

Train track expansions of measured foliations

Lee Mosher*

December 28, 2003

Contents

1	Introduction	6
1.1	Continued fraction expansions of real numbers	6
1.2	A dictionary	7
1.3	Continued fractions expansions and the torus	10
1.4	Measured foliations	12
1.5	Train tracks and train track expansions (§§3–4)	14
1.6	The expansion convergence theorem (§5)	15
1.7	The rational killing criterion (§6)	16
1.8	Expansions of arational measured foliations (§6 continued)	19
1.9	Stable equivalence (§7)	20
1.10	General splitting sequences (§8)	26
1.11	Construction of pseudo-Anosov mapping classes (§9)	27
1.12	Classification of pseudo-Anosov conjugacy classes (§10 in progress)	27
1.13	Enumeration of pseudo-Anosov conjugacy classes (§11 in progress)	29
1.14	The conjugacy problem in the mapping class group (§12 in progress)	29
1.15	Prerequisites	32
1.16	What’s not here.	32
1.17	Revision notes	33
1.18	Acknowledgements.	35
2	Measured foliations	36
2.1	Surfaces of finite type	36
2.2	Teichmüller space \mathcal{T}	36
2.3	The mapping class group	37
2.4	Quadratic differentials and their horizontal and vertical foliations	37

*Supported by a National Science Foundation grant.

2.5	Measured foliations	38
2.6	Partial measured foliations	42
2.7	Measured foliation space \mathcal{MF} , and its projectivization \mathcal{PMF}	46
2.8	Canonical models for partial measured foliations	47
2.9	Components of a measured foliation	48
2.10	Intersection number and joint filling	51
2.11	Thurston's classification of mapping classes	53
3	Train tracks	55
3.1	Pretracks	55
3.2	Train tracks	58
3.3	Train paths in the universal cover	64
3.4	Ties of a train track	72
3.5	Carrying maps between train tracks	74
3.6	Measured foliations carried on train tracks	77
3.7	Invariant weights on train tracks	84
3.8	Naturality of carrying	87
3.9	Filling properties of train tracks	88
3.10	Diagonal extension of train tracks	88
3.11	Recurrence	92
3.12	Slide moves and comb equivalence	96
3.13	Splitting of train tracks	101
3.14	Factoring a carrying map into splittings	110
3.15	Combinatorial types of train tracks.	113
4	Train track expansions of measured foliations	116
4.1	Definition and existence of train track expansions	116
4.2	Expansions of annular measured foliations	117
4.3	A unified construction of train track expansions	118
4.4	Universality of the construction.	125
5	Convergence of train track expansions	128
5.1	Convergence to an arational measured foliation	128
5.2	Proof of the expansion convergence theorem	129
5.3	Convergence to a partial arational measured foliation	134
6	Arational Expansion Theorem	138
6.1	Survival of subtracks	140
6.2	Survival of splitting cycles	144
6.3	Statement of the Arational Expansion Theorem	149

6.4	Application: Canonical train track expansions	152
6.5	Necessity of the iterated rational killing criterion	156
6.6	Sufficiency of the iterated rational killing criterion	159
6.7	The Stability Lemma.	162
6.8	Sufficiency of the iterated rational killing criterion: conclusion. . . .	165
7	Stable equivalence of expansion complexes	167
7.1	Stable equivalence of one cusp train track expansions	168
7.2	Expansion complexes of arational measured foliations	174
7.3	Proving the Stable Equivalence Theorem	182
7.4	Cells of the expansion complex	188
7.5	The face relation in the expansion complex.	197
7.6	Remarks on noncanonical expansion complexes.	207
7.7	Unmarked stable equivalence of expansion complexes	208
8	The structure of infinite splitting sequences	211
8.1	Twist splittings	212
8.2	Characterization of twist splitting sequences	217
8.3	Cutting one train track out of another	220
8.4	Proof of (1) \implies (2) in Theorem 8.2.1	226
8.5	General splitting sequences.	235
9	Constructing pseudo-Anosov homeomorphisms	238
9.1	Invariant train tracks	239
9.2	Automata for canonical splitting circuits (in progress).	248
9.3	Penner's recipe	253
9.4	One sink train track expansions	262
9.5	Remarks on the complexity of the Canonical Killing Criterion	269
10	Classifying pseudo-Anosov conjugacy classes	271
10.1	The torus	273
10.2	Pseudo-Anosov conjugacy using one cusp splitting circuits	278
10.3	Pseudo-Anosov conjugacy using circular expansion complexes	285
11	Enumeration of pseudo-Anosov conjugacy classes	294
11.1	The canonical killing criterion for one cusp splitting sequences	294
11.2	Combinatorial types of stable train tracks	317
11.3	Enumeration of one cusp splitting circuits	335
11.4	Example: the once-punctured torus (in progress)	343
11.5	Formal definition of the canonical digraph (in progress)	343

11.6	Proof of the Degeneracy Theorem (in progress)	343
11.7	Further examples (in progress)	343
12	The conjugacy problem in MCG (in progress)	344
12.1	Conjugacy invariants of finite order mapping classes, after Nielsen (in progress)	344
12.2	Conjugacy invariants of algebraically finite mapping classes, after Nielsen (in progress)	344
12.3	Conjugacy invariants of the Thurston reduction (in progress)	344
12.4	Complete conjugacy invariants (in progress)	344
12.5	Tailoring the Bestvina-Handel algorithm (in progress)	344
12.6	Two cusp expansion complexes.	344

List of Figures

1	Left and right splittings on a torus	12
2	Pinching a train track.	18
3	A k -pronged singularity	39
4	Unpinching a pinch point.	45
5	A nontaut measured foliation	50
6	Examples of train tracks	59
7	A nested union of train tracks in a cusp	70
8	Local models of a tie bundle over a train track	73
9	Train tracks and quadratic differentials	79
10	A dumbbell.	94
11	A slide.	97
12	Combing along a transition branch.	98
13	Elementary splittings	103
14	Wide splitting of a train track	107
15	Factoring a wide splitting	109
16	Events.	123
17	Shunts	124
18	Survival of C -splitting arcs	147
19	Induced carrying of subtracks	158
20	A one cusp splitting	169
21	A homotopic carrying injection	190
22	A cusp triangle	191
23	Equations associated to homotopic carrying injections.	193
24	Inequalities associated to homotopic carrying injections.	194

25	A collapsible tree of a train track	199
26	Twist splitting	213
27	Interactions of a surviving subtrack with a splitting	227
28	A C-splitting arc crossing a smooth nontwistable moat	234
29	The train track of a twist system	258
30	A fractional Dehn twist	259
31	Splitting a one sink train track	264
32	A stable train track	297
33	Sliding a type 1 stable train track	300
34	Type 2 moves	302
35	One cusp fold	304
36	A subtrack is killed	308
37	A source path	312
38	A transition path	315
39	Ends and prongs	321
40	The dual track of a cell division	322
41	The stable dual track of a stable cell division	324
42	A chord diagram in genus 2	328
43	A chord diagram in genus 1 with 3 punctures	329
44	Chord diagrams: genus 0 with 8 punctures	330
45	Some genus 0 chord diagrams and their dual train tracks.	331
46	Chord diagrams: genus 0 with 5 punctures.	331
47	Chord diagrams: genus 0 with 6 and 7 punctures.	333
48	Chord diagram: genus 1 with 1 puncture.	334
49	Chord diagram: genus 1 with 2 punctures.	334
50	Chord diagram: genus 1 with 2 punctures.	335
51	Chord diagram: genus 2 with 1 puncture.	336

1 Introduction

1.1 Continued fraction expansions of real numbers

Continued fraction expansions of real numbers encompass information about the action of the group $\mathrm{SL}(2, \mathbf{Z})$ on the extended number line $\mathbf{R}^* = \mathbf{R} \cup \infty$ by fractional linear transformations: $\begin{pmatrix} a & b \\ c & d \end{pmatrix}(r) = (ar + b)/(cr + d)$.

For example, the set of extended rationals $\mathbf{Q}^* = \mathbf{Q} \cup \infty$ and its complement the set of irrationals $\mathbf{I} = \mathbf{R} - \mathbf{Q}$ are each invariant under $\mathrm{SL}(2, \mathbf{Z})$, and membership of a real number r in either \mathbf{Q} or \mathbf{I} can be detected in terms of continued fractions expansions. Each finite continued fraction expansion

$$[n_0, n_1, n_2, \dots, n_k] = n_0 + \frac{1}{n_1 + \frac{1}{n_2 + \frac{1}{\dots + \frac{1}{n_k}}}}, \quad n_0 \in \mathbf{Z}, \quad n_1, \dots, n_k \in \mathbf{Z}_+$$

represents a rational number, and each rational number has exactly two such expansions, related by the formula

$$[n_0, n_1, n_2, \dots, n_k] = [n_0, n_1, n_2, \dots, n_k - 1, 1] \quad \text{if } n_k > 1$$

As a special case, ∞ has a finite continued fraction expansion $\infty = \frac{1}{0}$. Each infinite continued fraction expansion

$$[n_0, n_1, n_2, \dots] = n_0 + \frac{1}{n_1 + \frac{1}{n_2 + \dots}}, \quad n_0 \in \mathbf{Z}, \quad n_i \in \mathbf{Z}_+ \text{ for } i > 0$$

represents a unique real number in the sense that the sequence of finite continued fractions $[n_0], [n_0, n_1], [n_0, n_1, n_2], \dots$ necessarily converges, and the limit is irrational. Moreover, every irrational number has a unique such expansion. Thus, the rational–irrational dichotomy corresponds to the finite–infinite dichotomy of continued fraction expansions.

The set \mathbf{Q}^* forms a single orbit of the action of $\mathrm{SL}(2, \mathbf{Z})$ on \mathbf{R}^* . More generally, two irrational numbers $r = [m_0, m_1, m_2, \dots]$, $s = [n_0, n_1, n_2, \dots]$ are in the same orbit under the action of $\mathrm{SL}(2, \mathbf{Z})$ if and only if their continued fraction expansions are *stably equivalent*, meaning that they have the same infinite tail, that is, there exists $k, l \geq 0$ such that $m_{k+i} = n_{l+i}$ for all $i \geq 0$.

Quadratic irrationalities $(a + \sqrt{b})/c$, $(a, b, c \in \mathbf{Z})$, are precisely the fixed points of those elements $A \in \mathrm{SL}(2, \mathbf{Z})$ which are hyperbolic, meaning that $|\mathrm{Tr}(A)| > 2$. Moreover, quadratic irrationalities are characterized as having eventually periodic continued fraction expansion.

Every hyperbolic element $A \in \mathrm{SL}(2, \mathbf{Z})$ has exactly two fixed points in \mathbf{R}^* , an attracting fixed point r_A^+ and a repelling fixed point r_A^- . Moreover, A is completely

classified up to conjugacy by $\text{Tr}(A)$ and by the primitive period loop of the continued fraction expansion of r_A^+ . This gives a very satisfactory solution to the conjugacy problem in $\text{SL}(2, \mathbf{Z})$, by giving complete and easily computable conjugacy invariants, once one incorporates also the much simpler conjugacy classification of periodic elements ($|\text{Tr}(A)| \leq 1$) and parabolic elements ($|\text{Tr}(A)| = 2$).

1.2 A dictionary

Given a surface S of finite type, the action of the mapping class group $\mathcal{MCG}(S)$ on the space of projective measured laminations $\mathcal{PMF}(S)$ is a generalization of the action of $\text{SL}(2, \mathbf{Z})$ on \mathbf{R}^* : in the case of a torus T , the slope of a measured foliation on T defines a homeomorphism $\mathcal{PMF}(T) \approx \mathbf{R}^*$, and there is an isomorphism $\mathcal{MCG}(T) \approx \text{SL}(2, \mathbf{Z})$ which conjugates the action of $\mathcal{MCG}(T)$ on $\mathcal{PMF}(T)$ to the action of $\text{SL}(2, \mathbf{Z})$ on \mathbf{R}^* . This conjugacy leads to a precise dictionary relating theorems about continued fractions to theorems about foliations of the torus.

The goal of this work is to extend this dictionary so as to translate between theorems about continued fraction expansions of real numbers and theorems about train track expansions of measured foliations on an arbitrary surface of finite type. Our focus is on finding analogues of the basic “topological” theorems of continued fractions: convergence, irrationality, stable equivalence, and periodicity theorems.

The roots of this dictionary are found in Thurston’s classification of the elements of $\mathcal{MCG}(S)$ into finite order, reducible, and pseudo-Anosov ([Thu88], [FLP⁺79], [CB88]), generalizing the classification of elements of $\text{SL}(2, \mathbf{Z})$ into finite order, parabolic, and hyperbolic.

Another precursor of this dictionary occurs in the work of Keane [Kea75] [Kea77], Veech [Vee78], and Rauzy [Rau79] on interval exchange maps. An interval exchange map is determined by an interval $[a, b]$ subdivided at finitely many points $a = a_0 < a_1 < \dots < a_n = b$, and a permutation σ of $\{1, \dots, n\}$, which is then used to define a map $f: [a, b] \rightarrow [a, b]$ by reordering the subintervals $[a_0, a_1), [a_1, a_2), \dots, [a_{n-1}, a_n)$ according to the permutation σ . By choosing a certain subinterval and taking the first return map, one obtains an induced interval exchange map; iterating this induction process produces an expansion of an interval exchange map which generalizes continued fraction expansions. Such expansions were used, for example, to investigate ergodic properties of interval exchange maps. The works cited above contain, for example, analogues of the convergence of infinite continued fractions in the setting of expansions of interval exchange maps; see for example Lemma 1.5 of [Vee78].

Train tracks themselves were introduced by Thurston in the late 1970’s (see [Thu87]) as a tool for approximating measured geodesic laminations, and splitting of train tracks was used as a refinement of the approximation. In [Ker85] Kerckhoff

explained the connection between interval exchange maps and measured foliations, and how induction of interval exchange maps is related to splitting of train tracks. In particular this paper makes explicit the concept of a train track expansion of a measured foliation, and how this generalizes expansions of interval exchange maps and continued fraction expansions. In particular, Kerckhoff's paper contains the convergence result for train track expansions that is given here as Theorem 5.1.1; see [Ker85] p. 268.

The paper of Penner and Papadopolous [PP87] and Penner's book [Pen92] contain some entries in a general dictionary between continued fraction expansions and train track expansions. Convergence of train track expansions is considered in Proposition 3.3.2 of [Pen92] as well as in Theorem 3.1 of [PP87]; a general version of this result is given here in Theorem 5.1.1. Theorem 3.3.1 of [Pen92] is a special case of stable equivalence; the general stable equivalence theorem is given here in Section 7.

Starting from Thurston's classification of mapping classes, a correspondence has developed between theorems about $SL(2, \mathbf{Z})$ and $\mathcal{MCG}(S)$. The hyperbolic elements of $SL(2, \mathbf{Z})$ can be constructed in a completely explicit manner: they are the matrices with $|\text{trace}| > 2$. Generalizing this there are various constructions for pseudo-Anosov homeomorphisms [Mos86], [Fat87], [Pen88], [AF91], [Bau92], [Fat92], as well as algorithms for deciding whether a mapping class is pseudo-Anosov, reducible, or finite order [BH95], [HTC96]. Mosher virtually classified pseudo-Anosov homeomorphisms up to conjugacy in terms of periodicity data of a certain kind of expansion [Mos86], and gave an algorithm to compute the invariants [Mos83], thereby virtually solving the pseudo-Anosov part of the conjugacy problem; these results are contained here in Section 10. Papadopolous and Penner [PP87] characterized when an arational measured foliation is a fixed point of a pseudo-Anosov mapping class, in terms of existence of certain periodicity data of train track expansions. Takarajima [Tak94] characterized which periodic train track sequences are expansions of pseudo-Anosov fixed points.

The unfinished monograph [Mos93], an outgrowth of [Mos83], gives a theory of expansions of measured foliations, containing in particular theorems about continued fraction expansions mentioned above: convergence; characterization of arationality; stable equivalence; and classification of pseudo-Anosov mapping classes using periodicity data. This monograph had only a very narrow distribution, due to its long and unpolished state. Expansions of measured foliations as defined in [Mos93] involve train tracks only peripherally; instead expansions are defined in terms of so-called "CDPs" which are certain cell decompositions of the surface, decorated with extra combinatorial data. The CDP machinery was dictated by the applications, with an eye towards computational efficiency. For example, with these methods one obtains virtually complete conjugacy invariants of pseudo-Anosov

mapping classes — indeed, they are complete invariants for the relation of “almost conjugacy”, where two mapping classes Φ, Ψ are almost conjugate if there exists a positive integer m such that Φ^m, Ψ^m are conjugate (see Section ??). Moreover, these invariants are efficiently computable using pencil and paper, particularly on a once-punctured surface of low genus. The author’s experience with such computations led to the construction of an automatic structure for the mapping class group [Mos95], [Mos96]. The downside of the CDP machinery in [Mos93] is that effectiveness was achieved in these papers at the expense of clarity, comprehensiveness, and more general applicability. On the other hand, CDPs do occur naturally in the understanding of “one cusp” train track expansions, as explained in Section 11.2.

This monograph establishes a general theory of train track expansions of measured foliations on a surface of finite type S , in parallel with continued fraction expansions of real numbers. In effect, we have redone [Mos93] from the ground up, reformulating all the theorems in terms of train tracks, hopefully to maximize clarity, comprehensiveness and applicability without sacrificing too much effectiveness. Each of our main theorems specializes to the torus, where it directly translates into a theorem about continued fractions. For example, a well known folk theorem, proved here in Lemma 4.2.1, says that a measured foliation has a finite train track expansion if and only if all of its leaves are compact; this generalizes the fact that a real number has a finite continued fraction expansion if and only if it is rational. Theorem 5.1.1 shows that a train track expansion of an arational measured foliation converges to that measured foliation in the appropriate sense, just as a continued fraction expansion of an irrational number converges to that number.

