GEOMETRY OF THE MAPPING CLASS GROUPS III: QUASI-ISOMETRIC RIGIDITY

URSULA HAMENSTÄDT

ABSTRACT. Let S be an oriented surface of genus $g \geq 0$ with $m \geq 0$ punctures and $3g - 3 + m \geq 2$. We show that for every finitely generated group Γ which is quasi-isometric to the mapping class group $\mathcal{M}(S)$ of S there is a finite index subgroup Γ' of Γ and a homomorphism $\rho : \Gamma' \to \mathcal{M}(S)$ with finite kernel and finite index image. We also give a new proof of the following result of Behrstock and Minsky: The geometric rank of $\mathcal{M}(S)$ as well as the homological dimension of the asymptotic cone of $\mathcal{M}(S)$ equal $3g - 3 + m$.

CONTENTS

1. Introduction

Let S be an oriented surface of finite type, i.e. S is a closed surface of genus $g \geq 0$ from which $m \geq 0$ points, so-called *punctures*, have been deleted. We assume that $3g - 3 + m \geq 2$, i.e. that S is not a sphere with at most 4 punctures or a torus with at most 1 puncture. We then call the surface S *non-exceptional*. The *mapping class group* $\mathcal{M}(S)$ of all isotopy classes of orientation preserving self-homeomorphisms of S is finitely presented [I02], indeed it acts as a group of automorphisms on a contractible cell complex with finite cell stabilizers and compact quotient. In particular, it is finitely generated. We refer to the survey of Ivanov [I02] for more about the mapping class group and for references.

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

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

The map F is called an *L-quasi-isometry* if moreover for every $y \in Y$ there is some $x \in X$ with $d(Fx, y) \leq L$. The spaces $(X, d), (Y, d)$ are then called *quasi-isometric*.

Every finitely generated group G admits a natural family of metrics which are pairwise quasi-isometric. Namely, choose a finite symmetric set $\mathcal{G}$ of generators for G. Then every element $g \in G$ can be represented as a word in the alphabet $\mathcal{G}$. The minimal length $|g|$ of such a word defines the *word norm* of g. This word norm induces a metric on G which is invariant under left translation by defining $d(g, h) = |g^{-1}h|$. The word norm $|\,|'$ defined by a different set of generators is equivalent to $|\,|$ and hence the induced metrics d, d' are quasi-isometric. In particular, we can talk about quasi-isometric finitely generated groups. Note that a finitely generated group is quasi-isometric to each of its finite index subgroups and quasi-isometric to its image under a homomorphism with finite kernel. The main purpose of this note is to show.

Theorem A: *Let Γ be a finitely generated group which is quasi-isometric to $\mathcal{M}(S)$. Then there is a finite index subgroup Γ' of Γ and a homomorphism $\rho : \Gamma' \to \mathcal{M}(S)$ with finite kernel and finite index image.*

For surfaces with precisely one puncture (i.e. in the case $m = 1$), our theorem was earlier shown by Mosher and Whyte (see [M03b] for more details). Kida [K06] recently showed an analogous rigidity result in the context of *measure equivalence*. Namely, call two countable groups Γ, Λ *measure equivalent* if Γ, Λ admit commuting measure preserving actions on a standard Borel space X equipped with a Radon measure μ and with finite measure fundamental domains. Motivated by deep results of Zimmer and Furman, Kida showed that for a countable group Γ which is measure equivalent to the mapping class group the conclusion of our theorem holds true. Note however that measure equivalence for finitely generated groups is neither implied by nor implies quasi-isometry.

A choice of a word norm for the mapping class group and of a non-principal ultrafilter on $\mathbb{N}$ determines an *asymptotic cone* of $\mathcal{M}(S)$. The homological dimension of this cone, i.e. the maximal number $n \geq 0$ such that there are two open

subsets $V \subset U$ with $H_n(U, U - V) \neq 0$, is independent of the choices. We show the following version of a result of Behrstock and Minsky [BM05].

Theorem B [BM05]**:** *The homological dimension of an asymptotic cone of $\mathcal{M}(S)$ equals $3g - 3 + m$.*

The *geometric rank* of a metric space X is defined to be the maximal number $k \geq 0$ such that there is a quasi-isometric embedding $\mathbb{R}^k \rightarrow X$; it is not bigger than the homological dimension of an asymptotic cone for $\mathcal{M}(S)$. Farb, Lubotzky and Minsky [FLM01] showed that the geometric rank of $\mathcal{M}(S)$ is at least $3g - 3 + m$; thus as an immediate corollary of our theorem we obtain [BM05].

Corollary [BM05]**:** *The geometric rank of $\mathcal{M}(S)$ equals $3g - 3 + m$.*

The organization of this paper is as follows. In Section 2 we summarize those of the properties of the *train track complex* $\mathcal{TT}$ introduced in [H06a] which are needed for our purpose. Section 3 discusses train tracks which hit efficiently. Building on results from [H06b], we analyze in Section 4 in more detail the distance in the train track complex. This is used in Section 5 to single out a collection of infinite subsets of $\mathcal{M}(S)$ whose asymptotic cones are homeomorphic to euclidean cones of dimension at most $3g - 3 + m$. In Section 6 we establish a fairly precise description of the asymptotic cone of $\mathcal{M}(S)$ which leads to the proof of Theorem B. Section 7 then contains the proof of Theorem A.

2. The complex of train tracks

In this section we summarize some results and constructions from [PH92, H06a, H06b] which will be used throughout the paper (compare also [M03a]).

Let S be an oriented surface of genus $g \geq 0$ with $m \geq 0$ punctures and where $3g - 3 + m \geq 2$. A *train track* on S is an embedded 1-complex $\tau \subset S$ whose edges (called *branches*) are smooth arcs with well-defined tangent vectors at the endpoints. At any vertex (called a *switch*) the incident edges are mutually tangent. Through each switch there is a path of class C^1 which is embedded in τ and contains the switch in its interior. In particular, the branches which are incident on a fixed switch are divided into "incoming" and "outgoing" branches according to their 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. We always identify train tracks which are isotopic. A train track is called *maximal* if its complementary components are all trigons or once punctured monogons. A train track τ is called *large* if its complementary components are all topological discs and once punctured topological discs.

A *trainpath* on a train track τ is a C^1-immersion $\rho : [m, n] \to \tau \subset S$ which maps each interval $[k, k+1]$ ($m \le k \le n-1$) 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 *subtrack* of a train track τ is a subset σ of τ which itself is a train track. Thus every switch of σ is also a switch of τ, and every branch of σ is an embedded trainpath of τ. We write $\sigma < \tau$ if σ is a subtrack of τ.

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.

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 sum of the weights over all incoming branches at s is required to coincide with the sum of the weights over all outgoing branches at s. The train track is called *recurrent* if it admits a transverse measure which is positive on every branch. We call such a transverse measure μ *positive*, and we write $\mu > 0$. If μ is any transverse measure on a train track τ then the subset of τ consisting of all branches with positive μ-mass is a recurrent subtrack of τ. A train track τ is called *birecurrent* if τ is recurrent and transversely recurrent.

A *geodesic lamination* for a complete hyperbolic structure on S of finite volume is a *compact* subset of S which is foliated into simple geodesics. A geodesic lamination λ is called *minimal* if each of its half-leaves is dense in λ. Thus 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 consists of a disjoint union of finitely many minimal components 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 [CEG87, O96].

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 its complementary regions are all ideal triangles or once punctured discs with one cusp at the boundary. The space of all geodesic laminations on S equipped with the *Hausdorff topology* is a compact metrizable space. A geodesic lamination λ is called *complete* if λ is maximal and can be approximated in the Hausdorff topology by simple closed geodesics. The space $\mathcal{CL}$ of all complete geodesic laminations equipped with the Hausdorff topology is compact. Every geodesic lamination λ which is a disjoint union of finitely many minimal components is a *sublamination* of a complete geodesic lamination, i.e. there is a complete geodesic lamination which contains λ as a closed subset [H06a].

A train track or a geodesic lamination σ is *carried* by a transversely recurrent train track τ if there is a map $F : S \to S$ of class C^1 which is isotopic to the identity and maps σ into τ in such a way that the restriction of the differential of F to the

tangent space of σ vanishes nowhere; note that this makes sense since a train track has a tangent line everywhere. We call the restriction of F to σ a *carrying map* for σ. Write $\sigma \prec \tau$ if the train track or the geodesic lamination σ is carried by the train track τ. If σ is a train track which is carried by a train track τ then every geodesic lamination λ which is carried by σ is also carried by τ.

A train track τ is called *complete* if it is generic and transversely recurrent and if it carries a complete geodesic lamination. A complete train track is maximal and birecurrent. The space of all complete geodesic laminations which are carried by a fixed complete train track τ is open and closed in $\mathcal{CL}$. In particular, the space $\mathcal{CL}$ is totally disconnected [H06a].

A half-branch $\hat{b}$ in a generic train track τ incident on a switch v of τ is called *large* if every trainpath containing v in its interior passes through $\hat{b}$. A half-branch which is not large is called *small*. A branch b in a generic train track τ is called *large* if each of its two half-branches is large; in this case b is necessarily incident on two distinct switches, 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).

There are two simple ways to modify a train track τ to another train track. First, we can *shift* τ along a mixed branch to a train track τ' as shown in Figure A below. If τ is complete then the same is true for τ'. Moreover, a train track or a 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'.
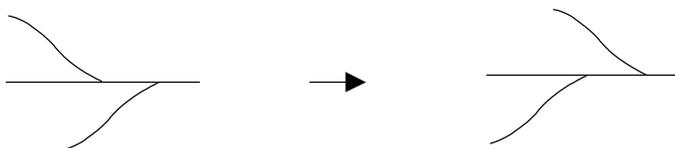
<p align="center">Figure A</p>

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 B. A right split at e is uniquely determined by the orientation of S and does not depend on the orientation of e. Using the labels in the figure, in the case of a right split we call the branches a and c *winners* of the split, and the branches b, d are *losers* of the split. If we perform a left split, then the branches b, d are winners of the split, and the branches a, c are losers of the split. The split τ' of a train track τ is carried by τ, and there is a natural choice of a carrying map which maps the switches of τ' to the switches of τ. The image of a branch of τ' is then a trainpath on τ whose length either equals one or two. There is a natural bijection 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 maximal transversely recurrent generic train track is maximal, transversely recurrent and generic. 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 such a split a λ-*split*. 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 train track is not recurrent any more (compare p.120 in [PH92]). The reverse of a split is called a *collapse*. Define moreover a *collision* of a train track τ at a large branch e to be a (right or left) split of τ at e followed by the removal of the diagonal of the split. The train track obtained from τ by a collision at e is carried by both train tracks obtained from τ by a split at e. The number of its branches equals the number of branches of τ minus one. If τ is complete then the collision of τ at e is recurrent if and only if both train tracks obtained from τ by a split at e are complete (Lemma 2.1.3 of [PH92]).

Figure B
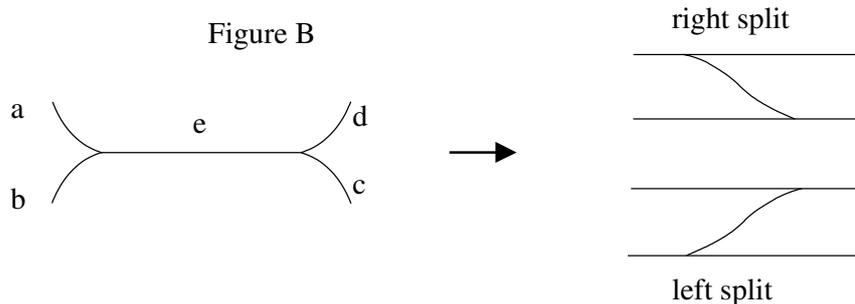
Denote by $\mathcal{TT}$ the directed graph whose vertices are the isotopy classes of complete train tracks on S and whose edges are determined as follows. The train track $\tau \in \mathcal{TT}$ is connected to the train track τ' by a directed edge if and only τ' can be obtained from τ by a single split. The graph $\mathcal{TT}$ is connected [H06a]. 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. The mapping class group $\mathcal{M}(S)$ of S acts properly and cocompactly on $\mathcal{TT}$ as a group of isometries. In particular, $\mathcal{TT}$ is $\mathcal{M}(S)$-equivariantly quasi-isometric to $\mathcal{M}(S)$ equipped with any word metric [H06a].

In the sequel we write $\tau \in \mathcal{V}(\mathcal{TT})$ if τ is a *vertex* of the graph $\mathcal{TT}$, i.e. if τ is a complete train track on S. Define a *splitting sequence* in $\mathcal{TT}$ to be a sequence $\{\alpha(i)\}_{0 \le i \le m} \subset \mathcal{V}(\mathcal{TT})$ with the property that for every $i \ge 0$ the train track $\alpha(i+1)$ can be obtained from $\alpha(i)$ by a single split. We view such a splitting sequence as a simplicial path in the graph $\mathcal{TT}$ which maps the interval $[i, i+1]$ onto the edge in $\mathcal{TT}$ connecting $\alpha(i)$ to $\alpha(i+1)$.

Recall from the introduction the definition of an L-quasi-isometric embedding of a metric space (X, d) into a metric space (Y, d). A *c-quasi-geodesic* in a metric space (X, d) is a c-quasi-isometric embedding of a closed connected subset of $\mathbb{R}$ into X. The following two results from [H06b] will be important in the sequel.

Proposition 2.1. *There is a number $c > 0$ such that every splitting sequence in $\mathcal{TT}$ is a c-quasi-geodesic.*

Proposition 2.2. *There is a number $d > 0$ with the following property. For arbitrary train tracks $\tau, \sigma \in \mathcal{V}(\mathcal{TT})$ there is a train track τ' contained in the*

d-neighborhood of τ which is splittable to a train track σ' contained in the d-neighborhood of σ.

3. Train tracks hitting efficiently

In this section we construct complete train tracks with some specific properties which are used to obtain a geometric control on the train track complex $\mathcal{TT}$. First, define a *bigon track* on S to be an embedded 1-complex on S which satisfies all the requirements of a train track except that we allow the existence of complementary bigons. Such a bigon track is called *maximal* if all complementary components are either bigons or trigons or once punctured monogons. Recurrence, transverse recurrence, birecurrence and carrying for bigon tracks are defined in the same way as they are defined for train tracks. Any complete train track is a maximal birecurrent bigon track in this sense. A *tangential measure* for a maximal bigon track ζ assigns to each branch b of ζ a nonnegative weight $\nu(b) \in [0, \infty)$ with the following properties. Each side of a complementary component of ζ can be parametrized as a trainpath ρ on ζ. Denote by $\nu(\rho)$ the sum of the weights of the branches contained in ρ counted with multiplicities. If ρ_1, ρ_2 are the two distinct sides of a complementary bigon then we require that $\nu(\rho_1) = \nu(\rho_2)$, and if ρ_1, ρ_2, ρ_3 are the three distinct sides of a complementary trigon then we require that $\nu(\rho_i) \leq \nu(\rho_{i+1}) + \nu(\rho_{i+2})$ where indices are taken modulo 3. A bigon track is transversely recurrent if and only if it admits a tangential measure which is positive on every branch [PH92].

A bigon track is called *generic* if all switches are at most trivalent. A bigon track τ which is not generic can be *combed* to a generic bigon track by successively modifying τ as shown in Figure C. By Proposition 1.4.1 of [PH92] (whose proof is also valid for bigon tracks), the combing of a recurrent bigon track is recurrent. However, the combing of a transversely recurrent bigon track need not be transversely recurrent (see the discussion on p.41 of [PH92]).

Figure C
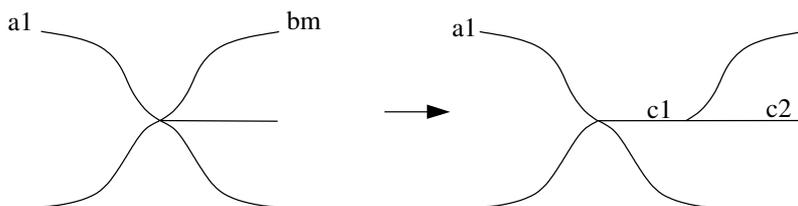
The next Lemma gives a criterion for a non-generic maximal transversely recurrent bigon track to be combable to a generic maximal transversely recurrent bigon track. For its formulation, we say that a positive tangential measure ν on a maximal bigon track σ satisfies the *strict triangle inequality for complementary trigons* if for every complementary trigon of σ with sides e_1, e_2, e_3 we have $\nu(e_i) < \nu(e_{i+1}) + \nu(e_{i+2})$. By Theorem 1.4.3 of [PH92], a *generic* maximal train track is transversely recurrent if and only if it admits a positive tangential measure satisfying the strict triangle inequality for complementary trigons. We have.

Lemma 3.1. *Let ζ be a maximal bigon track which admits a positive tangential measure satisfying the strict triangle inequality for complementary trigons. Then ζ can be combed to a generic transversely recurrent bigon track.*

Proof. Let σ be an arbitrary maximal bigon track. Then σ does not have any bivalent switches. For a switch s of σ denote the valence of s by $V(s)$ and define the excessive total valence $\mathcal{V}(\sigma)$ of σ to be $\sum_s (V(s) - 3)$ where the sum is taken over all switches s of σ; then $\mathcal{V}(\sigma) = 0$ if and only if σ is generic. By induction it is enough to show that a maximal non-generic bigon track σ which admits a positive tangential measure ν satisfying the strict triangle inequality for complementary trigons can be combed to a bigon track σ' which admits a positive tangential measure ν' satisfying the strict triangle inequality for complementary trigons and such that $\mathcal{V}(\sigma') < \mathcal{V}(\sigma)$.

For this let σ be such a non-generic maximal bigon track with tangential measure ν satisfying the strict triangle inequality for complementary trigons and let s be a switch of σ of valence at least 4. Assume that s has ℓ incoming and m outgoing branches where $1 \le \ell \le m$ and $\ell + m \ge 4$. We number the incoming branches in counter-clockwise order $a_1, \ldots, a_\ell$ (for the given orientation of S) and do the same for the outgoing branches $b_1, \ldots, b_m$. Then the branches b_m and a_1 are contained in the same side of a complementary component of σ, and the branches b_{m-1}, b_m are contained in adjacent (not necessarily distinct) sides e_1, e_2 of a complementary component T of σ. Assume first that T is a complementary trigon. Denote by e_3 the third side of T; by assumption, the total weight $\nu(e_3)$ is strictly smaller than $\nu(e_1) + \nu(e_2)$ and therefore there is a number $q \in (0, \min\{\nu(b_i) \mid 1 \le i \le m\})$ such that $\nu(e_3) < \nu(e_1) + \nu(e_2) - 2q$. Move the endpoint of the branch b_m to a point in the interior of b_{m-1} as shown in Figure C; we obtain a bigon track σ' with $\mathcal{V}(\sigma') < \mathcal{V}(\sigma)$.

The branch b_{m-1} decomposes in σ' into the union of two branches c_1, c_2 where c_1 is incident on s and on an endpoint of the image b'_m of b_m under our move. Assign the weight q to the branch c_1, the weight $\nu(b_m) - q$ to the branch b'_m and the weight $\nu(b_{m-1}) - q$ to the branch c_2. The remaining branches of σ' inherit their weight from the tangential measure ν on σ. This defines a positive weight function on the branches of σ' which by the choice of q is tangential measure satisfies the strict triangle inequality for complementary trigons.

Similarly, if the complementary component T containing b_m and b_{m-1} in its boundary is a bigon or a once punctured monogon, then we can shift b_m along b_{m-1} as before and modify our tangential measure to a positive tangential measure on the combed track with the desired properties. $\square$

Following [PH92] we say that a train track τ on our surface S *hits efficiently* a train track or a geodesic lamination σ if τ can be isotoped to a train track τ' which intersects σ transversely in such a way that $S - \tau' - \sigma$ does not contain any embedded bigon. As in [H06a] we define a *splitting and shifting sequence* to be a sequence $\{\tau_i\} \subset \mathcal{V}(\mathcal{TT})$ such that for every i the train track τ_{i+1} can be obtained from τ_i by a sequence of shifts and a single split. Denote by d the distance on $\mathcal{TT}$. We have.

Proposition 3.2. *There is a number $q > 0$ and for every $\tau \in \mathcal{V}(\mathcal{TT})$ and every complete geodesic lamination λ which hits τ efficiently there is a complete train track $\tau^* \in \mathcal{V}(\mathcal{TT})$ with the following properties.*

(1) *$d(\tau, \tau^*) \leq q$.*
(2) *τ^* carries λ.*
(3) *Let $\sigma \in \mathcal{V}(\mathcal{TT})$ be a train track which hits τ efficiently and carries λ; then τ^* carries a train track σ' which carries λ and can be obtained from σ by a splitting and shifting sequence of length at most q.*

Proof. By Lemma 3.4.4 and Proposition 3.4.5 of [PH92], for every complete train track τ there is a maximal birecurrent *dual bigon track* τ_b^* with the following property. A geodesic lamination or a train track σ hits τ efficiently if and only if σ is carried by τ_b^*. We construct the train track τ^* with the properties stated in the lemma from this dual bigon track and a complete geodesic lamination $\lambda \in \mathcal{CL}$ which hits τ efficiently and hence is carried by τ_b^*.

For this we recall from p.194 of [PH92] the precise construction of the dual bigon track τ_b^* of a complete train track τ. Namely, for each branch b of τ choose a short arc b^* meeting τ transversely in a single point in the interior of b and such that all these arcs are pairwise disjoint. Let $T \subset S - \tau$ be a complementary trigon of τ and let E be a side of T which is composed of the branches $b_1, \ldots, b_\ell$. Choose a point $p \in T$ and extend all the arcs $b_1^*, \ldots, b_\ell^*$ within T in such a way that they end at p, with the same inward pointing tangents at p. In the case $\ell \geq 2$ we then add an arc which connects p within T to a point $p' \in T$ and whose inward pointing tangent at p equals the outward pointing tangent at p of the arcs $b_1^*, \ldots, b_\ell^*$. We do this in such a way that the different configurations from the different sides of T are disjoint. If $q' \in T$ is the point in T arising in this way from a second side, then we connect p' (or p if $\ell = 1$) and q' by a smooth arc whose outward pointing tangent at p', q' coincides with the inward pointing tangents of the arcs constructed before which end at p', q'. In a similar way we construct the intersection of τ_b^* with a complementary once punctured monogon of τ. Note that the resulting graph τ_b^* is in general not generic, but its only vertices which are not trivalent arise from the sides of the complementary components of τ. Figure D shows the intersection of the dual bigon track τ_b^* with a neighborhood in S of a complementary trigon of τ and with a neighborhood in S of a complementary once punctured monogon.

Figure D

Following [PH92], τ_b^* is a maximal birecurrent bigon track, and the number of its branches is bounded from above by a constant only depending on the topological

type of S. Each complementary trigon of τ contains exactly one complementary trigon of τ_b^* in its interior, and these are the only complementary trigons. Each complementary once punctured monogon of τ contains exactly one complementary once punctured monogon of τ_b^* in its interior. All other complementary components of τ_b^* are bigons. The number of these bigons is uniformly bounded.

Now let μ be a positive integral transverse measure on τ with the additional property that the μ-weight of every branch of τ is at least 4. This weight then defines a *simple multicurve* c carried by τ in such a way that μ is just the counting measure for c (see [PH92]). Here a simple multicurve consists of a disjoint union of essential simple closed curves which can be realized disjointly; we allow that some of the curves are freely homotopic. For every side ρ of a complementary component of τ there are at least 4 connected subarcs of c which are mapped by the natural carrying map $c \to \tau$ *onto* ρ. Namely, the number of such arcs is just the minimum of the μ-weights of a branch contained in ρ.

Assign to a branch b^* of τ_b^* which is dual to the branch b of τ the weight $\nu(b^*) = \mu(b)$, and to a branch of τ_b^* which is contained in the interior of a complementary region of τ assign the weight 0. The resulting weight function ν is a tangential measure for τ_b^*, but it is not positive (this relation between transverse measures on τ and tangential measures on τ_b^* is discussed in detail in Section 3.4 of [PH92]). However by construction, every branch of vanishing ν-mass is contained in the interior of a complementary trigon or once punctured monogon of τ, and positive mass can be pushed onto these branches by "sneaking up" as described on p.39 and p.200 of [PH92]. Namely, the closed multicurve c defined by the positive integral transverse measure μ on τ hits the bigon track τ_b^* efficiently. For every branch b of τ the weight $\nu^*(b)$ equals the number of intersections between b^* and c. For each side of a complementary component T of τ there are at least 4 arcs from c which are mapped by the carrying map onto this side. If the side consists of more than one branch then we pull two of these arcs into T as shown in Figure E. If the side consists of a single branch then we pull a single arc into T in the same way.

Figure E
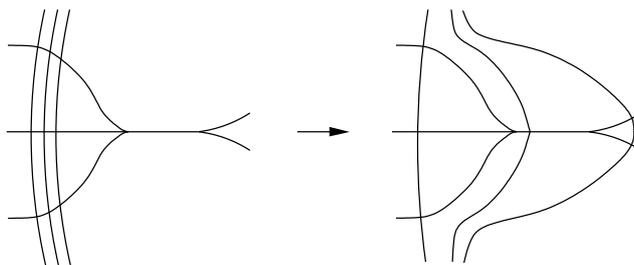
For a branch e of τ_b^* define $\mu^*(e)$ to be the number of intersections between e and the deformed multicurve. The resulting weight function μ^* is a positive integral tangential measure for τ_b^*. Note that the weight of each side of a complementary trigon in τ_b^* is exactly 2 by construction, and the weight of a side of a once punctured monogon is 2 as well. In particular, the tangential measure μ^* satisfies the strict triangle inequality for complementary trigons: If T is any complementary trigon

with sides e_1, e_2, e_3 then $\mu^*(e_i) < \mu^*(e_{i+1}) + \mu^*(e_{i+2})$ (compare the proof of Lemma 3.1).

We now modify our dual bigon track τ_b^* in a uniformly bounded number of steps to a complete train track τ^* as required in the lemma with a (non-deterministic) algorithm as follows.

The set of input data for our algorithm is the set $\mathcal{B}$ of quadruples (η, λ, ν, B) which consist of a maximal birecurrent bigon track η, a complete geodesic lamination λ carried by η, a positive tangential measure ν on η which satisfies the strict triangle inequality for complementary trigons and a complementary bigon B of η. If η does not have any complementary bigons, i.e. if η is a train track, then we put $B = \emptyset$. The algorithm modifies the quadruple $(\eta, \lambda, \nu, B) \in \mathcal{B}$ to a quadruple $(\eta', \lambda, \nu', B') \in \mathcal{B}$ with $B' = \emptyset$ as follows.

Step 1:

Let $(\eta, \lambda, \nu, B) \in \mathcal{B}$ be an input quadruple. If η does not contain a bigon, i.e. if $B = \emptyset$, then the algorithm stops. Otherwise proceed to Step 2.

Step 2:

Let $(\eta, \lambda, \nu, B) \in \mathcal{B}$ be an input quadruple with $B \neq \emptyset$. Let E, F be the sides of the complementary bigon B in the maximal birecurrent bigon track η. Check whether the boundary ∂B of B is embedded in η. If this is not the case then go to Step 3. Otherwise we construct from η and ν a maximal birecurrent bigon track $\tilde{\eta}$ which carries η and hence λ as follows. Assume that the side E of the bigon B consists of the ordered sequence of branches $e_1, \ldots, e_\ell$ and that the second side F of B consists of the branches $f_1, \ldots, f_k$. Assume also that the branches e_1, f_1 begin at a common cusp of the bigon B. We collapse the bigon B to a single arc in S with a map Ψ which identifies E and F as follows. If for some $p \geq 1$, $q \geq 1$ we have $\sum_{j=1}^{q-1} \nu(f_j) < \sum_{i=1}^{p} \nu(e_i) < \sum_{j=1}^{q} \nu(f_j)$ then Ψ maps the subarc $e_1 \cup \cdots \cup e_p$ of E homeomorphically onto a subarc of F which contains $f_1, \ldots, f_{q-1}$ and has its endpoint in the interior of the edge f_q. If $\sum_{i=1}^{p} \nu(e_i) = \sum_{j=1}^{q} \nu(f_j)$ then we map $e_1 \cup \cdots \cup e_p$ onto $f_1 \cup \cdots \cup f_q$, i.e. an endpoint of e_p is mapped to an endpoint of f_q. The resulting bigon track $\tilde{\eta}$ carries η and it is maximal. The natural carrying map $\Phi : \eta \to \tilde{\eta}$ maps each complementary trigon of η to a complementary trigon of $\tilde{\eta}$. By construction, the positive tangential measure ν on η induces a positive weight function $\tilde{\nu}$ on the branches of $\tilde{\eta}$. Note that the total weight of $\tilde{\nu}$ is strictly smaller than the total weight of ν and that the ν-weight of a side ρ of a complementary component $T \neq B$ in η coincides with the $\tilde{\nu}$-weight of the side $\Phi(\rho)$ of the complementary component $\Phi(T)$ in $\tilde{\eta}$. In particular, the weight function $\tilde{\nu}$ is a tangential measure $\tilde{\nu}$ on $\tilde{\eta}$ which satisfies the strict triangle inequality for complementary trigons. The number of complementary bigons in $\tilde{\eta}$ is strictly smaller than the number of complementary bigons in η. Namely, there is a one-to-one correspondence between the complementary bigons of $\tilde{\eta}$ and the complementary bigons of η distinct from B. The image of the bigon B under the map Φ is an embedded arc in $\tilde{\eta}$. The number of branches of $\tilde{\eta}$ does not exceed the number of branches of η. Every complete geodesic lamination which is carried by η is also carried by $\tilde{\eta}$. Choose an input quadruple of the form $(\tilde{\eta}, \lambda, \tilde{\nu}, \tilde{B}) \in \mathcal{B}$ for

a complementary bigon $\tilde{B}$ of $\tilde{\eta}$ (or $\tilde{B} = \emptyset$ if $\tilde{\eta}$ is a train track) and continue with Step 1 above for this input quadruple.

Step 3:

Let $(\eta, \lambda, \nu, B) \in \mathcal{B}$ be an input quadruple such that $B \neq \emptyset$ and that the boundary ∂B of B is *not* embedded. Then the sides E, F of B are immersed arcs of class C^1 on S which intersect or have self-intersections. Check whether the cusps of B coincide. If this is not the case, continue with Step 4.

Otherwise the two cusps of B are a common switch s of η which is necessarily at least 4-valent. By Lemma 3.1 and its proof, we can modify η with a sequence of combings near s to a maximal birecurrent bigon track $\tilde{\eta}$ in such a way that the two cusps of the complementary bigon $\tilde{B}$ in $\tilde{\eta}$ corresponding to B under the combing are distinct and such that the tangential measure ν on η induces a tangential measure $\tilde{\nu}$ on $\tilde{\eta}$ which satisfies the strict triangle inequality for complementary trigons. Note that $\tilde{\eta}$ carries the complete geodesic lamination λ. Now continue with Step 2 above with the input quadruple $(\tilde{\eta}, \lambda, \tilde{\nu}, \tilde{B}) \in \mathcal{B}$.

Step 4:

Let $(\eta, \lambda, \nu, B) \in \mathcal{B}$ be our input quadruple where B is a bigon in η with sides E, F and distinct cusps. Check whether the boundary ∂B of B contains any isolated self-intersection points. Such a point is a switch s contained in the interior of at least two distinct embedded subarcs ρ_1, ρ_2 of ∂B of class C^1 with the additional property that $\rho_1 - \{s\} \cap \rho_2 - \{s\} = \emptyset$. If ∂B does not contain such an isolated self-intersection point then continue with Step 5 below with the input quadruple $(\eta, \lambda, \nu, B) \in \mathcal{B}$.

Otherwise any such isolated self-intersection point s is a switch of η which is at least 4-valent. Thus we can modify η with a sequence combing to a complete birecurrent bigon track $\tilde{\eta}$ with the property that all self-intersection points of the boundary of the bigon $\tilde{B}$ in $\tilde{\eta}$ corresponding to B are non-isolated, i.e. they are contained in a self-intersection branch, and that the tangential measure ν on η induces a tangential measure $\tilde{\mu}$ on $\tilde{\eta}$ satisfying the strict triangle inequality for complementary trigons. Continue with Step 5 with the input quadruple $(\tilde{\eta}, \lambda, \tilde{\nu}, \tilde{B}) \in \mathcal{B}$.

Step 5:

Let $(\eta, \lambda, \nu, B) \in \mathcal{B}$ be an input quadruple where B is a complementary bigon for η whose boundary ∂B does not contain any isolated self-intersection points, whose cusps are distinct and such that the self-intersection of ∂B is not empty.

Check whether there is a branch e of η contained in the self-intersection of ∂B and which is not incident on any one of the two cusps s_1, s_2 of B. Since the interior of the bigon B is an embedded topological disc in S, such a branch e is necessarily a large branch. Now η carries λ and therefore there is a bigon track $\tilde{\eta}$ which is the image of η under a split at e and which carries λ. To each complementary region of $\tilde{\eta}$ naturally corresponds a complementary region of η of the same topological type (compare the discussion in Section 3 of [H06b] and in Section 4 of this paper). In particular, the number of complementary bigons in $\tilde{\eta}$ and η coincide, and the

bigon B in η corresponds to a bigon $\tilde{B}$ in $\tilde{\eta}$. The bigon track $\tilde{\eta}$ is recurrent, and it admits a positive tangential measure $\tilde{\nu}$ induced from the measure ν on η which satisfies the strict triangle inequality for complementary trigons. The number of branches contained in the self-intersection locus of the boundary $\partial\tilde{B}$ of the bigon $\tilde{B}$ is strictly smaller than the number of branches in the self-intersection locus of ∂B.

After a number of splits of this kind which is bounded from above by the number of branches of the bigon track η we obtain from η a bigon track η_1 which is maximal and birecurrent. There is a natural bijection from the collection of complementary bigons of η onto the collection of complementary bigons of η_1. If B_1 is the bigon of η_1 corresponding to B then the self-intersection locus of the boundary ∂B_1 of B_1 is a union of branches which are incident on one of the two cusps $s_1 \neq s_2$ of B_1. As before, η_1 admits a positive tangential measure ν_1 which satisfies the strict triangle inequality for complementary trigons and is induced from ν. Moreover, η_1 carries λ.

If the boundary ∂B_1 of B_1 is embedded then we proceed with Step 2 above for the input quadruple $(\eta_1, \lambda, \nu_1, B_1)$. Otherwise there is a self-intersection branch b of ∂B_1 which is incident on a cusp s_1 of B_1. Note that the branch b can *not* be large, so it is either small or mixed.

For a small branch b, there are again two possibilities which are shown in Figure F. A small branch b as shown on the left hand side of Figure F can be collapsed to a large branch. Using the fact that the boundary of an embedded bigon on a

Figure F
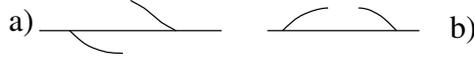
bigon track admits a natural orientation induced from the orientation of S, the small branch b is contained in the intersection of the two distinct sides of B_1. Since the tangential measure ν_1 on η_1 is positive by assumption, the construction in Step 2 above can be used to collapse the bigon B_1 in η_1 to a single simple closed curve. As in Step 2 above, we obtain a maximal birecurrent bigon track $\tilde{\eta}$ which carries λ and admits a positive tangential measure $\tilde{\nu}$ satisfying the strict triangle inequality for complementary trigons. The number of bigons of $\tilde{\eta}$ is strictly smaller than the number of bigons of η. Choose an arbitrary complementary bigon $\tilde{B}$ in $\tilde{\eta}$ or put $\tilde{B} = \emptyset$ if there is no such bigon and continue with Step 1 above and the input quadruple $(\tilde{\eta}, \lambda, \tilde{\nu}, \tilde{B}) \in \mathcal{B}$.

