)$. Also denote by R_{-1}, R_{-2} the (not necessarily distinct) rectangles which are mapped by F onto the branches of τ which are incident and small

at $F(x)$. These rectangles are properly contained in rectangles $A(R_{-1}), A(R_{-2})$ in $N_\delta - \beta[0, t + 2a)$. There is a short rectangle among the rectangles R_{-2}, R_{-1}, R_0 .

Assume that for some $q \geq 0$, the path $F(\beta[0, t + 2a))$ on τ passes through the switches $F(x), v_1, \dots, v_q = F(\beta(t))$ of τ in this order. If $q = 0$ then we have $\ell(A(R_{-i})) \geq \ell(R_{-i}) + a > a$ for $i = 1, 2$. Since every rectangle in N_δ different from R_i for $i = -2, -1, 0$ is also a rectangle in $N_\delta - \beta[0, t + 2a)$, in this case the claim is obvious.

In the case $q \geq 1$, we obtain the same conclusion as follows. For $i = 1, 2$ let as before $A(R_{-i})$ be the rectangle in $N_\delta - \beta_0[t, 2a)$ which contains R_{-i} . Note that by the definition of length and by convexity, we have

$$(1) \quad \ell(A(R_{-i})) \geq \ell(R_{-i}) + \ell(R_0).$$

We extend the map A to a bijection from the rectangles in N_δ to the rectangles in $N_\delta - \beta[0, t + 2a)$ and compare the length of a rectangle \tilde{R} in N_δ to the length of $A(\tilde{R})$ in $N_\delta - \beta_0[0, t + 2a)$ as follows. For $1 \leq i \leq q - 1$ let R_i be the rectangle in N_δ which is mapped by F to the branch of τ which is incident on the switches v_i and v_{i+1} and is crossed through by the path $F(\beta[0, t + 2a))$. For each $i \in \{1, \dots, q - 1\}$, the point $F^{-1}(v_i) \cap \sigma$ is a switch in σ . Thus there is a unique rectangle $A(R_i)$ in $N_\delta - \beta[0, t + 2a)$ which is mapped by F to a trainpath $\rho : [0, s] \rightarrow \tau$ with $\rho(0) = v_i$ and $\rho[0, 1] = F(R_i)$. Also let R_q be the rectangle \tilde{R} in N_δ which contains $\beta(t + 2a)$ in its interior, and let $A(R_q)$ be the rectangle in $N_\delta - \beta[0, t + 2a)$ which is mapped by F onto the arc $F(\beta[t, t + 2a])$. Finally, the rectangle R_0 is mapped by A to the rectangle which is contained in R_q and whose collapse to a branch of σ is large at the switch which is the collapse of $\beta_0(t + 2a)$.

Note that for $i \in 1 \leq i \leq q - 1$ the length of $A(R_i)$ is not smaller than the length of R_i , and the length of $A(R_0)$ and of $A(R_q)$ is at least a . Hence for every rectangle \tilde{R} in N_δ different from R_0, R_q the length of $A(\tilde{R})$ is not smaller than the length of \tilde{R} . Now if the rectangle R_0 is short, then a rectangle $A(\tilde{R})$ can only be short if \tilde{R} is short and different from R_0 . On the other hand, if the length of the rectangle R_0 is at least a , then at least one of the rectangles R_{-1}, R_{-2} is short. Moreover, by inequality (1), the length of each of the rectangles $A(R_i)$ for $i = -2, -1, 0, q$ is at least a and hence a rectangle $A(\tilde{R})$ can only be short if \tilde{R} is short and different from both R_{-1}, R_{-2} . This shows the above claim.

We can repeat this construction with the train track σ and the foliated set $N_\delta - \beta[0, t + 2a)$. Namely, a short rectangle in $N_\delta - \beta[0, t + 2a)$ does not contain the point $\beta(t + 2a)$ in its closure. Thus if R' is any short rectangle in $N_\delta - \beta[0, t + 2a)$, then its boundary contains a corner of a complementary component of $S - N_\delta$ as before, and this corner is the starting point of a one-sided infinite simple geodesic β' which is contained in N_δ and is disjoint from both λ and β . As a consequence, β' is contained in $N_\delta - \beta[0, t + 2a)$ and we can use the above construction for β' and $N_\delta - \beta[0, t + 2a)$ to reduce the number of short rectangles of $N_\delta - \beta[0, t + 2a)$.

In this way we construct inductively in a uniformly bounded number of steps a train track η which is embedded in N_δ , which carries λ and such that the length of the straightening of each branch of η is at least a . Then by the choice of δ , the train track η is a -long and ϵ -follows λ . This shows the lemma. \square

The next lemma shows how C^1 -approximation of geodesic laminations by train tracks relates to carrying and establishes together with Lemma 3.2 the second property of complete train tracks listed in the beginning of this section.

Lemma 3.3. *Let λ be a geodesic lamination and let τ be a train track which carries λ . Then there is a number $\epsilon > 0$ such that τ carries each train track σ which ϵ -follows λ .*

Proof. Let τ be a train track which carries the geodesic lamination λ . By Theorem 1.6.6 of [PH92], after possibly changing τ by an isotopy we may assume that λ is contained in the interior of a foliated neighborhood A of τ and is transverse to the ties. The restriction of the differential dF of the collapsing map $F : A \rightarrow \tau$ to each tangent line of λ is nonsingular. Since λ is a compact subset of S and the projectivized tangent bundle $PT\lambda$ of λ is a compact subset of the projectivized tangent bundle PTS of S , there is a number $\epsilon > 0$ such that A contains the ϵ -neighborhood of λ and that moreover the restriction of dF to each line $z \in PTS$ which is contained in the ϵ -neighborhood of $PT\lambda$ is nonsingular.

The straightening σ of train track ζ which $\epsilon/2$ -follows λ is embedded in A and is transverse to the ties. The collapsing map F restricts to a map on σ which maps σ to τ . Its differential maps each tangent line of σ onto a tangent line of τ . If $\rho \subset \sigma$ is the straightening of any trainpath on ζ , then the exterior angles at the breakpoints of ρ are at most $\epsilon/2$. Thus we can smoothen the graph σ near its vertices to a train track σ' which is isotopic to ζ and embedded in A and such that the restriction of dF to each tangent line of σ' is nonsingular. But this just means that σ' is carried by τ , with carrying map $F|_{\sigma'}$. This shows the lemma. \square

The space \mathcal{PML} of *projective measured geodesic laminations* on S is the quotient of the space \mathcal{ML} of measured geodesic laminations under the natural action of the multiplicative group $(0, \infty)$. The space \mathcal{PML} will be equipped with the quotient topology. With this topology, \mathcal{PML} is homeomorphic to a sphere [FLP91]. The mapping class group acts naturally on \mathcal{PML} as a group of homeomorphisms. We use this action to establish the third property from the beginning of this section.

Lemma 3.4. *Let λ, μ be any two complete geodesic laminations. Then there is a complete train track τ which carries both λ, μ .*

Proof. The mapping class group \mathcal{MCG} acts on the space of isotopy classes of complete train tracks, and it acts as a group of homeomorphisms on the space \mathcal{PML} of projective measured geodesic laminations. Every *pseudo-Anosov* element $g \in \mathcal{MCG}$ admits a pair of fixed points $\alpha_+ \neq \alpha_-$ in \mathcal{PML} and acts with respect to these fixed points with *south-sink dynamics*: For every neighborhood V of α_+ and every neighborhood W of α_- there is a number $k > 0$ such that $g^k(\mathcal{PML} - W) \subset V$ and $g^{-k}(\mathcal{PML} - V) \subset W$. Then α_+ is called the *attracting fixed point* of g .

The fixed points of pseudo-Anosov elements in \mathcal{MCG} are projective measured geodesic laminations whose support is minimal and *fills up* S . This means that the complementary components of this support are topological discs or once punctured

topological discs. As a consequence, every complete geodesic lamination which contains the support of such a fixed point as a sublamination consists of this support and finitely many isolated leaves which subdivide those complementary components which are ideal polygons with more than 3 sides or punctured polygons with more than one side. In particular, there are only finitely many complete geodesic laminations of this form.

For every complete train track τ on S there is an open subset U of \mathcal{PML} with the property that the support of each measured geodesic lamination $\lambda \in U$ is carried by τ [PH92]. The set U corresponds precisely to those projective measured geodesic laminations which are defined by positive projective transverse measures on τ (Theorem 1.7.12 in [PH92]).

Let $\lambda, \mu \in \mathcal{CL}$ be any two complete geodesic laminations. Since fixed points of pseudo-Anosov elements in \mathcal{MCG} are dense in \mathcal{PML} there is a pseudo-Anosov element $g \in \mathcal{MCG}$ with the following properties.

- (1) The attracting fixed point $\alpha_+ \in \mathcal{PML}$ of g is contained in the set U .
- (2) Every leaf of the support of the repelling fixed point $\alpha_- \in \mathcal{PML}$ of g intersects every leaf of λ, μ transversely.

By Lemma 2.2, there is a complete geodesic lamination which contains the support of α_+ as a sublamination. We claim that each such complete geodesic lamination ν is carried by τ . Namely, choose a sequence of simple closed geodesics $\{c_i\}$ which converge in the Hausdorff topology to ν . Then each of the curves c_i defines a projective measured geodesic lamination. After passing to a subsequence we may assume that these projective measured geodesic laminations converge as $i \rightarrow \infty$ in \mathcal{PML} to a projective measured geodesic lamination $\zeta \in \mathcal{PML}$ whose support is necessarily a sublamination of ν . Since the support of α_+ is the only minimal component of ν , the support of ζ coincides with the support of α_+ . Now the support of α_+ is *uniquely ergodic* which means that there is a single projective measured geodesic lamination whose support equals the support of α_+ [FLP91]. Thus $\zeta = \alpha_+$ and the projective measured geodesic laminations defined by c_i converge as $i \rightarrow \infty$ to α_+ . In particular, for sufficiently large i these projective measured geodesic laminations are contained in the open set U and the curves c_i are carried by τ . The set of geodesic laminations which are carried by τ is closed in the Hausdorff topology and hence ν is carried by τ as well. As a consequence, τ carries every complete geodesic lamination which contains the support of α_+ as a sublamination.

Up to passing to a subsequence, as $k \rightarrow \infty$ the complete geodesic laminations $g^k(\lambda), g^k(\mu)$ converge in the Hausdorff topology to complete geodesic laminations $\tilde{\lambda}, \tilde{\mu}$. Moreover, for any projective transverse measure β, ξ supported in λ, μ the projective measured geodesic laminations $g^k\beta, g^k\xi$ converge as $k \rightarrow \infty$ to one of the two fixed points α_+, α_- for the action of g on \mathcal{PML} which moreover is supported in $\tilde{\lambda}, \tilde{\mu}$. By (2) above, the geodesic laminations λ, μ intersect the support of the lamination α_- transversely and therefore $g^k\beta \rightarrow \alpha^+, g^k\xi \rightarrow \alpha^+$ as $k \rightarrow \infty$ [FLP91]. As a consequence, the laminations $\tilde{\lambda}, \tilde{\mu}$ contain the support of α_+ as a sublamination. Thus by the above consideration, $\tilde{\lambda}, \tilde{\mu}$ are carried by τ . By Lemma 2.3, the set of all complete geodesic laminations which are carried by τ is an open

subset of the space \mathcal{CL} of all complete geodesic laminations and hence there is some $k > 0$ such that the laminations $g^k\lambda, g^k\mu$ are both carried by τ as well. Then λ, μ are carried by the complete train track $g^{-k}(\tau)$. This completes the proof of the lemma. \square

Write $\sigma \prec \tau$ if the train track σ is carried by the train track τ . We have.

Corollary 3.5. *For complete train tracks τ, σ on S there are complete train tracks τ_0, σ_0, ζ such that $\tau_0 \prec \tau, \tau_0 \prec \zeta$ and $\sigma_0 \prec \sigma, \sigma_0 \prec \zeta$.*

Proof. For complete train tracks τ, σ on S choose geodesic laminations λ, μ such that λ is carried by τ and μ is carried by σ . By Lemma 3.4 there is a complete train track ζ which carries both λ and μ .

By Lemma 3.2, for every $\epsilon > 0$ there is a complete train track $\tau_{\lambda, \epsilon}, \sigma_{\mu, \epsilon}$ which ϵ -follows λ, μ . By Lemma 3.3, for sufficiently small ϵ the train tracks $\tau_{\lambda, \epsilon}, \sigma_{\mu, \epsilon}$ are carried by ζ and moreover $\tau_{\lambda, \epsilon}$ is carried by $\tau, \sigma_{\lambda, \epsilon}$ is carried by σ . \square

A map ρ which assigns to a positive integer k contained in an interval $[m, n] \subset \mathbb{R}$ a complete train track $\rho(k)$ such that $\rho(k+1)$ is obtained from $\rho(k)$ by a single split will be called a *splitting sequence*. Thus splitting sequences correspond precisely to directed edge-paths in the train track complex \mathcal{TT} . If there is a splitting sequence connecting a train track τ to a train track σ then we say that τ is *splittable* to σ . A map ρ which assigns to each integer k from an interval $[m, n] \subset \mathbb{R}$ a complete train track $\rho(k)$ such that either $\rho(k+1)$ is obtained from $\rho(k)$ by a single split or a single collapse is called a *splitting and collapsing sequence*. The following corollary is a consequence of Corollary 3.5 and Proposition 3.1.

Corollary 3.6. *Any two complete train tracks on S can be connected by a splitting and collapsing sequence.*

Proof. Let τ be a complete train track which is carried by a complete train track σ . Since τ is birecurrent and maximal it carries a geodesic lamination which is both minimal and maximal (see the discussion in the proof of Lemma 2.3). Thus by Proposition 3.1, there is a complete train track ρ which can be obtained from both τ and σ by a splitting sequence.

Let τ, σ be any complete train tracks on S , and let τ_0, σ_0, ζ be as in Corollary 3.5. Then the train tracks τ, τ_0 are both splittable to the same complete train track τ_1 , and the train tracks σ, σ_0 are both splittable to the same complete train track σ_1 . In particular, the train tracks τ_1, σ_1 are carried by ζ . Using Proposition 3.1 once more we deduce that the train tracks τ_1, ζ are splittable to the same train track τ_2 , and the train tracks σ_1, ζ are splittable to the same train track σ_2 . But this just means that there is a splitting and collapsing sequence from τ to σ which passes through $\tau_1, \tau_2, \zeta, \sigma_2, \sigma_1$. \square

The following corollary is now immediate from Corollary 3.6.

Corollary 3.7. *The train track complex \mathcal{TT} is connected.*

As a consequence, if we identify each edge in \mathcal{TT} with the unit interval $[0, 1]$ then this provides \mathcal{TT} with the structure of a connected locally finite metric graph. Thus \mathcal{TT} is a locally compact complete geodesic metric space. In the sequel we always assume that \mathcal{TT} is equipped with this metric without further comment.

4. TRAIN TRACKS AND THE MAPPING CLASS GROUP

The purpose of this section is to show that the connected graph \mathcal{TT} which we defined in Section 2 is quasi-isometric to the mapping class group of the surface S of genus g with m punctures where $3g - 3 + m \geq 2$. A quasi-isometry between two metric spaces was defined in the introduction.

Consider first the *Teichmüller space* $\mathcal{T}_{g,k}$ of complete marked hyperbolic metrics on S of finite volume. Let $\epsilon > 0$ be smaller than half of a *Margulis constant* for hyperbolic surfaces (see [B92]). Call a surface $S_0 \in \mathcal{T}_{g,k}$ *thick* if the systole of S_0 , i.e. the length of the shortest closed geodesic, is at least ϵ . For sufficiently small ϵ the set $\mathcal{T}_{\text{thick}}$ of thick surfaces is a path connected closed subset of $\mathcal{T}_{g,k}$ with dense interior which is invariant under the natural action of the mapping class group \mathcal{MCG} . The quotient of $\mathcal{T}_{\text{thick}}$ under the action of \mathcal{MCG} is compact.

The *Teichmüller metric* on $\mathcal{T}_{g,k}$ is a complete \mathcal{MCG} -invariant Finsler metric $\|\cdot\|$. We equip $\mathcal{T}_{\text{thick}}$ with the path metric defined by the restriction of this Finsler metric. In other words, the distance between two points in $\mathcal{T}_{\text{thick}}$ is defined to be the infimum of the $\|\cdot\|$ -lengths of paths in $\mathcal{T}_{\text{thick}}$ connecting these two points. With respect to this distance function, the mapping class group \mathcal{MCG} acts properly discontinuously, isometrically and cocompactly on $\mathcal{T}_{\text{thick}}$.

Recall from the introduction that a map $F : (X, d) \rightarrow (Y, d)$ between two metric spaces (X, d) and (Y, d) is an *L-quasi-isometric embedding* if

$$d(x, y)/L - L \leq d(Fx, Fy) \leq Ld(x, y) + L\forall x, y \in X.$$

If moreover $F(X)$ is *L-dense* in Y , i.e. if for every $y \in Y$ there is some $x \in X$ with $d(Fx, y) \leq L$, then F is called an *L-quasi-isometry*.

The mapping class group \mathcal{MCG} is finitely generated (see [I02]). A finite symmetric set \mathcal{G} of generators defines a word norm and hence an \mathcal{MCG} -invariant distance on \mathcal{MCG} . We always assume that \mathcal{MCG} is equipped with such a fixed distance. The next easy lemma is included here as an illustration of the various ways to understand the geometry of \mathcal{MCG} .

Lemma 4.1. *\mathcal{MCG} and $\mathcal{T}_{\text{thick}}$ are equivariantly quasi-isometric.*

Proof. The mapping class group \mathcal{MCG} acts on the length space $\mathcal{T}_{\text{thick}}$ isometrically, properly and cocompactly. By the well known lemma of Švarc-Milnor (Proposition I.8.19 in [BH99]) this implies that $\mathcal{T}_{\text{thick}}$ is equivariantly quasi-isometric to \mathcal{MCG} . \square

The mapping class group also acts naturally as a group of simplicial isometries on the train track complex \mathcal{TT} . Note that by definition, the train track complex is a locally finite directed metric graph and hence a complete locally compact geodesic metric space. We want to show that \mathcal{TT} is equivariantly quasi-isometric to \mathcal{MCG} . By the Švarc-Milnor lemma, for this it is enough to show that the action of \mathcal{MCG} on \mathcal{TT} is proper and cocompact. We show first that this action is cocompact. Recall that the vertices of \mathcal{TT} consist of *generic* train tracks.

Lemma 4.2. *\mathcal{MCG} acts cocompactly on \mathcal{TT} .*

Proof. The number k of switches of a complete train track τ on S is just the number of cusps of the complementary components of τ , and this number only depends on the topological type of S . There are only finitely many abstract trivalent graphs with k vertices. We assign to each half-edge of such a graph one of the three colors red, yellow, green in such a way that every vertex is incident on a half-edge of each color. We do not require that the two half-edges of a single edge have the same color. We call a trivalent graph with such a coloring a *colored graph*. Clearly there are only finitely many colored graphs with k vertices up to isomorphism preserving the coloring.

To every train track $\tau \in \mathcal{V}(\mathcal{TT})$ we associate a colored trivalent graph G with k vertices as follows. The underlying topological graph of G is just the abstract graph defined by τ . Thus the half-branches of τ are in one-one correspondence to the half-edges of G . We color a half-edge of G red if and only if the corresponding half-branch in τ is large. We orient such a large half-branch b in τ in such a way that it ends at the switch v of τ on which it is incident. A neighborhood of b in S is divided by b into two components. The orientation of b together with the orientation of S determine the component to the right and to the left of b . One of the two small half-branches incident on v lies to the left of the other, i.e. the union of this half-branch with the half-branch b is contained in the boundary of the component of $S - \tau$ to the left of b . This half-branch will be colored yellow, and the second small half-branch incident on v (which lies to the right of the yellow half-branch) will be colored green. In this way we obtain a map Ψ from $\mathcal{V}(\mathcal{TT})$ to a finite set of colored trivalent graphs with k vertices.

Let $\tau, \sigma \in \mathcal{V}(\mathcal{TT})$ be such that $\Psi(\tau) = \Psi(\sigma)$. Then the map Ψ determines a homeomorphism of τ onto σ which preserves the coloring of the half-branches. We claim that this homeomorphism maps the sides of the complementary components of τ to the sides of the complementary components of σ .

For this let b be any branch of τ and let v be a switch of τ on which b is incident. The branch b is contained in the boundary of two (not necessarily distinct) complementary components of τ . We orient b in such a way that it ends at v ; then we can distinguish the complementary component T of τ which is to the left of b . We have to show that the finite trainpath ρ on τ which defines the side of T containing b is uniquely determined by the colored graph $\Psi(\tau) = G$.

Assume first that the half-branch of b incident on v is colored red. By the definition of our coloring, the branch b is large at v and ρ necessarily contains the yellow half-branch incident on v . Similarly, if the half-branch of b incident on v is

green (and hence small) then ρ contains the red (large) half-branch incident on v . If the half-branch of b incident on v is yellow then T has a cusp at v and the trainpath ρ ends at v . But this just means that we can successively construct the trainpath ρ from the coloring. In other words, the sides of the complementary components of τ are uniquely determined by the colored graph G . As a consequence, whenever $\Psi(\tau) = \Psi(\sigma)$ then the homeomorphism of τ onto σ induced by Ψ maps the boundary of each complementary trigon of τ to the boundary of a complementary trigon of σ , and it maps the boundary of a complementary once punctured monogon of τ to the boundary of a complementary once punctured monogon of σ . Thus this homeomorphism of τ onto σ can be extended to the complementary components of τ , and this extension is an orientation preserving homeomorphism of S which maps τ to σ . In other words, σ and τ are contained in the same orbit of the action of \mathcal{MCG} on \mathcal{TT} .