The property of irrationality of a real number has several distinct analogues in measured foliations. The weakest analogue is the property that a measured foliation \mathcal{F} possesses a noncompact leaf and, by Lemma 4.2.1 as just mentioned, this occurs if and only if each train track expansion of \mathcal{F} is infinite. A more interesting and deeper analogue of irrationality says that \mathcal{F} has no closed leaf cycles. This is called *arationality* of \mathcal{F} , and it is equivalent to several other properties of \mathcal{F} , such as that \mathcal{F} has nonzero intersection number with every simple closed curve, and also that for any measured geodesic lamination λ equivalent to \mathcal{F} , all of the complementary components of λ are nonpunctured or once-punctured discs. The Arational Expansion Theorem 6.3.2 characterizes arationality of a measured foliation in terms of combinatorial properties of any train track expansion; these properties are a somewhat deeper analogue of the property of infiniteness of a continued fraction expansion.

The Stable Equivalence Theorem 7.2.3 characterizes arational measured foliations in terms of tails of train track expansions, although one must carefully bundle train track expansions into a more complicated structure called an “expansion complex” in order to formulate stable equivalence correctly.

We also obtain results on train track expansions of pseudo-Anosov unstable foliations, describing the periodicity of these expansions. We use this idea to describe general methods for construction, classification, and enumeration of pseudo-Anosov conjugacy classes.

Finally, we shall describe an algorithm for solving the conjugacy problem in the mapping class group, by algorithmically computing complete invariants of the conjugacy class of any mapping class.

1.3 Continued fractions expansions and the torus

In order to prepare for stating our theorems on a general surface, we first formulate them on a torus, and establish the dictionary with continued fraction theorems. A brief introduction to this dictionary is given in the author's article "What is a train track?" [Mos03].

Consider the torus $T = \mathbf{R}^2/\mathbf{Z}^2$. The group $\mathrm{SL}(2, \mathbf{Z})$ acts linearly on \mathbf{R}^2 preserving the integer lattice \mathbf{Z}^2 . This action descends to a faithful action of $\mathrm{SL}(2, \mathbf{Z})$ on T , and it induces an isomorphism $\mathrm{SL}(2, \mathbf{Z}) \approx \mathcal{MCG}(T)$. Given a matrix $A \in \mathrm{SL}(2, \mathbf{Z})$ let M_A denote the corresponding mapping class on T . For each extended real number $r \in \mathbf{R}^* = \mathbf{R} \cup \{\infty\}$ there is a foliation $\tilde{\mathcal{F}}_r$ of \mathbf{R}^2 by lines of slope r , which descends to a constant slope foliation \mathcal{F}_r of T . If r is irrational then \mathcal{F}_r supports a transverse measure which is unique up to positive scalar multiplication, whereas if r is rational then the leaves of \mathcal{F}_r are all compact, isotopic to a simple closed curve of slope r denoted c_r . By this method we obtain a homeomorphism $\phi: \mathbf{R}^* \rightarrow \mathcal{PMF}(T)$, which restricts to a bijection between \mathbf{Q}^* and the set \mathcal{C} of isotopy classes of essential simple closed curves on T . The action of $A \in \mathrm{SL}(2, \mathbf{Z})$ on \mathbf{R}^* by fractional linear transformations satisfies the property $M_A(\phi(r)) = \phi(A(r))$ for all $r \in \mathbf{R}^*$, and so the action of $\mathrm{SL}(2, \mathbf{Z})$ on \mathbf{R}^* is conjugate to the action of $\mathcal{MCG}(T)$ on $\mathcal{PMF}(T)$.

In \mathbf{R}^* we use the notation $r < s < t$ to mean that the ordered triple (r, s, t) is positively oriented with respect to the standard cyclic orientation on \mathbf{R}^* . This is the smallest ternary relation containing the usual order relation on triples in $\mathbf{R} \cup \{\infty\}$ subject to the rule that $r < s < t \implies s < t < r$. For $r \neq t \in \mathbf{R}^*$ we use interval notation $(r, t) = \{s \in \mathbf{R}^* \mid r < s < t\}$ and $[r, t] = \{r\} \cup (r, t) \cup \{t\}$.

The basic train track $\tau_{[0, \infty]}$ on T , shown in Figure 1, is obtained from $c_0 \cup c_\infty$ by flattening the angles at the transverse intersection point $c_0 \cap c_\infty$ until this point has a unique tangent line of positive slope. The train track $\tau_{[0, \infty]}$ has one bigon, and $\tau_{[0, \infty]}$ carries precisely those foliations \mathcal{F}_r with $r \in [0, \infty]$. More generally, consider integers a, b, c, d so that $ad - bc = 1$. The elements $p = \frac{c}{d}, q = \frac{a}{b} \in \mathbf{Q}^*$ determine simple closed curves c_p, c_q on T which intersect transversely in a single point. One can flatten the angles at this point to form a train track in one of two ways: one flattening produces a train track $\tau_{[p, q]}$ which carries \mathcal{F}_r if and only if $r \in [p, q]$,

and the other produces a train track $\tau_{[q,p]}$ carrying \mathcal{F}_r if and only if $r \in [q,p]$. If $a, b, c, d \geq 0$ then $p = \frac{c}{d} < \frac{a}{b} = q$, and if one flattens the intersection point to have positive slope then one obtains $\tau_{[p,q]}$.

Splittings affect train tracks on T as follows. As shown in Figure 1, a R splitting on $\tau_{[\frac{0}{1}, \frac{1}{0}]}$ produces the train track $\tau_{[\frac{0}{1}, \frac{1}{1}]}$, whereas a L splitting produces the train track $\tau_{[\frac{1}{1}, \frac{1}{0}]}$. More generally, for $a, b, c, d \geq 0$ and $p = \frac{c}{d} < \frac{a}{b} = q$, L and R splittings on the train track $\tau_{[p,q]}$ are determined by the Farey sum $r = \frac{c+a}{d+b}$: a R splitting on $\tau_{[p,q]}$ produces the train track $\tau_{[p,r]}$, and an L splitting produces the train track $\tau_{[r,q]}$. The extension to arbitrary $a, b, c, d \in \mathbf{Z}$ is left to the reader.

Continued fractions expansions of real numbers correlate precisely with train track expansions of constant slope foliations on the torus T . Consider a measured foliation \mathcal{F}_r with $r \in [0, \infty]$. We construct a train track expansion $\tau_0 \succ \tau_1 \succ \dots$ of \mathcal{F}_r , based at $\tau_0 = \tau_{[\frac{0}{1}, \frac{1}{0}]}$, as follows. If $r < 1$ then we do an R splitting obtaining $\tau_1 = \tau_{[\frac{0}{1}, \frac{1}{0}]}$, if $r > 1$ then we do a L splitting obtaining $\tau_1 = \tau_{[\frac{1}{1}, \frac{1}{0}]}$. If $r = 1$ then we do a C or Central splitting and the sequence stops at $\tau_1 = c_1$. Continuing in this manner one obtains a sequence D_1, D_2, \dots which is either an infinite sequence of Ls and Rs, or is a finite sequence of Ls and Rs terminating with a single C, and one obtains a splitting sequence whose parities are labelled by the sequence D_1, D_2, \dots :

$$\tau = \tau_0 \overset{D_1}{\succ} \tau_1 \overset{D_2}{\succ} \tau_2 \overset{D_3}{\succ} \dots$$

By construction, each train track in this sequence carries \mathcal{F}_r . This splitting sequence is an example of a *train track expansion* of \mathcal{F}_r .

For example, the train track expansion of the closed curve $c_{10/7}$ based at $\tau_{[\frac{0}{1}, \frac{1}{0}]}$ is given by

$$\tau_{[\frac{0}{1}, \frac{1}{0}]} \overset{L}{\succ} \tau_{[\frac{1}{1}, \frac{1}{0}]} \overset{R}{\succ} \tau_{[\frac{1}{1}, \frac{2}{1}]} \overset{R}{\succ} \tau_{[\frac{1}{1}, \frac{3}{2}]} \overset{L}{\succ} \tau_{[\frac{4}{3}, \frac{3}{2}]} \overset{L}{\succ} \tau_{[\frac{7}{5}, \frac{3}{2}]} \overset{L}{\succ} \tau_{[\frac{10}{7}, \frac{3}{2}]} \overset{C}{\succ} c_{\frac{10}{7}}$$

From the LR sequence of this train track expansion — LRLLLL, that is, 1 L, 2 Rs, 3 Ls — one can derive the continued fraction expansion $\frac{10}{7} = 1 + \frac{1}{2 + \frac{1}{3}}$.

For any real number $r \geq 0$, the continued fraction expansion

$$r = n_0 + \frac{1}{n_1 + \frac{1}{n_2 + \dots}}$$

may be read off by inspection from the RL sequence: $n_0 \geq 0$ is the length of the initial block of L's in D_i ; $n_1 \geq 1$ is the length of the following block of R's; $n_2 \geq 2$ is the length of the following block of L's, etc. If r is rational then the sequence D_i is finite and, ignoring the terminating C, we obtain a finite continued fraction expansion, as in the example of $c_{\frac{10}{7}}$ with LR sequence LRLLLL. If r is irrational then the sequence

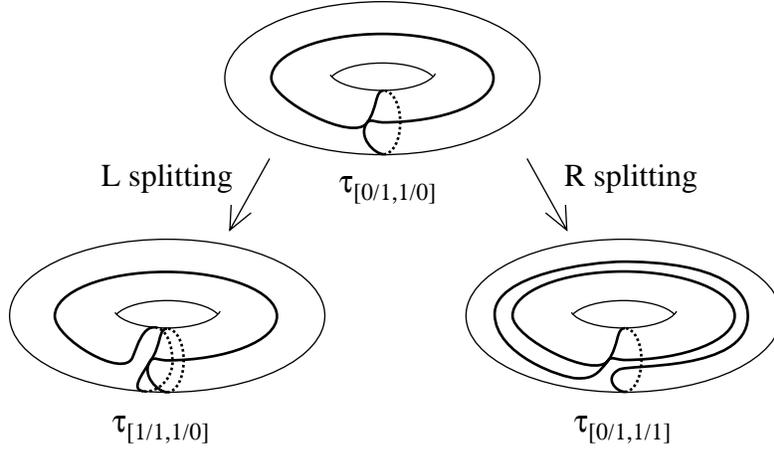


Figure 1: The train track $\tau_{[\frac{0}{1}, \frac{1}{0}]}$ carries those foliations \mathcal{F}_r whose slope r is in the interval $[0, \infty]$. A R splitting on τ yields the train track $\tau_{[\frac{0}{1}, \frac{1}{1}]}$, which carries foliations with slope in $[0, 1]$. A L splitting yields $\tau_{[\frac{1}{1}, \frac{1}{0}]}$ which carries foliations with slope in $[1, \infty]$. Not shown is a Central splitting, which yields a simple closed curve of slope 1.

D_i alternates infinitely often between L and R, and so the sequence n_0, n_1, \dots is infinite, and in this case we obtain the infinite continued fraction expansion of r .

Conversely, from the continued fraction expansion for r one can read off the train track expansion of \mathcal{F}_r . For example, the number $r = \frac{1+\sqrt{5}}{2}$ has a continued fraction expansion with $n_i = 1$ for all i , and so the train track expansion of \mathcal{F}_r based at $\tau_{[\frac{0}{1}, \frac{1}{0}]}$ consists of alternating Left and Right splittings.

The torus and the once-punctured torus have naturally isomorphic mapping class groups, Teichmüller spaces, measured foliation spaces, etc. This means that the dictionary is valid not only for the torus but also for the punctured torus, which has the advantage of being a hyperbolic surface, and hence one can use train tracks without bigons.

1.4 Measured foliations

In order to state our results, we briefly review the theory of measured foliations, initiated by Thurston [Thu88].

Let S be an oriented surface of finite type, that is, the complement of a finite set of punctures in a closed, connected, oriented surface. We assume throughout that S has negative Euler characteristic and that it is not homeomorphic to a 3-punctured

sphere. The set of nontrivial, nonperipheral simple closed curves on S up to isotopy is denoted \mathcal{C} . The Teichmüller space of S , denoted \mathcal{T} , is the set of isotopy classes of complete, finite area hyperbolic structures on S . For any given hyperbolic structure each element of \mathcal{C} is uniquely represented by a closed geodesic. Lengths of geodesics give a well-defined map $\mathcal{T} \times \mathcal{C} \rightarrow (0, \infty)$, producing an embedding $\mathcal{T} \rightarrow (0, \infty)^{\mathcal{C}}$ which imposes a natural topology on \mathcal{T} homeomorphic to an open ball.

We will adopt a broad notion of measured foliations, incorporating certain “partial” measured foliations supported on subsurfaces of S . A (total) measured foliation is a foliation with prong singularities, having 3 or more prongs at each nonpuncture singularity and 1 or more prongs at a puncture singularity, equipped with a positive transverse Borel measure. A partial measured foliation is what you get from a (total) measured foliation by slicing along finitely many finite leaf segments. Inverting this “slicing” operation we obtain a “fulfillment” operation, which from a partial measured foliation produces a (total) measured foliation by expanding the supporting subsurface to fill out all of S . The equivalence classes of measured foliations and partial measured foliations, up to isotopy, Whitehead equivalence, and fulfillment, form Thurston’s space of measured foliations $\mathcal{MF} = \mathcal{MF}(S)$. Multiplication of the transverse measure by a positive number determines a multiplicative action of $(0, \infty)$ on \mathcal{MF} , and the quotient of this action is the projective measured foliation space $\mathcal{PMF} = \mathcal{PMF}(S)$. The intersection number of a partial measured foliation \mathcal{F} with a simple closed curve c is denoted $\langle \mathcal{F}, c \rangle$, and taking the infimum over equivalence classes gives a well-defined map $\mathcal{MF} \times \mathcal{C} \rightarrow [0, \infty)$ which is homogenous in the first coordinate. This map induces embeddings $\mathcal{MF} \rightarrow [0, \infty)^{\mathcal{C}}$, $\mathcal{PMF} \rightarrow \mathcal{P}[0, \infty)^{\mathcal{C}}$ which are used to impose natural topologies on \mathcal{MF} and \mathcal{PMF} . Thurston’s Compactification Theorem says that the space \mathcal{PMF} is a sphere, and $\overline{\mathcal{T}} = \mathcal{T} \cup \mathcal{PMF}$ forms a closed ball embedded in $\mathcal{P}[0, \infty)^{\mathcal{C}}$, with interior \mathcal{T} and boundary \mathcal{PMF} .

If \mathcal{F} is a measured foliation, a *saddle connection* is a finite leaf segment each of whose endpoints is a singularity, and a *leaf cycle* is a union of saddle connections forming either a simple closed curve disjoint from the punctures or a simple arc intersecting the punctures at its endpoints. We say that \mathcal{F} is *arational* if it has no leaf cycles, or equivalently, if $\langle \mathcal{F}, c \rangle \neq 0$ for every essential, nonperipheral simple closed curve c . Arationality is invariant under the equivalence relation on measured foliations, and under positive scalar multiplication, and hence we can speak about an arational element of \mathcal{MF} or of \mathcal{PMF} . A general measured foliation can be decomposed into component foliations supported on disjoint subsurfaces, and each of these components is either arational on its support, or is an annulus foliated by parallel circles.

1.5 Train tracks and train track expansions (§§3–4)

Consider a train track τ on S which is generic, meaning that each switch has two branch ends on one side and one branch end on the other side. A measured foliation \mathcal{F} is carried on τ when \mathcal{F} can be sliced to produce a partial measured foliation which is contained in a regular neighborhood of τ , with leaves approximately parallel to branches of τ . Let $\mathcal{MF}(\tau)$ denote the subset of \mathcal{MF} represented by measured foliations carried by τ , and let $\mathcal{PMF}(\tau)$ denote the projective image of this subset in \mathcal{PMF} .

Given a generic train track τ on S , certain embedded train paths in τ are called *splitting arcs*, and the choice of a splitting arc α determines Left, Right, and Central splittings on τ . A splitting is elementary when the splitting arc is just a single branch of the train track (see Figure 13), but we allow a more general class of splittings $\tau \succ \tau'$ which can be factored as a composition of “slide moves” and a single elementary splitting (see Figure 14 and Proposition 3.13.4). Our most general definition of splittings works for arbitrary train tracks, not just ones which are generic, but for expository purposes we will stick with generic train tracks in this introduction.

Given a splitting $\tau \succ \tau'$, we have $\mathcal{MF}(\tau) \supset \mathcal{MF}(\tau')$ and $\mathcal{PMF}(\tau) \supset \mathcal{PMF}(\tau')$.

A *splitting sequence* is a finite or infinite sequence $\tau_0 \succ \tau_1 \succ \dots$. The sequence is *complete* if *either* the sequence is infinite, *or* the sequence is finite and ends in a train track τ_n each of whose components is a circle. Complete splitting sequences are the analogue, on a general surface, of continued fractions.

A *train track expansion* of a measured foliation \mathcal{F} is a complete splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ with the property that $[\mathcal{F}] \in \mathcal{MF}(\tau_n)$ for all τ_n in the sequence; equivalently, \mathcal{F} is carried by each τ_n . Thus, $\tau_0 \succ \tau_1 \succ \dots$ is an expansion of \mathcal{F} if and only if $[\mathcal{F}]$ is in the set $\bigcap_i \mathcal{MF}(\tau_i)$, or equivalently the projective class of $[\mathcal{F}]$ is in the set $\bigcap_i \mathcal{PMF}(\tau_i)$. Noting that the latter is a nonempty, compact subset of \mathcal{PMF} , it follows that every complete splitting sequence is an expansion of something; this is an analogue of the fact that every continued fraction represents some real number.