A small branch b as shown on the right hand side of Figure F can not be collapsed. In this case the branch b coincides with a side E_1 of the bigon B_1, and the second side F_1 of B_1 contains b as a proper subarc. Since the tangential measure ν_1 on η_1 is *positive* by assumption, this is impossible.

If the branch b is mixed then b and the cusp s_1 of B_1 are contained in the interior of a side E_1 of B_1. The bigon track η_1 can be modified with a single shift to a maximal birecurrent bigon track η_2 such that the switch s_1 is not contained any more in the interior of a side of the bigon B_2 corresponding to B_1 in η_1. The

tangential measure ν_1 on η_1 naturally induces a positive tangential measure ν_2 on η_2 which satisfies the strict triangle inequality for complementary trigons. We now proceed with Step 2 for the input quadruple $(\eta_2, \lambda, \nu_2, B_2)$. This completes the description of our algorithm.

We now apply our algorithm to the bigon track τ_b^*, the tangential measure μ^* for τ_b^* constructed from a suitably chosen transverse measure μ on τ and a complete geodesic lamination λ which hits τ efficiently and hence is carried by τ_b^*. Since the number of branches of τ_b^* is bounded from above by a constant only depending on the topological type of S, there is a universal upper bound $p > 0$ for the number of modifications of τ_b^* needed in our above algorithm to construct from these data a (possibly non-generic) birecurrent train track χ which carries λ and admits a positive tangential measure satisfying the strict triangle inequality for complementary trigons. By Lemma 3.1, this train track can be combed to a maximal birecurrent generic train track τ^* which satisfies property 2) stated in the proposition. The train track τ^* is not unique, and it depends on λ and μ. However, since our algorithm stops after a uniformly bounded number of steps and each step involves only a uniformly bounded number of choices, the number of such train tracks which can be obtained from τ_b^* by this procedure is uniformly bounded. Moroever, our algorithm is equivariant with respect to the action of the mapping class group $\mathcal{M}(S)$ and therefore by invariance under the action of $\mathcal{M}(S)$, the distance between τ and τ^* is uniformly bounded. In other words, τ^* has property 1) stated in the proposition as well.

To show property 3), let σ be a complete train track on S which hits τ efficiently and carries a complete geodesic lamination $\lambda \in \mathcal{CL}$. Then σ is carried by τ_b^* and hence it is carried by every bigon track which can be obtained from τ_b^* by a sequence of combings, shifts and collapses. On the other hand, if η_i $(0 \leq i \leq k)$ are the successive bigon tracks obtained from an application of our algorithm to $\eta_0 = \tau_b^*$ and if η_i is obtained from η_{i-1} by a split at a large branch e, then this split is a λ-split. By Lemma 4.5 of [H06a] and its proof (which is also valid for bigon tracks) there is a universal number $p > 0$ such that if σ is carried by η_{i-1} and carries λ then there is a train track $\tilde{\sigma}$ which carries λ, which is carried by η_i and which can be obtained from σ by a splitting and shifting sequence of length at most p. Since the number of splits occuring during our modification of τ_b^* to τ^* is uniformly bounded, this means that τ^* also satisfied the third requirement in our proposition. This completes the proof of our proposition. $\square$

Remark: We call the train track τ^* constructed in the proof of Proposition 3.2 from a complete train track τ and a complete geodesic lamination λ which hits τ efficiently a λ-*collapse* of τ_b^*. Note that a λ-collapse is not unique, and that in general it is neither carried by the dual bigon track of τ nor carries this bigon track. Thus in general a λ-collapse of τ does not hit τ efficiently. However, the number of different train tracks which can be obtained from our construction is bounded by a constant only depending on the topological type of S.

4. Flat strip projection

In this section we use the results from Sections 3 and from [H06b] to obtain a control on distances in the train track complex $\mathcal{TT}$.

Note first that a complementary component C of a train track σ on S is bounded by a finite number of arcs of class C^1, called *sides*. Each side either is a closed curve of class C^1 (i.e. the side does not contain any cusp) or an arc of class C^1 with endpoints at two not necessarily distinct cusps of the component. We call a side of C which does not contain cusps a *smooth side* of C. If C is a complementary component of σ whose boundary contains precisely $k \geq 0$ cusps, then the *Euler characteristic* $\chi(C)$ is defined by $\chi(C) = \chi_0(C) - k/2$ where $\chi_0(C)$ is the usual Euler characteristic of C viewed as a topological surface with boundary. Note that the sum of the Euler characteristics of the complementary components of σ is just the Euler characteristic of S. If T is a smooth side of a complementary component C of σ then we mark a point on T which is contained in the interior of a branch of σ. If T is a common smooth side of two distinct complementary components C_1, C_2 of σ then we assume that the marked points on T defined by C_1, C_2 coincide.

A *complete extension* of a train track σ is a complete train track τ containing σ as a subtrack and whose switches are distinct from the marked points in σ. Such a complete extension τ intersects each complementary component C of σ in an embedded graph. The closure τ_C of $\tau \cap C$ in S is a graph whose univalent vertices are contained in the complement of the cusps and the marked points of the boundary ∂C of C. We call two such graphs τ_C, τ_C' *equivalent* if there is an isotopy of C of class C^1 which fixes a neighborhood of the cusps and the marked points in ∂C and which maps τ_C onto τ_C'. The complete extensions τ, τ' of σ are called *σ-equivalent* if for each complementary component C of σ the graphs τ_C and τ_C' are equivalent in this sense.

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

  a) A switch v of η is also a switch of η' if and only if v is a switch of σ.
  b) For every intersection point $x \in \eta \cap \eta' \cap C$ there is a branch b of η containing x in its interior and a branch b' of η' containing x in its interior. Moreover, the branches b, b' intersect transversely at x.

We define $i_\sigma(\tau, \tau') = \sum_C i_C(\tau, \tau')$ to be the sum of the C-intersection numbers between τ and τ' where C runs through the complementary components of σ. For every number $m > 0$ there is a constant $q(m) > 0$ not depending on σ so that for every complete extension τ of σ the number of σ-equivalence classes of complete extensions τ' of σ with $i_\sigma(\tau, \tau') \leq m$ is bounded from above by $q(m)$ (see [H06b]).

Let again τ be a complete extension of a train track σ. To simplify our notation we do not distinguish between σ as a subset of τ (and hence containing switches of τ which are bivalent in σ) and σ viewed as a subtrack of τ, i.e. the graph from which the bivalent switches not contained in simple closed curve components have

been removed. Define the σ-*complexity* $\chi(\tau, \sigma)$ *of* τ to be the number of branches of τ contained in σ. Note that this complexity is not smaller than the number of branches of σ, and it coincides with this number if and only if σ is itself a complete train track. Similarly, if b is a branch in σ then we define the b-*complexity* $\chi(\tau, b)$ of τ to be the number of branches of τ contained in b. Note that $\chi(\tau, b) = 1$ if and only if b is a branch of τ.

Let b be any branch of σ which is not a branch of τ. Then b defines a trainpath $\rho : [0, m] \to \tau$ of length $m \geq 2$. For each $i < m$ the branch $\rho[i, i+1]$ of τ is a proper subset of b and will be called a *proper subbranch* of b in τ. A proper subbranch h of b in τ is incident on at least one switch v in τ which is bivalent in σ, and there is a unique branch $a \in \tau - \sigma$ which is incident on v. We call a a *neighbor* of b at v or simply a neighbor of σ at v. A neighbor of σ at an interior point v of a branch of σ is small at v. A *proper subbranch of* σ is a proper subbranch of some branch of σ. If e is a large proper subbranch of σ then τ can be split at e to a train track which contains σ as a subtrack. We call such a (not necessarily unique) split a σ-*split* of τ at e.

If $\sigma < \tau$ is a recurrent subtrack of a complete train track τ then there is a *simple geodesic multi-curve* ν on S, i.e. a collection of pairwise disjoint simple closed geodesics on S, which is carried by σ and such that a carrying map $\nu \to \sigma$ is surjective [PH92]. We call such a multi-curve *filling* for σ. More generally, we call a geodesic lamination ν carried by σ *filling* if the restriction to the minimal components of ν of a carrying map $\nu \to \sigma$ is surjective. It follows from the results of Section 2 in [H06a] that for every σ-filling lamination ν which is a disjoint union of minimal components there is a complete geodesic lamination λ which is carried by τ and contains ν as a sublamination. We call λ a *complete* τ-*extension* of ν. Whenever the existence of such a lamination λ is needed in the sequel, this lamination is given explicitly so we refrain from a more detailed discussion. We call the train track τ *tight* at a large branch e of σ if e is a large branch in τ, i.e. if $\chi(\tau, e) = 1$. In Lemma 4.3 of [H06b] we showed.

Lemma 4.1. *There is a number* $p > 0$ *with the following properties. Let* σ *be a recurrent train track, let* e *be a large branch of* σ *and let* ν *be a* σ-*filling geodesic lamination. Then there is an algorithm which associates to every complete extension* τ *of* σ *and every complete* τ-*extension* λ *of* ν *a complete train track* τ' *with the following properties.*

(1) τ' *can be obtained from* τ *by a sequence of* λ-*splits of length at most* p *at proper subbranches of* e, *and it contains* σ *as a subtrack.*
(2) τ' *is tight at* e.
(3) *If no marked point of* σ *is contained in the branch* e *then we have* $i_\sigma(\tau', \eta) \leq i_\sigma(\tau, \eta) + p(\chi(\tau, \sigma) - \chi(\tau', \sigma))$ *for every complete extension* η *of* σ.
(4) *If there is a marked point of* σ *contained in* e *then we have* $i_\sigma(\tau', \eta) \leq i_\sigma(\tau, \eta) + p$ *for every complete extension* η *of* σ.

For train tracks $\sigma < \tau$ as in Lemma 4.1 we call τ' the σ-*modification* of τ at the large branch e with respect to the complete τ-extension λ of ν.

Let σ be a recurrent train track on S without closed curve components. Let $\tau \in \mathcal{V}(\mathcal{TT})$ be a complete extension of σ and let $\{\sigma(i)\}_{0 \leq i \leq \ell}$ be a splitting sequence issuing from $\sigma(0) = \sigma$. We call this splitting sequence *recurrent* if each of the train tracks in the sequence is recurrent. Call a splitting sequence $\{\eta(j)\}_{0 \leq j \leq m} \subset \mathcal{V}(\mathcal{TT})$ beginning at $\eta(0) = \tau$ *induced* by the sequence $\{\sigma(i)\}_{0 \leq i \leq \ell}$ if there is an injective strictly increasing map $q : \{0, \ldots, \ell\} \to \{0, \ldots, m\}$ with the following properties.

a) $q(0) = 0, q(\ell) \leq m$ and for $q(i) \leq j < q(i+1)$ the train track $\eta(j)$ contains a subtrack isotopic to $\sigma(i)$.
b) Let $i < \ell$ and assume that the split $\sigma(i) \to \sigma(i+1)$ is a right (or left) split at a large branch e_i. Then there is a complete geodesic lamination $\lambda \in \mathcal{CL}$ which is a complete $\eta(q(i+1)-1)$-extension of a filling lamination for $\sigma(i)$ and such that $\eta(q(i+1)-1)$ is the $\sigma(i)$-modification of $\eta(q(i))$ at e_i with respect to λ. The train track $\eta(q(i+1)-1)$ is tight at e_i and the split $\eta(q(i+1)-1) \to \eta(q(i+1))$ is a right (or left) split at e_i.
c) For $q(\ell) \leq j < m$ the train track $\eta(j+1)$ is obtained from $\eta(j)$ by a split at a large proper subbranch of $\sigma(\ell)$.

In the next lemma we compare distances between complete train tracks which are obtained from splitting sequences induced by a splitting sequence of a common subtrack.

Lemma 4.2. *There is a number $R_0 > 0$ with the following property. Let $\sigma(0)$ be a train track on S without closed curve components and let $\{\sigma(i)\}_{0 \leq i \leq \ell}$ be a finite recurrent splitting sequence issuing from $\sigma(0)$. Let $\tau(0), \eta(0) \in \mathcal{V}(\mathcal{TT})$ be two complete extensions of $\sigma(0)$ and let $\{\tau(j)\}_{0 \leq j \leq m} \subset \mathcal{V}(\mathcal{TT}), \{\eta(p)\}_{0 \leq p \leq n} \subset \mathcal{V}(\mathcal{TT})$ be splitting sequences beginning at $\tau(0), \eta(0)$ which are induced by $\{\sigma(i)\}_{0 \leq i \leq \ell}$; then $d(\tau(m), \eta(n)) \leq d(\tau(0), \eta(0)) + R_0$.*

Proof. Let σ by a recurrent train track on S without closed curve components. Then σ decomposes S into a finite number of complementary components $C^1, \ldots, C^u$. Among these complementary components there are components $C^1, \ldots, C^s$ which contain essential simple closed curves not homotopic into a puncture, and there are components $C^{s+1}, \ldots, C^u$ which are topological discs or once punctured topological discs. As in the beginning of this section, we mark a point on each smooth boundary component of the sets C^i.

For $k \leq u$ the boundary ∂C^k of C^k consists of a finite number of connected components. Each of these components consists of finitely many sides. Such a side either is a simple closed curve or an arc terminating at two not necessarily distinct cusps of the component, and it can be represented as a not necessarily embedded trainpaths on σ. The union of these trainpaths determine a closed curve on S which is freely homotopic to a simple closed curve, and this simple closed curve is essential if and only if $k \leq s$, i.e. if the component C^k is different from a disc or a once punctured disc. In particular, for every $k \leq s$ there is a bordered oriented surface $S^k \subset C^k$ with smooth boundary ∂S^k consisting of a finite number of embedded circles and such that $C^k - S^k$ is a finite union of essential open annuli, one annulus for each boundary component of S^k. The Euler characteristic of S^k is negative. Define $S_0 = \cup_{k=1}^s S^k$; then S_0 is an embedded bordered subsurface of S.

Let $\mathcal{M}_0(S_0)$ be the subgroup of finite index of the mapping class group of the bordered surface S_0 consisting of all mapping classes which can be represented by diffeomorphisms preserving each component of S_0 and each of the boundary components of S_0. By convention, the group $\mathcal{M}_0(S_0)$ contains in its center the free abelian group generated by the Dehn twists about the boundary components of S_0. Then $\mathcal{M}_0(S_0)$ can naturally be viewed as a subgroup of $\mathcal{M}(S)$ consisting of mapping classes which can be represented by diffeomorphisms fixing the complement of a small neighborhood of S_0 pointwise. Since σ is contained in $S - S_0$, the group $\mathcal{M}_0(S_0)$ acts on the set $\mathcal{E}(\sigma) \subset \mathcal{V}(\mathcal{TT})$ of complete extensions of σ. Since each of the complementary components C^k ($s+1 \leq k \leq u$) of σ which is contained in $S - S_0$ is a topological disc or a once punctured topological disc, for $s+1 \leq k \leq u$ the boundary ∂C^k of C^k is connected and contains at least one cusp. Namely, otherwise this boundary defines a closed trainpath on σ which necessarily is an essential curve on S. There are only finitely many equivalence classes of graphs τ_{C^k} in the sense defined in the beginning of this section which occur as the closures of the intersection of a complete train track τ with C^k. As a consequence, there is a number $r > 0$ only depending on the topological type of S such that there are at most r orbits in $\mathcal{E}(\sigma)$ under the action of $\mathcal{M}_0(S_0)$. Moreover, there is a number $q > 0$ with the following property. For any $\tau, \eta \in \mathcal{E}(\sigma)$ there is an element $\Theta(\tau, \eta) \in \mathcal{M}_0(S_0)$ (see [H06a]) such that $i_\sigma(\tau, \Theta(\tau, \eta)\eta) \leq q$.

Let $\{\sigma(i)\}_{0 \leq i \leq \ell}$ be a recurrent sequence of splits of $\sigma = \sigma(0)$. Since we only allow right or left splits, for each i there is a natural diffeomorphism φ_i of $S - \sigma$ onto $S - \sigma(i)$ which can be chosen to be the identity on S_0. The restriction φ_i^k of this diffeomorphism to a complementary component C^k of σ maps C^k diffeomorphically onto a complementary component C_i^k of $\sigma(i)$ and extends continuously to a bijection of the cusps of ∂C^k onto the cusps of ∂C_i^k. Namely, if we split $\sigma(0)$ at a large branch e to the train track $\sigma(1)$, then for each complementary component C^k of $\sigma(0)$ there is a unique complementary component C_1^k of $\sigma(1)$ which is diffeomorphic to C^k with a diffeomorphism φ_1^k which fixes pointwise the complement of a neighborhood of e. For $k \leq s$ the diffeomorphism φ_1^k can be chosen to be the identity on the subsurface $S^k \subset C^k$ and to induce a natural (oriented) bijection between the sides of C^k and the sides of C_1^k which is the identity outside a neighborhood of e. For every smooth boundary component T of C^k there is a marked point x contained in the interior of a branch $b \subset T$. If $b \neq e$ then we require that the image of b under the natural bijection of the branches of σ onto the branches of $\sigma(1)$ contains a marked point in its interior, and if $b = e$ then we place a marked point in the interior of the unique losing branch of the split which is contained in the smooth boundary component of C_1^k corresponding to T. For $i \geq 2$ the maps φ_i^k are constructed inductively from the maps φ_{i-1}^k in this way.

Let τ, η be any complete extensions of σ and let $\Theta = \Theta(\tau, \eta) \in \mathcal{M}_0(S_0)$ be such that $i_\sigma(\tau, \Theta(\eta)) \leq q$. Since for every $i \leq \ell$ the map φ_i restricts to the identity on S_0, for every splitting sequence $\{\eta(j)\}_{0 \leq j \leq n}$ induced by a splitting sequence $\{\sigma(i)\}_{0 \leq i \leq \ell}$ of $\sigma = \sigma(0)$ and issuing from $\eta = \eta(0)$, the sequence $\{\Theta\eta(j)\}_{0 \leq j \leq n}$ is a splitting sequence issuing from $\Theta(\eta)$ and induced by the sequence $\{\sigma(i)\}_{0 \leq i \leq \ell}$. For every complete extension ξ of $\sigma(\ell)$ we have $d(\xi, \eta(n)) \leq d(\xi, \Theta\eta(n)) + d(\Theta\eta(n), \eta(n))$. Moreover, since the action of $\mathcal{M}_0(S_0)$ commutes with the action of the subgroup of the mapping class group of $S - S_0$ consisting of all mapping classes which preserve

every boundary component of $S - S_0$, by invariance we have $d(\Theta\eta(n), \eta(n)) \leq d(\Theta\eta(0), \eta(0)) + c$ where $c > 0$ is a universal constant. Thus for the proof of the lemma, it is enough to show the existence of a number $b > 0$ only depending on the topological type of S with the following properties. Let τ, η be any complete extensions of σ with $i_\sigma(\tau, \eta) \leq q$ and let $\{\tau(j)\}_{0 \leq j \leq m} \subset \mathcal{V}(\mathcal{TT}), \{\eta(k)\}_{0 \leq k \leq n} \subset \mathcal{V}(\mathcal{TT})$ be any two splitting sequences issuing from the train tracks $\tau(0) = \tau, \eta(0) = \eta$ which are induced by the splitting sequence $\{\sigma(i)\}_{0 \leq i \leq \ell}$; then $i_{\sigma(\ell)}(\tau(m), \eta(n)) \leq b$.

To show that this is indeed the case we use Lemma 4.1 inductively. Namely, let ν be a geodesic lamination which is carried by $\sigma(\ell)$ and fills $\sigma(\ell)$. Such a lamination exists since the splitting sequence $\{\sigma(i)\}_{0 \leq i \leq \ell}$ is recurrent by assumption, and it is carried by $\sigma = \sigma(0)$ and fills σ. Let λ be a complete $\tau(m)$-extension of ν and let μ be a complete $\eta(n)$-extension of ν. Then λ, μ are complete τ, η-extensions of ν. Assume that $\sigma(1)$ is obtained from σ by a split at a large branch e. Let $\tau', \eta' \in \mathcal{V}(\mathcal{TT})$ be the σ-modification of τ, η at e determined by λ, μ. By the definition of an induced splitting sequence, there are numbers $m(1) \geq 1, n(1) \geq 1$ such that $\tau(m(1)), \eta(n(1))$ can be obtained from τ', η' by a single split at e, and $\tau(m(1)), \eta(n(1))$ contain $\sigma(1)$ as a subtrack. We call the modification which transforms σ to $\sigma(1)$ and τ, η to $\tau(m(1)), \eta(n(1))$ a *move* and we distinguish two types of move.

Type 1: σ does not contain a marked point in the interior of e.

By Lemma 4.1 and the above discussion, with $p > 0$ as in Lemma 4.1 we have

$$i_{\sigma(1)}(\tau(m(1)), \eta(n(1))) = i_\sigma(\tau', \eta') \leq i_\sigma(\tau', \eta) + p(\chi(\eta, \sigma) - \chi(\eta', \sigma))$$
$$\leq i_\sigma(\tau, \eta) + p\big[\chi(\tau, \sigma) - \chi(\tau(m(1)), \sigma(1)) + \chi(\eta, \sigma) - \chi(\eta(n(1)), \sigma(1))\big]$$

and therefore $i_{\sigma(1)}(\tau(m(1)), \eta(n(1))) \leq i_\sigma(\tau, \eta) + 2r$ for a universal constant $r > 0$. Moreover, we have equality $i_{\sigma(1)}(\tau(m(1)), \eta(n(1))) = i_\sigma(\tau, \eta)$ if $\chi(\tau(m(1)), \sigma(1)) = \chi(\tau, \sigma)$ and $\chi(\eta(n(1)), \sigma(1)) = \chi(\eta, \sigma)$.

Type 2: There is a marked point of σ contained in the interior of the branch e.

Using again Lemma 4.1 and the above discussion, we conclude as before that $i_{\sigma(1)}(\tau(m(1)), \eta(n(1))) \leq i_\sigma(\tau, \eta) + 2p$.

Let $m(\ell) > 0, n(\ell) > 0$ be the smallest numbers such that the train tracks $\tau(m(\ell)), \eta(n(\ell))$ contain $\sigma(\ell)$ as a subtrack. By the definition of a splitting sequence induced from the sequence $\{\sigma(i)\}_{0 \leq i \leq \ell}$, the train tracks $\tau(m(\ell)), \eta(n(\ell))$ can be obtained from τ, η by ℓ moves. The train tracks $\tau(m), \eta(n)$ contain $\sigma(\ell)$ as a subtrack and are obtained from $\tau(m(\ell))$ by a sequence of splits at large proper subbranches of $\sigma(\ell)$. Since $\sigma(\ell)$ does not have any closed curve components, the length of a splitting sequence connecting $\tau(m(\ell)), \eta(n(\ell))$ to $\tau(m), \eta(n)$ is bounded from above by a universal constant (compare the detailed discussion in Section 4 of [H06b]). In particular, the distance between $\tau(m)$ and $\eta(n)$ is bounded from above by $d(\tau(m(\ell)), \eta(n(\ell)) + \tilde{c}$ where $\tilde{c} > 0$ is a universal constant. Since moreover $\chi(\tau, \sigma) + \chi(\eta, \sigma)$ is bounded from above by a universal constant, the lemma follows if we can show that the number of times a move of type 2 occurs in our splitting sequences is bounded from above by a universal constant.

However, if for some $i < \ell$ the split $\sigma(i) \to \sigma(i+1)$ is a split at a large branch e_i containing a marked point of σ in its interior then e_i is contained in a smooth component A of the boundary of a complementary component C_i^k of $\sigma(i)$. In particular, A defines a closed trainpath ρ on $\sigma(i)$ which is freely homotopic to a simple closed curve defining a boundary component of C_i^k. Thus the trainpath ρ passes through any branch of $\sigma(i)$ at most twice and hence its length is uniformly bounded. The train track $\sigma(i+1)$ obtained from $\sigma(i)$ by a single split at e_i contains A as an embedded trainpath ρ' whose length is strictly smaller than the length of ρ. As a consequence, the number of moves of type 2 is bounded from above by a universal constant. This completes the proof of the lemma. $\square$

Next we estimate distances in $\mathcal{T}\mathcal{T}$ between train tracks which do not carry a common geodesic lamination. We show.

Lemma 4.3. *For every $R > 0$ there is a number $\beta_0 = \beta_0(R) > 0$ with the following property. Let $\tau, \eta \in \mathcal{V}(\mathcal{T}\mathcal{T})$ with $d(\tau, \eta) \le R$ and let τ', η' be complete train tracks which can be obtained from τ, η by any splitting sequence. If τ, η do not carry any common geodesic lamination then there is a point $\tau'' \in \mathcal{V}(\mathcal{T}\mathcal{T})$ in the $\beta_0(R)$-neighborhood of τ' which can be connected to a point $\eta'' \in \mathcal{V}(\mathcal{T}\mathcal{T})$ contained in the $\beta_0(R)$-neighborhood of η' by a splitting sequence which passes through the $\beta_0(R)$-neighborhood of τ. Moreover, for any complete geodesic lamination ν which is carried by η', we can assume that η'' carries ν.*

Proof. For a fixed number $R > 0$ there are only finitely many orbits under the action of the mapping class group of pairs $(\tau, \eta) \in \mathcal{V}(\mathcal{T}\mathcal{T}) \times \mathcal{V}(\mathcal{T}\mathcal{T})$ where $d(\tau, \eta) \le R$ and such that τ, η do not carry any common geodesic lamination. Thus by invariance under the mapping class group it is enough to show the lemma for two fixed train tracks $\tau, \eta \in \mathcal{V}(\mathcal{T}\mathcal{T})$ which do not carry any common geodesic lamination and with a constant $\beta_0 > 0$ depending on τ, η.

For a complete train track ξ denote by $\mathcal{C}\mathcal{L}(\xi)$ the set of all complete geodesic laminations which are carried by ξ. By Lemma 2.4 of [H06a], the set $\mathcal{C}\mathcal{L}(\xi)$ is open and closed in $\mathcal{C}\mathcal{L}$. We first show that there are finitely many complete train tracks $\tau_1, \dots, \tau_\ell$ and $\eta_1, \dots, \eta_m$ with the following properties.

(1) For each $i \le \ell$ the train track τ_i is carried by τ and $\cup_i \mathcal{C}\mathcal{L}(\tau_i) = \mathcal{C}\mathcal{L}(\tau)$.
(2) For each $j \le m$ the train track η_j is carried by η and $\cup_j \mathcal{C}\mathcal{L}(\eta) = \mathcal{C}\mathcal{L}(\eta)$.
(3) For all $i \le \ell, j \le m$ the train tracks τ_i, η_j hit efficiently.

For this observe that since τ, η do not carry any common geodesic lamination, every lamination $\lambda \in \mathcal{C}\mathcal{L}(\tau)$ intersects every lamination $\mu \in \mathcal{C}\mathcal{L}(\eta)$ transversely. Namely, a complete geodesic lamination decomposes the surface S into ideal triangles and once punctured monogons. Thus if ℓ is any simple geodesic on S whose closure in S is compact and if ℓ does *not* intersect the complete geodesic lamination μ transversely then ℓ is contained in μ and hence the closure of ℓ is a sublamination of μ. Now if ℓ is a leaf of the complete geodesic lamination λ then the closure of ℓ is a sublamination of λ as well. Since λ is carried by τ and μ is carried by σ, this violates our assumption that τ, σ do not carry a common geodesic lamination.

As a consequence, for every lamination $\lambda \in \mathcal{CL}(\tau)$, every train track ξ which is sufficiently close to λ in the Hausdorff topology hits every lamination $\nu \in \mathcal{CL}(\eta)$ efficiently. Thus by Lemma 2.2 and Lemma 2.3 of [H06a], for every complete geodesic lamination $\lambda \in \mathcal{CL}(\tau)$ there is a train track $\tau(\lambda) \in \mathcal{V}(\mathcal{TT})$ which carries λ, which is carried by τ and which hits every lamination $\nu \in \mathcal{CL}(\eta)$ efficiently. The set $\mathcal{CL}(\tau(\lambda))$ is an open subset of $\mathcal{CL}(\tau)$ and therefore by compactness of $\mathcal{CL}(\tau)$ there are finitely many laminations $\lambda_1, \ldots, \lambda_\ell \in \mathcal{CL}(\tau)$ such that $\mathcal{CL}(\tau) = \cup_i \mathcal{CL}(\tau(\lambda_i))$. Write $\tau_i = \tau(\lambda_i)$.

Every lamination $\nu \in \mathcal{CL}(\eta)$ hits each of the train tracks τ_i efficiently. As a consequence, if ξ is a train track which is sufficiently close to ν in the Hausdorff topology then ξ hits each of the train tracks τ_i $(i \leq \ell)$ efficiently. As before, this implies that we can find a finite family $\eta_1, \ldots, \eta_m \in \mathcal{V}(\mathcal{TT})$ of train tracks which are carried by η, which hit each of the train tracks τ_i $(i \leq \ell)$ efficiently and such that $\cup_j \mathcal{CL}(\eta_j) = \mathcal{CL}(\eta)$. This shows our above claim.

Let $k = \max\{d(\tau, \tau_i), d(\eta, \eta_j) \mid i \leq \ell, j \leq m\}$. Let τ', η' be obtained from τ, η by a splitting sequence and let $\lambda \in \mathcal{CL}(\tau')$ be a complete geodesic lamination carried by τ'. Then $\lambda \in \mathcal{CL}(\tau)$ and hence by our above construction, there is some $i \leq \ell$ such that τ_i carries λ. By Corollary 4.6 of [H06a] there is a complete train track ξ which carries λ, is carried by both τ' and τ_i and whose distance to τ' is bounded from above by a universal constant $q > 0$ only depending on k. Similarly, for a complete train track η' which can be obtained from η by a splitting sequence and for some $\nu \in \mathcal{CL}(\eta')$ there is some $j \leq m$ and a train track ζ which is carried by both η' and η_j, which carries ν and whose distance to η' does not exceed q.

Since τ_i, η_j hit efficiently and τ_i carries ξ, the train tracks ξ, η_j hit efficiently. In particular, the geodesic lamination ν hits ξ efficiently. Therefore by Proposition 3.2, a ν-collapse ξ^* of the dual bigon track ξ_b^* of ξ carries a train track β which carries ν and can be obtained from η_j by a splitting and shifting sequence of length at most q. In particular, the distance between η_j and β and hence the distance between τ and β is uniformly bounded. Since η_j carries ζ and ζ carries ν, we conclude from Corollary 4.6 of [H06a] that β carries a train track σ contained in a uniformly bounded neighborhood of ζ and hence in a uniformly bounded neighborhood of η'.

The distance between τ' and ξ is uniformly bounded and therefore the distance between ξ^* and τ' is uniformly bounded as well (compare Proposition 3.2 and the following remark). On the other hand, since β is carried by ξ^* the train track ξ^* can be connected to β by a splitting and shifting sequence (Theorem 2.4.1 of [PH92]). We deduce from Lemma 5.4 of [H06a] that ξ^* is splittable to train tracks β', σ' which carry ν and are contained in a uniformly bounded neighborhood of β, σ. It then follows from Lemma 5.1 of [H06a] that β' is splittable to a train track σ'' which carries ν and is contained in a uniformly bounded neighborhood of σ'. Since the distance between σ' and η' is uniformly bounded, this implies the lemma.   □

Our next goal is to obtain a suitable extension of Lemma 4.3 to train tracks $\tau, \eta \in \mathcal{V}(\mathcal{TT})$ containing a common subtrack which is a union of simple closed curves. We first have to overcome a technical difficulty arising from the fact that the split of a complete train track need not be recurrent. For this call a large branch

e of a complete train track η *rigid* if only one of the two train tracks obtained from η by a split at e is complete. Note that e is a rigid large branch of η if and only if the (unique) complete train track obtained from η by a split at e carries *every* complete geodesic lamination which is carried by η. It follows from the results of [PH92] that a branch e in η is rigid if and only if the train track ζ obtained from η by a collision, i.e. a split followed by the removal of the diagonal of the split, is *not* recurrent. We have.

Lemma 4.4. *There is a number $a_1 > 0$ with the following property. For every $\eta \in \mathcal{V}(\mathcal{TT})$ there is a splitting sequence $\{\eta(i)\}_{0 \leq i \leq s}$ issuing from $\eta = \eta(0)$ of length $s \leq a_1$ such that for every i, $\eta(i+1)$ is obtained from $\eta(i)$ by a split at a rigid large branch and such that $\eta(s)$ does not contain any rigid large branches.*

Proof. We show the existence of a constant $a_1 > 0$ as in the lemma with an argument by contradiction. Assume that our claim does not hold. Then there is a sequence of pairs $(\beta_i, \xi_i) \in \mathcal{V}(\mathcal{TT})$ such that ξ_i can be obtained from β_i by a splitting sequence of length at least i consisting of splits at rigid large branches. Every complete geodesic lamination which is carried by β_i is also carried by ξ_i. By invariance under the action of the mapping class group $\mathcal{M}(S)$ and the fact that there are only finitely many $\mathcal{M}(S)$-orbits on $\mathcal{V}(\mathcal{TT})$, by passing to a subsequence we may assume that there is some $\eta \in \mathcal{V}(\mathcal{TT})$ such that $\beta_i = \eta$ for all i. Since η has only finitely many large branches, by a standard diagonal procedure we construct an *infinite* splitting sequence $\{\eta(i)\}_{i \geq 0}$ issuing from $\eta = \eta(0)$ such that for every i the train track $\eta(i+1)$ is obtained from $\eta(i)$ by a split at a rigid large branch. Then for every i, every complete geodesic lamination which is carried by η is also carried by $\eta(i)$. Now for every *projective measured geodesic lamination* whose support ν is carried by η there is a complete geodesic lamination λ carried by η which contains ν as a sublamination [H06a]. Thus for each i, the space $\mathcal{PM}(i)$ of projective measured geodesic laminations carried by $\eta(i)$ coincides with the space $\mathcal{PM}(0)$ of projective measured geodesic laminations carried by η. On the other hand, the set $\mathcal{PM}(0)$ contains an open subset of the space of all projective measured geodesic laminations since η is complete (see [PH92]). Therefore $\cap_i \mathcal{PM}(i) = \mathcal{PM}(0)$ contains an open subset of projective measured geodesic lamination space which contradicts Theorem 8.5.1 in [M03a]. This shows our lemma. $\square$

To each complete geodesic lamination λ on S and every simple closed curve component c of λ we associate a sign $\mathrm{sgn}(\lambda, c) \in \{+, -\}$ as follows. Let S_1 be the surface obtained from S by cutting S open along c. We view S_1 as a bordered surface with two boundary components c_1, c_2 corresponding to c. The orientation of S then induces a boundary orientation for c_1, c_2. For our complete geodesic lamination λ containing c as a minimal component and for $i = 1, 2$ there is at least one leaf of λ which is contained in S_1 and spirals about c_i. We associate to λ and c the sign $\mathrm{sgn}(\lambda, c) = +$ if the spiraling direction of such a leaf coincides with the boundary orientation of c_i for $i = 1, 2$, and we associate to λ and c the sign $\mathrm{sgn}(\lambda, c) = -$ otherwise. Since λ is complete by assumption, if $\mathrm{sgn}(\lambda, c) = -$ then the spiraling direction of a leaf of λ spiraling about c_i $(i = 1, 2)$ is opposite to the orientation of c_i as a component of the boundary of S_1 (compare Section 2 of [H06a]).