As a consequence, there is a *finite* subset A of $\mathcal{V}(\mathcal{TT})$ with the property that the translates of A under \mathcal{MCG} cover all of $\mathcal{V}(\mathcal{TT})$. But this is equivalent to saying that the action of \mathcal{MCG} on \mathcal{TT} is cocompact. \square

Lemma 4.3. *The action of \mathcal{MCG} on \mathcal{TT} is proper.*

Proof. Since \mathcal{TT} is a locally finite graph, every compact subset of \mathcal{TT} contains only finitely many edges. The action of \mathcal{MCG} on \mathcal{TT} is simplicial and isometric and therefore this action is proper if and only if the stabilizer in \mathcal{MCG} of every vertex $\tau \in \mathcal{V}(\mathcal{TT})$ is finite (compare [BH99]).

To show that this is the case, let $\tau \in \mathcal{V}(\mathcal{TT})$ and let $G < \mathcal{MCG}$ be the stabilizer of τ . Since the number of branches and switches of τ only depends on the topological type of S , the subgroup G_0 of G which fixes every branch and every switch of τ is of finite index in G . But a homeomorphism of S which preserves each branch and switch of τ is isotopic to a map which preserves the branches pointwise. Such a map then fixes pointwise the boundaries of the complementary regions of τ . However, every complementary region of τ is a topological disc or a once punctured topological disc, and every homeomorphism of such a disc which fixes the boundary pointwise is isotopic to the identity (see [I02]). This shows that G_0 is trivial and therefore the stabilizer of $\tau \in \mathcal{V}(\mathcal{TT})$ under the action of \mathcal{MCG} is finite. \square

Corollary 4.4. *\mathcal{MCG} and \mathcal{TT} are quasi-isometric.*

Proof. By Lemma 4.2 and Lemma 4.3, the mapping class group \mathcal{MCG} acts properly and cocompactly on the length space \mathcal{TT} . Thus the corollary follows from the lemma of Švarc-Milnor (Proposition I.8.19 in [BH99]). \square

5. CUBICAL EUCLIDEAN CONES

In this section we define a family of connected subgraphs of \mathcal{TT} , one for every complete train track τ and every complete geodesic lamination λ carried by τ . We show that these subspaces equipped with their *intrinsic* path-metric are isometric to cubical graphs contained in an euclidean space of fixed dimension. Thus these subgraphs can be viewed as “cubical Euclidean cones”. In Section 6 we show

Theorem 2 from the introduction which states that these cubical Euclidean cones are quasi-isometrically embedded in \mathcal{TT} .

For a complete train track $\tau \in \mathcal{V}(\mathcal{TT})$ and a complete geodesic lamination $\lambda \in \mathcal{CL}$ which is carried by τ , define the *cubical Euclidean cone* $E(\tau, \lambda) \subset \mathcal{TT}$ to be the full subgraph of \mathcal{TT} whose vertices consist of all complete train tracks which can be obtained from τ by any λ -splitting sequence, i.e. by a sequence of λ -splits. By construction, $E(\tau, \lambda)$ is a connected subgraph of \mathcal{TT} and hence it can be equipped with an intrinsic path metric d_E .

Let $m > 0$ be the number of branches of a complete train track on S . Let e_1, \dots, e_m be the standard basis of \mathbb{R}^m . Define a *cubical graph* in \mathbb{R}^m to be an embedded graph whose vertices are points with integer coordinates (i.e. points contained in \mathbb{Z}^m) and whose edges are line segments of length one connecting two of these vertices v_1, v_2 with $v_1 - v_2 = \pm e_i$ for some $i \leq m$. Note that each such edge has a natural direction.

As in the introduction, define a *splitting arc* in \mathcal{TT} to be a map $\gamma : [0, n] \rightarrow \mathcal{TT}$ with the property that for any integer $i \in [1, n]$, the arc $\gamma[i-1, i]$ is a directed edge in \mathcal{TT} . In other words, $\{\gamma(i)\}$ is a splitting sequence. We have.

Lemma 5.1. *For every cubical Euclidean cone $E(\tau, \lambda) \subset \mathcal{TT}$ 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 $\Phi(E(\tau, \lambda))$.*

Proof. Let $\lambda \in \mathcal{CL}$, let $\tau \in \mathcal{V}(\mathcal{TT})$ be a complete train track which carries λ , let $\sigma \in E(\tau, \lambda)$ and let e be a large branch in τ . We claim that whether or not a splitting sequence connecting τ to σ contains a split at the large branch e only depends on σ but not on the choice of the splitting sequence.

For this note first that σ is carried by τ . For a large branch e of τ define $\nu(e, \sigma)$ to be the minimal cardinality of the preimage of a point x in the interior of e under any carrying map $F : \sigma \rightarrow \tau$. Let τ' be obtained from τ by a λ -split at a large branch $e' \neq e$ and assume that τ' is splittable to σ . Then τ' carries σ , and the branch in τ' corresponding to e under the natural identification of the branches of τ with the branches of τ' is large. We denote it again by e . Let $\nu'(e, \sigma)$ be the minimal cardinality of the preimage of a point x in the interior of e under any carrying map $\sigma \rightarrow \tau'$; we claim that $\nu'(e, \sigma) \geq \nu(e, \sigma)$.

To see this, simply observe that there are disjoint neighborhoods U' of e' and U of e in S and there is a carrying map $G : \tau' \rightarrow \tau$ which equals the identity outside U' . Every carrying map $F' : \sigma \rightarrow \tau'$ can be composed with G to a carrying map $G \circ F' : \sigma \rightarrow \tau$. For a point x in the interior of e the cardinality of the preimage of x under the carrying map $G \circ F' : \sigma \rightarrow \tau$ coincides with the cardinality of the preimage of $x = G^{-1}(x)$ under the carrying map $F' : \sigma \rightarrow \tau'$. Thus we have $\nu(e, \sigma) \leq \nu'(e, \sigma)$.

As a consequence, if $e \subset \tau$ is a large branch with $\nu(e, \sigma) \geq 2$, then every splitting sequence connecting τ to σ has to contain a split at e , and the choice of a right or left split is determined by the requirement that the split track carries λ . On the

other hand, if $\nu(e, \sigma) = 1$ then by Lemma A.4 from the appendix, σ is not carried by a split at e . Thus whether or not a splitting sequence connecting τ to σ contains a split at e is independent of the splitting sequence.

As above, let m be the number of branches of a complete train track on S . Denote by e_1, \dots, e_m the standard basis of \mathbb{R}^m and choose any point $q \in \mathbb{Z}^m$. Number the branches of τ in an arbitrary way. Note that this numbering induces a natural numbering of the branches on any train track which can be obtained from τ by a single split. Let $\alpha : [0, \infty) \rightarrow E(\tau, \lambda)$ be any splitting arc with $\alpha(0) = \tau$. Define a map $\Phi_\alpha : \alpha[0, \infty) \rightarrow \mathbb{R}^m$ inductively as follows. Let $\Phi_\alpha(\tau) = q$ and assume by induction that Φ_α has been defined on $\alpha[0, \ell - 1]$ for some $\ell \geq 1$. Let $p \leq m$ be the number of the large branch e of $\alpha(\ell - 1)$ induced from the numbering of the branches of τ via α so that $\alpha(\ell)$ is obtained from $\alpha(\ell - 1)$ by a single λ -split at e . Define $\Phi_\alpha(\alpha[\ell - 1, \ell])$ to be the line segment in \mathbb{R}^m connecting $\Phi_\alpha(\alpha(\ell - 1))$ to $\Phi_\alpha(\alpha(\ell - 1)) + e_p$. In this way we obtain for every splitting arc $\alpha : [0, \infty) \rightarrow E(\tau, \lambda)$ a map $\Phi_\alpha : \alpha[0, \infty) \rightarrow \mathbb{R}^m$.

We claim that for every train track $\tau \in \mathcal{V}(\mathcal{TT})$, for every complete geodesic lamination $\lambda \in \mathcal{CL}$ carried by τ , for every $\sigma \in E(\tau, \lambda)$ and for every splitting arc $\alpha : [0, \infty) \rightarrow E(\tau, \lambda)$ issuing from $\alpha(0) = \tau$ and passing through $\alpha(\ell) = \sigma$, the image of σ under the map Φ_α which is determined as above by a numbering of the branches of τ , a point $q \in \mathbb{Z}^m$ and by α is in fact independent of the splitting arc α connecting τ to σ .

For this we proceed by induction on the length of the *shortest* splitting arc connecting τ to σ . The case that this length vanishes is trivial, so assume that for some $k \geq 1$ the above claim holds for all τ, λ and all $\sigma \in E(\tau, \lambda)$ which can be obtained from τ by a splitting arc of length at most $k - 1$. Let $\sigma \in E(\tau, \lambda)$ be such that there is a splitting arc $\alpha : [0, k] \rightarrow E(\tau, \lambda)$ of length k connecting $\alpha(0) = \tau$ to σ and let $\beta : [0, p] \rightarrow E(\tau, \lambda)$ be a splitting arc connecting $\beta(0) = \tau$ to σ of length $p \geq k$. For a fixed choice of a numbering of the branches of τ and a fixed point $q \in \mathbb{Z}^m$, these splitting arcs determine maps $\Phi_\alpha : \alpha[0, k] \rightarrow \mathbb{R}^m$, $\Phi_\beta : \beta[0, p] \rightarrow \mathbb{R}^m$ with $\Phi_\alpha(\tau) = \Phi_\beta(\tau) = q$.

Let $b \subset \tau$ be the large branch with the property that $\alpha(1)$ is obtained from τ by a split at b . Let $s \leq n$ be the number of b with respect to our numbering of the branches of τ ; then $\Phi_\alpha(\alpha(1)) = q + e_s$. By the discussion in the beginning of this proof, the splitting sequence $\{\beta(j)\}_{0 \leq j \leq p}$ also contains a split at the branch b . If this is the j -th split in this splitting sequence, then the first $j - 1$ splits of the sequence commute with the split at b . Let $\beta' : [0, p] \rightarrow E(\tau, \lambda)$ be the splitting arc which we obtain from β by exchanging the orders of the first j splits in such a way that $\beta'(0) = \tau$ and that for every $0 \leq i \leq j - 1$ the train track $\beta'(i + 1)$ is obtained from $\beta(i)$ by a λ -split at b . This splitting arc then determines a map $\Phi_{\beta'} : \beta'[0, p] \rightarrow \mathbb{R}^m$. By construction of the maps $\Phi_\beta, \Phi_{\beta'}$ we have $\Phi_\beta(\beta(j)) = \Phi_{\beta'}(\beta'(j))$ and $\Phi_\beta(\sigma) = \Phi_{\beta'}(\sigma)$, moreover $\alpha(1) = \beta'(1)$ and hence $\Phi_\alpha(\alpha(1)) = \Phi_{\beta'}(\beta'(1))$. Therefore we can apply the induction hypothesis to the splitting arcs $\alpha[1, k]$ and $\beta'[1, p]$ issuing from $\alpha(1)$, the numbering of the branches of $\alpha(1)$ inherited from the numbering of the branches of τ and the point $\Phi_\alpha(\alpha(1)) = q + e_s \in \mathbb{Z}^m$ to conclude that the images of σ under the maps Φ_α and Φ_β coincide.

By induction, this construction defines a path-isometric map Φ of $E(\tau, \lambda)$ into \mathbb{R}^m whose image is a cubical graph in \mathbb{R}^m . The above discussion shows that this map is uniquely determined by the choice of a numbering of the branches of τ and the choice of $\Phi(\tau) \in \mathbb{Z}^m$. This shows the lemma. \square

Since directed edge-paths in a cubical graph in \mathbb{R}^m are geodesics, we obtain as an immediate corollary.

Corollary 5.2. *Splitting arcs are geodesics in $(E(\tau, \lambda), d_E)$.*

In the remainder of this section we describe the intrinsic geometry of the cubical Euclidean cones $E(\tau, \lambda)$ more explicitly.

The *Hausdorff distance* between two subsets A, B of a metric space (X, d) is the infimum of all numbers $r > 0$ such that A is contained in the r -neighborhood of B and B is contained in the r -neighborhood of A . If the diameter of A, B is infinite then the Hausdorff distance between A and B may be infinite.

A connected subspace Y of a geodesic metric space (X, d) is called *strictly convex* if for any two points $y, z \in Y$, every geodesic in (X, d) connecting y to z is entirely contained in Y . The next lemma is a first easy step toward an understanding of the intrinsic geometry of a cubical Euclidean cone. For this note that for every vertex $\sigma \in E(\tau, \lambda)$ the cubical Euclidean cone $E(\sigma, \lambda)$ is a complete subgraph of $E(\tau, \lambda)$.

Lemma 5.3. *For $\sigma \in E(\tau, \lambda)$, the subspace $E(\sigma, \lambda)$ of $(E(\tau, \lambda), d_E)$ is strictly convex, and its Hausdorff distance to $E(\tau, \lambda)$ does not exceed $d_E(\tau, \sigma)$.*

Proof. Since a strictly convex subspace A of a strictly convex subspace B of a geodesic metric space X is strictly convex in X , it suffices to show the following. If $\sigma \in E(\tau, \lambda)$ is obtained from τ by a single split at a large branch e then $E(\sigma, \lambda)$ is a strictly convex subspace of $E(\tau, \lambda)$ whose Hausdorff distance to $E(\tau, \lambda)$ equals one.

For this note that by Lemma 5.1 and its proof, if $\sigma \in E(\tau, \lambda)$ is a train track which can be obtained from τ by a single split at a large branch e and if $\eta \in E(\tau, \lambda)$ is *not* contained in $E(\sigma, \lambda)$, then a splitting sequence connecting τ to η does not contain a split at e . Moreover, σ is splittable to a train track $\eta' \in E(\tau, \lambda)$ which can be obtained from η by a single split at e . In other words, there is a natural retraction $R : E(\tau, \lambda) \rightarrow E(\sigma, \lambda)$ which equals the identity on $E(\sigma, \lambda)$ and maps a train track $\eta \in E(\tau, \lambda) - E(\sigma, \lambda)$ to the train track obtained from η by a λ -split at e . This shows that the Hausdorff distance between $E(\tau, \lambda)$ and $E(\sigma, \lambda)$ does not exceed 1.

Now let ζ_1, ζ_2 be any vertices in $E(\tau, \lambda)$ which are connected by an edge. We may assume that ζ_2 can be obtained from ζ_1 by a single split at a large branch e' . If both ζ_1, ζ_2 are contained in $E(\tau, \lambda) - E(\sigma, \lambda)$ then since λ -splits at distinct large branches commute, the train tracks $R(\zeta_1), R(\zeta_2)$ are connected by an edge in $E(\sigma, \lambda)$. On the other hand, if $\zeta_1 \in E(\tau, \lambda) - E(\sigma, \lambda)$ and $\zeta_2 \in E(\sigma, \lambda)$ then $R(\zeta_1) = R(\zeta_2) = \zeta_2$. As a consequence, the retraction R is distance non-increasing. Moreover, any simplicial path in $E(\tau, \lambda)$ connecting two points $\sigma, \eta \in E(\sigma, \lambda)$ and

which is not entirely contained in $E(\sigma, \lambda)$ passes through an edge which is mapped by R to a single point. This implies strict convexity of $E(\sigma, \lambda) \subset E(\tau, \lambda)$. \square

The next lemma can be used to calculate distances in a cubical Euclidean cone $(E(\tau, \lambda), d_E)$ explicitly.

Lemma 5.4. *Let $\sigma, \eta \in E(\tau, \lambda)$ be any two vertices. Then there are unique train tracks $\Theta_-(\sigma, \eta), \Theta_+(\sigma, \eta) \in E(\tau, \lambda)$ with the following properties.*

- (1) $\sigma, \eta \in E(\Theta_-(\sigma, \eta), \lambda) \subset E(\tau, \lambda)$, and there is a geodesic in $(E(\tau, \lambda), d_E)$ connecting σ to η which passes through $\Theta_-(\sigma, \eta)$.
- (2) $E(\Theta_+(\sigma, \eta), \lambda) = E(\sigma, \lambda) \cap E(\eta, \lambda)$, and there is a geodesic in $(E(\tau, \lambda), d_E)$ connecting σ to η which passes through $\Theta_+(\sigma, \eta)$.
- (3) $d_E(\sigma, \Theta_-(\sigma, \eta)) = d_E(\eta, \Theta_+(\sigma, \eta))$.

Proof. Let σ, η be vertices in $E(\tau, \lambda)$. Then σ, η are complete train tracks, and τ is splittable to both σ and η . Let $A \subset E(\tau, \lambda)$ be the set of all complete train tracks which can be obtained from τ by a splitting sequence and which are splittable to both σ and η . Note that A is a *finite* set of vertices of $E(\tau, \lambda)$. For $\beta, \beta' \in A$ write $\beta < \beta'$ if β is splittable to β' . Then $<$ is a partial order on A .

Let $\Theta_-(\sigma, \eta)$ be a maximal element for this partial order. Then σ, η are both contained in $E(\Theta_-(\sigma, \eta), \lambda)$ and hence by Lemma 5.3, every geodesic in $E(\tau, \lambda)$ connecting σ to η is contained in $E(\Theta_-(\sigma, \eta), \lambda)$. If e is any large branch of $\Theta_-(\sigma, \eta)$ and if the train track obtained from $\Theta_-(\sigma, \eta)$ by a λ -split at e is splittable to σ then by maximality of $\Theta_-(\sigma, \eta)$, it is not splittable to η .

Let $\Phi : E(\Theta_-(\sigma, \eta), \lambda) \rightarrow \mathbb{R}^m$ be an isometry of $E(\Theta_-(\sigma, \eta), \lambda)$ onto a cubical graph in \mathbb{R}^m defined as in the proof of Lemma 5.1 by the choice of the basepoint $\Phi(\Theta_-(\sigma, \eta)) = 0$ and a numbering of the branches of $\Theta_-(\sigma, \eta)$.

We claim that up to a permutation of the standard basis of \mathbb{R}^m , there is a number $\ell \geq 0$ such that for the standard direct orthogonal decomposition $\mathbb{R}^m = \mathbb{R}^\ell \oplus \mathbb{R}^{m-\ell}$ we have $\Phi(\sigma) \in \mathbb{R}^\ell \times \{0\}$ and $\Phi(\eta) \in \{0\} \times \mathbb{R}^{m-\ell}$. Namely, by the choice of the train track $\Theta_-(\sigma, \eta)$ and the fact that σ, η both carry the complete geodesic lamination λ , the set of large branches of $\Theta_-(\sigma, \eta)$ can be partitioned into disjoint subsets $\mathcal{E}^+, \mathcal{E}^-$ such that a splitting sequence connecting $\Theta_-(\sigma, \eta)$ to σ does not contain any split at a large branch $e \in \mathcal{E}^+$ and that a splitting sequence connecting $\Theta_-(\sigma, \eta)$ to η does not contain any split at a large branch $e \in \mathcal{E}^-$.

Following [PH92], we call a trainpath $\rho : [0, p] \rightarrow \Theta_-(\sigma, \eta)$ *one-sided large* if for every $i < p$ the half-branch $\rho[i, i + 1/2]$ is large and if $\rho[p - 1, p]$ is a large branch. A one-sided large trainpath $\rho : [0, p] \rightarrow \Theta_-(\sigma, \eta)$ is embedded [PH92], and for every $i \in \{1, \dots, p - 1\}$ the branch $\rho[i - 1, i] \subset \Theta_-(\sigma, \eta)$ is mixed. For every large half-branch \hat{b} of $\Theta_-(\sigma, \eta)$ there is a unique one-sided large trainpath issuing from \hat{b} . Define $\mathcal{A}_0^+, \mathcal{A}_0^-$ to be the set of all branches of $\Theta_-(\sigma, \eta)$ contained in a one-sided large trainpath ending at a branch in $\mathcal{E}^+, \mathcal{E}^-$. Then the sets $\mathcal{A}_0^+, \mathcal{A}_0^-$ are disjoint, and a branch of $\Theta_-(\sigma, \eta)$ is *not* contained in $\mathcal{A}_0^+ \cup \mathcal{A}_0^-$ if and only if it is small. Each endpoint of a small branch is a starting point of a one-sided large trainpath. Define \mathcal{A}^\pm to be the union of \mathcal{A}_0^\pm with all small branches b of $\Theta_-(\sigma, \eta)$ with the

property that both large half-branches incident on the endpoints of b are contained in \mathcal{A}_0^\pm . If $b \notin \mathcal{A}^+ \cup \mathcal{A}^-$ then b is a small branch incident on two distinct switches, and one of these switches is the starting point of a one-sided large trainpath in \mathcal{A}_0^+ , the other is the starting point of a one-sided large trainpath in \mathcal{A}_0^- .