If a measured foliation \mathcal{F} is carried by a train track τ_0 , then there exists a train track expansion for \mathcal{F} starting from τ_0 ; this fact, proved in detail in Corollary 4.1.2, is the analogue of the fact that every real number has a continued fraction expansion. For example, choose any branch b of τ_0 which is a “sink branch” meaning that the switches at either end of b point “into” b , and use the weights that \mathcal{F} deposits on the branches incident to b to decide whether to do a Left, Right, or Central splitting $\tau_0 \succ \tau_1$ along b (see Figures 13 and 16, and also the “Splitting Inequalities” in Fact 3.13.1). Continuing inductively, we obtain the desired expansion $\tau_0 \succ \tau_1 \succ \tau_2 \succ \dots$. Unlike the situation on the torus, the base train track τ_0 does not

determine the entire splitting sequence, because of the choice of sink branch at each stage; on the torus, a train track has a unique sink branch, and so the only choice in splitting is the parity, Left, Right, or Central, which is determined by the measured foliation.

The statement “ $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion of \mathcal{F} ” is the analogue of the statement that a continued fraction converges to a certain number. Unlike on the torus, we generally do not bother to fix a base train track, and so there are no restrictions on the base train track τ_0 of an expansion of \mathcal{F} .

If one fixes a measured foliation \mathcal{F} and a train track τ_0 carrying \mathcal{F} , then as mentioned above there is not a uniquely determined train track expansion of \mathcal{F} based at τ_0 . However, there is a single object into which all train track expansions based at τ_0 can be packaged. This package is described in Sections 4.3 and 4.4, and the structure of this package is studied in more detail and exploited in Section 7 in order to generalize the phenomenon of stable equivalence of measured foliations.

1.6 The expansion convergence theorem (§5)

Our first result generalizes the theorem that for any infinite continued fraction

$$[n_0, n_1, n_2, \dots] = n_0 + \frac{1}{n_1 + \frac{1}{n_2 + \dots}}$$

the sequence of partial fractions $[n_0], [n_0, n_1], [n_0, n_1, n_2], \dots$ converges to a unique real number, and that number is irrational.

Given an arational measured foliation \mathcal{F} , the set of all non-negative transverse Borel measures on \mathcal{F} up to projective equivalence forms a finite dimensional simplex known as the *Choquet simplex* of \mathcal{F} , whose vertices consist precisely of the ergodic transverse measures on \mathcal{F} [Wal82]. The Choquet simplex of \mathcal{F} embeds in \mathcal{PMF} , and the image of this embedding will be denoted $\mathcal{PMF}(\mathcal{F})$. If there is a projectively unique transverse measure on \mathcal{F} , that is, if $\mathcal{PMF}(\mathcal{F})$ is a point, then \mathcal{F} is said to be *uniquely ergodic*. Almost every point in \mathcal{PMF} is uniquely ergodic [Vee82], [Mas82]. On the other hand there do exist arational points which are not uniquely ergodic, indeed such points form a subset of positive Hausdorff dimension in \mathcal{PMF} [MS91].

Another way to think of the set $\mathcal{PMF}(\mathcal{F})$ is as the set of projective classes of measured foliations \mathcal{F}' that are *unmeasured equivalent* to \mathcal{F}' , meaning that the underlying unmeasured partial foliations of S are equivalent up to the relations of isotopy, Whitehead moves, and fulfillment. Unmeasured equivalence of \mathcal{F} and \mathcal{F}' says roughly speaking that the leaves of \mathcal{F} and of \mathcal{F}' have the same asymptotic behavior; this can be made precise by lifting $\mathcal{F}, \mathcal{F}'$ to the universal cover \tilde{S} of S and studying ideal endpoints of leaves in the circle at infinity of \tilde{S} .

Theorem (Expansion Convergence Theorem 5.1.1). *If $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion of an arational measured foliation \mathcal{F} , then the set of all projective classes of measured foliations \mathcal{F}' for which $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion is equal to the set $\mathcal{PMF}(\mathcal{F})$.*

As mentioned above, this theorem appears in [Ker85], and special cases appear as Theorem 3.1 of [PP87] and as Proposition 3.3.2 of [Pen92].

Notice that the existence of nonuniquely ergodic measured foliations has the effect of weakening somewhat the analogy between the above theorem and the convergence theorem for infinite continued fractions: on a punctured torus the set $\bigcap_{i=0}^{\infty} \mathcal{PMF}(\tau_i)$ will be a point, but in general it is a Choquet simplex, possibly of positive dimension.

1.7 The rational killing criterion (§6)

The Expansion Convergence Theorem tells us that arationality of a measured foliation is an intrinsic property of a train track expansion. On a punctured torus, there is an intrinsic, combinatorial property of a train track expansion which tells us whether it is an expansion of an arational measured foliation of irrational slope: the splitting sequence has infinitely many Left splittings and infinitely many Right splittings. Motivated by these results, given a train track expansion $\tau_0 \succ \tau_1 \succ \dots$ of a measured foliation \mathcal{F} we now consider how, using only the combinatorics of a splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, one can tell whether \mathcal{F} is arational. We shall develop a “rational killing criterion” for a splitting sequence, which is necessary and sufficient for this purpose.

As an initial motivation for the rational killing criterion, we examine the situation on the punctured torus more closely. Consider a splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ on the punctured torus, which is an expansion of a measured foliation of slope r . Each train track τ_i has two subtracks, each of which is a simple loop. If $\tau_i \succ \tau_{i+1}$ is a Left splitting, then one of the two subtracks of τ_i , which we call the *Left loop*, survives and is isotopic to the Left loop of τ_{i+1} , whereas the Right loop of τ_i does not survive. The same statement holds with Left and Right reversed. It follows that a finite subsequence $\tau_i \succ \tau_{i+1} \succ \dots \succ \tau_j$ has splittings of both parities if and only if both the Left and Right subtracks of τ_i are killed before they can reach τ_j . As we have observed, the number r is irrational if and only if both Left and Right occur infinitely often, which occurs if and only if for each i there exists $j > i$ such that each subtrack of τ_i is killed before it can reach τ_j .

Our goal now is to develop an analogue of this “rational killing criterion” which holds for train track expansions on arbitrary surfaces. An earlier version of this criterion occurs in [Mos93], but the current formulation was motivated by joint work with Farb [FM99].

A train track τ is said to *fill* the surface S if every component of $S - \tau$ is a nonpunctured or once-punctured polygon. For example, if $\mathcal{PMF}(\tau)$ contains an arational element then τ fills, and hence if τ does not fill then $\mathcal{PMF}(\tau)$ consists entirely of nonarational elements.

To get the idea of a train track expansion on S that fails the rational killing criterion, start with a proper, essential subsurface S' and any splitting sequence $\sigma_0 \succ \sigma_1 \succ \cdots$ on S' , so that the train tracks σ_i fill S' , but of course they do *not* fill S . It follows that $\bigcap_i \mathcal{PMF}(\sigma_i)$ consists entirely of nonarational elements of $\mathcal{PMF}(S)$.

There are two constructions which convert a nonfilling train track on S into a filling train track: adding branches; and pinching. These lead to two constructions of train track expansions that fail the rational killing criterion.

For the first construction, let τ_0 be a filling train track obtained from σ_0 by adding branches. Think of the original branches of σ_0 as being black, and the added branches of $\tau_0 - \sigma_0$ as being gray. As you split σ_0 to σ_1 , the gray branches just go along for the ride, and you get a train track τ_1 obtained from σ_1 by adding branches. Moreover, τ_0 carries τ_1 , and this carrying can be factored into one or more splittings; it's not generally true that there is a single splitting from τ_0 to τ_1 , because the gray branches may have endpoints lying in the interior of the black splitting arc of σ_0 , and hence you may need more than one splitting to get from τ_0 to τ_1 . Now repeat the process, and you get a carrying sequence $\tau_0 \succ \tau_1 \succ \cdots$ on S , which can be factored to give a splitting sequence on S . Moreover, for each i there is a nonfilling subtrack $\sigma_i \subset \tau_i$ which survives the entire splitting sequence, in the sense that there are subtracks $\sigma_j \subset \tau_j$, $j \geq i$, such that σ_i splits to σ_j . The existence of a nonfilling subtrack of some τ_i which survives forever is a direct violation of the rational killing criterion for the splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$.

For the second construction, let τ_0 be a filling train track obtained from σ_0 by “pinching”. To give a precise example, suppose that S' is obtained from S by removing an essential, nonperipheral annulus. The train track σ_0 has one annular complementary component A_0 (see Figure 2). Let γ_0 be a properly embedded arc in A_0 connecting points on opposite boundary components, so that each endpoint of γ_0 is in the interior of some branch of σ_0 . Now pinch γ_0 down to a single point, dragging σ_0 along, creating a filling train track τ_0 with a new sink branch denoted α_0 , and note that τ_0 fills the surface; the annulus A_0 has been pinched off to form a polygonal component of $S - \tau_0$ which we shall denote B_0 . One can recover σ_0 from τ_0 in the following manner: from a core curve of the annulus A_0 , we obtain a closed curve called a “splitting cycle” of the train track τ_0 , of the form $\alpha_0 * \beta_0$ where α_0 is the newly created sink branch of τ_0 and β_0 is a properly embedded arc in B_0 connecting two cusps of B_0 . The train track σ_0 is obtained from τ_0 by splitting along the splitting cycle.

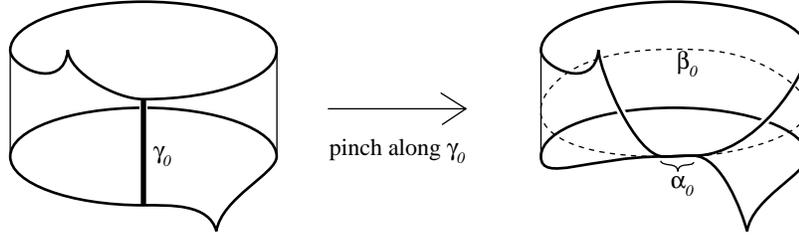


Figure 2: Pinching a train track.

Next, under the carrying map $\sigma_1 \rightarrow \sigma_0$, we can pull back the arc γ_0 of σ_0 to give an arc γ_1 of σ_1 , properly embedded in the annulus component A_1 of $S - \sigma_1$. Pinching γ_1 down to a point and dragging σ_1 along we obtain a filling train track τ_1 . Moreover, from the construction it follows that τ_0 carries τ_1 ; the carrying $\tau_0 \succ \tau_1$ factors into one or more splittings, depending on how γ_0 is situated with respect to the splitting arc of σ_0 . Also, we obtain a splitting cycle $\alpha_1 * \beta_1$ of τ_1 , so that splitting τ_1 along this splitting cycle we recover σ_1 . Continuing inductively we obtain a carrying sequence $\tau_0 \succ \tau_1 \succ \dots$, that factors to give a splitting sequence, and a splitting cycle $\alpha_i * \beta_i$ of τ_i , so that splitting τ_i along this splitting cycle recovers σ_i . Thus, for each i the splitting cycle $\alpha_i * \beta_i$ of τ_i survives the entire splitting sequence, in the sense that for each $j \geq i$ there is a splitting cycle $\alpha_j * \beta_j$ of τ_j which is isotopic to $\alpha_i * \beta_i$, both being isotopic to the core curve of the annulus $S - S'$. The existence of a splitting cycle of τ_i which survives forever is a direct violation of the rational killing criterion for the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$. More generally, if $\tau'_i \subset \tau_i$ is a filling subtrack which survives forever, the existence of a splitting cycle of τ'_i which survives forever is a direct violation of the rational killing criterion.

In each of the two constructions above, note also that $\bigcap_i \mathcal{PMF}(\tau_i)$ contains $\bigcap_i \mathcal{PMF}(\sigma_i)$, and hence contains nonrational elements of \mathcal{PMF} .

Motivated by the above discussion, we make the following definitions. Consider a splitting $\tau \succ \tau'$. Given a subtrack $\sigma \subset \tau$, we say that σ *survives* the splitting if there exists a subtrack $\sigma' \subset \tau'$ such that σ' is obtained from σ by a generalized splitting, which means a sequence of moves consisting of one splitting and the rest “slide moves”. If this is so then σ' is uniquely determined by σ , and we say that σ' is the *descendant* of σ in τ' . Given a splitting cycle c of τ , we say that c *survives* the splitting if there exists a splitting cycle c' of τ' such that c, c' are isotopic; if this happens then c' is uniquely determined by c and we say that c' is the *descendant* of c .

A finite splitting sequence $\tau_i \succ \tau_{i+1} \succ \dots \succ \tau_j$ satisfies the *rational killing criterion* if the following two statements are true:

- Each nonfilling subtrack of τ_i is killed before it can get to τ_j .
- If τ'_i is a filling subtrack of τ_i which survives to some filling subtrack τ'_j of τ_j , then each splitting cycle of τ'_i is killed before it can get to τ'_j .

An infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ satisfies the *iterated rational killing criterion* if for each $i \geq 0$ there exists $j > i$ such that the finite sequence $\tau_i \succ \dots \succ \tau_j$ satisfies the rational killing criterion.

1.8 Expansions of arational measured foliations (§6 continued)

The iterated rational killing criterion characterizes when a splitting sequence is an expansion of some arational measured foliation:

Theorem (Arational Expansion Theorem 6.3.2). *Suppose that $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion of a measured foliation \mathcal{F} . Then \mathcal{F} is arational if and only if the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ satisfies the iterated rational killing criterion.*

As remarked above, this theorem generalizes the situation on the torus T , because a splitting sequence on T satisfies the iterated rational killing criterion if and only if the parity of the sequence alternates infinitely often between L and R. According to the dictionary between continued fractions and train track expansions on T , this occurs if and only if the given splitting sequence is an expansion of a measured foliation on T of irrational slope.

We remark that RL sequences play an important role in earlier treatments of train track expansions such as [PP87] and [Mos93]. Nevertheless, obtaining an arationality criterion stated directly in terms of RL sequences seems to be something of a stretch. For example, it is not true that \mathcal{F} is a pseudo-Anosov stable foliation if and only if some train track expansion of \mathcal{F} has eventually periodic RL sequence—instead, the latter property is equivalent to the existence of a foliation component of \mathcal{F} which is the stable foliation of a pseudo-Anosov homeomorphism on some subsurface of S . Similarly, \mathcal{F} has an infinite train track expansion alternating infinitely often between L and R if and only if there is a foliation component of \mathcal{F} which is arational on its supporting subsurface. RL sequences therefore do not capture the full complexity of arationality, and for this reason we avoid RL sequences in the present work.

A specialized version of the Arational Expansion Theorem given in [Mos93] only deals with what are here called “one sink expansions”, as explained in Section 9.4 below. One sink expansions are particularly easy to work with from a combinatorial and computational point of view, not least because their train tracks have a very special, easily recognizable form.

The present, general version of the Arational Expansion Theorem was inspired by joint work of the author with Benson Farb [FM99], in which train track expansions of geodesics in \mathcal{T} are used to characterize precompactness of geodesics in the moduli space $\text{Mod}(S) = T/\mathcal{MCG}$, where \mathcal{MCG} is the mapping class group of S . Indeed, the Precompactness Theorem of [FM99] is a uniform version of the Arational Expansion Theorem:

Theorem (Precompactness Criterion [FM99]). *Suppose that $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion of a geodesic ray $g: [0, \infty) \rightarrow \mathcal{T}$. Then the projection of g to $\text{Mod}(S)$ is precompact if and only if there exists $K \geq 0$ such that for all $i \geq 0$, the pair τ_i, τ_{i+K} satisfies the rational killing criterion.*

Canonical expansions. Every arational measured foliation \mathcal{F} is equivalent to one which has no saddle connections: inductively collapse each saddle connection to a point. Starting from such a representative, if we then slice open along a short initial segment of each separatrix we obtain a partial measured foliation called a *canonical model* for \mathcal{F} , whose only singularities are 3-pronged boundary singularities. Canonical models are unique up to isotopy: any two in the same equivalence class are isotopic.

Suppose that we have a canonical model \mathcal{F} . If a train track τ carries \mathcal{F} so that \mathcal{F} maps onto a regular neighborhood of τ , then we say that τ *canonically carries* \mathcal{F} . A *canonical train track expansion* of \mathcal{F} is a train track expansion $\tau_0 \succ \tau_1 \succ \dots$ such that each τ_i canonically carries \mathcal{F} .

As a corollary of the Arational Expansion Theorem we shall prove the Canonical Expansion Theorem, in which we characterize Canonical Expansions by a combinatorial method similar to that by which Arational Expansions were characterized: we define a “canonical killing criterion” and show that $\tau_0 \succ \tau_1 \succ \dots$ is a canonical expansion of some measured foliation if and only if it satisfies the canonical killing criterion.

1.9 Stable equivalence (§7)

The Stable Equivalence Theorem for continued fractions says that given $r, s \in \mathbf{R}$, the numbers r, s are in the same orbit of the fractional linear action of $\text{SL}(2, \mathbf{R})$ on $\mathbf{R} \cup \{\infty\}$ if and only if the continued fraction expansions $[m_0, m_1, \dots], [n_0, n_1, \dots]$ of r, s are stably equivalent, meaning that for some integers $a, b \geq 0$ we have $m_{a+i} = n_{b+i}$ for all $i \geq 0$. One can view this theorem as classifying orbits of the $\text{SL}(2, \mathbf{R})$ action on $\mathbf{R} \cup \{\infty\}$ in terms of combinatorial invariants, namely, tails of sequences of positive integers.

