

On the asymptotic dimension of the curve complex

MLADEN BESTVINA
KEN BROMBERG

We give a bound, linear in the complexity of the surface, to the asymptotic dimension of the curve complex as well as the capacity dimension of the ending lamination space.

20F65

1. Introduction	2227
2. Good cell structures	2231
3. Train tracks	2234
4. Cell structures via splittings	2261
5. Capacity dimension of \mathcal{EL}	2265
Appendix. Train tracks and full dimension paths	2271
References	2274

1 Introduction

Let Σ be a closed orientable surface, possibly with punctures. The curve complex $\mathcal{C}(\Sigma)$ of Σ has played a fundamental role in recent work on the geometry of mapping class groups. Its hyperbolicity was established by Masur and Minsky [21], who also introduced many tools used to study its geometry. In [4] Bell and Fujiwara used the notion of tight geodesics of [21] and a finiteness theorem of Bowditch [7] to prove that $\mathcal{C}(\Sigma)$ has finite asymptotic dimension. This fact was then used by Bestvina, Bromberg and Fujiwara [5] to show that mapping class groups have finite asymptotic dimension.

Recall that a metric space X has asymptotic dimension $\leq n$ provided that for every $R > 0$ there exists a cover of X by uniformly bounded sets such that every metric R -ball in X intersects at most $n + 1$ elements of the cover.

Bowditch's finiteness theorem was nonconstructive and, as a result, Bell and Fujiwara were not able to derive any explicit upper bounds on the asymptotic dimension of $\mathcal{C}(\Sigma)$. More recently, Richard Webb [27] gave a constructive proof of Bowditch's theorem

and gave an explicit upper bound, exponential in the complexity of the surface, on the asymptotic dimension of $\mathcal{C}(\Sigma)$.

Asymptotic dimension of any visual δ -hyperbolic space X is closely related to the topology of its Gromov boundary ∂X . Buyalo [8] introduced the notion of the *capacity dimension* of a metric space and showed that $\text{asdim } X \leq \text{capdim } \partial X + 1$, where ∂X is equipped with a visual metric. (In the context of this paper, capacity dimension is the same as the Assouad–Nagata dimension.) Subsequently, Buyalo and Lebedeva [10] showed that when X is a hyperbolic group, equality holds above, and, moreover, $\text{capdim } \partial X = \dim \partial X$.

Klarreich [18] identified the boundary of the curve complex with the space \mathcal{EL} of ending laminations, which is a subquotient of the space \mathcal{PML} of projective measured laminations.

In his work on the topology of the ending lamination space, Gabai [11] produced upper bounds on the covering dimension of \mathcal{EL} : $\dim \mathcal{EL} \leq 4g + p - 4$ if Σ has genus g and $p > 0$ punctures, and $\dim \mathcal{EL} \leq 4g - 5$ if Σ is closed of genus g . We also note that the case of the 5-times punctured sphere was worked out earlier by Hensel and Przytycki [16].

Main Theorem $\text{capdim } \mathcal{EL} \leq 4g + p - 4$ if $p > 0$ and $\text{capdim } \mathcal{EL} \leq 4g - 5$ if $p = 0$.

Corollary 1.1 $\text{asdim } \mathcal{C}(\Sigma) \leq 4g + p - 3$ if $p > 0$ and $\text{asdim } \mathcal{C}(\Sigma) \leq 4g - 4$ if $p = 0$.

We note that these numbers are very close to the virtual cohomological dimension $\text{vcd } MCG(\Sigma)$ of the mapping class group, established by Harer [15]: if $p = 0$ then $\text{vcd} = 4g - 5$, if $p > 0$ and $g > 0$ then $\text{vcd} = 4g + p - 4$, and if $g = 0$ and $p \geq 3$ then $\text{vcd} = p - 3$.

Behrstock, Hagen and Sisto [3] used the Main Theorem to establish a quadratic bound on the asymptotic dimension of mapping class groups. It is an intriguing question whether asymptotic dimension for these groups is strictly bigger than the virtual cohomological dimension. There are groups — see eg Sapir [25] — that have finite cohomological but infinite asymptotic dimension. However, the authors are not aware of examples where both are finite but not equal.

Our method is to directly construct required covers of \mathcal{EL} via train track neighborhoods in \mathcal{PML} . Exactly such a strategy was employed by Gabai in proving his upper bounds on covering dimension but we will need to do extra work to gain more metric control

of the covers. Roughly speaking, train tracks give a cell structure on $\mathcal{PM}\mathcal{L}$ and a cell structure has a natural dual “handle decomposition” which gives an open cover of the space of multiplicity bounded by the dimension of the cell structure. By making the cell structure finer and showing that the multiplicity of the cover does not increase in \mathcal{EL} , Gabai obtains his upper bound. Note that cells of small dimension will not contain ending laminations, which is why in both Gabai’s work and ours the dimension bound is smaller than the dimension of $\mathcal{PM}\mathcal{L}$.

To bound the capacity dimension one needs to find, for any sufficiently small $\epsilon > 0$, covers that have bounded multiplicity and where all elements have diameter bounded above by ϵ while the Lebesgue number is bounded below by a fixed fraction of ϵ . This last property will not be satisfied by the family of covers constructed by Gabai.

The main motivation for this work is an attempt to find an alternative proof of the finiteness of asymptotic dimension of the curve complex, one that would generalize to the hyperbolic $\text{Out}(F_n)$ -complexes and provide an approach to proving $\text{asdim } \text{Out}(F_n) < \infty$. The notion of tight geodesics, used in the Bell–Fujiwara argument, does not seem to carry over to the $\text{Out}(F_n)$ -complexes, and we hope that the ideas in this paper will provide a new blueprint for attacking this question.

For readers familiar with train tracks we give a brief sketch of the construction of the cover which will highlight the difficulties in our approach. The set of laminations carried by a train track σ is naturally parametrized by a polyhedron $P(\sigma)$ in \mathbb{R}^n . (In what follows we will blur the distinction between a measured lamination and a projective measured lamination.) Note that σ carries both ending laminations and simple closed curves. We denote the former as $P_\infty(\sigma)$ and the latter as $S(\sigma)$. A basepoint $*$ in $\mathcal{C}(\Sigma)$ determines a visual metric ρ on \mathcal{EL} . To estimate the visual diameter of $P_\infty(\sigma)$ we take the curve $a \in S(\sigma)$ that is closest to $*$ in $\mathcal{C}(\Sigma)$ and then the diameter of $P_\infty(\sigma)$ is coarsely $A^{-d(a,*)}$ for some fixed constant A .

To construct our cover we will repeatedly *split* train tracks along large branches. The process of splitting σ gives two train tracks σ_+ and σ_- such that $P(\sigma_+) \cup P(\sigma_-) = P(\sigma)$ and $P(\sigma_+) \cap P(\sigma_-) = P(\tau)$, where $\tau = \sigma_+ \cap \sigma_-$ is a train track with $P(\tau)$ a codimension one face of both $P(\sigma_+)$ and $P(\sigma_-)$. To start the construction we take a cell structure on $\mathcal{PM}\mathcal{L}$ determined by a finite collection of train tracks. If the visual diameter of any of the top-dimensional cells is larger than a fixed $\epsilon > 0$ then we split. We continue this process and stop splitting a top-dimensional cell only when its diameter is $\leq \epsilon$.

At any finite stage of this construction we will obtain a cell structure on all of $\mathcal{PM}\mathcal{L}$. In particular every simple closed curve will be carried on some train track. For example, one of the cells must contain the basepoint $*$ and therefore will have large visual diameter. It immediately follows that we will need to split infinitely many times to get a collection of cells that have small visual size.

At the end of the construction we will have a countable collection of train tracks $\sigma_1, \sigma_2, \dots$, each determining a top-dimensional cell. The collection of these cells is locally finite and covers all filling laminations. To complete the proof we will need to establish the following facts:

- **Lemma 5.6** All cells $P_\infty(\sigma_i)$ have visual diameter bounded above by ϵ and bounded below by a fixed fraction of ϵ .
- **Proposition 3.21** The cells of dimension less than $\dim \mathcal{PM}\mathcal{L}$ obtained by intersecting $P(\sigma_i)$ also have the form $P(\sigma)$ and if $P_\infty(\sigma)$ is nonempty its visual diameter is also bounded below by a fraction of ϵ .
- **Proposition 4.4** If $a \in S(\sigma_i)$ and $b \in S(\sigma_j)$ are curves that are close in $\mathcal{C}(\Sigma)$ then both a and b are close to either
 - (i) a curve in $S(\tau) = S(\sigma_i) \cap S(\sigma_j)$ where τ is a subtrack of both σ_i and σ_j ,
or
 - (ii) the basepoint $*$ (when compared to $\max\{d(*, S(\sigma_i)), d(*, S(\sigma_j))\}$).

The key to proving the first bullet is the work of Masur and Minsky on splitting sequences (see Theorem 3.10). The second bullet follows from an adaptation of the work of Hamenstädt [14, Lemma 5.4] (see Propositions 3.19 and 3.20). The third bullet is the key technical advance of the paper and is proved using a version of Sela's shortening argument. (See Lemma 3.23.)

Plan of the paper In Section 2 we consider a subdivision process on polyhedral cell structures abstractly. In Section 3 we review train track theory, and prove our main technical result, Lemma 3.23. In Section 4 we apply this analysis and show that the visual size of the cover of $\mathcal{FPM}\mathcal{L}$ we produce is controlled. In Section 5 we finish the argument by producing the required “handle decomposition” from our cover and checking that it satisfies the definition of capacity dimension. Finally, in the appendix we prove a technical result (Corollary A.6) about train tracks that is presumably known to the experts. It was a surprise to us that there are nonorientable train tracks that carry only orientable laminations, and large birecurrent train tracks that do not carry filling laminations. These phenomena are discussed in the appendix.

Acknowledgements The authors thank the Warwick reading group for their careful reading of the paper and their many helpful comments: Federica Fanoni, Nicholas Gale, Francesca Iezzi, Ronja Kuhne, Beatrice Pozzetti, Saul Schleimer and Katie Vokes. We also thank the referee for very useful comments.

Both authors gratefully acknowledge the support by the National Science Foundation under the grant numbers DMS-1308178 and DMS-1207873 respectively.

2 Good cell structures

In this section we consider abstract cell structures obtained by successively subdividing cells in an initial cell structure.

2.1 Polytopes

A *polytope* in a finite-dimensional vector space $V \cong \mathbb{R}^n$ is a finite intersection of closed half-spaces.¹ The *dimension* of a polytope $U \subset V$ is the dimension of its affine span. A *face* of U is the intersection $U \cap H$ for a hyperplane $H \subset V$ such that U is contained in one of the two closed half-spaces of H . The *relative interior* of a face is its interior as a subspace of H . Faces of a polytope are also polytopes, a polytope has finitely many faces, and a face of a face is a face. The union of proper faces of a polytope is its boundary, and the complement of the boundary is the (relative) interior. See [13] or [28]. Our main example of a polytope is the set (a cone) $V(\sigma)$ of measured laminations carried by a train track σ on a surface Σ .

2.2 Cell structures

Definition 2.1 Let $U \subset V$ be a polytope. A finite collection \mathcal{C} of subsets of U which are also polytopes of various dimensions, called *cells*, is a *cell structure on U* if

- (C1) $\bigcup_{C \in \mathcal{C}} C = U$,
- (C2) when two cells intersect, their intersection is a union of cells,
- (C3) distinct cells have disjoint relative interiors,
- (C4) every face of every cell in \mathcal{C} is a union of cells.

Remark 2.2 We are really thinking about the filtration (into skeleta) $U^0 \subset U^1 \subset \dots \subset U^n = U$, so that the components of $U^i - U^{i-1}$ are open i -dimensional convex polytopes whose faces are subcomplexes.

¹Some authors require polytopes to be compact. Our polytopes will be cones on compact spaces.




Figure 1: A hexagon subdivided 3 times results in a good cell structure with twelve 0-cells, fifteen 1-cells and four 2-cells.

Definition 2.3 A cell structure \mathcal{C} on an n -dimensional polytope U is *good* if

(C5) for every $i < n$, every i -dimensional cell $C \in \mathcal{C}$ is the intersection of i -dimensional faces of $>i$ -dimensional cells in \mathcal{C} that contain C .

For example, a convex polygon with subdivided edges is not a good cell structure since (C5) fails. However, starting with a convex polygon and subdividing by line segments results in a good cell structure. See Figure 1.

2.3 Subdivision

Let \mathcal{C} be a good cell structure on a polytope U of dimension n and let W be the intersection of a codimension 0 cell $\Omega \in \mathcal{C}$ with a hyperplane (thus we are assuming $\dim \Omega = \dim U = n$). We will assume that the hyperplane intersects the relative interior of Ω . Construct a new collection \mathcal{C}' by “cutting by W ”. More precisely, replace each cell $E \in \mathcal{C}$ which is contained in Ω and with the property that $E - W$ is disconnected by the following three cells: $E \cap W$ and the closures E_1 and E_2 of the two complementary components of $E - W$. Thus W is a codimension 1 cell of \mathcal{C}' . The cells E_i have the same dimension as E , while $\dim(E \cap W) = \dim E - 1$. Figure 1 represents three consecutive subdivisions of a good cell structure consisting of a hexagon and its faces.

Lemma 2.4 *The collection \mathcal{C}' obtained from a good cell structure \mathcal{C} by subdividing is a good cell structure.*

Proof As in the notation of the definition of subdivision we subdivide a codimension 0 cell $\Omega \in \mathcal{C}$ by a codimension 1 cell W . We leave it as an exercise to prove that \mathcal{C}' is a cell structure and argue only that it is good. We show that an i -cell C' of \mathcal{C}'

with $i < n$ is the intersection of i -faces of $>i$ -cells containing C' . Let D be this intersection. Note that $D \supset C'$ so we only need to show that C' is not a proper subset of D .

Let $C \in \mathcal{C}$ be the smallest cell containing C' . Note that either $\dim C = \dim C'$ (and possibly $C = C'$) or $\dim C = \dim C' + 1$. Let $E \in \mathcal{C}$ be a cell that has a face F that contains C . Then there will be a cell $E' \subset E$ (possibly equal to E) in \mathcal{C}' with a face $F' \subset F$ and $F' \supset C'$. By letting E vary over all cells that have faces containing C we see that $D \subset C$. If $C = C'$ we are now done. If not then C is disconnected by W and in \mathcal{C} becomes three cells, C_1, C_2 and $C \cap W$, with C' being one of these three cells. Similarly, after subdivision Ω becomes three cells, Ω_1, Ω_2 and $\Omega \cap W = W$, with C_1 and C_2 contained in a face of Ω_1 and Ω_2 , respectively. In particular, if $C' = C_1$ (or $C' = C_2$) then C is contained in a face of Ω_1 (or Ω_2) but that face doesn't contain any points in $C \setminus C'$ so we must have that $C = D$. If $C' = C \cap W$ then C' is a face of W but since W doesn't contain any points in $C \setminus W$ we have that $C = D$ in this case also. □

Remark 2.5 When E has codimension 1 and U is a manifold (eg when U is a polytope), the intersection in (C5) consists of (at most) two elements. But when the codimension is > 1 the argument does not produce a uniform bound on the number of faces required.

Corollary 2.6 Suppose \mathcal{C} is a good cell structure. If a cell $E \in \mathcal{C}$ of dimension $i < n$ has m codimension 1 faces, then E can be written as the intersection of $\leq m$ i -dimensional faces of cells in \mathcal{C} of dimension $> i$.

Definition 2.7 A (finite or infinite) sequence $\mathcal{C}_0, \mathcal{C}_1, \dots$ of cell structures on U is excellent if

- (E1) \mathcal{C}_0 consists of U_p 's and their faces,
- (E2) for $i \geq 1, \mathcal{C}_i$ is obtained from \mathcal{C}_{i-1} by the subdivision process along codimension 0 cells described above, or else $\mathcal{C}_i = \mathcal{C}_{i-1}$.

By Lemma 2.4, the cell structures in an excellent sequence are good cell structures.

Remark 2.8 Easy examples in \mathbb{R}^3 show that it is *not* true in general that an i -cell is the intersection of i -faces of codimension 0 cells. eg consider the plane $x = 0$ and half-planes $z = 0, x \geq 0$ and $y = 0, x \leq 0$.

Remark 2.9 This lemma is where our cell structure differs from Gabai's. For our cell structure we only subdivide cells of positive codimension if they are induced by subdivisions of top-dimensional cells. The proposition insures that when doing this all cells are defined via train tracks (ie they are of the form $V(\theta)$ where θ is a train track; see Proposition 3.19). Gabai also needs this property but he achieves it by subdividing cells of positive codimension. We do not want to do this as the visual diameter of these cells may become arbitrarily small. See Figure 2.