The map Φ is determined by a numbering of the branches of $\Theta_-(\sigma, \eta)$. We may assume that this numbering is such that for the cardinality ℓ of \mathcal{A}^- , the set \mathcal{A}^- consists of the branches with numbers $1, \dots, \ell$. A splitting sequence connecting $\Theta_-(\sigma, \eta)$ to σ does not contain any split at a large branch $e \in \mathcal{E}^+$ by assumption. Therefore, such a splitting sequence only contains splits at the branches in \mathcal{A}^- . By the choice of our numbering, the image of any such splitting sequence under the map Φ is contained in the linear subspace $\mathbb{R}^\ell \subset \mathbb{R}^m$ spanned by the first ℓ vectors of the standard basis of \mathbb{R}^m . Similarly, the image under Φ of a splitting sequence connecting $\Theta_-(\sigma, \eta)$ to η is contained in the subspace $\mathbb{R}^{m-\ell} \subset \mathbb{R}^m$ spanned by the last $m - \ell$ vectors of the standard basis. This shows our claim.

The image under Φ of any splitting arc in $E(\Theta_-(\sigma, \eta), \lambda)$ is a geodesic edge path in the standard cubical graph $\mathcal{G} \subset \mathbb{R}^m$ with vertex set \mathbb{Z}^m and where two vertices are connected by a straight line segment if their distance in \mathbb{R}^m equals one. If $\gamma_\sigma, \gamma_\eta$ are such geodesic edge-paths connecting $0 = \Phi(\Theta_-(\sigma, \eta))$ to $\Phi(\sigma), \Phi(\eta)$ which are images of splitting arcs then $\gamma_\sigma \subset \mathbb{R}^\ell \times \{0\}, \gamma_\eta \subset \{0\} \times \mathbb{R}^{m-\ell}$ and hence $\gamma_\eta \circ \gamma_\sigma^{-1}$ is a geodesic in \mathcal{G} connecting $\Phi(\sigma)$ to $\Phi(\eta)$. Since Φ is an isometry of $E(\Theta_-(\sigma, \eta), \lambda)$ onto a connected subgraph of \mathcal{G} and $E(\Theta_-(\sigma, \eta), \lambda)$ is a strictly convex subset of $E(\tau, \lambda)$, we conclude that there is a geodesic in $E(\tau, \lambda)$ connecting σ to η which passes through $\Theta_-(\sigma, \eta)$. This shows the first statement of the lemma.

On the other hand, the considerations in the previous three paragraphs of this proof also show that there is a unique vertex $\Theta_+(\sigma, \eta) \in E(\Theta_-(\sigma, \eta), \lambda)$ such that $\Phi(\Theta_+(\sigma, \eta)) = \Phi(\sigma) + \Phi(\eta)$. The train track $\Theta_+(\sigma, \eta)$ satisfies the properties in statement 2) and 3) of the lemma. This completes the proof of the lemma. \square

In the sequel we call τ the *basepoint* of the cubical Euclidean cone $E(\tau, \lambda)$. Another immediate consequence of Lemma 5.1 is the following growth control.

Corollary 5.5. *For every $k > 0$ the number of vertices in $E(\tau, \lambda)$ whose intrinsic distance to the basepoint is at most k is not bigger than $(k + 1)^m$.*

Proof. Let $\Phi : E(\tau, \lambda) \rightarrow \mathbb{R}^m$ be an embedding of $E(\tau, \lambda)$ onto a cubical graph in \mathbb{R}^m as in Lemma 5.1. Assume that Φ maps the basepoint of $E(\tau, \lambda)$ to 0. By construction, if we denote by $||$ the norm on \mathbb{R}^m defined by $|x| = \sum_i |x_i|$ then Φ maps the ball of radius k about 0 in $E(\tau, \lambda)$ into the intersection of the ball of radius k about 0 in $(\mathbb{R}^m, ||)$ with the cone $\{x \mid x_i \geq 0\}$. Thus the image under Φ of the set of vertices in $E(\tau, \lambda)$ whose distance to the basepoint is at most k is not bigger than the cardinality of the set $\{x \in \mathbb{Z}^m \mid x_i \geq 0, |x| \leq k\}$, and this cardinality is not bigger than $(k + 1)^m$. \square

Call a train track $\sigma \in \mathcal{V}(TT)$ a *full split* of a train track $\tau \in \mathcal{V}(TT)$ if σ can be obtained from τ by splitting τ at *each* large branch precisely once. For a complete geodesic lamination $\lambda \in \mathcal{CL}$ which is carried by τ we call a full split σ of τ a *full*

λ -split if λ is carried by σ . Note that a full λ -split of τ is uniquely determined by λ and τ : There are no choices involved. A *full λ -splitting sequence* of length $k \geq 0$ is a sequence $\{\eta(i)\}_{0 \leq i \leq k} \subset \mathcal{V}(\mathcal{TT})$ of train tracks with the property that for every $i < k$ the train track $\eta(i+1)$ is a full λ -split of $\eta(i)$. A full λ -splitting sequence of length $k \geq 0$ issuing from τ is unique. We call its endpoint the *full k -fold λ -split* of τ . The following corollary is immediate from Lemma 5.1.

Corollary 5.6. *Let $\{\tau(i)\}_{0 \leq i \leq p}$ be a splitting sequence of length p . If $\lambda \in \mathcal{CL}$ is carried by $\tau(p)$ then $\tau(p)$ is splittable to the full p -fold λ -split of $\tau(0)$.*

Proof. If $\{\tau(i)\}_{0 \leq i \leq p}$ is any splitting sequence of length p and if $\lambda \in \mathcal{CL}$ is carried by $\tau(p)$ then $\tau(p) \in E(\tau(0), \lambda)$ and the same is true for the full p -fold λ -split σ of $\tau(0)$. Thus the corollary is immediate from Lemma 5.1. \square

6. QUASI-ISOMETRICALLY EMBEDDED SUBSPACES

In Section 5 we defined for a complete train track τ which carries a complete geodesic lamination λ the “cubical Euclidean cone” $E(\tau, \lambda) \subset \mathcal{TT}$, and we investigated its intrinsic path metric d_E . We showed that $(E(\tau, \lambda), d_E)$ is isometric to a cubical graph in \mathbb{R}^m where $m > 0$ is the number of branches of a complete train track on S . By definition, the inclusion $(E(\tau, \lambda), d_E) \rightarrow \mathcal{TT}$ is a one-Lipschitz map.

The goal of this section is to show Theorem 2 from the introduction which relates the intrinsic path metric d_E on $E(\tau, \lambda)$ to the restriction of the metric d on \mathcal{TT} . Since by Corollary 5.2 splitting arcs in $E(\tau, \lambda)$ are geodesics, Corollary 3 from the introduction is an immediate consequence of Theorem 2.

The proof of Theorem 2 is quite involved. The main technical difficulty arises from a lack of uniqueness and direct geometric control of the cubical Euclidean cones. As an example, there are cubical Euclidean cones whose Hausdorff distance is finite and which do not intersect. Moreover, even though these cubical Euclidean cones equipped with their intrinsic path metric are quasi-isometric, they may not be quasi-isometric with a multiplicative scaling constant close to one. A typical example of such a situation is a “standard model” τ in the sense of Penner and Harer [PH92] and a train track which can be obtained from τ by a shift.

We overcome this difficulty by using a second geometric model for the mapping class group. This model is a metric graph \mathcal{ST} whose vertex set is again the set $\mathcal{V}(\mathcal{TT})$ of all isotopy classes of complete train tracks on S and where two such vertices τ, σ are connected by an edge of length one if and only if either they are connected in \mathcal{TT} or if σ can be obtained from τ by a single shift. Since \mathcal{TT} is connected, the same is true for \mathcal{ST} . Then \mathcal{ST} is a locally finite metric graph which admits a properly discontinuous cocompact isometric action of \mathcal{MCG} . Thus we have.

Lemma 6.1. *The natural \mathcal{MCG} -equivariant inclusion $\mathcal{TT} \rightarrow \mathcal{ST}$ is a one-Lipschitz quasi-isometry.*

As a consequence, it suffices to show that the inclusion $E(\tau, \lambda) \rightarrow \mathcal{ST}$ is an L_0 -quasi-isometric embedding for a number $L_0 > 1$ not depending on τ and λ . This in turn is carried out in three steps.

The first step consists in reducing the problem further as follows. For a complete geodesic lamination λ on S let $\mathcal{S}(\lambda)$ be the full subgraph of \mathcal{ST} whose vertex set $\mathcal{V}(\mathcal{S}(\lambda)) \subset \mathcal{V}(\mathcal{TT})$ is the set of all complete train tracks on S which carry λ . Lemma 3.2 and Lemma 3.3 together with Proposition 6.2 below immediately imply that $\mathcal{S}(\lambda)$ is connected. In other words, $\mathcal{S}(\lambda)$ with its intrinsic path metric d_λ is a geodesic metric space. We show that there is a number $q_1 > 1$ not depending on λ such that the inclusion $(\mathcal{S}(\lambda), d_\lambda) \rightarrow \mathcal{ST}$ is a q_1 -quasi-isometry. As a consequence, it suffices to show that the inclusion $(E(\tau, \lambda), d_E) \rightarrow (\mathcal{S}(\lambda), d_\lambda)$ is a uniform quasi-isometric embedding.

Now the image of a cubical Euclidean cone $E(\tau, \lambda)$ under this inclusion is contained in the full subgraph $C(\tau, \lambda)$ of \mathcal{ST} whose vertex set is the set of all complete train tracks on S which are carried by τ and carry λ . In a second step, we show that the inclusion $E(\tau, \lambda) \rightarrow C(\tau, \lambda)$ is a uniform quasi-isometry for the intrinsic path metrics d_E, d_C . The third and final step then consists in showing that the inclusion $(C(\tau, \lambda), d_C) \rightarrow (\mathcal{S}(\lambda), d_\lambda)$ is a uniform quasi-isometric embedding. Together this then yields the proof of Theorem 2.

In the remainder of this section, we exclusively work with the metric graph \mathcal{ST} . Distances between complete train tracks are measured in this graph (and are as before denoted by d). By Lemma 6.1, this is of no significance for our purpose. We emphasize that the space \mathcal{ST} is a purely auxiliary object which was introduced exclusively for the proof of Theorem 2 in view of the following important result of Penner and Harer. Define a *splitting and shifting sequence* to be a sequence $\{\alpha(i)\} \subset \mathcal{V}(\mathcal{TT})$ such that for each i , the train track $\alpha(i+1)$ is obtained from $\alpha(i)$ either by a single split or a single shift. By definition, such a splitting and shifting sequence defines an edge path in \mathcal{ST} . Theorem 2.4.1 of [PH92] relates splitting and shifting of train tracks to carrying.

Proposition 6.2. [PH92] *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.*

We also need the following local version of Proposition 6.2.

Lemma 6.3. *For every $k > 0$ there is a number $p_1(k) > 0$ with the following property. Let $\sigma \in \mathcal{V}(\mathcal{TT})$ be carried by $\tau \in \mathcal{V}(\mathcal{TT})$. If $d(\tau, \sigma) \leq k$ then τ can be connected to σ by a splitting and shifting sequence of length at most $p_1(k)$.*

Proof. Up to the action of the mapping class group, for every $k > 0$ there are only finitely many pairs $\sigma \prec \tau$ of complete train tracks whose distance is at most k . Thus the lemma follows from Proposition 6.2 and invariance under the action of the mapping class group. \square

To compare the intrinsic geometry of the metric $(\mathcal{S}(\lambda), d_\lambda)$ to the geometry of \mathcal{ST} , we use the following uniform local control on the graphs $\mathcal{S}(\lambda)$.

Lemma 6.4. *For every $k > 0$ there is a number $p_2(k) > 0$ with the following property. Let $\tau, \eta \in \mathcal{V}(\mathcal{TT})$ be two complete train tracks of distance at most k which carry a common complete geodesic lamination λ . Then there is a train track σ which carries λ , which is carried by both τ, η and whose distance to τ, η is at most $p_2(k)$.*

Proof. Let α, β be two complete train tracks on S and let $\mathcal{CL}(\alpha, \beta)$ be the set of all complete geodesic laminations which are carried by both α and β . By Lemma 2.3, $\mathcal{CL}(\alpha, \beta)$ is an open and closed subset of \mathcal{CL} . By Lemma 3.2 and Lemma 3.3, for every $\nu \in \mathcal{CL}(\alpha, \beta)$ there is a complete train track η which carries ν and is carried by both α, β .

By Lemma 2.3, for every complete train track η , the set $\mathcal{CL}(\eta)$ of all complete geodesic laminations which are carried by η is an open subset of \mathcal{CL} . By compactness of $\mathcal{CL}(\alpha, \beta)$, there are *finitely* many train tracks η_1, \dots, η_k carried by both α and β and such that $\mathcal{CL}(\alpha, \beta) = \cup_{i=1}^k \mathcal{CL}(\eta_i)$. In other words, there is a number $\ell(\alpha, \beta) > 0$ with the following property. For every $\nu \in \mathcal{CL}(\alpha, \beta)$ there is a complete train track η on S which carries ν , which is carried by both α and β and whose distance to α, β is at most $\ell(\alpha, \beta)$.

For $k > 0$, up to the action of the mapping class group there are only finitely many pairs (α, β) of complete train tracks whose distance in \mathcal{ST} is at most k and which carry a common complete geodesic lamination. By invariance under the action of the mapping class group, this implies that there is a number $p_2(k) > 0$ such that the following holds true. For every complete geodesic lamination λ and for every pair $\alpha, \beta \in \mathcal{V}(\mathcal{S}(\lambda))$ with $d(\alpha, \beta) \leq k$ there is a train track $\eta \in \mathcal{V}(\mathcal{S}(\lambda))$ with $\max\{d(\eta, \alpha), d(\eta, \beta)\} \leq p_2(k)$ and which is carried by both α, β . This shows the lemma. \square

Now we are ready to complete the first step in the proof of Theorem 2.

Proposition 6.5. *There is a number $q_1 > 0$ such that for every $\lambda \in \mathcal{CL}$ the inclusion $(\mathcal{S}(\lambda), d_\lambda) \rightarrow \mathcal{ST}$ is a q_1 -quasi-isometry.*

Proof. By Lemma 2.3, for every $\tau \in \mathcal{V}(\mathcal{TT})$ the set of all complete geodesic laminations which are carried by τ is an open subset of \mathcal{CL} . By Lemma 3.2, every complete geodesic lamination is carried by some $\tau \in \mathcal{V}(\mathcal{TT})$. Since \mathcal{CL} is a compact space, there is a finite set $\mathcal{E} \subset \mathcal{V}(\mathcal{TT})$ so that every complete geodesic lamination is carried by a train track $\tau \in \mathcal{E}$. On the other hand, the mapping class group acts cocompactly on \mathcal{ST} , and it acts as a group of homeomorphisms on \mathcal{CL} . Thus by equivariance under the action of the mapping class group, there is a number $D > 0$ and for every $\sigma \in \mathcal{ST}$ and every $\lambda \in \mathcal{CL}$ there is some $\tau \in \mathcal{V}(\mathcal{S}(\lambda))$ with $d(\tau, \sigma) \leq D$. Since $\mathcal{S}(\lambda)$ is a complete subgraph of \mathcal{ST} this means that the inclusion $\mathcal{S}(\lambda) \rightarrow \mathcal{ST}$ is a 1-Lipschitz map with D -dense image.

Let again d_λ be the intrinsic distance on $\mathcal{S}(\lambda)$. We have to show that there is a number $L > 1$ not depending on λ such that $d_\lambda(\tau, \sigma) \leq Ld(\tau, \sigma)$ for all $\tau, \sigma \in \mathcal{V}(\mathcal{S}(\lambda))$. For this let τ, σ be any two vertices of $\mathcal{S}(\lambda)$. Let $\gamma : [0, m] \rightarrow \mathcal{ST}$ be a simplicial geodesic connecting $\tau = \gamma(0)$ to $\sigma = \gamma(m)$ (i.e. γ maps integer points in

\mathbb{R} to vertices of \mathcal{ST}). By the above consideration, for every $i \leq m$ there is a train track $\zeta(i) \in \mathcal{V}(\mathcal{S}(\lambda))$ with $d(\zeta(i), \gamma(i)) \leq D$ and where $\zeta(0) = \tau, \zeta(m) = \sigma$. Then the distance in \mathcal{ST} between $\zeta(i)$ and $\zeta(i+1)$ is at most $2D+1$.

By Lemma 6.4 and Lemma 6.3, there is a constant $\kappa = p_1(p_2(2D+1)) > 0$ only depending on D and for every $i \leq m$ there is a train track $\beta(i) \in \mathcal{V}(\mathcal{S}(\lambda))$ which carries λ , which is carried by both $\zeta(i)$ and $\zeta(i+1)$ and which can be obtained from both $\zeta(i)$ and $\zeta(i+1)$ by a splitting and shifting sequence of length at most κ . As a consequence, we have $d_\lambda(\zeta(i-1), \zeta(i)) \leq 2\kappa$ for all $i \leq m$. But this just means that the distance $d_\lambda(\tau, \sigma)$ in $\mathcal{S}(\lambda)$ between τ and σ is not bigger than $2\kappa d(\tau, \sigma)$. In other words, the inclusion $(\mathcal{S}(\lambda), d_\lambda) \rightarrow \mathcal{ST}$ is a 2κ -quasi-isometry for the constant $2\kappa > 1$ not depending on λ . This shows the proposition. \square

The second step in the proof of Theorem 2 is technically the most involved part of the paper. As before, consider the full subgraph $C(\tau, \lambda) \subset \mathcal{S}(\lambda)$ of $\mathcal{S}(\lambda)$ whose vertex set is the set of all complete train tracks η which are carried by τ and which carry λ . By Proposition 6.2, $C(\tau, \lambda)$ is connected and hence its intrinsic path metric d_C is defined. Moreover, $C(\tau, \lambda)$ contains the cubical Euclidean cone $E(\tau, \lambda)$ as a connected subgraph. Using Proposition A.6 from the appendix, we show that the natural inclusion $(E(\tau, \lambda), d_E) \rightarrow (C(\tau, \lambda), d_C)$ is a uniform quasi-isometry.

Lemma 6.6. *There is a number $q_2 > 0$ such that for every complete geodesic lamination $\lambda \in \mathcal{CL}$ and every complete train track τ which carries λ the inclusion $(E(\tau, \lambda), d_E) \rightarrow (C(\tau, \lambda), d_C)$ is a q_2 -quasi-isometry.*

Proof. We begin with showing that there is a constant $k_1 > 0$ such that the subgraph $E(\tau, \lambda)$ is k_1 -dense in $C(\tau, \lambda)$ with respect to the intrinsic metric d_C on $C(\tau, \lambda)$.

For this let $\chi > 0$ be as in Proposition A.6. Then by the definition of the graph $\mathcal{S}(\lambda)$, the distance in $\mathcal{S}(\lambda)$ between a vertex η of the graph $C(\tau, \lambda)$ and its subgraph $E(\tau, \lambda)$ is at most χ . Thus by Lemma 6.4 and Lemma 6.3, there is a complete train track $\sigma \in E(\tau, \lambda)$ and there is a complete train track $\zeta \in C(\tau, \lambda)$ which is carried by both σ and η and such that moreover σ, η can be connected to ζ by a splitting and shifting sequence of uniformly bounded length. As a consequence, the distance in $C(\tau, \lambda)$ between η and σ is uniformly bounded. This shows that there is a constant $k_1 > 0$ such that $E(\tau, \lambda)$ is k_1 -dense in $(C(\tau, \lambda), d_C)$.

Since the inclusion $(E(\tau, \lambda), d_E) \rightarrow (C(\tau, \lambda), d_C)$ is clearly one-Lipschitz, for the proof of the lemma we are left with showing the existence of a universal constant $k_2 > 0$ such that $d_E(\xi, \eta) \leq k_2 d_C(\xi, \eta)$ for all vertices $\xi, \eta \in E(\tau, \lambda)$.

Now both $E(\tau, \lambda)$ and $C(\tau, \lambda)$ are geodesic metric spaces and $E(\tau, \lambda)$ is k_1 -dense in $C(\tau, \lambda)$. Therefore it is enough to show the existence of a constant $k_3 > 0$ with the following property. If $\xi, \eta \in E(\tau, \lambda)$ are any two vertices with $d_C(\xi, \eta) \leq 3k_1$ then $d_E(\xi, \eta) \leq k_3$.

Namely, assume that this property holds true. Let $\xi, \eta \in E(\tau, \lambda)$ be arbitrary vertices with $d_C(\xi, \eta) = d > 0$. Let $\gamma : [0, d] \rightarrow C(\tau, \lambda)$ be a simplicial geodesic connecting $\gamma(0) = \xi$ to $\gamma(d) = \eta$. Since $E(\tau, \lambda)$ is k_1 -dense in $C(\tau, \lambda)$ we can

replace γ by a simplicial path $\tilde{\gamma} : [0, \tilde{d}] \rightarrow C(\tau, \lambda)$ of length $\tilde{d} \leq 3d$ with the same endpoints and the additional property that $\tilde{\gamma}(3\ell k_1) \in E(\tau, \lambda)$ for all integers $\ell \leq d/k_1$. The arc $\tilde{\gamma}[0, 3k_1]$ is obtained by concatenation of $\gamma[0, 2k_1]$ with an arc of length at most k_1 which connects $\gamma(2k_1)$ with a point in $E(\tau, \lambda)$. Inductively, for each $\ell \leq d/k_1$ the arc $\tilde{\gamma}[3(\ell-1)k_1, 3\ell k_1]$ is up to parametrization a concatenation of a segment of length at most k_1 connecting a point in $E(\tau, \lambda)$ to $\gamma((\ell-1)k_1)$, the arc $\gamma[(\ell-1)k_1, \ell k_1]$ and an arc of length at most k_1 connecting $\gamma(\ell k_1)$ to a point in $E(\tau, \lambda)$. Replace each of the arcs $\tilde{\gamma}[3(\ell-1)k_1, 3\ell k_1]$ of length at most $3k_1$ with endpoints in $E(\tau, \lambda)$ by an arc of length at most k_3 which is contained in $E(\tau, \lambda)$. The resulting path is contained in $E(\tau, \lambda)$, it connects ξ to η and its length does not exceed $k_3 d_C(\xi, \eta)/k_1$.