Next we look at a complete train track τ containing a subtrack β which is a union of $k \geq 1$ simple closed curves embedded in τ. In other words, β is a subtrack of τ without large branches. Such a train track τ carries a complete geodesic lamination λ which contains the k simple closed curve components $c_1, \ldots, c_k$ of β as minimal components. By the above, λ determines for each of the curves c_i a sign. Recall the definition of a *Dehn twist* about an essential simple closed curve c in S. We define the twist to be *positive* if the direction of the twist coincides with the direction given by the boundary orientation of c in the surface S_1 obtained by cutting S open along c. We use these signs to estimate distances in $\mathcal{TT}$ obtained by splitting complete train tracks along simple closed curve subtracks. We show.

Lemma 4.5. *There is a constant $a_2 > 0$ with the following property. Let $\tau \in \mathcal{V}(\mathcal{TT})$, let $c < \tau$ be an embedded simple closed curve and let φ_c be the positive Dehn twist about c. Let $\{\tau_i\}_{0 \leq i \leq m}$ be a splitting sequence issuing from $\tau_0 = \tau$ which consists of c-splits at large proper subbranches of c. Let $\lambda \in \mathcal{CL}$ be a complete geodesic lamination which is carried by τ and contains c as a minimal component. Then there is some $i \geq 0$ such that the distance between τ_m and the train track $\varphi_c^{\mathrm{sgn}(\lambda,c)i}\tau$ is at most a_2.*

Proof. A *standard twist connector* in a complete train track ξ is an embedded closed curve α in ξ which consists of a large branch and a small branch, connected at two switches. If ξ' is obtained from ξ by an α-split at the large branch in α, then ξ' is obtained from ξ by a half-Dehn twist about α whose sign is determined by the neighbors of the subtrack $\alpha < \xi$ (compare the discussion in Section 2 of [H06a]). In particular, for every complete geodesic lamination λ which is carried by ξ and contains α as a minimal component, the sign $\mathrm{sgn}(\lambda, \alpha)$ is determined by the twist connector and hence does not depend on λ.

By the considerations in Section 4 of [H06b], there is a number $a_0 > 0$ only depending on the topological type of S such that for every complete train track τ containing an embedded simple closed curve c the image η of τ under a sequence of c-splits of length at most a_0 contains c as a *simple vertex cycle*, i.e. η can be shifted to a train track η' which contains c as a standard twist connector. Thus any sequence of c-splits modifies η to a train track η_1 which is contained in uniformly bounded neighborhood of the image of η under a multiple of the Dehn twist along c whose direction is determined by $\mathrm{sgn}(\lambda, c)$ where λ is a complete extension of c carried by η. From this the lemma is immediate. □

A *simple multicurve* c on S consists of a finite collection $c_1, \ldots, c_k$ of essential simple closed curves which are not mutually freely homotopic and which can be realized disjointly. A simple multicurve c can be viewed as a train track without large branches by adding a single switch to each component of c. A *Dehn multitwist* of a multicurve $c = \cup_i c_i$ is an element $\varphi \in \mathcal{M}(S)$ which can be represented in the form $\varphi = \varphi_{c_1}^{m_1} \circ \cdots \circ \varphi_{c_k}^{m_k}$ for some $m_i \in \mathbb{Z}$ and where as before, φ_{c_i} is the positive Dehn twist about c_i. The next lemma is an extension of Lemma 4.3. For its formulation (and later use), for a train track $\tau \in \mathcal{V}(\mathcal{TT})$ which is splittable to a train track $\tau' \in \mathcal{V}(\mathcal{TT})$ define the *flat strip* $E(\tau, \tau')$ to be the maximal subgraph of $\mathcal{TT}$ whose vertices consist of the collection of all train tracks which can be obtained from τ by a splitting sequence and are splittable to τ'. We have.

Lemma 4.6. *For every $R > 0$ there is a number $\beta_1(R) > 0$ with the following property. Let $\tau, \eta \in \mathcal{V}(\mathcal{TT})$ with $d(\tau, \eta) \leq R$. Assume that τ, η contain a common subtrack σ which is a multicurve and that the components of this multicurve are precisely the minimal geodesic laminations which are carried by both τ and η. Let τ', η' be complete train tracks which can be obtained from τ, η by a splitting sequence. Then there is a Dehn multitwist φ about c and a point τ'' in the $\beta_1(R)$-neighborhood of τ' which can be connected to a point η'' in the $\beta_1(R)$-neighborhood of η' by a splitting sequence passing through the $\beta_1(R)$-neighborhood of $\varphi\tau$. Moreover, the distance between $\varphi(\tau)$ and $E(\tau, \tau'), E(\eta, \eta')$ is at most $\beta_1(R)$.*

Proof. As in the proof of Lemma 4.3, by invariance under the mapping class group it suffices to show the lemma for some fixed train tracks $\tau, \eta \in \mathcal{V}(\mathcal{TT})$ which satisfy the assumptions in the lemma. In particular, τ, η contain a common subtrack c which is a simple multicurve and such that every minimal geodesic lamination which is carried by both τ and η is a component of c. By Lemma 4.3, we may assume that c is not empty.

Let $c_1, \ldots, c_k$ be the components of c. Then for $i \neq j$, a c-split of τ at a large proper subbranch of c_i commutes with a c-split of τ at a large proper subbranch of c_j. Assume that τ is splittable to a train track τ'. By Lemma 4.4, via replacing τ' by its image under a splitting sequence of uniformly bounded length we may assume that τ' does not contain any rigid large branches. Define inductively a sequence $\{\tau(i)\}_{0 \leq i \leq k} \subset E(\tau, \tau')$ consisting of train tracks $\tau(i)$ which contain c as a subtrack as follows. Put $\tau(0) = \tau$ and assume that $\tau(i-1)$ has been defined for some $i \in \{1, \ldots, k\}$. Let $\tau(i) \in E(\tau, \tau')$ be the train track which can be obtained from $\tau(i-1)$ by a sequence of c_i-splits of maximal length at proper large subbranches of c_i. Since splits at large subbranches of c_i, c_j for $i \neq j$ commute, the train track $\tau(k)$ only depends on τ, τ', c but not on the ordering of the components c_i of c. By Lemma 4.5, there are numbers $b_i \in \mathbb{Z}$ such that for $\varphi_\tau = \varphi_{c_1}^{b_1} \circ \cdots \circ \varphi_{c_k}^{b_k} \in \mathcal{M}(S)$ the distance between $\tau(k)$ and the train track $\varphi_\tau(\tau)$ is bounded from above by a_2.

Similarly, if η is splittable to a train track $\eta' \in \mathcal{V}(\mathcal{TT})$ without rigid large branches, then there are numbers $p_i \in \mathbb{Z}$ such that for $\varphi_\eta = \varphi_{c_1}^{p_1} \circ \cdots \circ \varphi_{c_k}^{p_k} \in \mathcal{M}(S)$ the distance between the train track $\eta(k)$ obtained from η and η' by the above procedure and $\varphi_\eta(\eta)$ is bounded from above by a_2.

For $i \leq k$ define $m(i) = 0$ if the signs of b_i, p_i are distinct, and if the signs of b_i, p_i coincide then define $m(i) = \text{sgn}(a_i) \min\{|a_i|, |b_i|\}$. Write $\varphi = \varphi_{c_1}^{m(1)} \circ \cdots \circ \varphi_{c_k}^{m(k)}$. By the choice of the multi-twist φ and by invariance of the distance function on $\mathcal{TT}$ under the action of the mapping class group, there is a number $\chi > 0$ only depending on $d(\tau, \eta)$ and there are train tracks $\tau_1 \in E(\tau, \tau'), \eta_1 \in E(\eta, \eta')$ contained in the χ-neighborhood of $\varphi(\tau), \varphi(\eta)$. Moreover, by the choice of φ, we may assume that τ_1, η_1 contain the simple multicurve c as a subtrack and that for each i one of the following two possibilities is satisfied.

(1) $m(i) = 0$ and a splitting sequence connecting τ, η to τ_1, η_1 does not contain any split at a large proper subbranch of c_i.

(2) If $|b_i| \leq |p_i|$ then the flat strip $E(\tau, \tau')$ does not contain any train track which can be obtained from τ_1 by a c_i-split at a large proper subbranch of τ_1 and similarly for η_1 in the case that $|p_i| \leq |b_i|$.

After reordering we may assume that there is some $s \leq k$ such that $m(i) = 0$ for $i \leq s$ and that $m(i) \neq 0$ for $i \geq s + 1$. Let $i \geq s + 1$ and assume without loss of generality that $|b_i| \leq |p_i|$. Since the train tracks τ_1, η_1 are complete and contain the curve c_i as a subtrack, both τ_1 and η_1 contain a large proper subbranch of c_i. By the choice of τ_1, for each such large subbranch e of c_i, the train track obtained from τ_1 by a c_i-split at e is *not* contained in the flat strip $E(\tau, \tau')$. Let $\hat{\tau}_1$ be the train track obtained from τ_1 by the split at e which is *not* a c_i-split. Then either we have $\hat{\tau}_1 \in E(\tau, \tau')$, in particular $\hat{\tau}_1$ is complete, or no train track which can be obtained from τ_1 by a split at e is splittable to τ'. In the second case, e is a large branch of τ' by uniqueness of splitting sequences (Lemma 5.1 of [H06a]). Since τ' does not contain any rigid large branch by assumption, both train tracks which can be obtained from τ' by a single split at e are complete. As a consequence, the train track $\hat{\tau}_1$ is complete and is splittable to a complete train track which can be obtained from τ' by a single split at e. The train track $\hat{\tau}_1$ does not carry the simple closed curve c_i.

Since splits at large proper subbranches of the distinct components of c commute, we construct in this way successively in $k - s$ steps from the train tracks τ_1, η_1 complete train tracks τ_2, η_2 which can be obtained from τ_1, η_1 by a splitting sequence of length at most $k - s$. The train tracks τ_2, η_2 are splittable to train tracks τ_2', η_2' which can be obtained from τ', η' by splitting sequences of length at most $k - s$. A minimal component of a geodesic lamination which is carried by both τ_2, η_2 coincides with one of the curves c_i for $i \leq s$. Moreover, the train tracks τ_2, η_2 contain the simple closed curves $c_1, \ldots, c_s$ as embedded subtracks, and their distance in $\mathcal{TT}$ is bounded from above by a universal constant only depending on $d(\tau, \eta)$ and the topological type of S.

By the considerations in Section 4 of [H06b], after reordering of the components c_i and after possibly replacing τ_2, η_2 by their images under a splitting sequence of uniformly bounded length and which are splittable to τ_2', η_2' we may assume that there is a number $u \leq s$ with the following property. For each $i > u$, either the train track τ_2 or the train track η_2 contains a large proper subbranch e of c_i with the property that a c_i-split of τ_2 (or η_2) is not splittable to τ_2' (or η_2'). For every $i \leq u$, the curve c_i is a simple vertex cycle in both τ_2, η_2, i.e. τ_2 and η_2 can be shifted to train tracks which contain c_i as a twist connector. Apply the above construction to the train tracks τ_2, η_2 and the simple closed curves c_i ($u+1 \leq i \leq s$). We obtain train tracks τ_3, η_3 of uniformly bounded distance which are splittable to train tracks τ_3', η_3' obtained from τ_2', η_2' by a splitting sequence of uniformly bounded length and such that a minimal component of a geodesic lamination carried by both τ_3, η_3 coincides with one of the curves $c_1, \ldots, c_u$, and these curves are simple vertex cycles of τ_3, η_3.

Let τ_4 be the train track which can be obtained from τ_3 by a sequence of $\cup_{i=1}^{u} c_i$-splits of maximal length and which is splittable to τ_3'. For $\psi = \varphi_{c_1}^{b_1} \circ \cdots \circ \varphi_{c_u}^{b_u}$, the distance between τ_4 and $\psi(\tau_3)$ is uniformly bounded. Moreover, since the curves

c_i are simple vertex cycles in η_3 and since for each $i \leq u$ the signs of b_i, p_i do *not* coincide, the train track $\psi\eta_3 = \eta_4$ is splittable to η_3. Replace τ_4 by its image τ_5 under a splitting sequence of uniformly bounded length which is splittable to the image of τ_3' under a splitting sequence of uniformly bounded length and which does not carry any of the curves c_i. Then the distance between τ_5 and η_4 is bounded from above by a constant only depending on $d(\tau, \eta)$, and η_4 and τ_5 do not carry any common geodesic lamination. Moreover, η_4 is splittable to η_3' with a splitting sequence which passes through η_3.

Let ν be a complete geodesic lamination which is carried by η_3'. By Lemma 4.3, applied to the train tracks τ_5, η_4 which are splittable to the train tracks τ_3', η_3' in uniformly bounded neighborhoods of τ', η', there is a splitting sequence $\{\alpha(i)\}_{0 \leq i \leq n}$ which connects a train track $\alpha(0)$ in a uniformly bounded neighborhood of τ' to a train track $\alpha(n)$ in a uniformly bounded neighborhood of η' which carries ν, and such that this splitting sequence passes through a uniformly bounded neighborhood of η_4. Now by the considerations in Section 5 of [H06a], this splitting sequence can be chosen to pass through a uniformly bounded neighborhood of η_3 as well. Thus our lemma follows with the Dehn-multisplit φ as above which maps τ, η into a uniformly bounded neighborhood of τ_1, η_1. $\qquad\square$

For a train track $\tau \in \mathcal{V}(\mathcal{T}\mathcal{T})$ containing a subtrack σ, recall from the beginning of this section the definition of a splitting sequence of τ induced by a splitting sequence of σ. We have.

Lemma 4.7. *For every train track $\tau \in \mathcal{V}(\mathcal{T}\mathcal{T})$ which is splittable to a train track η and every recurrent subtrack σ of τ without closed curve components there is a unique train track $\tau' \in E(\tau, \eta)$ with the following properties.*

(1) *There is a recurrent splitting sequence $\{\sigma(i)\}_{0 \leq i \leq p}$ issuing from $\sigma(0) = \sigma$ such that τ' can be obtained from τ by a splitting sequence induced by $\{\sigma(i)\}$.*

(2) *If $\tilde{\tau} \in E(\tau, \eta)$ can be obtained from τ by a recurrent splitting sequence induced by a sequence of splits of σ then $\tilde{\tau}$ is splittable to τ'.*

Proof. Let $\tau \in \mathcal{V}(\mathcal{T}\mathcal{T})$ be a complete train track which is splittable to a train track $\eta \in \mathcal{V}(\mathcal{T}\mathcal{T})$ and let σ be a subtrack of τ. Recall that a σ-*split* of τ is a split of τ at a large proper subbranch of σ with the property that the split track contains σ as a subtrack. Similar to the procedure in the proof of Lemma 4.3 of [H06b] we define a (non-deterministic) finite algorithm which takes τ as input and yields a finite sequence $\{\zeta(i)\}_i \subset E(\tau, \eta)$ of train tracks with $\tau = \zeta(0)$ and such that for each i, either $\zeta(i+1)$ is obtained from $\zeta(i)$ by a single σ-split at a large proper subbranch of σ or $\zeta(i+1)$ is obtained from $\zeta(i)$ by putting a mark on a large proper subbranch of σ. The initial train track $\tau = \zeta(0)$ does not have marked branches.

In the i-th step ($i \geq 1$) the algorithm begins with checking for the existence of an unmarked large proper subbranch e of σ in the train track $\zeta(i-1)$. If there is no such branch then the algorithm stops. Otherwise the algorithm chooses such a large proper subbranch e of σ and proceeds as follows.

If there is a train track $\tilde{\zeta}(i) \in E(\tau, \eta)$ which can be obtained from $\zeta(i-1)$ by a single split at e and which contains σ as a subtrack, then define $\zeta(i)$ to be the

train track $\tilde{\zeta}(i)$ equipped with the markings obtained from the markings of the branches of $\zeta(i-1)$ via the natural identification of the branches of $\zeta(i-1)$ with the branches of $\tilde{\zeta}(i)$. Otherwise define $\zeta(i)$ to be the train track $\zeta(i-1)$ equipped with an additional mark on the branch e (see the proof of Lemma 4.3 of [H06b]).

It follows from the discussion in Section 4 of [H06b] that there is a universal constant $q > 0$ only depending on the topological type of S such that our algorithm stops after at most q steps (see the proof of Lemma 4.3 in [H06b]). It produces a train track $\tilde{\tau}(1) \in E(\tau, \eta)$ which contains σ as a subtrack. For each large branch e of $\tilde{\tau}(1)$ contained in σ, either e is a large branch in σ, i.e. $\tilde{\tau}(1)$ is tight at e, or e is marked. Since splits at distinct large branches of a train track τ commute, the marked train track $\tilde{\tau}(1)$ only depends on τ, η, σ but not on any choices made. Moreover, if $e_1, \ldots, e_s$ are the large branches of σ which correspond to the unmarked large branches of $\tilde{\tau}(1)$ then there is a complete geodesic lamination carried by $\tilde{\tau}(1)$ such that the train track obtained from τ by a successive σ-modification at the branches $e_1, \ldots, e_s$ with respect to λ is splittable to $\tilde{\tau}(1)$ with a splitting sequence of uniformly bounded length.

After reordering, there is a number $\hat{k} \in \{0, \ldots, s\}$ such that for every $i \leq \hat{k}$ the flat strip $E(\tau, \eta)$ contains a train track ξ_i which can be obtained from $\tilde{\tau}(1)$ by a single right or left split at e_i. For each i, there is a unique choice of a right or left split of σ at e_i such that the resulting train track $\tilde{\sigma}_i$ is a subtrack of ξ_i. After reordering, we may assume that there is some $k \leq \hat{k}$ such that for every $i \leq k$ the train track $\tilde{\sigma}_i$ is recurrent but that this is not the case for the train tracks σ_j for $k < j \leq \hat{k}$. Define $\tau(1)$ to be the train track obtained from $\tilde{\tau}(1)$ by splitting $\tilde{\tau}(1)$ with a single split at each of the branches e_i for $i \leq k$ in such a way that the resulting train track is splittable to η and by putting a mark on each of the branches $e_{k+1}, \ldots, e_s$. The marked train track $\tau(1)$ contains a subtrack $\sigma(1)$ which can be obtained from σ by a single split at each of the branches $e_1, \ldots, e_k$. By our construction, the subtrack $\sigma(1)$ of $\tau(1)$ is recurrent (compare [PH92]). Moreover, since splits of a complete train track at distinct large branches commute, the train track $\tau(1)$ can be obtained from τ by a splitting sequence which is induced from a splitting sequence connecting σ to $\sigma(1)$ as defined in the beginning of this section. Note that $\tau(1), \sigma(1)$ only depend on τ, η, σ but not on any choices made.

Repeat the above procedure with the train track $\tau(1)$ and the subtrack $\sigma(1)$ of $\tau(1)$. After finitely many steps we obtain a train track $\tau' \in E(\tau, \eta)$ which clearly satisfies the requirements in the lemma. $\square$

For a complete train track τ and a complete geodesic lamination λ carried by τ define the *flat strip* $E(\tau, \lambda)$ to be the maximal subgraph of $\mathcal{TT}$ whose vertices consist of the set of all complete train tracks which can be obtained from τ by a splitting sequence and which carry λ. For a convenient formulation of the following lemma, we say that a train track η is splittable to a complete geodesic lamination λ if η carries λ. We have.

Lemma 4.8. *Let $\tau \in \mathcal{V}(\mathcal{TT})$ and let $\mathcal{S}(\tau) \subset \mathcal{V}(\mathcal{TT})$ be the set of all complete train tracks which can be obtained from τ by a splitting sequence. Let $E(\tau, \eta)$ be a flat strip where either $\eta \in \mathcal{S}(\tau)$ or η is a complete geodesic lamination carried*

by τ. Then there is a projection $\Pi^1_{E(\tau,\eta)} : \mathcal{S}(\tau) \to E(\tau,\eta)$ such that for every $\zeta \in \mathcal{S}(\tau)$ the train track $\Pi^1_{E(\tau,\eta)}(\zeta)$ is splittable to both ζ, η and such that there is no $\chi \in E(\Pi^1_{E(\tau,\eta)}(\zeta),\eta) - \Pi^1_{E(\tau,\eta)}(\zeta)$ with this property.

Proof. Let $\tau \in \mathcal{V}(\mathcal{T}\mathcal{T})$ be a train track which is splittable to a train track η. Define $\mathcal{S}(\tau) \subset \mathcal{V}(\mathcal{T}\mathcal{T})$ to be the set of all train tracks ζ which can be obtained from τ by a splitting sequence. For the proof of our lemma, we construct by induction on the length m of a splitting sequence connecting τ to η a projection $\Pi^1_{E(\tau,\eta)} = \Pi^1_\eta : \mathcal{S}(\tau) \to E(\tau,\eta)$. If $m = 0$, i.e. if $\tau = \eta$, then we define $\Pi^1_\eta(\zeta) = \tau$ for every $\zeta \in \mathcal{S}(\tau)$. By induction, assume that for some $m \geq 1$ we determined such a projection of $\mathcal{S}(\tau)$ into $E(\tau,\eta)$ for each pair (τ,η) with the property that τ is splittable to η with a splitting sequence of length at most $m - 1$. Let $\{\alpha(i)\}_{0 \leq i \leq m} \subset \mathcal{V}(\mathcal{T}\mathcal{T})$ be a splitting sequence of length m connecting the train track $\tau = \alpha(0)$ to $\eta = \alpha(m)$ and let $\{e_1, \ldots, e_\ell\}$ be the collection of all large branches of τ with the property that the splitting sequence $\{\alpha(i)\}_{0 \leq i \leq m}$ contains a split at e_i. Note that $\ell \geq 1$ since $m \geq 1$. For each i, the choice of a right or left split at e_i is determined by the requirement that the split track carries η.

Let $\zeta \in \mathcal{S}(\tau)$ and assume that there is a large branch $e \in \{e_1, \ldots, e_\ell\}$ with the property that the train track $\tilde{\alpha}(1)$ obtained from τ by a split at e and which is splittable to η is also splittable to ζ. There is then a splitting sequence $\{\tilde{\alpha}(i)\}_{1 \leq i \leq m}$ of length $m - 1$ connecting $\tilde{\alpha}(1)$ to $\tilde{\alpha}(m) = \alpha(m) = \eta$ (compare Lemma 5.1 of [H06a]). The flat strip $E(\tilde{\alpha}(1), \eta)$ is contained in the flat strip $E(\tau,\eta)$, and we have $\zeta \in \mathcal{S}(\tilde{\alpha}(1))$. By induction hypothesis, there is a unique projection point $\Pi^1_{\tilde{\alpha}(1)}(\zeta) \in E(\tilde{\alpha}(1),\eta) \subset E(\tau,\eta)$ with the property that $\Pi^1_{\tilde{\alpha}(1)}(\zeta)$ is splittable to ζ but that this is not the case for any point $\rho \in E(\Pi^1_{\tilde{\alpha}(1)}(\zeta),\eta) - \Pi^1_{\tilde{\alpha}(1)}(\zeta)$. We define $\Pi^1_\eta(\zeta) = \Pi^1_{\tilde{\alpha}(1)}(\zeta)$. Then $\Pi^1_\eta(\zeta)$ is splittable to ζ and this is not the case for any train track in $E(\Pi^1_\eta(\zeta),\eta) - \Pi^1_\eta(\zeta)$. On the other hand, a splitting sequence connecting τ to ζ is unique up to order (see Lemma 5.1 of [H06a] for a detailed discussion of this fact) and therefore if $\xi \in E(\tau,\eta)$ is such that ξ is splittable to ζ and such that a splitting sequence connecting τ to ξ does not contain a split at e, then ξ contains e as a large branch, and there is a train track ξ' which can be obtained from ξ by a split at e and which is splittable to ζ. This just means that the point $\Pi^1_\eta(\zeta)$ does not depend on the above choice of the large branch e.

If none of the train tracks $\xi \in E(\tau,\eta)$ obtained from τ by a split at one of the branches $e_1, \ldots, e_\ell$ is splittable to ζ, then no train track $\beta \in E(\tau,\eta) - \tau$ is splittable to ζ and we define $\Pi^1_\eta(\zeta) = \tau$. This completes our inductive construction of the map $\Pi^1_\eta : \mathcal{S}(\tau) \to E(\tau,\eta)$. Note that we have $\Pi^1_\eta(\zeta) = \Pi^1_\zeta(\eta)$ for all $\zeta, \eta \in \mathcal{S}(\tau)$. Namely, $\Pi^1_\eta(\zeta)$ is splittable to both ζ, η, but this is not the case for any train track which can be obtained from $\Pi^1_\eta(\zeta)$ by a split. This shows the lemma. $\square$

Let F be a *framing* for S (or *marking* in the terminology of [MM99]), i.e. F consists of a pants decomposition P for S and a system of $3g - 3 + m$ *spanning curves*. The framing determines a family $\mathcal{P}(F)$ of train tracks *in standard form for F* ([PH92] and [H06a]). Let $X \subset \mathcal{V}(\mathcal{T}\mathcal{T})$ be the set of all train tracks which can be obtained from a train track in standard form for F by a splitting sequence. By

Proposition 2.2 (in the slightly more precise version which is immediate from the proof given in [H06b]), there is a number $q > 0$ such that the q-neighborhood of X in $\mathcal{TT}$ is all of $\mathcal{TT}$. Thus if we equip X with the restriction of the metric on $\mathcal{TT}$ then the inclusion $X \to \mathcal{TT}$ is a quasi-isometry.

For $\eta \in X$ there is a *unique* train track τ in standard form for F which is splittable to η. Define the *flat strip* $E(F, \eta) = E(\tau, \eta)$ to be the maximal subgraph of $\mathcal{TT}$ whose vertices are the train tracks which can be obtained from τ by a splitting sequence and which are splittable to η. If λ is any complete geodesic lamination then λ is carried by a unique train track τ in standard form for F, and we write $E(F, \lambda) = E(\tau, \lambda)$.

For a complete train track τ and a complete geodesic lamination λ carried by τ define a subset A of $E(\tau, \lambda)$ to be *combinatorially convex* if A can be written in the form $A = \cup_i E(\tau, \sigma_i)$ where for each i we have $\sigma_i \in E(\tau, \sigma_{i+1})$. The next result is the key to a geometric understanding of the train track complex. For its formulation, if $\lambda \in \mathcal{CL}$ is a complete geodesic lamination and if $\{\eta(i)\}_{0 \leq i}$ is an infinite splitting sequence then we say that the sequence *connects* $\eta(0)$ *to* λ (or to a point in the δ-neighborhood of λ for some $\delta > 0$) if $\cap \mathcal{CL}(\eta(i)) = \{\lambda\}$ where as before, we denote for $\eta \in \mathcal{V}(\mathcal{TT})$ by $\mathcal{CL}(\eta)$ the set of all complete geodesic laminations on S which are carried by η (we refer to [M03a] for a discussion of a related construction). We have.

Proposition 4.9. *There is a number $\kappa > 0$ with the following property. Let F be a framing for S and let $X \subset \mathcal{V}(\mathcal{TT})$ be the set of all train complete train tracks which can be obtained from a train track in standard form for F by a splitting sequence. Then for every $\eta \in X \cup \mathcal{CL}$ there is a map $\Pi_{E(F,\eta)} : X \to E(F, \eta)$ such that for every $\zeta \in X$ the following is satisfied.*

1. *There is a splitting sequence connecting a train track τ' in standard form for F to ζ which passes through the κ-neighborhood of $\Pi_{E(F,\eta)}(\zeta)$.*
2. *There is a splitting sequence connecting a point in the κ-neighborhood of ζ to a point in the κ-neighborhood of η which passes through the κ-neighborhood of $\Pi_{E(F,\eta)}(\zeta)$.*
3. $d(\Pi_{E(F,\eta)}(\zeta), \Pi_{E(F,\zeta)}(\eta)) \leq \kappa$ *for all $\eta, \zeta \in X$.*
4. *For all $\lambda, \nu \in \mathcal{CL}$ the set $\Pi_{E(F,\lambda)} E(F, \nu) \subset E(F, \lambda)$ is combinatorially convex.*
5. $d(\Pi_{E(F,\eta)}(\zeta), \Pi_{E(F,\eta)}(\xi)) \leq \kappa d(\zeta, \xi) + \kappa$ *for all $\xi, \zeta \in X$, all $\eta \in X \cup \mathcal{CL}$.*
6. *If $\eta, \zeta \in E(\tau, \lambda)$ for a train track τ in standard form for F which carries the complete geodesic lamination $\lambda \in \mathcal{CL}$ then $\Pi_{E(F,\eta)}(\zeta) = \Pi^1_{E(\tau,\eta)}(\zeta) = \Pi^1_{E(\tau,\zeta)}(\eta)$.*

Proof. Let $\tau \in \mathcal{V}(\mathcal{TT})$ be a complete train track and let $\mathcal{S}(\tau)$ be the set of all complete train tracks which can be obtained from τ by a splitting sequence. Let $\eta \in \mathcal{S}(\tau)$ and let $\zeta \in \mathcal{S}(\tau)$. By Lemma 4.4, via replacing ζ, η by their images under a splitting sequence of uniformly bounded length we may assume that ζ, η do not contain any rigid large branch. Recall from Lemma 4.8 the definition of the "minimal distance" projection $\Pi^1_{E(\tau,\eta)} = \Pi^1_\eta : \mathcal{S}(\tau) \to E(\tau, \eta)$. The projection point $\Pi^1_\eta(\zeta) = \zeta_1$ is uniquely determined by the requirement that $\Pi^1_\eta(\zeta)$ is splittable

to ζ but that no $\chi \in E(\Pi^1_\eta(\zeta), \eta) - \Pi^1_\eta(\zeta)$ has this property. The train track $\Pi^1_\eta(\zeta)$ determines flat strips $E(\Pi^1_\eta(\zeta), \eta), E(\Pi^1_\eta(\zeta), \zeta)$ which intersect in the unique point $\Pi^1_\eta(\zeta)$. Let $\mathcal{E}(\zeta), \mathcal{E}(\eta)$ be the set of large branches e of the train track $\Pi^1_\eta(\zeta)$ with the property that a splitting sequence connecting $\Pi^1_\eta(\zeta)$ to ζ, η contains a split at e. If $\mathcal{E}(\zeta) = \emptyset$ then $\zeta = \Pi^1_\eta(\zeta) \in E(\tau, \eta)$ and we define $\Pi_\eta(\zeta) = \zeta$. Similarly, if $\mathcal{E}(\eta) = \emptyset$ then $\Pi^1_\eta(\zeta) = \eta$, the train track η is splittable to ζ and we define $\Pi_\eta(\zeta) = \eta$.

Now assume that the sets $\mathcal{E}(\zeta), \mathcal{E}(\eta)$ are both non-empty. If $\mathcal{E}(\zeta) \cap \mathcal{E}(\eta) = \emptyset$ then define $\Pi_\eta(\zeta) = \Pi^1_\eta(\zeta)$; note that this is in particular the case if there is some $\lambda \in \mathcal{CL}$ such that $\zeta, \eta \in E(\tau, \lambda)$. Otherwise let $\{e_1, \ldots, e_s\} = \mathcal{E}(\zeta) \cap \mathcal{E}(\eta)$; then by the definition of the map Π^1_η, for each $i \leq s$ a splitting sequence connecting $\Pi^1_\eta(\zeta)$ to ζ contains a right (or left) split at the branch e_i and a splitting sequence connecting $\Pi^1_\eta(\zeta)$ to η contains a left (or right) split at e_i. Let $\zeta_2 \in E(\Pi^1_\eta(\zeta), \zeta), \eta_2 \in E(\Pi^1_\eta(\zeta), \eta)$ be the train track obtained from $\Pi^1_\eta(\zeta) = \zeta_1 = \eta_1$ by a split at each of the large branches $e_1, \ldots, e_s$. Then ζ_2, η_2 contain a common subtrack $\hat{\chi}$ which is obtained from $\Pi^1_\eta(\zeta)$ by a collision at each of the large branches $e_1, \ldots, e_s$, i.e. a split followed by the removal of the diagonal of the split. Note that every geodesic lamination which is carried by both ζ, η is carried by $\hat{\chi}$ and that the number of branches of $\hat{\chi}$ is strictly smaller than the number of branches of ζ, η. Moreover, by a successive application of Lemma 2.3.1 of [PH92], the train track $\hat{\chi}$ is recurrent since the train tracks ζ_2, η_2 are both complete and hence recurrent. Denote by χ the subtrack of ζ_2, η_2 obtained from $\hat{\chi}$ by removing all simple closed curve components of $\tilde{\chi}$.

By Lemma 4.7, there is a recurrent splitting sequence $\{\chi(i)\}_{0 \leq i \leq p} \subset E(\tau, \eta)$ issuing from $\chi = \chi(0)$ which induces a splitting sequence $\{\alpha(i)\}_{0 \leq i \leq k} \subset E(\eta_2, \eta) \subset E(\tau, \eta)$ of maximal length issuing from $\eta_2 = \alpha(0)$. The train track $\chi(p)$ is a recurrent subtrack of $\alpha(k)$, and $\chi(p)$ and $\alpha(k)$ only depend on η_2, η, χ but not on any choices made for the construction of the splitting sequences. Similarly there is a recurrent splitting sequence $\{\tilde{\chi}(i)\}_{0 \leq i \leq q}$ issuing from $\chi = \tilde{\chi}(0)$ which induces a splitting sequence $\{\beta(j)\}_{0 \leq j \leq \ell} \subset E(\zeta_2, \zeta) \subset E(\tau, \zeta)$ of maximal length issuing from ζ_2. The pairs of train tracks $(\chi, \chi(p))$ and $(\tilde{\chi}, \tilde{\chi}(q))$ define flat strips $E(\chi, \chi(p)), E(\chi, \tilde{\chi}(q))$. These flat strips contain all train tracks which can be obtained from χ by a splitting sequence and which are splittable to $\chi(p), \tilde{\chi}(q)$. Apply Lemma 4.8 to these flat strips $E(\chi, \chi(p))$ and $E(\chi, \tilde{\chi}(q))$; this is possible since the construction in the proof of Lemma 4.8 does not use the assumption of completeness for our train tracks. We find a train track $\sigma = \Pi^1_{\chi(p)} \tilde{\chi}(q) = \Pi^1_{\tilde{\chi}(q)} \chi(p) \in E(\chi, \chi(p)) \cap E(\chi, \tilde{\chi}(q))$ with the property that σ is splittable to both $\chi(p), \tilde{\chi}(q)$ but that this is not the case for any train track which can be obtained from σ by a split. By Lemma 4.7, a splitting sequence in $E(\chi, \sigma)$ connecting χ to σ induces splitting sequences $\{\tilde{\zeta}(i)\} \subset E(\tau, \zeta), \{\tilde{\eta}(j)\} \subset E(\tau, \eta)$ of maximal length issuing from $\tilde{\zeta}(0) = \zeta_2, \tilde{\eta}(0) = \eta_2$ and connecting ζ_2, η_2 to train tracks ζ_3, η_3 which contain σ as a subtrack and which are splittable to ζ, η. By Lemma 4.2, the distance between ζ_3, η_3 in $\mathcal{TT}$ is uniformly bounded. Moreover, a geodesic lamination which is carried by both ζ_3, η_3 is carried by the union $\hat{\sigma}$ of σ with the simple closed curve components of $\hat{\chi}$.

Repeat this construction with the train track σ instead of $\Pi^1_\eta(\zeta)$ and the flat strips $E(\sigma, \chi(p))$ and $E(\sigma, \tilde{\chi}(q))$. By Lemma 4.7 we obtain a recurrent splitting

sequences contained in $E(\sigma, \chi(p)), E(\sigma, \tilde{\chi}(q))$ which then induce splitting sequences in $E(\tau, \eta), E(\tau, \zeta)$.

After a uniformly bounded number of steps we obtain a pair of train tracks $\eta', \zeta' \in \mathcal{V}(\mathcal{T}\mathcal{T})$ with the following properties.

(1) The distance between η', ζ' is bounded from above by a universal constant.

(2) ζ', η' contain a common recurrent subtrack β (which possibly is a union of simple closed curves) which carries every geodesic lamination carried by both ζ', η'.