Figure 2: The cell drawn in a thick line arises as the intersection of top-dimensional cells. We do not want to subdivide it further as this would make the visual size too small.

3 Train tracks

3.1 Notation and background

Fix a surface Σ of finite type. In what follows all constants will depend on the topology of Σ . We will assume the reader is familiar with the theory of train tracks. The standard reference is [24]. See also [21; 14] for introductions to the theory. A quick definition is that a train track in a surface Σ is a smooth graph with a well-defined tangent line at every point, including at the vertices, such that no complementary component is a (smooth) disk, a monogon, a bigon or a punctured disk, and so that every edge can be extended in both directions to a smoothly immersed path (these are called legal paths or train paths). All our train tracks will always be generic (ie all vertices have valence 3) and in general they will be recurrent and transversely recurrent (birecurrent). However, there will be occasions when nonrecurrent tracks will appear. A train track $\sigma \subset \Sigma$ is *large* if each complementary component is homeomorphic to a disk or a once-punctured disk. A train track is *maximal* if all complementary components are triangles or punctured monogons, with the exception of the punctured torus, where a maximal train track contains a single punctured bigon in its complement.

3.1.1 Transverse measures The edges of the train track are *branches* and the vertices are *switches*. At each switch of a generic train track $\sigma \subset \Sigma$ there are three incident half-branches. Two of these are tangent (ie determine the same unit tangent vector) and are called *small*, while the third is a *large* half-branch. A branch whose both half-branches are large is called *large*. If both half-branches are small then the branch is *small*. Otherwise the branch is *mixed*.

A *transverse measure* on a (generic) train track is an assignment of nonnegative weights to each branch that satisfy the *switch equations*. That is, at each switch the sum of the weights of the two small half-branches should be equal to the weight of the large half-branch. A transverse measure determines a unique measured lamination on Σ . These are the laminations *carried* by τ .

A train track is *recurrent* if it admits a transverse measure which is positive on every branch. All of our train tracks are going to be *transversally recurrent*; see [24] for the definition. We will not use this property directly, but most results in the literature assume it, and further there is no harm doing so as transverse recurrence persists under splits and subtrack moves. A train track is *birecurrent* if it is both recurrent and transversally recurrent.

The set of all measured laminations on Σ is denoted by \mathcal{ML} and the set of measured laminations carried by σ is denoted by $V(\sigma)$. Thus $V(\sigma)$ is the closed positive cone in the vector space of real weights on the branches of σ satisfying the (linear) switch equations; in particular, $V(\sigma)$ is a polytope. We denote by \mathcal{PML} the projective space of measured laminations and for a train track σ we let $P(\sigma) \subset \mathcal{PML}$ be the set of projective measured laminations carried by σ . Then $P(\sigma)$ can be identified with the projectivization of $V(\sigma) - \{0\}$. We will often blur the distinction between a measured lamination and its projective class.

We also denote by $\mathcal{FPM}\mathcal{L} \subset \mathcal{PML}$ the subset of those laminations that are *filling*, ie whose complementary components are disks or punctured disks. Given a measured lamination $\lambda \in \mathcal{ML}$ (or \mathcal{PML}) we let $[\lambda]$ be the underlying geodesic lamination.

We have a quotient map $\mathcal{FPM}\mathcal{L} \rightarrow \mathcal{EL}$ to the space of *ending laminations*, defined by $\lambda \mapsto [\lambda]$. Recall that Klarreich [18] showed that \mathcal{EL} is the Gromov boundary of the curve complex $\mathcal{C}(\Sigma)$. Note that in general if $\lambda_i \in \mathcal{FPM}\mathcal{L}$ is a sequence with limit λ then $[\lambda]$ may be a proper subset of the Hausdorff limit of $[\lambda_i]$.

For a train track σ let $P_\infty(\sigma) = P(\sigma) \cap \mathcal{FPM}\mathcal{L}$.

At each switch the tangent direction gives a way to compare the orientation of each branch adjacent to the switch. A train track is *orientable* if each branch can be given an orientation that is consistent at each switch.

When σ is a generic birecurrent train track we have $\frac{1}{3}|b| = \frac{1}{2}|v| = -\chi(\sigma)$, where $|b|$ and $|v|$ denote the numbers of branches and switches, respectively.

Lemma 3.1 [24, Lemma 2.1.1] *Let σ be a connected recurrent train track. Then the dimension of $V(\sigma)$ is $\frac{1}{3}|b|$ if σ is nonorientable and $\frac{1}{3}|b| + 1$ if σ is orientable.*

Sketch of the proof Suppose first that σ is nonorientable. Given a switch v , there is a train path that starts and ends at v , and the initial and terminal half-branches are the two small half-branches at v . This path assigns weights to the branches of σ that satisfy all switch equations except at v . This shows that the switch equations are linearly independent, proving the assertion.

Now suppose σ is orientable. Choose an orientation and write each switch equation as the sum of incoming branch(es) equals the sum of outgoing branch(es). Then summing all switch equations yields an identity, with each branch occurring once on both sides. Thus one switch equation is redundant, and we need to argue that the others are independent. Let v and w be two distinct switches. Choose a train path that connects v to w . This path assigns weights to all edges, and the switch equations are satisfied except at v and w . This proves the claim. \square

3.1.2 Faces of $V(\sigma)$ There is a bijection between faces of $V(\sigma)$ and recurrent subtracks of σ . (Here we allow train tracks to be disconnected and to contain components that are simple closed curves.) A subtrack of σ may not be recurrent but any track has a unique maximal recurrent subtrack.

3.1.3 Splitting Starting with a maximal, birecurrent train track σ we will describe a splitting operation on train tracks that will us to subdivide $V(\sigma)$ and produce an excellent sequence of cell structures on $V(\sigma)$. We describe this now.

If b is a large branch of σ , one can produce two new train tracks σ_1 and σ_2 by *splitting* b . See Figure 3. We say that σ_1 is obtained by the *left split* and σ_2 by the *right split*.

Every lamination that is carried by σ will be carried by either σ_1 or σ_2 . If a lamination is carried by both σ_1 and σ_2 then it will be carried by the *central split* $\tau = \sigma_1 \cap \sigma_2$, obtained from either σ_1 or σ_2 by removing the diagonally drawn branch.




Figure 3: A large branch e in the middle is split in two ways to give train tracks σ_1 and σ_2 .

We have the following facts:

- [24, Lemma 1.3.3(b)] If σ is transversely recurrent, so are σ_1 and σ_2 and τ .
- [24, Lemma 2.1.3] If σ is recurrent, then either all three of σ_1 , σ_2 and τ are recurrent, or exactly one is recurrent.

It is also easy to see that σ_1 and σ_2 are orientable if and only if σ is.

3.1.4 Subdivision Now suppose σ is a birecurrent train track and b a large branch of σ . We describe a process that subdivides $V(\sigma)$. There are several cases. Denote by σ_1 , σ_2 and τ the left, right and central splits of σ along b .

- (S1) If all three of σ_1 , σ_2 and τ are recurrent, the cell $V(\tau)$ is a codimension 1 hyperplane in $V(\sigma)$ and cuts it into $V(\sigma_1)$ and $V(\sigma_2)$. Thus $\dim V(\sigma) = \dim V(\sigma_i) = \dim V(\tau) + 1$. In this case we are subdividing $V(\sigma)$ as in Section 2.3.
- (S2) If σ_1 is recurrent but σ_2 and τ are not recurrent then $V(\sigma) = V(\sigma_1)$ while $V(\sigma_2) = V(\tau)$ will be a proper face of $V(\sigma)$ (possibly empty).
- (S3) Suppose τ is recurrent, but σ_1 and σ_2 are not. Then τ is the maximal recurrent subtrack of both σ_i and $V(\sigma) = V(\sigma_i) = V(\tau)$. Since $\dim V(\sigma) = \dim V(\tau)$, Lemma 3.1 implies that σ is nonorientable while τ is orientable. Note that if this case occurs every lamination carried by σ is orientable. It may also happen that σ is large while τ is not, so we have a situation that a large birecurrent train track does not carry any filling laminations.

3.2 Carrying maps, stationary and active sets

If σ and τ are train tracks then a map $\sigma \rightarrow \tau$ is a *carrying map* if it is locally injective on each edge and takes legal train paths to legal train paths. We also say σ is *carried* by τ and we are implicitly assuming some explicit carrying map has been chosen. We say that a carrying map $\sigma \rightarrow \tau$ is *fully carrying* if it is a homotopy equivalence, and we then write $\sigma \twoheadrightarrow \tau$. If λ is a lamination carried by τ , we write $\lambda \rightarrow \tau$ for the carrying map. If moreover this map induces a bijection between complementary components that preserves the topology and numbers of sides and punctures, we say that τ *fully carries* λ and we write $\lambda \twoheadrightarrow \tau$. Thus in this case splitting τ according to λ always produces train tracks that fully carry λ .

Our definition of a track fully carrying a lamination is stronger than what is used in [11], where it is only assumed that any realization of λ as a measured lamination will be in the relative interior of $V(\tau)$.

If σ_1 is a splitting of σ there is a unique (up to homotopy rel vertices) full carrying map $\sigma_1 \twoheadrightarrow \sigma$ that is a bijection on vertices and is a homeomorphism outside a small neighborhood of the large branch where the split occurs. If σ_1 is obtained from a finite sequence of splittings of σ , we will always assume that the carrying map $\sigma_1 \twoheadrightarrow \sigma$ is a composition of such maps.

If τ is obtained from σ by some finite combination of splits and central splits, we write $\tau \xrightarrow{s} \sigma$. If τ is obtained by a finite sequences of splits only then τ is fully carried by σ , and we write $\tau \xrightarrow{s} \sigma$.

We also use the notation $\sigma \xrightarrow{ss} \tau$ to mean that σ is obtained from τ by a sequence of splits, central splits and passing to subtracks. A single *move* is either a split, a central split or passing to a subtrack. The number of *splitting moves* in $\sigma \xrightarrow{ss} \tau$ is the number of splits and central splits in the sequence. When we write $\sigma \xrightarrow{ss} \tau$ we will be implicitly assuming that some sequence of splits and subtracks has been chosen. However, the choice of a sequence is not unique and different choices of sequences may have a different number of moves.

Given two sequences $\sigma_1 \xrightarrow{ss} \tau$ and $\sigma_2 \xrightarrow{ss} \tau$, we would like to find a new train track σ with $\sigma \xrightarrow{ss} \sigma_i$ for $i = 1, 2$ and $V(\sigma) = V(\sigma_1) \cap V(\sigma_2)$. To accomplish this we need to develop some machinery about train tracks. The main technical result we need is Proposition 3.19.

Given a sequence of $\sigma \xrightarrow{ss} \tau$ we now want to define the set of active and stationary branches. To do so we first make some general comments about sets of branches

and half-branches and their complements. Let \mathcal{S} be a collection of branches and half-branches of a train track τ such that if \mathcal{S} contains a branch then it contains both half-branches, and if it contains a half-branch then it contains both other half-branches at the same switch. Then the *complementary branch set* \mathcal{A} contains a branch b if neither b nor any of its half-branches are in \mathcal{S} and contains a half-branch h if h is not in \mathcal{S} . Note that \mathcal{A} will also have the property that if a branch is in \mathcal{A} then both half-branches will be in \mathcal{A} but will also have the stronger property that if \mathcal{A} contains both half-branches of a branch then it will contain the branch. We also note that $\mathcal{S} \cup \mathcal{A}$ may not contain all branches of τ but it will contain all half-branches. Let $|\mathcal{S}|$ be the union of branches and half-branches in \mathcal{S} . We think of half-branches as germs, so if both half-branches of a branch b are in \mathcal{S} but b is not in \mathcal{S} , then $|\mathcal{S}|$ will be missing an interval in the interior of b .

A convenient way to visualize the set \mathcal{S} is to view the train track τ as a graph. Then switches with incident half-branches in \mathcal{S} correspond to some vertices, and branches in \mathcal{S} to some edges in τ . These vertices and edges define a subgraph $\tau_{\mathcal{S}}$ of τ . The complementary set \mathcal{A} similarly corresponds to the maximal subgraph of τ disjoint from $\tau_{\mathcal{S}}$.

Given train tracks σ and τ with $\sigma \rightarrow \tau$ a branch b in σ is *stationary* if the carrying map is a homeomorphism from a neighborhood of b to its image in τ . We similarly define a half-branch to be stationary and let $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)$ be the set of stationary branches and half-branches in σ . Note that a half-branch is contained in $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)$ if and only if the carrying map is a homeomorphism on a neighborhood of the switch adjacent to the half-branch to its image. We emphasize that the stationary set depends on the choice of carrying map and two homotopic carrying maps may have different stationary sets. In particular, a choice of sequence $\sigma \xrightarrow{ss} \tau$ determines the carrying map and hence the stationary set but a different choice of sequence may determine a different stationary set.

The image of the stationary set in τ will be a collection of branches and half-branches, which we denote by $\mathcal{S}(\sigma \rightarrow \tau; \tau)$. The carrying map $\sigma \rightarrow \tau$ factors through a train track τ' if $\sigma \rightarrow \tau$ is the composition of carrying maps $\sigma \rightarrow \tau'$ and $\tau' \rightarrow \tau$ and we define $\mathcal{S}(\sigma \rightarrow \tau; \tau')$ to be the image of $\mathcal{S}(\sigma \rightarrow \tau; \sigma)$ in τ' . The main example for us is when we have a sequence $\sigma \xrightarrow{ss} \tau$ and τ' is a track in the sequence.

The carrying map will restrict to a homeomorphism from $|\mathcal{S}(\sigma \rightarrow \tau; \sigma)|$ to $|\mathcal{S}(\sigma \rightarrow \tau; \tau)|$. However, for a general carrying map, the preimage of $|\mathcal{S}(\sigma \rightarrow \tau; \tau)|$ in σ may be

larger than the carrying set. For carrying maps that come from sequences $\sigma \xrightarrow{ss} \tau$, this does not happen.

Lemma 3.2 *Let σ and τ be train tracks with $\sigma \xrightarrow{ss} \tau$. The carrying map $\sigma \xrightarrow{ss} \tau$ restricts to a homeomorphism from $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)|$ to $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)|$ and the preimage of $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)|$ in σ is $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)|$.*

Proof We induct on the number of moves in $\sigma \xrightarrow{ss} \tau$. If $\sigma \xrightarrow{ss} \tau$ is a single move then the lemma follows by direct examination. If $\sigma \xrightarrow{ss} \tau$ has m moves then we choose a train track τ' such that $\sigma \xrightarrow{ss} \tau' \xrightarrow{ss} \tau$ with $\sigma \xrightarrow{ss} \tau'$ having $m - 1$ moves and $\tau' \xrightarrow{ss} \tau$ a single move. As

$$\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau') = \mathcal{S}(\sigma \xrightarrow{ss} \tau'; \tau') \cap \mathcal{S}(\tau' \xrightarrow{ss} \tau; \tau')$$

and by the induction hypothesis, the carrying map $\sigma \xrightarrow{ss} \tau'$ restricts to a homeomorphism from $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)|$ to $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau')|$ and the carrying map $\tau' \xrightarrow{ss} \tau$ restricts to a homeomorphism from $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau')|$ to $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)|$. Therefore $\sigma \xrightarrow{ss} \tau$ restricts to a homeomorphism from $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)|$ to $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)|$. A similar argument show that the preimage of $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)|$ in σ is $|\mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)|$. \square

Given train tracks τ_1 and τ_2 in a sequence $\sigma \xrightarrow{ss} \tau$ and a collection of branches and half-branches $\mathcal{B}_i \subset \mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau_i)$ for $i = 1, 2$, we write $\mathcal{B}_1 = \mathcal{B}_2$ if the bijection from $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau_1)$ to $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau_2)$ takes \mathcal{B}_1 to \mathcal{B}_2 .

We can define the set of *active branches* $\mathcal{A}(\sigma \xrightarrow{ss} \tau; \tau')$ to be the complementary branch set of the stationary branches $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau')$, where τ' is a train track in the sequence $\sigma \xrightarrow{ss} \tau$.

Recall that in general if two half-branches of a track are in the stationary set, the full branch may not be. However there is one special case where this does hold.

Lemma 3.3 *Let σ and τ be train tracks with $\sigma \xrightarrow{ss} \tau$. If b is a branch in τ such that both of its half-branches are contained in $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)$, then $b \in \mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)$.*

Proof We first observe how the lemma can fail for $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)$. Let b' be a branch in σ with both half-branches in $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)$. Under the carrying map $\sigma \xrightarrow{ss} \tau$ the branch b' will map to a legal path that starts and ends at a switch. (Here we are using that $\sigma \xrightarrow{ss} \tau$ takes switches to switches by construction.) Then $b' \in \mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)$ if and only if the legal path is a single branch in τ .