Stable equivalence on the torus. Translating this result in a straightforward manner to the torus, we obtain the following equivalent statement, which one can view as giving a complete classification of constant slope foliations up to topological equivalence, stated in terms of combinatorial invariants, namely, tails of train track expansions:

Theorem 1.9.1 (Torus Stable Equivalence, version 1). *Given $\mathcal{F}, \mathcal{F}'$ two constant slope foliations on the torus, if $\tau_0 \succ \tau_1 \succ \dots$ and $\tau'_0 \succ \tau'_1 \succ \dots$ are the expansions of $\mathcal{F}, \mathcal{F}'$ starting at the standard base train track $\tau_0 = \tau'_0$, then $\mathcal{F}, \mathcal{F}'$ are in the same orbit under the action of the mapping class group of the torus if and only if there exist integers $c, d \geq 0$ such that for all $j \geq 1$ the splittings $\tau_{c+j-1} \succ \tau_{c+j}$ and $\tau'_{d+j-1} \succ \tau'_{d+j}$ have the same parity.*

The “if” direction in the above theorem is straightforward to prove: if c, d exist as described, then by choosing a mapping class ϕ such that $\phi(\tau_c) = \tau_d$, one can show that $\phi(\mathcal{F}) = \mathcal{F}'$.

The “only if” direction of the above theorem is a consequence of the following result. Two splitting sequences $\tau_0 \succ \tau_1 \succ \dots$ and $\tau'_0 \succ \tau'_1 \succ \dots$ are said to be *stably equivalent* if they eventually coincide up to isotopy, meaning that there exist integers $c, d \geq 0$ such that for all $i \geq 0$ the train tracks τ_{c+i}, τ'_{d+i} are isotopic.

Theorem 1.9.2 (Torus Stable Equivalence, version 2). *If \mathcal{F} is a foliation on the torus of constant, irrational slope, then any two train track expansions of \mathcal{F} are stably equivalent.*

To see the relation between the above two theorems, we give here the proof that version 2 implies the “only if” direction of version 1 as follows. Define a weaker equivalence relation among torus splitting sequences: two splitting sequences $\tau_0 \succ \tau_1 \succ \dots$ and $\tau'_0 \succ \tau'_1 \succ \dots$ are *unmarked stably equivalent* if there exist integers $c, d \geq 0$ such that for all $i \geq 1$ the splittings $\tau_{c+i-1} \succ \tau_{c+i}$ and $\tau'_{d+i-1} \succ \tau'_{d+i}$ have the same parity. It is an easy observation that stable equivalence implies unmarked stable equivalence. Moreover, for any splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ and any mapping class ϕ , the sequences $\tau_0 \succ \tau_1 \succ \dots$ and $\phi(\tau_0) \succ \phi(\tau_1) \succ \dots$ are unmarked stably equivalent: a mapping class, since it preserves orientation, preserves parity. To prove the “only if” statement in version 1 of Torus Stable Equivalence, supposing that $\mathcal{F}' = \phi(\mathcal{F})$, it follows that $\tau_0 \succ \tau_1 \succ \dots$ and $\phi(\tau_0) \succ \phi(\tau_1) \succ \dots$ are unmarked stably equivalent, but the latter sequence is an expansion of \mathcal{F}' and so it is stably equivalent, in both the marked and unmarked sense, to $\tau'_0 \succ \tau'_1 \succ \dots$.

As the proof shows, the difference between versions 1 and 2 is that in version 1 we use unmarked stable equivalence to classify measured foliations up to unmarked equivalence—that is, up to the action of the mapping class group—whereas in ver-

sion 2 we use (marked) stable equivalence to classify measured foliations up to equivalence in the measured foliation space \mathcal{MF} .

Stable equivalence on a finite type surface. We wish to formulate stable equivalence results for any finite type surface S that generalize both the marked and unmarked versions of Torus Stable Equivalence. Marked stable equivalence will classify measured foliations up to unmeasured equivalence. Unmarked stable equivalence will classify measured foliations up to the coarser relation of *topological equivalence*, where $\mathcal{F}, \mathcal{F}'$ are topologically equivalent if they may be chosen in their equivalence classes so that there exists some homeomorphism taking \mathcal{F} to \mathcal{F}' ; this homeomorphism need not be isotopic to the identity. To put it another way, unmarked stable equivalence will classify the \mathcal{MCG} orbits of the sets $\mathcal{PMF}(\mathcal{F})$ for arational measured foliations \mathcal{F} , generalizing the classification of $\mathrm{PSL}(2, \mathbf{Z})$ orbits of irrational numbers by stable equivalence of continued fraction expansions.

Considering first the marked version (version 2), if one attempts a straightforward generalization of Theorem 1.9.2 then one immediately encounters a well-known weakness in the analogy between continued fraction expansions and train track expansions: as remarked earlier, a train track expansion of a measured foliation \mathcal{F} based at a given train track is not at all unique, for if τ is any train track fully carrying \mathcal{F} then one can choose any splitting arc of τ along which to split consistently with \mathcal{F} . To put it another way, there is no *natural* train track expansion of \mathcal{F} starting with any given train track τ_0 . The lack of a natural choice is evident in other treatments of stable equivalence, for example Theorem 3.3.1 of [Pen92]. In particular, on any oriented, finite type surface which has genus 0 with ≥ 5 punctures, genus 1 with ≥ 2 punctures, or genus ≥ 2 , any arational measured foliation on S actually has *uncountably many* distinct train track expansions, and so certainly there exist two which are not stably equivalent.

We use two methods for resolving this lack of naturality. One method is to restrict attention to a small subclass of train track expansions; Theorems 7.1.7 and 7.1.5 show that by restricting to “one cusp” train track expansions we obtain a good stable equivalence theory. Another method, which has the advantage of producing stable equivalence results about arbitrary train track expansions, is to combine all expansions of a given measured foliation \mathcal{F} based at a given train track τ_0 into one structure, which we call the *expansion complex* of \mathcal{F} based at τ_0 ; Theorems 7.7.1 and 7.2.3 describe stable equivalence theory for expansion complexes.

One cusp train track expansions. To describe these, consider an ordered pair (τ, v) consisting of a generic train track τ and cusp v of τ . Associated to v there exists a unique minimal splitting arc α having v as one endpoint. Fixing a choice

of $d \in \{L, R\}$ consider the splitting $\tau \succ \tau'$ of parity d along the arc α . Under the natural correspondence between cusps of τ and cusps of τ' , let v' be the cusp of τ' corresponding to v . We obtain an ordered pair (τ', v') , and we say that $(\tau, v) \succ (\tau', v')$ is a *one cusp splitting* of parity d . Consider now an arational measured foliation \mathcal{F} , a train track τ_0 canonically carrying \mathcal{F} , and a cusp v_0 of τ_0 ; associated to this data there is a naturally defined one cusp train track expansion of \mathcal{F} of the form $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$. Thus, we do not obtain a single natural expansion of \mathcal{F} based at τ_0 , but we do obtain a natural finite set of expansions, namely, one for each cusp of τ_0 .

We can state a marked and an unmarked version of stable equivalence for one cusp train track expansions. For the marked version, we use the same definition as before of stable equivalence: two splitting sequences on S are stably equivalent if they eventually coincide up to isotopy.

Theorem (One cusp stable equivalence, marked version (Theorem 7.1.5)).

Given an arational measured foliation \mathcal{F} and two train tracks τ, τ' canonically carrying \mathcal{F} , there is a bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$, such that if $v \in \text{cusps}(\tau)$ and $v' \in \text{cusps}(\tau')$ correspond then the one cusp train track expansion of \mathcal{F} based at (τ, v) is stably equivalent to the one cusps train track expansion of \mathcal{F} based at (τ', v') .

For the unmarked version, we say that two train tracks τ, τ' are *combinatorially equivalent* if they are pieced together out of switches, branches, and oriented 2-cells in the same combinatorial pattern. An equivalent way to say this is simply that there exists an orientation preserving homeomorphism f of the surface such that $f(\tau) = \tau'$. Similarly, a pair of cusped train tracks $(\tau, v), (\tau', v')$ are combinatorially equivalent if there exists a homeomorphism f such that $f(\tau, v) = (\tau', v')$. We say that two one-cusp splitting sequences $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ and $(\tau'_0, v'_0) \succ (\tau'_1, v'_1) \succ \dots$ are *unmarked stably equivalent* if there exist integers $c, d \geq 0$ such that for all $i \geq 0$, the one cusped train tracks (τ_{c+i}, v_{c+i}) and (τ'_{d+i}, v'_{d+i}) are combinatorially equivalent, and the splittings $\tau_{c+i} \succ \tau_{c+i+1}$ and $\tau'_{d+i} \succ \tau'_{d+i+1}$ have the same parity.

Theorem (One cusp stable equivalence, unmarked version (Theorem 7.1.7)).

Given two arational measured foliations $\mathcal{F}, \mathcal{F}'$, the following are equivalent:

- *There exists a mapping class ϕ such that $\mathcal{PMF}(\phi(\mathcal{F})) = \mathcal{PMF}(\mathcal{F}')$.*
- *For any train tracks τ, τ' that canonically carry $\mathcal{F}, \mathcal{F}'$, there exists cusps v, v' of τ, τ' , respectively, such that the expansion of \mathcal{F} based at (τ, v) is unmarked stably equivalent to the expansion of \mathcal{F}' based at (τ', v') .*
- *For any train tracks τ, τ' that canonically carry $\mathcal{F}, \mathcal{F}'$, respectively, there is a bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$, such that if $v \in \text{cusps}(\tau)$ and $v' \in \text{cusps}(\tau')$*

correspond, the expansion of \mathcal{F} based at (τ, v) is unmarked stably equivalent to the expansion of \mathcal{F}' based at (τ', v') .

Expansion complexes. Next we describe another approach to stable equivalence. Consider an arational measured foliation \mathcal{F} and a train track τ that canonically carries \mathcal{F} . Instead of focussing on a special class of expansions of \mathcal{F} based at τ , we shall describe a stable equivalence phenomenon which encompasses *all* expansions of \mathcal{F} based at τ . To do this we define an object called the *expansion complex* of \mathcal{F} based at τ , which incorporates all train track expansions of \mathcal{F} based at τ , and that depends naturally on τ . For any other train track τ'_0 carrying \mathcal{F} , the expansion complexes based at τ_0 and at τ'_0 exhibit the appropriate form of (marked) stable equivalence. Also, we will have an appropriate version of unmarked stable equivalence.

To describe expansion the expansion complex of \mathcal{F} based at τ , we may choose \mathcal{F} within its equivalence class so that it is a canonical model, and so that it is contained in a regular neighborhood of τ_0 with leaves approximately parallel to branches of τ_0 . Since \mathcal{F} is a canonical model, each of its singularities s is a 3-pronged boundary singularities with two of the prongs on the boundary of the support of \mathcal{F} , and one of the prongs being the start of an infinite half leaf of \mathcal{F} based at s , called an *infinite separatrix*. Let N be the number of infinite separatrices. Each infinite separatrix can be parameterized by the half-line $[0, \infty)$, and a choice of parameter determines a finite separatrix based at s . If one chooses a point in each infinite separatrix, one obtains a collection of finite separatrices called a *separatrix family*. The set of all separatrix families of \mathcal{F} is therefore parameterized by the positive orthant in N -dimensional Euclidean space $[0, \infty)^N$.

As we shall describe in Sections 4.3 and 4.4, one can use the space $[0, \infty)^N$ of separatrix families to give a unified description of all train track expansions of \mathcal{F} based at τ_0 . By using a carrying map from \mathcal{F} to τ_0 , we can associated to any separatrix family $\xi \in [0, \infty)^N$ a train track denoted $\sigma(\xi)$. As the coordinates of the point ξ grow, the train track $\sigma(\xi)$ undergoes splittings. Starting from the origin of the orthant $[0, \infty)^N$ and moving along some path out to infinity so that each coordinate grows monotonically, one obtains an “expanding separatrix family”, which generates a train track expansion of \mathcal{F} based at τ_0 . In Proposition 4.4.1 we prove that every train track expansion of \mathcal{F} based at τ_0 arises in this manner from some expanding separatrix family.

In Section 7 we use the train tracks $\sigma(\xi)$ to define a decomposition of the parameter space $\Xi = [0, \infty)^N$, where two points $\xi, \eta \in \Xi$ are in the same decomposition element if, roughly speaking, the train tracks $\sigma(\xi), \sigma(\eta)$ are isotopic. Each decomposition element is thus labelled with the isotopy class of some train track. In

Theorem 7.2.2 we prove that the decomposition elements form the open cells of a linear cell decomposition on $\Xi = [0, \infty)^N$, and we prove that the face relation among cells of this decomposition can be described in terms of a natural relation among train tracks which is similar in nature to the splitting relation. The labelled cell decomposition Ξ is called the *expansion complex* of \mathcal{F} based at τ_0 . There is a slight lie in this discussion which we are overlooking for sake of exposition: the train track $\sigma(\xi)$ has some additional structure imposed on it in the form of certain decorations on the cusps of the train track, and hence the labels on the cells of Ξ are isotopy class of *decorated* train tracks.

Theorem (Stable Equivalence Theorem, marked version 7.2.3). *Given an arational measured foliation \mathcal{F} and two train tracks τ_0, τ'_0 that canonically carry \mathcal{F} , if Ξ, Ξ' are the expansion complexes of \mathcal{F} based at τ_0, τ'_0 respectively, then Ξ, Ξ' are stably equivalent in the sense that there exist subcomplexes $\hat{\Xi}, \hat{\Xi}'$, each a neighborhood of infinity in the one-ended topological space $[0, \infty)^N$, and there exists a label preserving cellular isomorphism $\Phi: \hat{\Xi} \rightarrow \hat{\Xi}'$.*

To say that the cellular isomorphism Φ preserves labels simply means that if $\Phi(c(\sigma)) = c'(\sigma')$, where $c(\sigma)$ is a cell of $\hat{\Xi}$ and $c'(\sigma')$ is a cell of $\hat{\Xi}'$ labelled by train tracks σ, σ' , respectively, then σ, σ' are isotopic.

Next we sketch the unmarked version of stable equivalence for expansion complexes, at the expense of a lack of precision and slight lie; the precise and true statement of the theorem is given in Section 7.7. As in one cusp stable equivalence, the statement is based on combinatorial equivalence of train tracks. Two expansion complexes Ξ, Ξ' are said to be *unmarked stably equivalent* if there exist subcomplexes $\hat{\Xi} \subset \Xi, \hat{\Xi}' \subset \Xi'$, each a neighborhood of infinity, and there exists a cellular isomorphism $\Phi: \hat{\Xi} \rightarrow \hat{\Xi}'$ with the following properties:

- For any cell $c(\sigma)$ of $\hat{\Xi}$, setting $c'(\sigma') = \Phi(c(\sigma))$, the train tracks σ, σ' are combinatorially equivalent.
- Φ preserves the face relation among train tracks labelling cells (for a precise statement see Section 7.7).

Theorem (Unmarked Stable Equivalence Theorem 7.7.1). *For $i = 1, 2$ let \mathcal{F}_i be an arational measured foliation canonically carried on a train track τ_i and let Ξ_i be the expansion complex of \mathcal{F}_i based at τ_i . The following are equivalent:*

- There exists $\Phi \in \mathcal{MCG}$ such that $\mathcal{PMF}(\Phi(\mathcal{F}_1)) = \mathcal{PMF}(\mathcal{F}_2)$.
- The expansion complexes Ξ_1, Ξ_2 are unmarked stably equivalent.

The reason this statement is a slight lie is that we have not defined unmarked stable equivalence correctly. Given a train track τ that fills the surface, if we pick one distinguished cusp v then the pair (τ, v) has trivial stabilizer in \mathcal{MCG} , however the train track τ itself may have a nontrivial finite stabilizer in \mathcal{MCG} . Because of this, it is not enough to define an unmarked stable equivalence $\Phi: \hat{\Xi} \rightarrow \hat{\Xi}'$ in terms of combinatorial equivalence of train tracks σ, σ' labelling corresponding cells: one must pick a particular mapping class that realizes the combinatorial equivalence; this mapping class is needed in order to make sense out of the second requirement that “ Φ preserves the face relation among train tracks labelling cells”.

1.10 General splitting sequences (§8)

Infinite splitting sequences on a torus T satisfy a simple dichotomy corresponding to the irrational–rational dichotomy for real numbers. A splitting sequence on T in which each parity L, R occurs infinitely often is an expansion of a foliation of irrational slope on T . On the other hand, if a splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ has eventually constant parity $d \in \{L, R\}$ then $\tau_0 \succ \tau_1 \succ \dots$ is an expansion of some essential closed curve c , and moreover $\tau_0 \succ \tau_1 \succ \dots$ is an expansion of a measured foliation \mathcal{F} if and only if \mathcal{F} is a foliation by closed curves isotopic to c . In this case, for i sufficiently large that the parity is constant, and the train track τ_{i+1} is obtained from τ_i by a d Dehn twist along the curve c .

On a general surface S , infinite splitting sequences satisfy a more complicated classification scheme, blending the two extremes of irrationality and rationality in much the same way that Thurston’s classification of mapping classes blends pseudo-Anosov and finite order phenomena. At one extreme, an infinite splitting sequence on S that satisfies the iterated rational killing criterion is an expansion of an arational measured foliation. At the other extreme is a kind of “Dehn twist” behavior, described in Sections 8.1 and 8.2 as a “twist splitting sequence” on S , which means that asymptotically the splitting sequence behaves like iterated Dehn twists along some essential curve system on S .