To show the existence of a constant $k_3 > 0$ with the above properties, let $\xi, \eta \in E(\tau, \lambda)$ be vertices such that $d_C(\xi, \eta) \leq 3k_1$ and let $\zeta \in C(\tau, \lambda)$ be a complete train track which carries λ and which can be obtained from both ξ, η by a splitting and shifting sequence whose length is bounded from above by a universal constant $p > 0$. Such a complete train track exists by Lemma 6.4 and Lemma 6.3 and the fact that the distance in \mathcal{ST} between ξ, η does not exceed $3k_1$.

Since ζ is carried by τ , there is a train track $\beta \in E(\tau, \lambda)$ which carries ζ and such that no split of β carries ζ . It follows from Lemma A.3 in the appendix that β is unique. Now $\xi \in E(\tau, \lambda)$ carries ζ and therefore by Lemma 5.1, ξ is splittable to β . The same argument also shows that η is splittable to β . To summarize, both ξ and η are splittable to β , β carries ζ and both ξ and η can be connected to ζ by a splitting and shifting sequence of length at most p .

There are only finitely many orbits for the action of the mapping class group on $\mathcal{V}(\mathcal{TT})$. For a given $\alpha \in \mathcal{V}(\mathcal{TT})$, there are only finitely many complete train tracks $\rho \in \mathcal{V}(\mathcal{TT})$ which can be obtained from α by a splitting and shifting sequence of length at most p . Thus up to the action of the mapping class group, there are only finitely many triples of train tracks (α, ρ, ζ) such that ζ can be obtained from α by a splitting and shifting sequence of length at most p and that α is splittable to ρ , ρ carries ζ and not split of ρ carries ζ (note as before that ρ is uniquely determined by these requirements).

As a consequence, using the above notations, the train tracks $\xi, \eta \in E(\tau, \lambda)$ can be connected to $\beta \in E(\tau, \lambda)$ by a splitting sequence of uniformly bounded length, and ξ can be connected to η by a path in $E(\tau, \lambda)$ of uniformly bounded length which is the concatenation of a splitting sequence connecting ξ to β and a collapsing sequence connecting β to η . This completes the proof of the lemma. \square

Lemma 6.4, Lemma 6.6 and Lemma 5.3 are used to show the following.

Lemma 6.7. *For every $k > 0$ there is a number $p_3(k) > 0$ with the following property. Let τ, σ be complete train tracks which carry a common complete geodesic lamination λ and such that $d(\tau, \sigma) \leq k$. Then for every vertex $\tau' \in C(\tau, \lambda)$ there is a complete train track $\sigma' \in C(\tau', \lambda) \cap C(\sigma, \lambda)$ with $d(\tau', \sigma') \leq p_3(k)$.*

Proof. Let $k > 0$, let $\lambda \in \mathcal{CL}$ and let $\tau, \sigma \in \mathcal{V}(\mathcal{TT})$ be complete train tracks which carry λ and such that $d(\tau, \sigma) \leq k$. By Lemma 6.4 and Lemma 6.3 there is a complete train track $\xi \in C(\tau, \lambda) \cap C(\sigma, \lambda)$ which can be obtained from τ by a splitting and shifting sequence of length at most $p_1(p_2(k))$. In particular, the distance in \mathcal{ST} between τ and ξ is not bigger than $p_1(p_2(k))$. Since $C(\xi, \lambda) \subset C(\sigma, \lambda)$ it is therefore enough to show the statement of the lemma for τ and the vertex ξ of $C(\tau, \lambda)$. In other words, it suffices to show the lemma under the additional assumption that the complete train track σ is a vertex of $C(\tau, \lambda)$.

If $\eta, \eta' \in \mathcal{V}(\mathcal{TT})$ are shift equivalent train tracks which carry λ then we have $C(\eta, \lambda) = C(\eta', \lambda)$. This implies that it is enough to show the lemma in the particular case that σ is a vertex of $E(\tau, \lambda) \subset C(\tau, \lambda)$ which can be obtained from τ by a single split.

Now by Proposition A.6 and Lemma 6.3, the Hausdorff distance in \mathcal{TT} between $C(\tau, \lambda)$ and $E(\tau, \lambda)$ is at most $p_1(\chi)$ where $\chi > 0$ is as in Proposition A.6, and similarly for $C(\sigma, \lambda)$ and $E(\sigma, \lambda)$. Moreover, by Lemma 5.3, the Hausdorff distance between $E(\tau, \lambda)$ and $E(\sigma, \lambda)$ equals 1. Together we conclude that the Hausdorff distance between $C(\tau, \lambda)$ and $C(\sigma, \lambda) \subset C(\tau, \lambda)$ does not exceed $2p_1(\chi) + 1$.

As a consequence, for $\tau' \in C(\tau, \lambda)$ there is some vertex $\sigma' \in C(\sigma, \lambda)$ with $d(\tau', \rho) \leq 2p_1(\chi) + 1$. Since by assumption we have $C(\sigma, \lambda) \subset C(\tau, \lambda)$, the train track σ' satisfies the requirement in the lemma. \square

The following lemma is the main remaining step for the proof of Theorem 2. For its formulation, denote as before by d_E the intrinsic path metric on a cubical Euclidean cone $E(\tau, \lambda)$. Recall moreover the definition of the metric graph $\mathcal{S}(\lambda)$ for a complete geodesic lamination λ .

Lemma 6.8. *There is a number $p_4 > 0$ with the following property. Let λ be a complete geodesic lamination, let τ be a complete train track which carries λ and let $\sigma, \eta \in E(\tau, \lambda)$. If $\gamma : [0, m] \rightarrow \mathcal{S}(\lambda)$ is any simplicial path connecting $\gamma(0) = \sigma$ to $\gamma(m) = \eta$ then the length of γ is not smaller than $d_E(\sigma, \eta)/p_4$.*

Proof. Let $\sigma, \eta \in E(\tau, \lambda)$ and let $\xi = \Theta_-(\sigma, \eta)$ be as in Lemma 5.4. Let $\ell_1 = d_E(\xi, \sigma)$, $\ell_2 = d_E(\xi, \eta)$ and assume that $\ell_1 \leq \ell_2$. By Lemma 5.4, there is a train track $\Theta_+(\sigma, \eta)$ which can be obtained from σ by a splitting sequence of length ℓ_2 and which can be obtained from η by a splitting sequence of length ℓ_1 .

Let $p_3(1) > 0$ be as in Lemma 6.7 and let $p = p_1(p_3(1))$ be as in Lemma 6.3. Let $\gamma : [0, n] \rightarrow \mathcal{S}(\lambda)$ be a simplicial path connecting $\sigma = \gamma(0)$ to $\eta = \gamma(n)$. We construct inductively a sequence $\{\sigma(i)\}_{0 \leq i \leq n} \subset \mathcal{V}(\mathcal{S}(\lambda))$ with the following properties.

- (1) For every $i \leq n$, $\sigma(i)$ is carried by $\gamma(i)$ and carries λ .
- (2) For every $i < n$ the train track $\sigma(i+1)$ can be obtained from $\sigma(i)$ by a splitting and shifting sequence of length at most p .

For the construction of the sequence $\{\sigma(i)\}$, we first define $\sigma(0) = \sigma$. Assume by induction that we constructed already the train tracks $\sigma(i)$ for all $i < i_0$ and some $i_0 > 0$. Consider the train track $\gamma(i_0 - 1)$; by assumption, it carries the train track $\sigma(i_0 - 1)$. If the train track $\gamma(i_0)$ is obtained from $\gamma(i_0 - 1)$ by a collapse or a shift then $\gamma(i_0)$ carries $\sigma(i_0 - 1)$ and we define $\sigma(i_0) = \sigma(i_0 - 1)$. Otherwise $\gamma(i_0)$ is obtained from $\gamma(i_0 - 1)$ by a single λ -split. By Lemma 6.7 and Lemma 6.3, there is a train track $\sigma(i_0)$ which carries λ , is carried by $\gamma(i_0)$ and which can be obtained from $\sigma(i_0 - 1)$ by a splitting and shifting sequence of length at most p . The inductively defined sequence $\{\sigma(i)\}$ has the required properties.

By construction, the train track $\sigma = \sigma(0)$ can be connected to $\sigma(n)$ by a splitting and shifting sequence of length at most pn . In particular, we have $\sigma(n) \in C(\sigma, \lambda)$. By Lemma 6.6, Lemma 5.1 and Corollary 5.2, the train track σ is splittable with a sequence of λ -splits of length at most $q_2pn + q_2$ to a train track $\nu \in E(\sigma, \lambda) \subset E(\tau, \lambda)$ which is contained in the q_2 -neighborhood of $\sigma(n)$ in $C(\tau, \lambda)$. Since $\sigma(n) \in C(\gamma(n), \lambda) = C(\eta, \lambda) \subset C(\tau, \lambda)$, another application of Lemma 6.6 shows that via replacing ν by a train track in a uniformly bounded neighborhood we may assume that η is splittable to ν as well. Lemma 5.4 then shows that the train track $\Theta_+(\sigma, \eta) \in E(\tau, \lambda)$ is splittable to ν .

From Lemma 5.1 we deduce that the length $\ell_2 \geq \ell_1$ of a splitting sequence connecting σ to $\Theta_+(\sigma, \eta)$ is not bigger than $q_2pn + L$ where $L > q_2$ is a universal constant. Since the length of a geodesic in $E(\tau, \lambda)$ connecting σ to η equals $\ell_1 + \ell_2 \leq 2\ell_2$, we conclude that the intrinsic distance in $E(\tau, \lambda)$ between σ and η does not exceed $2(q_2pn + L)$. This is just the statement of the lemma. \square

Theorem 2 from the introduction is now an immediate consequence of Lemma 6.8. Namely, by Lemma 6.1 and Proposition 6.5 we only have to show that there is a number $L > 0$ such that for every $\lambda \in \mathcal{CL}$ and every train track $\tau \in \mathcal{V}(\mathcal{S}(\lambda))$, if $\sigma, \eta \in E(\tau, \lambda)$ can be connected in $\mathcal{S}(\lambda)$ by a simplicial path of length $n \geq 0$ then there is an arc in $E(\tau, \lambda)$ connecting σ to η whose length does not exceed $nL + L$. However, this was shown in Lemma 6.8. Theorem 2 is proven. \square

As an immediated consequence of Theorem 2 we obtain the result of Farb, Lubotzky and Minsky [FLM01] (Corollary 4 from the introduction).

Corollary 6.9. *If P is a pants decomposition for S and if Γ is the free abelian group of rank $3g - 3 + m$ generated by the Dehn twists about the pants curves of P then $\Gamma < \mathcal{MCG}$ is undistorted.*

Proof. Let $P = \{\gamma_1, \dots, \gamma_{3g-3+m}\}$ be any pants decomposition for S . For each $i \leq 3g - 3 + m$ let $\varphi_i \in \mathcal{MCG}$ be a simple (positive or negative) Dehn twist about γ_i . The elements $\varphi_1, \dots, \varphi_{3g-3+m}$ generate a free abelian subgroup Γ of \mathcal{MCG} . We equip Γ with the word norm $||$ defined by the generators $\varphi_i, \varphi_i^{-1}$. To show that Γ is undistorted, it suffices to show that there is a number $L > 1$ and that there is a Γ -equivariant L -quasi-isometric embedding of the semi-group $\Gamma_+ = \{\varphi_1^{\ell_1} \circ \dots \circ \varphi_{3g-3+m}^{\ell_{3g-3+m}} \mid \ell_i \geq 0\}$ into a cubical Euclidean cone $E(\tau, \lambda)$ for some $\lambda \in \mathcal{CL}$. Namely, if this holds true then the corollary follows from Theorem 2 and the fact that \mathcal{TT} is \mathcal{MCG} -equivariantly quasi-isometric to \mathcal{MCG} .

For this choose a train track $\tau \in \mathcal{V}(\mathcal{TT})$ which is obtained by collapsing a small tubular neighborhood of a finite complete geodesic lamination λ on S as in Section 2 whose minimal components are precisely the components of P . Up to replacing τ by a shift equivalent train track, every component γ_i of P is carried by an embedded trainpath of length 2 in τ which consists of a large branch and a small branch (note that τ is just a complete train track *in standard form* for the pants decomposition P of S as defined in [PH92]). Moreover, every large branch of τ is of this form. For a suitable choice of the spiraling directions of λ about the components of P , for every $i \in \{1, \dots, 3g - 3 + m\}$ the train track τ is splittable to $\varphi_i\tau$ with a splitting sequence of length 2 (with two splits at a large branch contained in the embedded trainpath γ_i , with the small branch in γ_i as a winner). Moreover, for $i \neq j$ these splitting sequences commute. As a consequence, the cubical Euclidean cone $E(\tau, \lambda)$ is invariant under the semi-group Γ_+ , and the map which associates to an element $\varphi \in \Gamma_+$ the train track $\varphi(\tau) \in E(\tau, \lambda)$ is an equivariant 2-quasi-isometry between Γ_+ and $E(\tau, \lambda)$. \square

7. BOUNDARY AMENABILITY

The mapping class group naturally acts on the compact Hausdorff space \mathcal{CL} of all complete geodesic laminations on S as a group of homeomorphisms. In this section we use the results from Section 6 to show that this action is topologically amenable.

For this we use the assumptions and notations from the previous sections. In particular, for $\tau \in \mathcal{V}(\mathcal{TT})$ and a complete geodesic lamination λ carried by τ let as before $C(\tau, \lambda)$ be the graph whose vertices are the complete train tracks which are carried by τ and carry λ and where two such vertices σ, η are connected by an edge of length one if and only if either σ can be obtained from η by a single shift or σ, η are connected in \mathcal{TT} by an edge of length one. The intrinsic path metric on $C(\tau, \lambda)$ is denoted as before by d_C .

Similarly, the cubical Euclidean cone $E(\tau, \lambda)$ is the full subgraph of \mathcal{TT} whose vertices are the complete train tracks which carry λ and which can be obtained from τ by a splitting sequence. We call τ the *basepoint* of $E(\tau, \lambda)$. The intrinsic path metric on $E(\tau, \lambda)$ is denoted as before by d_E .

Choose a finite subset \mathcal{E} of $\mathcal{V}(\mathcal{TT})$ such that $\cup_{\varphi \in \mathcal{MCG}} \varphi\mathcal{E} = \mathcal{V}(\mathcal{TT})$ and that moreover for every $\lambda \in \mathcal{CL}$ there is a train track $\eta \in \mathcal{E}$ which carries λ . By equivariance under the action of the mapping class group, this implies that for every $\varphi \in \mathcal{MCG}$ and every $\lambda \in \mathcal{CL}$ there is a train track $\eta \in \varphi\mathcal{E}$ which carries λ . As in the proof of Proposition 6.5, such a set exists since \mathcal{CL} is compact and since the set of all complete geodesic laminations which are carried by a complete train track τ is open in \mathcal{CL} .

Let \mathcal{G} be a finite symmetric set of generators for \mathcal{MCG} containing the identity. For an element $\varphi \in \mathcal{MCG}$ let $|\varphi|$ be the word norm of φ with respect to the generating set \mathcal{G} . Recall from the end of Section 5 the definition of a full k -fold λ -split of a complete train track which carries λ . We have.

Lemma 7.1. *There are numbers $0 < \kappa_1 < \kappa_2$ with the following property. Let $k \geq 1$ and let $\varphi \in \mathcal{MCG}$ be such that $|\varphi| = k$. Let $\lambda \in \mathcal{CL}$ and let $\sigma \in \varphi\mathcal{E}, \tau \in \mathcal{E}$ be train tracks which carry λ . Let $n \geq k$ and let ξ be the full $n\kappa_1$ -fold λ -split of σ . Then there is a vertex $\zeta \in E(\tau, \lambda)$ with $d_E(\tau, \zeta) \leq n\kappa_2$ and $d(\zeta, \xi) \leq \kappa_2$.*

Proof. Since $\cup_{g \in \mathcal{G}} g\mathcal{E}$ is finite, its diameter D in \mathcal{TT} is finite as well. Now \mathcal{MCG} acts on \mathcal{TT} as group of isometries and consequently the following holds true. Let $\varphi \in \mathcal{MCG}$ and let $\psi \in \mathcal{G}$. If $\sigma \in \varphi\mathcal{E}$ and if $\eta \in \varphi\psi\mathcal{E}$ then we have $d(\sigma, \eta) \leq D$.

Let $\lambda \in \mathcal{CL}$ be a complete geodesic lamination on S , let $\tau \in \mathcal{E}$ be a train track which carries λ and let $\varphi \in \mathcal{MCG}$ with $|\varphi| = k$. Then φ can be represented in the form $\varphi = \varphi_1 \dots \varphi_k$ with $\varphi_i \in \mathcal{G}$. Let $\sigma \in \varphi\mathcal{E}$ be such that σ carries λ . For $0 \leq i \leq k$ let $\tau_i \in \varphi_1 \dots \varphi_{k-i}\mathcal{E}$ be a train track which carries λ and such that $\tau_0 = \sigma, \tau_k = \tau$. Then for each i , the distance between τ_i and τ_{i-1} is bounded from above by D .

Let $p_3(D) > 0$ be as in Lemma 6.7. We construct inductively a sequence $(\eta_i)_{0 \leq i \leq k}$ of complete train tracks with the following properties.

- (1) $\eta_0 = \sigma$.
- (2) For each i , the train track η_i carries λ .
- (3) For each i , η_i is carried by η_{i-1} and τ_i .
- (4) The distance between η_i and η_{i-1} does not exceed $p_3(D)$.

For the construction, assume by induction that the train tracks η_j are already determined for all $j \leq i-1$ and some $i \geq 1$. To construct η_i , note that since η_{i-1} is carried by τ_{i-1} and since $d(\tau_i, \tau_{i-1}) \leq D$, by Lemma 6.7 there is a complete train track η_i which carries λ , which is carried by both η_{i-1} and τ_i and whose distance to η_{i-1} is at most $p_3(D)$. This yields the construction.

Since η_k is carried by σ , i.e. $\eta_k \in C(\sigma, \lambda)$, by Lemma 6.6 there is train track $\sigma' \in E(\sigma, \lambda)$ whose distance to η_k is at most q_2 . Then the distance between σ' and σ does not exceed $p_3(D)k + q_2$ and therefore by Theorem 2 and Corollary 5.2, the length of a splitting sequence connecting σ to σ' is at most $Lp_3(D)k + L(q_2 + 1)$ where $L > 1$ is a universal constant. By Corollary 5.6, σ' is splittable to the train track σ'' obtained from σ by a full splitting sequence of length $Lp_3(D)k + L(q_2 + 1)$. If $m > 0$ is the number of branches of a complete train track on S , then a full split consists of at most m single splits and hence σ'' is obtained from σ by a splitting sequence in the usual sense of length at most $mLp_3(D)k + m(Lq_2 + 1)$. Moreover, by Lemma 6.7, the distance between σ'' and $C(\eta_k, \lambda) \subset C(\tau, \lambda)$ is at most $p_3(q_2)$. Then Lemma 6.6 shows that the train track σ'' is contained in the $p_3(q_2) + q_2$ -neighborhood of the cubical Euclidean cone $E(\tau, \lambda)$.

On the other hand, we have $d(\tau, \sigma'') \leq d(\tau, \sigma) + d(\sigma, \sigma'') \leq kD + mLp_3(D)k + m(Lq_2 + 1)$ by the above consideration. In other words, the distance between τ and σ'' is bounded from above by $\nu k + \nu$ for a universal constant $\nu > 0$. Then Theorem 2 shows the existence of a universal constant $\kappa_2 > 0$ such that σ'' is contained in the κ_2 -neighborhood of the ball of radius $\kappa_2 k$ in $E(\tau, \lambda)$. This shows the lemma. \square

We are now ready to complete the proof of Theorem 1 from the introduction.

Proposition 7.2. *The action of \mathcal{MCG} on \mathcal{CL} is topologically amenable.*

Proof. Let $\mathcal{E} \subset \mathcal{V}(\mathcal{TT})$ be as in Lemma 7.1, and for $\lambda \in \mathcal{CL}$ and $\varphi \in \mathcal{MCG}$ let $\tau(\lambda, \varphi) \in \varphi\mathcal{E}$ be a train track which carries λ . Following Kaimanovich [Ka04], for $n \geq 1, k \leq 2n$ and $\varphi \in \mathcal{MCG}$ define $Y(\varphi, \lambda, n, k)$ to be the set of all complete train tracks $\sigma \in \mathcal{V}(\mathcal{TT})$ with the following property. There is an element $\psi \in \mathcal{MCG}$ with $|\psi\varphi^{-1}| \leq k$ and there is a complete train track $\sigma_0 \in \psi\mathcal{E}$ which carries λ and such that σ can be obtained from σ_0 by a full $2n\kappa_1$ -fold λ -split where $\kappa_1 > 0$ is as in Lemma 7.1.