(3) For every large branch e of β, one of the following (not mutually exclusive) possibilities holds.

    a) One of the two train tracks ζ' or η' is not tight at e and for every large proper subbranch e' of e in ζ' (or η') the β-split of ζ' (or η') at e' is not splittable to ζ (or η).

    b) One of the train tracks ζ' (or η') is tight at e and no train track which can be obtained from ζ' (or η') by a single split at e is splittable to ζ (or η).

Let $\beta_0 \subset \beta$ be the union of the simple closed curve components of β. We claim that there is a universal number $r > 0$ with the following properties.

    a) The train tracks ζ', η' can be split with a splitting sequence of length at most r to train tracks $\hat{\zeta}, \hat{\eta}$ which contain a simple multicurve $c \supset \beta_0$ as a subtrack and such that every minimal geodesic lamination carried by both $\hat{\zeta}, \hat{\eta}$ is one of the simple closed curve components of c.

    b) $\hat{\zeta}, \hat{\eta}$ are splittable to train tracks which can be obtained from ζ, η by a splitting sequence of length at most r.

If $\beta_0 = \beta$ then there is nothing to show, so assume that $\beta - \beta_0 = \beta' \neq \emptyset$. Define a β'-*fake collision branch* of the train track ζ' to be a large branch e in ζ' which is a proper subbranch of β' and such that every train track obtained from ζ' by a split at e contains β as a subtrack. If $\tilde{\zeta}$ is obtained from ζ' by any split at e then the number of branches of $\tilde{\zeta}$ contained in β' is strictly smaller than the number of branches of ζ' contained in β'. By the definition of ζ', if e is a β-fake collision branch of ζ' then e is a branch of ζ, i.e. no train track obtained from ζ' by a split at e is splittable to ζ. Namely, otherwise a splitting sequence connecting ζ' to ζ contains a split at e which is necessarily a β'-split. However, this violates our choice of ζ'. Thus via replacing ζ, η by their images under a splitting sequence whose length does not exceed the number q of branches of a complete train track on S, we may assume that the train tracks ζ', η' do not have any β'-fake collision branches.

Let b be a large branch of β' and let e be any large branch of ζ' which is a proper subbranch of b. Note that if ζ' is not tight at b, such a branch always exists. By assumption, e is not a β-fake collision branch. Since β is recurrent by assumption, there is a simple closed multicurve ν which is carried by β and which fills β. The train track ζ' carries a complete extension λ of ν, and the β-split of ζ' at e is necessarily a λ-split. This means that the train track ξ obtained from ζ'

by a β'-split at e is necessarily complete. However, by the construction of ζ', the train track ξ is *not* splittable to ζ. Thus if ξ is obtained from ζ' by say a right split (for convenience of notation), then either the train track $\zeta'(1)$ obtained from ζ' by a left split at e is splittable to ζ or no train track which can be obtained from ζ' by a split at e is splittable to ζ. In the first case we define $\zeta(1) = \zeta$. In the second case the branch e can naturally be viewed as a large branch in ζ. By assumption on ζ, this branch is not rigid and the train track $\zeta(1)$ obtained from ζ by a left split at e is complete. Now the train track $\zeta'(1)$ obtained from ζ' by a left split at e is splittable to $\zeta(1)$ and therefore the train track $\zeta'(1)$ is complete as well.

By construction, a geodesic lamination which is carried by both $\zeta'(1), \eta'$ is carried by the largest recurrent subtrack $\beta(1)$ of β which does *not* contain the branch b. In other words, the number of branches of $\beta(1)$ is strictly smaller than the number of branches of β. Every large proper subbranch a of a large branch of $\beta(1)$ contained in $\zeta'(1)$ is a large branch in ζ' and therefore if a is contained in $\zeta(1)$ then a is not rigid. Thus we can repeat the above construction with the train tracks $\beta(1)$ and $\zeta(1)$. In a number $s \geq 0$ of steps which is bounded from above by the number q of branches of a complete train track we obtain in this way a train track $\zeta'(s)$ containing a recurrent subtrack $\beta(s)$ of β as a subtrack with the additional property that $\zeta'(s)$ does not contain any proper subbranches of large branches of $\beta(s)$. Moreover, a geodesic lamination which is carried by both $\zeta'(s)$ and η is carried by $\beta(s)$, and $\beta(s)$ contains β_0 as a subtrack. The train track $\zeta'(s)$ is splittable to a complete train track $\zeta(s)$ which can be obtained from ζ by a splitting sequence of uniformly bounded length.

If $\beta(s)$ contains components which are not simple closed curves then $\beta(s)$ contains large branches $e_1, \ldots, e_\ell$, and each such branch is tight in $\zeta'(s)$. There is a number $k \leq \ell$ such that for each $i \leq k$ a splitting sequence connecting $\zeta'(s)$ to $\zeta(s)$ contains a split at e_i. Let $\zeta'(s+1)$ be the train track which is splittable to $\zeta(s)$ and which can be obtained from $\zeta'(s)$ by a single split at each of the large branches e_i $(1 \leq i \leq k)$. Then $\zeta'(s + 1)$ contains a subtrack $\beta(s + 1)$ which can be obtained from $\beta(s)$ by a collision at each of the branches $e_1, \ldots, e_\ell$. Every geodesic lamination which is carried by $\zeta'(s + 1)$ and η' is carried by $\beta(s + 1)$.

Repeat the above construction with the train track η' and its subtrack $\beta(s)$. We obtain a subtrack $\beta(s + 1)$ of $\beta(s)$ containing β_0 and a train track $\eta'(s + 1)$ which can be obtained from η' by a splitting sequence of uniformly bounded length. Every large branch e of $\beta(s + 1)$ is tight in both $\zeta'(s)$ and $\eta'(s + 1)$. After finitely many steps we obtain train tracks $\hat{\zeta}, \hat{\eta}$ which satisfy the requirements a),b).

Now the distance between $\hat{\zeta}, \hat{\eta}$ is uniformly bounded, and every minimal geodesic lamination which is carried by both $\hat{\zeta}, \hat{\eta}$ is one of the simple closed curves which form the components of c. Lemma 4.6 applied to $\hat{\zeta}, \hat{\eta}$ then yields a train track $\tilde{\zeta} = \Pi_{E(F,\eta)}(\zeta)$ which satisfy the properties 1),2),3) stated in the proposition, and property 4) follows immediately from our construction. If η is a complete geodesic lamination carried by τ then choose an infinite splitting sequence $\{\tau(i)\}$ issuing from $\tau(0) = \tau$ with $\cap \mathcal{CL}(\tau(i)) = \{\eta\}$. By construction, for every i the train track $\Pi_{E(F,\tau(i))}\zeta$ is splittable to $\Pi_{E(F,\tau(i+1))}\zeta$ and there is a number $i_0 > 0$ such that

$\Pi_{E(F,\tau(i))}\zeta = \Pi_{E(F,\tau(j))}\zeta = \Pi_{E(F,\eta)}\zeta$ for all $i, j \geq i_0$. The train track $\Pi_{E(F,\eta)}\zeta$ satisfies properties 1)-4) in the proposition.

Now let $\zeta, \eta \in X$ be arbitrary; then there are unique train tracks τ, τ' in standard form for F so that τ is splittable to ζ and τ' is splittable to η. Let $\mathcal{M}(\tau), \mathcal{M}(\tau')$ be the set of all measured geodesic laminations carried by τ, τ'. If $\mathcal{M}(\tau) \cap \mathcal{M}(\tau') = \{0\}$ then we define $\Pi_{E(F,\zeta)}(\eta) = \tau$ and $\Pi_{E(F,\eta)}(\zeta) = \tau'$. On the other hand, if $\mathcal{M}(\tau) \cap \mathcal{M}(\tau') \neq \{0\}$ then τ, τ' contain a common maximal recurrent subtrack χ which carries the support of every lamination in $\mathcal{M}(\tau) \cap \mathcal{M}(\tau')$ (see Lemma 4.4 of [H06b]). We apply our above construction to the train tracks τ, τ', which are splittable to ζ, η and the common subtrack χ of ζ, η and extend in this way the maps $\Pi_{E(F,\zeta)}, \Pi_{E(F,\eta)}$ to all of X in such a way that properties 1)-4) stated in the lemma are satisfied.

To show property 5) in the proposition, let $\sigma, \zeta_1, \zeta_2 \in X$ and let τ be a train track in standard form for F which is splittable to σ. Write $\chi_i = \Pi_{E(F,\sigma)}(\zeta_i)$ $(i = 1, 2)$ and let $\chi = \Pi_{E(F,\chi_1)}(\chi_2)$. Since χ_1, χ_2 are both contained in the same flat strip $E(\tau, \sigma)$, by our above construction the train track χ is splittable to both χ_1, χ_2 and there are disjoint sets $\mathcal{E}_1, \mathcal{E}_2$ of large branches of χ such that a splitting sequence connecting χ to χ_i contains a split at a large branch e if and only if $e \in \mathcal{E}_i$ $(i = 1, 2)$. Let $\ell_i \geq 0$ be the length of a splitting sequence connecting χ to χ_1, χ_2 $(i = 1, 2)$ and let $\ell = \max\{\ell_1, \ell_2\}$. Then the distance between χ_1, χ_2 is not bigger than 2ℓ.

Let τ_1, τ_2 be train tracks in standard form for F such that τ_i is splittable to ζ_i $(i = 1, 2)$. It follows from our above construction that there is a universal constant $\kappa > 0$ and there is a splitting sequence connecting τ_i to ζ_i which passes through the κ-neighborhood of χ_i $(i = 1, 2)$. Since by Proposition 2.1 splitting sequences are uniform quasi-geodesics, we conclude that the distance between ζ_1 and ζ_2 is bounded from below by $c\ell$ for a universal constant $c > 0$. This finishes the proof of the lemma. $\qquad\square$

For every $\tau \in \mathcal{V}(\mathcal{TT})$ and every complete geodesic lamination λ carried by τ, the flat strip $E(\tau, \lambda) \subset \mathcal{TT}$ is connected and hence can be equipped with the intrinsic metric d_λ. The following observation is a consequence of Lemma 4.9.

Corollary 4.10. *There is a number $c > 0$ with the following property. For every $\tau \in \mathcal{V}(\mathcal{TT})$ and every complete geodesic lamination λ carried by τ, the natural inclusion $(E(\tau, \lambda), d_\lambda) \to \mathcal{TT}$ is a c-quasi-isometric embedding.*

Proof. Since splitting sequences are uniform quasi-geodesics in $\mathcal{TT}$ which define geodesics in $E(\tau, \lambda)$ (see Lemma 5.1 of [H06a]), we only have to show the existence of a number $c > 0$ with the following property. Let $\tau \in \mathcal{V}(\mathcal{TT})$, let $\lambda \in \mathcal{CL}$ and let $\sigma_1, \sigma_2 \in E(\tau, \lambda)$. Using the notations from Lemma 4.8, let $\nu = \Pi^1_{E(\tau,\sigma_1)}(\sigma_2) = \Pi^1_{E(\tau,\sigma_2)}(\sigma_1) \in E(\tau, \lambda)$ be the unique train track which is splittable to σ_1, σ_2 and such that no train track which can be obtained from ν by a single split has this property. Let $\ell_1, \ell_2 \geq 0$ be the length of a splitting sequence connecting ν to σ_1, σ_2; then $d(\sigma_1, \sigma_2) \geq (\ell_1 + \ell_2)/c - c$ where d is the distance of $\mathcal{TT}$.

By Proposition 4.9, there is a splitting sequence connecting a point in the κ-neighborhood of σ_1 to a point contained in the κ-neighborhood of σ_2 which passes through the κ-neighborhood of ν. Now by Proposition 2.1, splitting sequences are L-quasi-geodesics in $\mathcal{T}\mathcal{T}$ for a universal number $L > 1$ and therefore the distance between σ_1, σ_2 is not smaller than $d(\sigma_1, \nu)/L - L - 2\kappa + d(\sigma_2, \nu)/L - L - 2\kappa$. On the other hand, the distance in $E(\tau, \lambda)$ between σ_1 and σ_2 is not bigger than $Ld(\sigma_1, \nu) + Ld(\sigma_2, \nu) + 2L$ from which the corollary follows. $\square$

5. The large-scale geometry of flat strips

In this section we have a closer look at the geometry of flat strips in $\mathcal{T}\mathcal{T}$. In particular, we compute the *asymptotic cone* of such a flat strip. Here a flat strip $E(\tau, \lambda)$ is determined by a complete train track $\tau \in \mathcal{V}(\mathcal{T}\mathcal{T})$ and a complete geodesic lamination λ carried by τ, and it is the maximal subgraph of $\mathcal{T}\mathcal{T}$ whose set of vertices consists of all train tracks $\sigma \in \mathcal{V}(\mathcal{T}\mathcal{T})$ which carry λ and can be obtained from τ by a splitting sequence. The flat strip $E(\tau, \lambda)$ is connected and can be equipped with the intrinsic path-metric d_λ. By Corollary 4.10, there is a number $c > 1$ not depending on τ, λ such that the natural inclusion $(E(\tau, \lambda), d_\lambda) \to \mathcal{T}\mathcal{T}$ is a c-quasi-isometric embedding.

By Lemma 5.1 of [H06a], there is an isometry of $(E(\tau, \lambda), d_\lambda)$ onto a connected *cubical graph* in $\mathbb{R}^q$ where $q > 0$ is the number of branches of the complete train track τ. Such an isometry Φ is determined by the choice of a point $\Phi(\tau) \in \mathbb{Z}^q$ and the choice of a numbering of the branches of τ and has the following property. Let $x_1, \ldots, x_q$ be the standard basis of $\mathbb{R}^q$. If $\sigma \in E(\tau, \lambda)$ is a complete train track then the numbering of the branches of τ induces a numbering of the branches of σ. If the branch i in σ is large, then the train track $\sigma' \in E(\tau, \lambda)$ obtained from σ by a single split at i is mapped by Φ to $\Phi(\sigma) + x_i$. We call such an isometry Φ of $E(\tau, \lambda)$ onto the cubical graph $\Phi(E(\tau, \lambda)) \subset \mathbb{R}^q$ *standard*.

To obtain an understanding of the intrinsic geometry of the graph $\Phi(E(\tau, \lambda))$, consider for the moment an arbitrary connected *cubical complex* K as defined on p.111-112 in [BH99] which is isometrically embedded in the euclidean space $\mathbb{R}^q$. Such a complex K is a closed subset of $\mathbb{R}^q$ which is the union of an at most countable number of *standard cubes*, i.e. subsets of $\mathbb{R}^q$ which are isometric to a cube $[0, 1]^\ell$ for some $\ell \leq q$. The intersection of any two such cubes is either empty or is again a standard cube. If the vertices of K are points in the standard integer lattice $\mathbb{Z}^q$ then we call the cubical complex *standard*.

Following Definition II.5.15 of [BH99], call an abstract *simplicial* complex L with vertex set V a *flag complex* if every finite subset A of V with the property that any two distinct points in A are connected by an edge spans a simplex. By Theorem II.5.4 and Theorem II.5.18 of [BH99], a standard cubical complex K in $\mathbb{R}^q$ has *non-positive curvature* if and only if for every vertex v of K the link complex $L(v)$ of v is a flag complex. Moreover, $L(v)$ is a flag complex if and only if $L(v)$ equipped with the path metric induced from the round metric on the $(q-1)$-dimensional unit sphere in $\mathbb{R}^q$ is a Cat(1)-space.

Let again $\tau \in \mathcal{V}(\mathcal{TT})$ and let $\lambda \in \mathcal{CL}$ be a complete geodesic lamination carried by τ. Let Φ be a standard isometry of the flat strip $E(\tau, \lambda) \subset \mathcal{TT}$ equipped with its intrinsic metric d_λ onto an embedded standard cubical graph in $\mathbb{R}^q$. Define the *maximal extension* of the graph $E(\tau, \lambda)$ to be the maximal cubical subcomplex $C(\tau, \lambda)$ of $\mathbb{R}^q$ whose one-skeleton equals $\Phi(E(\tau, \lambda))$. This complex is uniquely determined by $E(\tau, \lambda)$ up to permutations of vectors from the standard basis of $\mathbb{R}^q$ and translation by a vector in $\mathbb{Z}^q$. In particular, it is uniquely determined by $E(\tau, \lambda)$ up to cubical isometry. The two-skeleton $C^2(\tau, \lambda)$ of the complex $C(\tau, \lambda)$ is determined as follows. Let $x_1, \ldots, x_q$ be the standard basis of $\mathbb{R}^q$; then for some $v \in \mathbb{Z}^q$ the two-dimensional cube in $\mathbb{R}^q$ with vertices $v, v + x_i, v + x_j, v + x_i + x_j$ is a cube in $C^2(\tau, \lambda)$ if and only if each of its four sides is contained in $\Phi(E(\tau, \lambda))$. For $k \geq 3$ the k-skeleton $C^k(\tau, \lambda)$ of $C(\tau, \lambda)$ is constructed in the same way by induction: If Q is any k-cube in $\mathbb{R}^q$ all of whose sides are contained in $C^{k-1}(\tau, \lambda)$ then we require that Q is contained in $C^k(\tau, \lambda)$. We have.

Lemma 5.1. *The maximal extension $C(\tau, \lambda)$ of a flat strip $E(\tau, \lambda) \subset \mathcal{TT}$ is a complete $\mathrm{Cat}(0)$-space.*

Proof. We show first that the maximal extension $C(\tau, \lambda)$ of the graph $E(\tau, \lambda)$ is of non-positive curvature. For this we have to show that the link complex $L(v)$ of every vertex v of $C(\tau, \lambda)$ is a flag complex.

Let $x_1, \ldots, x_q$ be the standard basis of $\mathbb{R}^q$. By construction of the map Φ (see above and compare Lemma 5.1 of [H06a]), if $v \in \Phi(E(\tau, \lambda)) \cap \mathbb{Z}^q$ and if $1 \leq i \leq q$ is such that $v + x_i \in \Phi(E(\tau, \lambda))$ then the line segment in $\mathbb{R}^q$ connecting v to $v + x_i$ is contained in $\Phi(E(\tau, \lambda))$ as well.

Now assume that $v \in \Phi(E(\tau, \lambda)) \cap \mathbb{Z}^q$ and that $1 \leq i < j \leq q$ are such that $v, v + x_i, v + x_j \in \Phi(E(\tau, \lambda))$. Let $\sigma \in E(\tau, \lambda)$ be such that $\Phi(\sigma) = v$. By construction of the map Φ, σ is a complete train track equipped with a numbering of its branches such that the branches with numbers i, j are large. The train track σ' obtained from σ by the λ-split at the branch i is mapped by Φ to $v + x_i$, and the train track σ'' obtained from σ by the λ-split at the branch j is mapped by Φ to $v + x_j$. By definition, the line segments connecting v to $v + x_i, v + x_j$ are contained in $\Phi(E(\tau, \lambda))$. Since λ-splits at distinct large branches in σ commute, the branch j in the train track σ' (with respect to the numbering inherited from the numbering of the branches of σ) is large and the train track $\tilde{\sigma}$ obtained from σ' by the λ-split at j is mapped by Φ to $v + x_i + x_j$. The same consideration also shows that $\tilde{\sigma}$ can be obtained from σ'' by the λ-split at the branch i. This implies that the boundary of the two-dimensional cube Q in $\mathbb{R}^q$ with vertices $v, v + x_i, v + x_j, v + x_i + x_j$ is contained in $C(\tau, \lambda)$ and hence the cube Q is contained in $C(\tau, \lambda)$ as well. In other words, if v is a vertex in $\Phi(E(\tau, \lambda))$ and if the points x_i, x_j (viewed as directions in the unit sphere at v) are contained in the link complex $L(v)$ of v then the spherical edge connecting x_i to x_j is contained in $L(v)$ as well. The obvious extension of this discussion to more than two of the standard basis vectors $x_1, \ldots, x_q$ shows the following. If v is a vertex in $E(\tau, \lambda)$, if $k \geq 1$ and if $1 \leq i_1 < \cdots < i_k \leq q$ are such that $v + x_{i_j} \in \Phi(E(\tau, \lambda))$ for every $j \leq k$ then the k-dimensional standard cube $Q \subset \mathbb{R}^q$ which is determined by the vertices $v, v + x_{i_j}$ $(j \leq k)$ is contained in $C(\tau, \lambda)$. Thus the vertices in the link complex $L(v)$ of v defined by the directions

x_{i_j} are pairwise joined by edges, and their closed convex hull is a spherical simplex contained in $L(v)$.

Let $i, j \leq q$ be such that $v, v + x_i, v - x_j$ are vertices in $\Phi(E(\tau, \lambda))$ and let $\sigma \in E(\tau, \lambda)$ be such that $\Phi(\sigma) = v$. Then the branch i in σ is large and the branch j is small, in particular we have $i \neq j$. The small branch j can be collapsed in σ, and the branch with number j in the train track σ' obtained from σ by this collapse is large. There are now two possibilities. The first possibility is that the branch i in σ' is a large branch which is equivalent to saying that the branches i and j in σ are not incident on a common switch. Then the train track σ'' obtained from σ' by a λ-split at i is mapped by Φ to $v + x_i - x_j$ and hence as above, the vertices $v - x_j, v, v + x_i - x_j, v + x_i$ are contained in $\Phi(E(\tau, \lambda))$ and span a two-dimensional cube in $C(\tau, \lambda)$. However, if the branch i in σ' is *not* large then the two-dimensional cube with vertices $v - x_j, v, v - x_j + x_i, v + x_i$ is *not* contained in $C(\tau, \lambda)$ and the vertices $x_i, -x_j$ in the link complex $L(v)$ are *not* connected by an edge. As a consequence, if the points $v, v - x_j, v + x_{i_1}, \ldots, v + x_{i_k}$ are vertices in $\Phi(E(\tau, \lambda))$ and if $\sigma' \in E(\tau, \lambda)$ is mapped by Φ to $v - x_j$ then the spherical edge in the space of directions at v connecting the vertices $-x_j, x_{i_\ell}$ ($\ell \leq k$) is contained in $L(v)$ if and only if the branch i_ℓ in σ' is large. This shows the following. Assume that the edges in $\Phi(E(\tau, \eta))$ which connect v to $v - x_j, v + x_{i_\ell}$ ($\ell \leq k$), viewed as vertices in the link complex $L(v)$ of v, are pairwise connected in $L(v)$ by an edge. Then the branches $j, i_1, \ldots, i_k$ in σ' are all large and the $k + 1$-dimensional standard cube in $\mathbb{R}^q$ determined by these vertices is contained in $C(\tau, \lambda)$. As a consequence, the k-dimensional spherical simplex in $L(v)$ spanned by the vertices $-x_j, x_{i_1}, \ldots, x_{i_k}$ is contained in $L(v)$ if and only if any two of these vertices are connected by an edge.

Now consider a triple of vertices in $\Phi(E(\tau, \lambda))$ of the form $v, v - x_i, v - x_j$ for some $v \in \mathbb{Z}^q$. If $\sigma, \sigma_i, \sigma_j$ are the preimages of $v, v - x_i, v - x_j$ under the map Φ then σ_i, σ_j is obtained from σ by a collapse of the small branch i, j. However, both train tracks σ_i, σ_j can be obtained from the same train track τ by a splitting sequence. Since a splitting sequence connecting τ to σ is unique up to the order of the splits (see the discussion in the proof of Lemma 5.1 of [H06a]), the branch i in σ_j is a small branch and the train track η which can be obtained from σ_j by a collapse of the branch i is splittable to both σ_i, σ_j. In particular, as before the two-dimensional cube in $\mathbb{R}^q$ with vertices $v, v - x_i, v - x_j, v - x_i - x_j$ is contained in $C(\tau, \lambda)$. The obvious extension of this consideration to more than two of the standard basis vectors $x_1, \ldots, x_q$ shows that if $x_{i_1}, \ldots, x_{i_k}$ are such that $v, v - x_{i_1}, \ldots, v - x_{i_k}$ are contained in $C(\tau, \lambda)$ then the same is true for the k-dimensional cube determined by these points. Together with our above discussion we conclude the following. Let $v = \Phi(\sigma)$ be a vertex of $C(\tau, \lambda)$ and let $x_{i_1}, \ldots, x_{i_k}, x_{j_1}, \ldots, x_{j_\ell}$ be such that for each $p \leq k, q \leq \ell$ the points $v - x_{i_p}, v + x_{j_q}$ are vertices of $C(\tau, \lambda)$. If any two of the directions x_{i_p}, x_{j_q} defined by these vertices are connected in $L(v)$ by an edge then for every $p \leq k, q \leq \ell$ the branch j_q is large in σ as well as in the train track obtained from σ by a single collapse at i_p. Equivalently, the small branch i_p in σ does not have a switch in common with the large branch j_q. However, we observed above that in this case the train track ν obtained from σ by a collapse of each of the branches $i_1, \ldots, i_p$ contains the large branches $j_1, \ldots, j_q$ and the cube of dimension

$p + q$ determined by the vertices $v, v - x_{i_p}, v + x_{i_q}$ for $p \leq k, q \leq \ell$ is contained in $C(\tau, \lambda)$. This shows that the link complex $L(v)$ of v is a flag complex as claimed.

Since the map $\Phi : E(\tau, \lambda) \to \mathbb{R}^q$ is proper by construction, the cubical complex $C(\tau, \lambda)$ is a complete geodesic metric space. Therefore to show that $C(\tau, \lambda)$ is indeed a complete CAT(0)-space it is now enough to establish that $C(\tau, \lambda)$ is simply connected. This in turn follows if we can show that every closed edge-path in $C(\tau, \lambda)$ which begins and ends at $\Phi(\tau)$ is contractible. Note that via the isometry Φ such an edge-path can be identified with a path in the graph $E(\tau, \lambda)$.

To show that this is indeed the case we proceed by induction on the combinatorial length of the path. If this length vanishes then the claim is trivial, so assume that the claim holds for all flat strips $E(\sigma, \zeta)$ where $\sigma \in \mathcal{V(TT)}$ and where $\zeta \in \mathcal{CL}$ is carried by σ and all closed edge-paths of combinatorial length at most $m - 1$ for some $m \geq 0$ which begin and end at σ. Let $\gamma : [0, m] \to E(\tau, \lambda)$ be a closed edge-path of combinatorial length m beginning and ending at τ. Then $\gamma(1)$ is a train track which can be obtained from $\gamma(0) = \tau$ by a single split at a large branch e. The branch e_0 in $\gamma(1)$ corresponding to e is small. Assume without loss of generality that the standard isometric embedding $\Phi : E(\tau, \lambda) \to \mathbb{R}^q$ satisfies $\Phi(\tau) = 0$ and $\Phi(\gamma(1)) = x_1$ where $x_1, \ldots, x_q$ is the standard basis of $\mathbb{R}^q$. Let $\alpha_1, \ldots, \alpha_q$ be the basis of $(\mathbb{R}^q)^*$ which is dual to $x_1, \ldots, x_q$. Then $\alpha_1(\gamma(1)) = 1$ and if $\alpha_1(\gamma(i)) > 0$ for some $i \in \{1, \ldots, m - 1\}$ then $\gamma(i) \in E(\gamma(1), \lambda)$. In particular, if $\alpha_1(\gamma(i)) > 0$ for every $i \in \{1, \ldots, m - 1\}$ then $\gamma[1, m - 1]$ is a loop in $E(\gamma(1), \lambda)$ beginning and ending at $\gamma(1)$ of combinatorial length $m - 2$. By our induction hypothesis, this loop is contractible in $C(\gamma(1), \lambda) \subset C(\tau, \lambda)$ and therefore γ is contractible in $C(\tau, \lambda)$.

Otherwise there is a first number $i_0 \in \{2, \ldots, m - 1\}$ such that $\alpha_1(\gamma(i_0)) = 0$. Then $\gamma(i_0)$ can be obtained from τ by some splitting sequence not containing a split at e. In particular, $\gamma(i_0)$ contains the large branch e and $\gamma(i_0 - 1)$ is obtained from $\gamma(i_0)$ by a single λ-split at e. Let i_1 be the minimum of all numbers $i > i_0$ such that $\alpha_1(\gamma(i)) > 0$; if there is not such i then define $i_1 = m$. Note that $\gamma(i_1) = \gamma(i_0 - 1)$ if $i_1 = i_0 + 1$.

If $i_1 < m$ then $\gamma(i_1)$ can be obtained from $\gamma(i_1 - 1)$ by a single λ-split at e. Put $\tilde{\gamma}(j) = \gamma(j)$ for $j \leq i_0 - 1$, $\tilde{\gamma}(j) = \gamma(j + 2)$ for $j \geq i_1 - 1$ and for $i_0 \leq j \leq i_1 - 2$ define $\tilde{\gamma}(j)$ to be the train track which can be obtained from $\gamma(j + 1)$ by a single λ-split at e. Then the assignment $j \to \tilde{\gamma}(j)$ $(i_0 - 1 \leq j \leq i_1 - 2)$ determines an edge path contained in $C(\gamma(1), \lambda)$ connecting $\tilde{\gamma}(i_0 - 1) = \gamma(i_0 - 1)$ to $\tilde{\gamma}(i_1 - 2) = \gamma(i_1)$. For every $j \in \{i_0, \ldots, i_1 - 2\}$ the vertices $\gamma(j), \gamma(j + 1), \tilde{\gamma}(j - 1), \tilde{\gamma}(j)$ are the vertices of a 2-dimensional cube embedded in $C(\tau, \lambda)$. Thus by the definition of the maximal extension $C(\tau, \lambda)$ of the flat strip $E(\tau, \lambda)$, this edge path is homotopic with fixed endpoints to the edge path $\gamma[i_0 - 1, i_1]$. Then $\tilde{\gamma}$ is homotopic to γ with fixed endpoints. Since the combinatorial length of $\tilde{\gamma}$ equals $m - 2$, by induction hypothesis the edge-path $\tilde{\gamma}$ is contractible in $C(\tau, \lambda)$ and hence the same holds true for the edge-path γ.

If $i_1 = m$ then we let $\tilde{\gamma}(j) = \gamma(j)$ for $j \leq i_0 - 1$ and for $i_0 \leq j \leq m - 1$ define $\tilde{\gamma}(j)$ to be the unique train track which can be obtained from $\gamma(j + 1)$ by a single split at e. Also put $\tilde{\gamma}(m) = \tau$. By the above consideration, the loop $\tilde{\gamma}$ is homotopic to γ. On the other hand, the curve $\tilde{\gamma}[1, m - 1]$ is a loop contained in

$E(\gamma(1), \lambda)$ of combinatorial length $m - 2 < m$ and hence this loop is contractible in $C(\gamma(1), \lambda) \subset C(\tau, \lambda)$ by induction hypothesis. But then $\tilde{\gamma}$ is contractible in $C(\tau, \lambda)$ and hence the same holds true for γ. This completes the proof of our lemma.   □

Let $\tau \in \mathcal{V}(\mathcal{TT})$ be a complete train track which is splittable to a train track $\eta \in \mathcal{V}(\mathcal{TT})$. Then the flat strip $E(\tau, \eta)$ is defined. The proof of Lemma 5.1 can be applied without modification to $E(\tau, \eta)$ and shows that $E(\tau, \eta)$ admits a natural CAT(0)-cubical extension $C(\tau, \eta)$. For every complete geodesic lamination λ which is carried by η, this extension is naturally a subspace of the extension $C(\tau, \lambda)$ of the flat strip $E(\tau, \lambda)$. Moreover, the maximal extension $C(\eta, \lambda)$ of the flat strip $E(\eta, \lambda)$ is a subspace of $C(\tau, \lambda)$ as well. We have.

Lemma 5.2. *Let $\tau \in \mathcal{V}(\mathcal{TT})$ and let λ be a complete geodesic lamination which is carried by τ. Let $\Phi : E(\tau, \lambda) \to \mathbb{R}^q$ be a standard isometric embedding and let d be the intrinsic metric on the cubical complex $C(\tau, \lambda)$.*

(1) *For every vertex $\eta \in E(\tau, \lambda)$ the maximal extensions $C(\tau, \eta), C(\eta, \lambda)$ are convex subspaces of $C(\tau, \lambda)$, and $d(C(\eta, \lambda), \Phi(\tau)) = d(\Phi(\eta), \Phi(\tau))$.*
(2) *The restriction of every coordinate function α^i of $\mathbb{R}^q$ to a geodesic ray $\gamma : [0, \infty) \to C(\tau, \lambda)$ issuing from $\gamma(0) = \Phi(\tau)$ is non-increasing.*

Proof. Let $\tau \in \mathcal{V}(\mathcal{TT})$, let λ be a complete geodesic lamination carried by τ and let $\Phi : E(\tau, \lambda) \to \mathbb{R}^q$ be a standard isometric embedding with $\Phi(\tau) = 0$.

Let $\eta \in \mathcal{V}(\mathcal{TT})$ be a vertex in $E(\tau, \lambda)$. We show first that $C(\eta, \lambda) \subset C(\tau, \lambda)$ is convex. For this note that if X is a any complete Cat(0)-space and if $A \subset X$ is a closed convex subset, then A is a complete Cat(0)-space. Moreover, every closed convex subset $B \subset A$ is convex in X. Using this fact inductively, we conclude that it is enough to show that $C(\eta, \lambda)$ is convex subspace of $C(\tau, \lambda)$ for every complete train track $\eta \in E(\tau, \lambda)$ which can be obtained from τ by a single split.

Let as before $\alpha_1, \ldots, \alpha_q$ be the basis of $(\mathbb{R}^q)^*$ which is dual to the standard basis $x_1, \ldots, x_q$ of $\mathbb{R}^q$. Assume without loss of generality that $\Phi(\eta) = x_1$; then a point $z \in C(\tau, \lambda)$ is contained in $C(\eta, \lambda)$ if and only if $\alpha_1(z) \geq 1$. By the construction of the map Φ, for every point $(z_1, z_2, \ldots, z_q) \in C(\tau, \lambda)$ with $z_1 < 1$ the point $(1, z_2, \ldots, z_q)$ is contained in $C(\tau, \lambda)$ as well. Thus the cubical complex $C(\tau, \lambda) \subset \mathbb{R}^q$ is invariant under the natural distance-non-increasing shortest distance projection ρ of $\mathbb{R}^q$ onto the closed half-space $\{\alpha_1 \geq 1\}$ which maps a point $z = (z_1, \ldots, z_q)$ with $z_1 < 1$ to $\rho(z) = (1, z_2, \ldots, z_q)$. Since $C(\tau, \lambda)$ is equipped with the complete path metric induced from the euclidean metric, the restriction to $C(\tau, \lambda)$ of the retraction ρ is distance non-increasing as well. Since the image of $C(\tau, \lambda)$ under ρ is just the cubical complex $C(\eta, \lambda)$, the subcomplex $C(\eta, \lambda) \subset C(\tau, \lambda)$ is convex. The same argument also shows that $d(C(\eta, \lambda), \Phi(\tau)) = d(\Phi(\eta), \Phi(\tau))$.

To show that $C(\tau, \eta) \subset C(\tau, \lambda)$ is convex for every complete train track $\eta \in E(\tau, \lambda)$ we argue in the same way. Namely, by Lemma 5.1 and its proof, for each $\eta \in E(\tau, \lambda)$ the space $C(\tau, \eta)$ is a complete Cat(0)-space. Now if $X_1 \subset X_2 \subset \ldots$ is a nested sequence of complete locally compact Cat(0)-spaces with complete locally compact Cat(0)-union $\cup_i X_i = X$ and if for each i the space X_i is a convex subspace of X_{i+1} then for each i, the space X_i is convex in X as well. Thus as above, it is

enough to show that $C(\tau, \eta) \subset C(\tau, \zeta)$ is convex whenever ζ can be obtained from η by a single split at a large branch e.