To describe the results, suppose that $\tau_0 \succ \tau_1 \succ \dots$ is an expansion of a measured foliation \mathcal{F} on S . We say that \mathcal{F} is *infinitely split* by the sequence $\tau_0 \succ \tau_1 \succ \dots$ if the sequence of subtracks that fully carry \mathcal{F} itself forms an infinite splitting sequence.

Theorem 8.2.1 shows that a splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ is eventually a twist splitting sequence if and only if it does not infinitely split any measured foliation of which it is an expansion. Moreover, in this case there exists a single train track μ such that $\tau_0 \succ \tau_1 \succ \dots$ is an expansion of a measured foliation \mathcal{F} if and only if \mathcal{F} is carried by μ ; that is, $\cap_i \mathcal{MF}(\tau_i) = \mathcal{MF}(\mu)$.

Theorem 8.5.1 gives the general structure of an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, blending the two extremes. On the one hand, there is a (possibly empty)

collection of partial arational measured foliations $\mathcal{F}_1, \dots, \mathcal{F}_J$ with pairwise disjoint supports, such that a measured foliation $\mathcal{F} \in \cap_i \mathcal{MF}(\tau_i)$ is infinitely split by $\tau_0 \succ \tau_1 \succ \dots$ if and only if \mathcal{F} has a component which is topologically equivalent to one of $\mathcal{F}_1, \dots, \mathcal{F}_J$. On the other hand there is a (possibly empty) train track μ , with support disjoint from the supports of $\mathcal{F}_1, \dots, \mathcal{F}_J$, such that a measured foliation $\mathcal{F} \in \cap_i \mathcal{MF}(\tau_i)$ is not infinitely split if and only if \mathcal{F} is carried by μ . The entire set $\cap_i \mathcal{MF}(\tau_i)$ can be described as the join of the sets $\mathcal{MF}(\mathcal{F}_1), \dots, \mathcal{MF}(\mathcal{F}_J), \mathcal{MF}(\mu)$. The corollary to Theorem 8.5.1 also gives information on how to examine the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ and extract from it canonical expansions of $\mathcal{F}_1, \dots, \mathcal{F}_J$, the train track μ , and the “twist splitting” phenomenon.

1.11 Construction of pseudo-Anosov mapping classes (§9)

There is an extensive literature on methods of construction of pseudo-Anosov mapping classes. Many of these methods depend at their root on the following decision problem. Given a mapping class ϕ and a train track τ which is an “invariant train track” for ϕ , meaning that ϕ carries the image $\phi(\tau)$, use the structure of the map $\tau \rightarrow \phi(\tau) \rightarrow \tau$ to decide whether ϕ is pseudo-Anosov. We shall describe a solution to this decision problem due to Casson [CB], as well as some variations on the solution designed to make it more computationally efficient. Then we shall present proofs of several pseudo-Anosov constructions from the literature, in particular a recipe of Penner [Pen88], as well a method of the author in [Mos86] that produces a power of every pseudo-Anosov mapping class.

1.12 Classification of pseudo-Anosov conjugacy classes (§10 in progress)

In this section we will describe complete invariants for pseudo-Anosov conjugacy classes. Here is an outline.

There is a folk theorem, known for a long time, which gives a straightforward method of classifying pseudo-Anosov conjugacy classes using a finite amount of data, and this data can be compared in a straightforward manner, giving a solution to the pseudo-Anosov part of the conjugacy problem. To describe this method, suppose that $f: S \rightarrow S$ is a pseudo-Anosov homeomorphism. Let $\tau \subset S$ be an invariant train track for f , meaning that $f(\tau)$ is carried by τ . There is a map $h: S \rightarrow S$ homotopic to the identity such that $h \mid f(\tau)$ is a smooth map onto τ . One can arrange so that h takes switches to switches, and so the composed map $(S, \tau) \xrightarrow{f} (S, f(\tau)) \xrightarrow{h} (S, \tau)$ can be encoded finitistically; let this map be called a *train track representation* of f . Conjugacy between two train track representations can be decided algorithmically, by comparing the finite encodings. Although $f: S \rightarrow S$ has infinitely many train track representations up to homotopy, it has only finitely many up to conjugacy,

because the mapping class of f acts on train track representations with finitely many orbits: if $hf: (S, \tau) \rightarrow (S, \tau)$ is one train track representation.

One can compute the finite set of conjugacy types of train track representations as follows. First find a single train track representation, using the algorithm in [BH95], say; one can also decide whether the mapping class is pseudo-Anosov using any single train track representation. Then one can generate new train track representations up to conjugacy, by the folding operations of [BH95]. Eventually all possible folding operations have been tested and no new train track representations will be produced. Since any two orbits of train track representations are connected by a sequence of foldings, this produces the entire list of train track representations.

Although this does indeed give a solution to the pseudo-Anosov part of the conjugacy problem, there are some serious efficiency issues to consider. By taking into account the natural structure of train track expansions and expansion complexes, one obtains a much more efficient solution. We will not attempt to quantify the computational complexity of this solution, but it seems likely that this can be done.

The method is illustrated first for the punctured torus, which amounts to classifying elements of $\mathrm{SL}(2, \mathbf{Z})$ whose trace has absolute value > 2 . In this setting, the classification can be derived from the classical theorems of continued fractions. Given $A \in \mathrm{SL}(2, \mathbf{Z})$ with $|\mathrm{Tr}(A)| > 2$, the slope s of the eigenvector of A has eventually periodic continued fraction expansion. Three pieces of data determine the conjugacy class of A : the primitive period loop of the continued fraction expansion of s ; the largest integer k for which A is a k th power; and the sign of the trace of A . This data can be conveniently encoded using train track expansions on the once punctured torus.

The general classification is obtained by using our results that parallel continued fraction theory, in particular the Stable Equivalence Theorem. A special case of this was carried out in [Mos83], see also [Mos86], enough to give a finite-to-one conjugacy pseudo-Anosov conjugacy classification, and a method for computing the invariants. The reason one only obtained a finite-to-one classification in that case can be traced to the difference between expansion sequences and expansion complexes. Here is a complete conjugacy classification.

If $f: S \rightarrow S$ is a pseudo-Anosov mapping class with unstable measured foliation \mathcal{F}^u , then any expansion complex of \mathcal{F}^u is eventually f -periodic; this is a consequence of the Stable Equivalence Theorem and the fact that $f(\mathcal{F}^u)$ is topologically equivalent to \mathcal{F}^u . The quotient by f of the periodic part of the expansion complex gives a complex homotopy equivalent to the circle, whose cells are labelled by combinatorial types of train tracks, and whose face relations are labelled by simple combinatorial operations between train tracks. This finite amount of data is called the *circular expansion complex* of f , and it is a complete invariant of the conjugacy class of f . Any two circular expansion complexes can be compared for isomorphism,

thus deciding the pseudo-Anosov part of the conjugacy problem.

1.13 Enumeration of pseudo-Anosov conjugacy classes (§11 in progress)

We describe how to use the classification described in §10 to get a (virtual) enumeration of pseudo-Anosov conjugacy classes.

1.14 The conjugacy problem in the mapping class group (§12 in progress)

This section describes complete invariants of all conjugacy classes in the mapping class group, in terms of the Thurston decomposition of a mapping class. By tailoring methods of Bestvina and Handel [BH95] we shall also describe an algorithm for computing the invariants. Here is an outline.

Recall that a mapping class Φ is *reducible* if there exists a *reducing family*, which is a finite family $\{a_j\}$ of essential closed curves, no two of which are isotopic, such that the family $\{a_j\}$ is Φ -invariant up to isotopy.

Thurston's classification of mapping classes says that each mapping class Φ has a reducing family $\{a_j\}$ so that for each component S' of $S - \cup_j a_j$, the first return of Φ to S' is either pseudo-Anosov or finite order; moreover, with the additional requirement that $\{a_j\}$ has minimal cardinality, it follows that $\{a_j\}$ is unique up to isotopy. Minimal cardinality of $\{a_j\}$ is equivalent to saying that for each a_j , if S' is the essential subsurface obtained by deleting a_j and conglomerating the complementary subsurfaces incident to a_j , then the first return to S' is neither pseudo-Anosov nor finite order.

Let A_j be a regular neighborhood of a_j , let $\{S_i\}$ be the set of components of $S - \cup_j \text{int}(A_j)$, and let $\{\gamma_k\}$ be the collection of boundary curves of the A_j , so each γ_k bounds an S_i on one side and an A_j on the other side. There is a representative $f: S \rightarrow S$ of the mapping class which has the following properties. First, f preserves the decomposition $S = (S_1 \cup \dots \cup S_k) \cup (A_1 \cup \dots \cup A_l)$. For each γ_k , if f^n is the lowest power which maps γ_k onto itself by an orientation reversing homeomorphism then the restriction of f^n to γ_k has finite rotation number, and in particular has a periodic orbit. No subsurface S_i is a nonpunctured annulus, and each A_j is a nonpunctured annulus. The first return map of f to each S_i is either a finite order homeomorphism, or a pseudo-Anosov homeomorphism of a surface-with-boundary, which means that the induced homeomorphism obtained by collapsing each boundary component to a puncture is pseudo-Anosov and the first return map of f to S_i is obtained by blowing up such punctures. Finally, the decomposition is minimal, in the sense that for each A_j , if the first return of f to the S_i 's incident to each side of A_j are finite order, then the first return of f to the union of A_j with the S_i 's incident to each

side is *not* of finite order.

From this decomposition one writes down a list of complete conjugacy invariants, as follows.

The first invariant is a graph with one vertex for each S_i and one edge for each A_j , together with the induced homeomorphism on this graph. Let the graph be denoted Γ .

The second invariant is the pseudo-Anosov data. For each pseudo-Anosov subsurface S_i , label the corresponding vertex of Γ with the finite data that classifies the conjugacy class of the first return of f to S_i . We must keep track of an extra finite amount of data in which boundary components of S_i are colored according to which edge of Γ they are attached to; thus, the data classifying f keeps track of the color of a boundary component collapsed to a puncture.

The third invariant is the finite order data. Nielsen described complete invariants of finite order mapping classes. Each finite order mapping class is represented by a finite order homeomorphism $f: S \rightarrow S$. The conjugacy class of f is determined by the orders and rotation numbers of the periodic points of f of nonmaximal period; we call this data the *Nielsen invariant* of f (see e.g. [Gil76] for details). It is convenient to record this data in terms of an f -invariant cell decomposition of S such that each periodic point of nonmaximal order is a vertex; we refer to this as an f -decomposition of S . From any f -decomposition one can read off the Nielsen invariants. For each finite order subsurface S_i of the Thurston decomposition, label the corresponding vertex of Γ with the Nielsen data, again using additional “coloring” data to keep track of which boundary components of S_i are attached to which edges of Γ .

The fourth data is the twist data associated to an annulus A_j . If f^n is the lowest power that maps A_j to itself preserving each component of ∂A_j , then the restriction $f^n | \partial A_j$ has rational rotation number on each component. Pick an arc α connecting a periodic point in one component to a periodic point in the other component. The arc $f^n(\alpha)$ differs from α by a rational number called the *twist* of $f^n | A_j$, and this number is well-defined independent of the choice of α . To compute this number, the path $f^n(\alpha)$ is obtained, up to path homotopy, by going around a rational number of times on one boundary component, then travelling along α , then going around a rational number of times on the other boundary component; the sum of these two rational numbers is the desired twist. Going around a component of ∂A_j a rational number of times means travelling along a path whose initial and terminal points are on the same periodic orbit.

This completes the conjugacy data. This data is finitistic, and it is complete. The data for any two mapping classes is finitistically comparable: one checks for an isomorphism between the two graphs which preserves the conjugacy data associated to each edge and vertex, respecting colors. The data is moreover computable, which

completes the solution to the conjugacy problem.

To see how the data is computed, the main tool is an algorithm for finding an invariant train track, say, the algorithm of [BH95]. Given a mapping class Φ , what this algorithm actually does is to produce one of two outcomes: either an invariant train track for Φ ; or a finite order homeomorphism $f: S \rightarrow S$ representing Φ together with an f -decomposition. In the latter case one can compute the finite order invariant for Φ and we are done.

Suppose then that one produces an invariant train track τ for Φ . As we will describe in Section 9, one can use τ to decide whether Φ is pseudo-Anosov; if so, then following the method described in Section 1.12 we are done.

If one discovers, using τ , that Φ is not pseudo-Anosov, the method by which this is discovered produces a way to alter τ by a sequence of central splittings, to produce a new invariant train track for Φ which does not fill the surface. One can continue this process, simplifying τ by central splittings, as long as the result is still a Φ -invariant train track. Let τ continue to denote the end product.

For each component τ' of τ , let $\text{Supp}(\tau')$ be a minimal essential subsurface of S containing τ' , obtained from a regular neighborhood of τ by taking the union with any complementary component of the regular neighborhood that is a nonpunctured or once-punctured disc. Because we simplified τ as much as possible using central splittings, it follows that the first return of Φ to each subsurface $\text{Supp}(\tau')$ is irreducible. If τ' is a circle then we take $\text{Supp}(\tau')$ to be one of the A_j . If τ' is not a circle, then the first return of Φ to $\text{Supp}(\tau')$ is either pseudo-Anosov or finite order; in either case, we take $\text{Supp}(\tau')$ to be one of the S_i . By then restricting the mapping class to the complement of $S - \cup'_\tau \text{Supp}(\tau')$, one continues by induction.

In the end one obtains a decomposition $S = (\cup_i S_i) \cup (\cup_j A_j)$ as desired, and one obtains the pseudo-Anosov and finite order data for each S_i . It remains to describe the computation of the twist data for each A_j . From the computation so far, for each pseudo-Anosov subsurface S_i we have produced an invariant train track, and for each finite order subsurface S_i we have produced an invariant cell decomposition. This data determines the behavior of the f on the boundary of A_j , in particular we obtain periodic orbits on each boundary component. From this we are now able to compute the twist data, from any arc connecting a periodic point in one boundary component to a periodic point in the other boundary component.

There is one last detail: if the surfaces on either side of A_j are of finite order, and if the twist number of A_j equals zero, then one must remove A_j and conglomerate the surfaces on either side of it, for that surface will have a finite order first return.

1.15 Prerequisites

We will assume that the reader is familiar with foundational material about measured foliations and/or measured geodesic laminations, in particular: the topology on \mathcal{MF} and \mathcal{PMF} ; the embeddings $\mathcal{MF} \rightarrow [0, \infty)^{\mathcal{C}}$, $\mathcal{PMF} \rightarrow \mathcal{P}[0, \infty)^{\mathcal{C}}$; continuity of intersection number on $\mathcal{MF} \times \mathcal{MF}$; compactness of \mathcal{PMF} ; Thurston’s boundary of Teichmüller space using \mathcal{PMF} . Most of these items will be reviewed without proof in Section 2. Excellent resources for these materials include [FLP⁺79], [CB88], [CB].

1.16 What’s not here.

Geodesic laminations. A good understanding of measured geodesic laminations on hyperbolic surfaces is extremely useful for understanding measured foliations and train track expansions. The connection between geodesic laminations and measured foliations is explained in [CB88], [CB]. For a survey on geodesic laminations see [Bon01]. Geodesic laminations with transverse Hölder structures are closely related to the structure of the “tangent bundle” of $\mathcal{PMF}(S)$ [Bon97b], [Bon97a].

Despite the close connection between measured geodesic laminations and measured foliations, we rarely make use of this connection, except as occasionally needed to prove some simple results (e.g. Lemma 2.10.1). We have found that by using “partial measured foliations” instead of geodesic laminations, the connection with train track expansions is much tighter and more easily explained.

Structures on \mathcal{MF} and \mathcal{PMF} . Besides not giving a comprehensive account of the theory of measured foliations, other aspects not covered here are the usage of train tracks to describe various important structures on $\mathcal{MF}(S)$ and $\mathcal{PMF}(S)$, e.g. that $\mathcal{PMF}(S)$ is homeomorphic to a sphere of dimension $3|\chi(S)| - \#(\text{punctures})$, and that this sphere carries a natural piecewise integral projective structure and a symplectic structure ([Pen92], [Bon99], [Gol84]).

Dynamics on moduli space. There is a large literature relating measured foliations to the dynamics of interval exchange maps, with emphasis on diophantine properties and unique ergodicity, and with applications to billiard dynamics and to the dynamics of the geodesic flow on moduli space. See for example [Vee82], [Mas82], [Mas92], [Ker85], [KN76]. Explicit connections to train tracks are explained in several of these papers, particularly [Ker85], which was indeed the first paper to clearly describe train track expansions of measured foliations. It seems likely that train track expansions can be used effectively to unify and extend results on the

dynamics of moduli space, particularly in analogy to results on the dynamics of the modular surface $\mathbf{H}^2/\mathrm{PSL}(2, \mathbf{Z})$. One recent such result is joint work of the author and Benson Farb, characterizing bounded trajectories in the moduli space of S in terms of train track expansions [FM99], in analogy to the fact that a bounded geodesic in the modular surface corresponds to a continued fraction with bounded entries.

Teichmüller curves. Given a quadratic differential q with respect to a conformal structure on S , the affine deformations of q define an isometric embedding of the hyperbolic plane into the Teichmüller space $\mathcal{T}(S)$, called the *complex geodesic* $C(q)$ determined by q . The stabilizer of the complex geodesic $C(q)$ in \mathcal{MCG} can be identified with the subgroup $\mathrm{Aff}(q)$ of \mathcal{MCG} that acts affinely on q . The action of $\mathrm{Aff}(q)$ on $C(q)$ is properly discontinuous, and the study of these actions is a rich and interesting pursuit [Vee89], [McM02a], [McM02b]. In particular, Veech studied the situation when $\mathrm{Aff}(q)$ is a lattice acting on $C(q)$, proving many remarkable properties of such actions, and giving new examples. When $\mathrm{Aff}(Q)$ is a lattice, the quotient $C(q)/\mathrm{Aff}(q)$, a finite type Riemann surface, maps holomorphically to Mod with image called a *Teichmüller curve*.