In our case the half-branches of b are in $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)$ and as the carrying maps are good the preimage of each will be a single half-branch in σ and therefore the preimage of b will be a single branch b' in σ . Then, by the above paragraph, $b' \in \mathcal{S}(\sigma \xrightarrow{ss} \tau; \sigma)$ and its image, b , will be in $\mathcal{S}(\sigma \xrightarrow{ss} \tau; \tau)$. \square

Corollary 3.4 *Let σ_1, σ_2 and τ be train tracks with $\sigma_i \xrightarrow{ss} \tau$ for $i = 1, 2$. Then $\mathcal{A}(\sigma_1 \xrightarrow{ss} \tau; \tau) \subset \mathcal{S}(\sigma_2 \xrightarrow{ss} \tau; \tau)$ if and only if $\mathcal{A}(\sigma_2 \xrightarrow{ss} \tau; \tau) \subset \mathcal{S}(\sigma_1 \xrightarrow{ss} \tau; \tau)$.*

Proof As the set of half-branches of τ is the disjoint union of the half-branches in $\mathcal{S}(\sigma_i \xrightarrow{ss} \tau; \tau)$ and $\mathcal{A}(\sigma_i \xrightarrow{ss} \tau; \tau)$, we only need to check full branches. In particular, if $\mathcal{A}(\sigma_1 \xrightarrow{ss} \tau; \tau) \subset \mathcal{S}(\sigma_2 \xrightarrow{ss} \tau; \tau)$ and b is a full branch in $\mathcal{A}(\sigma_2 \xrightarrow{ss} \tau; \tau)$, then we need to show that b is in $\mathcal{S}(\sigma_1 \xrightarrow{ss} \tau; \tau)$. If b is not in $\mathcal{S}(\sigma_1 \xrightarrow{ss} \tau; \tau)$ then by Lemma 3.3 a half-branch h of b is not in $\mathcal{S}(\sigma_1 \xrightarrow{ss} \tau; \tau)$ and therefore $h \in \mathcal{A}(\sigma_1 \xrightarrow{ss} \tau; \tau) \subset \mathcal{S}(\sigma_2 \xrightarrow{ss} \tau; \tau)$. However, if $h \in \mathcal{S}(\sigma_2 \xrightarrow{ss} \tau; \tau)$ then $b \notin \mathcal{A}(\sigma_2 \xrightarrow{ss} \tau; \tau)$, contradicting our assumption. \square

We say that $\sigma_1 \xrightarrow{ss} \tau$ and $\sigma_2 \xrightarrow{ss} \tau$ are *disjoint* if either of the conditions of Corollary 3.4 hold.

Lemma 3.5 *Let $\tau_n \xrightarrow{ss} \tau_{n-1} \xrightarrow{ss} \dots \xrightarrow{ss} \tau_0$ be a sequence of moves and σ_0 another train track such that $\sigma_0 \xrightarrow{ss} \tau_0$ with $\tau_n \xrightarrow{ss} \tau_0$ and $\sigma_0 \xrightarrow{ss} \tau_0$ disjoint. Then there exists a sequence $\sigma_n \xrightarrow{ss} \sigma_{n-1} \xrightarrow{ss} \dots \xrightarrow{ss} \sigma_0$ such that*

- (a) $\sigma_n \xrightarrow{ss} \sigma_0$ has the same number of moves and splitting moves as $\tau_n \xrightarrow{ss} \tau_0$;
- (b) $\tau_i \xrightarrow{ss} \sigma_i$ where the sequence has the same number of moves and splitting moves as $\sigma_0 \xrightarrow{ss} \tau_0$;
- (c) $\mathcal{A}(\sigma_i \xrightarrow{ss} \tau_i; \tau_i) \subset \mathcal{S}(\tau_n \xrightarrow{ss} \tau_0; \tau_i)$ and $\mathcal{A}(\tau_i \xrightarrow{ss} \tau_0; \tau_i) \subset \mathcal{S}(\sigma_i \xrightarrow{ss} \tau_i; \tau_i)$;
- (d) $\sigma_i \xrightarrow{ss} \tau_i$ and $\tau_n \xrightarrow{ss} \tau_i$ are disjoint;
- (e) $V(\sigma_i) = V(\tau_i) \cap V(\sigma_0)$.

Proof Assume that $\sigma_i \xrightarrow{ss} \tau_i$ has been constructed. We will first construct a track σ_{i+1} with $\sigma_{i+1} \rightarrow \tau_{i+1}$ and then show that it can be realized as a sequence of moves. The move $\tau_{i+1} \xrightarrow{ss} \tau_i$ is either a splitting move or a subtrack move on a branch b of τ_i . As $\sigma_i \rightarrow \tau_i$ and $\tau_n \xrightarrow{ss} \tau_i$ are disjoint we have $b \in \mathcal{S}(\sigma_i \rightarrow \tau_i; \tau_i)$, so the preimage of b in σ_i is a branch b' of the same type and we can perform the same move on b' to form σ_{i+1} . The carrying map $\sigma_i \rightarrow \tau_i$ gives a map from $|\mathcal{S}(\sigma_{i+1} \xrightarrow{ss} \sigma_i; \sigma_{i+1})|$ to

$|\mathcal{S}(\tau_{i+1} \xrightarrow{ss} \tau_i; \tau_{i+1})|$. If the move is a right or left split then the complement of the stationary set (for both $\sigma_{i+1} \xrightarrow{ss} \sigma_i$ and $\tau_{i+1} \xrightarrow{ss} \tau_i$) is the neighborhood of a small branch. If it is a central split or a subtrack move then the complement will be the interior of two branches. In all cases the map from $|\mathcal{S}(\sigma_{i+1} \xrightarrow{ss} \sigma_i; \sigma_{i+1})|$ to $|\mathcal{S}(\tau_{i+1} \xrightarrow{ss} \tau_i; \tau_{i+1})|$ extends to a carrying map $\sigma_{i+1} \xrightarrow{ss} \tau_{i+1}$ that is a homeomorphism in the complement of the two stationary sets. In particular, the active set $\mathcal{A}(\sigma_{i+1} \xrightarrow{ss} \tau_{i+1}; \tau_{i+1})$ is contained in the stationary set $\mathcal{S}(\tau_{i+1} \xrightarrow{ss} \tau_i)$ and the carrying map $\tau_{i+1} \xrightarrow{ss} \tau_i$ takes it homeomorphically to $\mathcal{A}(\sigma_i \xrightarrow{ss} \tau_i; \tau_i)$. Therefore, as $\mathcal{A}(\sigma_i \xrightarrow{ss} \tau_i; \tau_i) \subset \mathcal{S}(\tau_n \xrightarrow{ss} \tau_0; \tau_i)$, we have $\mathcal{A}(\sigma_{i+1} \xrightarrow{ss} \tau_{i+1}; \tau_{i+1}) \subset \mathcal{S}(\tau_n \xrightarrow{ss} \tau_0; \tau_{i+1})$. The second inclusion in (c) follows from the first exactly as in Corollary 3.4. The first inclusion in (c) implies that $\mathcal{A}(\sigma_{i+1} \xrightarrow{ss} \tau_{i+1}; \tau_{i+1}) \subset \mathcal{S}(\tau_n \xrightarrow{ss} \tau_{i+1}; \tau_{i+1})$ and therefore (d) holds.

To see that $\sigma_{i+1} \rightarrow \tau_{i+1}$ can be realized as a sequence we observe that if $\sigma_0 \xrightarrow{ss} \tau_0$ is a single move then so is $\sigma_{i+1} \rightarrow \tau_{i+1}$. In general we induct on the number of moves in $\sigma_0 \xrightarrow{ss} \tau_0$.

For (e) we observe that $V(\sigma_{i+1}) \subset V(\tau_{i+1}) \cap V(\sigma_i)$. Let λ be a lamination in $V(\tau_{i+1}) \cap V(\sigma_i) \subset V(\tau_i)$. Then λ is realized by transverse measures m_i, m_{i+1} and m'_i on τ_i, τ_{i+1} and σ_i . Then m_i and m_{i+1} will agree on the stationary set of $\tau_{i+1} \xrightarrow{ss} \tau_i$ and m'_i and m_i will agree on the stationary set of $\sigma_i \xrightarrow{ss} \tau_i$. By examining the various cases we see that there is a transverse measure m'_{i+1} on τ_{i+1} such that m'_{i+1} agrees with m'_i on the stationary set of $\tau_{i+1} \xrightarrow{ss} \tau_i$ and m'_{i+1} agrees with m_{i+1} on the stationary set of $\sigma_{i+1} \xrightarrow{ss} \tau_i$. For any single move transverse measures on each of the tracks that agree on the stationary set will determine the same lamination. Therefore, m'_{i+1} realizes λ , so $V(\sigma_{i+1}) = V(\sigma_i) \cap V(\tau_{i+1})$. As $V(\tau_{i+1}) \subset V(\tau_i)$ and $V(\sigma_i) = V(\tau_i) \cap V(\sigma_0)$, this implies that

$$V(\sigma_{i+1}) = V(\tau_{i+1}) \cap V(\sigma_i) = V(\tau_{i+1}) \cap V(\tau_i) \cap V(\sigma_0) = V(\tau_{i+1}) \cap V(\sigma_0). \quad \square$$

Lemma 3.6 *Let $b \in \mathcal{A}(\sigma \xrightarrow{s} \tau; \tau)$ be a large branch in τ . Then there exists a train track σ' with $\sigma' \xrightarrow{s} \tau$ a single move on b and $\sigma \xrightarrow{s} \sigma'$ with the sequence having at most the same number of moves as $\sigma \xrightarrow{s} \tau$.*

Proof Assume that the sequence $\sigma \xrightarrow{s} \tau$ has been chosen so that the move on b occurs as early as possible. More concretely, given any sequence $\sigma \xrightarrow{s} \tau$ there exist tracks τ_1 and τ_2 in the sequence such that $\tau_1 \xrightarrow{s} \tau_2$ is a single move, and $b \in \mathcal{S}(\tau_1 \xrightarrow{s} \tau; \tau)$ but $b \notin \mathcal{S}(\tau_2 \xrightarrow{s} \tau; \tau)$. We assume that the sequence has been chosen minimizing the number of moves in $\tau_2 \xrightarrow{s} \tau$.

Let $\tau_2 \xrightarrow{s} \tau_3$ be the next move in the sequence. This will be a move on a large branch b' in τ_3 . As $b \in \mathcal{S}(\tau_2 \xrightarrow{s} \tau; \tau_2)$ we also have $b \in \mathcal{S}(\tau_2 \xrightarrow{s} \tau_3; \tau_2)$. In particular, b is also a large branch in τ_3 and it is distinct from b' . We then let $\tau'_2 \xrightarrow{s} \tau_3$ be the same move on b as $\tau_1 \xrightarrow{s} \tau_2$ and note that $b' \in \mathcal{S}(\tau'_2 \xrightarrow{s} \tau_3; \tau_3)$, so b' is a large branch in τ'_2 and we can choose $\tau'_1 \xrightarrow{s} \tau'_2$ to be the same move as $\tau_2 \xrightarrow{s} \tau_3$. By direct examination we see that $\tau'_1 = \tau_1$, so we have made a new sequence $\sigma \xrightarrow{s} \tau$, where the move on b occurs earlier, a contradiction. \square

Lemma 3.7 *Let σ and τ be train tracks with τ recurrent. If $\sigma \xrightarrow{ss} \tau$ and $V(\sigma)$ intersects the relative interior of $V(\tau)$ then $\sigma \xrightarrow{s} \tau$.*

Proof If $V(\sigma)$ intersects the relative interior of $V(\tau)$ then the carrying map $\sigma \xrightarrow{ss} \tau$ must be surjective. Let $\tau' \xrightarrow{ss} \tau$ be the first move. This map must be surjective and for a single move this can only happen for a split or central split. If $\sigma \rightarrow \tau'$ is not surjective, then $\tau' \xrightarrow{s} \tau$ is a split and the image of σ in τ' includes all edges except the diagonal. Thus we can replace the first split with the central split and proceed by induction. \square

Lemma 3.8 *Let τ be a train track and b a branch. Then there exists a nonempty collection of large branches \mathcal{B} such that if σ is a train track and $\sigma \xrightarrow{s} \tau$ with $b \in \mathcal{A}(\sigma \xrightarrow{s} \tau; \tau)$, then every branch in \mathcal{B} is in $\mathcal{A}(\sigma \xrightarrow{s} \tau; \tau)$.*

Proof If b is large then $\mathcal{B} = \{b\}$. If not, consider a small half-branch b_1 of b . There is a unique large half-branch b'_2 adjacent to b_1 and let b_2 be the other half-branch of the branch B_2 that contains b'_2 . If B_2 (ie b_2) is large then we note that B_2 must be split before b becomes active. If b_2 is small, we continue inductively and construct half-branches b_3, b_4, \dots, b_k ending in a large half-branch b_k (see [24, page 127; 14, page 574]) and note that the associated large branch B_k must split before b does. The inductive process must terminate with a large half-branch for otherwise some half-branch will repeat and, by the same argument, none of the branches listed will ever be active. Thus \mathcal{B} can be taken to have cardinality 1 or 2. \square

3.3 Splitting sequences and excellent cell structures

Given a maximal birecurrent train track σ we describe a construction of an excellent sequence of cell structures \mathcal{C}_j for $j = 0, 1, \dots$ on the polytope $V(\sigma)$.

We start by defining \mathcal{C}_0 to consist of $V(\sigma)$ and its faces. Inductively, each top-dimensional cell E of \mathcal{C}_j will correspond to a birecurrent track θ_E such that $E = V(\theta_E)$.

To define \mathcal{C}_{j+1} , choose a top-dimensional cell E of \mathcal{C}_j and a large branch b of θ_E . Let θ_1 , θ_2 and τ be the left, right and central splits of θ_E along b . We now consider the three cases (S1)–(S3) as in Section 3.1.4.

If all three of θ_1 , θ_2 and τ are recurrent, we split $E = V(\theta_E)$ along the hyperplane $V(\tau)$, yielding new top-dimensional cells $V(\theta_1)$ and $V(\theta_2)$, and we subdivide all cells that are cut by this hyperplane as described in Section 2.3.

If θ_1 is recurrent but θ_2 and τ are not, then $V(\theta_1) = V(\theta_E)$ and we define $\mathcal{C}_{j+1} = \mathcal{C}_j$ and $\theta_E = \theta_1$. We proceed similarly if θ_2 is recurrent but θ_1 and τ are not.

The last case is when τ is recurrent, but θ_1 and θ_2 are not. However, this would imply that τ is not maximal and therefore not a top-dimensional cell by Lemma 3.1.

A sequence \mathcal{C}_j obtained in this way is said to be obtained by a *splitting process* from σ . Note that if $E = V(\theta_E)$ is a top-dimensional cell in \mathcal{C}_i and if $E' = V(\theta_{E'})$ is a top-dimensional cell in \mathcal{C}_j such that $j > i$ and $E \subsetneq E'$, then $\theta_{E'} \xrightarrow{s} \theta_E$ and the sequence of splits and central splits contains at most one central split.

We have two goals for the next few sections:

- We will show that every *every* cell of \mathcal{C}_j has the form $V(\theta)$ for a suitable birecurrent train track θ . Here the key is to show (under suitable restrictions) that if σ_1 and σ_2 are train tracks then $V(\sigma_1) \cap V(\sigma_2) = V(\sigma)$ for a train track σ . One difficulty is that the dimension of the intersection may be less than the dimension of the original cells.
- We also need to control the “size” of the individual cells. We need to both show that for any ending lamination $\lambda \in V(\sigma)$ we can subdivide so that the cell containing λ is small but also that the size of any proper face of cell is comparable to the size of the cell.

The main result we need is Proposition 3.21.

3.3.1 The curve graph and vertex cycles We denote by $\mathcal{C}(\Sigma)$ the curve graph of Σ . Its vertices are isotopy classes of essential simple closed curves on Σ , and two vertices are connected by an edge if the corresponding classes have disjoint representatives. When Σ has low complexity, $\mathcal{C}(\Sigma)$ can be empty or discrete, and in the sequel we will always assume that $\mathcal{C}(\Sigma)$ contains edges. In that case, $\mathcal{C}(\Sigma)$ is connected and the edge–path metric is δ –hyperbolic [21].

The train track σ carries a curve that crosses each branch at most twice, and if it crosses a branch twice, it does so with opposite orientations. Such curves are the *vertex cycles* of σ . To a train track $\sigma \subset \Sigma$ we associate the sets $B(\sigma) \subset \mathcal{C}(\Sigma)$ consisting of all vertex cycles for σ , and the set $S(\sigma) \subset \mathcal{C}(\Sigma)$ of all curves carried by σ . We think of $B(\sigma) \subset S(\sigma)$ as a thick basepoint of $S(\sigma)$. It is a nonempty uniformly bounded subset of $S(\sigma)$.