Assume without loss of generality that the number of e in η with respect to the numbering of the branches of τ defining our standard isometry Φ equals one. Then using our above notation we have $\alpha_1(\zeta) = \alpha_1(\eta) + 1, \alpha_i(\zeta) = \alpha_i(\eta)$ for $i \geq 2$ and therefore the distance between $C(\tau, \eta)$ and $\Phi(\zeta)$ with respect to the restriction of the euclidean metric on $\mathbb{R}^q$ equals one. Moreover, by construction of the map Φ, for every point $(z_1, z_2, \ldots, z_q) \in C(\tau, \zeta)$ with $z_1 > \alpha_1(\eta)$ the point $(\alpha_1(\eta), z_2, \ldots, z_q)$ is contained in $C(\tau, \eta)$. As above, this implies that $C(\tau, \eta) \subset C(\tau, \zeta)$ is convex and completes the proof of the first part of the lemma.

For the second part of the lemma it is enough to show that for every geodesic $c : [0, b] \to C(\tau, \lambda)$ issuing from $c(0) = \Phi(\tau)$ and every $i \geq 0$ the function $t \to \alpha_i(c(t))$ is non-decreasing. Namely, in this case there is an edge-path $\rho : [0, r] \to C(\tau, \lambda)$ in the cubical complex $C(\tau, \lambda)$ whose Hausdorff distance to $c[0, b]$ is uniformly bounded and which has the same property. However by construction, the successive vertices met by such an edge-path are the image under Φ of a splitting sequence in $E(\tau, \lambda)$. However, by our above consideration, for every $i \geq 1$ and every $s \in \mathbb{R}$ the set $\{\alpha_i \geq s\} \cap C(\tau, \lambda)$ is convex in $C(\tau, \lambda)$ and hence the function $t \to \alpha_i(c(t))$ is necessarily non-decreasing. The lemma follows. $\square$

Next we justify the notion "flat strip" for the sets $E(\tau, \lambda)$. Namely, recall that the natural inclusion $(E(\tau, \lambda), d_\lambda) \to \mathcal{TT}$ is a quasi-isometric embedding and that $C(\tau, \lambda)$ is quasi-isometric to its one-skeleton $E(\tau, \lambda)$. The next lemma shows that the path-metric on $C(\tau, \lambda)$ is quasi-isometric to the restriction of the euclidean metric.

Lemma 5.3.   (1) *There is a number $c > 0$ such that for every flat strip $E(\tau, \lambda)$ the inclusion $C(\tau, \lambda) \to \mathbb{R}^q$ is a c-quasi-isometric embedding.*
   (2) *If $\lambda \in \mathcal{CL}$ is carried by τ and if $\zeta, \eta \in E(\tau, \lambda)$ then there is a unique train track $\Theta(\zeta, \eta) \in E(\tau, \lambda)$ such that $\zeta, \eta \in E(\tau, \Theta(\zeta, \eta))$ and that $\Theta(\zeta, \eta) \in E(\tau, \xi)$ for every train track $\xi \in \mathcal{V}(\mathcal{TT})$ which can be obtained from both ζ, η by a splitting sequence.*

Proof. Let $E(\tau, \lambda)$ be a flat strip and let $\Phi : E(\tau, \lambda) \to \mathbb{R}^q$ be a standard isometric embedding. Since the inclusion $\iota : \Phi(E(\tau, \lambda)) \to \mathbb{R}^q$ is a one-Lipschitz map, for the first part of the lemma it is enough to show the existence of a universal constant $c > 0$ such that for all $\sigma, \eta \in E(\tau, \lambda)$ we have $d_\lambda(\sigma, \eta) \leq c\|\Phi(\sigma) - \Phi(\eta)\|$ where $\| \|$ is the euclidean norm on $\mathbb{R}^q$ and d_λ is the intrinsic path metric on $E(\tau, \lambda)$.

For this let $\sigma, \eta \in E(\tau, \lambda)$. Using the notations from Lemma 4.8, write $\zeta = \Pi^1_{E(\tau, \sigma)}(\eta) = \Pi^1_{E(\tau, \eta)}(\sigma)$. By construction, ζ is splittable to both σ, η and this is not the case for any train track which can be obtained from ζ by a single split. Via replacing Φ by the composition of Φ with a translation by a vector in $\mathbb{Z}^q$ we may assume that $\Phi(\zeta) = 0$.

We claim that up to a permutation of the standard basis of $\mathbb{R}^q$, there is a number $\ell \geq 1$ such that for the standard direct orthogonal decomposition $\mathbb{R}^q = \mathbb{R}^\ell \oplus \mathbb{R}^{q-\ell}$ we have $\Phi(\sigma) \in \mathbb{R}^\ell$ and $\Phi(\eta) \in \mathbb{R}^{q-\ell}$. Namely, by the choice of the train track ζ

and the fact that σ, η both carry the complete geodesic lamination λ, the set of large branches $\mathcal{E}(\zeta)$ of ζ can be partitioned into disjoint subsets $\mathcal{E}^+, \mathcal{E}^-$ such that a splitting sequence connecting ζ to σ does not contain any split at a large branch branch $e \in \mathcal{E}^+$ and a splitting sequence connecting ζ to η does not contain any split at a large branch $e \in \mathcal{E}^-$.

Following [PH92], we call a trainpath $\rho : [0, m] \to \zeta$ *one-sided large* if for every $i < m$ the half-branch $\rho[i, i + 1/2]$ is large and if $\rho[m - 1, m]$ is a large branch. A one-sided large trainpath $\rho : [0, m] \to \zeta$ is embedded [PH92], and for every $i \in \{1, \ldots, m - 1\}$ the branch $\rho[i - 1, i] \subset \zeta$ is mixed. For every large half-branch $\hat{b}$ of ζ 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 ζ 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 ζ 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 ζ 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 ζ (compare Lemma 5.1 of [H06a]). 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 $1, \ldots, \ell$. A splitting sequence connecting ζ 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 spanned by the first ℓ vectors of the standard basis of $\mathbb{R}^q$. Similarly, the image under Φ of a splitting sequence connecting ζ to η is contained in the subspace $\mathbb{R}^{q-\ell} \subset \mathbb{R}^q$ spanned by the last $q - \ell$ vectors of the standard basis. This shows our claim.

The image under Φ of any edge-path in $E(\tau, \lambda)$ defined by a splitting sequence is an edge-path in the standard cubical subgraph $\mathcal{G}$ of $\mathbb{R}^q$ whose vertices are the points $\mathbb{Z}^q$ in $\mathbb{R}^q$ with integral coordinates and whose edges are the integral translates of the line segments connecting 0 to the standard basis vectors. Such a path is without backtracking, i.e. if $\alpha_1, \ldots, \alpha_q$ is the basis of $(\mathbb{R}^q)^*$ dual to the standard basis of $\mathbb{R}^q$ then for each i the restriction of the function α_i to such a path is non-decreasing. As a consequence, such a path is a geodesic in the graph $\mathcal{G}$ equipped with the intrinsic path metric and hence a uniform quasi-geodesic in $\mathbb{R}^q$. If $\gamma_\sigma, \gamma_\eta$ are such edge-paths connecting $0 = \Phi(\zeta)$ to $\Phi(\sigma), \Phi(\eta)$ induced by a splitting sequence then $\gamma_\sigma \subset \mathbb{R}^\ell, \gamma_\eta \subset \mathbb{R}^{q-\ell}$ and hence $\gamma_\eta \circ \gamma_\sigma^{-1}$ is a uniform quasi-geodesic in $\mathbb{R}^q$ connecting $\Phi(\sigma)$ to $\Phi(\eta)$. But this just means that the distance in $\mathbb{R}^q$ between $\Phi(\sigma), \Phi(\eta)$ is bounded from below by a universal multiple of the distance of σ, η in $E(\tau, \lambda)$ and shows the first part of our lemma.

To show the second part, let again $\sigma, \eta \in E(\tau, \lambda)$ and let $\zeta = \Pi_{E(\tau, \eta)}^1(\sigma)$. Then ζ is splittable to both σ, η. By our above consideration, a splitting sequence connecting ζ to σ commutes with a splitting sequence connecting ζ to η. Using our

above notation, if $\Phi(\zeta) = 0$ then there is a train track $\Theta(\eta, \sigma) \in E(\tau, \lambda)$ with $\Phi(\Theta(\eta, \sigma)) = \Phi(\sigma) + \Phi(\eta)$. This train track has the property stated in the second part of the lemma.   □

A *nonprincipal ultrafilter* is a finitely additive probability measure ω on the natural numbers $\mathbb{N}$ such that $\omega(S) = 0$ or 1 for every $S \subset \mathbb{N}$ and $\omega(S) = 0$ for every finite subset $S \subset \mathbb{N}$. Given a compact metric space X and a sequence $(a_i) \subset X$ ($i \in \mathbb{N}$), there is a unique element $\omega - \lim a_i \in X$ such that for every neighborhood U of $\omega - \lim a_i$ we have $\omega\{i \mid a_i \in U\} = 1$. In particular, given any bounded sequence $(a_i) \subset \mathbb{R}$, $\omega - \lim a_i$ is a point selected by ω.

Let (X, d) be any metric space and let $(z_i) \subset X$. Write $X_\infty = \{(x_i) \in \prod_{i \in \mathbb{N}} X \mid d(x_i, z_i)/i$ is bounded$\}$. For $x = (x_i), y = (y_i) \in X_\infty$ the sequence $d(x_i, y_i)/i$ is bounded and hence we can define $\tilde{d}_\omega(x, y) = \omega - \lim d(x_i, y_i)/i$. Then $\tilde{d}_\omega$ is a pseudodistance on X_∞, and the quotient metric space X_ω equipped with the projection d_ω of the pseudodistance $\tilde{d}_\omega$ is called the *asymptotic cone* of X with respect to the non-principal ultrafilter ω and with basepoint defined by the sequence (z_i). If $z_i = x_0$ for all i and some fixed $x_0 \in X$ then we denote this basepoint by $*$. Note that neither the asymptotic cone defined by X and the constant sequence (x_0) nor the basepoint $*$ depend on the choice of $x_0 \in X$. In the sequel we always assume that the basepoint in the construction of an asymptotic cone of a metric space X is defined by a constant sequence unless explicitly stated otherwise. The cone $(X_\omega, *)$ with basepoint $*$ may depend on the choice of ω. If the isometry group of X acts cocompactly then an asymptotic cone with respect to the ultrafilter ω admits a transitive group of isometries whose elements can be represented by sequences in $\mathrm{Iso}(X)$. The asymptotic cone of a CAT(0)-space is a CAT(0)-space. We refer to [K99] for a careful discussion of asymptotic cones of CAT(0)-spaces.

Our next goal is to determine the asymptotic cone $C(\tau, \lambda)_\omega$ with basepoint the constant sequence (τ) of the maximal extension $C(\tau, \lambda)$ of a flat strip $E(\tau, \lambda) \subset \mathcal{TT}$ where λ is a complete geodesic lamination carried by a train track $\tau \in \mathcal{V}(\mathcal{TT})$. For this define a *metric cone* over a metric space $(\partial Y, \angle)$ to be a metric space (Y, d) of the form $Y = [0, \infty) \times \partial Y / \sim$ where the equivalence relation $\sim$ identifies the set $\{0\} \times \partial Y$ with a single point. The metric d on Y is given by $d((a, \xi), (b, \eta)) = \sqrt{a^2 + b^2 - 2ab \cos \angle(\xi, \eta)}$. The space (Y, d) is a Cat(0)-space if and only if $(\partial Y, \angle)$ is Cat(1) [BH99]. The metric cone is a proper Cat(0)-space if $(\partial Y, \angle)$ is a compact Cat(1)-space. We have.

Lemma 5.4. *The asymptotic cone $C(\tau, \lambda)_\omega$ with respect to a non-principal ultrafilter ω of the maximal extension $C(\tau, \lambda)$ of a flat strip $E(\tau, \lambda) \subset \mathcal{TT}$ is a proper Cat(0)-metric cone and does not depend on ω.*

Proof. Let for the moment Y be an arbitrary proper complete Cat(0)-space with basepoint y_0 and distance function d. Then the *ideal boundary* ∂Y of Y is defined; equipped with the *cone topology*, ∂Y is compact. The boundary ∂Y can also be equipped with the *angular metric* $\angle$; however, the topology defined by this metric need not coincide with the cone topology. The metric space $(\partial Y, \angle)$ is a complete CAT(1)-space (Theorem II.9.13 of [BH99]) which may consist of uncountably many distinct connected components; we call it the *angular boundary* of Y. If $\xi_0 \neq \xi_1 \in$

∂Y are such that $\angle(\xi_0, \xi_1) < \pi$ then there is a geodesic in ∂Y connecting ξ_0 to ξ_1 (Proposition II.9.21 of [BH99]).

Assume that for some sequence $\{i(j)\}$ going to infinity the pointed Cat(0)-spaces $(Y, y_0, d/i(j))$ converge as $j \to \infty$ in the *pointed Gromov Hausdorff topology* to a locally compact pointed metric space $(Y_\infty, y_0, d_\infty)$. Then (Y_∞, d_∞) is a complete Cat(0)-space. By the discussion on p.38 of [B95], (Y_∞, d_∞) is the quotient $[0, \infty) \times (\partial Y, \angle)/ \sim$ where $\{0\} \times \partial Y$ is identified with a single point (the basepoint y_0) and where the metric d_∞ is defined by $d_\infty((a, \xi), (b, \eta)) = \sqrt{a^2 + b^2 - 2ab\cos\angle(\xi, \eta)}$. In other words, (Y_∞, d_∞) equals the metric cone defined by the Cat(1)-space $(\partial Y, \angle)$. Since (Y_∞, d_∞) is locally compact, its ideal boundary equipped with the cone topology is compact and hence the metric space $(\partial Y, \angle)$ is compact. In particular, it consists of only finitely many connected components. The limit space (Y_∞, d_∞) is independent of the sequence $\{i(j)\}$ used to define it, and it is uniquely determined up to isometry by a closed metric ball of positive radius about the basepoint y_0 in Y_∞.

Now let $\lambda \in \mathcal{CL}$ be a complete geodesic lamination and let τ be a train track which carries λ. Let ω be a non-principal ultrafilter. Let $\Phi : E(\tau, \lambda) \to \mathbb{R}^q$ be a standard isometric embedding which maps τ to $\Phi(\tau) = 0$ and determines the maximal extension $C(\tau, \lambda)$ of $E(\tau, \lambda)$. Let $(X_\omega, *)$ be the asymptotic cone of $C(\tau, \lambda)$ defined by the non-principal ultrafilter ω whose basepoint $*$ is given by the constant sequence $(\Phi(\tau))$. By Lemma 5.1, $C(\tau, \lambda)$ is a complete Cat(0)-space and hence the same is true for X_ω. By Lemma 5.3 the inclusion $C(\tau, \lambda) \to \mathbb{R}^q$ is a quasi-isometric embedding and therefore there is a natural bilipschitz embedding of X_ω into $\mathbb{R}^q$, the asymptotic cone of $\mathbb{R}^q$ (see e.g. [KL97]). Since X_ω is complete, the image of this embedding is a closed subset of $\mathbb{R}^q$ and hence X_ω is proper. In particular, X_ω is the limit of a sequence of scaled pointed metric spaces $(C(\tau, \lambda), \Phi(\tau), \frac{1}{i(j)})$ in the pointed Gromov-Hausdorff topology where $\{i(j)\} \subset \mathbb{N}$ is a sequence with $\omega\{i(j) \mid j\} = 1$ (see [K99]). Thus by our above observation, the asymptotic cone X_ω is just the euclidean cone over the ideal boundary $\partial C(\tau, \lambda)$ of $C(\tau, \lambda)$ equipped with the angular metric $\angle$, and it does not depend on the sequence $\{i(j)\}$. This shows the lemma. $\square$

Let again λ be a complete geodesic lamination. Then λ consists of a finite number $\lambda_1, \ldots, \lambda_k$ $(1 \le k \le 3g - 3 + m)$ of minimal components which are connected by a finite number of isolated leaves. The components λ_i are either simple closed curves or minimal arational laminations. If λ_i is a minimal arational component then λ_i *fills* a unique bordered connected subsurface $S_i \subset S$ of S. This means that λ_i is contained in S_i, and every essential simple closed curve c on S which has an essential intersection with S_i, i.e. which is not freely homotopic to a curve contained in $S - S_i$, has an essential intersection with λ_i as well. Up to homotopy, the subsurfaces S_i of S are pairwise disjoint. We call S_i the *characteristic subsurface* of S for λ_i.

If we replace each boundary component of S_i by a puncture then we obtain a surface of finite type, again denoted by S_i, and of negative Euler characteristic which we call the *characteristic surface* of λ_i (recall that we assumed that λ_i is minimal arational). Sometimes we do not distinguish between the characteristic

surface of λ_i and the characteristic subsurface of S for λ_i. The surface S_i may be a four times punctured sphere or a once punctured torus. The lamination λ_i can be viewed as a geodesic lamination on the surface S_i which is minimal and fills S_i, i.e. every complementary component of $\lambda_i \subset S_i$ either is a topological disc or a once punctured topological disc. Hence every train track ζ on S_i which carries λ_i defines a flat strip $E(\zeta, \lambda_i)$ with maximal extension $C(\zeta, \lambda_i)$; note that this also makes sense if S_i is a forth times punctured sphere or a once punctured torus, see [PH92]. For simplicity we denote the asymptotic cone of $C(\zeta, \lambda_i)$ with respect to the non-principal ultrafilter ω and basepoint the constant sequence (ζ) by $A(\lambda_i)$. Note however that $A(\lambda_i)$ may depend on ζ. If λ_i is a simple closed curve then we define $A(\lambda_i)$ to be a single ray $[0, \infty)$.

Call a complete geodesic lamination *spread-out* if it contains precisely $3g - 3 + m$ minimal components. Examples of spread out geodesic laminations are complete geodesic laminations whose minimal components form a pants decomposition of S. There are also other types of spread out geodesic laminations. For example, let $g \geq 1$ and let P be a pants decomposition of S containing a separating pants curve c such that the surface obtained from S by cutting along c is the union of a bordered torus S_0 and a surface S_1 of genus $g - 1$ with one boundary component and m punctures. The surface S_0 contains a pants curve c_0 from the decomposition P in its interior. There is a spread-out geodesic lamination λ on S which contains the components of the simple geodesic multi-curve $P - c_0$ as minimal components and whose intersection with S_0 is the union of a minimal arational geodesic lamination λ_0 and two isolated leaves which connect λ_0 to the boundary circle of S_0.

Given two euclidean cones Y_1, Y_2 with basepoints y_1, y_2, the product $Y_1 \times Y_2$ can be equipped with a product metric in such a way that the resulting metric space is an euclidean cone with basepoint (y_1, y_2). With respect to this metric, the cones $Y_1 \times \{y_2\}$ and $\{y_1\} \times Y_2$ are convex subspaces of $Y_1 \times Y_2$. Any two non-constant geodesic rays $\gamma_1 : [0, \infty) \to Y_1 \times \{y_2\}$ and $\gamma_2 : [0, \infty) \to \{y_1\} \times Y_2$ issuing from the basepoint $\gamma_1(0) = \gamma_2(0) = (y_1, y_2)$ bound a flat convex subspace in $Y_1 \times Y_2$ which is isometric to the closed quadrant $\{(x_1, x_2) \in \mathbb{R}^2 \mid x_i \geq 0\}$. We call $Y_1 \times Y_2$ equipped with this metric the *conical product* of Y_1 and Y_2. The angular boundary of $Y_1 \times Y_2$ equals the *spherical join* $Y_1 * Y_2$ of Y_1 and Y_2 (see [BH99] Chapter I.5).

Call the cone $Z = \{(x_1, \ldots, x_{3g-3+m}) \in \mathbb{R}^{3g-3+m} \mid x_i \geq 0\}$ the *standard partition cone* of dimension $3g - 3 + m$; it equals the iterated conical product of $3g - 3 + m$ single rays, viewed as euclidean cones over single points. We have.

Lemma 5.5. *For a complete geodesic lamination λ on S with minimal components $\lambda_1, \ldots, \lambda_k$ and a complete train track τ which carries λ, the asymptotic cone $C(\tau, \lambda)_\omega$ with basepoint (τ) equals the conical product of k metric cones which are bilipschitz equivalent to the cones $A(\lambda_i)$. If λ is spread out then $C(\tau, \lambda)_\omega$ is isometric to a standard partition cone of dimension $3g - 3 + m$.*

Proof. Let λ be a complete geodesic lamination and let $\lambda_1, \ldots, \lambda_k$ be the minimal components of λ. Let $s \leq k$ be such that (after reordering) the components $\lambda_1, \ldots, \lambda_s$ of λ are minimal arational and that the components $\lambda_{s+1}, \ldots, \lambda_k$ are simple closed curves. We say that a train track η which carries λ *separates* λ if η contains disjoint subtracks $\zeta_1, \ldots, \zeta_k$ with the following property. For each i, the

train track ζ_i carries λ_i. If $i \leq s$ then ζ_i is contained in the characteristic subsurface S_i of S for λ_i, and complementary components of ζ_i on S_i are in one-to-one correspondence with the complementary components of λ_i. If $i \geq s+1$ then ζ_i is a simple closed curve. Moreover, every large branch of η is a subbranch of $\cup_i \zeta_i$. We claim that for every train track $\tau \in \mathcal{V}(\mathcal{TT})$ which carries λ there is a finite splitting sequence $\{\tau(i)\}_{0 \leq i \leq m} \subset E(\tau, \lambda)$ issuing from $\tau(0) = \tau$ such that $\tau(m)$ separates λ.

For this we use the results of [PH92]. Recall that a collision of a train track η at a large branch e is a split of η at e followed by the removal of the diagonal of the split. A collision strictly decreases the number of branches of our train track η. Moreover, the train track obtained from η by a collision at e is a subtrack of a train track obtained from η by a split at e. A *degenerate splitting sequence* is a sequence $\{\eta(i)\}$ of train tracks such that for every i the train track $\eta(i+1)$ can be obtained from $\eta(i)$ by a split or a collision.

For $i \leq s$ let S_i be the characteristic surface of λ_i. As in Section 2, call a train track ξ on a surface $\tilde{S}$ *large* if the complementary components of ξ are all topological discs and once punctured topological discs. For each $i \leq s$ choose a large train track σ_i on S_i which carries λ_i and such that there is a one-to-one correspondence between the complementary components of σ_i and the complementary components of λ_i on S_i. For $i > s$ let σ_i be the train track which is just the simple closed curve λ_i together with the choice of one switch. We may assume that the train tracks σ_i are in fact train tracks on S which are pairwise disjoint. Then $\sigma = \cup_i \sigma_i$ is a train track on S which carries $\cup_i \lambda_i$. Moreover, if ν is a train track obtained from σ by a splitting sequence, if ν carries $\cup_i \lambda_i$ and if ν is a subtrack of a complete train track η which carries λ then η separates λ provided that η does not contain any large branch in $\eta - \nu$.

Choose a transverse measure μ on $\cup_i \lambda_i$ with full support. By Theorem 2.3.1 of [PH92] there is a degenerate splitting sequence $\{\tilde{\tau}(i)\}_{0 \leq i \leq \ell}$ issuing from $\tilde{\tau}(0) = \tau$ with the following properties. The train track $\tilde{\tau}(\ell)$ carries $\cup_i \lambda_i$ and can be obtained from σ by a splitting sequence. Moreover, for each i, the measure μ induces a *positive* transverse measure on $\tilde{\tau}(i)$. In particular, the train tracks $\tilde{\tau}(i)$ are recurrent. The sequence $\{\tilde{\tau}(i)\}$ contains a uniformly bounded number of collisions. Say that there is a sequence $0 \leq i_1 < \cdots < i_p < \ell$ (where p is bounded from above by the number of branches of a complete train track on S) such that for each $j \leq p$ the train track $\tilde{\tau}(i_j + 1)$ is obtained from $\tilde{\tau}(i_j)$ by a collision at a large branch e_j and that for $j \notin \{i_1, \ldots, i_p\}$ the train track $\tilde{\tau}(j+1)$ is obtained from $\tilde{\tau}(j)$ by a split.

Since τ carries the complete geodesic lamination $\lambda \supset \cup_i \lambda_i$ by assumption and since μ define a positive transverse measure on $\tilde{\tau}(i)$ for each i, we may assume that $\tilde{\tau}(j)$ carries λ for every $j \leq i_1$ (compare the discussion in the proof of Lemma 4.3 of [H06b]). The train track $\tilde{\tau}(i_1 + 1)$ is obtained from $\tilde{\tau}(i_1)$ by a collision at the large branch e_1. Let $\tau(i_1 + 1)$ be the unique train track which carries λ and which can be obtained from $\tilde{\tau}(i_1)$ by a *split* at e_1. Then $\tau(i_1 + 1)$ carries λ and contains $\tilde{\tau}(i_1 + 1)$ as a subtrack. As a consequence, the splitting sequence connecting $\tilde{\tau}(i_1 + 1)$ to $\tilde{\tau}(i_2)$ induces as in Section 4 a splitting sequence connecting $\tau(i_1 + 1)$ to a train track $\tau(q)$ which carries λ and contains $\tilde{\tau}(i_2)$ as a subtrack. Inductively in finitely many steps we obtain in this way a splitting sequence $\{\tau(i)\}_{0 \leq i \leq m} \subset E(\tau, \lambda)$ with

the property that $\tau(m)$ contains $\tilde{\tau}(\ell)$ as a subtrack. In particular, $\tau(m)$ separates λ provided that $\tau(m) - \tilde{\tau}(\ell)$ does not contain any large branch.

Now if e is a large branch of $\tau(m)$ which is not contained in $\tilde{\tau}(\ell)$ then e is not incident on a switch contained in $\tilde{\tau}(\ell)$ and the preimage of e under a carrying map $\lambda \to \tau(m)$ does not intersect a minimal component of λ. Therefore this preimage consists of *finitely many* arcs. Let $\hat{\tau}$ be the train track obtained from $\tau(m)$ by a single λ-split at e. The branch $\hat{e}$ of $\hat{\tau}$ corresponding to the branch e in τ is small. Since λ is complete by assumption, a carrying map $\lambda \to \hat{\tau}$ is surjective and hence the number of components of the preimage in λ of the branch $\hat{e}$ of $\hat{\tau}$ under such a carrying map is strictly smaller than the number of components of the preimage in λ of the branch e of $\tau(m)$. As a consequence, after possibly replacing $\tau(m)$ by a train track which can be obtained from $\tau(m)$ by a finite splitting sequence we may assume that $\tau(m)$ separates λ.

By Lemma 5.2, $C(\tau(m), \lambda)$ is a *convex* subspace of the CAT(0)-space $C(\tau, \lambda)$ whose m-neighborhood in $C(\tau, \lambda)$ is all of $C(\tau, \lambda)$ (compare the discussion in the proof of Lemma 5.2). Therefore the asymptotic cone $C(\tau(m), \lambda)_\omega$ with basepoint $(\tau(m))$ is isometric to the asymptotic cone $C(\tau, \lambda)_\omega$ with basepoint (τ). Thus for the purpose of our lemma we may assume without loss of generality that τ separates λ. In particular, for a transverse measure μ on $\cup_i \lambda_i$ with full support, the subtrack σ of τ of all branches of τ with *positive* μ-weight decomposes into k connected components $\sigma_1, \ldots, \sigma_k$ where σ_i carries λ_i for each i. If $i \leq k$ is such that the component λ_i is a simple closed curve then σ_i is an embedded simple closed curve in τ, and if i is such that λ_i is minimal arational then σ_i is a large train track on the characteristic subsurface S_i of S for λ_i.

Let again $s \leq k$ be the such that for $i \leq s$ the component λ_i is minimal arational and that for $i > s$ the component λ_i is a simple closed curve. By the discussion in Section 4, for every $i \leq s$, every splitting sequence issuing from $\sigma_i \subset S_i$ which consists of train tracks carrying λ_i induces a splitting sequence issuing from τ which is contained in the flat strip $E(\tau, \lambda)$. Moreover, for $i \neq j$ a splitting sequence in $E(\tau, \lambda)$ induced by a sequence of λ_i-splits issuing from σ_i commutes with a splitting sequence in $E(\tau, \lambda)$ induced by a sequence of λ_j-splits issuing from σ_j. Since τ separates λ by assumption, up to reordering and composing the isometric embedding $\Phi : E(\tau, \lambda) \to \mathbb{R}^q$ with a translation by an element in $\mathbb{Z}^q$, the maximal extension $C(\tau, \lambda)$ of the flat strip $E(\tau, \lambda)$ is of the form $C(\sigma_1, \lambda_1) \times \cdots \times C(\sigma_k, \lambda_k) \subset \mathbb{R}^{n_1} \times \cdots \times \mathbb{R}^{n_k} \times \mathbb{R}^u = \mathbb{R}^q$ where for each $i \leq k$, $C(\sigma_i, \lambda_i)$ is the convex intersection of $C(\tau, \lambda)$ with the euclidean subspace $\mathbb{R}^{n_i}$ of $\mathbb{R}^q$ spanned by all standard basis vectors which correspond to subbranches of σ_i in τ and where $\mathbb{R}^u$ is spanned by all standard basis vectors which correspond to branches of τ not contained in any of the subtracks σ_i. The convex subspace $C(\sigma_i, \lambda)$ of $C(\tau, \lambda)$ just equals the maximal extension of the subgraph of $E(\tau, \lambda)$ of all train tracks which carry λ and can be obtained from τ by a splitting sequence induced by a sequence of λ_i-splits of σ_i. In particular, the space $C(\sigma_i, \lambda_i)$ is bilipschitz equivalent to the maximal extension of the flat strip $E(\sigma_i, \lambda_i)$ (in general, however, it is not isometric to this extension). As a consequence, the asymptotic cone $C(\tau, \lambda)_\omega$ is a product cone of the form stated in the lemma.

We are left with showing that for a spread out complete geodesic lamination λ the asymptotic cone $C(\tau, \lambda)_\omega$ is a standard partition cone of dimension $3g - 3 + m$. Thus let $\lambda \in \mathcal{CL}$ be a complete geodesic lamination which contains $3g - 3 + m$ minimal components. We claim that λ contains a sublamination Q which is a union of simple closed curves dividing S into pairs of pants, borderd tori with one boundary circle and X-*pieces*, i.e. bordered punctured spheres of Euler characteristic -2. Namely, if λ does not contain any minimal arational component then the minimal components of λ consist of a collection of $3g - 3 + m$ simple closed curves. In other words, these components form a pants decomposition for S and our claim is immediate. Otherwise let λ_0 be a minimal arational component of λ with characteristic subsurface S_0 of S. Then every essential simple closed curve on S which has an essential intersection with S_0 (i.e. which can not be freely homotoped to a curve contained in $S - S_0$) intersects λ_0 transversely. A boundary component of S_0 has vanishing intersection number with λ and hence since the number of minimal components contained in λ equals $3g - 3 + m$, the boundary circles of S_0 are necessarily minimal components of λ. Moreover, since λ_0 is the only minimal component of λ which intersects S_0, either S_0 is a bordered torus with one boundary component or an X-piece as claimed above.

Let $\tau \in \mathcal{V}(\mathcal{TT})$ be a complete train track which carries λ. By our above consideration, for the identification of the asymptotic cone of $C(\tau, \lambda)$ we may assume without loss of generality that τ separates λ. In particular, the images of the minimal components $\lambda_1, \ldots, \lambda_{3g-3+m}$ of λ under a carrying map $\lambda \to \tau$ are disjoint subtracks σ_i of τ. If λ_i is a minimal arational component then σ_i is a train track contained in the interior of a bordered subsurface S_i of S which either is a one-holed torus of Euler characteristic -1 or a four holed sphere of Euler characteristic -2 (where some of the holes may be punctures) and whose boundary consists of simple closed embedded curves in τ. As a consequence, σ_i consists of at most six branches and four switches (Corollary 1.1.3 of [PH92]), and it contains a single large branch since otherwise S_i contains two disjoint simple closed not mutually freely homotopic essential curves. The mapping class group of the surface S_i contains the free group with two generators as a subgroup of finite index, and an infinite splitting sequence of a complete train track on S_i corresponds to choosing an infinite word in these generators. As a consequence, for *every* $i \in \{1, \ldots, 3g - 3 + m\}$ a sequence of λ_i-splits issuing from σ_i is unique, and the asymptotic cone of $C(\sigma_i, \lambda_i)$ is just the single ray $[0, \infty)$. A sequence of λ-splits issuing from τ then consists in choosing in each step one of the subtracks $\tilde{\sigma}_i$ which are filled by the laminations λ_i and performing either a $\tilde{\sigma}_i$-split at a proper large subbranch of $\tilde{\sigma}_i$ or a split which is induced by a λ_i-split of $\tilde{\sigma}_i$. Together with the above, this shows that the asymptotic cone $C(\tau, \lambda)_\omega$ is isometric to the standard $3g - 3 + m$-dimensional partition cone. This completes the proof of the lemma. $\qquad\square$

Call a proper complete CAT(0)-metric cone Y *standard* if its defining Cat(1)-space ∂Y is of diameter strictly smaller than π. Since Y is a Cat(0)-space, this then implies that the angular boundary $(\partial Y, \angle)$ of Y is arcwise connected [BH99]. The following lemma gives additional information on the asymptotic cones of all maximal extensions of flat strips in $\mathcal{TT}$. We always equip the boundary $\partial C(\tau, \lambda)$ of $C(\tau, \lambda)$ with the angular metric.

Lemma 5.6. *Let $\tau \in \mathcal{V}(\mathcal{TT})$, let λ be a complete geodesic lamination which is carried by τ and let ω be non-principal ultrafilter. Then the asymptotic cone $C(\tau,\lambda)_\omega$ of the $\mathrm{Cat}(0)$-space $C(\tau,\lambda)$ is a standard proper $\mathrm{CAT}(0)$ cone with boundary $\partial C(\tau,\lambda)$ of diameter not bigger than $\pi/2$. There is a number $b \in (0,1)$ and an embedding of $\partial C(\tau,\lambda)$ onto a compact arcwise connected subset of a spherical shell $\{x = (x_1,\ldots,x_q) \in \mathbb{R}^q \mid 0 \le x_i \le 1, b \le \|x\| \le 1\}$.*

Proof. Let $\tau \in \mathcal{V}(\mathcal{TT})$ and let λ be a complete geodesic lamination carried by τ. We show that the diameter of the angular boundary $(\partial C(\tau,\lambda), \angle)$ of $C(\tau,\lambda)$ is at most $\pi/2$.

If $\partial C(\tau,\lambda)$ consists of a single point then there is nothing to show, so assume that $\partial C(\tau,\lambda)$ contains at least two points. Let $\xi \ne \xi' \in \partial C(\tau,\lambda)$ be points such that the angle $\angle(\xi,\xi')$ between ξ, ξ' is maximal. Such points exist since by Lemma 5.4, the space $(\partial C(\tau,\lambda), \angle)$ is compact. Let $\gamma \ne \gamma' : [0,\infty) \to C(\tau,\lambda)$ be geodesic rays issuing from $\gamma(0) = \gamma'(0) = \tau$ which define the points ξ, ξ' in $\partial C(\tau,\lambda)$.

Assume that $C(\tau,\lambda)$ is defind by a standard isometric embedding $\Phi : E(\tau,\lambda) \to \mathbb{R}^q$. Let $\alpha^1,\ldots,\alpha^q \subset (\mathbb{R}^q)^*$ be the dual basis of the standard basis of $\mathbb{R}^q$, i.e. the functions α^i are the standard coordinate functions on $\mathbb{R}^q$. By Lemma 5.2, the restriction of each of the euclidean coordinate functions α^i to any geodesic arc in $C(\tau,\lambda)$ issuing from $\Phi(\tau)$ is non-decreasing. More precisely, there is a number $p > 0$ and there is a splitting sequence $\{\tau(i)\}$ in $E(\tau,\lambda)$ issuing from τ such that the Hausdorff distance between $\{\Phi(\tau(i))\}$ and γ does not exceed p. Similarly, there is a splitting sequence $\{\eta(i)\} \subset E(\tau,\lambda)$ such that the Hausdorff distance between $\gamma'[0,\infty)$ and $\Phi(\{\eta(i)\})$ does not exceed p.