By Lemma 7.1 and invariance under the action of the mapping class group, the set $Y(\varphi, \lambda, n, k)$ is contained in the κ_2 -neighborhood in \mathcal{TT} of the ball of radius $2n\kappa_2$ about the basepoint in the cubical Euclidean cone $E(\tau(\lambda, \varphi), \lambda)$ with respect to the intrinsic metric. Since the number of vertices of \mathcal{TT} contained in any ball of radius κ_2 is uniformly bounded, by Corollary 5.5 there is a number $C > 0$ such that the number of points in the set $Y(\varphi, \lambda, n, k)$ is bounded from above by Cn^C . In particular, the cardinality of the sets $Y(\varphi, \lambda, n, k)$ is bounded from above by a universal polynomial in n .

By construction, for any $\varphi, \varphi' \in \mathcal{MCG}$ with $q = |\varphi'\varphi^{-1}|$ and for every $\lambda \in \mathcal{CL}$ the sets $Y(\varphi, \lambda, n, k), Y(\varphi', \lambda, n, k')$ satisfy the following *nesting condition* from Lemma 1.35 of [Ka04] (see the proof of Corollary 1.37 in [Ka04]): For every $n > q$ and every $k \leq 2n - q$ we have

$$(2) \quad Y(\varphi, \lambda, n, k) \subset Y(\varphi', \lambda, n, k + q) \text{ and } Y(\varphi', \lambda, n, k) \subset Y(\varphi, \lambda, n, k + q).$$

Let $\mathcal{P}(\mathcal{TT})$ be the space of all probability measures on the set $\mathcal{V}(\mathcal{TT})$ of all vertices of \mathcal{TT} . A probability measure on a countable set Z can be viewed as a non-negative function f on Z with $\sum_{z \in Z} f(z) = 1$. In other words, such a probability measure is a point in the space $\ell^1(Z)$ of integrable functions on Z and hence the space of probability measures on Z can be equipped with the usual Banach norm on $\ell^1(Z)$. Thus $\mathcal{P}(\mathcal{TT})$ is equipped with a natural norm $\|\cdot\|$.

Let $m_{Y(\varphi, \lambda, n, k)} \in \mathcal{P}(\mathcal{TT})$ be the normalized counting measure for $Y(\varphi, \lambda, n, k)$, i.e. the probability measure which is the normalization of the sum of the Dirac measures on the points in $Y(\varphi, \lambda, n, k)$. Define

$$\nu_n(\varphi, \lambda) = \frac{1}{n} \sum_{k=1}^n m_{Y(\varphi, \lambda, n, k)}.$$

Denote by $|B|$ the cardinality of a finite subset of $\mathcal{V}(\mathcal{TT})$. Since for each fixed n the family of sets $Y(\varphi, \lambda, n, k)$ satisfies the nesting condition (2) above, Lemma 1.35 of [Ka04] shows that whenever $|\varphi'\varphi^{-1}| \leq q$ and $n > q$ then we have

$$\|\nu_n(\varphi, \lambda) - \nu_n(\varphi', \lambda)\| \leq \frac{2q}{n} + \frac{4(n-q)}{n} \left[1 - \left(\frac{\text{const}}{|Y(\varphi, \lambda, n, n+q)|} \right)^{2q/(n-q)} \right].$$

The consideration in the second paragraph of this proof shows that for $n > q$ the cardinality of the set $Y(\varphi, \lambda, n, n+q)$ is bounded from above by a fixed polynomial in n . As a consequence, we obtain that

$$(3) \quad \|\nu_n(\varphi, \lambda) - \nu_n(\varphi', \lambda)\| \rightarrow 0$$

as $n \rightarrow \infty$ and locally uniformly on \mathcal{MCG} .

For $n \geq 0$ and $\lambda \in \mathcal{CL}$ define

$$\mu_n(\lambda) = \nu_n(e, \lambda).$$

Then $\varphi\mu_n(\lambda) = \nu_n(\varphi, \varphi\lambda)$ for every $\varphi \in \mathcal{MCG}$ and hence by (3) above, the measures $\mu_n(\lambda) \in \mathcal{P}(\mathcal{TT})$ satisfy $\|\mu_n(\varphi\lambda) - \varphi\mu_n(\lambda)\| \rightarrow 0$ ($n \rightarrow \infty$) uniformly on compact subsets of $\mathcal{MCG} \times \mathcal{CL}$. Since the action of \mathcal{MCG} on \mathcal{TT} is isometric and properly discontinuous, the measures $\mu_n(\lambda)$ on \mathcal{TT} can be lifted to measures on \mathcal{MCG} with the same property (see [AR00] and compare [Ka04]). As a consequence, the action of \mathcal{MCG} on \mathcal{CL} is topologically amenable. \square

8. A STRONG BOUNDARY FOR \mathcal{MCG}

A *strong boundary* for a locally compact second countable topological group Γ is a standard probability space (X, μ) with a measure class preserving action of Γ and the following two additional properties.

- (1) The action of Γ on (X, μ) is amenable.
- (2) *Double ergodicity*: Let (E, π) be any separable coefficient module for Γ ; then every Γ -equivariant weak*-measurable map $f : (X \times X, \mu \times \mu) \rightarrow E$ is constant almost everywhere.

Kaimanovich [Ka03] showed that for every locally compact second countable topological group Γ , the Poisson boundary of every spread-out non-degenerate symmetric probability measure on Γ is a strong boundary for Γ .

The mapping class group acts on the Teichmüller space equipped with the *Teichmüller metric* as a group of isometries. Hence the restriction of the Teichmüller distance to an orbit of \mathcal{MCG} defines a \mathcal{MCG} -invariant distance function on \mathcal{MCG} (which however is not quasi-isometric to the distance function defined by a word norm). In [KM96], Kaimanovich and Masur investigate the Poisson boundary for a symmetric probability measure μ of finite entropy on the mapping class group \mathcal{MCG} . They show that if the first logarithmic moment of μ with respect to the Teichmüller distance is finite and if the support of μ generates \mathcal{MCG} , then the Poisson boundary of μ can be viewed as a measure on the space \mathcal{PML} of all projective measured geodesic laminations on S . The measure class of this measure is \mathcal{MCG} -invariant and gives full mass to the subset of \mathcal{PML} of projective measured geodesic laminations whose support is uniquely ergodic.

There is a particular \mathcal{MCG} -invariant measure class on \mathcal{PML} , the *Lebesgue measure class*, which can be obtained from the family of local linear structures on \mathcal{PML} defined by complete train tracks. Namely, the transverse measure of every measured geodesic lamination whose support is carried by a complete train track τ defines a nonnegative weight function on the branches of τ which satisfies the switch conditions. Vice versa, every nonnegative weight function on τ satisfying the switch conditions determines a measured geodesic lamination whose support is carried by τ [PH92].

The switch conditions are a system of linear equations with integer coefficients and hence the solutions of these equations have the structure of a linear space of dimension $6g - 6 + 2k$. The standard Lebesgue measure on $\mathbb{R}^{6g-6+2k}$ then induces via the thus defined coordinate system a Lebesgue measure on the closure of an open subset of the space of all measured geodesic laminations which is invariant under scaling. This measure projects to a measure class on the space of projective measured geodesic laminations. Since the transformations of weight functions induced by carrying maps are linear, these locally defined measure classes do not depend on the choice of the train track used to define the coordinates and hence they define a \mathcal{MCG} -invariant measure class on \mathcal{PML} .

If ν is a measured geodesic lamination whose support is not maximal, then this support is carried by a birecurrent generic train track which is not maximal. Thus the set of all measured geodesic laminations whose support is not maximal is contained in a countable union of "hyperplanes" in \mathcal{PML} and hence has vanishing Lebesgue measure. In other words, the Lebesgue measure class gives full measure to the set of projective measured geodesic laminations whose support is maximal.

Now the support of a measured geodesic lamination is a disjoint union of minimal components. Such a geodesic lamination can only be maximal if it consists of a single minimal component and hence if it is minimal as well. Moreover, by a result of Masur [M82], the Lebesgue measure class gives full mass to the set of projective measured geodesic laminations whose support is *uniquely ergodic*, i.e. it admits a unique transverse measure up to scale. In other words, the measure-forgetting map restricted to a Borel subset \mathcal{D} of \mathcal{PML} of full Lebesgue measure defines an \mathcal{MCG} -equivariant measurable bijection of \mathcal{D} onto a Borel subset of \mathcal{CL} and hence determines an \mathcal{MCG} -invariant measure class λ on \mathcal{CL} which we call the *Lebesgue measure class*. Since by Proposition 7.2 the action of \mathcal{MCG} on \mathcal{CL} is topologically amenable, the action of \mathcal{MCG} on \mathcal{CL} with respect to the Lebesgue measure class is amenable [AR00]. Then the action of \mathcal{MCG} on \mathcal{PML} with respect to the Lebesgue measure class is amenable as well. This is used to show.

Proposition 8.1. *\mathcal{PML} equipped with the Lebesgue measure class is a strong boundary for \mathcal{MCG} .*

Proof. By our above discussion, we only have to show double ergodicity for the action of \mathcal{MCG} with respect to the Lebesgue measure class. This in turn follows immediately from a celebrated result of Masur [M82] and Veech [V86]. For completeness, we show how their result translates into our context.

Let again $\mathcal{T}_{g,k}$ be the Teichmüller space of all marked complete hyperbolic metrics on S of finite volume. The *Teichmüller geodesic flow* is a flow acting on the bundle $Q^1 \rightarrow \mathcal{T}_{g,k}$ of area one *quadratic differentials* over Teichmüller space. This flow projects to a flow on the quotient $Q^1/\mathcal{MCG} = Q$. The flow on Q preserves a Borel probability measure μ in the Lebesgue measure class which lifts to an \mathcal{MCG} -invariant Lebesgue measure λ_0 on Q^1 [M82].

By the Hubbard-Masur theorem [HM79], for every point $z \in \mathcal{T}_{g,k}$ and every projective measured geodesic lamination $[\nu]$ on S there is a unique quadratic differential $q \in Q_z^1$ whose *vertical measured geodesic lamination* is contained in the class $[\nu]$

(where we identify a projective measured geodesic lamination with an equivalence class of projective measured foliations on S in the usual way). As z varies over $\mathcal{T}_{g,k}$, the set of all these quadratic differentials defines a topological submanifold $W^s([\nu])$ of \mathcal{Q}^1 which projects homeomorphically onto $\mathcal{T}_{g,k}$.

If the support of $[\nu]$ is uniquely ergodic and maximal then more can be said. Namely, for every projective measured geodesic lamination $[\nu]' \neq [\nu]$ there is a unique Teichmüller geodesic line in $\mathcal{T}_{g,k}$ generated by a quadratic differential whose vertical measured geodesic lamination is contained in the class $[\nu]$ and whose horizontal measured geodesic lamination is contained in the class $[\nu]'$. Moreover, if $\gamma, \gamma' : \mathbb{R} \rightarrow \mathcal{T}_{g,k}$ are Teichmüller geodesics with vertical measured geodesic lamination contained in the class $[\nu]$, then there is a unique number $a \in \mathbb{R}$ such that $\lim_{t \rightarrow \infty} d(\gamma(t), \gamma'(t+a)) = 0$ [M82].

Thus if we denote by \mathcal{UL} the set of all projective measured geodesic laminations on S whose support is uniquely ergodic and maximal, then the set of unit area quadratic differentials in \mathcal{Q}^1 with horizontal and vertical measured geodesic lamination whose support is uniquely ergodic and maximal is (non-canonically) homeomorphic to $(\mathcal{UL} \times \mathcal{UL} - \Delta) \times \mathbb{R}$ (where Δ is the diagonal) by associating to a pair of distinct points in \mathcal{UL} and the origin in \mathbb{R} a point on the geodesic determined by the pair of projective measured geodesic laminations and extending this map to $(\mathcal{UL} \times \mathcal{UL} - \Delta) \times \mathbb{R}$ in such a way that the Teichmüller geodesic flow acts on $(\mathcal{UL} \times \mathcal{UL} - \Delta) \times \mathbb{R}$ via $\Phi^t([\nu], [\nu'], s) = ([\nu], [\nu'], t+s)$.

The Lebesgue measure λ_0 is locally of the form $d\lambda_0 = d\mu^+ \times d\mu^- \times dt$ where μ^\pm is a measure on the space of projective measured geodesic laminations which gives full mass to the projective measured geodesic laminations whose support is uniquely ergodic and maximal and which is uniformly exponentially expanding (contracting) under the Teichmüller geodesic flow. The Lebesgue measure class on $\mathcal{PML} \times \mathcal{PML}$ can be obtained by disintegration of the Lebesgue measure on $(\mathcal{UL} \times \mathcal{UL} - \Delta) \times \mathbb{R}$. As a consequence, double ergodicity for the action of \mathcal{MCG} on $(\mathcal{UL} \times \mathcal{UL}) - \Delta$ with respect to the Lebesgue measure class is equivalent to (strong) ergodicity of the measure μ on the moduli space \mathcal{Q} of area one quadratic differentials under the Teichmüller geodesic flow. This ergodicity was established by Masur and Veech [M82, V86]. This shows the proposition. \square

The *curve graph* $\mathcal{C}(S)$ of S is the metric graph whose vertices are the free homotopy classes of essential simple closed curves on S and where two such curves are joined by an edge of length one if and only if they can be realized disjointly. The curve graph is a hyperbolic geodesic metric graph [MM99], and the mapping class group acts on $\mathcal{C}(S)$ as a group of simplicial isometries.

The *Gromov boundary* $\partial\mathcal{C}(S)$ of the curve graph of S can be identified with the space of minimal geodesic laminations which fill up S (i.e. whose complementary components are topological discs or once punctured topological discs) equipped with the *coarse Hausdorff topology*: A sequence $\{\lambda_i\} \subset \partial\mathcal{C}(S)$ converges to a lamination λ if and only if λ is the minimal component of every accumulation point of this sequence with respect to the usual Hausdorff topology on the space of geodesic

laminations [Kl99, H06]. This description of $\partial\mathcal{C}(S)$ is used to derive the second corollary from the introduction.

Corollary 8.2. *The action of \mathcal{MCG} on the Gromov boundary of the curve graph is universally amenable.*

Proof. Let $\mathcal{A} \subset \mathcal{CL}$ be the set of all complete geodesic laminations which contain a minimal component filling up S . Then \mathcal{A} is a Borel subset of \mathcal{CL} . Namely, for a simple closed curve c on S let $B(c)$ be the set of all complete geodesic laminations which either contain c as a minimal component or are such that they intersect c transversely in finitely many points. Clearly $B(c)$ is a countable union of closed subsets of \mathcal{CL} and therefore $B(c)$ is a Borel set. On the other hand, the complement of \mathcal{A} in \mathcal{CL} is the countable union of the Borel sets $B(c)$ where c ranges over the simple closed curves on S . Then the set \mathcal{A} is a Borel set as well. Moreover, \mathcal{A} is invariant under the action of the mapping class group.

There is a natural continuous \mathcal{MCG} -equivariant map $\varphi : \mathcal{A} \rightarrow \partial\mathcal{C}(S)$ which maps a lamination $\lambda \in \mathcal{A}$ to its (unique) minimal component. By Lemma 2.2, the map φ is surjective. Moreover, φ is finite-to-one, which means that the cardinality of the preimage of a point in $\partial\mathcal{C}(S)$ is bounded from above by a universal number $\ell > 0$. Since the action of \mathcal{MCG} on \mathcal{CL} is topologically amenable, the action of \mathcal{MCG} on \mathcal{A} is universally amenable [AR00]. Now the map φ is finite-to-one and therefore by Lemma 3.6 of [A96] the action of \mathcal{MCG} on $\partial\mathcal{C}(S)$ is universally amenable as well. This shows the corollary. \square

Since the set of projective measured geodesic laminations whose supports are uniquely ergodic, maximal and minimal embeds into the Gromov boundary of the curve graph [Kl99, H06], the Lebesgue measure class on \mathcal{PML} induces an \mathcal{MCG} -invariant measure class on $\partial\mathcal{C}(S)$. The following corollary is immediate from Proposition 8.1.

Corollary 8.3. *The Gromov boundary of the curve graph equipped with the Lebesgue measure class is a strong boundary for \mathcal{MCG} .*

9. SUPER-RIGIDITY OF COCYCLES

In this section we show Theorem 3 from the introduction. We begin with having a closer look at the action of \mathcal{MCG} on the curve graph $\mathcal{C}(S)$ of S .

An action of a group Γ on a Borel space X is called *tame* if there exists a countable collection Θ of Γ -invariant Borel subsets of X which *separates orbits*. This means that if X_1, X_2 are Γ -orbits and if $X_1 \neq X_2$ then there is some $X_0 \in \Theta$ such that $X_1 \subset X_0$ and $X_2 \cap X_0 = \emptyset$ [Z84, A96].

Let $\mathcal{MC}(S)$ be the countable set of multi-curves on S equipped with the discrete topology and let $\mathcal{PMC}(S)$ be the space of probability measures on $\mathcal{MC}(S)$. The mapping class group acts naturally on $\mathcal{MC}(S)$ and hence on $\mathcal{PMC}(S)$. We have.

Lemma 9.1. *The action of \mathcal{MCG} on \mathcal{PMCS} is tame, and the stabilizer of a point either is a finite subgroup of \mathcal{MCG} or it is contained in the stabilizer of a multi-curve.*

Proof. Since the set $\mathcal{MC}(S)$ of all multi-curves on S is countable, the action of \mathcal{MCG} on $\mathcal{MC}(S)$ is tame. Let $\mathcal{Q} = \{(v, q_1, q_2) \in \mathcal{MC}(S) \times \mathbb{Q} \times \mathbb{Q} \mid 0 \leq q_1 < q_2 \leq 1\}$. The mapping class group acts on \mathcal{Q} , and this action is tame. Moreover, each point $(v, q_1, q_2) \in \mathcal{Q}$ defines a Borel subset $P(q)$ of \mathcal{PMCS} by $P(q) = \{\mu \mid \mu(v) \in [q_1, q_2]\}$. The \mathcal{MCG} -orbits of these sets together with their finite intersections define a countable collection of \mathcal{MCG} -invariant Borel subsets of \mathcal{PMCS} which clearly separates orbits. This yields that the action of \mathcal{MCG} on \mathcal{PMCS} is tame.

We are left with showing that the stabilizer of a probability measure $\mu \in \mathcal{PMCS}$ either is a finite subgroup of \mathcal{MCG} or is contained in the stabilizer of a multi-curve. For this note that a mapping class which preserves the measure μ preserves the subset C of $\mathcal{MC}(S)$ on which the value of μ is maximal. Since this set is finite and since a *pseudo-Anosov* element of \mathcal{MCG} acts on the curve graph with unbounded orbits, the stabilizer of μ does not contain any pseudo-Anosov element. By a structural result of McCarthy and Papadopoulos [MP89] for subgroups of \mathcal{MCG} , this implies that either this stabilizer is finite or that it is contained in the stabilizer of a multi-curve as claimed. \square

The Gromov boundary $\partial\mathcal{C}(S)$ of the curve graph $\mathcal{C}(S)$ of S can be identified with the space of all minimal geodesic laminations on S which fill up S , equipped with the coarse Hausdorff topology (see [K199, H06] and Section 8). In particular, $\partial\mathcal{C}(S)$ is non-compact. Denote by $\mathcal{P}(\partial\mathcal{C}(S))$ the space of probability measures on the Gromov boundary $\partial\mathcal{C}(S)$ of $\mathcal{C}(S)$ and let $\mathcal{P}_{\geq 3}$ be the subspace of all measures whose support contains at least three points. The mapping class group acts on $\mathcal{P}_{\geq 3}$ as a group of automorphisms. We have.

Lemma 9.2. *The action of \mathcal{MCG} on $\mathcal{P}_{\geq 3}$ is tame, and the stabilizer of a point is a finite subgroup of \mathcal{MCG} .*

Proof. Let T be the space of triples of pairwise distinct points in $\partial\mathcal{C}(S)$. Every measure $\mu \in \mathcal{P}_{\geq 3}$ induces a non-trivial finite Borel measure $\mu^3 = \mu \times \mu \times \mu$ on T . The mapping class group acts diagonally on T as a group of homeomorphisms. Every point $x \in T$ has a neighborhood N in T such that $g(N) \cap N \neq \emptyset$ only for finitely many $g \in \mathcal{MCG}$ (we refer to [H08] for a discussion how this follows from a result of Bowditch [Bw08]). Since T is second countable, this immediately implies that the action of \mathcal{MCG} on the space of probability measures on T is tame. But then the action of \mathcal{MCG} on $\mathcal{P}_{\geq 3}$ is tame as well.