Recently McMullen has shown that if S is closed and of genus 2, and if the quadratic differential q is the square of a holomorphic differential (meaning that the horizontal and vertical foliations of q are orientable), then the limit set of $\mathrm{Aff}(q)$ acting on $C(q)$ is the entire circle at infinity of $C(q)$. Moreover, if q has a unique singular point, of angle 6π , then $\mathrm{Aff}(q)$ is a lattice; whereas if q has two singular points, each of angle 4π , then there are examples where $\mathrm{Aff}(q)$ is not a lattice and hence is infinitely generated.

It would be extremely interesting to find an explicit connection between train track expansions and Teichmüller curves. Consider for example the following problem. Closed geodesics in moduli space are in one-to-one correspondence with pseudo-Anosov conjugacy classes in \mathcal{MCG} , and these are classified in terms of finitistic data, namely “circular expansion complexes” as described in Section 10. Find an algorithm which, given the data describing two pseudo-Anosov conjugacy classes, decides whether the corresponding closed geodesics in moduli space lie on the same Teichmüller curve. The subtlety of this question is evidenced by the existence of complex geodesics whose stabilizer is infinitely generated, that is, Teichmüller curves whose fundamental group is infinitely generated.

1.17 Revision notes

Predecessors. The conjugacy classification of pseudo-Anosov mapping classes that fix a separatrix was given in [Mos83] and in [Mos86]. This work was expanded in

the unfinished monograph “Topological invariants of measured foliations” [Mos93], which contains versions of the Expansion Convergence Theorem, the Arational Expansion Theorem, and the Stable Equivalence Theorem, as well as application to the classification of arational measured foliations up to topological equivalence, and application to the conjugacy classification of pseudo-Anosov mapping classes that fix a separatrix.

Version 1. The first version of this monograph was entitled “Train track expansions of arational measured foliations”.

Version 2. Section 7, on stable equivalence of expansion complexes, is new to this version, although I had it in the back of my mind for many years.

Section 8 on the structure of general splitting sequences is entirely new—all the theorems in this section were proved during the writing. They was motivated by the desire to remove the word “arational” from the title of the monograph, but even more by the desire to close a gap in the dictionary, as explained at the beginning of the section.

Section 9.4, on constructing pseudo-Anosov mapping classes using one switch train tracks, is a recasting in train track language of results from [Mos86].

There were several changes in terminology from the previous version. Most significant is the terminology for the criterion that tells when a splitting sequence is an expansion of an arational measured foliation. This was called the “strong filling criterion” in the previous version, and is now called the “iterated rational killing criterion”; more of a mouthful, but hopefully more descriptive. Also, a train track whose complementary components are nonpunctured and once-punctured discs is now called “filling” rather than “weakly filling”; I had in mind to describe a stronger version of the filling property, but eventually the stronger property seemed not sufficiently important to warrant an extra adjective in the terminology for the weaker property, and I decided to restrict the discussion of the stronger property to some remarks after the statement of Proposition 3.13.2.

A few sections from Version 1 were removed: one on transverse recurrence; and one on a folk construction of pseudo-Anosov mapping classes.

Version 3. This version contains substantial rewritings of Section 7.

Section 10 is new. The result on using one cusp splitting circuits to classify pseudo-Anosov mapping classes fixing a separatrix is a recasting in train track language of results from [Mos86]. The result on using circular expansion complexes to give a complete pseudo-Anosov conjugacy classification is new, although as with Section 7 I have had this in the back of my mind for years.

Future versions planned. Further revision of Section 10 will incorporate some of the detailed information on enumeration of pseudo-Anosov conjugacy classes which is contained in the unfinished monograph [Mos93]. Section 12 will give a practical solution to the conjugacy problem in \mathcal{MCG} .

1.18 Acknowledgements.

The immediate need for this paper, as well as the form of the Arational Expansion Theorem 6.3.2 and certain other results, were inspired by the author's joint work with Benson Farb on precompact geodesics in moduli space, which convinced me of the need to reformulate the combinatorial approach of [Mos86] and [Mos93] into a more directly accessible approach using train tracks.

I would also like to thank readers of the first version for comments which were very useful in preparing the second version, in particular Moon Duchin and Pallavi Dani.

2 Measured foliations

In this preliminary section we review the theory of measured foliations, often referring the reader to other sources for proofs, in particular: [FLP⁺79], [CB88], [CB].

2.1 Surfaces of finite type

A *surface of finite type* $S = \bar{S} - P$ is a closed, oriented, connected surface \bar{S} minus a finite set $P \subset \bar{S}$ called the *punctures*. The surface \bar{S} is called the *filled in surface*. Isotopies of S are identified with isotopies of \bar{S} leaving P stationary. Although we almost never use the notation \bar{S} , this will often entail an abuse of notation. For example, when we speak of a “measured foliation on S ” we really mean a measured foliation on the closed surface \bar{S} with certain allowance for singularities at the punctures.

Consider a simple closed curve c on S . We say c is *trivial* if c bounds a non-punctured disc in S , and c is *peripheral* if c bounds a once-punctured disc in S . An *essential curve on S* is a nontrivial, nonperipheral simple closed curve in S . Let $\mathcal{C} = \mathcal{C}(S)$ be the set of isotopy classes of essential curves in S .

Throughout this paper we shall fix a connected surface of finite type $S = \bar{S} - P$ such that $\chi(S) < 0$ and S is not a thrice punctured sphere. With these conditions it follows that \mathcal{C} is countably infinite.

2.2 Teichmüller space \mathcal{T}

A *hyperbolic structure* on S means a complete, finite area Riemannian metric on S of constant curvature -1 . A *conformal structure* on S means a conformal structure with removable singularities at the punctures, thereby extending to a conformal structure on the filled in surface \bar{S} . With the assumption that $\chi(S) < 0$, the Uniformization Theorem produces a one-to-one correspondence between hyperbolic structures and conformal structures, natural with respect to isotopy. Let $\mathcal{T} = \mathcal{T}(S)$ be the *Teichmüller space* of S , the set of conformal structures modulo isotopy, or equivalently the set of hyperbolic structures modulo isotopy. The assumption that S is not a thrice punctured sphere implies that \mathcal{T} is not a single point.

Given a hyperbolic structure on S , every essential curve is isotopic to a unique geodesic whose length depends only on the isotopy classes of the curve and the hyperbolic structure, thereby defining a map $\mathcal{L}: \mathcal{T} \rightarrow [0, \infty)^{\mathcal{C}}$. Letting $\mathcal{P}: [0, \infty)^{\mathcal{C}} \rightarrow \mathcal{P}[0, \infty)^{\mathcal{C}}$ be projectivization, the composition $\mathcal{P} \circ \mathcal{L}: \mathcal{T} \rightarrow \mathcal{P}[0, \infty)^{\mathcal{C}}$ is an embedding, and it induces a topology and smooth structure making \mathcal{T} diffeomorphic to Euclidean space of dimension $6g - 6 + 2p$.

The Thurston Compactification Theorem for \mathcal{T} says that the closure of $\mathcal{P} \circ \mathcal{L}(\mathcal{T})$

in $\mathcal{P}[0, \infty)^{\mathcal{C}}$ is a closed ball of dimension $6g - 6 + 2p$ whose interior is $\mathcal{P} \circ \mathcal{L}(\mathcal{T})$. The boundary points of this ball are usually described in one of two ways: using measured geodesic laminations, or using measured foliations. We shall adopt the latter method, which is better designed for working with Teichmüller geodesics. However, we shall extend the class of measured foliations to include “partial measured foliations”, which have some of the flexibility of geodesic laminations.

2.3 The mapping class group

Let $\text{Homeo}(S)$ be the group of homeomorphisms of S , and let $\text{Homeo}_0(S)$ be the normal subgroup of homeomorphisms isotopic to the identity on S . The *mapping class group* is $\mathcal{MCG} = \mathcal{MCG}(S) = \text{Homeo}(S)/\text{Homeo}_0(S)$. There is a natural action of \mathcal{MCG} on the Teichmüller space \mathcal{T} . Indeed, as we define other natural topological and geometric structures on surfaces on which $\text{Homeo}(S)$ acts, the action of $\text{Homeo}_0(S)$ induces an equivalence relation on such structures which coincides with isotopy, and \mathcal{MCG} acts naturally on the set of isotopy classes. Examples include the actions of \mathcal{MCG} on the isotopy classes of measured foliations \mathcal{MF} , and on the isotopy classes of quadratic differentials QD.

The action of \mathcal{MCG} on \mathcal{T} is properly discontinuous, and the quotient space is naturally a smooth orbifold called the *moduli space* of S , denoted $\mathcal{M} = \mathcal{M}(S) = \mathcal{T}/\mathcal{MCG}$.

2.4 Quadratic differentials and their horizontal and vertical foliations

The singularities of measured foliations are locally modelled on singularities of horizontal foliations of quadratic differentials. We shall therefore pause for a brief foray into quadratic differentials, which will be covered in more detail later.

If we give S a conformal structure, a *meromorphic quadratic differential* q on S is an expression $q = f(z) dz^2$ for each local coordinate z , such that $f(z)$ is meromorphic, and such that for any overlap map $z' \mapsto z$ between local coordinates, if $q = f(z) dz^2$ in the z coordinate, then in the z' coordinate we have $q = f(z(z')) \left(\frac{dz}{dz'}\right)^2 dz'^2$. Zeroes and poles of q are well-defined, as is the order of a zero or pole. q determines an area form, expressed in a coordinate z as $|f(z)| |dz|^2$, and the total area is denoted $\|q\|$. Note that q is integrable, meaning that $\|q\|$ is finite, if and only if q has a meromorphic extension to the filled in surface \bar{S} so that each pole has degree 1. In studying the Teichmüller space of a Riemann surface S of finite type, the most relevant quadratic differentials on S are those which are holomorphic and integrable. For the rest of the paper, therefore, a *quadratic differential* on S is assumed to have these properties.

Consider a quadratic differential q . Near any regular point p there is a local coordinate z so that $q = dz^2$; the germ of this coordinate is unique up to transformations of the form $z \mapsto \pm z + c$. Near any singular point p there is a local coordinate z , taking p to the origin of the z -plane, in which $q = z^n dz^2$, such that either $n = -1$ and p is a pole of q , or $n \geq 1$ and p is a zero of order n ; the coordinate z is unique up to multiplication by an $(n+2)^{\text{nd}}$ root of unity (completely unique in the case $n = -1$). In any of these cases we refer to z as a *canonical coordinate* for q near p .

A quadratic differential q determines two singular transversely measured foliations, called the *horizontal* and *vertical* foliations, denoted $\mathcal{F}^h(q)$ and $\mathcal{F}^v(q)$. These foliations have the same singular set, namely the set of zeroes and poles of q . Near a nonsingular point p with canonical coordinate $z = x + iy$, horizontal leaf segments are parallel to the x -axis and the transverse measure on $\mathcal{F}^h(q)$ is defined by integration of $|dy|$, while vertical leaf segments are parallel to the y -axis with transverse measure defined by integrating $|dx|$. These measured foliations are well-defined, because the canonical coordinate near a regular point is well defined up to transformations of the form $z \mapsto \pm z + c$.

To visualize the measured foliations $\mathcal{F}^h(q)$ and $\mathcal{F}^v(q)$ near a singular point p with canonical coordinate z , we have $q = z^{k-2} dz^2$ for some $k \geq 1$; either $k \geq 3$ and p is a zero of order $k-2$, or $k = 1$ and p is a pole of order 1. Near any $z \neq 0$ an easy calculation shows that $z' = z^{k/2}$ is a canonical coordinate for $z^{k-1} dz^2$, using either choice of the square root in the case where k is odd. Horizontal and vertical leaves of $z^{k-2} dz^2$ can therefore be described as follows. For each $i \in \mathbf{Z}/k\mathbf{Z}$ let S_i be the angular sector of \mathbf{C} defined in polar coordinates by $2\pi i/n \leq \theta \leq 2\pi(i+1)/n$. Let $z' = x' + iy'$ vary over the closed upper half plane of \mathbf{C} , and let $z'^{2/k}$ denote the $2/k$ th root of z' taking values in the angular sector S_0 . For each $i \in \mathbf{Z}/k\mathbf{Z}$, the map $z' \xrightarrow{\phi_i} z = \exp(2\pi i/k) z'^{2/k}$ takes values in the sector S_i , the horizontal leaves of $z^{k-2} dz^2$ in S_i are the images under the map ϕ_i of lines parallel to the x' -axis in the closed upper half z' plane, and the transverse measure is obtained by pushing forward $|dy'|$; a similar discussion holds for vertical leaves. The ray $R_i = S_{i-1} \cap S_i$, whose angle is $\theta = 2\pi i/n$, is a leaf of the horizontal foliation called a *separatrix attached to the origin*. The set of separatrices, indexed by $\mathbf{Z}/k\mathbf{Z}$, has a natural circular ordering. By convention, a regular point of a foliation will sometimes be called a *2-pronged singularity*, especially when that point is a puncture on a surface of finite type.

2.5 Measured foliations

The general concept of a measured foliation on a surface of finite type is obtained by using the horizontal measured foliation of a quadratic differential as a local model.



Figure 3: A k -pronged singularity ($k = 3$), modelled on the horizontal foliation of $z^{k-2} dz^2$ near the origin of \mathbf{C} . Leaf segments are pre-images, under a transformation $z' = z^{k/2}$ (with any choice of the root), of horizontal segments in the z' plane. Each ray passing through a k^{th} root of unity is a leaf, called a *separatrix*. The region between two separatrices is foliated like the horizontal foliation of the upper half-plane.

For surfaces with boundary, we need two other local models to define measured foliations tangent to the boundary: a *regular tangential boundary point* is locally modelled on the horizontal foliation of the closed upper half plane of \mathbf{C} , near the origin; and a *k -pronged tangential boundary singularity* with $k \geq 3$ is locally modelled on the horizontal foliation of $z^{2k-2} dz^2$ near the origin in the closed upper half plane of \mathbf{C} , with 2 separatrices on the boundary and $k - 2$ separatrices pointing into the interior.

Consider now a finite type surface-with-boundary $F = \bar{F} - P(F)$, meaning that \bar{F} is a compact, oriented surface-with-boundary and $P(F)$ is a finite subset of $\text{int}(F)$. A *measured foliation* on F is a foliation \mathcal{F} on \bar{F} with finitely many prong singularities $\text{sing}(\mathcal{F})$ and with a positive transverse Borel measure, such that each singularity of \mathcal{F} in $\text{int}(\bar{F}) - P$ is a k -pronged singularity for some $k \geq 3$, each puncture is a k -pronged singularity for some $k \geq 1$, and each point of ∂F is either a regular tangential boundary point or a k -pronged tangential boundary singularity for some $k \geq 3$. By convention every puncture of \bar{F} is considered to be a singularity, even if it is a 2-pronged singularity, so $P \subset \text{sing}(\mathcal{F})$. The restriction $\mathcal{F} | F - \text{sing}(\mathcal{F})$ is a true foliation; it is locally modelled on the horizontal foliation of \mathbf{R}^2 with overlap maps of the form $(x, y) \mapsto (f(x, y), \pm y + c)$, and so the transverse measure is defined in local models as integration of $|dy|$. The interested reader can work out the form of overlap maps between regular and singular models for \mathcal{F} , and between two singular models.

Given a leaf ℓ of $\mathcal{F} | \bar{F} - \text{sing}(\mathcal{F})$, let $\bar{\ell}$ be obtained from ℓ by adding any singularity of \mathcal{F} which is the limit point of some end of ℓ ; we define $\bar{\ell}$ to be a *nonsingular leaf* of \mathcal{F} . There are several types of nonsingular leaves: a *bi-infinite leaf* is homeomorphic to \mathbf{R} ; a *closed leaf* is a circle containing no singularities; an

infinite separatrix is homeomorphic to $[0, \infty)$ with one endpoint at a singularity; and a *saddle connection* is a compact leaf containing a singularity, either an arc with endpoints at distinct singularities, or a circle containing one singularity. A *nonsingular leaf segment* is a compact segment contained in a nonsingular leaf.

We still have not defined what we mean by a leaf of \mathcal{F} , unadorned with any qualifiers. Consider a compact segment α which is a union of nonsingular leaf segments. Given a transverse orientation V on α , we say that α is *perturbable* in the direction of V if there exists an embedding $f: \alpha \times [0, 1] \rightarrow \bar{F}$ such that $f(\alpha \times 0) = \alpha$, the oriented segment $f(x \times [0, 1])$ represents the transverse orientation V at the point $f(x) \in \alpha$, and for each $t \in (0, 1]$ the set $f(\alpha \times t)$ is a nonsingular leaf segment. We say that α is a *leaf segment* if it has a perturbable transverse orientation. For example, if α is a nonsingular leaf segment then both transverse orientations are perturbable and in particular α is a leaf segment. More generally, perturbability means that at a singularity $s \in \text{int}(\alpha)$, α contains two separatrices incident to the same sector at s , and V points into that sector. Finally, a *leaf* ℓ of F is the embedded image of either \mathbf{R} or S^1 , whose image is a union of nonsingular leaves, such that for some transverse orientation V on ℓ , the restriction of V to every compact subsegment of ℓ is perturbable. Notice that if $\ell \approx \mathbf{R}$ then ℓ is *never* globally perturbable in either direction; this uses the existence of a positive transverse measure on \mathcal{F} .