For $k > 0$ let $\ell(k) \ge k, \ell'(k) \ge k$ be such that the distance between $\gamma(k)$ and $\Phi(\tau(\ell(k)))$ and the distance between $\gamma'(k)$ and $\Phi(\eta(\ell'(k)))$ is bounded from above by p. Using the notations from Lemma 4.8, for $k \ge 0$ define $\zeta(k) = \Pi^1_{E(\tau,\tau(\ell(k)))}(\eta(\ell'(k)))$. Then the train track $\zeta(k)$ is splittable to both $\tau(\ell(k))$ and $\eta(\ell'(k))$ but this is not true for any train track in $E(\zeta(k),\lambda) - \zeta(k)$.

Denote by d the $\mathrm{Cat}(0)$-metric on $C(\tau,\lambda)$. We claim that there is a number $\alpha \in (0, \angle(\xi,\xi'))$ such that $\min\{d(\Phi(\zeta(k)),\Phi(\tau(\ell(k)))), d(\Phi(\zeta(k)),\Phi(\eta(\ell'(k))))\} \ge \alpha k$ for every sufficiently large $k > 0$. To show this claim, consider the triangle Δ in $C(\tau,\lambda)$ with vertices $\Phi(\tau), \Phi(\zeta(k)), \Phi(\tau(\ell(k)))$ and the triangle Δ' with vertices $\Phi(\tau), \Phi(\zeta(k)), \eta(\ell'((k)))$. The triangles Δ, Δ' have a common side which consists of the geodesic arc connecting $\Phi(\tau)$ to $\Phi(\zeta(k))$. Let Δ_0, Δ'_0 be comparison triangles in the Euclidean plane; we may assume that Δ_0, Δ'_0 have a common side with vertices A, C correponding to the points $\Phi(\tau), \Phi(\zeta(k))$. Let $c : [0,b] \to C(\tau,\lambda)$ be the geodesic arc connecting $c(0) = \Phi(\zeta(k))$ to $c(b) = \Phi(\tau(\ell(k)))$. By Lemma 5.2, the geodesic arc c is contained in the convex subset $C(\zeta(k),\lambda)$ of $C(\tau,\lambda)$ whose distance to $\Phi(\tau)$ equals $d(\Phi(\tau),\Phi(\zeta(k)))$. Thus by convexity of the distance function on a $\mathrm{CAT}(0)$-space, the distance between $\Phi(\tau)$ and $c(s)$ is non-decreasing with s. By comparison, this implies that the angles of the triangles Δ_0, Δ'_0 at the vertex C are not smaller than $\pi/2$.

Since $C(\tau,\lambda)$ is a $\mathrm{Cat}(0)$-space, there is a number $a \in (0, \angle(\xi,\xi'))$ and there is a number $t(a) > 0$ such that $d(\gamma(t),\gamma'(t)) \ge at+2p$ for all $t \ge t(a)$ (compare [BH99]).

Thus if for some $\epsilon > 0$ and large enough k the distance between $\Phi(\zeta(k))$ and $\Phi(\tau(\ell(k)))$ is smaller than $\epsilon a k$ then the distance between $\Phi(\zeta(k))$ and $\Phi(\tau)$ is not smaller than $(1 - a\epsilon)k$, and the distance between $\Phi(\zeta(k))$ and $\Phi(\eta(\ell'(k)))$ is at least $(1-\epsilon)ak$. Since the angle at C of the triangle Δ_0' is not smaller than $\pi/2$, comparison shows that the distance between $\Phi(\eta(\ell'(k)))$ and $\Phi(\tau)$ is not smaller than the length of the side opposite to the right angle of an euclidean right-angled triangle whose sides adjacent to the right angle have length not smaller than $(1 - a\epsilon)k, (a - a\epsilon)k$. Therefore this distance is not smaller than $k\sqrt{1 + a^2 - 2a\epsilon - 2a^2\epsilon + 2a^2\epsilon^2}$ which is strictly bigger than $k + 2p$ provided that $k > 0$ is sufficiently large and $\epsilon > 0$ is sufficiently small compared to a. But the distance between $\Phi(\tau)$ and $\Phi(\eta(\ell'(k)))$ is at most $k + p$ by the choice of $\eta(\ell'(k))$ which is a contradiction. This shows the existence of a number $\alpha > 0$ as claimed above.

As in the proof of Lemma 5.3, observe that the set $\mathcal{E}$ of large branches of $\zeta(k)$ can be partitioned into two disjoint subsets $\mathcal{E} = \mathcal{E}^+ \cup \mathcal{E}^-$ so that a splitting sequence connecting $\zeta(k)$ to $\tau(\ell(k)), \eta(\ell'(k))$ does not contain any split at a branch $e \in \mathcal{E}^+, e' \in \mathcal{E}^-$. Using the notations from the proof of Lemma 5.3, let $\mathcal{A}_0^+, \mathcal{A}_0^-$ be the set of all branches of $\zeta(k)$ contained in a one-sided large trainpath on $\zeta(k)$ terminating at a large branch in $\mathcal{E}^+, \mathcal{E}^-$ and let $\mathcal{A}^\pm$ be the union of $\mathcal{A}_0^\pm$ with those small branches whose endpoints are both starting points of a one-sided large trainpath in $\mathcal{A}_0^\pm$. If we denote by $\mathcal{A}^0$ the collection of all small branches not contained in $\mathcal{A}^+ \cup \mathcal{A}^-$ then we obtain a partition of the set $\mathcal{A}$ of all branches of $\zeta(k)$ into the disjoint sets $\mathcal{A}^+, \mathcal{A}^-, \mathcal{A}^0$ (compare the proof of Lemma 5.3). Normalize the map Φ by a composition with a translation in such a way that $\Phi(\zeta(k)) = 0$. After possibly a permutation of the standard basis of $\mathbb{R}^q$, the partition of the branches of $\zeta(k)$ into the disjoint sets $\mathcal{A}^+, \mathcal{A}^-, \mathcal{A}^0$ determines a direct decomposition $\mathbb{R}^q = \mathbb{R}^{q_1} \times \mathbb{R}^{q_2} \times \mathbb{R}^{q_3}$ with $q_1 > 0, q_2 > 0$ and $q_3 \geq 0$ such that the image under Φ of a splitting sequence $\{\beta(i)\}$ connecting $\zeta(k)$ to $\tau(\ell(k))$ is contained in $\mathbb{R}^{q_1}$, and the image under Φ of a splitting sequence $\{\xi(i)\}$ connecting $\zeta(k)$ to $\eta(\ell'(k))$ is contained in $\mathbb{R}^{q_2}$ (with the obvious interpretation as linear subspaces of $\mathbb{R}^q$). Since splits at large branches in $\mathcal{A}^+, \mathcal{A}^-$ commute and since both train tracks $\tau(\ell(k)), \eta(\ell'(k))$ are contained in the flat strip $E(\tau, \lambda)$, if we denote by C^+ and C^- the maximal extensions of the flat strips $E(\zeta(k), \tau(\ell(k))), E(\zeta(k), \eta(\ell'(k)))$, viewed as convex subsets of $C(\tau, \lambda) \subset \mathbb{R}^q$ (see Lemma 5.2), then for every $x \in C^+$ and $y \in C^-$ we have $x + y \in C(\tau, \lambda)$.

Let again $c : [0, b] \to C(\tau, \lambda)$ be the geodesic connecting $c(0) = \Phi(\zeta(k))$ to $c(b) = \Phi(\tau(\ell(k))))$ and let $c' : [0, b'] \to C(\tau, \lambda)$ be the geodesic connecting $c'(0) = \Phi(\zeta(k))$ to $c'(b') = \Phi(\eta(\ell'(k))))$. Then c, c' are curves in $\mathbb{R}^q$ which are parametrized by arc length. Let $a_1 : [0, \infty) \to \mathbb{R}^2$ be two rays in the euclidean plane parametrized by arc length and issuing from $a_1(0) = a_2(0) = 0$ which enclose a right angle at 0. We may assume that the tangent vectors of a_i at 0 are the standard basis vectors e_1, e_2. Then for $s \in [0, b]$ and $t \in [0, b']$ the line segment $\ell(s, t)$ in $\mathbb{R}^2$ connecting $a_1(s)$ to $a_2(t)$ and parametrized proportional to arc length on $[0, 1]$ can uniquely be represented in the form $\ell(s, t)(u) = a_1(\rho_1(s, t)(u)) + a_2(\rho_2(s, t)(u))$ for functions $\rho_1(s, t), \rho_2(s, t)$ on $[0, 1]$ with values in $[0, s], [0, t]$ and depending continuously on s, t. By our above consideration, for all $s \in [0, b], t \in [0, b']$ and all $u \in [0, 1]$ the point $c(\rho_1(s, t)(u)) + c'(\rho_2(s, t)(u))$ is contained in $C(\tau, \lambda)$. The curve $u \to c(\rho_1(s, t)(u)) + c'(\rho_2(s, t)(u))$ connects $c(s)$ to $c'(t)$, and its length coincides with the length of the curve $\ell(s, t)$. By comparison, this implies that the triangle in $C(\tau, \lambda)$

with vertices $\Phi(\zeta(k)), \Phi(\tau(\ell(k))), \Phi(\eta(\ell'(k))))$ is flat, and its angles at the vertices $\Phi(\tau(\ell(k))), \Phi(\eta(\ell'(k)))$ sum up to $\pi/2$. By comparison and our above discussion, the distance between $\Phi(\tau)$ and $\Phi(\tau(\ell(k))), \Phi(\eta(\ell'(k)))$ is not smaller than the distance between $\Phi(\zeta(k))$ and $\Phi(\tau(\ell(k))), \Phi(\eta(\ell'(k)))$. Therefore if we denote by $\tilde{\Delta}(k)$ a comparison triangle in the euclidean plane for the triangle $\Delta(k)$ in $C(\tau, \lambda)$ with vertices $\Phi(\tau), \Phi(\tau(\ell(k))), \Phi(\eta(\ell'(k)))$, then by comparison, the angle at the point corresponding to τ in $\tilde{\Delta}(k)$ is not bigger than $\pi/2$. Since $k > 0$ was arbitrary and the distance between $\gamma(k), \gamma'(k)$ and $\Phi(\tau(\ell(k))), \Phi(\eta(\ell'(k)))$ is uniformly bounded, by the definition of the angle between ξ, ξ' and the results in Chapter II.9 of [BH99], this means that $\angle(\xi, \xi') \le \pi/2$. However, we chose ξ, ξ' in such a way that their angular distance is maximal among all distances in the angular boundary of $C(\tau, \lambda)$ and therefore the diameter of $\partial C(\tau, \lambda)$ with respect to the angular metric is at most $\pi/2$. This completes the proof of the first part of our lemma.

To show the second part, let again $\alpha^1, \ldots, \alpha^q$ be the basis of $\mathbb{R}^q$ which is dual to the standard basis of $\mathbb{R}^q$. For every $z \in \partial C(\tau, \lambda)$ there is a unique geodesic ray $\gamma_z : [0, \infty) \to C(\tau, \lambda)$ issuing from $\gamma_z(0) = \Phi(\tau)$ which is asymptotic to z. For $j \le q$ let $\alpha_\omega^j(z) = \omega - \lim_{k \to \infty} \alpha^j(\gamma_z(k))/k$. Since $\alpha^j(\gamma_z(k)) \le k$ for all k, this limit exists. By our explicit construction, the point $\rho(z) = (\alpha_\omega^1(z), \ldots, \alpha_\omega^q(z))$ has non-negative entries, has norm bounded in $[b, 1]$ for universal constant $b > 0$ and depends continuously on z. Moreover the map $z \to \rho(z)$ is injective and hence the assignment $\rho : z \to \rho(z)$ defines an embedding of the boundary of $C(\tau, \lambda)$ onto a compact path-connected subset of the spherical shell $\{x = (x_1, \ldots, x_q) \in \mathbb{R}^q \mid x_i \in [0, 1], b \le \|x\| \le 1\}$. $\qquad\square$

Finally we are able to estimate from above the topological dimensions of the asymptotic cones $C(\tau, \lambda)_\omega$.

Lemma 5.7. *The topological dimension of the cones $C(\tau, \lambda)_\omega$ is bounded from above by $3g - 3 + m$.*

Proof. Let $\tau \in \mathcal{V}(\mathcal{T}\mathcal{T})$ and let λ be a complete geodesic lamination carried by τ. Choose a non-principal ultrafilter ω. We have to show that the topological dimension of $C(\tau, \lambda)_\omega$ does not exceed $3g - 3 + m$. For this note first that by Lemma 5.5, this holds true for spread-out complete geodesic laminations. In particular, it holds true for an exceptional surface S, i.e. a one-punctured torus or a forth punctured sphere with the obvious interpretation of flat strips for these exceptional surfaces (see the proof of Lemma 5.5). By induction, we therefore may assume that our dimension estimate is valid for all proper subsurfaces of S. By Lemma 5.5, it then also holds for every geodesic lamination which does not contain a minimal component which fills up S.

Thus let λ be a complete geodesic lamination which contains a minimal component λ_0 which fills up S. Let $\tau \in \mathcal{V}(\mathcal{T}\mathcal{T})$ by a train track which carries λ. Let $\Phi : E(\tau, \lambda) \to \mathbb{R}^q$ be a standard isometric embedding with $\Phi(\tau) = 0$ which defines the maximal extension $C(\tau, \lambda)$. By Lemma 5.5, the asmyptotic cone $C(\tau, \lambda)_\omega$ is the metric cone over the angular boundary $(\partial C(\tau, \lambda), \angle)$, and by Lemma 5.6, this boundary is a compact CAT(1) geodesic metric space of diameter at most $\pi/2$. By Lemma 5.6, there is a number $b > 0$ such that the ultrafilter ω defines an embedding

ρ of the boundary $\partial C(\tau, \lambda)$ onto a compact connected subset C of a spherical shell $\{z = (z_1, \ldots, z_q) \in \mathbb{R}^q \mid z_i \geq 0, b \leq \|z\| \leq 1\}$ for some $b > 0$.

We have to show that the topological dimension of the set C is at most $3g-4+m$. For this let again $\alpha^1, \ldots, \alpha^q$ be the basis of $(\mathbb{R}^q)^*$ which is dual to the standard basis. Let $\gamma_1, \ldots, \gamma_n$ be geodesic rays in $C(\tau, \lambda)$ issuing from $\Phi(\tau)$ with corresponding points $z_1, \ldots, z_n$ in the compact set C. We choose the rays γ_i in such a way that for each of the points $z_j \in C$ there is a linear function α^{i_j} which assumes a maximum at z_j. We may also assume that there is some $\epsilon > 0$ with the property that for $j \neq k$ the value of α^{i_j} on z_k is smaller than $\alpha^{i_j}(z_j)/(1 + 2\epsilon)$. After reordering of the standard basis vectors, we may assume that $i_j = j$ for all $j \leq n$.

By Lemma 5.2 and its proof, for each i there is a splitting sequence $\{\tau_i(j)\}$ whose image under the map Φ is of Hausdorff distance to $\gamma_i[0, \infty)$ bounded from above by a universal constant $p > 0$. For $k > 0$ let $\ell_i(k)$ be the such that the distance between $\Phi(\tau_i(\ell_i(k)))$ and $\gamma_i(k)$ is at most p. We may assume that for ω-all k we have $\alpha^1(\Phi(\tau_1(\ell_1(k)))/\alpha^1(\Phi(\tau_i(\ell_i(k))) \geq 1 + \epsilon$ for all $i \geq 2$.

Using the notation from Lemma 4.8, for $i \geq 2$ let $\eta(i) = \Pi^1_{E(\tau, \tau_1(\ell_1(k)))}\tau_i(\ell_i(k)) \in E(\tau, \tau_1(\ell_1(k)))$. As in Lemma 4.8, there is a train track $\eta \in E(\tau, \tau_1(\ell_1(k)))$ such that for each i, $\eta(i)$ is splittable to η and that moreover if $\zeta \in E(\tau, \tau_1(\ell_1(k)))$ is such that $\eta(i)$ is splittable to ζ for each i then η is splittable to ζ. The coordinate functions of $\Phi(\eta)$ satisfy $\alpha^j(\Phi(\eta)) = \max\{\alpha^j(\Phi(\eta(i))) \mid i \geq 1\}$. Since $\alpha^1(\tau_1(\ell_1(k)) > (1 + \epsilon) \max_{i \geq 2} \alpha^1(\tau_i(\ell_i(k)))$ for $i \geq 2$ there is a partition of the set $\mathcal{E}$ of large branches of η into disjoint sets $\mathcal{E} = \mathcal{E}_1 \cup \mathcal{E}_2$ with the property that a splitting sequence connecting η to $\tau_1(\ell_1(k))$ does not contain any split at a large branch $e \in \mathcal{E}_2$ and that moreover the following holds. For $i \geq 2$ and using the notation from Lemma 5.2, let $\zeta_i = \Theta(\eta, \tau_i(\ell_i(k))) \in E(\tau, \lambda)$. Then both η and $\tau_i(\ell_i(k)))$ are splittable to ζ_i, and a splitting sequence connecting η to ζ_i does not contain any split at a large branch in $\mathcal{E}_1$. Define moreover inductively $\nu_2 = \zeta_2$ and $\nu_i = \Theta(\zeta_i, \nu_{i-1})$ for $i \geq 3$ and write $\zeta = \nu_k$. Then each of the train tracks ζ_i is splittable to ζ, and every train track with this property can be obtained from ζ by a splitting sequence.

By construction, we have $\alpha^1(\tau_1(\ell_1(k)))/\alpha^1(\tau_i(\ell_i(k))) \geq 1 + \epsilon$ for all i. Now by monotonicity of the coordinate functions on geodesic rays issuing from $\Phi(\tau)$ we necessarily have $\alpha^1(\tau_1(\ell(k)) \to \infty$ $(k \to \infty)$ and therefore for sufficiently large k the union $\mathcal{A}(\eta)$ of the set $\mathcal{A}_0(\eta)$ of all branches which are contained in a one-sided large trainpath on η terminating at a branch in $\mathcal{E}_1$ with the set of all small branches whose endpoints are both contained in $\mathcal{A}_0(\eta)$ is a subgraph of η which contains a simple closed curve. The train tracks $\tau_i(\ell_i(k))$ $(i \geq 2)$ are contained in a flat strip $E(\tau, \zeta)$ where ζ can be obtained from η by a splitting sequence which does not contain a split at any of the branches in $\mathcal{A}(\eta)$. Hence a splitting sequence connecting τ to ζ is induced from a splitting sequence of a subtrack of η contained in a proper subsurface of S of strictly bigger Euler characteristic. Thus by induction hypothesis, the images under the map ρ of the rays $\gamma_2, \ldots, \gamma_n$ are contained in a compact subset of C of dimension at most $3g - 5 + m$. Then the dimension of the convex hull in $\partial C(\tau, \lambda)$ with respect to the angular distance of the points $z_1, \ldots, z_n$ is at most $3g - 4 + m$. Since the points of C were arbitrarily chosen with the above properties this shows the lemma. $\square$

6. A QUASI-CONVEX BICOMBING OF THE TRAIN TRACK COMPLEX

A *bicombing* of a metric space (X, d) assigns to every pair of points $x, y \in X$ a curve $c_{x,y} : [0,1] \to X$ connecting $x = c_{x,y}(0)$ to $y = c_{x,y}(1)$. The curve $c_{x,y}$ is called the *combing line* connecting x to y. We call the bicombing *symmetric* if $c_{x,y}(t) = c_{y,x}(1-t)$ for all x, y and all $t \in [0,1]$, *reflexive* if $c_{x,x}(t) = x$ for all $x \in X$ and all $t \in [0,1]$ and *L-Lipschitz* for some $L \geq 1$ if for all $x, y \in X$ the curve $t \to c_{x,y}(t/d(x,y))$ $(t \in [0, d(x,y)])$ is L-Lipschitz. Call moreover the bicombing *L-quasi-convex* for some $L > 0$ if for all $x, y, x', y' \in X$ and all $t > 0$ we have $d(c_{x,y}(t), c_{x',y'}(t)) \leq L(d(x, x') + d(y, y')) + L$. As an example, if X is a Cat(0)-space then any two points can be connected by a unique geodesic parametrized proportional to arc length, and these geodesics define a reflexive symmetric 1-Lipschitz 1-quasi-convex bicombing of X which we call the *geodesic bicombing*.

The purpose of this section is to construct a reflexive symmetric L-Lipschitz L-quasi-convex bicombing for the train track complex $\mathcal{TT}$. For this fix a framing F of S and let X be the set of *all* complete train tracks which can be obtained from a train track in standard form for F by a splitting sequence. Then X is r-dense in $\mathcal{TT}$ for some $r > 0$. As a consequence, it is sufficient to construct such a reflexive symmetric L-Lipschitz L-quasi-convex bicombing for the set X equipped with the restriction of the metric on $\mathcal{TT}$.

We begin with constructing for a train track τ in standard form for F and for a complete geodesic lamination λ carried by τ a bicombing for the flat strip $E(\tau, \lambda)$. The combing path connecting τ to a train track $\eta \in E(\tau, \lambda)$ is obtained from a particular splitting sequence connecting τ to η. First we establish some suitable notations. Namely, let $\rho : [0, m] \to \tau$ be any trainpath on τ. Then for every $i \in \{1, \ldots, m-1\}$ there is a single branch of τ which is incident on $\rho(i)$ and not contained in ρ. We call such a branch a *neighbor* of ρ at $\rho(i)$. The switch is called a *left* switch (or a *right* switch) if the neighbor of ρ at $\rho(i)$ is to the left (or to the right) of ρ with respect to the orientation of ρ and the orientation of S. Define a *special trainpath* on a train track σ to be a trainpath $\rho : [0, 2k-1] \to \sigma$ 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 G shows a special trainpath of length 5.

Figure G

split   shift
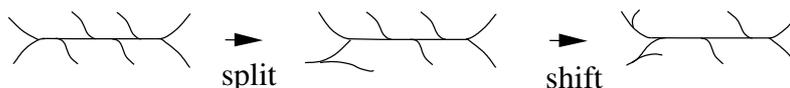
A *special circle* in a train track σ is a trainpath $\rho : [0, 2k-1] \to \sigma$ with $\rho[0,1] = \rho[2k-2, 2k-1]$ (as oriented arcs) for some $k \geq 2$ such that $\rho[0, 2k-2)$ is embedded in σ and which satisfies the requirements 2), 3) in the definition of a special trainpath.

If τ' is any train track containing a special trainpath or a special circle ρ' of length $2k-1$ and if τ is shift equivalent to τ', then there is a natural bijection of the branches of τ' onto the branches of τ, and this bijection preserves the type of branches. In particular, there is a trainpath ρ on τ of length at least $2k-1$ which contains the same large and small branches as ρ' under our identification of branches. We say that ρ *corresponds* to ρ'. We call a trainpath $\rho : [0, m] \to \tau$ on a train track τ *symmetric large* if ρ corresponds to a special trainpath in this way, and we call ρ a *symmetric circle* if it corresponds to a special circle. Finally, a *symmetric trainpath* is a trainpath which either is a symmetric large trainpath or a symmetric circle.

If $\rho : [0, m] \to \tau$ is a special trainpath on τ of length $m \geq 2$ and if $i \geq 0$ is such that $\rho[i-1, i]$ is a large branch then there is a unique choice of a right or left split of τ at ρ for which the branches $\rho[i-2, i-1]$ and $\rho[i, i+1]$ are winners. We call such a split the *ρ-split* of τ at $\rho[i-1, i]$ (note that one of the two branches $\rho[i-2, i-1]$ or $\rho[i, i+1]$ may be empty). If the length m of ρ equals one, then every split of τ at $\rho[0, 1]$ is a ρ-split by definition. If τ' is obtained from τ by a ρ-split at a large branch $\rho[i-1, i]$ then there is a natural bijection $\varphi(\tau, \tau')$ from the branches of τ onto the branches of τ' (compare the discussion in [H06a]). The image of ρ under the map $\varphi(\tau, \tau')$ is a trainpath ρ' on τ' of the same length as ρ. Thus for every symmetric trainpath ρ on τ it makes sense to talk about a splitting sequence consisting of a single ρ-split at *every* branch of ρ. We first observe.

Lemma 6.1.   (1) *Let $\tau \in \mathcal{V}(\mathcal{TT})$ and let $\rho : [0, m] \to \tau$ be a symmetric large trainpath. Then there is a unique train track τ' which can be obtained from τ by a splitting sequence of length m consisting of a single ρ-split at each branch $\rho[i-1, i]$ $(i \leq m)$. The train track τ' contains a symmetric large trainpath ρ' of maximal length $n \in \{0, \ldots, m-2\}$ which is contained in the image of ρ under $\varphi(\tau, \tau')$. If ρ corresponds to a special trainpath of length one then ρ' is trivial.*

(2) *Let $\rho : [0, m] \to \tau$ be a symmetric circle. Then there is a unique train track τ' which can be obtained from τ by a splitting sequence of length $m-1$ consisting of a single ρ-split at each branch $\rho[i-1, i]$ $(i \leq m-1)$. The image of ρ in τ' under the map $\varphi(\tau, \tau')$ is a symmetric circle of the same length. Moreover, there is a number $k < m/2$ such that the train track obtained from τ by k such modifications is the image of τ under a simple Dehn twist along the circle $\rho[0, m-1]$.*

Proof. We show the first part of the lemma by induction of the length m of a symmetric large trainpath ρ on a complete train track τ. The case that this length equals one is trivial, so assume that we showed the claim for all symmetric large trainpaths of length at most $m-1$ for some $m \geq 2$.

Let $\rho : [0, m] \to \tau$ be a symmetric large trainpath of length m on a complete train track τ. If $\rho[0, 1]$ is a large branch then by assumption, there is a unique choice of a right or left split of τ at $\rho[0, 1]$ so that the branch $\rho[1, 2]$ is a winner of the split.

Let $\tilde{\tau}$ be the train track obtained from τ by this choice of a split. Then the image of ρ under the map $\varphi(\tau, \tilde{\tau})$ is a trainpath in $\tilde{\tau}$ which begins with a small branch, and its subpath $\varphi(\tau, \tilde{\tau})\rho[1, m]$ is a symmetric large trainpath $\tilde{\rho}$ of length $m - 1$. Clearly $\tilde{\rho}$ is the maximal symmetric large trainpath on τ' which is contained in the image of ρ under the map $\varphi(\tau, \tilde{\tau})$. We then can apply our induction hypothesis to ρ' and deduce in this way the statement of the lemma.

If the branch $\rho[0, 1]$ is *not* large then it is mixed and the half-branch $\rho[0, 1/2]$ is large. Thus $\rho(0)$ is the starting point of a unique one-sided large trainpath $\zeta : [0, k] \to \tau$ which terminates at a large branch $\zeta[k - 1, k]$. Note that we have $2 \leq k \leq m$, moreover necessarily $\zeta[0, k] = \rho[0, k]$. There is a unique choice of a split of τ at $\rho[k - 1, k]$ with the branch $\rho[k - 2, k - 1]$ as a winner. Let τ_1 be the split track. The image of ρ in τ_1 under the map $\varphi(\tau, \tau_1)$ is a symmetric large trainpath $\rho_1 : [0, m] \to \tau_1$ with the additional property that the subpaths $\rho_1[0, k - 1]$ and $\rho_1[k, m]$ are both symmetric large, however the trainpath $\rho_1[k, m]$ may be trivial. Since the lengths of these trainpaths are strictly smaller than m, we can apply our induction hypothesis and obtain the statement of the lemma for symmetric large trainpaths.

If $\rho : [0, m] \to \tau$ is a symmetric circle then we may assume that $\rho[0, 1]$ is a large branch. Let $\hat{\tau}$ be the train track obtained from τ by a single ρ-split at $\rho[0, 1]$. The image of ρ under the map $\varphi(\tau, \hat{\tau})$ is a special circle $\hat{\rho} : [0, m] \to \hat{\tau}$ of the same length as ρ. Moreover, $\hat{\rho}[1, m - 1]$ is a symmetric large trainpath. By our above consideration, there is a splitting sequence consisting of a single split at each branch of $\hat{\rho}$. The image of $\hat{\rho}$ under this sequence is a trainpath ρ' on a train track τ' beginning and ending with a small half-branch which combines with $\varphi(\tau, \tau')\hat{\rho}[0, 1]$ to a symmetric circle. Moreover, the map $\varphi(\tau, \tau')$ of ρ onto ρ' preserves the type of the branches. By our explicit construction, if $2k + 1$ is the length of the special circle which is shift equivalent to ρ then after k such steps we obtain a train track $\tilde{\tau}$ which is the image of τ under a single Dehn twist along the circle $\rho[0, m - 1]$, with the twist direction determined by ρ. This shows the lemma. $\square$

We say that the train track τ' as in Lemma 6.1 is obtained from the symmetric large trainpath $\rho : [0, m] \to \tau$ on τ by a *level-one ρ-multi-split*.

Now let $\tau \in \mathcal{V}(\mathcal{TT})$ be a complete train track which is splittable to a complete train track σ. Define a *level-one splittable σ-configuration* to be a symmetric large trainpath $\rho : [0, m] \to \tau$ of *maximal length* with the property that a level-one ρ-multi-split of τ is splittable to σ. A *level-one non-splittable σ-configuration* consists of a single large branch e of τ so that no train track obtained from τ by a split at e is splittable to σ. We have.

Lemma 6.2. *Let τ be a train track which is splittable to σ. Then every large branch e of τ is contained in a unique level-one σ-configuration, and two such level-one σ-configurations either coincide or are disjoint.*

Proof. By definition of a level-one σ-configuration and by uniqueness of splitting sequences, a large branch e of τ such that no split of τ at e is splittable to σ

is contained in a unique level-one σ-configuration, and this configuration is non-splittable.

Now assume that there is a train track $\tilde{\tau}$ obtained from τ by a split at e which is splittable to σ. Let v be a switch of τ on which e is incident and let a be the branch which is incident and small at v and which is a winner of the split connecting τ to $\tilde{\tau}$. Let v' be the second vertex on which a is incident. Assume first that a is a small branch. Then there is a unique trainpath $\rho : [0, m] \to \tau$ with $\rho[0, 1] = e, \rho[1, 2] = a$ and such that $\rho[2, m]$ is the one-sided large trainpath on τ starting at $v' = \rho(2)$. Write $a' = \rho[m - 1, m]$.

If $a' = a$ and if $\rho(1) = \rho(m - 1)$ then $\rho[0, m]$ is a loop which has a cusp at $\rho(1)$ and therefore the branch a is *not* contained in the level-one σ-configuration containing e. In other words, in this case the switch v is contained in the boundary of any level-one σ-configuration containing e. If $a' = a$ and if $\rho(1) = \rho(m)$ then $\rho[0, m - 1]$ is a C^1-embedded circle in τ. By construction, this circle c is clearly symmetric. Thus c equals the level-one σ-configuration containing e if and only if the train track obtained from τ by a c-multi-split is splittable to σ. If this is not the case then the switch v is contained in the boundary of any level-one σ-configuration containing e.

If $a' \neq a$ then the trainpath $\rho[0, m]$ on τ is symmetric large, and the branch a is contained in a level-one σ-configuration containing e if and only if the level-one ρ-multi-split of τ is splittable to σ.

Now if the branch a is mixed then the branch $\tilde{a}$ corresponding to a in the train track $\tilde{\tau}$ obtained from τ by a single σ-split at e is large. The branch a belongs to a level-one σ-configuration if and only if the train track obtained from $\tilde{\tau}$ by a σ-split at $\tilde{a}$ with the small branch corresponding to e in $\tilde{\tau}$ as a winner is splittable to σ. Moreover, this condition chooses uniquely one of the two neighbors of a which can possibly be contained in a level-one σ-configuration containing e.

In finitely many steps of this form determined by σ we extend in this way our trainpath starting at e beyond the switch v and passing through a until no further such extension is possible. If the resulting trainpath is not closed then we repeat this construction with the second switch w on which e is incident. By uniqueness of our procedure, in finitely many steps we construct in this way a maximal symmetric trainpath ρ so that the train track obtained from τ by a level-one ρ-multi-split is splittable to σ. $\square$

Now let $\tau \in \mathcal{V}(\mathcal{TT})$ and let $\rho : [0, m] \to \tau$ be a symmetric large trainpath. Let τ_1 be the train track obtained from τ by the level-one multi-split along ρ. Then τ_1 contains a symmetric large trainpath $\rho_1 : [0, n] \to \tau_1$ of length $n \leq m - 2$ which is contained in the image of ρ under the map $\varphi(\tau, \tau_1)$. Define the *level-two ρ-multi-split* to be the train track τ_2 obtained from τ_1 by a level-one ρ_1-multi-split. Then τ_2 contains a symmetric large trainpath ρ_2 of length at most $m - 4$ which is contained in the image of ρ_1 under the natural bijection $\varphi(\tau_1, \tau_2)$. Inductively in at most $m/2$ steps we repeat this construction until the length of our trainpath vanishes. The train track obtained from τ by a *ρ-multi-split* is by definition the train track defined inductively in this way.

Similarly, if $\rho : [0, m] \to \tau$ is a symmetric circle then we define the ρ-*multi-split* to be the train track obtained from τ by a sequence of ρ-splits and which is the image of τ under a simple Dehn twist along the circle $\rho[0, m - 1]$ as described in Lemma 6.1.

Now assume that τ is splittable to σ and let ρ be a level-one σ-configuration as defined in Lemma 6.2. We define the train track obtained from τ by a σ-*move at* ρ to be the unique train track τ' with the following two properties.

(1) τ' is splittable to a train track obtained from τ by a ρ-multi-split.
(2) If η is splittabe to both σ and the train track obtained from τ by a ρ-multi-split then η is splittable to τ'.

For a train track τ which is splittable to a train track σ we define a σ-*move* to be the following modification of τ. Let $\rho_1, \ldots, \rho_k$ be the splittable level-one σ-configurations of τ. By Lemma 6.2, these are uniquely defined pairwise disjoint symmetric trainpaths on τ. We define the train track τ' obtained from τ by a σ-move to be the train track resulting from σ-moves at each of the symmetric trainpaths ρ_i.

For each complete train track τ which is splittable to a complete train track σ define now inductively a sequence $\{\tau(i)\}_{0 \leq i \leq m} \subset E(\tau, \sigma)$ beginning at τ and ending at σ by requiring that for each i the train track $\tau(i + 1)$ is obtained from $\tau(i)$ by a σ-move. We call the sequence the *tight multi-sequence* connecting τ to σ, and we denote it by $\gamma(\tau, \sigma)$. Note that a tight multi-sequence is uniquely determined by τ and σ. Moreover, since $\tau(i + 1)$ can be obtained from $\tau(i)$ by a non-trivial splitting sequence of uniformly bounded length, there is a universal constant $\kappa > 0$ such that every tight multi-sequence defines a κ-quasi-geodesic in $\mathcal{TT}$.

We call two curves $c_1 : [0, a_1] \to \mathcal{TT}, c_2 : [0, a_2] \to \mathcal{TT}$ for some $0 \leq a_1 \leq a_2 < \infty$ *weight-L fellow travellers* if the following holds.

(1) $d(c_1(t), c_2(t)) \leq L(d(c_1(0), c_2(0)) + d(c_1(a_1), c_2(a_2)))$ for every $t \in [0, a_1]$.
(2) $d(c_1(a_1), c_2(t)) \leq Ld(c_1(a_1), c_2(a_2))$ for all $t \in [a_1, a_2]$.