To show that the stabilizer of $\mu \in \mathcal{P}_{\geq 3}$ under the action of \mathcal{MCG} is finite, choose a subset V of T whose closure in $\partial\mathcal{C}(S) \times \partial\mathcal{C}(S) \times \partial\mathcal{C}(S)$ does not contain points in the complement of T and such that $\mu^3(V) > 3\mu^3(T)/4$. Such a set exists since μ is Borel and T is an open subset of the metrizable space $\partial\mathcal{C}(S) \times \partial\mathcal{C}(S) \times \partial\mathcal{C}(S)$. Now if $g \in \mathcal{MCG}$ stabilizes μ then $gV \cap V \neq \emptyset$. On the other hand, there are only finitely many elements $g \in \mathcal{MCG}$ such that $gV \cap V \neq \emptyset$ [H08]. Therefore the stabilizer of μ is finite. \square

As in Section 8, let $\mathcal{A} \subset \mathcal{CL}$ be the Borel subset of all complete geodesic laminations which contain a minimal component filling up S . Write $\mathcal{B} = \mathcal{CL} - \mathcal{A}$. We have.

Lemma 9.3. *There is a MCG -equivariant Borel map $\Phi : \mathcal{B} \rightarrow \mathcal{MC}(S)$.*

Proof. For every $\lambda \in \mathcal{B}$, the minimal arational components of λ do not fill up S . Thus a minimal arational component $\tilde{\lambda}$ of λ fills a proper connected subsurface \tilde{S} of S which is determined by requiring that it contains $\tilde{\lambda}$ and that every non-peripheral simple closed curve on \tilde{S} intersects $\tilde{\lambda}$ transversely. The boundary of \tilde{S} consists of a non-empty collection of essential simple closed curves on S which up to free homotopy are disjoint from all minimal components of λ . As a consequence, the union of these (geodesic) boundary curves for all minimal arational components of λ with the closed curve components of λ defines a non-trivial multi-curve on S . We obtain a natural MCG -equivariant map $\Phi : \mathcal{B} \rightarrow \mathcal{MC}(S)$ by associating to $\lambda \in \mathcal{B}$ this multi-curve.

We claim that the map Φ is Borel. For this let c be any multi-curve and let S_1, \dots, S_k be the components of $S - c$, ordered in such a way that the pairs of pants among these components are precisely the surfaces $S_{\ell+1}, \dots, S_k$. Let J be any (possibly empty) subset of $\{1, \dots, \ell\}$ and let $C_0(J)$ be the set of all geodesic laminations which consist of $|J|$ minimal arational components and such that each of these components fills one of the subsurfaces S_j for $j \in J$. Let $C(J) \subset \Phi^{-1}(c)$ be the set of all complete geodesic laminations $\zeta \in \Phi^{-1}(c)$ whose minimal arational components fill the surfaces S_j for $j \in J$. Then $\Phi^{-1}(c)$ is a finite disjoint union of the sets $C(J)$ and hence Φ is Borel if each of the sets $C(J)$ is a Borel set.

To show that this is the case, note that there is a natural map $\Phi_0 : C(J) \rightarrow C_0(J)$. For a fixed geodesic lamination $\nu \in C_0(J)$, the set $\Phi_0^{-1}(\nu) \subset C(J) \subset \Phi^{-1}(c)$ is countable. Namely, every lamination $\zeta \in \Phi_0^{-1}(\nu)$ is determined by the union of its minimal components not contained in ν which are components of c and by additional finitely many isolated leaves which spiral about one or two of the minimal components. Each of the spiraling leaves is determined by the minimal components of ζ contained in its closure, the spiraling direction and by the homotopy class relative to the boundary of its intersections with the components S_i .

On the other hand, every complete geodesic lamination $\xi \in C(J)$ can be obtained from a point $\zeta \in \Phi_0^{-1}(\nu)$ by replacing some of the components of ν by another such minimal geodesic lamination filling the same component of $S - c$. As a consequence, the set $C(J)$ is a countable union $\cup_{\zeta \in \Phi_0^{-1}(\nu)} Q(\zeta)$ where for $\zeta \in \Phi_0^{-1}(\nu)$ the set $Q(\zeta)$ consists of all complete geodesic laminations which can be obtained from ζ by exchanging some of the minimal arational components by some other minimal arational components filling up the same subsurface. Now to show that $\Phi^{-1}(c)$ is a Borel set, it suffices to show that each of the sets $Q(\zeta)$ ($\zeta \in \Phi_0^{-1}(\nu)$) is Borel.

Thus let $\zeta \in \Phi_0^{-1}(\nu)$ for $\nu \in C_0(J)$ be fixed. The set $Q'(\zeta)$ of all complete geodesic laminations whose restriction to $S - \cup_{j \in J} S_j$ coincides with ζ is a closed subset of \mathcal{CL} . Hence by the consideration in the proof of Corollary 8.2, the subset $Q(\zeta)$ of $Q'(\zeta)$ is a Borel subset of \mathcal{CL} . This shows that indeed $\Phi^{-1}(c)$ is a Borel subset of \mathcal{B} and completes the proof of the lemma. \square

Now let $n \geq 2$ and for $i \leq n$ let G_i be a locally compact second countable topological group. Let $G = G_1 \times \cdots \times G_n$ and let $\Gamma < G$ be an irreducible lattice. This means that Γ is a discrete subgroup of G such that the volume of G/Γ with respect to a Haar measure λ on G is finite (in particular, G is unimodular), and that for every i the projection of Γ to G_i is dense. The group G acts on G/Γ by left translation preserving the projection of the Haar measure.

Let (X, μ) be a standard probability space with a measure preserving mildly mixing action of Γ . An \mathcal{MCG} -valued *cocycle* for the action of Γ on (X, μ) is a measurable map $\alpha : \Gamma \times X \rightarrow \mathcal{MCG}$ such that $\alpha(gh, x) = \alpha(g, hx)\alpha(h, x)$ for almost all $x \in X$, all $g, h \in \Gamma$. The cocycle α is cohomologous to a cocycle α' if there is a measurable map $\varphi : X \rightarrow \mathcal{MCG}$ such that $\varphi(gx)\alpha(g, x) = \alpha'(g, x)\varphi(x)$ for every $g \in \Gamma$ and almost every $x \in X$. We show.

Proposition 9.4. *Let α be an \mathcal{MCG} -valued cocycle for the action of Γ on (X, μ) . Then α is cohomologous to a cocycle α' which satisfies one of the following four possibilities.*

- (1) *There is some $i \leq n$ and there is a finite subgroup K of \mathcal{MCG} with normalizer N such that α' is induced from a continuous homomorphism $G \rightarrow N/K$ which factors through G_i .*
- (2) *The image of α' is contained in a finite subgroup of \mathcal{MCG} .*
- (3) *The image of α' is contained in a virtually abelian subgroup of \mathcal{MCG} .*
- (4) *The image of α' is contained in the stabilizer of a multi-curve.*

Proof. Let Γ be an irreducible lattice in $G = G_1 \times \cdots \times G_n$ and let (X, μ) be a standard probability space with a mildly mixing measure preserving action of Γ . Since $\Gamma < G$ is a lattice, the quotient space $(G \times X)/\Gamma$ can be viewed as a bundle over G/Γ with fiber X .

Let $\Omega \subset G$ be a Borel fundamental domain for the right action of Γ on G . Then $\Omega \times X \subset G \times X$ is a finite measure Borel fundamental domain for the action of Γ on $G \times X$. Thus up to normalization, the product $\lambda \times \mu$ of a Haar measure λ on G and the measure μ on X projects to a probability measure ν on $(G \times X)/\Gamma$. The action of G on $G \times X$ by left translation commutes with the right action of Γ and hence it projects to an action on $(G \times X)/\Gamma$ preserving the measure ν (see p.75 of [Z84]). This action is ergodic.

For $z \in \Omega$ and $g \in G$ let $\eta(g, z) \in \Gamma$ be the unique element such that $gz \in \Omega\eta(g, z)^{-1}$. Clearly η is a Borel map. Let $\alpha : \Gamma \times X \rightarrow \mathcal{MCG}$ be a cocycle. We then obtain a $\lambda \times \nu$ -measurable map $\beta : G \times (G \times X)/\Gamma \rightarrow \mathcal{MCG}$ by defining $\beta(g, (z, \sigma)) = \alpha(\eta(g, z), \sigma)$. By construction, β satisfies the cocycle equation for the action of G on $(G \times X)/\Gamma$.

Let (B, ζ) be a strong boundary for the group G (see the beginning of Section 8). We may assume that B is also a strong boundary for Γ . By Proposition 4.3.4 of [Z84] the action of G on $(G \times X)/\Gamma \times B$ is amenable. Let \mathcal{PCL} be the space of all Borel probability measures on \mathcal{CL} . Since \mathcal{CL} is a compact metric \mathcal{MCG} -space, by Proposition 4.3.9 of [Z84] there is a β -equivariant Furstenberg map $f : (G \times X)/\Gamma \times B \rightarrow \mathcal{PCL}$. Now (B, ζ) is a strong boundary for G and hence the

diagonal action of G on $(G \times X)/\Gamma \times B$ is ergodic (Proposition 2.4 of [MS04]). Therefore if there is a subset D of $(G \times X)/\Gamma \times B$ of positive measure such that for every $x \in D$ the support of $f(x)$ meets the \mathcal{MCG} -invariant Borel subset \mathcal{A} of \mathcal{CL} of all complete geodesic laminations which contain a minimal component filling up S , then this is the case for almost every $x \in (G \times X)/\Gamma \times B$. Since every measure on \mathcal{CL} can be written as a sum of a measure on \mathcal{A} and a measure on $\mathcal{B} = \mathcal{CL} - \mathcal{A}$, we may assume without loss of generality that either for almost every x the measure $f(x)$ gives full mass to the set \mathcal{A} or that for almost all x the measure $f(x)$ gives full mass to the set \mathcal{B} .

We distinguish now two cases.

Case 1: For almost all $x \in (G \times X)/\Gamma \times B$, the measure $f(x)$ gives full mass to the set \mathcal{B} .

By Lemma 9.3, there is an equivariant Borel map $\Phi : \mathcal{B} \rightarrow \mathcal{MC}(S)$. The map Φ induces an \mathcal{MCG} -equivariant map Φ_* of the space of Borel probability measures on \mathcal{B} into the space $\mathcal{PMC}(S)$ of probability measures on $\mathcal{MC}(S)$. Since by Lemma 9.1 the action of \mathcal{MCG} on $\mathcal{PMC}(S)$ is tame, with stabilizers which either are finite or preserve a multi-curve, there is a β -equivariant map $\tilde{f} : (G \times X)/\Gamma \times B \rightarrow \mathcal{MCG}/\mathcal{M}_0$ where $\mathcal{M}_0 < \mathcal{MCG}$ either is a finite group or is the stabilizer of a multi-curve.

Since the diagonal action of G on $(G \times X)/\Gamma \times B \times B$ is also ergodic (Proposition 2.4 of [MS04]), if \mathcal{M}_0 is a finite group then Lemma 5.2.10 of [Z84] implies that β and hence α is cohomologous to a cocycle ranging in a finite subgroup of \mathcal{MCG} . Thus in this case the second alternative in the proposition holds.

If \mathcal{M}_0 is the stabilizer of a multi-curve c , then as before, α is cohomologous to a cocycle with image in the stabilizer of a multi-curve, i.e. the fourth alternative in the proposition holds true. To see that this is indeed the case, it is enough to inspect the proof of Lemma 5.2.10 of [Z84]. Namely, in this case the quotient space $\mathcal{MCG}/\mathcal{M}_0$ can be identified with the space of multi-curves which are topologically equivalent to c , i.e. which have the same number of components as c and which decompose the surface S into connected components of the same topological types. Since \mathcal{MCG} is countable, the left action of \mathcal{MCG} on $\mathcal{MCG}/\mathcal{M}_0 \times \mathcal{MCG}/\mathcal{M}_0$ is tame. For $s \in (G \times X)/\Gamma$ let $\tilde{f}_s : B \times B \rightarrow \mathcal{MCG}/\mathcal{M}_0 \times \mathcal{MCG}/\mathcal{M}_0$ be given by $\tilde{f}_s(y_1, y_2) = (f(s, y_1), f(s, y_2))$. Since the action of G on $(G \times X)/\Gamma \times B \times B$ is ergodic, there is an \mathcal{MCG} -orbit \mathcal{O} for the action of \mathcal{MCG} on $\mathcal{MCG}/\mathcal{M}_0 \times \mathcal{MCG}/\mathcal{M}_0$ such that for ν -almost every $s \in (G \times X)/\Gamma$ and for ζ -almost all $y_1, y_2 \in B$ we have $\tilde{f}_s(y_1, y_2) \in \mathcal{O}$. On the other hand, since $\mathcal{MCG}/\mathcal{M}_0$ is countable, the measure $(\tilde{f}_s)_*(\zeta \times \zeta)$ is a purely atomic product measure and therefore it does not vanish on the diagonal of $\mathcal{MCG}/\mathcal{M}_0 \times \mathcal{MCG}/\mathcal{M}_0$. This diagonal is invariant under the action of \mathcal{MCG} and hence it coincides with the orbit \mathcal{O} . Therefore for ν -almost every $s \in (G \times X)/\Gamma$ the measure $(\tilde{f}_s)_*\zeta$ is supported on a single point. This implies that \tilde{f} induces a measurable β -equivariant map $(G \times X)/\Gamma \rightarrow \mathcal{MCG}/\mathcal{M}_0$. By the cocycle reduction lemma (Lemma 5.2.11 of [Z84]), this just means that the cocycle β is cohomologous to a cocycle into \mathcal{M}_0 and hence the same is true for α .

Case 2: For almost all $x \in (G \times X)/\Gamma \times B$ the measure $f(x)$ gives full mass to the invariant Borel set \mathcal{A} .

Via composing f with the natural \mathcal{MCG} -equivariant map $\mathcal{A} \rightarrow \partial\mathcal{C}(S)$ (compare Section 8) we may assume that f maps $(G \times X)/\Gamma \times B$ into the space of probability measures on $\partial\mathcal{C}(S)$. By Lemma 9.2 the action of \mathcal{MCG} on $\mathcal{P}_{\geq 3}$ is tame, with finite point stabilizers. Thus if there is a set of positive measure in $(G \times X)/\Gamma \times B$ such that for every point x in this set the support of $f(x)$ contains at least 3 points, then we conclude as in Case 1 (and following [Z84]) that α is cohomologous to a cocycle ranging in a finite subgroup of \mathcal{MCG} , i.e. the second possibility in the proposition is satisfied.

In the case that for almost all x the measure $f(x)$ is supported on two distinct points, we follow Lemma 23 in [MMS04]. Namely, let \mathcal{D}_2 be the space of subsets of cardinality two in $\partial\mathcal{C}(S)$ and let \mathcal{D} be the set of pairs of distinct (not necessarily disjoint) points $(x, y) \in \mathcal{D}_2 \times \mathcal{D}_2$. The mapping class group naturally acts on \mathcal{D} as a group of transformations.

By hyperbolicity, there is a family of distance functions δ_z on $\partial\mathcal{C}(S)$ ($z \in \mathcal{C}(S)$) of uniformly bounded diameter [BH99]. For two disjoint elements $x = \{x_1, x_2\}$ and $y = \{y_1, y_2\}$ of \mathcal{D}_2 let

$$\rho'(x, y) = \left| \log \frac{\delta_z(x_1, y_1)\delta_z(x_2, y_2)}{\delta_z(x_1, y_2)\delta_z(y_2, x_1)} \right|^{-1}.$$

This function extends to a continuous function $\rho' : \mathcal{D} \rightarrow [0, \infty)$ such that $\rho'(x, y) = 0$ if and only if $x \cap y \neq \emptyset$ (see [MMS04]). The metric space $(\partial\mathcal{C}(S), \delta_z)$ is not locally compact, but the action of \mathcal{MCG} on the space of triples of pairwise distinct points is metrically proper [H08] (this means that for every open subset V of T whose distance to $\partial\mathcal{C}(S) \times \partial\mathcal{C}(S) \times \partial\mathcal{C}(S) - T$ with respect to a product metric is positive, we have $gV \cap V \neq \emptyset$ only for finitely many $g \in \mathcal{MCG}$). Therefore there is a continuous function $h : \mathcal{D} \rightarrow [0, \infty)$ with the following property.

- (1) For all $(x, y) \in \mathcal{D}$ we have $\sum_{g \in \mathcal{MCG}} h(g^{-1}(x, y)) = 1$.
- (2) For every closed set $L \subset \mathcal{D}$ whose distance to $(\rho')^{-1}(0)$ with respect to the product metric is positive, the same is true for the intersection of the support of h with $\cup_{g \in \mathcal{MCG}} gL$.

The function $\rho(x, y) = \sum_{g \in \mathcal{MCG}} h(g^{-1}(x, y))\rho'(g^{-1}(x, y))$ is \mathcal{MCG} -invariant and continuous, and $\rho(x, y) > 0$ if x, y are disjoint. Using the function ρ we argue as in Lemma 3.4 of [MS04] and its proof. We conclude that either the image of the cocycle β stabilizes a pair of distinct points in $\partial\mathcal{C}(S)$, or the image of f is contained in the space of Dirac measures on $\partial\mathcal{C}(S)$, i.e. f is an equivariant map $(G \times X)/\Gamma \times B \rightarrow \partial\mathcal{C}(S)$. On the other hand, if the image of β stabilizes a pair of distinct points in $\partial\mathcal{C}(S)$, then this image is contained in a virtually abelian subgroup of \mathcal{MCG} and therefore the third alternative of the proposition is satisfied.

Now we are left with the case that for almost every $x \in (G \times X)/\Gamma \times B$ the measure $f(x)$ is supported in a single point, i.e. that f is a \mathcal{MCG} -equivariant map $(G \times X)/\Gamma \times B \rightarrow \partial\mathcal{C}(S)$. Let $H < \mathcal{MCG}$ be the subgroup of \mathcal{MCG} generated by the image of the cocycle β and let $\Lambda \subset \partial\mathcal{C}(S)$ be the *limit set* for the action of H on $\mathcal{C}(S)$, i.e. Λ is the set of accumulation points of an H -orbit on $\mathcal{C}(S)$. If this limit set consists of at most two points, then either H is virtually abelian and the third alternative in the proposition holds true, or an H -orbit in $\mathcal{C}(S)$ is bounded.

However, if an H -orbit on $\mathcal{C}(S)$ is bounded then the structure result of [MP89] implies that either H is finite or H stabilizes a multi-curve (compare the proof of Lemma 9.1). Thus for the purpose of the proposition, we may assume that Λ contains more than two points; then Λ is a closed H -invariant subset of $\partial\mathcal{C}(S)$

Let T be the space of triples of pairwise distinct points in Λ . We showed in [H08] that there is a continuous non-trivial cocycle $\rho : T \rightarrow L^2(H)$ for the action of $H < \mathcal{MCG}$ on T . Via β , this cocycle then induces a non-trivial measurable $L^2(H)$ -valued cocycle on $(G \times X)/\Gamma \times B$. Since the action of G on B is doubly ergodic, such a cocycle defines a non-trivial cohomology class for G with coefficients in the separable Hilbert space $L^{[2]}((G \times X)/\Gamma, L^2(H))$ of all measurable maps $(G \times X)/\Gamma \rightarrow L^2(H)$ with the additional property that for each such map α , the function $x \rightarrow \|\alpha(x)\|$ is square integrable on $(G \times X)/\Gamma$.

For every $i \leq n$ define $G'_i = \prod_{j \neq i} G_j$. By assumption, the projection of Γ to $G_i = G'_i \backslash G$ is dense and therefore by Lemma 2.2.13 of [Z84], the action of Γ on G_i by right translation is ergodic. Then by Moore's ergodicity theorem, the action of G'_i on G/Γ is ergodic as well. Since by assumption the action of Γ on (X, μ) is mildly mixing, the results of [SW82] imply that the action of G_i on $((G \times X)/\Gamma, \nu)$ is ergodic. As in [MMS04], we conclude from the results of Burger and Monod [BM99, BM02] that there is some $i \leq n$ and there is an equivariant map $(G \times X)/\Gamma \rightarrow L^2(H)$ for the restriction of β to $G'_i \times (G \times X)/\Gamma$. Since the action of G'_i on $(G \times X)/\Gamma$ is ergodic, the restriction of β to G'_i is equivalent to a cocycle into the stabilizer of a non-zero element of $L^2(H)$ (see Lemma 5.2.11 of [Z84]). Since the action of H on itself by left translation is simply transitive, this stabilizer is a finite subgroup of H .

We now follow the proof of Theorem 1.2 of [MS04] and find a minimal such finite subgroup K of \mathcal{MCG} . The cocycle β is cohomologous to a cocycle β' into the normalizer N of K in \mathcal{MCG} , and the same is true for the cocycle α . Proposition 3.7 of [MS04] then shows that β' is induced from a continuous homomorphism $G \rightarrow N/K$ which factors through G_i . In other words, the first alternative in the statement of the proposition is satisfied. This completes the proof of the proposition. \square

Theorem 3 from the introduction follows from Proposition 9.4 and an analysis of cocycles α which are cohomologous to a cocycle, again denoted by α , with values in the stabilizer H of a multi-curve c . If c has $d \leq 3g - 3 + m$ components, then the group H is a direct product of \mathbb{Z}^d with the mapping class group \mathcal{M}_0 of the (possibly disconnected) surface S_0 which we obtain from $S - c$ by pinching each boundary circle to a point, and where the group \mathbb{Z}^d is generated by the Dehn twists about the components of c . Note that \mathcal{M}_0 contains the direct product of the mapping class groups of the connected components of S_0 as a subgroup of finite index. Let $p : H \rightarrow \mathcal{M}_0$ be the natural projection and let \mathcal{C}_0 be the space of all complete geodesic laminations on S_0 . Then \mathcal{C}_0 is a compact \mathcal{M}_0 -space and hence we can apply our above procedure to the cocycle $p \circ \alpha$, the compact space \mathcal{C}_0 and the group \mathcal{M}_0 . After at most $3g - 3 + k$ steps we conclude that indeed the cocycle α is cohomologous to a cocycle with values in subgroup of \mathcal{MCG} which is a finite extension of a group of the form $\mathbb{Z}^\ell \times N/K$ where $\ell \leq 3g - 3 + k$, where $K < \mathcal{MCG}$

is finite, N is the normalizer of K in \mathcal{MCG} and such that there is a continuous homomorphism $G \rightarrow N/K$. This completes the proof of Theorem 3.