A *finite separatrix* of \mathcal{F} is a nonsingular leaf segment ℓ with one end at a singularity $s \in \text{sing}(\mathcal{F})$. More precisely we say that ℓ is a *separatrix located at* s . If s is an n -pronged singularity then there are n distinct separatrix germs at s . Each infinite separatrix represents a unique separatrix germ, each saddle connection represents exactly two separatrix germs, and every separatrix germ is uniquely represented either by an infinite separatrix or by a saddle connection.

A *leaf cycle* of \mathcal{F} is a union of saddle connections which forms either an embedded circle in \bar{F} , or an embedded arc in \bar{F} not lying in ∂F each of whose endpoints is either a puncture or a boundary singularity. Note that if \mathcal{F} has a closed leaf then \mathcal{F} has a leaf cycle; the converse is not true, because a leaf cycle need not have a perturbable transverse orientation.

Suppose that \mathcal{F} is a measured foliation on F and α is a saddle connection which is not a leaf cycle (so α is an embedded arc, and either $\alpha \subset \partial F$ or at least one endpoint is neither a puncture nor a boundary singularity). Collapsing α to a point induces a new measured foliation on F which we say is obtained from \mathcal{F} by a *Whitehead collapse*. The equivalence relation on the set of measured foliations on F generated by Whitehead collapse and isotopy is called *Whitehead equivalence*.

The term “measured foliation” without any other qualifiers will always imply, as defined above, that each boundary point is a regular tangential boundary point or a tangential boundary singularity. We shall have a limited need for *boundary transverse measured foliations* of F as well. These have the same local models in

$F - \partial F$, but on the boundary the local models for *regular* and *k-pronged transverse* boundary points are the vertical foliations of dz^2 and $z^{2k-2} dz^2$, respectively, restricted to the closed upper half plane, near the origin. The only place where we make significant use of boundary transverse foliations is in the tie bundle over a train track, in Section 3.

The definition of a leaf of \mathcal{F} that we have adopted in this section behaves well under Whitehead equivalence. That is, if $\mathcal{F}, \mathcal{F}'$ are Whitehead equivalent measured foliations then there is a natural bijection between leaves of \mathcal{F} and leaves of \mathcal{F}' . This is obvious when $\mathcal{F}, \mathcal{F}'$ are isotopic. It is not hard to check for a Whitehead collapse from \mathcal{F} to \mathcal{F}' : it follows immediately from the observation that if α is a segment consisting of a union of leaf segments of \mathcal{F} , and if α' is the image of α , then α' is a union of leaf segments of \mathcal{F}' , and α has a perturbable transverse orientation if and only if α' does.

Also, the definition of a leaf of \mathcal{F} is consistent with the concept of a leaf of an equivalent measured geodesic lamination λ , in the following sense. Recall that equivalence means that there exists a map $\pi: (S, \lambda) \rightarrow (S, \mathcal{F})$ that collapses components of $S - \lambda$ to leaf cycles of \mathcal{F} , takes leaf segments of λ to leaf cycles of \mathcal{F} , and pushes the transverse measure on λ forward to the transverse measure on \mathcal{F} . Under these circumstances, π takes a leaf segment of λ to segment which is a union of leaf segments of \mathcal{F} and which has a perturbable transverse orientation, and hence π takes leaves of λ bijectively to leaves of \mathcal{F} .

Arational measured foliations. A measured foliation \mathcal{F} on F is *arational* if every leaf cycle is a component of ∂F . Equivalently, for every essential, nonperipheral simple closed curve c , the intersection number of \mathcal{F} with c is nonzero. Arationality is invariant under Whitehead equivalence.

In the case where F has no punctures and no boundary, arationality of \mathcal{F} means that there are no leaf cycles at all; equivalently, the union of saddle connections of \mathcal{F} is a disjoint union of trees. When F has punctures, arationality of \mathcal{F} means that the union of saddle connections is a disjoint union of trees, *and* each of these trees contains at most one puncture. When F has punctures and boundary, arationality of \mathcal{F} means that for each component C of the union of saddle connections, either C is a tree containing at most one puncture, or C contains a component of ∂F as a deformation retract and C contains no puncture at all.

A property related to arationality is *minimality*, which means that each leaf of \mathcal{F} is dense. Each arational measured foliation \mathcal{F} is minimal: if there is a nondense leaf ℓ then its closure is a proper 2-complex in S , whose frontier is a union of leaf segments of \mathcal{F} . On a torus with ≤ 1 puncture the converse is true: the arational measured foliations are the same as the minimal measured foliations, and they are

precisely the ones isotopic to a constant irrational slope foliation. However, if S is a more complicated surface than a punctured torus, then S has a minimal measured foliation which is not arational: starting with an essential, nonperipheral, nonseparating simple closed curve c on S , let S' be the connected surface obtained by cutting S open along c , let \mathcal{F}' be any minimal measured foliation on S' , and by regluing we obtain a minimal, nonarational measured foliation \mathcal{F} on S . This construction works because S' is not a three-punctured sphere and therefore supports an arational measured foliation.

In general, as we shall see in §2.9 below, every measured foliation \mathcal{F} on S breaks naturally into foliation components, each component being supported on a canonical subsurface. \mathcal{F} is minimal if and only if it has a single foliation component which is arational; furthermore \mathcal{F} is arational if and only if it has a single foliation component which fills the entire surface S . A general measured foliation can have a mixture of “arational” and “rational” foliation components.

Quadratic differentials and measured foliations. We have defined measured foliations based on a local analogy with quadratic differentials: the local models for a general measured foliation are the same as for the horizontal measured foliation of a quadratic differential. However, there is a global analogy as well, due to a result of Gardiner and Masur [GM91] which we recall. It is not quite true that every measured foliation is equal to $\mathcal{F}^h(q)$ for some quadratic differential q , but it is almost true.

A measured foliation \mathcal{F} is *taut* if for every nonsingular point x there exists a simple closed curve c transverse to \mathcal{F} and passing through x . For instance, if \mathcal{F} is arational then it is taut.

Theorem 2.5.1 ([GM91]). *A measured foliation \mathcal{F} is taut if and only if there exists a quadratic differential q such that \mathcal{F} is isotopic to $\mathcal{F}^h(q)$. Moreover, every measured foliation is Whitehead equivalent to a taut measured foliation.*

The proof of this theorem in [GM91] gives an explicit procedure for performing a sequence of Whitehead collapses and their inverses, to produce a taut measured foliation. We give another proof at the end of Section 2.9, using ideas of partial measured foliations.

2.6 Partial measured foliations

For various reasons it is useful to study measured foliations supported on proper subsurfaces of the filled in surface \bar{S} . Such “partial” measured foliations serve many of the same purposes as measured geodesic laminations. Partial measured foliations have some advantages over geodesic laminations, in particular by being more closely

related to quadratic differentials and train tracks. There are disadvantages as well, namely the complexity of the equivalence relation among partial measured foliations.

Consider $S = \bar{S} - P$, our fixed surface of finite type. Consider a finite type subsurface $\bar{F} \subset \bar{S}$ such that $P(F) = \bar{F} \cap P \subset \text{int}(F)$. Let \mathcal{F} be a measured foliation on F . Given a boundary singularity s of \mathcal{F} , a separatrix of \mathcal{F} at s is *proper* if it is not contained in ∂F . Define the *local index* of s relative to \mathcal{F} to be

$$\iota_{\mathcal{F}}(s) = -\frac{1}{2} \cdot \#\{\text{proper separatrices at } s\}$$

Given a component C of $\text{Cl}(S - F)$, the *index* of C relative to \mathcal{F} , denoted $\iota_{\mathcal{F}}(C)$, is defined to be

$$\iota_{\mathcal{F}}(C) = \chi(C) + \sum_{s \in \partial C} \iota(s)$$

We say that \mathcal{F} is a *partial measured foliation on S with support* $\text{Supp}(\mathcal{F}) = F$ if for each component C of $\text{Cl}(S - F)$ we have $\iota_{\mathcal{F}}(C) \leq 0$, and if C contains a puncture then $\iota_{\mathcal{F}}(C) < 0$. When C is a once-punctured disc it follows that \mathcal{F} has at least one boundary singularity on ∂C , and so at least one proper separatrix on ∂C ; whereas if C is a nonpunctured disc then \mathcal{F} has at least two proper separatrices at boundary singularities on ∂C . This requirement on C is analogous to the requirement that for a measured foliation on S , there must be at least 1 prong at each puncture and at least 2 prongs at each regular point.

We will very occasionally use the term *total measured foliation* on S to refer to a partial measured foliation \mathcal{F} for which $\text{Supp}(\mathcal{F}) = S$.

We now describe a “fulfillment” relation which from a partial measured foliation produces a measured foliation, and an inverse “slicing” relation which from a measured foliation produces a partial measured foliation.

A map $\rho: \gamma \rightarrow \Theta$ from a 1-manifold γ to a 1-complex Θ is called a *folded immersion* if each $x \in \gamma$ has a neighborhood $U \subset \gamma$ such one of the following holds: $\rho|_U$ is an embedding; or, $\rho(x)$ is a valence 1 vertex of Θ , $\rho(U)$ is a neighborhood of $\rho(x)$ homeomorphic to $[0, \infty)$, and $\rho|_U: (U, x) \rightarrow (\rho(U), \rho(x))$ is conjugate to the absolute value function $|\cdot|: (\mathbf{R}, 0) \rightarrow ([0, \infty), 0)$; or, $\rho(x)$ is an isolated point of Θ and hence $\rho(U) = \rho(x)$.

Consider two partial measured foliations $\mathcal{F}, \mathcal{F}'$ on S . A *partial fulfillment map* from \mathcal{F}' to \mathcal{F} , denoted $r: (S, \mathcal{F}') \rightarrow (S, \mathcal{F})$ or simply $r: \mathcal{F}' \rightarrow \mathcal{F}$ when S is understood, is a continuous map $r: (\bar{S}, P) \rightarrow (\bar{S}, P)$ such that, for some 1-complex $\Theta \subset \bar{S}$ which is a finite union of isolated points and nonboundary leaf segments of \mathcal{F} , we have:

- (1) $r: (\bar{S}, P) \rightarrow (\bar{S}, P)$ is homotopic to the identity rel P .

- (2) The restriction $r \mid (\text{Supp}(\mathcal{F}') - \partial \text{Supp}(\mathcal{F}'))$ is a homeomorphism onto $\text{Supp}(\mathcal{F}) - (\partial \text{Supp}(\mathcal{F}) \cup \Theta)$.
- (3) The pushforward via r of the measured foliation $\mathcal{F}' \mid (\text{Supp}(\mathcal{F}') - \partial \text{Supp}(\mathcal{F}'))$ is $\mathcal{F} \mid (\text{Supp}(\mathcal{F}) - (\Theta \cup \partial \text{Supp}(\mathcal{F})))$.
- (4) The restriction $r \mid \text{Cl}(S - \text{Supp}(\mathcal{F}'))$ is a homotopy equivalence onto $\Theta \cup \text{Cl}(S - \text{Supp}(\mathcal{F}))$.
- (5) The restriction $r \mid \partial \text{Supp}(\mathcal{F}')$ is a folded immersion onto $\Theta \cup \partial \text{Supp}(\mathcal{F})$.

When \mathcal{F} is a measured foliation then we say that \mathcal{F} is a *fulfillment* of \mathcal{F}' , and $r: (S, \mathcal{F}') \rightarrow (S, \mathcal{F})$ is a *fulfillment map*.

Fact 2.6.1. *Every partial measured foliation \mathcal{F}' has a fulfillment \mathcal{F} .*

Proof. Choose Θ to be a 1-dimensional spine of $\text{Cl}(S - \text{Supp}(\mathcal{F}'))$ containing every puncture of $\text{Cl}(S - \text{Supp}(\mathcal{F}'))$. Choose $r: \text{Cl}(S - \text{Supp}(\mathcal{F}')) \rightarrow \Theta$ to be a deformation retraction whose restriction to $\partial \text{Cl}(S - \text{Supp}(\mathcal{F}')) = \partial \text{Supp}(\mathcal{F}')$ satisfies item 5, and make the choices so that for each $x \in \partial \text{Supp}(\mathcal{F}')$, if r folds x around the valence 1 vertex $f(x)$ of Θ then either $f(x)$ is a puncture or x is a boundary singularity of \mathcal{F}' , and if $f(x)$ is an isolated point of Θ then $f(x)$ is a puncture. Extend r over S to satisfy items 1, 2. Now define \mathcal{F} by pushing forward \mathcal{F}' , so as to satisfy item 3. \diamond

Conversely,

Fact 2.6.2 (and definition of Slicing). *Given a measured foliation \mathcal{F} and a finite union Θ of leaf segments of \mathcal{F} , there exists a partial measured foliation \mathcal{F}' and a fulfillment map $r: (S, \mathcal{F}') \rightarrow (S, \mathcal{F})$ such that $r(\text{Cl}(S - \text{Supp}(\mathcal{F}'))) = \Theta$. We say that \mathcal{F}' is obtained from \mathcal{F} by slicing along Θ .*

Proof. Choose N to be a closed regular neighborhood of Θ . Choose $r: N \rightarrow \Theta$ to be a deformation retraction whose restriction to ∂N is a folded immersion. Extend continuously to a map $r: (S, P) \rightarrow (S, P)$ which is homotopic to the identity rel P and which restricts to a homeomorphism from $S - N$ to $S - \Theta$. Define \mathcal{F}' to be the unique partial measured foliation supported on $\text{Cl}(S - N)$ whose restriction to $S - N$ pushes forward via r to $\mathcal{F} \mid S - \Theta$. \diamond

Note in Fact 2.6.2 that the restriction of r to $\text{Supp}(\mathcal{F}')$ is uniquely determined up to isotopy, and so slicing \mathcal{F} along Θ determines a unique partial measured foliation \mathcal{F}' up to isotopy.

As we've defined these concepts, fulfillment and slicing are not quite inverses. Under slicing, Θ is a union of leaf segments, whereas under fulfillment, Θ is a union

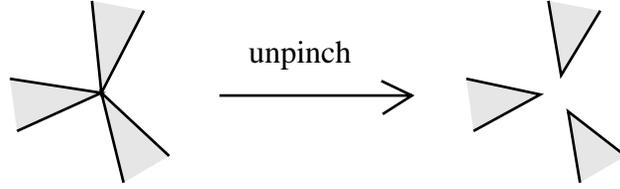


Figure 4: Unpinching a pinch point.

of isolated points and leaf segments. So if \mathcal{F} is obtained from \mathcal{F}' by fulfillment then it seems that \mathcal{F}' might not be obtained from \mathcal{F} by slicing. However, it is not hard to show that if $r: (S, \mathcal{F}') \rightarrow (S, \mathcal{F})$ is a fulfillment map with $\Theta = r(S - \text{Supp}(\mathcal{F}'))$, then the isolated components of Θ can be augmented by adding on short separatrices, and r altered by homotopy rel P , so that \mathcal{F}' is indeed obtained from \mathcal{F} by slicing.

We will also occasionally need a variation on partial measured foliations which allows for “pinched” supports (see Figure 4). A *closed angle* in the complex plane \mathbf{C} is a subset of the form

$$A[\theta_1, \theta_2] = \{re^{2\pi i\theta} \mid \theta_1 \leq \theta \leq \theta_2\}$$

for real numbers θ_1, θ_2 with $0 < \theta_2 - \theta_1 < \pi$. A *pinched subsurface* of S is a smooth, closed 2-complex $F \subset S$, such that $F \cap P \subset \text{int}(F)$, and F is a 2-manifold with boundary except at finitely many *pinch points* at which F is locally modelled on a finite union of closed angles in \mathbf{C} which are pairwise disjoint except at the origin; each such angle is called a *sector* of the pinch point. There is a subsurface F' obtained from F by *unpinching*, which is defined locally near a pinch point by pulling the sector $A[\theta_1, \theta_2]$ away from the origin, replacing it by a set of the form

$$A'[\theta_1, \theta_2] = \{re^{2\pi i\theta} \mid \theta_1 \leq \theta \leq \theta_2, \quad r \geq f(\theta)\}$$

where $f: [\theta_1, \theta_2] \rightarrow (0, \infty)$ is a C^1 concave function (see Figure 4). There is a *pinching map* $r: F' \rightarrow F$ homotopic to inclusion, one-to-one except that over a pinch point p the map r is n -to-one where n is the number of sectors of p .

Consider now a pinched subsurface F , an unpinching F' , a pinch map $r: F' \rightarrow F$, and a partial measured foliation \mathcal{F} with support F' . The pushforward $r(\mathcal{F})$ is defined to be a *pinched partial measured foliation* with support F , and \mathcal{F} is an *unpinching* of $r(\mathcal{F})$. We will generally consider a pinched partial measured foliation to be equivalent to any unpinching; note that the unpinching is well-defined up to isotopy.

Earlier we discussed that the concept of a leaf of a measured foliation behaves well under Whitehead equivalence. A similar discussion shows that leaves of partial

measured foliations behave well under slicing and fulfillment, as the reader may easily check. It follows that if the partial measured foliations $\mathcal{F}, \mathcal{F}'$ are related by a sequence of isotopies, Whitehead moves, and fulfillment operations and their inverses, then there is a natural bijection between the leaves of \mathcal{F} and the leaves of \mathcal{F}' .