If c_1, c_2 are weight-L fellow travellers then the Hausdorff distance in $\mathcal{TT}$ between the images $c_1[0, a_1]$ and $c_2[0, a_2]$ is bounded from above by $L(d(c_1(0), c_2(0)) + d(c_1(a_1), c_2(a_2)))$.

Lemma 6.3. *There is a number $L > 0$ with the following property. Let λ be a complete geodesic lamination carried by a complete train track τ and let $\sigma, \eta \in E(\tau, \lambda)$. Then the tight multi-sequences $\gamma(\tau, \sigma)$ and $\gamma(\tau, \eta)$ are weight-L fellow travellers.*

Proof. By Corollary 4.10, it suffices to show the existence of a number $a > 0$ with the following property. Let d_λ be the intrinsic path-metric on $E(\tau, \lambda)$. Then for $\sigma, \eta \in E(\tau, \lambda)$ the curves $\gamma(\tau, \sigma)$ and $\gamma(\tau, \eta)$ are weight-L fellow travellers. Moreover, by the explicit description of the intrinsic distance function on $E(\tau, \lambda)$, for our purpose it is in fact enough to show this property for train tracks $\eta, \sigma \in E(\tau, \lambda)$ such that σ can be obtained from η by a single split at a large branch e.

Thus let $\{\tau(i)\}_{0 \le i \le k}$ be the tight multi-sequence connecting τ to σ and let $\{\zeta(i)\}_{0 \le i \le \ell}$ be the tight multi-sequence connecting τ to η; then there is a largest number $i \le k$ such that $\zeta(i-1) = \tau(i-1)$. If $i-1 = \ell$ then we have $\tau(\ell) = \eta$. Since σ can be obtained from η by a single split at a large branch e we obtain $k = \ell + 1$, and the distance between corresponding points on the tight multi-sequences connecting τ to σ, η is at most one which shows our claim.

Now consider the case that $i - 1 < \ell$. Since $\zeta(i) \ne \tau(i)$ the train track $\tau(i-1)$ contains a level-one σ-configuration $\rho : [0, m] \to \tau(i-1)$ so that the train track τ_1 obtained from the σ-move at ρ is not splittable to η.

Assume first that the level-one ρ-multi-split is not splittable to η. Then there is some $j \le m$ such that a splitting sequence connecting $\tau(i-1)$ to η does not contain a split at the branch $\rho[j-1, j]$. We distinguish three cases.

Case 1: $\rho[j-1, j]$ is a large branch.

By the fact that σ can be obtained from η by a single split at a large branch e and uniqueness of splitting sequences, the branch $\rho[j-1, j]$ coincides with e via the map $\varphi(\tau(i-1), \eta)$. Now the branch $\rho[j, j+1]$ is incident and small at $\rho(j)$ and hence no splitting sequence issuing from $\tau(i-1)$ which does not contain a split at $\rho[j-1, j]$ contains a split at $\rho[j, j+1]$. This implies that a splitting sequence connecting $\tau(i-1)$ to η does not contain a split at $\rho[j, j+1]$, and the same is true for a splitting sequence connecting $\tau(i-1)$ to σ. By our assumption on η, σ and by the definition of a level-one σ-configuration, we conclude that ρ is not a symmetric circle and that $m = j$. The same argument also shows that $j = 1$ and hence ρ consists of a single large branch e. As a consequence, the train track $\tau(i)$ can be obtained from $\eta(i)$ by a single split at e. Inductively we conclude that for every $u \in \{i, \ldots, k\}$ the train track $\tau(u)$ can be obtained from $\zeta(u)$ by a single split at e. In other words, the distance between cooresponding points on the tight multi-sequences $\gamma(\tau, \sigma)$ and $\gamma(\tau, \eta)$ is at most one.

Case 2: $\rho[j-1, j]$ is a mixed branch.

Assume without loss of generality that $\rho[j-1, j]$ is large at $\rho(j-1)$, i.e. that the one-sided large trainpath issuing from $\rho[j-1, j]$ is the path $\rho[j-1, q]$ for some $q \le m$. If $j \ne 1$ then the branch $\rho[j-2, j-1]$ is small at $\varphi(j-1)$ and a splitting sequence which does not contain a split at $\rho[j-1, j]$ can not contain a split at $\rho[j-2, j-1]$. It now follows as in Case 1 above from the definition of a level-one σ-configuration that we necessarily have $j = 1$. Then the trainpath $\rho[1, m]$ is symmetric large and defines a level-one η-configuration. By construction, this implies that for every $u \in \{i, \ldots, k\}$ the train track $\tau(u)$ can be obtained from $\zeta(u)$ by a single split at the branch e and hence the distance between corresponding points on the tight multi-sequences $\gamma(\tau, \sigma)$ and $\gamma(\tau, \eta)$ is at most one.

Case 3: $\rho[j-1, j]$ is a small branch.

Assume first that ρ is a symmetric large trainpath. Then by definition, we necessarily have $2 \le j \le m - 1$ and the trainpaths $\rho[0, j-1]$ and $\rho[j, m]$ are level-one η-configurations. It follows from our explicit construction that for every $u \in \{i, \ldots, k\}$ the train track $\tau(u)$ can be obtained from $\eta(u)$ by a single split at e.

If ρ is a symmetric circle then the trainpath $\zeta : [0, m - 2] \to \tau(i - 1)$ defined by $\zeta[k, k + 1] = \rho[k + j - 1, k + j]$ (indices are taken modulo $m - 1$) is symmetric large and an η configuration. By our definition, the ζ-multi-split is splittable to the ρ-multi-split and therefore as before, for every $u \in \{i, \ldots, k\}$ the the train track $\tau(u)$ can be obtained from $\eta(u)$ by a single split at e.

In the case that the level-one σ-configuration ρ is also a level-one η-configuration we can apply the above consideration to the train tracks obtained from $\tau(i - 1)$ and $\sigma(i - 1)$ by the level-one ρ-multi-split. Our control on the distance between corresponding points on $\gamma(\tau, \eta)$ and $\gamma(\tau, \sigma)$ follows. This shows the lemma. $\square$

In the following proposition, we denote by $E(F, \lambda)$ the flat strip defined by a train track in standard form for F which carries the complete geodesic lamination λ.

Proposition 6.4. *There is a number $L > 0$ with the following property. Let F be any framing of S and let $X \subset \mathcal{V}(\mathcal{TT})$ to be the set of all train tracks which can be obtained from a train track in standard form for F by a splitting sequence. Then there is a reflexive symmetric L-Lipschitz L-quasi-convex bicombing of X equipped with the restriction of the metric on $\mathcal{TT}$. If $x \in E(F, \lambda), y \in E(F, \nu)$ for some $\lambda, \nu \in \mathcal{CL}$ then the combing line connecting x to y is contained in the L-neighborhood of $E(F, \lambda) \cup E(F, \nu)$.*

Proof. Let F be a framing for S and let τ be a complete train track in standard form for F. Let λ be a complete geodesic lamination carried by τ. We construct first a reflexive symmetric L-Lipschitz and L-quasi-convex bicombing of the flat strip $E(\tau, \lambda)$ as follows.

Let $\eta, \sigma \in E(\tau, \lambda)$. Using the notations from Section 4, let $\zeta = \Pi^1_{E(\tau, \eta)}(\sigma)$. Define γ to be the composition of the inverse of the tight multi-sequence connecting ζ to η with the tight multi-sequence connecting ζ to σ. Define $c_{\eta, \sigma}$ to be the constant speed reparametrization of γ on $[0, 1]$. Note that $c_{\sigma, \eta}$ is just the inverse of $c_{\eta, \sigma}$ and hence this defines a symmetric reflexive L-Lipschitz bicombing of $E(\tau, \lambda)$. By the results of Section 4 and Lemma 6.3, this bicombing is moreover L-quasi-convex for a universal constant $L > 1$. If τ is a train track in standard form for the framing F then we also write $c_{F, \sigma}$ instead of $c_{\tau, \sigma}$.

Now let $\eta \in E(F, \lambda), \beta \in E(F, \mu)$ and write $\zeta = \Pi_{E(F, \eta)}(\beta), \tilde{\zeta} = \Pi_{E(F, \beta)}(\eta)$. We claim that the distance between corresponding points on the curves $c_{F, \zeta}$ and $c_{F, \tilde{\zeta}}$ is uniformly bounded.

For this let σ be a subtrack of a train track τ and let σ' be a train track obtained from σ by a single split at a large branch e. Using the terminology from Section 4, let τ' be the train track which contains σ' as a subtrack and is obtained from τ as follows. Modify τ to a train track $\tilde{\tau}$ obtained from τ by a σ-move at e and let τ' be the train track obtained from $\tilde{\tau}$ by a single split at the tight branch e and which contains σ' as a subtrack. If the σ-complexity $\chi(\tau, \sigma)$ of τ coincides with the σ'-complexity $\chi(\tau', \sigma')$ of τ' then the large branch e of σ defines an embedded

trainpath in τ which is just the τ'-configuration of τ as defined above. Moreover, τ' is obtained from τ by a τ'-move.

Together this shows the following. Let σ be any birecurrent generic train track on S. Assume that σ is splittable to a train track σ'. Then we can define as above a tight multi-sequence $\{\sigma(i)\}_{0 \leq i \leq p}$ connecting $\sigma = \sigma(0)$ to a train track $\sigma' = \sigma(p)$. Let τ be a complete train track which contains σ as a subtrack and let τ' be obtained from τ by a splitting sequence induced from a splitting sequence connecting σ to σ'. Assume that $\chi(\tau, \sigma) = \chi(\tau', \sigma')$ and let $\{\tau(j)\}$ be the tight multi-sequence connecting τ to τ'. Then for every $i \leq p$, the train track $\tau(i)$ contains $\sigma(i)$ as a subtrack. In particular, if η is another complete train track containing σ as a subtrack, if η' is obtained from η by a splitting sequence induced from a splitting sequence connecting σ to σ' and if $\{\eta(j)\}$ is the tight multi-sequence connecting $\eta = \eta(0)$ to $\eta' = \eta(p)$ then the distance between corresponding points on $\{\tau(i)\}, \{\eta(j)\}$ is bounded from above by $Ld(\tau, \eta) + L$ for a universal constant $L > 0$.

Now for every splitting sequence $\{\tau(i)\} \subset \mathcal{V}(\mathcal{TT})$ induced by a splitting sequence $\sigma(j(i))$ of subtracks $\sigma(j(i)) < \tau(i)$ the number of splits $\tau(i) \to \tau(i+1)$ which reduce the complexity, i.e. such that $\chi(\tau(i), \sigma(j(i))) > \chi(\tau(i+1), \sigma(j(i+1)))$, is uniformly bounded. Together with Lemma 6.3 and using the explicit construction of the maps $\Pi_{E(\tau, \eta)}$ and $\Pi_{E(\tau, \beta)}$ we deduce that the distance between corresponding points on the tight splitting sequences connecting τ to β, η is uniformly bounded.

Now define γ to be the composition of the inverse of the tight multi-sequence in $E(\tau, \beta)$ connecting $\tilde{\zeta}$ to β with the tight multi-sequence connecting ζ to η. The curve γ is not continuous but can be made continuous by inserting an arc of uniformly bounded length parametrized on $[0, 1]$ which connects $\tilde{\zeta}$ to ζ. Let $c_{\sigma, \eta}$ be the constant speed reparametrization of γ on $[0, 1]$. By the considerations in Section 4, this defines indeed a reflexive symmetric L-Lipschitz L-quasi-convex bicombing of X. $\qquad\square$

7. The geometric rank

This section is devoted to the proof of Theorem B from the introduction. We use an argument which is motivated by the work of Kleiner and Leeb [KL97].

Choose again a non-principal ultrafilter ω and consider the asymptotic cone $\mathcal{TT}_\omega$ of $\mathcal{TT}$ with respect to ω and basepoint the constant sequence (τ_0) where τ_0 is a train track in standard form for a framing F of S. Then $\mathcal{TT}_\omega$ is a complete geodesic metric space (Lemma 2.5.2 of [KL97]). Since the mapping class group $\mathcal{M}(S)$ acts properly and cocompactly as a group of isometries on $\mathcal{TT}$, the asymptotic cone $\mathcal{TT}_\omega$ is independent of the point τ_0 and admits a transitive group of isometries (Proposition 2.5.6 of [KL97]). If we denote by $X \subset \mathcal{V}(\mathcal{TT})$ the set of all complete train tracks which can be obtained from a train track in standard form for F by a splitting sequence, equipped with the restriction of the metric on $\mathcal{TT}$, then the cone $\mathcal{TT}_\omega$ is bilipschitz equivalent to the asymptotic cone X_ω of X with respect to ω and the basepoint (τ_0). By Corollary 4.10, for every complete geodesic lamination λ carried by a train track $\tau \in \mathcal{V}(\mathcal{TT})$ in standard form for F the inclusion $E(\tau, \lambda) \to$

$\mathcal{TT}$ is a quasi-isometric embedding and hence the cone $\mathcal{TT}_\omega$ contains a uniform bilipschitz image $C(\lambda)$ of the asymptotic cone of the flat strip $E(\tau, \lambda)$ with basepoint the constant sequence (τ). Since $E(\tau, \lambda)$ is uniformly quasi-isometric to its maximal extension $C(\tau, \lambda)$, the set $C(\lambda)$ is uniformly bilipschitz equivalent to the asymptotic cone $C(\tau, \lambda)_\omega$ of $C(\tau, \lambda)$. Therefore by Lemma 5.7, $C(\lambda)$ is locally compact and its topological dimension is bounded from above by $3g - 3 + m$. We call the image of a cone $C(\lambda)$ of this form under an isometry of $\mathcal{TT}_\omega$ a *cone*. Note that each cone in $\mathcal{TT}_\omega$ is an ultralimit of a sequence of flat strips in $\mathcal{TT}$.

For $k \geq 0$ let $\Delta^k = \{(x_1, \ldots, x_{k+1}) \in \mathbb{R}^{k+1} \mid x_i \geq 0, \sum_i x_i = 1\}$ be the standard k-simplex in $\mathbb{R}^{k+1}$. For $i \geq 0$ let Δ^i be the subsimplex of Δ^k which is the standard face of dimension i obtained by intersecting Δ^k with $\mathbb{R}^{i+1} \subset \mathbb{R}^{k+1}$; we have $\Delta^0 \subset \Delta^1 \subset \cdots \subset \Delta^k$. A singular k-simplex in a topological space Y is a continuous map $\sigma : \Delta^k \to Y$. Denote by $C_*(Y)$ the chain complex of singular chains in Y. For a subset V of Y and a number $k > 0$ let $C_k(Y, V)$ be the set of all singular k-chains whose boundaries are contained in V. We use the results from Section 6 to show.

Lemma 7.1. *Let $\emptyset \neq V \subset U \subset \mathcal{TT}_\omega$ be open sets, let $k \geq 1$ and let $\sigma \in C_k(U, V)$ be a singular k-chain with boundary in V. Then there is a singular k-chain $Str(\sigma) \in C_k(U, V)$ with the following properties.*

 (1) *$Str(\sigma)$ and σ define the same class in $H_k(U, V)$.*
 (2) *$Str(\sigma)$ is contained in a finite union of cones.*

Proof. As before, denote by $X \subset \mathcal{TT}$ the set of all train tracks which can be obtained from a train track in standard form for some fixed framing F of S by a splitting sequence. We equip X with the restriction of the metric on $\mathcal{TT}$. Let τ_0 be a train track in standard form for F. Since X is r-dense in $\mathcal{TT}$ for some $r > 0$, for every non-principal ultrafilter ω the ω-asymptotic cone (X_ω, d_ω) of X with basepoint the constant sequence (τ_0) is bilipschitz equivalent and hence homeomorphic to the ω-asymptotic cone of $\mathcal{TT}$ with basepoint (τ_0). In Proposition 6.4 we constructed for some $L > 1$ a reflexive symmetric L-Lipschitz L-quasi-convex bicombing of X. Taking the ω-limits of the combing lines defines a reflexive symmetric L-Lipschitz bicombing of X_ω and hence $\mathcal{TT}_\omega$ for a possibly different constant $L > 0$. If we denote for $x, y \in X_\omega$ by $c_{x,y}$ the combing line connecting x to y, then this bicombing is moreover L-convex in the following sense. For every quadruple x, y, x', y' of points in X_ω and all $t > 0$, we have $d_\omega(c_{x,y}(t), c_{x',y'}(t)) \leq L(d_\omega(x, x') + d_\omega(y, y'))$. In particular, the combing lines depend continuously on their endpoints.

To every singular simplex $\sigma : \Delta^k \to X_\omega$ we associate a *straightened simplex* $Str(\sigma)$ inductively as follows. First let $S^0(\Delta^k)$ be the 0-skeleton of Δ^k consisting of $k + 1$ vertices and define $Str(\sigma)(S^0(\Delta^k)) = \sigma(S^0(\Delta^k))$. Assume by induction that the restriction of $Str(\sigma)$ to the subsimplex Δ^j has been defined for some $j \in \{0, \ldots, k-1\}$. Let x be the vertex of Δ^{j+1} which is not contained in Δ^j and extend $Str(\sigma)$ to Δ^{j+1} by connecting $Str(\sigma)(x)$ to each point in $Str(\sigma)(\Delta^j)$ by the combing line with the same endpoints. By the above observation, this defines a continuous map $Str(\sigma) : \Delta^k \to X_\omega$ which coincides with σ on the vertices of Δ^k. Since our bicombing is symmetric, straightening commutes with the boundary maps. In particular, the boundary of the straightening of σ is the straightening of

the boundary of σ. This means that Str defines a chain map of the chain complex $C_*(Y)$. Moreover, there is a number $L(k) > L$ such that if the diameter of the vertex set of a singular simplex σ is smaller than some $r > 0$ then the diameter of $Str(\sigma)$ is smaller than $L(k)r$.

Now let $\emptyset \neq V \subset U$ be open sets in X_ω and let $\sigma \in C_k(U, V)$ be a singular chain. Since the image of σ is compact, there is a positive lower bound $\delta > 0$ for the distance betwen the image of σ and $X_\omega - U$ and for the distance between the boundary of σ and $X_\omega - V$. After a sufficiently fine barycentric subdivision we may assume that the diameter of each simplex in our chain is at most $\delta/4L^2(k)$. Then the diameter of each singular simplex in the straightened chain $Str(\sigma)$ is at most $\delta/4L(k)$. Since the 0-skeleton of σ and $Str(\sigma)$ coincide, the distance between a point $z \in \sigma$ and the corresponding point in the straightening $Str(\sigma)$ is bounded from above by the sum of the diameter of σ and $Str(\sigma)$ and hence this distance is at most $\delta/2L(k)$. In particular, the singular chain $Str(\sigma)$ is contained in $C_k(U, V)$.

We claim that $Str(\sigma)$ and σ define the same class in $H_k(U, V)$. Namely, connect each point in σ to the corresponding point in its straightening by the combing line connecting these two points. Since the length of a combing line is bounded from above by L times the distance between its endpoints, these combing lines are entirely contained in U, and the combing lines which connect a boundary point of σ to a boundary point of $Str(\sigma)$ are entirely contained in V. Thus the collection of these combing lines define a $k + 1$-chain in $C_{k+1}(U, V)$ with boundary $\sigma - Str(\sigma)$. In other words, the relative cocycles σ, $Str(\sigma)$ are homologous.

Now by construction, each straightened simplex of the chain $Str(\sigma)$ is contained in finitely many cones. More precisely, a vertex v of a singular simplex σ can be represented by a sequence $(x_i) \subset X$. If w is another vertex which is represented by the sequence (y_i) then for each i there is a combing line c_i connecting x_i to y_i, and the combing line $c_{v,w}$ is the ω-limit of the sequence (c_i). In particular, this line is contained in the union of the two cones containing $(x_i), (y_i)$. As a consequence, the straightened chain $Str(\sigma)$ is contained in finitely many cones as well. This completes the proof of the lemma.   □

The following proposition completes the proof of Theorem B and of the corollary from the introduction.

Proposition 7.2. *If $k > 3g - 3 + m$ then $H_k(U, V) = 0$ for all pairs of open sets $V \subset U$ in $\mathcal{TT}_\omega$.*

Proof. Let $\emptyset \neq V \subset U \subset \mathcal{TT}_\omega$ and assume that $H_k(U, V) \neq 0$ for some $k \geq 1$. Then there is a singular k-chain $c = \sum_i a_i c_i$ for some $a_i \in \mathbb{Z}$ and for continuous maps $c_i : \Delta^k \to \mathcal{TT}_\omega$ whose boundary is contained in V and such that this chain is not homologous to a chain in V. By Lemma 7.1 we may assume without loss of generality that the chain c is straightened. This means in particular that c is contained in a finite union of cones. These cones are embedded in $\mathcal{TT}_\omega$ and are glued along closed subsets. By Lemma 5.7 the topological dimension of a cone in X_ω is not bigger than $3g-3+m$. In other words, σ defines a nontrivial relative homology class in $H_k(U', V')$ where U' is an open subset in a topological space obtained from

glueing a finite disjoint collection of standard proper cones of dimension at most $3g - 3 + m$ along a finite collection of closed subsets. But this just means that the topological dimension of the set U' is at most $3g - 3 + m$ and hence we necessarily have $k \leq 3g - 3 + m$. This completes the proof of the proposition. $\qquad \square$

Now let $k \geq 1$, let $c > 1$ and let $\eta : \mathbb{R}^k \to \mathcal{TT}$ be a c-quasi-isometric embedding with $\eta(0) = \tau_0$ for our basepoint τ_0. Then the ω-asymptotic cone of $\mathbb{R}^k$ admits a bi-Lipschitz embedding into the asymptotic cone $\mathcal{TT}_\omega$ of $\mathcal{TT}$. Thus there is a bilipschitz embedding of $\mathbb{R}^k$ into $\mathcal{TT}_\omega$. Since $\mathbb{R}^k$ is an absolute retract, there are open subsets $U \supset V$ in $\mathcal{TT}_\omega$ such that the relative homology group $H_k(U, V)$ is non-trivial. By Proposition 7.2 this means that $k \leq 3g - 3 + m$ which shows the corollary from the introduction.

Corollary 7.3. *The geometric rank of $\mathcal{M}(S)$ equals $3g - 3 + m$.*

8. Quasi-isometric rigidity

This section is devoted to the proof of Theorem A from the introduction. We call a finitely generated group Γ *quasi-isometrically rigid* [M03b] if for every finitely generated group H which is quasi-isometric to Γ there is a finite index subgroup H' of H and a homomorphism of H' with finite kernel onto a subgroup of Γ of finite index. Our goal is to show that $\mathcal{M}(S)$ is quasi-isometrically rigid; in the case of once-punctured surfaces (i.e. if $m = 1$) this result is due to Mosher and Whyte (see [M03b]).

The *curve graph* $\mathcal{C}(S)$ of S is the locally infinite metric graph whose vertices are the free homotopy classes of essential simple closed curves on S, i.e. simple closed curves which are neither contractible nor freely homotopic into a puncture, and where two such vertices c_1, c_2 are connected by an edge if and only if the curves c_1, c_2 can be realized disjointly. There is a natural homomorphism ρ from the *extended mapping class group* of *all* isotopy classes of homeomorphisms of S into the group $\mathrm{Aut}(\mathcal{C}(S))$ of simplicial automorphisms of $\mathcal{C}(S)$. By a result of Ivanov (see [I02]) and Luo [L00], if S is different from the closed surface of genus 2 and the twice punctured torus, then ρ is an isomorphism. If S is a closed surface of genus 2 then ρ is surjective, with kernel the group $\mathbb{Z}/2\mathbb{Z}$ generated by the hyperelliptic involution. If S is the twice punctured torus, then the kernel of ρ is again the subgroup $\mathbb{Z}/2\mathbb{Z}$ generated by the hyperelliptic involution, and the image of ρ is a subgroup of index 5 in $\mathrm{Aut}(\mathcal{C}(S))$. As a consequence, for the purpose of our theorem it suffices to construct for every finitely generated group Γ which is quasi-isometric to $\mathcal{M}(S)$ a homomorphism $\rho : \Gamma \to \mathrm{Aut}(\mathcal{C}(S))$ with finite kernel and finite index image.

We begin with constructing a *Tits boundary* $\mathcal{TB}$ for $\mathcal{M}(S)$. For this let λ be a complete geodesic lamination with $k \geq 1$ minimal components $\lambda_1, \ldots, \lambda_k$. After reordering we may assume that for some $s \leq k$ the laminations $\lambda_1, \ldots, \lambda_s$ are minimal arational and the laminations $\lambda_{s+1}, \ldots, \lambda_k$ are simple closed geodesics. For $i \leq k$ the complete lamination λ determines a sign $\mathrm{sgn}_\lambda(\lambda_i) \in \{+, -\}$ for λ_i as follows. If λ_i is minimal arational then we define the sign to be positive. If λ_i is a simple closed curve then for a given orientation of λ_i, the orientation of

S determines the right and the left side of λ_i in a tubular annulus about λ_i. The complete lamination λ contains at least one leaf which spirals about λ_i from the left side. If this spiraling leaf approaches λ_i in the direction given by the orientation of λ_i then we choose the sign to be positive, otherwise the sign is chosen to be negative. Since λ is complete by assumption, this choice of sign does not depend on the orientation of λ_i used to define it (see the discussion in [H06a]). We simply write sgn_λ for this collection of signs or also sgn if no confusion is possible.

Using the notations from Section 5, for $i \leq s$ denote by $A(\lambda_i)$ the asymptotic cone of a flat strip $C(\zeta, \lambda_i)$ where ζ is a complete train track on the characteristic surface S_i for λ_i which carries λ_i. Recall that $A(\lambda_i)$ is a proper CAT(0) cone defined by a compact CAT(1)-space $\partial A(\lambda_i)$, and it does not depend on ζ up to uniform bilipschitz identification. Let again $\Delta^{k-1} = \{(x_1, \ldots, x_k) \in \mathbb{R}^k \mid x_i \geq 0, \sum_i x_i = 1\}$ be the standard $k-1$-dimensional simplex in $\mathbb{R}^k$ and write $\Delta(\lambda) = \{\sum_i \mathrm{sgn}_\lambda(\lambda_i) x_i \mu_i \mid \mu_i \in \partial A(\lambda_i), (x_1, \ldots, x_k) \in \Delta^{k-1}\}$. The space $\Delta(\lambda)$ is uniquely determined by λ up to permutations of the minimal components of λ and therefore the topology on $\Delta(\lambda)$ induced from the topology on Δ^{k-1} and the compact Cat(1)-spaces $\partial A(\lambda_i)$ is independent of any choices. More precisely, with this topology the space $\Delta(\lambda)$ is homeomorphic to the spherical join $\partial A(\lambda_1) * \cdots * \partial A(\lambda_k)$ of the spaces $\partial A(\lambda_i)$ $(i \leq k)$ and hence $\Delta(\lambda)$ is homeomorphic to the boundary $\partial C(\tau, \lambda)_\omega$ of the asymptotic cone of $C(\tau, \lambda)$. We equip the collection $\tilde{\mathcal{T}\mathcal{B}} = \{\Delta(\lambda) \mid \lambda \in \mathcal{C}\mathcal{L}\}$ with the topology as a disjoint union of the spaces $\Delta(\lambda)$, i.e. $\tilde{\mathcal{T}\mathcal{B}}$ is naturally a disjoint union of compact Cat(1)-spaces of dimension at most $3g - 4 + m$ (compare Section 5).

Define an equivalence relation $\sim$ on $\tilde{\mathcal{T}\mathcal{B}}$ as follows. Let again $\lambda_1, \ldots, \lambda_k$ be the minimal components of the complete geodesic lamination λ and let $\mu_1, \ldots, \mu_m$ be the minimal components of the complete geodesic lamination μ. After reordering there is some $\ell \leq \min\{m, k\}$ such that the following holds.

(1) $\mu_i = \lambda_i$ for all $i \leq \ell$.
(2) If λ_i is a simple closed curve for some $i \leq \ell$ then $\mathrm{sgn}_\lambda(\lambda_i) = \mathrm{sgn}_\mu(\lambda_i)$.
(3) If there is some $i > \ell$ and a component λ_i of λ which coincides with a component μ_j of μ for $j > \ell$ then $\mathrm{sgn}_\lambda(\lambda_i) \neq \mathrm{sgn}_\mu(\mu_j)$.

Then both $\Delta(\lambda)$ and $\Delta(\mu)$ contain the (signed) spherical join $\partial A(\lambda_1) * \cdots * \partial A(\lambda_\ell)$ as a closed subspace, and we define $x \in \Delta(\lambda)$ to be equivalent to $y \in \Delta(\mu)$ if and only if x, y are both contained in this signed spherical join and define the same point there.

By the definition of our topology on $\tilde{\mathcal{T}\mathcal{B}}$, the equivalence relation $\sim$ is closed. It identifies the spaces $\Delta(\lambda), \Delta(\mu)$ for all complete geodesic laminations λ, μ which contain the same unique minimal component λ_0 which fills up S, i.e. which is such that every simple closed curve on S intersects λ_0 transversely. Note that the number of such spaces which are identified in this way with a fixed space $\Delta(\lambda)$ is bounded from above by a universal constant. The quotient space $\mathcal{T}\mathcal{B} = \tilde{\mathcal{T}\mathcal{B}} / \sim$ has naturally the structure of a (locally infinite) complex of dimension $3g - 4 + m$; it is closely related to the curve graph on S (compare [MM99] for a discussion of the curve complex). Note that we do not claim that $\mathcal{T}\mathcal{B}$ is a cell complex in the usual

sense. Moreover, $\mathcal{TB}$ has infinitely many connected components. Namely, every minimal geodesic lamination ν which fills up S defines a connected component of $\mathcal{TB}$ which is homeomorphic to a single compact Cat(1)-space. However, since the curve graph is connected, the components of $\tilde{\mathcal{TB}}$ defined by geodesic laminations with a minimal component which does not fill up S all map to the same connected component $\mathcal{TB}_0$ of $\mathcal{TB}$.

We call the image in $\mathcal{TB}$ of a set $\Delta(\lambda) \subset \tilde{\mathcal{TB}}$ a *cell*. Every spread out geodesic lamination on S defines a cell which is naturally homeomorphic to a standard $3g - 4 + m$-dimensional simplex, and we call these cells *chambers*. The vertices of a chamber in $\mathcal{TB}$ either correspond to signed simple closed geodesics or to minimal arational geodesic laminations which fill a bordered subtorus or an X-piece of S, i.e. a bordered punctured sphere of Euler characteristic -2. Denote by $\mathcal{TB}_1$ the subcomplex of $\mathcal{TB}$ which is the closure of the chambers. The complex $\mathcal{TB}_1$ is a connected simplicial complex in the usual sense.

If S is not a closed surface of genus 2 or a twice punctured torus then we define $\mathcal{M}_0(S)$ to be the extended mapping class group of all isotopy classes of *any* homeomorphism of S including the orientation reversing ones. For a closed surface S of genus 2 or a twice punctured torus we define $\mathcal{M}_0(S)$ to be the quotient of the extended mapping class group under the hyperelliptic involution. Then $\mathcal{M}_0(S)$ naturally acts on $\mathcal{TB}$ as a group homeomorphisms preserving the cell structure and the subcomplex $\mathcal{TB}_1$.

As before, denote by $\mathrm{Aut}(\mathcal{C}(S))$ the group of simplicial automorphisms of the curve graph $\mathcal{C}(S)$ of S. The automorphism group $\mathrm{Aut}(\mathcal{C}(S))$ naturally contains the group $\mathcal{M}_0(S)$ as a subgroup. In fact, equality holds if S is not a twice punctured torus. Using this fact, the next lemma gives a description of the group of isotopy classes of homeomorphisms of $\mathcal{TB}$.

Lemma 8.1.   (1) *Every homeomorphism of $\mathcal{TB}$ preserves the subcomplex $\mathcal{TB}_1$.*
   (2) *There is an injective homomorphism of the group of isotopy classes of homeomorphisms of $\mathcal{TB}_1$ into $\mathrm{Aut}(\mathcal{C}(S))$ whose restriction to $\mathcal{M}_0(S)$ is the identity.*

Proof. Let φ by an arbitrary homeomorphism of $\mathcal{TB}$. Since $\mathcal{TB}$ is the disjoint union of the locally infinite connected subcomplex $\mathcal{TB}_0$ and infinitely many compact components, φ preserves $\mathcal{TB}_0$. We claim that φ also preserves $\mathcal{TB}_1$.

For this let λ be a complete geodesic lamination with minimal components $\lambda_1, \ldots, \lambda_k$ which defines a cell $\Delta(\lambda)$ in $\mathcal{TB}$ of dimension $3g - 4 + m$ (by abuse of notation we use now the same symbol for a the space $\Delta(\lambda) \subset \tilde{\mathcal{TB}}$ and its image cell in $\mathcal{TB}$). After reordering, there is a number $s \geq 0$ such that the components $\lambda_1, \ldots, \lambda_s$ are precisely those simple closed curve components of λ which are contained in the boundary of a characteristic subsurface S_j of a minimal arational component λ_j of λ. Then for every $i \leq s$, *every* complete geodesic lamination μ which defines a cell $\Delta(\mu)$ of dimension $3g - 4 + m$ and which contains each of the minimal components λ_j for $j \neq i$ also contains λ_i.

Recall that a point in $\Delta(\lambda)$ defines a tuple $(s_1, \ldots, s_k) \in \Delta^{k-1}$ and a tuple $(x_1, \ldots, x_k) \in \partial A(\lambda_1) \times \cdots \times \partial A(\lambda_k)$. By the above consideration, for each $i \leq s$ the set of all points in $\Delta(\lambda)$ corresponding to a tuple $(s_1, \ldots, s_k) \in \Delta^{k-1}$ with $s_i = 0$ and $s_j > 0$ for $j \neq i$ is contained in the boundary of precisely two cells of maximal dimension, namely one cell for each choice of a sign for λ_i. If x has a neighborhood in $\Delta(\lambda)$ which is homeomorphic to a closed half-space in $\mathbb{R}^{3g-4+m}$ containing x in its boundary (which is always the case if $\Delta(\lambda)$ is a chamber), then x has a neighborhood in $\mathcal{TB}_0$ which is homeomorphic to $\mathbb{R}^{3g-4+m}$. Call the union of all cells in $\mathcal{TB}_0$ of maximal dimension whose intersection with $\Delta(\lambda)$ is of this form the *star* of $\Delta(\lambda)$. Define moreover inductively the *multi-cell* $R(\lambda)$ containing $\Delta(\lambda)$ to be the smallest subcomplex of $\mathcal{TB}_0$ which contains $\Delta(\lambda)$ and which contains with each cell its star. Note that a multi-cell consists of a uniformly bounded number of cells. Namely, if $\mu \in \mathcal{CL}$ defines a cell contained in the multi-cell $R(\lambda)$ then the minimal components of μ coincide with the minimal components of λ. If the star of a cell coincides with the cell itself then we also call this single cell a multi-cell.

We claim that a point $x \in \mathcal{TB}_0$ which is not contained in the interior of a multi-cell does not have a neighborhood in $\mathcal{TB}_0$ which is homeomorphic to an open subset of $\mathbb{R}^{3g-4+m}$. Namely, let $x \in \mathcal{TB}_0$ be any point admitting a neighborhood in $\mathcal{TB}_0$ which is homeomorphic to an open subset of $\mathbb{R}^{3g-4+m}$. If x is *not* an interior point of a multi-cell in $\mathcal{TB}_0$, i.e. if no neighborhood of x in $\mathcal{TL}_0$ is entirely contained in a multi-cell, then x is necessarily contained in the boundary of a cell of maximal dimension. Thus assume that x is contained in the boundary of the cell $\Delta(\lambda)$ for some $\lambda \in \mathcal{CL}$. Let $\lambda_1, \ldots, \lambda_k$ be the minimal components of λ. Then up to reordering, the point x defines a tuple $(s_1, \ldots, s_k) \in \Delta^{k-1}$ and a tuple $(x_1, \ldots, x_k) \in \partial A(\lambda_1) \times \cdots \times \partial A(\lambda_k)$ where $s_1 = 0$.