3.3.2 Splitting sequences and the geometry of the curve graph We begin with an elementary lemma relating a single splitting to the geometry of the curve graph.

Lemma 3.9 *Suppose $\sigma \xrightarrow{ss} \tau$ is a single move. Then*

$$d(B(\sigma), B(\tau))$$

is uniformly bounded.

Proof Vertex cycles in subtracks are also vertex cycles in the track. In the case of splittings, the intersection number between a vertex cycle of σ and a vertex cycle of σ_i is uniformly bounded, and so is the distance in $\mathcal{C}(\Sigma)$. \square

Given a sequence $\sigma \xrightarrow{ss} \tau$, the previous lemma implies that the corresponding sequence of vertex cycles is a coarse path in $\mathcal{C}(\Sigma)$. It is a theorem of Masur and Minsky that given a sequence of carrying maps of birecurrent train tracks whose vertex cycles are a coarse path in $\mathcal{C}(\Sigma)$, the sequence of vertex cycles are an unparametrized quasigeodesic. In our case we know that if τ is transversely recurrent then every track in the sequence will also be transversely recurrent. However, even if τ is recurrent, the other tracks in the sequence need not be, so we don't automatically get a sequence of birecurrent tracks. On the other hand, for any carrying map $\sigma \rightarrow \tau$, if σ is recurrent then its image in τ will be contained in the largest recurrent subtrack and, furthermore, the largest recurrent subtrack has the same vertex cycles as the original track. In particular, if we replace each track in a sequence of moves with its largest recurrent subtrack, we have the following:

Theorem 3.10 [23, Theorem 1.3; 2, Theorem 1.1] *Let σ_i be a sequence of transversely recurrent train tracks such that $\sigma_{i+1} \xrightarrow{ss} \sigma_i$ is a single move. Then the sequence $B(\sigma_i)$ is a reparametrized quasigeodesic in $\mathcal{C}(\Sigma)$ with constants depending only on Σ .*

Lemma 3.11 *Let τ be a train track. Then $S(\tau)$ is quasiconvex, with uniform constants.*

Proof Let $a \in S(\tau)$. Split τ towards a . This gives a nested sequence of tracks and thus a quasigeodesic g_a from $B(\tau)$ to a that remains in $S(\tau)$.

If $a, b \in S(\tau)$ then, by hyperbolicity, $[a, b]$ is coarsely contained in $g_a \cup g_b \subset S(\tau)$. \square

The proof of the following lemma uses a technical result (Corollary A.6) whose proof is deferred to the appendix.

Lemma 3.12 *Assume that $P_\infty(\tau) \neq \emptyset$. Then $S(\tau)$ is the coarse convex hull of the set of ending laminations carried by τ .*

Proof As $\mathcal{C}(\Sigma)$ is hyperbolic any quasiconvex subset contains the coarse convex hull of its Gromov boundary. By Klarreich's theorem, the Gromov boundary of $\mathcal{C}(\Sigma)$ is the space of ending laminations. If $\gamma_i \in S(\tau)$ converge to the boundary then there exist $\lambda_i \in P(\tau)$ with $\lambda_i \rightarrow \lambda \in P(\tau)$ such that γ_i is a component of $[\lambda_i]$ and the Hausdorff limit of the γ_i contains the ending lamination $[\lambda]$. In particular, the Gromov boundary of $S(\tau)$ is exactly the ending laminations in $P_\infty(\tau)$, so $S(\tau)$ coarsely contains its convex hull.

By Corollary A.6, for any $a \in S(\tau)$, either a is uniformly close to $B(\tau)$ or there exists a sequence of ending laminations $\lambda_i \in P_\infty(\tau)$ such that the Hausdorff limit of $[\lambda_i]$ contains a . Then the projections of λ_i to the curve complex of the annulus around a go to infinity and so, by the bounded geodesic image theorem [22], when $j \gg i$ the geodesic between $[\lambda_i]$ and $[\lambda_j]$ passes within distance 1 of a . Therefore, either a is distance at most 1 from the convex hull of $S(\tau)$ or it is a bounded distance from $B(\tau)$. However, as $S(\tau)$ is quasiconvex, it is coarsely connected. Therefore, $S(\tau)$ is the coarse convex hull of the ending laminations carried by τ . \square

Lemma 3.13 *Let τ and σ be birecurrent train tracks with $\sigma \xrightarrow{ss} \tau$. Then $B(\sigma)$ is coarsely the closest point within $S(\sigma)$ to $B(\tau)$.*

Proof Consider a splitting sequence from τ to σ . It determines a quasigeodesic from $B(\tau)$ to $B(\sigma)$. Now, if $a \in S(\sigma)$ is any curve, the splitting sequence and the quasigeodesic can be continued until a crosses every branch at most once. This extended quasigeodesic ends at a and this proves the claim. \square

Lemma 3.14 *Let τ and σ be train tracks with $\sigma \xrightarrow{ss} \tau$. There exists a constant $C = C(\Sigma)$ such that if $\sigma \xrightarrow{ss} \tau$ has C or more moves then $\mathcal{A}(\sigma \xrightarrow{ss} \tau; \tau)$ contains a vertex cycle.*

Proof There is a bound, depending only on Σ , on the number of moves that are central splits and passing to subtracks. Therefore there will be tracks σ' and τ' in the

sequence $\sigma \xrightarrow{ss} \tau$ with $\sigma \xrightarrow{ss} \sigma' \xrightarrow{s} \tau' \xrightarrow{ss} \tau$ and $\sigma' \xrightarrow{s} \tau'$ having as many moves as we want, provided C is made large. For each right/left split there will be two branches that are each mapped to the union of two branches. Similarly, for each central split there will be two branches that are mapped to the union of three branches. Therefore, by increasing the number of moves we can guarantee that there is a branch b in σ' that is mapped to a legal path in τ' , and hence τ , of arbitrary length. Any legal path in τ that is sufficiently long will contain a subpath that closes up and that does not cross any branch exactly once. Thus all branches it crosses are in the active set. There is a further subpath that closes up and crosses each branch at most once. This gives a vertex cycle contained in the active set.

As all constants will only depend on Σ , this implies the lemma. □


Lemma 3.15 *Let τ, σ_1 and σ_2 be train tracks with $\sigma_1 \xrightarrow{ss} \tau$ and $\sigma_2 \xrightarrow{ss} \tau$ disjoint sequences. Then there exists a train track $\sigma \xrightarrow{ss} \sigma_i$ for $i = 1, 2$ with $V(\sigma) = V(\sigma_1) \cap V(\sigma_2)$ and both*

$$\min\{d(B(\sigma), B(\sigma_1)), d(B(\sigma), B(\sigma_2))\}, \quad \min\{d(B(\tau), B(\sigma_1)), d(B(\tau), B(\sigma_2))\}$$

uniformly bounded.

Proof We apply Lemma 3.5 to $\sigma_1 \xrightarrow{ss} \tau$ and $\sigma_2 \xrightarrow{ss} \tau$. In particular, we have a train track σ and a sequence $\sigma \xrightarrow{ss} \sigma_2$ that has the same number of moves and splitting moves as $\sigma_1 \xrightarrow{ss} \tau$ and $V(\sigma) = V(\sigma_1) \cap V(\sigma_2)$. Let C be the constant from Lemma 3.14. If $\sigma_1 \xrightarrow{ss} \tau$ has less than C moves then the distance bound follows from Lemma 3.9. If $\sigma_2 \xrightarrow{ss} \tau$ has less than C moves, we swap the roles of σ_1 and σ_2 and again the lemma follows. Therefore we can assume that both $\sigma_1 \xrightarrow{ss} \tau$ and $\sigma_2 \xrightarrow{ss} \tau$ have at least C moves.

Let σ' and τ' be the tracks in the sequences $\sigma \xrightarrow{ss} \sigma_2$ and $\sigma_1 \xrightarrow{ss} \tau$ that are C moves from σ and σ_1 . In particular, by Lemma 3.5, $\sigma' \xrightarrow{ss} \tau'$ with the same number of moves and splitting moves as $\sigma \xrightarrow{s} \tau$:



To bound $d(B(\sigma_1), B(\tau))$ we observe that as $\sigma_2 \xrightarrow{ss} \tau$ has more than C moves so by Lemma 3.14 there is a vertex cycle c in $|\mathcal{A}(\sigma_2 \xrightarrow{ss} \tau; \tau)| \subset \tau$. As $\sigma_1 \xrightarrow{ss} \tau$ and $\sigma_2 \xrightarrow{ss} \tau$ are disjoint it follows that c is in $\mathcal{S}(\sigma_1 \xrightarrow{ss} \tau; \tau)$ and hence is a vertex cycle in σ_1 . This gives our bound on $d(B(\sigma_1), B(\tau))$.

More generally, exactly the same argument works on any diamond-shaped diagram when the arrows represent $\geq C$ disjoint moves to show that the distance in $\mathcal{C}(\Sigma)$ between the vertex cycles of the bottom train track and the train tracks on the sides is uniformly bounded. Using the upper diamond in the diagram plus symmetry between σ_1 and σ_2 we conclude that $B(\sigma_1)$, $B(\tau')$, $B(\tau)$, $B(\sigma_2)$ and $B(\sigma')$ are all within uniform distance of each other. Finally we observe that as $\sigma \xrightarrow{ss} \sigma'$ is exactly C moves, we have a uniform bound on $d(B(\sigma), B(\sigma'))$ by Lemma 3.9. \square

When A is a geodesic lamination on Σ , we denote by $M(A)$ the lamination obtained from A by removing all isolated nonclosed leaves. Thus $M(A)$ consists of closed leaves and of minimal components and it is the maximal sublamination of A that supports a transverse measure. We call $M(A)$ the *measurable part* of A .

Lemma 3.16 *Suppose sequences a_i and b_i of closed geodesics converge to geodesic laminations A and B , respectively, in the Hausdorff topology. Assume*

- (i) *both sequences go to infinity in the curve complex $\mathcal{C}(\Sigma)$, and*
- (ii) *$d(a_i, b_i)$ is uniformly bounded.*

Then A and B have equal measurable parts, ie $M(A) = M(B)$.

Proof It suffices to prove the claim when $d(a_i, b_i) \leq 1$ for all i . Then A and B have no transverse intersections. If C is a minimal component of $M(A)$ that does not belong to $M(B)$, then it does not belong to B either, and so for large i the curve b_i is disjoint from the subsurface supporting C (which may be an annulus), contradicting (i). \square

3.4 Train tracks for cells

Given train tracks τ , σ_1 and σ_2 with $\sigma_i \xrightarrow{ss} \tau$, we would like to find a fourth track σ with $V(\sigma) = V(\sigma_1) \cap V(\sigma_2)$. If all three tracks are maximal and the relative interior of $V(\sigma_1) \cap V(\sigma_2)$ is open in $V(\tau)$, then this is due to Hamenstädt [14]. We begin with two preliminary results.

Lemma 3.17 Suppose $\sigma \xrightarrow{ss} \tau$ and $\tau' \subset \tau$ is a subtrack. Then there exists a subtrack σ' of σ with $V(\sigma') = V(\sigma) \cap V(\tau')$, $\sigma' \xrightarrow{ss} \tau'$ and the number of splitting moves not exceeding the number of splitting moves in $\sigma \xrightarrow{ss} \tau$.

Proof We first assume that $\sigma \xrightarrow{ss} \tau$ is a single move. The general case will follow by induction.

The intersection $V(\sigma) \cap V(\tau')$ will be a face of $V(\sigma)$ and hence there will be a subtrack $\sigma' \subset \sigma$ with $V(\sigma') = V(\sigma) \cap V(\tau')$. To show that $\sigma' \xrightarrow{ss} \tau'$ there are several cases for each type of move in $\sigma \xrightarrow{ss} \tau$:

- (1) $\sigma \xrightarrow{ss} \tau$ is a subtrack move. Then $\sigma' = \sigma \cap \tau'$ so σ' is a subtrack of τ' .
- (2) $\sigma \xrightarrow{s} \tau$ is a split or central split along a large branch b and τ' contains b and all its adjacent branches. Then $\sigma' \xrightarrow{s} \tau'$ is a single move on the same branch b .
- (3) $\sigma \xrightarrow{s} \tau$ is a split along b and one or more of the two large half-branches adjacent to b in σ is not in τ' . Then the restriction of the carrying map $\sigma \xrightarrow{s} \tau$ to σ' will be a switch-preserving homeomorphism, so σ' is a subtrack of τ' . See Figure 4.




Figure 4: If the branch a is removed in τ' then it also must be removed in σ' . However, then both small branches adjacent to a must be removed and σ' will be a subtrack of τ' .

- (4) $\sigma \xrightarrow{s} \tau$ is a central split and one or more of the half-branches adjacent to b in τ is not in τ' . Then, as in (3), σ' is a subtrack of τ' .
- (5) $\sigma \xrightarrow{s} \tau$ is a split and τ' contains both of the large half-branches adjacent to b in σ and does not contain one or more of the two adjacent small half-branches. In this case τ' is isotopic to a subtrack of σ and $\sigma' = \tau'$. See Figure 5. □





Figure 5: Case where τ' is a subtrack of σ .

Lemma 3.18 *Let τ, σ_1 and σ_2 be train tracks such that $\sigma_1 \xrightarrow{s} \tau$ and $\sigma_2 \xrightarrow{s} \tau$ are not disjoint. Then there exist train tracks $\tau_1 \subset \sigma_1, \tau_2 \subset \sigma_2$ and τ' such that $\tau' \xrightarrow{s} \tau$ and $\tau_i \xrightarrow{ss} \tau'$ with each sequence $\tau_i \xrightarrow{ss} \tau'$ having less splitting moves than $\sigma_i \xrightarrow{s} \tau$:*



Furthermore, $V(\tau_1) \cap V(\tau_2) = V(\sigma_1) \cap V(\sigma_2)$.

Proof If there is branch that is active in both sequences then by Lemma 3.8 there must be a large branch b that is active in both sequences. By Lemma 3.6 we can assume that the first move in both sequences is along b . If it is the same move then τ' is the track obtained from this first move and $\tau_i = \sigma_i$. If not then we let τ' be the central split on b . First suppose that both $\sigma_i \rightarrow \tau$ consist of a single move. For at least one of them, say $\sigma_1 \xrightarrow{s} \tau$, this move will be a right (or left) split on b and τ' will be obtained from σ_1 by removing the diagonal, and we set $\tau_1 = \tau'$. If $\sigma_2 \rightarrow \tau$ is a left (or right) split, we similarly put $\tau_2 = \tau'$. Finally, if $\sigma_2 \rightarrow \tau$ is the central split on b , we have $\tau_2 = \tau' = \sigma_2$.


In general, when $\sigma_i \rightarrow \tau$ have more than one move, we use the above paragraph for the first move and then apply Lemma 3.17.

Note that in all cases $V(\tau') \supset V(\sigma_1) \cap V(\sigma_2)$ and $V(\tau_i) = V(\sigma_i) \cap V(\tau')$. It follows that $V(\tau_1) \cap V(\tau_2) = V(\sigma_1) \cap V(\sigma_2)$. □

Proposition 3.19 *Let τ , σ_1 and σ_2 be train tracks such that $\sigma_i \xrightarrow{ss} \tau$ for $i = 1, 2$. Assume that $V(\sigma_1) \cap V(\sigma_2) \neq \emptyset$. Then there exist train tracks σ^\pm and subtracks $\sigma'_i \subset \sigma_i$ such that*

- (a) $\sigma^- \xrightarrow{ss} \sigma'_i \xrightarrow{ss} \sigma^+ \xrightarrow{ss} \tau$ for $i = 1, 2$;
- (b) $\sigma'_1 \xrightarrow{ss} \sigma^+$ and $\sigma'_2 \xrightarrow{ss} \sigma^+$ are disjoint;
- (c) $V(\sigma^-) = V(\sigma_1) \cap V(\sigma_2)$;
- (d) $\min\{d(B(\sigma_1), B(\sigma^-)), d(B(\sigma_2), B(\sigma^-))\}$ is uniformly bounded;
- (e) $\min\{d(B(\sigma_1), B(\sigma^+)), d(B(\sigma_2), B(\sigma^+))\}$ is uniformly bounded.

See



Proof If $\sigma_1 \xrightarrow{ss} \tau$ and $\sigma_2 \xrightarrow{ss} \tau$ are disjoint then the proposition follows from Lemma 3.15 with $\sigma^+ = \tau$ and σ^- the track given by Lemma 3.15. If not, we describe an algorithm that replaces σ_1 and σ_2 with subtracks τ_1 and τ_2 and τ with a train track τ' such that $\tau_i \xrightarrow{ss} \tau'$ and $\tau' \xrightarrow{ss} \tau$. Furthermore one of the following will hold:

- (i) $\dim V(\tau') < \dim V(\tau)$;
- (ii) the number of splitting moves in $\tau_i \xrightarrow{ss} \tau'$ is less than in $\sigma_i \xrightarrow{ss} \tau$.