APPENDIX A. SPLITTING, SHIFTING AND CARRYING

In this appendix we establish some technical properties of complete train tracks on S . This appendix only uses results from the literature which are summarized in Section 2, and is independent from the rest of the paper. Its main goal is the proof of Proposition A.6 which is used in an essential way in Section 6. Throughout we use the notations from the introduction and from Section 2.

By the definition given in Section 2, a complete train track on the non-exceptional surface S of finite type is a train track which is maximal, generic and birecurrent. By Lemma 2.3, such a train track τ carries a complete geodesic lamination λ . It can be split at any large branch to a train track σ which carries λ as well, and σ is complete and is carried by τ . Moreover, the choice of a right or left split at e is uniquely determined by λ . On the other hand, τ can not be split to *every* complete train track η which is carried by τ . Namely, to pass from τ to η it may be necessary to do some shifting moves as well (see Chapter 2 of [PH92] for a detailed discussion of this fact). Shifting involves making choices, and these choices are hard to control combinatorially.

We overcome this difficulty by obtaining a better quantitative understanding of the relation between splitting, shifting and carrying of complete train tracks. This is done by putting a train track σ which is carried by a train track τ in a standard position with respect to τ and use this standard position to define a numerical invariant which can be controlled under splitting moves. Part of the material presented here is motivated by the results in Section 2.3 and Section 2.4 of [PH92]. Recall from Section 2 the definition of a foliated neighborhood A of a generic train track τ [PH92].

Definition A.1. Let τ be a generic train track with foliated neighborhood A and collapsing map $F : A \rightarrow \tau$. A generic train track σ is *in general position with respect to A* if the following three conditions are satisfied.

- (1) σ is contained in the interior of A and is transverse to the ties.
- (2) No switch of σ is mapped by F to a switch of τ .
- (3) For every $x \in \tau$ the tie $F^{-1}(x)$ contains at most one switch of σ .

Note that by definition, a train track in general position with respect to a foliated neighborhood A of τ is always carried by τ , with the restriction of the collapsing map $F : A \rightarrow \tau$ as a carrying map. Moreover, every generic train track σ which is carried by τ can be isotoped to a train track in general position with respect to A .

Let σ be a generic train track in general position with respect to the standard neighborhood A of the complete train track τ , with collapsing map $F : A \rightarrow \tau$. Let b be a branch of τ which is incident and large on a switch v of τ . A *cutting arc* for σ and v is an embedded arc $\gamma : [0, d] \rightarrow F^{-1}(b)$ which is transverse to the ties of A , which is disjoint from σ except possibly at its endpoint and such that the length

of $F(\gamma) \subset b$ is maximal among all arcs with these properties. The maximality condition implies that either $F(\gamma[0, d]) = b$, i.e. that γ crosses through the foliated rectangle $F^{-1}(b)$, or that $\gamma(d)$ is a switch of σ contained in the interior of $F^{-1}(b)$, and the component of $S - \sigma$ which contains v has a cusp at $\gamma(d)$.

We use cutting arcs to investigate whether or not σ is carried by a split of τ at some large branch e .

Lemma A.2. *Let σ be in general position with respect to the foliated neighborhood A of τ , with collapsing map $F : A \rightarrow \tau$. Let e be a large branch which is incident on the switches v, v' of τ . Let $\gamma : [0, d] \rightarrow F^{-1}(e), \gamma' : [0, d] \rightarrow F^{-1}(e)$ be cutting arcs for σ and v, v' . If σ is not carried by a split of τ at e then $\gamma(d), \gamma'(d)$ are switches of σ contained in the interior of $F^{-1}(e)$, and there is a trainpath $\rho : [0, m] \rightarrow \sigma \cap F^{-1}(e)$ connecting $\rho(0) = \gamma(d)$ to $\rho(m) = \gamma'(d)$.*

Proof. Let e be a large branch of τ incident on the switches v and v' . Let σ be in general position with respect to the foliated neighborhood A of τ with collapsing map $F : A \rightarrow \tau$. Assume that σ is not carried by any split of τ at e and let $\gamma : [0, d] \rightarrow F^{-1}(e), \gamma' : [0, d] \rightarrow F^{-1}(e)$ be cutting arcs for σ and for v, v' .

We show first that $\gamma(d)$ is a switch of σ contained in the interior of $F^{-1}(e)$. For this we argue by contradiction and we assume that $\gamma(d) \in F^{-1}(v')$. Since σ is in general position with respect to A and v' is contained in the boundary of A , we may assume that $\gamma(d)$ is contained in $F^{-1}(v') - v' - \sigma$. Then $\gamma(d)$ is contained in the boundary of a unique foliated rectangle $R \subset A$ whose image under the collapsing map F is a branch of τ which is incident and small at the switch v' . Extend γ a bit beyond $\gamma(d)$ to an arc $\tilde{\gamma}$ transverse to the ties of A and disjoint from σ whose endpoint is contained in the interior of R . The train track obtained by cutting the foliated neighborhood A open along $\tilde{\gamma}$ as shown in Figure D and by collapsing the ties of the induced foliation of $A - \tilde{\gamma}$ carries σ and is a split of τ at e . This means that σ is carried by a split of τ at e which is a contradiction. As a consequence, $w = \gamma(d)$ and $w' = \gamma'(d)$ are switches of σ contained in the interior of $F^{-1}(e)$ which are cusps for the complementary components of σ containing v, v' .

Let $\nu : [0, \infty) \rightarrow \sigma$ be a trainpath on σ issuing from $\nu(0) = w$ such that $\nu[0, 1/2]$ is the half-branch of σ which is incident on w and large at w . Since σ is transverse to the ties of A , there is a smallest number $t > 0$ such that $F(\nu(t)) = v'$. The arc $\nu[0, t]$ is embedded in $\sigma \cap F^{-1}(e - F(\gamma[0, d]))$. Now if every such embedded arc of class C^1 in $F^{-1}(e) \cap \sigma$ connects w to the same component c of $F^{-1}(v') - v'$ then σ is carried by a split of τ at e (and with the branch of τ determined by c as a winner). Thus there are two embedded arcs ν, ν' of class C^1 in $\sigma \cap F^{-1}(e - F(\gamma[0, d]))$ whose endpoints are contained in the two different components of $F^{-1}(v') - v'$.

Let C be the connected component of $F^{-1}(e) - (\nu \cup \nu')$ which contains the switch v' of τ . The closure of C does not intersect $F^{-1}(v)$, and it contains the switch w' of σ . Every trainpath on σ issuing from w' which intersects $F^{-1}(v)$ necessarily intersects one of the arcs ν or ν' . Since the restriction of the differential dF of F to σ vanishes nowhere, this implies that there is an embedded trainpath $\rho : [0, m] \rightarrow \sigma$ which is entirely contained in $F^{-1}(e)$ and connects $w = \rho(0)$ to $w' = \rho(m)$. This shows the lemma. \square


Next we show that a complete train track σ on S which is carried by a complete train track τ and is not carried by any split of τ can be shifted and isotoped to a train track σ' which is in standard position relative to τ . For this we have to introduce some more terminology.

Let σ be any generic train track on S and let $\rho : [0, n] \rightarrow \sigma$ be an embedded trainpath. For $i \in \{1, \dots, n-1\}$ the point $\rho(i)$ is a switch of σ contained in the interior of $\rho[0, n]$. Since every switch is trivalent, there is a unique branch of τ which is incident on $\rho(i)$ and not contained in $\rho[0, n]$. We call $\rho(i)$ a *right switch* if this branch lies to the right of ρ with respect to the orientation of S and the orientation of ρ ; otherwise $\rho(i)$ is called a *left switch*. Moreover, the switch $\rho(i)$ is called *incoming* if the half-branch $\rho[i-1/2, i]$ is small at $\rho(i)$, and we call the switch $\rho(i)$ *outgoing* otherwise.

Define a *special trainpath* on a generic train track σ to be a trainpath $\rho : [0, 2k-1] \rightarrow \sigma$ of length $2k-1$ for some $k \geq 1$ with the following properties.

- (1) $\rho[0, 2k-1]$ is embedded in σ .
- (2) For each $j \leq k-1$ the branch $\rho[2j, 2j+1]$ is large and the branch $\rho[2j+1, 2j+2]$ is small.
- (3) With respect to the orientation of S and the orientation of ρ , right and left switches in $\rho[1, 2k-2]$ alternate.

The left part of Figure E shows a special trainpath of length 5.



Let again σ be a generic train track and let $\rho : [0, k] \rightarrow \sigma$ be any embedded trainpath on σ . Define a *standard neighborhood* of ρ in σ to be a closed connected neighborhood of $\rho[0, k]$ in the graph σ which does not contain a switch distinct from the switches contained in $\rho[0, k]$. As an example, a standard neighborhood of a single branch e of σ incident on two distinct switches is the union of e with four neighboring half-branches.

Let τ be a complete train track on S and let A be a foliated neighborhood of τ with collapsing map $F : A \rightarrow \tau$. Let $\sigma \subset A$ be a complete train track which is in general position with respect to A . Then the restriction of the collapsing map F to σ is surjective. Thus if we denote for a branch b of τ by $\nu(b, \sigma)$ the minimal cardinality of $F^{-1}(x) \cap \sigma$ where x varies through the points in b then $\nu(b, \sigma) > 0$ for every branch b . Since no switch of σ is mapped by F to a switch of τ , there is a point x contained in the interior of b for which the cardinality of $F^{-1}(x) \cap \sigma$ equals $\nu(b, \sigma)$. Moreover, since any tie of A contains at most one switch of σ by assumption, we may assume that $F^{-1}(x)$ does not contain any switch of σ . Note that if we deform σ with an isotopy of train tracks in general position relative to

A , then the values $\nu(b, \sigma)$ do not change along the isotopy. However, this may not be the case for general deformations of σ transverse to the ties of A .

The following technical lemma shows that a complete train track σ which is carried by a complete train track τ and is not carried by any split of τ can be shifted and isotoped to a train track σ' which is in standard position relative to τ .

Lemma A.3. *Let τ be a complete train track on S and let A be a foliated neighborhood of τ with collapsing map $F : A \rightarrow \tau$. Let σ be a complete train track in general position with respect to A , let e be a large branch of τ and let $V \subset U$ be standard neighborhoods of e in τ so that V is contained in the interior of U . If σ is not carried by a split of τ at e , then σ can be shifted and isotoped to a train track σ' with the following properties.*

- (1) σ' is in general position with respect to A , and we have $\sigma' \cap F^{-1}(\tau - U) = \sigma \cap F^{-1}(\tau - U)$.
- (2) $F^{-1}(V) \cap \sigma'$ is the disjoint union of a standard neighborhood V' of a special trainpath ρ on σ' and a (possibly empty) collection of simple arcs not containing any switch of σ' . The set V' is mapped by F onto V .
- (3) If $\nu(e, \sigma) = 1$ then $F^{-1}(V) \cap \sigma'$ is a standard neighborhood of a single large branch in σ' .

Proof. Let τ be a complete train track with foliated neighborhood A and collapsing map $F : A \rightarrow \tau$. Let $\sigma \subset A$ be a complete train track in general position with respect to A . Then σ is carried by τ with the restriction of F as a carrying map. Let e be a large branch of τ and assume that σ is not carried by any split of τ at e .

Let v, v' be the switches of τ on which the large branch e is incident. The points v, v' are contained in complementary regions of σ . Let $\gamma : [0, d] \rightarrow F^{-1}(e)$ and $\gamma' : [0, d'] \rightarrow F^{-1}(e)$ be cutting arcs for σ and v, v' . By Lemma A.2 the points $w = \gamma(d)$ and $w' = \gamma'(d')$ are switches of σ contained in the interior of $F^{-1}(e)$. Moreover, there is a trainpath $\rho : [0, m] \rightarrow \sigma \cap F^{-1}(e)$ which connects $\rho(0) = \gamma(d)$ to $\rho(m) = \gamma'(d')$. The half-branch $\rho[0, 1/2]$ is large at $\rho(0)$, and $\rho[m - 1/2, m]$ is large at $\rho(m)$. Since σ does not have embedded bigons, ρ is unique. We call ρ the *cutting connector* for σ and e .

Recall the definition of right and left and incoming and outgoing switches along the trainpath ρ . Define the switch $\rho(0)$ to be incoming and the switch $\rho(m)$ to be outgoing. Let $j \in \{1, \dots, m\}$ be the smallest number with the property that $\rho(j)$ is outgoing. We claim that if $\rho(j)$ is right outgoing, then none of the switches $\rho(k)$ for $k > j$ is right incoming. For this we assume to the contrary that for some $k > j$ the switch $\rho(k)$ is right incoming. Now every subarc of a trainpath on σ which is entirely contained in $F^{-1}(e)$ is mapped C^1 -diffeomorphically onto a subarc of e by the map F . Let b_j, b_k be the branches of σ which are incident on $\rho(j), \rho(k)$ and whose interiors are disjoint from ρ . Then there is a subarc of a trainpath on σ which begins at $\rho(k)$, connects $\rho(k)$ to a point in $F^{-1}(F\rho(0))$ and whose initial segment is contained in the branch b_k . Since σ does not have embedded bigons, the interior of this arc is entirely contained in the interior of the component R of $F^{-1}(F(\rho[0, m])) - \rho[0, m]$ to the right of ρ with respect to the orientation of ρ and the orientation of S . Similarly, there is an embedded subarc of a trainpath on σ

which begins at $\rho(j)$, connects $\rho(j)$ to a point in $F^{-1}(F\rho(m))$ and whose initial segment is contained in the branch b_j . The interior of this arc is contained in R as well. Therefore these arcs have to intersect in the interior of R . However, σ does not have embedded bigons and hence this is impossible. Note also that if there is a point $x \in e$ such that $F^{-1}(x) \cap \sigma$ consists of a single point, then $x \in F(\rho[j-1, j])$ and none of the switches $\rho(i)$ for $i \geq j$ is incoming.

We can modify the train track σ with a sequence of shifts and isotopies in $m-1$ steps as follows. Let again $j \leq m$ be the smallest number such that the switch $\rho(j)$ is outgoing. The branches incident on the switches $\rho(1), \dots, \rho(j-1)$ can successively be moved backward along ρ with a sequence of shifts past the switch w to a switch contained in $F^{-1}(e - F\rho[0, m])$ as shown on the right hand side of Figure E. This operation can be done in such a way that the resulting train track $\tilde{\sigma}$ is in general position with respect to A , that $\nu(e, \sigma) = 1$ if and only if $\nu(e, \tilde{\sigma}) = 1$ and that $\tilde{\sigma}$ coincides with σ in $A - F^{-1}(e)$. The length of the cutting connector $\tilde{\rho}$ for $\tilde{\sigma}$ and e equals $m-j+1$, and $\tilde{\rho}[k, k+1] = \rho[k+j-1, k+j]$ for $k \geq 1$. The branch $\tilde{\rho}[0, 1]$ in $\tilde{\sigma}$ which is incident and large at w is a large branch. Moreover, if $\nu(e, \sigma) = 1$ then there is a point $x \in F(\tilde{\rho}[0, 1])$ such that the cardinality of $\tilde{\sigma} \cap F^{-1}(x)$ equals one.

If each of the switches $\tilde{\rho}(i)$ for $i \geq 1$ is outgoing then we can modify $\tilde{\sigma}$ with a sequence of shifts so that the cutting connector for the modified train track and e is a single large branch. Note that this is the case if $\nu(e, \sigma) = 1$. Otherwise assume that $\tilde{\rho}(1)$ is right outgoing and let $k > 1$ be the smallest number such that $\tilde{\rho}(k)$ is left incoming. Then each of the switches $\tilde{\rho}(i)$ for $1 \leq i \leq k-1$ is right outgoing and hence $\tilde{\sigma}$ can be modified with a sequence of shifts in such a way that the path $\tilde{\rho}[1, k]$ corresponds to a single small branch in the cutting connector for the modified train track and e . Inductively, this just means that with a sequence of shifts and an isotopy we can modify σ to a train track σ_1 which is in general position with respect to e , which coincides with σ in $A - F^{-1}(e)$ and such that the cutting connector ρ_1 for σ_1 and e is a special trainpath. If $\nu(e, \sigma) = 1$ then ρ_1 consists of a single large branch, and there is a point z in the interior of this branch such that the cardinality of $F^{-1}(F(z)) \cap \sigma_1$ equals one.

Since σ_1 is mapped by F onto τ , with respect to the orientation of S and the orientation of a cutting arc γ for σ_1 and v , for each point $u \in F(\gamma)$ there is a unique point $r(u) \in \sigma_1 \cap F^{-1}(u)$ (or a unique point $\ell(u) \in \sigma_1 \cap F^{-1}(u)$) which is to the right (or to the left) of γ and which is closest to $\gamma \cap F^{-1}(u)$ with this property as measured along the arc $F^{-1}(u)$. Then $u \rightarrow r(u)$ and $u \rightarrow \ell(u)$ are embedded arcs of class C^1 in σ_1 which are disjoint from γ except at their endpoints.

Let $V \subset U$ be standard neighborhoods of e so that the closure of V is contained in the interior of U . We can isotope σ_1 by sliding the switches contained in the interior of the arcs r and ℓ backward along r, ℓ outside of $F^{-1}(V)$. This isotopy can be chosen to be supported in $F^{-1}(U) - F^{-1}(e - F(\gamma[0, d]))$ and to be transverse to the ties of A . Moreover, we may assume that the image of σ_1 under this isotopy is in general position with respect to A . Since for every $u \in e - F(\gamma[0, d])$ the isotopy does not change the cardinality of $F^{-1}(u) \cap \sigma_1$, if $\nu(e, \sigma_1) = 1$ then the same is true for the image of σ_1 under this isotopy.

As a consequence, after modifying σ_1 with an isotopy supported in $F^{-1}(U)$ we obtain a train track σ_2 whose intersection with $F^{-1}(V)$ contains a connected component which is a standard neighborhood V' of a special trainpath ρ_2 , and ρ_2 is the cutting connector for σ_2 and e .

Now if there is a connected component $\beta \neq V'$ of $F^{-1}(V) \cap \sigma_2$, then $\nu(e, \sigma) > 1$ and $\beta \cap F^{-1}(e)$ is contained in a connected component of $F^{-1}(e) - \rho_2$ whose boundary intersects both $F^{-1}(v)$ and $F^{-1}(v')$. It is then immediate that we can deform σ_2 with an isotopy supported in $F^{-1}(U)$ to a train track σ' which is in general position with respect to A and has the property that every connected component of $\sigma' \cap F^{-1}(e)$ which does not coincide with the standard neighborhood V' of ρ_2 is a single arc not containing any switches. In other words, σ' satisfies the requirements in the lemma. \square

The following observation is a partial converse of Lemma A.3. It is used to determine whether or not a given complete train track σ on S which is carried by a complete train track τ is carried by a split of τ at a given large branch e .

Lemma A.4. *Let A be a foliated neighborhood of a complete train track τ with collapsing map $F : A \rightarrow \tau$. Let σ be a complete train track in general position with respect to A . Let e be a large branch of τ ; if $\nu(e, \sigma) = 1$ then σ is not carried by any split of τ at e .*

Proof. Let τ be a complete train track on S with foliated neighborhood A and collapsing map $F : A \rightarrow \tau$. Let σ be a complete train track in general position with respect to A . Let e be a large branch of τ and assume that $\nu(e, \sigma) = 1$. Then there is a point x in the interior of e with the property that $F^{-1}(x) \cap \sigma$ is a single point z contained in the interior of a branch of σ .

Let v be a switch on which the branch e is incident and let $\gamma : [0, d] \rightarrow F^{-1}(e)$ be a cutting arc for σ and v . Since σ is complete, the restriction of the map F to σ is surjective. Thus there is an embedded arc $\nu : [0, 1] \rightarrow \sigma \cap F^{-1}(e)$ (or $\nu' : [0, 1] \rightarrow F^{-1}(e)$) which is mapped by F onto e and which begins at a point $\nu(0) \in F^{-1}(v)$ (or $\nu'(0) \in F^{-1}(v)$) to the right (or to the left) of v with respect to the orientation of γ and the orientation of S . Since $F^{-1}(x) \cap \sigma$ consists of the single point z , the arcs ν, ν' have to intersect. Thus $\gamma(d)$ is a switch of σ contained in the interior of e , and $x \in e - F(\gamma[0, d])$. Similarly, a cutting arc $\gamma' : [0, d] \rightarrow F^{-1}(e)$ for σ and v' ends at a switch of σ contained in the interior of $F^{-1}(e)$ and such that $x \in e - F(\gamma'[0, d])$.