Next assume that λ_1 is a minimal arational geodesic lamination with characteristic surface S_1. Since the dimension of $\Delta(\lambda)$ equals $3g - 4 + m$ by assumption, there are infinitely many pants decompositions of S_1 whose union with $\cup_{j \geq 2} \lambda_j$ determine a cell in $\mathcal{TB}_0$ of maximal dimension. Then x is contained in the boundary of infinitely many such cells and hence a neighborhood of x in $\mathcal{TB}_0$ is not locally compact. As a consequence, necessarily λ_1 is a simple closed geodesic.

If λ_1 is a simple closed geodesic which is not contained in the boundary of the characteristic surface of a minimal arational component $\lambda_i \neq \lambda_1$ of λ then there is a bordered subsurface S_1 of S of negative Euler characteristic containing λ_1 which does not have an essential intersection with $\cup_{i \geq 2} \lambda_i$. Thus there are infinitely many pairwise distinct simple closed geodesics which do not intersect $\cup_{i \geq 2} \lambda_i$ and which define together with $\lambda_2, \ldots, \lambda_k$ a cell of maximal dimension. As a consequence, x is contained in infinitely many distinct such cells which contradicts our assumption that there is a neighborhood of x in $\mathcal{TB}_0$ which is homeomorphic to a ball in $\mathbb{R}^{3g-4+m}$. This shows that indeed a point $x \in \mathcal{TB}_0$ has a neighborhood in $\mathcal{TB}_0$ which is homeomorphic to $\mathbb{R}^{3g-4+m}$ only if x is an interior point of a multi-cell. In particular, every homeomorphism of $\mathcal{TB}_0$ preserves the union of all interior points of all such multi-cells in $\mathcal{TB}_0$.

Our next goal is to give a topological characterization of multi-cells defined by spread out geodesic laminations. Thus let λ be a spread out geodesic lamination with minimal components $\lambda_1, \ldots, \lambda_{3g-3+m}$ which defines the multi-cell $R(\lambda)$. Then

$R(\lambda)$ has a canonical structure of a finite simplicial complex. The boundary of $R(\lambda)$ is partitioned into finitely many sides, i.e. simplices of codimension one. Such a side either is defined by a minimal arational component of λ or by a simple closed geodesic which is not a boundary component of the characteristic surface of a minimal arational minimal component of λ. Thus a neighborhood in $\mathcal{TB}$ of an interior point x of such a side contains a set which is homeomorphic to a countable collection of closed half-spaces in $\mathbb{R}^{3g-4+m}$ glued along their boundary, and x is contained in the interior of the boundary of these half-spaces. In particular, such a neighborhood is not homeomorphic to the neighborhood of a point in a side of a chamber of codimension at least 2. In other words, there is an open dense subset V of the boundary of the multi-cell $R(\lambda)$ which consists of $p > 0$ open disjoint sides, whose closure is the whole boundary of $R(\lambda)$ and which admits a purely topological characterization. On the other hand, if $R(\lambda)$ is a multi-cell and if λ is *not* spread out, then $R(\lambda)$ contains points which admit a compact neighborhood in $R(\lambda)$ not containing any neighborhood which is homeomorphic to an open subset of $\mathbb{R}^{3g-3+m}$. But this just means that multi-cells defined by chambers in $\mathcal{TB}$ admit a topological characterization and hence the image of a multi-cell defined by a chamber under a homeomorphism φ of $\mathcal{TB}$ is again a multi-cell defined by a chamber. As a consequence, every homeomorphism φ of $\mathcal{TB}$ restricts to a homeomorphism of the simplicial complex $\mathcal{TB}_1$ which is homotopic to a simplicial map.

Call a vertex v of $\mathcal{TB}$ *simple* if v is defined by a signed simple closed curve. For a vertex a of $\mathcal{TB}_1$ let $\mathcal{F}(a)$ be the collection of all multi-cells containing a; we claim that a is simple if and only if for every $k \geq 2$ there are multi-cells $F_1, \ldots, F_k \in \mathcal{F}(a)$ with $F_i \cap F_j = \{a\}$ for all $i \neq j$. Namely, if a is simple then for every $k \geq 1$ we can find k pants decompositions $P_1, \ldots, P_k$ of S containing a as one of their pants curves and such that every pants curve $c \in P_i - \{a\}$ intersects at least one of the curves in the collection $P_j - \{a\}$ transversely. As a consequence, if C_i, C_j are two chambers in $\mathcal{TB}_1$ defined by the pants decompositions P_i, P_j and a system of signs which coincide for the vertex a then we have $C_i \cap C_j = \{a\}$ by the definition of our complex $\mathcal{TB}$.

On the other hand, if a is a vertex of the simplicial complex $\mathcal{TB}_1$ which is not simple then a is defined by a minimal arational geodesic lamination which fills a non-trivial subsurface S_0 of S. More precisely, there is a unique bordered subsurface S_0 of S with boundary $\partial S_0 \neq \emptyset$ and with the additional property that every simple closed curve on S which has an essential intersection with ∂S_0 also has an essential intersection with a. Since a is a vertex of a chamber, the surface S_0 either is a once punctured torus or a forth punctured sphere. Therefore every multi-cell of $\mathcal{TB}_1$ which contains a as a vertex also contains each component of the boundary of S_0 equipped with a choice of a sign. Since the boundary of S_0 consists of at most 4 components, if $F_1, \ldots, F_{2^4+1}$ are $2^4 + 1$ multi-cells containing the vertex a then at least two of them, say the multi-cells C_1, C_2, contain the boundary ∂S_0 of S_0 equipped with the same collection of signs and hence the intersection $C_1 \cap C_2$ contains an edge. As a consequence, simple vertices can be distinguished from non-simple ones by the multi-cells they are contained in and therefore a simplicial automorphism of $\mathcal{TB}_1$ has to preserve the set of simple vertices.

Recall that every simple vertex of $\mathcal{TB}_1$ consists of a simple closed curve on S together with a sign. In other words, a simple closed curve a gives rise to two distinct vertices $(a,+),(a,-)$ which differ by their sign. We claim that a homeomorphism of $\mathcal{TB}_1$ preserves the set of pairs $((a,+),(a,-))$ of such vertices. For this we argue as before. Namely, let a be a simple closed geodesic and let $(a,+)$ be any simple vertex of $\mathcal{TB}$ defined by a (here we mean that $+$ is the sign of our vertex). Let $C \subset \mathcal{TB}_1$ be any chamber containing $(a,+)$. Then the vertices of the side F of C opposite to a correspond to a geodesic lamination with $3g - 4 + m$ minimal components. The side F is also a side of a chamber which contains the simple vertex $(a,-)$. In other words, for *every* chamber C_+ containing $(a,+)$ as a vertex, the point $(a,-)$ is a vertex of a chamber C_- with the same opposite side.

Now let C be a chamber which contains a vertex x defined by a minimal arational component filling up a subsurface which contains the simple closed curve a in its boundary; note that such a chamber always exists. Then a is contained in the boundary of a multi-cell which is not a chamber, and this multi-cell contains both vertices $(a,+),(a,-)$. The multi-cell is distinguished from a chamber by the number of its sides, and the intersection of all these multi-cells consists precisely of the two points $(a,+),(a,-)$. This shows that the pair of vertices $(a,+),(a,-)$ is characterized by its intersection pattern with the chambers of $\mathcal{TB}_1$. Therefore every simplicial automorphism of $\mathcal{TB}_1$ preserves the pairs of simple vertices determined by a simple closed geodesic on S. Thus such an automorphism φ induces a bijection of the 0-skeleton of the curve graph $\mathcal{C}(S)$ of S.

This bijection preserves disjointness for simple closed curves. Namely, two simple closed curves on S can be realized disjointly if and only if they define two distinct vertices of a common chamber. In other words, a homeomorphism φ of $\mathcal{TB}_1$ defines a simplicial automorphism of $\mathcal{C}(S)$. This completes the proof of our lemma. $\square$

Let $\mathcal{TC}$ be the (simplicial) cone over the Tits boundary $\mathcal{TB}$ (compare [KL97]). Define a *chamber* in $\mathcal{TC}$ to be the cone over a chamber in $\mathcal{TB}$. Note that a chamber is a $3g - 3 + m$-dimensional standard proper cone. Define moreover an *apartment* in $\mathcal{TC}$ to be a union of chambers which is homeomorphic to $\mathbb{R}^{3g-3+m}$. Every pants decomposition P of S defines 2^{3g-3+m} chambers, one for each choice of sign combination for our pants curves, and the union of all these chambers defines an apartment. The cone over a cell in $\mathcal{TB}$ is contained in an apartment if and only if it is a chamber. Namely, we saw in the proof of Lemma 8.1 that a cell in $\mathcal{TB}$ of maximal dimension which is not a chamber contains points which are not contained in subsets of $\mathcal{TB}$ homeomorphic to $\mathbb{R}^{3g-4+m}$ and hence such a cell can not be contained in an apartment. On the other hand, if λ is a spread out complete geodesic lamination with minimal components $\lambda_1,\ldots,\lambda_{3g-3+m}$ and minimal arational components $\lambda_1,\ldots,\lambda_s$ then for each $i \leq s$ we can find a simple closed curve c_i which does not intersect λ_j for $j \neq i$ and such that the curves $c_1,\ldots,c_s$ are disjoint. In particular, for every subset $\{\lambda_{i_1},\ldots,\lambda_{i_\ell}\}$ of the set $\{\lambda_1,\ldots,\lambda_s\}$ for some $\ell \leq s$ we obtain a new chamber by replacing each of the components λ_{i_j} of λ by the signed simple closed curve $(c_{i_j},+)$ $(j \leq \ell)$. A signed simple closed curve component λ_i of λ $(i \geq s+1)$ can be replaced by the same curve with the opposite sign. The union of the chambers defined by λ with all those chambers obtained from all possible combinations of such replacements defines an apartment by construction.

Choose again a non-principal ultrafilter ω. Fix a framing F for S and let X be the set of all train tracks which can be obtained from a train track in standard form for F by a splitting sequence. Choose a train track τ_0 in standard form for F and use the constant sequence (τ_0) as a basepoint for the asymptotic cone $\mathcal{TT}_\omega$. Recall that for every complete geodesic lamination $\lambda \in \mathcal{CL}$ the flat strip $E(F, \lambda)$ is quasi-isometric to its maximal extension $C(F, \lambda)$ which is a Cat(0)-space, moreover $E(F, \lambda)$ is quasi-isometrically embedded in $\mathcal{TT}$. As a consequence, the ω-asymptotic cone $C(F, \lambda)_\omega$ of $C(F, \lambda)$ with basepoint a constant sequence is topologically embedded in the ω-asymptotic cone $\mathcal{TT}_\omega$ of $\mathcal{TT}$. The cone $C(F, \lambda)_\omega$ is homeomorphic to the euclidean cone over the cell $\Delta(\lambda)$. If $\lambda_1, \ldots, \lambda_k$ are the minimal components of λ then $C(F, \lambda)_\omega = \partial A(\lambda_1) * \cdots * \partial A(\lambda_k)$ where for each i, $\partial A(\lambda_i)$ is a compact connected Cat(1)-space.

If $\Pi_{E(F,\lambda)} : X \to E(F, \lambda)$ is the projection as in Proposition 4.9 then for every complete geodesic lamination ν the projection $\Pi_{E(F,\lambda)}E(F, \nu)$ of $E(F, \nu)$ is a combinatorially convex subset of $E(F, \lambda)$ and hence its maximal extension is a complete Cat(0)-space. The Hausdorff distance in $\mathcal{TT}$ between the graph $\Pi_{E(F,\nu)}E(F, \lambda)$ and the graph $\Pi_{E(F,\lambda)}E(F, \nu)$ is uniformly bounded. The asymptotic cone $E(F, \lambda)_\omega$ of $E(F, \lambda)$ contains the asymptotic cone of the projection $\Pi_{E(F,\lambda)}E(F, \nu)$. Let $\lambda_1, \ldots, \lambda_s$ be those minimal components of λ which are minimal components of ν as well and with the additional property that for every minimal component λ_i which is a simple closed curve, the sign $\mathrm{sgn}_\lambda(\lambda_i)$ defined as above by λ for λ_i coincides with the sign $\mathrm{sgn}_\nu(\lambda_i)$ defined by ν for λ_i. By Lemma 5.7, the asymptotic cone of the projection $\Pi_{E(F,\lambda)}E(F, \nu)$ is the subset of $E(F, \lambda)_\omega$ which is the cone over $\partial A(\lambda_1) * \cdots * \partial A(\lambda_s)$. As a consequence, our Tits cone $\mathcal{TC}$ is topologically embedded in the asymptotic cone $\mathcal{TT}_\omega$ of $\mathcal{TT}$.

As in Section 6.3 of [KL97], for a point $z \in Z$ we say that two subsets S_1, S_2 of Z *have the same germ at z* if $S_1 \cap U = S_2 \cap U$ for some neighborhood U of z. The equivalence classes of subsets with the same germ at z will be denoted $\mathrm{Germ}_z Z$. Write $k = 3g - 3 + m$ for short. For $x \in \mathcal{TT}_\omega$ consider the collection $\mathcal{S}_1(x)$ of all germs of topological embeddings of $\mathbb{R}^k$ into $\mathcal{TT}_\omega$ passing through x. Note that each such germ determines a local homology class of degree k whose support contains x. Let $\mathcal{S}_2(x)$ be the lattice of germs generated by $\mathcal{S}_1(x)$ under finite intersection and union. The following lemma is the analog of Lemma 6.3.1 of [KL97].

Lemma 8.2. *The lattice $\mathcal{S}_2(x)$ admits a natural embedding into the lattice $\mathcal{KTC}$ of the Tits cone $\mathcal{TC}$ for $\mathcal{TT}_\omega$ generated by the cells of maximal dimension under finite union. The image of this embedding contains the sublattice of the Tits cone $\mathcal{TC}$ generated by the chambers under finite intersection and union.*

Proof. Following Kleiner and Leeb (Section 6 of [KL97]), for a subset Y of a topological space Z and for $[c] \in H_k(Z, Y)$ define $\mathrm{Supp}(Z, Y, [c]) \subset Z - Y$ to be the set of points $z \in Z - Y$ such that the image of $[c]$ in the local homology group $H_k(Z, Z - \{z\})$ is nonzero. Then $\mathrm{Supp}(Z, Y, [c])$ is a closed subset of $Z - Y$ contained in the image of any chain c representing the relative class $[c]$.

Since the isometry group of $\mathcal{TT}_\omega$ acts transitively, it is sufficient to show the claim of the lemma for the lattice $\mathcal{S}_2(*)$ defined by the basepoint $*$ of $\mathcal{TT}_\omega$.

By Lemma 6.2.1 of [KL97], every germ of a topological embedding of $\mathbb{R}^k$ through $*$ defines a nontrivial class in $H_k(\mathcal{TT}_\omega, \mathcal{TT}_\omega) - \{*\}$, i.e. there is an open subset U of $\mathcal{TT}_\omega - \{*\}$ and a singular chain $c \in C_k(\mathcal{TT}_\omega, U)$ whose support contains $*$.

Recall from Section 7 the construction of straightening which associates to a singular chain c representing $[c]$ the straightened chain $Str(c)$. Via passing to a sufficiently small barycentric subdivision we may assume that the boundary of the straightening $Str(c)$ of c is contained in U. By Lemma 7.1, the chain $Str(c)$ is contained in a finite union $\mathcal{P} = \cup_{i=1}^\ell P_i$ of cones P_i which intersect along their boundaries. Now the support of the class $[c] \in H_k(\mathcal{TT}_\omega, U)$ is contained in the image of $Str(c)$, on the other hand it is defined by the germ of our embedding of $\mathbb{R}^k$ into $\mathcal{TT}_\omega$. Since the dimension of each cone is at most k, elements of $\mathcal{S}_1(*)$ define finite unions of cells of maximal dimension in $\mathcal{TC}$.

On the other hand, by the discussion in the beginning of this section, each chamber is a finite intersection of apartments and hence of elements of $\mathcal{S}_1(x)$. Intersections of chambers yield sides of the Tits cone, so we have a well defined inclusion of lattices $\Theta : \mathcal{S}_2(x) \to \mathcal{KTC}$ containing the sublattice generated by the apartments under finite union and intersection. $\square$

If $\varphi : \mathcal{TT}_\omega \to \mathcal{TT}_\omega$ is any homeomorphism which fixes the basepoint $* = (\tau_0)$ then φ induces a homeomorphism of lattices $\mathcal{S}_2(*) \to \mathcal{S}_2(*)$ and therefore by Lemma 8.1 and Lemma 8.2, this homeomorphism defines an element of $\text{Aut}(\mathcal{C}(S))$. We state this fact as a corollary.

Corollary 8.3. *There is a homomorphism from the group of homeomorphisms of $\mathcal{TT}_\omega$ which fix the basepoint $*$ into the group $\text{Aut}(\mathcal{C}(S))$.*

Proof. If φ is a homeomorphism of $\mathcal{TT}_\omega$ fixing $*$ then by our above discussion, φ determines a homeomorphism $\tilde{\varphi}$ of subcone of the Tits cone $\mathcal{TC}$ containing the simplicial cone over the complex $\mathcal{TB}_1$. We then associate to φ the element $\rho(\varphi) \in \text{Aut}(\mathcal{C}(S))$ whose restriction to the cone over $\mathcal{TB}_1$ coincides with the restriction of $\tilde{\varphi}$ up to an isotopy preserving the simplicial structure of $\mathcal{TB}_1$. $\square$

Now we are ready to complete the proof of Theorem A from the introduction. Namely, let Γ be a finitely generated group with a word norm $||$ defined by a finite symmetric set of generators. The norm $||$ defines a distance function d on Γ via $d(g, h) = |g^{-1}h|$ which is invariant under left translation. With respect to this distance, the group Γ acts on itself as a group of isometries by left translation. Assume that Γ is quasi-isometric to $\mathcal{M}(S)$, i.e. that there is a quasi-isometry $\Theta_0 : \Gamma \to \mathcal{M}(S)$. Since $\mathcal{M}(S)$ is equivariantly quasi-isometric to $\mathcal{TT}$ [H06a] and hence it is quasi-isometric to X, there is a quasi-isometry $\Theta : \Gamma \to X \subset \mathcal{V}(\mathcal{TT})$ with $\Theta(e) = \tau_0$ (here e denotes the unit) and inverse $\Lambda : X \to \Gamma$ where τ_0 is a train track in standard form for the framing F defining X. Via this quasi-isometry, the group Γ induces a quasi-action as a group of uniform quasi-isometries on X as follows. The quasi-isometry determined by $g \in \Gamma$ is the map $\varphi(g)$ defined by $\varphi(g)(\eta) = \Theta \circ g \circ \Lambda(\eta)$ where g acts on Γ by left translation. By construction, there is a universal constant $L > 0$ such that $d(\varphi(g)\varphi(h)(\eta), \varphi(gh)(x)) \leq L$ for all $g, h \in \Gamma$ and all $x \in X$ (compare the discussion in [M03b]). The quasi-action of Γ on

X then induces an action of Γ as a group of uniformly bilipschitz homeomorphisms on the asymptotic cone $\mathcal{TT}_\omega = X_\omega$. Since the basepoint $* = (\tau_0)$ is the ultralimit of both the constant sequence (τ_0) and the constant sequence $(\Theta \circ g \circ \Lambda(\tau_0))$, this action preserves the basepoint $* = (\tau_0)$. By Corollary 8.3 there is a homomorphism of the group of homeomorphisms of $\mathcal{TT}_\omega$ preserving $*$ into the group $\mathrm{Aut}(\mathcal{C}(S))$ and therefore we obtain a homomorphism $\Gamma \to \mathrm{Aut}(\mathcal{C}(S))$. We summarize our discussion as follows.

Lemma 8.4. *Let* Γ *be quasi-isometric to* $\mathcal{M}(S)$*; then there is a homomorphism* $\rho : \Gamma \to \mathrm{Aut}(\mathcal{C}(S))$.

For the proof of the theorem in the introduction we are left with showing that the kernel of our homomorphism is finite and that its image is of finite index.

We follow again [KL97]. Namely, the following result is the analog of Proposition 7.1.1 of [KL97]. For its formulation, for a constant $D > 0$ define a *D-Hausdorff envelope* of a set $A \subset X$ to be a set B containing A in its D-neighborhood. Define a *maximal quasi-flat* in $\mathcal{TT}$ to be the image under an element of $\mathcal{M}(S)$ of a finite union of flat strips which is uniformly quasi-isometric to $\mathbb{R}^{3g-3+m}$. Then the asymptotic cone of such a maximal quasi-flat with respect to a basepoint defined by a constant sequence is an apartment in the Tits cone $\mathcal{TC}$. Define more generally an *apartment* (or a *chamber*) in the asymptotic cone $\mathcal{TT}_\omega$ to be the image of an apartment (or a chamber) in the Tits cone $\mathcal{TC}$ under an isometry of $\mathcal{TT}_\omega$. Then an apartment consists of finitely many chambers. A chamber is bilipschitz equivalent to a standard partition cone of dimension $3g - 3 + m$. This cone is determined by its boundary which admits a natural identification with a compact convex subset in the standard unit sphere S^{3g-4+m} with dense interior. For a number $\delta > 0$ define a *δ-truncated chamber* to be the closed subcone of a chamber A defined by the compact convex subset of the boundary ∂A of A consisting of all points whose spherical distance to the boundary of ∂A in S^{3g-4+m} is at least δ. A *δ-truncated apartment* is obtained from an apartment by replacing a chamber by the δ-truncated chamber contained in its interior. We define moreover a *δ-truncated maximal quasi-flat* in $\mathcal{TT}$ to be a subset of a maximal quasi-flat Y whose asymptotic cone is the δ-truncated apartment contained in the asymptotic cone of Y. We have (compare [KL97]).

Proposition 8.5. *Let* $\mathcal{Q}$ *be a family of subsets of* $\mathcal{TT}$ *which are uniformly quasi-isometric to* $\mathbb{R}^{3g-3+m}$. *Then for every* $\delta > 0$ *there is a constant* $D = D(\delta) > 0$ *so that any set* $Q \in \mathcal{Q}$ *is a* D-*Hausdorff envelope of a* δ-*truncated maximal quasi-flat.*

Before we show the proposition we use it to derive the theorem from the introduction.

Corollary 8.6. *If* Γ *is quasi-isometric to* $\mathcal{M}(S)$ *then there is a homomorphism* $\Gamma \to \mathrm{Aut}(\mathcal{C}(S))$ *with finite kernel and finite index image.*

Proof. We observed above that a quasi-isometry $\Theta_0 : \Gamma \to \mathcal{M}(S)$ of a finitely generated group Γ into $\mathcal{M}(S)$ induces a quasi-isometry $\Theta : \Gamma \to X$ with inverse $\Lambda : X \to \Gamma$. We may assume that Θ maps the identity e in Γ to a train track τ_0 in standard form for the framing F as above. We showed above that there is a

homomorphism $\rho : \Gamma \to \mathrm{Aut}(\mathcal{C}(S))$ obtained from the fact that Γ acts as a group of homeomorphisms on the asymptotic cone $\mathcal{TT}_\omega$ preserving the basepoint (τ_0). The homeomorphism induced by an element $h \in \Gamma$ is the ultralimit of the map $\varphi(h) = \Theta \circ h \circ \Lambda$.

We have to show that the kernel of this homomorphism is finite. For this assume to the contrary that the kernel is an infinite subgroup H of Γ. Let d be the distance in Γ induced by a word norm; then there is a sequence of elements $h_i \in H$ $(i \geq 0)$ with $d(e, h_i) \to \infty$ $(i \to \infty)$. Let $\mathcal{Q}_0 \subset \mathcal{TT}$ be a finite collection of maximal quasi-flats in $\mathcal{TT}$ defined by a finite collection $\mathcal{P} = \{P_1, \ldots, P_s\}$ of pants decompositions of S. By this we mean that for each $Q \in \mathcal{Q}_0$ there is some $j \leq s$ such that Q is a union of flat strips $E(\tau_j^\ell, \lambda_j^\ell)$ where $\ell \leq 2^{3g-3+m}$, where τ_j^ℓ is a complete train track in standard form for P_j, where the complete geodesic laminations λ_j^ℓ have P_j as the union of their minimal components and where the signs defined by λ_j^ℓ for the pants curves of P_j run through all possible sign combinations. We require that for every i and every pants curve $\gamma \in P_i$ there is some $j \neq i$ and some $\tilde{\gamma} \in P_j$ which intersects γ transversely. We may assume that for every $D > 0$ the intersection of the D-neighborhoods of these quasi-flats is a compact neighborhood of τ_0 in $\mathcal{TT}$.

Let $\mathcal{Q} = \cup\varphi(h_i)\mathcal{Q}_0$; since the maps $\varphi(h_i) = \Theta \circ h_i \circ \Lambda$ are uniform quasi-isometries of $\mathcal{TT}$, the family $\mathcal{Q}$ satisfies the assumptions in Proposition 8.5. Thus by Proposition 8.5, for a small number $\delta > 0$ there is a constant $D > 0$ such that each of the sets $\varphi(h_i)Q$ $(Q \in \mathcal{Q}_0, i > 0)$ is a D-Hausdorff envelope of a δ-truncated maximal quasi-flat $F(h_i, Q)$ in $\mathcal{TT}$.

Let $A(h_i, Q)$ be the asymptotic cone of $\varphi(h_i)Q$ with basepoint the constant sequence $\varphi(h_i)\tau_0 = *$. Then $A(h_i, Q)$ is a bilipschitz-embedded euclidean space of dimension $3g - 3 + m$ passing through $*$ which contains the asymptotic cone of the δ-truncated maximal quasi-flat $F(h_i, Q)$ with basepoint $(\varphi(h_i)\tau_0) = *$. As a consequence, $A(h_i, Q)$ equals the unique apartment which contains the δ-truncated quasi-flat $F(h_i, Q)$. In particular, this apartment contains the basepoint $*$ and hence it is an apartment in the Tits cone $\mathcal{TC}$. However, by assumption the elements h_i are contained in the kernel of the homomorphism ρ and therefore $A(h_i, Q) = Q_\omega$ for all i where Q_ω is the asymptotic cone of Q with basepoint a constant sequence. As a consequence, $\varphi(h_i)$ maps each of the maximal quasi-flats $Q \in \mathcal{Q}_0$ to a set containing the δ-truncated maximal quasi-flat $F(h_i, Q == Q_\delta \subset Q$ in its D-neighborhood for a universal constant $D > 0$.

After possibly increasing D, the D-neighborhoods of the maximal truncated quasi-flats $Q_\delta \subset Q \in \mathcal{Q}_0$ intersect, and this intersection is contained in a uniformly bounded neighborhood of e. Therefore for each i the $2D$-neighborhoods of the sets $\varphi(h_i)(Q)$ $(Q \in \mathcal{Q}_0)$ contain intersection points in a uniformly bounded neighborhood of e which is independent of i. On the other hand, the maps $\varphi(h_i)$ are L-quasi-isometries for a universal constant $L > 0$ and hence the intersection of the $2D$-neighborhoods of the images $\varphi(h_i)(Q)$ $(Q \in \mathcal{Q}_0)$ is contained in a uniformly bounded neighborhood of the image under $\varphi(h_i)$ of the intersections of the $2LD$-neighborhoods of the sets $Q \in \mathcal{Q}_0$. By the choice of our family $\mathcal{Q}_0$, this implies that the distance between $\varphi(h_i)(\tau_0)$ and τ_0 is bounded from above by a universal constant not depending on i. Now the map $\Theta : \Gamma \to X$ is a quasi-isometry with

inverse Λ and therefore we conclude that the distance between $h_i = h_i(\Lambda\tau_0)$ and $e = \Lambda(\tau_0)$ in Γ is uniformly bounded from above as well. This is a contradiction and shows that the kernel H of our homomorphism ρ is a finite subgroup of Γ.

Our above argument shows that for every $h \in \Gamma$, the distance between $\varphi(h)(\tau_0)$ and $\Theta(h)$ is uniformly bounded. From this we deduce that the image of the homomorphism ρ is of finite index in $\mathcal{M}(S)$. Namely, if this is not the case then there is a sequence $\{g_i\} \subset \mathcal{M}(S)$ with $d(g_i, \rho(\Gamma)) \to \infty$. Since the quasi-isometry $\Theta_0 : \Gamma \to \mathcal{M}(S)$ is coarsely surjective there is a sequence $\{h_i\} \subset \Gamma$ with $d(\Theta_0(h_i), g_i) \leq c$ for a fixed constant $c > 0$. This means that $d(\Theta_0(h_i), \varphi(h_i)(\tau_0)) \to \infty$ which is impossible by our above argument. This completes the proof of our corollary and hence the proof of the theorem from the introduction. $\square$

We are left with the proof of Proposition 8.5. For this we follow Section 7 of [KL97]. Namely, let $\mathcal{Q}$ be a family of uniformly quasi-isometrically embedded euclidean spaces $\mathbb{R}^{3g-3+m}$ in X as in Proposition 8.5. Consider first a single set Q from the family $\mathcal{Q}$ and choose a basepoint $q \in Q$. By assumption, the ultralimit $\omega - \lim(\frac{1}{n}Q, q)$ is an apartment $A(Q)$ in the asymptotic cone $\omega - \lim(\frac{1}{n}X, q)$ which contains the basepoint $* = (q)$. Such an apartment consists of a finite union of chambers which meet at the chamber walls.

Let $\overline{*x_\omega}$ be a line segment contained in the δ-truncated apartment $A_\delta \subset A(Q)$ and issuing from the basepoint $*$. Note that a chamber has naturally the structure of a standard cone, so this is well defined. Note also that by the definition of a chamber, the line segment $\overline{*x_\omega}$ coincides with the combing line connecting $*$ to x_ω. There is a sequence $(x_n) \subset Q$ of points such that $\omega - \lim(x_n) = x_\omega$. Since $x_n \in X$ for all n, for every n there is a train track in standard form for F which is splittable to x_n. For ω-almost all n this train track coincides with a fixed track τ_0, so we may assume that τ_0 is splittable to x_n for all n. Then for each n the flat strip $E(\tau_0, x_n)$ is defined. We view $E(\tau_0, x_n)$ as a subset of its Cat(0)-extension $C(\tau_0, x_n)$.

Let $d_n \leq cn$ for all n and a fixed number $c > 0$ and assume that $d_n \to \infty$ $(n \to \infty)$. Then the ultralimit $\omega - \lim(\frac{1}{d_n}Q, q)$ is an apartment in $\mathcal{T}\mathcal{T}_\omega$ which necessarily coincides with $A(Q)$. If $\omega - \lim d_n/n \to 0$ then the ultralimit of the geodesic arcs in $C(\tau_0, x_n)$ connecting τ_0 to x_n is a geodesic ray in the asymptotic cone $C = \omega - \lim_{n\to\infty}\frac{1}{d_n}C(\tau_0, x_n)$ which defines a point ζ in the Tits boundary of $\mathcal{T}\mathcal{T}_\omega$. This point coincides with the point in the boundary of the apartment $A(Q)$ defined by the unique geodesic extension of the line segment $\overline{*x_\omega}$.

Sublemma: There is a number $r > 0$ so that for ω-all n the sets $E(\tau_0, x_n)$ are contained in the tubular r-neighborhood of Q.

Choose a point $z_n \in E(\tau_0, x_n)$ at maximal distance d_n from Q. Note that there is a universal constant $c > 0$ such that $d_n \leq cn$. We argue by contradiction and we assume that $\omega - \lim d_n = \infty$. Since $E(\tau_0, x_n)$ is connected, via possibly changing z_n we may assume that $\omega - \lim d_n/n = 0$. The asymptotic cone $\omega - \lim(\frac{1}{d_n}\mathcal{T}\mathcal{T}, \tau_0)$ contains the cell $E = \omega - \lim_n \frac{1}{d_n}E(\tau_0, x_n)$ and the apartment $\omega - \lim_n \frac{1}{d_n}Q = A(Q)$. The point $z_\omega = (z_n)$ is contained in E but not in $A(Q)$ and therefore E is not contained in $A(Q)$. We may assume that there is a point $x'_\omega \in \partial_\infty A(Q)$

in the asymptotic boundary of $A(Q)$ obtained from an ω-limit of combing lines in Q connecting $*$ to x_n. However, this limit of lines is a combing segment in $\mathcal{TT}_\omega$ connecting $*$ to x_ω and hence is necessarily contained in E as well. As a consequence, the flat strips $E, A(Q)$ intersect in the ray. Since the ray is regular, it is contained in the interior of a unique regular chamber and hence the chamber E is contained in $A(Q)$, a contradiction which shows the sublemma.

Now since the point x_ω in the δ-truncated apartment A_δ was arbitrary we conclude as in [KL97] that there is a number $R > 0$ such that each $Q \in \mathcal{Q}$ contains the truncated maximal quasi-flat defining the truncated apartment A_δ in its R-neighborhood. This completes the proof of the proposition. $\qquad\square$

References

[B95]   W. Ballmann, *Lectures on spaces of nonpositive curvature*, DMV Seminar, Band 25, Birkhäuser, Basel 1995.

[BM05]   J. Behrstock, Y. Minsky, *Dimension and rank for mapping class groups*, arXiv:math.GT/0512352.

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

[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.

[G87]   M. Gromov, *Hyperbolic groups*, Essays in group theory (S.M. Gersten, editor), MSRI Publications 8, Springer 1987, 75–263.

[G93]   M. Gromov, *Asymptotic invariants of infinite groups*, in "Geometric group theory", London Math. Soc. Lecture Notes 182 (1993).

[H06a]   U. Hamenstädt, *Geometry of the mapping class groups I: Boundary amenability*, revised version, March 2006, arXiv:math.GR/0510116.

[H06b]   U. Hamenstädt, *Geometry of the mapping class groups II: (Quasi)-geodesics*, revised version, May 2006, arXiv:math.GT/0511349.

[Har86]   J. Harer, *The virtual cohomological dimension of the mapping class group of an oriented surface*, Invent. Math. 84 (1986), 157-176.

[Hat91]   A. Hatcher, *On triangulations of surfaces*, Topology Appl. 40 (1991), 189-194.

[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.

[K06]   Y. Kida, *Measure equivalence rigidity for the mapping class group*, preprint 2006.

[K99]   B. Kleiner, *The local structure of length spaces with curvature bounded above*, Math. Zeit. 231 (1999), 409–456.

[KL97]   B. Kleiner, B. Leeb, *Rigidity of quasi-isometries for symmetric spaces and Euclidean buildings*, Inst. Hautes Etudes Sci. Publ. Math. No. 86 (1997), 115–197.

[L00]   F. Luo, *Automorphisms of the complex of curves*, Topology 39 (2000), 283–298.

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

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

[M03a]   L. Mosher, *Train track expansions of measured foliations*, unpublished manuscript.

[M03b]   L. Mosher, *Homology and dynamics in quasi-isometric rigidity of once-punctured mapping class groups*, preprint 2003.

[O96]   J. P. Otal, *Le Théorème d'hyperbolisation pour les variétés fibrées de dimension 3*, Astérisque 235, Soc. Math. Fr. 1996.

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

[T79]   W. Thurston, *Three-dimensional geometry and topology*, unpublished manuscript.

MATHEMATISCHES INSTITUT DER UNIVERSITÄT BONN
BERINGSTRASSE 1, D-53115 BONN, GERMANY