In addition, neither the dimension nor the number of moves ever increases. If $\tau_1 \xrightarrow{ss} \tau'$ and $\tau_2 \xrightarrow{ss} \tau'$ are disjoint then, as above, the proposition follows from Lemma 3.15. If (i) or (ii) hold, we apply the algorithm to the three new tracks. Both (i) and (ii) can only happen a finite number of times, so we must eventually have that the two sequences are disjoint.

We now describe the algorithm. Let $\tau' \subset \tau$ be the smallest birecurrent subtrack such that $V(\sigma_1) \cap V(\sigma_2) \subset V(\tau')$.

- (1) If τ' is a proper subtrack of τ , we let $\tau_i \subset \sigma_i$ be the subtracks given by Lemma 3.17. In this case, (i) holds.

- (2) If $\tau' = \tau$ then $V(\sigma_i)$ intersects the relative interior of $V(\tau)$, so by Lemma 3.7 we can assume that $\sigma_i \xrightarrow{s} \tau$. By assumption, $\sigma_1 \xrightarrow{s} \tau$ and $\sigma_2 \xrightarrow{s} \tau$ are not disjoint and we replace σ_1, σ_2 and τ with the tracks τ_1, τ_2 and τ' given by Lemma 3.18. In this case, (ii) holds. \square

Proposition 3.20 *Assume that $\sigma \xrightarrow{ss} \tau$ and that $\theta \xrightarrow{s} \tau$ is a central split such that $V(\sigma) \cap V(\theta)$ is the intersection of $V(\sigma)$ with a hyperplane that intersects the relative interior of $V(\sigma)$. Then there is a central split $\theta' \xrightarrow{s} \sigma$ such that $V(\theta') = V(\sigma) \cap V(\theta)$.*

Proof Say $\theta \xrightarrow{s} \tau$ is the central split on the large branch b . If b is not in the stationary set for $\sigma \xrightarrow{ss} \tau$, then by Lemma 3.6 we can assume that the first move in $\sigma \xrightarrow{ss} \tau$ is on b . But then the hyperplane assumption cannot hold. Thus b is in the stationary set and is a large branch in σ . We define θ to be the central split in b . The conclusion now follows from Lemma 3.5. \square

Proposition 3.21 *Let \mathcal{C}_j be an excellent sequence of cell structures obtained by splitting a train track τ . To every cell $E \in \mathcal{C}_j$ one can assign a birecurrent train track θ_E satisfying the following:*

- (1) $E = V(\theta_E)$.
- (2) If E is a top-dimensional cell, then θ_E is the track associated to E in the definition of the splitting sequence.
- (3) If $F \subset E$ are cells then $\theta_F \xrightarrow{ss} \theta_E$.
- (4) There is a constant $C = C(\Sigma)$ such that for each cell $F \in \mathcal{C}_k$ there is a top-dimensional cell $E \in \mathcal{C}_k$ with $F \subset E$ and $d(B(\theta_E), B(\theta_F)) \leq C$.

In particular, if all top-dimensional cells in \mathcal{C}_j have vertex cycles distance at most B from $B(\tau)$ then $d(B(\tau), B(\theta_E)) \leq B + C$ while if $E \subset F$ for a cell $F \in \mathcal{C}_j$ with $d(B(\tau), B(\theta_F)) \geq A$ then $d(B(\tau), B(\theta_E)) \geq A - C$.

Proof We define θ_E for $E \in \mathcal{C}_j$ by induction on j . When $j = 0$ each cell E is naturally associated to a subtrack of τ and we define θ_E to be this subtrack. Now suppose that θ_E has been defined for all cells in \mathcal{C}_j of dimension $> i$ for a certain $i < n$. Let $F \in \mathcal{C}_j$ with $\dim F = i$. By property (C5) of an excellent sequence, if E_1, \dots, E_ℓ are all i -dimensional cells in \mathcal{C}_j with $F \subset E_s$ then $F = \cap E_s$. Let $F_k = E_1 \cap \dots \cap E_k$. Via induction we have tracks θ_{F_k} with $V(\theta_{F_k}) = F_k$ and if $E \in \mathcal{C}_j$ with $E_s \subset E$ for some $s = 1, \dots, k$ then $\theta_{F_k} \xrightarrow{ss} \theta_E$. The track θ_{F_k} is defined

by applying Proposition 3.19 to $\theta_{F_{k-1}}$ and θ_{E_k} . If θ_F is not recurrent, we can replace it with its largest recurrent subtrack. We then set $\theta_F = \theta_{F_\ell}$ and this track will satisfy properties (1)–(3).

To get the distance bound in (4) we observe that Proposition 3.19(c) gives a bound that is linear in ℓ . While we cannot a priori control the size of ℓ , once we know that $F = V(\theta_F)$ for the train track θ_F we observe that the number of codimension one faces of F is bounded by the number of small branches of θ_F and hence a constant only depending on Σ . In particular there is a subcollection of the E_1, \dots, E_ℓ of uniformly bounded size whose intersection gives F by Corollary 2.6. Applying the argument of the previous paragraph to this subcollection, we get a track θ'_F with $V(\theta'_F) = F$ and the distance bound in (4).

Finally we note that while θ_F and θ'_F may not be the same track (and θ'_F may not satisfy (3)), since $V(\theta_F) = V(\theta'_F)$ the two tracks have the same vertex cycles and therefore (4) holds for θ_F also. □

Given a lamination $\lambda \in P_\infty(\tau)$ let τ_i be a sequence of tracks such that $\tau_0 = \tau$, $\tau_{i+1} \xrightarrow{ss} \tau_i$ is a single move and $\lambda \in P_\infty(\tau_i)$ for all i . We say that the sequence is a *full splitting sequence* if for every i and every large branch b in τ_i there exists an i_n such that $b \in \mathcal{A}(\tau_{i_n} \xrightarrow{ss} \tau_i; \tau_i)$.

Proposition 3.22 *Assume that λ is fully carried by τ . Then there exists a full splitting sequence $\tau = \tau_1, \tau_2, \dots$ such that λ is fully carried by every τ_i . Moreover, any infinite splitting sequence starting at τ and carrying λ is a full splitting sequence. Furthermore, if λ' is carried by every τ_i then $[\lambda] = [\lambda']$.*

Proof The first statement follows from [1, Lemma 2.1]. In fact, the proof of [1, Lemma 2.1] proves the stronger second statement. The third statement is probably well known but as we could not find a proof we provide one here. Assume that λ' is carried by all τ_i but $[\lambda] \neq [\lambda']$. By [24, Corollary 1.7.13] we can find a birecurrent train track τ' that carries λ , does not carry λ' and is carried by τ . Hence it will fully carry λ , but it may not come from a sequence of splits and central splits of τ . Instead we use [24, Theorem 2.3.1] to find a track σ with $\sigma \xrightarrow{ss} \tau'$, $\sigma \xrightarrow{ss} \tau$ and λ carried by σ . As all three tracks fully carry λ , we in fact have $\sigma \xrightarrow{s} \tau'$ and $\sigma \xrightarrow{s} \tau$.

We will show that for sufficiently large i we have $\tau_i \rightarrow \sigma$. As τ_i carries λ' but σ does not this will be a contradiction. We repeatedly apply Proposition 3.19. Let $\sigma_1^+ = \tau$ and

assume that we have constructed tracks $\sigma_1^+, \dots, \sigma_{j-1}^+$ with $\sigma_i^+ \xrightarrow{s} \sigma_{i-1}^+$, $\sigma \xrightarrow{s} \sigma_i$, $\tau_i \xrightarrow{s} \sigma_i^+$ and $\sigma \xrightarrow{s} \sigma_i^+$ and $\tau_i \xrightarrow{s} \sigma_i^+$ are disjoint. As $\tau_j \xrightarrow{ss} \tau_{j-1}$ we have $\tau_j \xrightarrow{s} \sigma_{j-1}^+$ and we can apply Proposition 3.19 to $\tau_j \xrightarrow{ss} \sigma_{j-1}^+$ and $\sigma \xrightarrow{s} \sigma_{j-1}^+$ and let $\sigma_j^+ = \sigma^+$, where σ^+ is as given in the proposition. Note that since λ is fully carried by all of the tracks, all the carrying maps given by Proposition 3.19 are fully carrying. This also implies that λ is in the relative interior of the associated cells so we also never need to pass to subtracks. In particular, $\sigma_j^+ \xrightarrow{s} \sigma_{j-1}^+$, $\sigma \xrightarrow{s} \sigma_j$ and $\tau_i \xrightarrow{s} \sigma_j^+$, so the induction step is complete.

When we apply Proposition 3.19, if $\tau_j \xrightarrow{s} \sigma_{j-1}^+$ and $\sigma \xrightarrow{s} \sigma_{j-1}^+$ are disjoint then $\sigma_j^+ = \sigma_{j-1}^+$ and $\sigma \xrightarrow{s} \sigma_{j-1}^+$ and $\sigma \xrightarrow{s} \sigma_j^+$ have the same number of moves. If not, then as $\tau_j \xrightarrow{s} \sigma_{j-1}^+$ factors as $\tau_j \xrightarrow{s} \tau_{j-1} \xrightarrow{s} \sigma_{j-1}^+$ with $\tau_j \xrightarrow{s} \tau_{j-1}$ a single move and $\tau_{j-1} \xrightarrow{s} \sigma_{j-1}^+$ and $\sigma \xrightarrow{s} \sigma_{j-1}^+$ disjoint, we have that $\sigma_j^+ \xrightarrow{s} \sigma_{j-1}^+$ is a single move and $\sigma \xrightarrow{s} \sigma_j^+$ has one less move than $\sigma \xrightarrow{s} \sigma_{j-1}^+$. This implies that the composition of sequences $\sigma \xrightarrow{s} \sigma_j^+ \xrightarrow{s} \tau$ has the same number of moves as the original sequence $\sigma \xrightarrow{s} \tau$. In particular, the number of times that $\sigma_j^+ \neq \sigma_{j-1}^+$ is bounded by the number of moves in $\sigma \xrightarrow{s} \tau$ and there must exist an N such that if $j > N$ then $\sigma_j^+ = \sigma_N^+$. The sequence $\sigma_N^+, \tau_N, \tau_{N+1}, \dots$ is a full splitting sequence so, for i sufficiently large, $\mathcal{A}(\tau_i \xrightarrow{s} \sigma_N; \sigma_N)$ is all of σ_N . The active branches for $\sigma \xrightarrow{s} \sigma_N$ must be disjoint from $\mathcal{A}(\tau_i \xrightarrow{s} \sigma_N; \sigma_N)$ so we must have $\sigma = \sigma_N$ and $\tau_N \xrightarrow{s} \sigma$, as desired. □

3.5 A shortening argument

In this section we assume that σ , τ and ρ are *partial train tracks*, ie each is a subgraph of a train track. We allow valence 2 vertices with the turn illegal, or even valence 1 vertices. Even though the main result is used only when τ and ρ are train tracks, the extra flexibility of passing to subgraphs will make the proof easier. More precisely, we assume

- τ and ρ are two partial train tracks on Σ ,
- σ is the graph that consists of edges that τ and ρ have in common,
- branches of $\tau - \sigma$ and $\rho - \sigma$ intersect transversally and any vertex in common to τ and ρ is also a vertex of σ ,
- any lamination carried by both τ and ρ is carried by σ , ie

(*)
$$V(\tau) \cap V(\rho) = V(\sigma).$$

Given a triple $\sigma = (\sigma; \tau, \rho)$ as above, define the complexity $\chi(\sigma)$ to be the pair $(e(\sigma) + I(\tau, \rho), e(\sigma))$, ordered lexicographically, where $e(\sigma)$ is the number of edges of σ and $I(\tau, \rho)$ is the number of transverse intersections between the branches of τ and ρ . Note that for a given complexity there are only finitely many σ up to the action of the mapping class group.

The number of branches of σ is uniformly bounded depending only on the surface Σ , so the bound on $\chi(\sigma)$ really only amounts to the bound on the intersection between the branches of τ and ρ .

As an example of the extra flexibility, note that if we remove an edge of σ from all three graphs σ , τ and ρ , the listed conditions continue to hold, but the new triple has smaller complexity. In the proof below, the intersection number $I(\tau, \rho)$ will increase only if $e(\sigma)$ decreases by at least as much.

Denote by $\text{Supp}(\sigma)$ the support of σ , ie the smallest subsurface that contains σ (possibly \emptyset , or disconnected, or all of Σ). Thus $\text{Supp}(\sigma) = \emptyset$ if and only if σ is contained in a disk.

Lemma 3.23 *For every C and every χ there is $C' = C'(\Sigma, C, \chi)$ such that if $\chi(\sigma) \leq \chi$, and $a \in S(\tau)$ and $b \in S(\rho)$ with $d(a, b) \leq C$, then:*

(i) *If $\text{Supp}(\sigma) = \emptyset$ then*

$$d(a, B(\tau)) \leq C' \quad \text{and} \quad d(b, B(\rho)) \leq C'.$$

(ii) *If $\text{Supp}(\sigma) \neq \emptyset, \Sigma$ then*

$$d(a, \mathcal{C}(\text{Supp}(\sigma))) \leq C' \quad \text{and} \quad d(b, \mathcal{C}(\text{Supp}(\sigma))) \leq C'.$$

(iii) *If $\text{Supp}(\sigma) = \Sigma$ then*

$$d(a, S(\sigma)) \leq C' \quad \text{and} \quad d(b, S(\sigma)) \leq C'.$$

In (ii), by $\mathcal{C}(\text{Supp}(\sigma))$ we mean the set of curves carried by $\text{Supp}(\sigma)$, even when $\text{Supp}(\sigma)$ is disconnected.

Most of the time when we apply Lemma 3.23, we will have that τ and ρ are subtracks of some large track ω and $\sigma = \tau \cap \rho$, and then the condition (*) is standard and quickly follows from the fact that legal paths in the universal cover are quasigeodesics and that they are uniquely determined by their endpoints on the circle at infinity. The proof of Lemma 3.23 is by modifying the tracks and then τ and ρ may develop intersecting branches.

If σ is a train track or a partial train track and a is carried by σ , then the *combinatorial length* $\ell_\sigma(a)$ is the sum of the weights of a .

Proof of Lemma 3.23 We will suppose such C' does not exist and obtain a contradiction.

(i) If the lemma fails for a particular σ , there are sequences of curves $a_n \in S(\tau)$ and $b_n \in S(\rho)$ such that $d(a_n, b_n) \leq C$, $d(a_n, B(\tau)) > n$ and $d(b_n, B(\rho)) > n$. After passing to a subsequence n_j , we may assume that $a_{n_j} \rightarrow A$ and $b_{n_j} \rightarrow B$ in the Hausdorff topology, where A and B are geodesic laminations. By Lemma 3.16, A and B have the same (nonempty) measurable part Λ , which must be carried by σ by assumption (*). This contradicts the assumption that σ is contained in a disk.

(ii) We induct on the complexity.

For each σ with $\chi(\sigma) \leq \chi$ where the lemma fails, there are curves $a_n^\sigma \in S(\tau)$ and $b_n^\sigma \in S(\rho)$ with $d(a_n^\sigma, b_n^\sigma) \leq C$, $d(a_n^\sigma, \mathcal{C}(\text{Supp}(\sigma))) > n$ and $d(b_n^\sigma, \mathcal{C}(\text{Supp}(\sigma))) > n$ for every n . We will assume that, subject to these conditions,

$$\ell_\tau(a_n^\sigma) + \ell_\rho(b_n^\sigma)$$

is minimal possible.

To obtain a contradiction we will find a sequence of triples $\sigma_i = (\sigma_i; \tau_i, \rho_i)$ where the lemma fails with $\sigma = \sigma_1$ and for each σ_i an infinite sequence $\{n_{i,j}\}$ such that

- (1) $n_{1,j} = j$ and $\{n_{i,j}\}$ is a subsequence of $\{n_{i-1,j}\}$ for $i > 1$;
- (2) $a_{n_{i,j}}^{\sigma_i} \in S(\tau_i)$, $b_{n_{i,j}}^{\sigma_i} \in S(\rho_i)$;
- (3) $d(a_{n_{i,j}}^{\sigma_i}, b_{n_{i,j}}^{\sigma_i}) \leq C$;
- (4) $d(a_{n_{i,j}}^{\sigma_i}, \mathcal{C}(\text{Supp}(\sigma_i))) > n_{i,j}$ and $d(b_{n_{i,j}}^{\sigma_i}, \mathcal{C}(\text{Supp}(\sigma_i))) > n_{i,j}$;
- (5) $\ell_{\tau_i}(a_{n_{i,j}}^{\sigma_i}) + \ell_{\rho_i}(b_{n_{i,j}}^{\sigma_i}) < \ell_{\tau_i}(a_{n_{i,j}}^{\sigma_{i-1}}) + \ell_{\rho_i}(b_{n_{i,j}}^{\sigma_{i-1}})$;
- (6) for every i and j , $\ell_{\tau_i}(a_{n_{i,j}}^{\sigma_i}) + \ell_{\rho_i}(b_{n_{i,j}}^{\sigma_i})$ is minimal possible subject to (2)–(4);
- (7) $\chi(\sigma_i) \leq \chi(\sigma_{i-1})$;
- (8) σ_i satisfies (*), ie $V(\tau_i) \cap V(\rho_i) = V(\sigma_i)$.