Now every embedded arc $\alpha : [0, 1] \rightarrow \sigma$ of class C^1 with $\alpha(0) = \gamma(d)$ and $F\alpha[0, 1] = e - F\gamma[0, d]$ passes through z , and the same holds true for an embedded arc $\alpha' : [0, 1] \rightarrow \sigma$ of class C^1 with $\alpha'(0) = \gamma'(d)$ and $F\alpha'[0, 1] = e - \gamma'[0, d]$. But this implies that there is an embedded trainpath in σ which is contained in $F^{-1}(e)$ and connects $\gamma(d)$ to $\gamma'(d)$. Then σ is not carried by a split of τ at e . \square

Call two complete train tracks τ, σ on S *shift equivalent* if up to isotopy, σ can be obtained from τ by a sequence of shifts. This clearly defines an equivalence relation on $\mathcal{V}(TT)$. The number of points in each equivalence class is bounded from above

by a universal constant only depending on the topological type of the surface S . The following lemma characterizes complete train tracks σ on S which are shift equivalent to a complete train track τ .

Lemma A.5. *Let τ be a complete train track on S and let A be a foliated neighborhood of τ with collapsing map $F : A \rightarrow \tau$. Let σ be a complete train track in general position with respect to A . If $\nu(e, \sigma) = 1$ for every large branch e of τ , then σ is shift equivalent to τ .*

Proof. Let τ be a complete train track on S with foliated neighborhood A and collapsing map $F : A \rightarrow \tau$. Let σ be a complete train track in general position with respect to the A . Assume that $\nu(e, \sigma) = 1$ for every large branch of τ ; we have to show that σ is shift equivalent to τ .

By Lemma A.4, σ is not carried by any split of τ . Lemma A.3 then shows that after modifying σ by a sequence of shifts and an isotopy we may assume that for every large branch e of τ there is a standard neighborhood U of e such that $F^{-1}(U) \cap \sigma$ is a standard neighborhood of a large branch in σ .

Following [PH92], define a *large one-way trainpath* on τ to be a trainpath $\rho : [0, m] \rightarrow \tau$ such that $\rho[0, 1/2]$ and $\rho[m-1/2, m]$ are large half-branches and that for every $i < m$ the half-branch $\rho[i-1/2, i]$ is small. Every switch v in τ is the starting point of a unique large one-way trainpath, and this path is embedded. Define the *height* $h(v)$ of a switch v of τ to be the length of the large one-way trainpath starting at v . The switches of height 1 are precisely the switches on which a large branch is incident.

Write $\sigma_1 = \sigma$. We construct inductively for each $m \geq 1$ a complete train track σ_m which can be connected to σ_{m-1} by a sequence of shifts and an isotopy and with the following properties.

- (1) σ_m is in general position with respect to A .
- (2) Every large one-way trainpath ρ on τ of length at most m admits a standard neighborhood U with the property that $F^{-1}(U)$ is a standard neighborhood V of a large one-way trainpath $b(\rho)$ in σ_m which is isotopic to ρ .

Assume that for some $m-1 \geq 1$ we constructed the train track σ_{m-1} . Let v be a switch in τ of height m and let $\rho : [0, m] \rightarrow \tau$ be the large one-way trainpath issuing from v . Let $\gamma : [0, d] \rightarrow F^{-1}(\rho[0, 1])$ be a cutting arc for σ and v . We claim that $\gamma(d)$ is a switch of σ contained in the interior of $F^{-1}(\rho[0, 1])$. For this assume to the contrary that this is not the case. Then $\gamma(d)$ is contained in the singular tie $F^{-1}(\rho(1))$. Since by assumption singular ties in A do not contain switches of σ_{m-1} , the point $\gamma(d)$ is not a switch of σ_{m-1} . Thus by possibly modifying γ with a small isotopy we may assume that $\gamma(d) \notin \sigma_{m-1}$.

Since the height of v equals $m \geq 2$, $\rho[0, 1]$ is a mixed branch in τ and therefore the switch $\rho(1)$ divides the tie $F^{-1}(\rho(1))$ into two closed connected subarcs with intersection $\{\rho(1)\}$, where one of the subarcs is the intersection of $F^{-1}(\rho(1))$ with the boundary of the rectangle $F^{-1}(\rho[0, 1])$; we denote this arc by c . Since σ_{m-1} is contained in the interior of the foliated neighborhood A of τ and $\gamma(d) \in c$, by

possibly modifying γ with another isotopy we may assume that $\gamma(d)$ is contained in the interior of c .

Now F maps σ_{m-1} onto τ and hence with respect to the orientation of S and the orientation of γ there is a point in $\sigma_{m-1} \cap F^{-1}(\rho(0))$ to the right of $\rho(0) = \gamma(0)$, and there is a point in $\sigma_{m-1} \cap F^{-1}(\rho(0))$ to the left of $\gamma(0)$. Since γ is disjoint from σ_{m-1} this implies that there is a point $y \in c \cap \sigma_{m-1}$ to the right of $\gamma(d)$, and there is a point $z \neq y \in c \cap \sigma_{m-1}$ to the left of $\gamma(d)$. However, by the induction hypothesis, the intersection of σ_{m-1} with c consists of a single point and hence we arrive at a contradiction.

As a consequence, the cutting arc γ terminates at a switch w of σ_{m-1} which is contained in the interior of $F^{-1}(\rho[0, 1])$. As in the proof of Lemma A.3 we can modify σ_{m-1} by an isotopy in such a way that for the resulting train track, again denoted by σ_{m-1} , there are two small half-branches which are incident on w and are mapped by F onto $F(\gamma[0, d])$.

Let $\xi : [0, n] \rightarrow \sigma_{m-1}$ be the large one-way trainpath issuing from the switch w . Since σ_{m-1} is transverse to the ties of A and its intersection with the boundary component c of the rectangle $F^{-1}(\rho[0, 1])$ consists of a unique point, the arc $\xi[0, n]$ can not be contained in $F^{-1}(\rho[0, 1])$. Then there is a minimal number $t > 0$ such that $\xi(t) \in c$. If $\ell \geq 0$ is the nonnegative integer such that $t \in [\ell, \ell + 1)$, then for $0 < i \leq \ell$ the half-branch $\xi[i - 1/2, i]$ is small. Thus as in the proof of Lemma A.3 we can modify σ_{m-1} by a sequence of shifts and an isotopy in such a way that in the modified train track σ' the arc $\xi[0, \ell + 1]$ consists of a unique branch. Moreover, there is a standard neighborhood W of the trainpath ρ in τ with the property that $F^{-1}(\rho) \cap \sigma'$ is a standard neighborhood of a large trainpath in σ' which is isotopic to W .

We can now repeat this construction for all large one-way trainpaths on τ of length m . The resulting train track σ_m satisfies properties 1) and 2) above for m .

If m_0 denotes the maximal height of any large trainpath on τ , then after m_0 steps we obtain a train track σ_{m_0} which can be connected to σ by a sequence of shifts and an isotopy, is contained in A and transverse to the ties. Let $B \subset \tau$ be a closed subset which we obtain from τ by removing from each small branch b of τ an open subset whose closure is contained in the interior of b . By construction, the train track σ_{m_0} contains a closed subset D which is isotopic to B and coincides with $F^{-1}(B) \cap \sigma_{m_0}$. Since σ_{m_0} and τ are both complete, they have the same number of switches. Thus every switch of σ_{m_0} is contained in D and therefore σ_{m_0} is isotopic to τ . This finishes the proof of the lemma. \square

Now we are ready to show the main result of this appendix which relates splitting to carrying in a quantitative way. For this define a *splitting, shifting and collapsing sequence* to be a sequence $\{\alpha_i\}_{0 \leq i \leq m}$ of complete train tracks such that for every i the train track $\alpha(i + 1)$ can be obtained from $\alpha(i)$ by a single split, a single shift or a single collapse. Recall from the introduction or from Subsection 2.1 the definition of a complete geodesic lamination.

Proposition A.6. *There is a number $\chi > 0$ with the following property. Let $\sigma \prec \tau$ be complete train tracks on S and let λ be a complete geodesic lamination carried by σ . Then τ is splittable to a complete train track τ' which carries λ and which can be obtained from σ by a splitting, shifting and collapsing sequence of length at most χ consisting of complete train tracks which carry λ .*

Proof. The number of complete train tracks which can be obtained from a fixed complete train track by a sequence of shifts is uniformly bounded. Hence it is enough to show that whenever $\sigma \prec \tau$ and the complete geodesic lamination λ is carried by σ , then τ is splittable to a train track σ' which carries λ and can be obtained from σ by a uniformly bounded number of modifications where each modification consists in a number of shifts and a single split or a single collapse.

Thus let τ be a complete train track and let A be a foliated neighborhood of τ with collapsing map $F : A \rightarrow \tau$. Let $\sigma \subset A$ be a complete train track in general position with respect to A . Let λ be a complete geodesic lamination which is carried by σ . We use the collapsing map F to define the quantities $\nu(b, \sigma)$ for a branch b of τ as in the beginning of this section. If e is a large branch of τ with $\nu(e, \sigma) \geq 2$ and if σ is carried by a split $\tilde{\tau}$ of τ at e then for every standard neighborhood V of e in τ there is a foliated neighborhood \tilde{A} of $\tilde{\tau}$ which coincides with A outside of $F^{-1}(V)$ and such that σ is in general position with respect to \tilde{A} . Thus if $b \neq e$ is any branch different from e with $\nu(b, \sigma) = 1$ then we may assume that $\nu(\tilde{b}, \sigma) = 1$ for the branch \tilde{b} of $\tilde{\tau}$ corresponding to b under the natural identification of the branches of τ with the branches of $\tilde{\tau}$.

Let $m \geq 0$ be the number of branches b of τ with $\nu(b, \sigma) = 1$. Let τ_0 be a train track which can be obtained from τ by a splitting sequence, which carries σ and such that no split of τ_0 carries σ . By the above observation, there is a foliated neighborhood A_0 of τ_0 with collapsing map $F_0 : A_0 \rightarrow \tau_0$ such that σ is in general position with respect to A_0 and that the number m_0 of branches b_0 of τ_0 with $\nu(b_0, \sigma) = 1$ is not smaller than m .

By Lemma A.3, σ can be modified with a sequence of shifts and an isotopy to a train track σ_0 which is in general position with respect to A_0 and such that the following properties hold.

- (1) Every large branch e of τ_0 admits a standard neighborhood $V(e)$ in τ_0 such that $F_0^{-1}(V(e)) \cap \sigma_0$ is the disjoint union of a standard neighborhood V_0 of a special trainpath ρ on σ_0 and a (possibly empty) collection of simple arcs not containing any switch of σ_0 . The set V_0 is mapped by F_0 onto $V(e)$.
- (2) If $b \in \tau_0$ is any branch with $\nu(b, \sigma) = 1$ then $\nu(b, \sigma_0) = 1$.

If $\nu(e, \sigma_0) = 1$ for every large branch e of τ_0 then by Lemma A.5, σ_0 is shift equivalent to τ_0 and we are done. Otherwise we choose a large branch e of τ_0 with $\nu(e, \sigma_0) \geq 2$. We modify τ_0, σ_0 to train tracks τ_1, σ_1 with the following properties.

- (a) τ_1 can be obtained from τ_0 by at most one split at e .
- (b) σ_1 carries λ and can be obtained from σ_0 by a sequence of collapses, shifts and splits of uniformly bounded length.

- (c) There is a foliated neighborhood A_1 of τ_1 such that σ_1 is in general position with respect to A_1 , is transverse to the ties and such that the number of branches b of τ_1 with $\nu(b, \sigma_1) = 1$ is not smaller than $m_0 + 1$.

Namely, using the above notations, let $2\ell - 1 \geq 1$ be the length of the special trainpath $\rho \subset F_0^{-1}(V(e)) \cap \sigma_0$. Note that $\ell \geq 1$ is bounded from above by a constant only depending on the topological type of S . If $\ell = 1$ then ρ consists of a single large branch e' in σ_0 , and $F_0^{-1}(V(e)) \cap \sigma_0$ is the disjoint union of a standard neighborhood of the large branch e' and a collection of simple arcs. There is a unique choice of a right or left split of σ_0 at e' such that the split track σ_1 carries λ . If this split is a right (or left) split, then σ_1 is carried by the train track τ_1 which can be obtained from τ_1 by a right (or left) split at e . If \tilde{e} is the diagonal branch of the split in τ_1 , then for a suitable choice of a foliated neighborhood of τ_1 and a suitable collapsing map we have $\nu(\tilde{e}, \sigma_1) = 1$. Moreover, via the natural identification of the branches of τ_0 with the branches of τ_1 we have $\nu(\tilde{b}, \sigma_1) \leq \nu(b, \sigma_0)$ for every branch b in τ_0 and the corresponding branch \tilde{b} in τ_1 . In particular, the number of branches b of τ_1 with $\nu(b, \sigma_1) = 1$ is at least $m_0 + 1$ and the train tracks $\sigma_1 \prec \tau_1$ satisfy the properties (a)-(c) above.

If the length $2\ell - 1$ of the special trainpath ρ on σ_0 is at least 3 then $\rho[1, 2]$ is a small branch which can be collapsed as shown in Figure E. The train track σ'_0 which is obtained from σ_0 by this collapse is carried by τ_0 . In particular, this train track is transversely recurrent and hence complete. Moreover, for a suitable choice of a carrying map we have $\nu(b, \sigma'_0) \leq \nu(b, \sigma_0)$ for every branch b of τ_0 .

The train track σ'_0 is not carried by any split of τ_0 . Thus by Lemma A.3 and as shown in Figure E, σ'_0 can be shifted and isotoped to a train track η which satisfies $\nu(b, \eta) = 1$ if and only if $\nu(b, \sigma'_0) = 1$ for every branch b of τ_0 and such that there is a standard neighborhood V of e with the property that $F_0^{-1}(V) \cap \eta$ is the disjoint union of a special trainpath of length $2\ell - 3$ and a collection of simple arcs. Repeat this procedure with $\eta \prec \tau_0$. After $\ell - 1$ modifications of σ_0 consisting each of a single collapse and a sequence of shifts and isotopies we obtain a train track η' which is in general position with respect to the foliated neighborhood A_0 and such that the preimage under the collapsing map F_0 of a standard neighborhood $V(e)$ of e intersects η' in a single large branch e' and a collection of simple arcs. Moreover, we have $\nu(b, \eta') \leq \nu(b, \sigma_0)$ for every branch b of τ_0 . If $\nu(e, \eta') = 1$ then the number of branches b of τ_0 with $\nu(b, \eta') = 1$ is strictly bigger than the number of branches of τ_0 with $\nu(b, \sigma_0) = 1$ and the pair $\sigma_1 = \eta' \prec \tau_1 = \tau_0$ satisfies the above properties. Otherwise modify η' and τ_0 with a single split at e as above to obtain a pair of train tracks which fulfill the requirements (a)-(c) of the inductive construction.

Reapply the above procedure with the train tracks $\sigma_1 \prec \tau_1$. After at most p steps where p is the number of branches of a complete train track on S we obtain a pair of train tracks $\sigma' \prec \tau'$ with the following properties.

- (1) τ' carries λ and is obtained from τ by a splitting sequence.
- (2) σ' is in general position with respect to a foliated neighborhood A' of τ' .
- (3) $\nu(e, \sigma') = 1$ for every large branch e of τ' .

- (4) σ' carries λ and can be obtained from σ by a splitting, shifting and collapsing sequence of uniformly bounded length consisting of train tracks which carry λ .

By Lemma A.5, σ' and τ' are shift equivalent. This completes the proof of the proposition. \square

Acknowledgement: Part of this work was carried out during a visit of the MSRI in Berkeley. I thank the Institute for its hospitality. I am also grateful to Curt McMullen for useful comments.

REFERENCES

- [A96] S. Adams, *Reduction of cocycles with hyperbolic target*, Erg. Th. & Dyn. Syst. 16 (1996), 1111–1145.
- [AR00] C. Anantharaman-Delaroche, J. Renault, *Amenable groupoids*, Monographie 36 de l'Enseignement Math., Genève 2000.
- [BCH94] P. Baum, A. Connes, N. Higson, *Classifying space for proper G -actions and K -theory of group C^* -algebras*, Contemp. Math. 167 (1994), 241–291.
- [Bo97] F. Bonahon, *Geodesic laminations on surfaces*, Contemp. Math. 269 (1997), 1–37.
- [Bw08] B. Bowditch, *Tight geodesics in the curve complex*, Invent. Math. 171 (2008), 281–300.
- [BH99] M. Bridson, A. Haefliger, *Metric spaces of non-positive curvature*, Springer Grundlehren 319, Springer, Berlin 1999.
- [BM99] M. Burger, N. Monod, *Bounded cohomology of lattices in higher rank Lie groups*, J. Eur. Math. Soc. 1 (1999), 199–235.
- [BM02] M. Burger, N. Monod, *Continuous bounded cohomology and applications to rigidity theory*, Geom. Funct. Anal. 12 (2002), 219–280.
- [B92] P. Buser, *Geometry and spectra of compact Riemann surfaces*, Birkhäuser, Boston 1992.
- [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.
- [CB88] A. Casson with S. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, Cambridge University Press, Cambridge 1988.
- [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.
- [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.
- [H08] U. Hamenstädt, *Bounded cohomology and isometry groups of hyperbolic spaces*, J. Eur. Math. Soc. 10 (2008), 315–349.
- [H00] N. Higson, *Bivariant K -theory and the Novikov conjecture*, Geom. Funct. Anal. 10 (2000), 563–581.
- [HR00] N. Higson, J. Roe, *Amenable group actions and the Novikov conjecture*, J. reine angew. Math. 519 (2000), 143–153.
- [HM79] J. Hubbard, H. Masur, *Quadratic differentials and foliations*, Acta Math. 142 (1979), 221–274.
- [I02] N. V. Ivanov, *Mapping class groups*, Chapter 12 in Handbook of Geometric Topology (Editors R.J. Daverman and R.B. Sher), Elsevier Science (2002), 523–633.
- [Ka03] V. Kaimanovich, *Double ergodicity of the Poisson boundary and applications to bounded cohomology*, Geom. Funct. Anal. 13 (2003), 852–861.
- [Ka04] V. Kaimanovich, *Boundary amenability of hyperbolic spaces*, Contemp. Math 347 (2004), 83–114.

- [KM96] V. Kaimanovich, H. Masur, *The Poisson boundary of the mapping class group*, Invent. Math. 125 (1996), 221–264.
- [KMS86] S. Kerckhoff, H. Masur, J. Smillie, *Ergodicity of billiard flows and quadratic differentials*, Ann. Math. 124 (1986), 293–311.
- [Ki05] Y. Kida, *The mapping class group from the viewpoint of measure equivalence theory*, arXiv:math.GR/0512230.
- [Kl99] E. Klarreich, *The boundary at infinity of the curve complex and the relative Teichmüller space*, unpublished manuscript, Ann Arbor 1999.
- [M82] H. Masur, *Interval exchange transformations and measured foliations*, Ann. Math. 115 (1982), 169–201.
- [MM99] H. Masur, Y. Minsky, *Geometry of the complex of curves I: Hyperbolicity*, Invent. Math. 138 (1999), 103–149.
- [MP89] J. McCarthy, A. Papadopoulos, *Dynamics on Thurston’s sphere of projective measured foliations*, Comm. Math. Helv. 64 (1989), 133–166.
- [MMS04] I. Mineyev, N. Monod, Y. Shalom, *Ideal bicombings for hyperbolic groups and applications*, Topology 43 (2004), 1319–1344.
- [MV03] G. Mislin, A. Valette, *Proper group actions and the Baum-Connes conjecture*, Birkhäuser, Basel 2003.
- [MS04] N. Monod, Y. Shalom, *Cocycle superrigidity and bounded cohomology for negatively curved spaces*, J. Diff. Geom. 67 (2004), 395–456.
- [M03] L. Mosher, *Train track expansions of measured foliations*, unpublished manuscript, 2003.
- [PH92] R. Penner with J. Harer, *Combinatorics of train tracks*, Ann. Math. Studies 125, Princeton University Press, Princeton 1992.
- [SW82] K. Schmidt, P. Walters, *Mildly mixing actions of locally compact groups*, Proc. London Math. Soc. 45 (1982), 506–518.
- [T79] W. Thurston, *Three-dimensional geometry and topology*, unpublished manuscript, 1979.
- [V86] W. Veech, *The Teichmüller geodesic flow*, Ann. Math. 124 (1986), 441–530.
- [Y98] G. Yu, *The coarse Baum-Connes conjecture for groups with finite asymptotic dimension*, Ann. Math. 147 (1998), 325–355.
- [Y00] G. Yu, *The coarse Baum-Connes conjecture for spaces which admit a uniform embedding into Hilbert space*, Invent. Math. 139 (2000), 201–240.
- [ZB04] X. Zhu, F. Bonahon, *The metric space of geodesic laminations on a surface I*, Geom. Top. 8 (2004), 539–564.
- [Z84] R. Zimmer, *Ergodic theory and semisimple groups*, Birkhäuser, Boston 1984.

MATHEMATISCHES INSTITUT DER UNIVERSITÄT BONN
 BERLINGSTRASSE 1, 53115 BONN, GERMANY
 e-mail: ursula@math.uni-bonn.de