By (7) our sequence σ_i must eventually repeat (up to $\text{Mod}(\Sigma)$), so there are $k < l$ with $\sigma_k = \phi(\sigma_l)$ for some mapping class ϕ . By repeated applications of (5), we have

$$\ell_{\tau_l}(a_{n_{l,j}}^{\sigma_l}) + \ell_{\rho_l}(b_{n_{l,j}}^{\sigma_l}) < \ell_{\tau_k}(a_{n_{k,j}}^{\sigma_k}) + \ell_{\rho_k}(b_{n_{k,j}}^{\sigma_k}),$$

obtaining our contradiction to (6), since for large j the curves $\phi(a_{n_{i,j}}^{\sigma_i})$ and $\phi(b_{n_{i,j}}^{\sigma_i})$ satisfy (2)–(4) (for $i = k$) and have smaller total combinatorial length than $a_{n_{k,j}}^{\sigma_k}$ and $b_{n_{k,j}}^{\sigma_k}$.

We will construct the sequence σ_i inductively. Assume that σ_i and the sequence $\{n_{i,j}\}$ have been defined satisfying the above conditions. We then define a subsequence $\{n_{i+1,j}\}$ of $\{n_{i,j}\}$ and show that there exists a σ_{i+1} such that (1)–(8) hold with suitable choices of curves. We first choose the subsequence $\{n_{i+1,j}\}$ so that $a_{n_{i+1,j}}^{\sigma_i} \rightarrow A$ and $b_{n_{i+1,j}}^{\sigma_i} \rightarrow B$, where A and B are two geodesic laminations and convergence is with respect to Hausdorff topology. The construction of σ_{i+1} is more involved.

Lemma 3.16 implies that A and B have the same measurable part Λ and differ only in isolated nonclosed leaves. By assumption (8), Λ is carried by σ_i . Let $\bar{\sigma} \subset \sigma_i$ be the union of the branches crossed by Λ . Thus $\bar{\sigma}$ is a train track.

Case 1 ($\bar{\sigma}$ has at least one illegal turn) Note that Λ supports a transverse measure of full support and in particular $\bar{\sigma}$ has a large branch (one with maximal transverse measure). Split along this branch so that Λ is still carried to obtain a new track σ_{i+1} .

Case 1a The nondegenerate case that such a split is unique (ie the pairs of weights at the two ends are distinct) is pictured in Figure 6. The vertical segment represents a large branch of $\bar{\sigma}$ and the two branches at the top and at the bottom are also in $\bar{\sigma}$. The branches pictured on the sides are branches of $\sigma_i - \bar{\sigma}$, $\tau_i - \sigma_i$ or $\rho_i - \sigma_i$. The splitting operation consists of cutting along the large branch thus producing two vertical branches of the split $\bar{\sigma}$, adding the suitable diagonal branch so that Λ is carried, and attaching the side branches at exactly the same point, to either the left or the right vertical branch.




Figure 6

We define σ_{i+1} to be the split version of σ_i . Thus σ_{i+1} includes the two vertical branches, the two branches at the top, the two branches at the bottom, the diagonal branch and any side branches that came from $\sigma_i - \bar{\sigma}$. The track τ_{i+1} contains σ_{i+1} and includes side branches that came from $\tau_i - \bar{\sigma}_i$, and similarly for ρ_{i+1} . Observe that $\chi(\sigma_{i+1}) = \chi(\sigma_i)$, so (7) holds.

Claim For large j , $a_{n_{i+1},j} \in S(\tau_{i+1})$ and $b_{n_{i+1},j} \in S(\rho_{i+1})$.

Indeed, there are leaves of Λ that cross from the upper left (right) to the lower left (right) branch on the left diagram in Figure 6, and likewise from upper left to lower right. The same is therefore true for segments of $a_{n_{i+1},j}$ for large j . This prevents $a_{n_{i+1},j}$ from entering the vertical segment say from a side branch on the left and exiting through a side branch on the right, or the top or bottom right branch. Since such configurations do not occur, the claim holds.

Thus, after discarding an initial portion of each sequence, properties (2)–(5) hold (for (4) note that $\text{Supp}(\sigma_{i+1}) = \text{Supp}(\sigma_i)$ and for (5) note that since $a_{n_{i+1},j}$ contains segments that cross from upper left to lower left, from upper right to lower right and from upper left to lower right, the combinatorial length strictly decreases after the split). Now define $a_{n_{i+1},j}^{\sigma_{i+1}}$ and $b_{n_{i+1},j}^{\sigma_{i+1}}$ to be a pair of curves that minimize the sum of the combinatorial lengths, subject to (2)–(5).

It remains to prove (8). Let Ω be a lamination carried by τ_{i+1} and by ρ_{i+1} . It is therefore carried by τ_i and ρ_i , so by (8) for σ_i it is carried by σ_i . Now we again have to argue that certain configurations do not occur, eg that leaves of Ω do not enter on a left side branch and exit on a right side branch. If this occurs then Ω would not be carried by τ_{i+1} or ρ_{i+1} .

Case 1b In the degenerate case when both splits carry Λ (ie when Λ does not cross the diagonally drawn branches in Figure 7), we define σ_{i+1} to be the track obtained from σ_i by cutting open along the vertical segment. Thus σ_{i+1} does not include either of the diagonal branches. See Figure 7.

Next, we observe that for large j the curves $a_{n_{i+1},j}^{\sigma_{i+1}}$ cannot cross both from top left to bottom right and from top right to bottom left, and the same is true for $b_{n_{i+1},j}^{\sigma_{i+1}}$. Thus after passing to a further subsequence we can add one of the two diagonal branches to τ_{i+1} and ensure that $a_{n_{i+1},j}^{\sigma_{i+1}} \in S(\tau_{i+1})$, and likewise $b_{n_{i+1},j}^{\sigma_{i+1}} \in S(\rho_{i+1})$ after including one of the two diagonals. It is possible that one diagonal is added to τ_{i+1} and the other to ρ_{i+1} and then the intersection number increases by 1. But the number




Figure 7

of branches of σ_{i+1} decreased, so we still have $\chi(\sigma_{i+1}) < \chi(\sigma_i)$ and we are done by induction. If the same diagonal is added to both τ_{i+1} and to ρ_{i+1} , we will also add it to σ_{i+1} . The rest of the argument is similar to the nondegenerate case.

Case 2 ($\bar{\sigma}$ does not have any illegal turns) Thus $\bar{\sigma}$ is a collection of legal simple closed curves and so is Λ . In A and B there must be isolated leaves spiraling towards each component of Λ , in opposite directions on the two sides. The spiraling directions are the same for both A and B , since otherwise the projection distance on the curve complex of the annulus would be large. In other words, both $a_{n_{i,j}}^{\sigma_i}$ and $b_{n_{i,j}}^{\sigma_i}$ wind around the same annulus and in the same direction a large number of times. Applying the Dehn twist (left, or right, as appropriate) shortens both curves as they wind around the annulus one less time. At the same time this operation does not change the distance to $\mathcal{C}(\text{Supp}(\sigma))$. This contradicts the minimality and we are done.

(iii) Again the proof is by induction on the complexity. We will inductively assume (1)–(8) except that (4) is replaced with

$$(4') \quad d(a_{n_{i,j}}^{\sigma_i}, S(\sigma_i)) > n_{i,j} \text{ and } d(b_{n_{i,j}}^{\sigma_i}, S(\sigma_i)) > n_{i,j}.$$

The proof follows closely our proof of (ii). As in that proof, we pass to a further subsequence and construct limiting laminations A and B that have a common measurable part Λ which is carried by σ_i , and $\bar{\sigma}$ is the union of the edges of σ crossed by Λ .

There are two cases, as in (ii).

Case 1 ($\bar{\sigma}$ contains an illegal turn) We split along a large branch of $\bar{\sigma}$ as before and define σ_{i+1} in the same way (in both subcases, whether the split is degenerate or

nondegenerate). The only change is that now we have to argue that (4') holds, instead of (4). The reason now is that $S(\sigma_{i+1}) \subset S(\sigma_i)$.

Case 2 ($\bar{\sigma}$ is a collection of legal loops) Now we cannot simply apply a Dehn twist since this does not necessarily preserve $S(\sigma_i)$. Note that there must be branches of σ_i attached to both sides of $\bar{\sigma}$ for otherwise we would be in situation (ii).

Case 2a (all branches of σ_i attached to a component of $\bar{\sigma}$ are attached in the same direction) See Figure 8.




Figure 8

Thus there is a unique curve in $S(\sigma_i)$ that crosses an edge of this component of $\bar{\sigma}$ (and it equals the component). Therefore the Dehn twist preserves $S(\sigma_i)$ and we proceed as before.

Case 2b (there are branches of σ_i attached to a component of $\bar{\sigma}$ in opposite directions) We will assume here that every branch of σ_i is crossed by either $a_{n_i,j}^{\sigma_i}$ or by $b_{n_i,j}^{\sigma_i}$ (or both) for every j , for otherwise we can remove this edge from all three σ_i , τ_i and ρ_i and use induction.

Then we can find two branches of σ_i attached in opposite directions and on opposite sides of this component of $\bar{\sigma}$ (the curves $a_{n_i,j}^{\sigma_i}$ or $b_{n_i,j}^{\sigma_i}$ spiral and cannot escape on the same side). In other words, we have a picture as in Figure 6, where the vertical segment as well as top left and lower right branches (or top right and lower left branches) belong to $\bar{\sigma}$, and the top right and the lower left branches (or top left and lower right branches) belong to $\sigma_i - \bar{\sigma}$. Perform the split as in Figure 6 so that Λ is carried. If there are any side branches attached to the vertical segment, then after the split the number of side branches attached to $\bar{\sigma}$ is strictly smaller and we may induct on this number.

If there are no such side branches, then the combinatorial lengths of $a_{n_i,j}^{\sigma_i}$ and $b_{n_i,j}^{\sigma_i}$ strictly decrease after the split (eg consider a piece of $a_{n_i,j}^{\sigma_i}$ that enters $\bar{\sigma}$ through the top branch which is not part of $\bar{\sigma}$). Then proceed as before, by defining $a_{n_{i+1},j}^{\sigma_{i+1}}$ and $b_{n_{i+1},j}^{\sigma_{i+1}}$ to be curves that minimize combinatorial length subject to (2)–(4'). \square

We will only use two special cases of Lemma 3.23, and we state them below.

Corollary 3.24 *For every $C > 0$ there is $C' > 0$ depending only on the surface Σ such that the following holds. Let τ be a large track on Σ . Assume one of the following:*

- (I) τ_1 and τ_2 are large subtracks of τ . Let $\sigma = \tau_1 \cap \tau_2$. After pruning dead ends, σ becomes a track (possibly empty) and $V(\sigma) = V(\tau_1) \cap V(\tau_2)$.
- (II) τ_1 and τ_2 are the two tracks obtained from τ by splitting along a large branch e , and σ the track obtained by a central split at e . Thus $P(\tau) = P(\tau_1) \cup P(\tau_2)$ and $P(\sigma) = P(\tau_1) \cap P(\tau_2)$.

Then one of the following holds:

- (1) σ is not large (possibly it is empty), and for any two curves $a_i \in P(\tau_i)$ with $d(a_1, a_2) \leq C$ it follows that $d(a_i, B(\tau_i)) \leq C'$, or
- (2) σ is large and for any two curves $a_i \in S(\tau_i)$ with $d(a_1, a_2) \leq C$ there is a curve $c \in S(\sigma)$ such that $d(a_i, c) \leq C'$.

4 Cell structures via splittings

Now we take $U = V(\tau)$ for a recurrent, transversely recurrent, maximal train track τ .

Let \mathcal{C}_j be the excellent sequence obtained by repeating the subdivision process, at every step choosing one of the codimension 0 cells $V(\sigma)$, with σ a recurrent, transversely recurrent, maximal train track, and splitting in a selected large branch. Thus, inductively, codimension 0 cells in \mathcal{C}_j are in one-to-one correspondence with a set of train tracks, each obtained from τ by a splitting sequence, and all tracks in the splitting sequence correspond to codimension 0 cells in \mathcal{C}_k for $k \leq j$.

4.1 Interpolating curves process

In this section we set the groundwork for proving that the distance between disjoint cells is not too small. This follows easily from Lemma 3.23 when the associated train

tracks have bounded intersection number. To handle the general case we define a certain iterative procedure that constructs sequences of curves relating different cells in \mathcal{C}_j .

We start by defining a sequence C_0, C_1, \dots inductively. Here $C_0 > 0$ is a fixed constant, and C_{i+1} is defined as C' in Corollary 3.24 for the constant $C = 2^n C_i$, where $n = \dim \mathcal{ML}$.

When a is a simple closed curve we denote by $\text{Carr}_j(a)$ the *carrier* of a in \mathcal{C}_j , ie the smallest cell of \mathcal{C}_j that contains a .

Definition 4.1 A sequence of curves $\mathbf{a} = a_0, a_2, \dots, a_m$ in Σ is *good* with respect to the cell structure \mathcal{C}_j (or \mathcal{C}_j -good) if for any two adjacent curves a_i and a_{i+1} in the sequence the carriers $\text{Carr}_j(a_i)$ and $\text{Carr}_j(a_{i+1})$ are nested (or possibly equal), ie $\text{Carr}_j(a_i) \subseteq \text{Carr}_j(a_{i+1})$ or $\text{Carr}_j(a_{i+1}) \subseteq \text{Carr}_j(a_i)$.

A sequence which is \mathcal{C}_j -good may not be \mathcal{C}_{j+1} -good. We now describe an inductive procedure that consists of inserting curves to produce \mathcal{C}_k -good sequences with k large.

We start with a \mathcal{C}_0 -good sequence of bounded length. For example, we might start with a sequence a_0, a_1 of length 2 consisting of two curves in the interior of the same cell in \mathcal{C}_0 . Inductively assume that we inserted some curves in the sequence and obtained a \mathcal{C}_j -good sequence \mathbf{a} .




Figure 9: Interpolating points to achieve goodness at the next stage.

Suppose a and b are two consecutive curves in \mathbf{a} that fail to satisfy the definition of \mathcal{C}_{j+1} -good, that is, the carriers $A = \text{Carr}_{j+1}(a)$ and $B = \text{Carr}_{j+1}(b)$ are not nested. There are several cases:

- (i) $\text{Carr}_j(a) = \text{Carr}_j(b)$; we call this cell C . Thus the subdivision operation splits C into A and B (if the cut contained either point, A and B would be nested) and $A \cap B = W$ is the codimension 1 cut. See the left diagram in Figure 9. We now apply Corollary 3.24(II) to the train tracks τ_1, τ_2 and σ such that $V(\tau_1)$ and $V(\tau_2)$ are the two splits of $V(\tau) = C$ (so τ_1 and τ_2 have branches e_1 and e_2

that intersect) and σ is a common subtrack of τ_1 and τ_2 obtained by deleting e_1 from τ_1 or e_2 from τ_2 (these tracks exist by Proposition 3.20). Therefore we obtain a curve $w \in P(\sigma) = W$, and if $d(a, b) \leq C_i$, we have in addition that $d(a, w) \leq C_{i+1}$ and $d(b, w) \leq C_{i+1}$. We insert w in the sequence between a and b . The consecutive curves in a , w and b satisfy the C_{j+1} -goodness condition.

- (ii) $\text{Carr}_j(a) \not\subseteq \text{Carr}_j(b)$. This is depicted in the other two diagrams in Figure 9. Notice that the cut W cannot contain a , or else the goodness condition would hold in C_{j+1} . There are two further subcases. If b does not belong to W either, we are in the situation of the middle diagram. First apply Corollary 3.24(II) as in (i) above to find $w \in W$. Then apply Corollary 3.24(I) to curves a and w to find a curve v carried by the intersection of the C_{j+1} -carriers of a and w . Finally, interpolate to get the sequence a, v, b . The other subcase is that $b \in W$, depicted in the right diagram in Figure 9. We again interpolate v in the intersection of C_{j+1} -carriers of a and b .

Whenever we apply Corollary 3.24 it may happen that conclusion (1) occurs. In that case we stop the process and do not attempt to define a C_{j+1} -good sequence.

To the C_j -good sequence $\mathbf{a}_j = a_0, a_1, \dots, a_m$ constructed in this way we will associate a *dimension sequence* $D(\mathbf{a}_j)$ inductively. This is a sequence of nonnegative integers d_0, d_1, \dots, d_m with the requirement that the dimension of $\text{Carr}_j(a_i)$ is $\leq d_i$. It is also constructed inductively. For the initial sequence we take the dimensions of the C_0 -carriers. Inductively, we extend the dimension sequence. For each curve x that is inserted when extending the sequence from \mathbf{a}_j to \mathbf{a}_{j+1} define the corresponding integer as the dimension of $\text{Carr}_{j+1}(x)$. For curves that were part of the sequence \mathbf{a}_j leave the value unchanged. Thus the number associated to a curve in the sequence is the dimension of its carrier when the curve first appeared. The dimension of the carrier of a curve may decrease, but the value in the dimension sequence is unchanged.

The following proposition summarizes the essential features of the construction.

Proposition 4.2 *Suppose that a curve x got inserted between the curves a and b in a C_j -good sequence \mathbf{a}_j .*

- (i) *The value of the dimension sequence at x is strictly less than at both a and b .*
- (ii) *If $d(a, b) \leq C_i$ for some i then $d(a, x) \leq C_{i+2}$ and $d(b, x) \leq C_{i+2}$.*

In (ii) we may be applying Corollary 3.24 twice, and this is why the conclusion involves C_{i+2} .

The following lemma can be proved by a straightforward induction on n . We will apply it to dimension sequences.

Lemma 4.3 *Let $\mathbf{D}_i = (x_{i0}, x_{i1}, \dots, x_{ij_i})$ for $i = 0, 1, 2, \dots, N$ be a sequence of finite sequences of nonnegative integers. Assume*

- (a) $\mathbf{D}_0 = (x_{00}, x_{01})$ has length 2 and $x_{00}, x_{01} \leq n$,
- (b) for $i \geq 0$ the sequence \mathbf{D}_{i+1} is obtained from \mathbf{D}_i by inserting between some consecutive terms a nonnegative integer strictly smaller than each of the two terms.

Then $j_N \leq 2^n$.

For example, 33, 323, 31213, 301020103 is such a sequence with $n = 3$, $N = 3$ and $j_3 = 8$.

Proposition 4.4 *For every $C > 0$ there is $C' = C'(C, \Sigma)$ such that the following holds. Let \mathcal{C}_j be an excellent sequence of cell structures with all cells (ie their vertex cycles) at distance $\leq K$ from $*$. Suppose A and B are two cells in \mathcal{C}_j . If $a \in \text{int}(A)$ and $b \in \text{int}(B)$ and $d(a, b) \leq C$ then either*

- $d(*, a), d(*, b) \leq K + C'$, or
- there is a curve c contained in a cell of \mathcal{C}_j which is contained in a face of each of A and B such that $d(a, c), d(b, c) \leq C'$.

Proof First assume that a, b is a C_0 -good sequence. We set $C_0 = C$ and define C_i inductively as above. Run the process starting with a, b . There are now two possibilities.

Case 1 (the process produces a C_j -good sequence $a = a_0, a_1, \dots, a_N$) From Lemma 4.3 we see that $N \leq 2^n$, where $n = \dim \mathcal{ML}$. Thus there were at most $2^n - 1$ insertions and this implies that $d(a_i, a_{i+1}) \leq C_{2(2^n-1)}$ for any two consecutive curves a_i and a_{i+1} .

The sequence of C_j -carriers $\text{Carr}_j(a_i)$ for $i = 0, 1, \dots$ either increases or decreases (or stays the same) at every step. We now modify the sequence, by “pushing the peaks down” so that an initial part of the sequence of carriers is nonincreasing, and the rest is nondecreasing. Let a_i, a_{i+1}, \dots, a_k be a subsequence of consecutive curves such that

$$\text{Carr}_j(a_i) \subsetneq \text{Carr}_j(a_{i+1}) = \text{Carr}_j(a_{i+2}) = \dots = \text{Carr}_j(a_{k-1}) \supsetneq \text{Carr}_j(a_k).$$

First we pass to the length 3 subsequence a_i, a_{i+1}, a_k . We have $d(a_{i+1}, a_k) \leq 2^n C_{2(2^n-1)}$ (we are happy with very crude estimates), so applying Corollary 3.24(I) we find a curve x with $\text{Supp}_j(x) \subset \text{Carr}_j(a_i) \cap \text{Carr}_j(a_k)$ and with $d(a_i, x), d(a_k, x) \leq C_{2(2^n-1)+2}$. (If conclusion (1) occurs, see Case 2.) Continuing in this way produces the desired sequence. The number of steps that consist of pushing the peaks is bounded, eg by $n2^n$, so at the end the distance between any two consecutive curves is bounded by $C_{2(2^n-1)+2n2^n}$. Finally, pass to a length 3 sequence a, c, b where c has the minimal carrier, and set $C' = 2^n C_{2(2^n-1)+2n2^n}$.

Case 2 (at some stage in the process, when applying Corollary 3.24, conclusion (1) occurs) This applies also to the part of the procedure when we push the peaks down. Thus we have a sequence $a = a_0, a_1, \dots, a_N = b$ with $N \leq 2^n$, $d(a_i, a_{i+1}) \leq C_{2(2^n-1)+2n2^n}$, and for some i we have $d(*, a_i) \leq C_{2(2^n-1)+2n2^n+1}$. This implies by the triangle inequality that

$$d(*, a), d(*, b) \leq C_{2(2^n-1)+2n2^n+1} + 2^n C_{2(2^n-1)+2n2^n}$$

and we may take C' to be this bound.

Finally, consider the general case when a, b is not a C_0 -good sequence. Let $A = \text{Carr}_0(a)$ and $B = \text{Carr}_0(b)$. Thus $A = V(\alpha)$ and $B = V(\beta)$ for certain tracks α and β . Lemma 3.23 gives that either $d(*, a)$ and $d(*, b)$ are uniformly bounded, as functions of Σ and C , or there is a curve $c \in A \cap B$ within uniform distance — call it C_0 — from a and b . (Note here that since C_0 is a fixed cell structure, the intersection number between any two tracks defining it is uniformly bounded, so Lemma 3.23 applies uniformly.) Thus a, c, b is a C_0 -good sequence and the procedure above proves the statement. □

5 Capacity dimension of \mathcal{EL}

5.1 Capacity dimension

Let (Z, ρ) be a metric space. The notion of capacity dimension of Z was introduced by Buyalo in [8]. One of several possible equivalent definitions is the following; see [8, Proposition 3.2]. We also note that for bounded metric spaces, such as the boundary of a hyperbolic space with visual metric, capacity dimension agrees with the Assouad–Nagata dimension. See [19].

Definition 5.1 The capacity dimension of a metric space Z is the infimum of all integers m with the following property: there exists a constant $c > 0$ such that for all sufficiently small $s > 0$, Z has a cs -bounded covering with s -multiplicity at most $m + 1$.

The covering \mathcal{L} is cs -bounded if all elements have diameter $< cs$ and the s -multiplicity of \mathcal{L} is $\leq m + 1$ if every $z \in Z$ is at distance $< s$ from at most $m + 1$ elements of \mathcal{L} .

We will produce covers that resemble cell structures and whose thickenings resemble handle decompositions. It is more convenient here to index the handles starting with 1, rather than with 0. We will use following form of the definition of capacity dimension.

Proposition 5.2 Suppose that there is a constant $c > 10$ such that for all sufficiently small $s > 0$ there is a cover \mathcal{K} of Z with the following properties:

- The collection \mathcal{K} is the disjoint union of subcollections $\mathcal{K}_1, \mathcal{K}_2, \dots, \mathcal{K}_{m+1}$.
- The diameter of any set in \mathcal{K} is $\leq s$.
- If $A, B \in \mathcal{K}_i$ are distinct elements in the same subcollection and if $a \in A, b \in B$ and $\rho(a, b) < s/c^{3i-1}$, then there is some $e \in E \in \mathcal{K}_k$ for some $k < i$ such that $\rho(a, e) < s/c^{3i-2}$ and $\rho(b, e) < s/c^{3i-2}$.

Then the capacity dimension of Z is at most m .

Proof Inductively on i , for each $K \in \mathcal{K}_i$ define the associated “handle”

$$H(K) = N_{s/c^{3i}}(K) - \bigcup_{\substack{K' \in \mathcal{K}_k \\ k < i}} H(K').$$

It is clear that the collection of all handles forms a cover of Z and that the diameter of each element is bounded by $s + 2s/c^3$. We will argue that the s/c^{3m+4} -multiplicity of the cover is $\leq m + 1$. Suppose $z \in Z$ is at distance $< s/c^{3m+4}$ from $m + 2$ handles. Then two of the handles have the same index, say $H(A)$ and $H(B)$ with $A, B \in \mathcal{K}_i$. Thus we have $a_0 \in H(A)$ and $b_0 \in H(B)$ with $\rho(a_0, b_0) < 2s/c^{3m+4}$. Choose $a \in A$ and $b \in B$ with $\rho(a, a_0) < s/c^{3i}$ and $\rho(b, b_0) < s/c^{3i}$. Then $\rho(a, b) < 2s/c^{3m+4} + 2s/c^{3i} < s/c^{3i-1}$. By assumption there is $e \in K \in \mathcal{K}_k$ with $k < i$ and with $\rho(a, e) < s/c^{3i-2}$ and $\rho(b, e) < s/c^{3i-2}$. Thus $\rho(a_0, e) < s/c^{3i-2} + s/c^{3i} < s/c^{3i-3} \leq s/c^{3k}$, so $a_0 \in H(K)$ or it belongs to a lower-index handle, and similarly for b_0 . This is a contradiction since $H(A)$ and $H(B)$ are disjoint from lower-index handles. □




Figure 10: The handle decomposition associated to a triangulation. Cells shaped like the one pictured on the right will result in handle decompositions with bad Lebesgue number.

Remark 5.3 In Section 5.3 we will show that the cover of \mathcal{EL} induced by $P_\infty(\sigma)$ as $V(\sigma)$ ranges over cells in \mathcal{C}_j with σ large satisfy the assumptions of Proposition 5.2. For example, shapes like in Figure 10 are ruled out by Proposition 4.4.

We also have the following general fact. We thank Vera Tonic for pointing it out to us.

Proposition 5.4 [19] *Suppose Z is written as a finite union of (closed) subsets Z_i . If $\text{capdim } Z_i \leq n$ for all i , then $\text{capdim } Z \leq n$.*

5.2 Visual size and distance

Recall that a metric ρ on the boundary Z of a δ -hyperbolic space X is said to be *visual* if there is a basepoint $*$ $\in X$ and constants $a > 1$ and $c_1, c_2 > 0$ such that

$$c_1 a^{-(z|z')} \leq \rho(z, z') \leq c_2 a^{-(z|z')}$$

for all $z, z' \in Z$, where $(\cdot | \cdot)$ denotes the Gromov product

$$(x | x') = \frac{1}{2}(d(*, x) + d(*, x') - d(x, x'))$$

on X , extended naturally to Z . See [12, Chapitre 7] for more details and for the construction of visual metrics.

Also recall that, in a δ -hyperbolic space, the Gromov product $(a | b)$ is, to within a uniform bound that depends on δ , the distance between the basepoint $*$ and any geodesic $[a, b]$. The same is true when a and b are distinct points at infinity. We may also replace geodesics $[a, b]$ with quasigeodesics, but then the uniform bound depends also on quasigeodesic constants.

A δ -hyperbolic space X is *visual* [6] if for some (every) basepoint x_0 there exists $C > 0$ such that for every $x \in X$ there is a (C, C) -quasigeodesic ray in X based at x_0

and passing through x . Equivalently, X is the coarse convex hull of the boundary ∂X . Any δ -hyperbolic space whose isometry group acts coboundedly and that contains a biinfinite quasigeodesic line is visual. Thus a curve complex is visual.

Theorem 5.5 [8] *Let X be a visual δ -hyperbolic metric space and Z its Gromov boundary endowed with a visual metric. Then*

$$\text{asdim}(X) \leq \text{capdim}(Z) + 1.$$

Lemma 5.6 *Assume that σ is a large track obtained from τ by a sequence of splittings. Using $B(\tau)$ as a basepoint, the visual diameter of $P_\infty(\sigma)$ is $a^{-d(*, B(\sigma))}$, to within a bounded factor.*

Proof This follows from Lemmas 3.12 and 3.13. □

Proposition 5.7 *Let \mathcal{C}_j be an excellent sequence of cell structures obtained by splitting tracks. For all sufficiently large constants c (depending only on Σ), the following holds for all sufficiently small $s > 0$. Suppose for a certain j the visual diameter of each $P_\infty(\sigma)$ is $> s$, where σ ranges over all maximal train tracks such that $V(\sigma) \in \mathcal{C}_j$. Then:*

- *The visual diameter of each $P_\infty(\sigma)$ is $> s/c$ for every large track σ determining a cell $V(\sigma)$ in \mathcal{C}_j .*
- *Suppose cells $V(\sigma)$ and $V(\sigma')$ in \mathcal{C}_j are distinct with both σ and σ' large. Suppose $a \in P_\infty(\sigma)$, $b \in P_\infty(\sigma')$ and $\rho(a, b) < s/c^2$. Then there is a cell $V(\mu) \subset V(\sigma) \cap V(\sigma')$ and there is a $e \in P_\infty(\mu)$ such that $\rho(a, e) < s/c$ and $\rho(b, e) < s/c$.*

Proof The first bullet follows from Lemma 5.6 and Proposition 3.21.

To prove the second bullet we use Proposition 4.4. Consider a quasigeodesic ray from $*$ to a that passes through $B(\sigma)$ and between $B(\sigma)$ and a stays in $S(\sigma)$. By Lemmas 3.11 and 3.12 we may assume that these rays are uniformly quasigeodesic. Likewise, construct such a ray from $*$ to b . Also choose a uniform quasigeodesic from a to b . We now have a uniformly thin triangle with two vertices at infinity. Choose a' on the first ray and b' on the second ray, just past the thick part viewed from $*$. Thus $d(a', b')$ is uniformly bounded, say by C . Also note that $d(*, a')$ and $d(*, b')$ are definite amounts larger than $d(*, B(\sigma))$ and $d(*, B(\sigma'))$ by the assumption that $\rho(a, b) < s/c$ with c sufficiently large. In particular, $a' \in S(\sigma)$ and $b' \in S(\sigma')$. For this C , Proposition 4.4

provides a constant C' . Now the first bullet in the conclusion of Proposition 4.4 cannot hold if c is sufficiently large. Therefore, there is some $e' \in P(\mu) \in \mathcal{C}_j$ with $d(a', e'), d(b', e') \leq C'$ and with $\dim V(\mu) < \dim V(\sigma) = \dim V(\sigma')$. Again using Lemmas 3.11 and 3.12, construct a uniform quasigeodesic ray from $*$ through e' to some $e \in P_\infty(\mu)$. This e satisfies the conclusions. \square

5.3 The cover

It is well known that \mathcal{PML} can be covered by finitely many sets of the form $P(\tau)$ for τ a large train track (for a concrete cover see [24]). Thus \mathcal{EL} is finitely covered by sets of the form $P_\infty(\tau)$. In view of Proposition 5.4 we need to find an upper bound to $\text{capdim } P_\infty(\tau)$.

Here we fix a large birecurrent train track τ and describe a cover of $Z = P_\infty(\tau)$ that satisfies Proposition 5.2 for a certain m depending on Σ and for small $s > 0$.

The dimensions of cones $V(\sigma)$ for large train tracks $\sigma \subset \Sigma$ belong to a certain interval $[A, A + K]$ that depends on Σ . We put $m = K$. We also fix a large constant c .

Now fix a small $s > 0$ and start with the standard cell structure on $V(\tau)$. This is \mathcal{C}_0 . Now suppose \mathcal{C}_j has been constructed and the visual size of each $P_\infty(\sigma)$ for a top-dimensional cell $V(\sigma) \in \mathcal{C}_j$ is $> s/c$. Enumerate all top-dimensional cells $V(\sigma) \in \mathcal{C}_j$ such that the visual size of $P_\infty(\sigma)$ is $\geq s$ and also enumerate all large branches in the corresponding tracks σ . Then construct $\mathcal{C}_{j+1}, \dots, \mathcal{C}_k$ by splitting along these branches, in any order. We call this collection of splits a *multisplit*.

This gives an infinite excellent sequence. Note that once some $P_\infty(\sigma)$ (with $V(\sigma)$ maximal) reaches visual size $< s$ at the end of a multisplit, it never gets subdivided again (see Lemma 3.9). Coarsely, reaching a certain visual size is equivalent to $B(\sigma)$ reaching a certain distance from the basepoint $*$ (see Lemma 5.6).

Lemma 5.8 *Let λ be a filling lamination, and for every j let $E_j = V(\sigma_j)$ be the cell in \mathcal{C}_j that contains λ in its interior. Then the sequence E_j eventually stabilizes.*

Proof We argue by contradiction. From Proposition 3.21 we have that $\sigma_{j+1} \xrightarrow{\text{ss}} \sigma_j$. Let a_j be a vertex cycle of σ_j . Then we may assume, perhaps for a subsequence, that $a_j \rightarrow \lambda'$ in \mathcal{PML} and the lamination λ' is necessarily carried by all σ_j . By Lemma A.4, for large j , σ_j will fully carry λ , so by Proposition 3.22 we have $[\lambda'] = [\lambda]$. By an argument of Kobayashi (see [21, page 124]) the sequence a_j goes to infinity in

the curve complex, so the visual size of E_j goes to 0 by Lemma 5.6. But in the construction of \mathcal{C}_j the visual diameter of all top-dimensional cells is bounded below and by Proposition 3.21 this bounds below the visual diameter of all cells, giving a contradiction. \square

We let the cover \mathcal{K} consist of the sets of the form $P_\infty(\sigma)$ such that $V(\sigma)$ is a stable cell. We partition the sets in \mathcal{K} according to the dimension of the cell.

Theorem 5.9 *We have*

$$\text{capdim}(EL) \leq K(\Sigma),$$

where $K = K(\Sigma)$ is the smallest integer such that every recurrent, transversely recurrent, large track σ on Σ has $\dim P(\sigma) \in [A, A + K]$ for some $A = A(\Sigma)$.

Proof We only need to argue that the cover \mathcal{K} satisfies the conditions of Proposition 5.2. This is clear from the construction and Propositions 3.21 and 4.4 applied to \mathcal{C}_j for large j . \square

Corollary 5.10 $\text{asdim}(\mathcal{C}(\Sigma)) \leq K(\Sigma) + 1.$

Example 5.11 One can see easily what happens in the case of the punctured torus. Then $\mathcal{FPM}\mathcal{L} = \mathcal{EL}$ is homeomorphic to the set of irrational numbers, or equivalently to \mathbb{Z}^∞ . The visual metric is complete, and the cover \mathcal{K} constructed above will be infinite and will consist of pairwise disjoint sets, all of the same index, and all of comparable sizes. For example, consider a standard track that supports laminations whose slope is in the interval $[1, \infty]$. Splitting produces two tracks, one carrying laminations in the interval $[1, 2]$ and the other in the interval $[2, \infty]$. We can take the curve with slope ∞ as the basepoint and agree to stop subdividing when the distance from ∞ to $B(\sigma)$ is > 0 , ie when ∞ is no longer carried by σ . Thus we stop splitting the track carrying $[1, 2]$ and we split the other track. We get tracks carrying $[2, 3]$ and $[3, \infty]$. Continuing in this way we get an infinite cover $P_\infty(\sigma_n)$ of \mathcal{EL} where σ_n carries ending laminations with slope in $[n, n + 1]$.

Remark 5.12 There are two other closely related notions to asymptotic dimension. In the *linearly controlled* asymptotic dimension, or the *asymptotic Assouad–Nagata* dimension $\ell\text{-asdim}$, one insists on the linear control on the size of the cover. Also,

say that $\text{ecodim}(X) = n$ if X quasiisometrically embeds in a product of n trees and n is smallest possible. Then there is a chain of inequalities

$$\text{asdim} \leq \ell\text{-asdim} \leq \text{ecodim}$$

for any metric space, and Buyalo shows in [9] that, when X is δ -hyperbolic,

$$\text{ecodim}(X) \leq \text{capdim } \partial X + 1.$$

See also the discussion in [20]. Therefore our arguments also give the same bound on ℓ -asdim and ecodim for $\mathcal{C}(\Sigma)$. Previously, Hume has shown $\text{ecodim}(\mathcal{C}(\Sigma)) < \infty$ [17].

Appendix Train tracks and full dimension paths

A *splitting path* is a legal embedded path in a thickened train track that begins and ends at a cusp. See Figure 11.

If τ is a recurrent train track and θ is the central split along the splitting path then θ will have one or two connected components and a total of three less branches and two less switches than τ . By Lemma 3.1 we then have:

Lemma A.1 *Either $\dim V(\theta) < \dim V(\tau)$ or $\dim V(\theta) = \dim V(\tau)$. If $\dim V(\theta) = \dim V(\tau)$ then one of the following holds:*

- (1) τ is nonorientable and θ is connected and orientable,
- (2) τ is orientable and θ has two components (both necessarily orientable), or
- (3) τ is nonorientable and θ has one orientable component and one nonorientable component.




Figure 11: The dashed line is the splitting path. It must start and end at a cusp in the thickened track.

If $\dim V(\theta) = \dim V(\tau)$, we say that the splitting path is a *full-dimension splitting path* or *fd-path*. While a splitting path will be embedded in the thickened train track, in the actual train track it may cross a single branch multiple times. However, an fd-path can cross any branch at most twice and this strong restriction implies that there is a uniform bound on the number of fd-paths in a given track.

Lemma A.2 *An fd-path of types (1) or (2) crosses each branch of τ at most once and an fd-path of type (3) crosses each branch at most twice.*

Proof Orient the splitting path. If the path crosses a branch more than once, we examine two consecutive strands of the path in the branch, as seen crossing the branch transversally.

- If the two strands have the same orientation then θ will be connected and we must be in case (1). Then τ will be nonorientable and θ will be orientable, so we can choose an orientation for θ . On opposite sides of the splitting path the orientation of θ cannot agree or τ would be orientable. Using the orientation of the splitting path (and the surface) we can assume that to the right of the splitting path the orientation of θ and the splitting path agree while to the left they are opposite of each other. However, this is not possible if there are two consecutive strands in the same branch with the same orientation. Therefore we can never have consecutive branches with the same orientation.
- If consecutive branches (seen transversally) in the same branch have opposite orientation then the component of θ between the two strands will be nonorientable and we must be in case (3). If the splitting path crosses the same branch three or more times then orientation of consecutive branches will always be opposite so both components of θ will be nonorientable, a contradiction. \square

Lemma A.3 *Let τ_1 and τ_2 be recurrent train tracks with $\tau_2 \xrightarrow{ss} \tau_1$ and $\dim V(\tau_1) = \dim V(\tau_2)$. Let θ_2 be the central split of τ_2 along a splitting path. Then there is a splitting path in τ_1 with central split θ_1 such that $V(\theta_2) = V(\theta_1) \cap V(\tau_2)$.*

Proof As usual, via induction we can reduce this to the case when $\tau_2 \xrightarrow{ss} \tau_1$ is a single move. In τ_2 the splitting path starts and ends at large half-branches. If these large half-branches and their adjacent half-branches are not active in the move then the composition of the splitting path with the carrying map is a splitting path. If not, the carrying map is a “fold” along the switch and, in τ_1 , we extend the path along the fold.

If the move is a right or left split along a large branch b in τ_1 then $\theta_1 = \theta_2$ if the path crosses b in τ_2 . If not, that θ_2 is a split of θ_1 along b . If the move is a central split then θ_2 is a central split of θ_1 . \square

One consequence of the existence of fd-paths is that there are large train tracks that do not fully carry any lamination. We will show that this is the only obstruction.

Lemma A.4 *Let τ be a recurrent train track and let λ be a lamination in the interior of $V(\tau)$. Then there exists a recurrent train track σ with $\sigma \xrightarrow{s} \tau$ and λ is fully carried by σ .*

Proof By [26, Proposition 8.9.2] there exists a recurrent (in fact birecurrent) train track τ' that fully carries λ . Note that while fully carrying is not discussed in this proposition, one sees that if the ϵ in the construction is chosen to be sufficiently small then the track will be fully carrying. Then, as in Proposition 3.22, we use [24, Theorem 2.3.1] to find a train track σ that carries λ with $\sigma \xrightarrow{s} \tau'$ and $\sigma \xrightarrow{s} \tau$. As τ' fully carries λ , so must σ . \square

Observe that if $\sigma \xrightarrow{ss} \tau$ and $\dim V(\sigma) < \dim V(\tau)$ then there will be a hyperplane P defined by equations that have rational coefficients and such that P has positive codimension and $V(\sigma) \subset P \cap V(\tau)$.

Proposition A.5 *Let τ be a recurrent large train track and assume that $\lambda \in V(\tau)$ is not contained in a rational hyperplane of positive codimension. Then either τ contains an fd-path with central split θ and $\lambda \in V(\theta)$ or λ is an ending lamination.*

Proof By Lemma A.4 we can find a recurrent train σ such that $\sigma \xrightarrow{s} \tau$ and σ fully carries λ . If τ fully carries σ then it also fully carries λ and λ must be an ending lamination. If not, a central split must occur in the sequence $\sigma \xrightarrow{s} \tau$. Let τ_1 be the track in the sequence that occurs just before the first central split and let b be the large branch where the central split occurs. Then $\dim V(\sigma) = \dim V(\tau)$ for, if not, λ would be contained in a rational hyperplane of positive codimension. Therefore the large branch b is an fd-path of length one. Let θ_1 the central split of τ_1 along the fd-path b . By Lemma A.3 there exists an fd-path in τ with central split θ such that $V(\theta_1) = V(\tau) \cap V(\theta)$, so $\lambda \in V(\theta)$. \square

To summarize, in the presence of fd-paths it is generally not true that for a large track σ generic points of $V(\sigma)$ (those in the complement of rational hyperplanes) represent

ending laminations, but this will be true in the complement of a finite collection of subcells of $V(\sigma)$. The subcells could cover $V(\sigma)$, but for example if $V(\sigma)$ contains one ending lamination, it contains infinitely many.

Corollary A.6 *Let τ be a birecurrent train track and $a \in S(\tau)$ a simple closed curve. There exists a $C = C(\Sigma)$ such that either $d(B(\tau), a) \leq C$ or there exists a sequence $\lambda_i \in P_\infty(\tau)$ such that a is contained in the Hausdorff limit of the $[\lambda_i]$.*

Proof By Lemma A.2, τ contains finitely many fd-paths. Assume there are $k \geq 0$ such paths. We begin by splitting on each of these paths to obtain k new tracks, which we label $\theta_1, \dots, \theta_k$. If a is in the complement of $\bigcup V(\theta_i)$ then the corollary follows from Proposition A.5 applied to a sequence of laminations in $V(\tau)$ converging to a and not contained in proper rational planes. If not, $a \in S(\theta_i)$ for some i . If θ_i is small then $d(B(\theta_i), a) \leq 2$ and $d(B(\theta_i), B(\tau))$ is uniformly bounded since an fd-path is at worst two-to-one by Lemma A.2. Therefore $d(B(\tau), a)$ is uniformly bounded.

If θ_i is large then it is connected and, by Lemma A.1, orientable. As in the previous paragraph we split along all fd-paths to get a collection of tracks $\theta_1^i, \dots, \theta_{k_i}^i$. Since θ_i is orientable, Lemma A.1 implies that the θ_j^i are disconnected and hence small. If $a \in \theta_j^i$ for some j then a is uniformly close to $B(\theta_j^i)$, and therefore to $B(\tau)$, as above. If not, we again apply Proposition A.5. \square

References

- [1] **I Agol**, *Ideal triangulations of pseudo-Anosov mapping tori*, from “Topology and geometry in dimension three” (W Li, L Bartolini, J Johnson, F Luo, R Myers, JH Rubinstein, editors), Contemp. Math. 560, Amer. Math. Soc., Providence, RI (2011) 1–17 MR
- [2] **T Aougab**, *Quadratic bounds on the quasiconvexity of nested train track sequences*, Topology Proc. 44 (2014) 365–388 MR
- [3] **J Behrstock**, **M F Hagen**, **A Sisto**, *Asymptotic dimension and small-cancellation for hierarchically hyperbolic spaces and groups*, Proc. Lond. Math. Soc. 114 (2017) 890–926 MR
- [4] **G C Bell**, **K Fujiwara**, *The asymptotic dimension of a curve graph is finite*, J. Lond. Math. Soc. 77 (2008) 33–50 MR
- [5] **M Bestvina**, **K Bromberg**, **K Fujiwara**, *Constructing group actions on quasi-trees and applications to mapping class groups*, Publ. Math. Inst. Hautes Études Sci. 122 (2015) 1–64 MR

- [6] **M Bonk, O Schramm**, *Embeddings of Gromov hyperbolic spaces*, Geom. Funct. Anal. 10 (2000) 266–306 MR
- [7] **B H Bowditch**, *Tight geodesics in the curve complex*, Invent. Math. 171 (2008) 281–300 MR
- [8] **S V Buyalo**, *Asymptotic dimension of a hyperbolic space and the capacity dimension of its boundary at infinity*, Algebra i Analiz 17 (2005) 70–95 MR In Russian; translated in St. Petersburg Math. J. 17 (2006) 267–283
- [9] **S V Buyalo**, *Capacity dimension and embedding of hyperbolic spaces into the product of trees*, Algebra i Analiz 17 (2005) 42–58 MR In Russian; translated in St. Petersburg Math. J. 17 (2006) 581–591
- [10] **S V Buyalo, N D Lebedeva**, *Dimensions of locally and asymptotically self-similar spaces*, Algebra i Analiz 19 (2007) 60–92 MR In Russian; translated in St. Petersburg Math. J. 19 (2008) 45–65
- [11] **D Gabai**, *On the topology of ending lamination space*, Geom. Topol. 18 (2014) 2683–2745 MR
- [12] **E Ghys, P de la Harpe**, *Sur les groupes hyperboliques d’après Mikhael Gromov*, Progr. Math. 83, Birkhäuser, Boston, MA (1990) MR
- [13] **B Grünbaum**, *Convex polytopes*, 2nd edition, Graduate Texts in Mathematics 221, Springer (2003) MR
- [14] **U Hamenstädt**, *Geometry of the mapping class groups, I: Boundary amenability*, Invent. Math. 175 (2009) 545–609 MR
- [15] **J L Harer**, *The virtual cohomological dimension of the mapping class group of an orientable surface*, Invent. Math. 84 (1986) 157–176 MR
- [16] **S Hensel, P Przytycki**, *The ending lamination space of the five-punctured sphere is the Nöbeling curve*, J. Lond. Math. Soc. 84 (2011) 103–119 MR
- [17] **D Hume**, *Embedding mapping class groups into a finite product of trees*, Groups Geom. Dyn. 11 (2017) 613–647 MR
- [18] **E Klarreich**, *The boundary at infinity of the curve complex and the relative Teichmüller space*, preprint (2018) arXiv
- [19] **U Lang, T Schlichenmaier**, *Nagata dimension, quasisymmetric embeddings, and Lipschitz extensions*, Int. Math. Res. Not. 2005 (2005) 3625–3655 MR
- [20] **J M Mackay, A Sisto**, *Embedding relatively hyperbolic groups in products of trees*, Algebr. Geom. Topol. 13 (2013) 2261–2282 MR
- [21] **H A Masur, Y N Minsky**, *Geometry of the complex of curves, I: Hyperbolicity*, Invent. Math. 138 (1999) 103–149 MR
- [22] **H A Masur, Y N Minsky**, *Geometry of the complex of curves, II: Hierarchical structure*, Geom. Funct. Anal. 10 (2000) 902–974 MR

- [23] **H A Masur, Y N Minsky**, *Quasiconvexity in the curve complex*, from “In the tradition of Ahlfors and Bers, III” (W Abikoff, A Haas, editors), *Contemp. Math.* 355, Amer. Math. Soc., Providence, RI (2004) 309–320 MR
- [24] **R C Penner, J L Harer**, *Combinatorics of train tracks*, *Annals of Mathematics Studies* 125, Princeton Univ. Press (1992) MR
- [25] **M Sapir**, *A Higman embedding preserving asphericity*, *J. Amer. Math. Soc.* 27 (2014) 1–42 MR
- [26] **W P Thurston**, *The geometry and topology of three-manifolds*, lecture notes, Princeton University (1979) Available at <http://msri.org/publications/books/gt3m>
- [27] **R C H Webb**, *Combinatorics of tight geodesics and stable lengths*, *Trans. Amer. Math. Soc.* 367 (2015) 7323–7342 MR
- [28] **G M Ziegler**, *Lectures on polytopes*, *Graduate Texts in Mathematics* 152, Springer (1995) MR

*Department of Mathematics, University of Utah
Salt Lake City, UT, United States*

*Department of Mathematics, University of Utah
Salt Lake City, UT, United States*

bestvina@math.utah.edu, ken.w.bromberg@gmail.com

Proposed: Bruce Kleiner

Received: 15 September 2015

Seconded: Jean-Pierre Otal, Martin Bridson

Revised: 19 July 2018