2.7 Measured foliation space \mathcal{MF} , and its projectivization \mathcal{PMF}

Consider the set consisting of all measured foliations and partial measured foliations on S , with equivalence relation generated by isotopy, Whitehead equivalence, and fulfillment. This is identical to the equivalence relation generated by isotopy, Whitehead equivalence, and partial fulfillment, because a partial fulfillment $\mathcal{F}_1 \rightarrow \mathcal{F}_2$ can be replaced by two fulfillments $\mathcal{F}_1 \rightarrow \mathcal{F} \leftarrow \mathcal{F}_2$. The set of equivalence classes is denoted \mathcal{MF} , the *space of measured foliations* on S .

We also define *unmeasured equivalence* of measured foliations, where $\mathcal{F}, \mathcal{F}'$ are unmeasured equivalent if the underlying unmeasured partial foliations are equivalent up to isotopy, Whitehead equivalence, and fulfillment.

Our definition of \mathcal{MF} is equivalent to the traditional definition, given in [FLP⁺79], which is the set of equivalence classes of total measured foliations up to isotopy and Whitehead equivalence. To prove this equivalence one must show that if two total measured foliations $\mathcal{F}, \mathcal{F}'$ are related by a sequence of isotopies, Whitehead moves, partial fulfillments and their inverses, then $\mathcal{F}, \mathcal{F}'$ are related by a sequence of isotopies and Whitehead moves. One easily reduces to the case of a sequence $\mathcal{F} = \mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_3 = \mathcal{F}'$ where $\mathcal{F}_1, \mathcal{F}_2$ are partial measured foliations, there are partial fulfillments $\mathcal{F}_1 \rightarrow \mathcal{F}_0$ and $\mathcal{F}_2 \rightarrow \mathcal{F}_3$, and there is a Whitehead move from \mathcal{F}_1 to \mathcal{F}_2 . By slicing open $\mathcal{F}_1, \mathcal{F}_2$ even further one can obtain isotopic partial measured foliations $\mathcal{F}'_1, \mathcal{F}'_2$, thereby making a further reduction to the case of a sequence $\mathcal{F} = \mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2 = \mathcal{F}'$ where there are partial fulfillments $\mathcal{F}_1 \rightarrow \mathcal{F}_0$ and $\mathcal{F}_1 \rightarrow \mathcal{F}_2$. These partial fulfillments are defined by choosing two different spines of the surface $\mathcal{C}(S - \text{Supp}(\mathcal{F}_1))$ and collapse maps onto these two spines. It is not hard to see that two spines of the same surface are connected by a sequence of Whitehead moves, and that two collapse maps onto the same spine are similarly connected.

Scalar multiplication of a transverse measure by a positive constant induces a well-defined scalar multiplication $(0, \infty) \times \mathcal{MF} \rightarrow \mathcal{MF}$; the quotient of \mathcal{MF} by the action of scalar multiplication is denoted \mathcal{PMF} .

Let \mathcal{F} be a partial measured foliation and let c be an essential closed curve. The intersection number $\langle \mathcal{F}, c \rangle$ is defined by integrating the transverse measure along c . If $[\mathcal{F}]$ is the class of \mathcal{F} in \mathcal{MF} and $[c]$ is the class of c in \mathcal{C} , define $\langle [\mathcal{F}], [c] \rangle$ to be the infimum of $\langle \mathcal{F}', c' \rangle$ taken over $\mathcal{F}' \in [\mathcal{F}]$ and $c' \in [c]$. The infimum is always realized by any representatives $\mathcal{F}' \in [\mathcal{F}], c' \in [c]$ such that c' intersects \mathcal{F}' transversely and

efficiently, meaning that there is no closed disc in $\text{Cl}(S - (P \cup \text{Supp}(\mathcal{F}')))$ whose boundary is a union of a segment of c' and a nonsingular leaf segment of $\text{Supp}(\mathcal{F}')$; for such representatives, $\langle [\mathcal{F}], [c] \rangle$ is nonzero when $c' \cap \text{Supp}(\mathcal{F}') \neq \emptyset$, and it is zero otherwise. We obtain a map $I: \mathcal{MF} \rightarrow [0, \infty)^{\mathcal{C}}$, where $I(\mathcal{F})(c) = \langle \mathcal{F}, c \rangle$. This map is an injection, and it induces an injection $\mathcal{PMF} \rightarrow \mathcal{P}[0, \infty)^{\mathcal{C}}$.

The images of \mathcal{T} and \mathcal{PMF} in $\mathcal{P}[0, \infty)^{\mathcal{C}}$ are disjoint, and we have the following theorem of Thurston (see [Thu88], [FLP⁺79], or [CB88]):

Theorem 2.7.1. *The embedding $\overline{\mathcal{T}} = \mathcal{T} \cup \mathcal{PMF} \rightarrow \mathcal{P}[0, \infty)^{\mathcal{C}}$ has image homeomorphic to a closed ball of dimension $6g - 6 + 2p$, with interior \mathcal{T} and with boundary \mathcal{PMF} .*

2.8 Canonical models for partial measured foliations

A partial measured foliation \mathcal{F} has *canonical singularities* if every singularity is a 3-pronged boundary singularity; in particular, $P \cap \text{Supp}(\mathcal{F}) = \emptyset$. If \mathcal{F} has canonical singularities, and if C is a component of $\text{Cl}(S - \text{Supp}(\mathcal{F}))$, we say that C is a *complementary bigon* if C is a nonpunctured disc with exactly two singularities on its boundary, and C is a *complementary smooth annulus* if C is a nonpunctured annulus with no singularities on its boundary.

A partial measured foliation \mathcal{F} is *canonical*, also called a *canonical model*, if \mathcal{F} has canonical singularities, no proper saddle connections, no complementary bigons, and no complementary smooth annuli.

Proposition 2.8.1 (Canonical models). *Every class in \mathcal{MF} is represented by a canonical partial measured foliation, unique up to isotopy. Moreover, isotopies of canonical models are canonical, in the following sense: if $\mathcal{F}, \mathcal{F}'$ are canonical models, if $f, g: S \rightarrow S$ are isotopic homeomorphisms, and if $f(\mathcal{F}) = g(\mathcal{F}) = \mathcal{F}'$, then for each leaf ℓ of \mathcal{F} we have $f(\ell) = g(\ell)$.*

Proof. To prove existence, for each partial measured foliation \mathcal{F} let $\Theta(\mathcal{F})$ denote a family of nonsingular leaf segments consisting of the union of all nonboundary saddle connections and one finite separatrix representing the germ of every infinite separatrix. Slicing \mathcal{F} along $\Theta(\mathcal{F})$ produces a partial measured foliation with canonical singularities and no proper saddle connections; in addition it also has no complementary bigons or smooth annuli, as long as \mathcal{F} was a measured foliation to begin with. In particular, every measured foliation can be sliced to produce a canonical model.

To prove uniqueness, suppose $\mathcal{F}, \mathcal{F}'$ are two equivalent canonical models. There is a sequence of partial measured foliations $\mathcal{F} = \mathcal{F}_0, \mathcal{F}_1, \dots, \mathcal{F}_n = \mathcal{F}'$ with the property that for each $i = 1, \dots, n$ either $\mathcal{F}_i \rightarrow \mathcal{F}_{i+1}$ or $\mathcal{F}_{i+1} \rightarrow \mathcal{F}_i$ is a Whitehead

collapse, partial fulfillment or isotopy. For each i , first slice \mathcal{F}_i along $\Theta(\mathcal{F}_i)$, then collapse any complementary bigons and smooth annuli, producing a canonical model \mathcal{F}_i^c . Then check that the relation between \mathcal{F}_i and \mathcal{F}_{i+1} , a Whitehead collapse, partial fulfillment, or isotopy, is converted into an isotopy relation between \mathcal{F}_i^c and \mathcal{F}_{i+1}^c . It follows that $\mathcal{F}_0 = \mathcal{F}_0^c$ is isotopic to $\mathcal{F}_n^c = \mathcal{F}'$.

To prove the last sentence, let $\mathbf{H}^2 \rightarrow S$ be a universal covering map, and let $\tilde{f}, \tilde{g}: \mathbf{H}^2 \rightarrow \mathbf{H}^2$ be lifts of f, g that restrict to the the same map on $\partial\mathbf{H}^2$. For each leaf ℓ of \mathcal{F} choose a lift $\tilde{\ell}$, which is a quasigeodesic with distinct endpoints $\partial\tilde{\ell} \in \partial\mathbf{H}^2$. It follows that $\partial\tilde{f}(\tilde{\ell}) = \partial\tilde{g}(\tilde{\ell})$ which implies $\tilde{f}(\tilde{\ell}) = \tilde{g}(\tilde{\ell})$ and so $f(\ell) = g(\ell)$. The fact that a lifted leaf is quasigeodesic comes from the correspondence between leaves of a partial measured foliation and leaves of an equivalent geodesic lamination, or one can apply Proposition 3.3.3 and its corollary together with Proposition 3.6.1. \diamond

Given a partial measured foliation \mathcal{F} , the subsurface that supports a canonical model for \mathcal{F} is well defined up to isotopy. As indicated in the proof of Proposition 2.8.1, this subsurface obtained by slicing along nonboundary saddle connections and representative finite separatrices of \mathcal{F} .

Suppose that \mathcal{F} is an arational measured foliation. The saddle connections of \mathcal{F} form a forest, each component of which contains at most one puncture, and therefore each saddle connection can be collapsed to a point by a sequence of Whitehead moves, producing an arational measured foliation \mathcal{F}' with no saddle connections. Let i_n (resp. p_n) be the number of nonpuncture (resp. puncture) n -pronged singularities of \mathcal{F}' . The pair of sequences $(i_3, i_4, \dots; p_1, p_2, \dots)$ is called the *singularity type* of \mathcal{F} . Note that the singularity type can be described alternatively by choosing a canonical partial measured foliation \mathcal{F}'' equivalent to \mathcal{F} , and letting i_n (resp. p_n) be the number of nonpunctured (resp. punctured) components of $S - \text{int}(\text{Supp}(\mathcal{F}''))$ with n boundary singularities. It follows from Proposition 2.8.1 that the singularity type is an invariant of an arational measured foliation class.

2.9 Components of a measured foliation

Let \mathcal{F} be a partial measured foliation on S . For each leaf ℓ of \mathcal{F} , we define the *foliation component* of \mathcal{F} containing ℓ as follows. When ℓ is not a circle, let $F = \text{Cl}(\ell)$, a pinched subsurface of the filled in surface \bar{S} . When ℓ is a circle, define F to be the closure in \bar{S} of the union of all circular leaves of \mathcal{F} isotopic to ℓ , also a pinched subsurface. The restriction $\mathcal{F} \upharpoonright F$ is defined to be the component of \mathcal{F} containing ℓ . For any two leaves ℓ, ℓ' , their components intersect only along the union of saddle connections of \mathcal{F} . It follows that \mathcal{F} is the union of finitely many components, and any two foliation components have pairwise disjoint interiors.

When \mathcal{F} is a canonical model, the foliation components of \mathcal{F} are precisely the

restrictions of \mathcal{F} to the connected components of $\text{Supp}(\mathcal{F})$, and each component is either arational or an annulus. More generally, there are two types of foliation components of a partial measured foliation \mathcal{F} : annular and partial arational. In a *partial arational* component, every half-infinite leaf is dense in the support of the component. In an *annular* component every leaf is a circle. If \mathcal{F} is actually a measured foliation then each annular foliation component has a support surface whose interior is connected; however this need not be true if \mathcal{F} is only a partial measured foliation, for example the union of two isotopic but disjoint annuli, foliated by circles, forms a single annular foliation component, despite the fact that the support is disconnected.

Partial fulfillment respects foliation components: any partial fulfillment map $r: (S, \mathcal{F}') \rightarrow (S, \mathcal{F})$ induces a bijection between foliation components of \mathcal{F}' and foliation components of \mathcal{F} . Also, distinct foliation components of a partial measured foliation are inequivalent in \mathcal{MF} . It follows that each element of \mathcal{MF} has a well-defined set of foliation components, forming a finite subset of \mathcal{MF} .

Remark. Components clarify the question of when a total measured foliation \mathcal{F} on a surface S is “minimal” in the sense that every leaf is dense. This always happens when f is arational, or equivalently, if the support of a canonical model for f has complement equal to a union of nonpunctured or once-punctured discs. However, minimality also occurs when the support of a canonical model for \mathcal{F} has exactly one component, and that component is partial arational. Arationality is therefore a strictly stronger property than minimality.

The Gardiner–Masur Theorem 2.5.1: a proof and an example. We give a proof using canonical models and their components. We want to show that every measured foliation \mathcal{F} is equivalent to a taut measured foliation. We noted earlier that if \mathcal{F} is arational then it is taut, and so we assume that \mathcal{F} is not arational. Choose a canonical model for \mathcal{F} , and collapse to a point each nonpunctured disc component of the complement of the canonical model, to obtain a partial measured foliation \mathcal{F}' . Each component of \mathcal{F}' is either arational or annular. Let $\mathcal{F}'_1, \dots, \mathcal{F}'_n$ be the components of \mathcal{F}' . Choose a pairwise disjoint collection of arcs transverse to \mathcal{F}' , so that each arc α is properly embedded in some component \mathcal{F}'_i , with at least one such arc for each \mathcal{F}'_i . Choose these arcs moreover so that each boundary component of each \mathcal{F}'_i contains an even number of endpoints of the arcs α ; this can be achieved if necessary by doubling each arc α . Now choose a spine for $\mathcal{C}(S - \text{Supp}(\mathcal{F}'))$ and a collapse map onto that spine subject to the condition that the collapse map identifies the endpoints of the arcs α in pairs. Let \mathcal{F}'' be the total measured foliation produced by this collapse map, and so $\mathcal{F}, \mathcal{F}''$ are equivalent; as remarked in Section 2.7, it

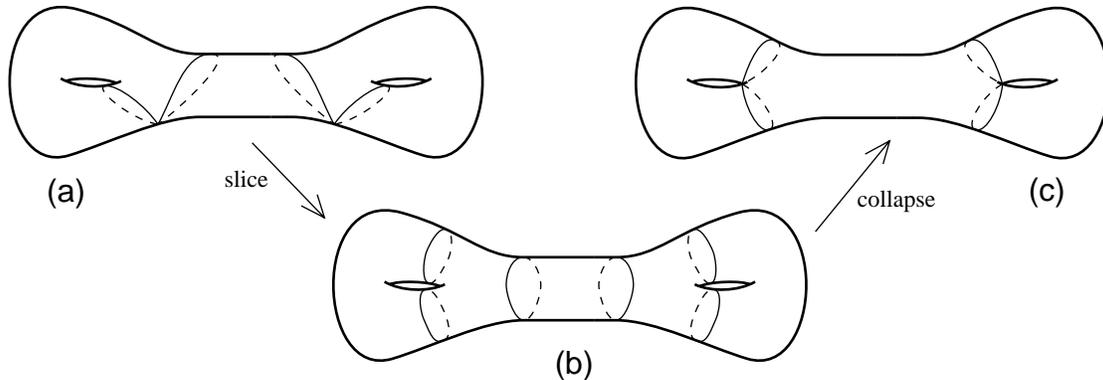


Figure 5: The measured foliation (a) is not taut. Its canonical model (b) collapses to a measured foliation (c) which is taut.

follows that $\mathcal{F}, \mathcal{F}''$ are connected by a finite sequence of isotopies and Whitehead moves. The union of the images of the arcs α gives a collection of transverse simple closed curves intersecting every leaf of \mathcal{F}'' , that is, \mathcal{F}'' is taut. This completes the proof of Theorem 2.5.1.

Figure 5(a) shows an example which I learned from Dennis Sullivan. This example was originally conceived as a generic 1-form which is not harmonic with respect to any Riemannian metric. But a generic 1-form ω has level sets determining a transversely oriented measured foliation \mathcal{F} with 4-pronged singularities, and ω is harmonic with respect to some metric if and only if \mathcal{F} is taut. The measured foliation \mathcal{F} of Figure 5(a) is not taut, but we shall show how to use the method of the previous paragraph to construct an equivalent measured foliation \mathcal{F}' which is taut.

Let S be a surface of genus 2 with the measured foliation shown in Figure 5(a). The union of saddle connections of \mathcal{F} has two components, each a figure eight. The foliation \mathcal{F} has three foliation components, each annular. If E is either of the figure eights, there is an annulus component of \mathcal{F} with one boundary component mapped surjectively to E and the other boundary component mapped to one loop of the E , from which it follows easily that \mathcal{F} is not taut. Cutting open along the saddle connections we obtain a canonical model with three annular components and two complementary components each of which is a pair of pants, shown in Figure 5(b). Collapsing each of these pair of pants onto a spine in a different way we obtain a taut measured foliation, shown in Figure 5(c).

2.10 Intersection number and joint filling

The intersection number between a measured foliation and an essential curve generalizes to an intersection number between partial measured foliations as follows.

Two partial measured foliations $\mathcal{F}, \mathcal{F}'$ are *transverse* if away from $\text{sing}(\mathcal{F}) \cup \text{sing}(\mathcal{F}')$ the foliations $\mathcal{F}, \mathcal{F}'$ are transverse, and near a singularity the two foliations are locally modelled on the horizontal and vertical foliations of a singular quadratic differential $z^n dz^2$, possibly restricted to certain sectors. On the subsurface $\text{Supp}(\mathcal{F}) \cap \text{Supp}(\mathcal{F}')$ the Fubini product of transverse measures produces a measure whose integral is defined to be $\langle \mathcal{F}, \mathcal{F}' \rangle$; this integral is always finite. The infimum of this intersection number over the equivalence classes of \mathcal{F} and \mathcal{F}' is realized in a certain sense, as we now explain.

We say that partial measured foliations $\mathcal{F}, \mathcal{F}'$ are *topologically equivalent* if they have the same underlying (nonmeasured) singular foliation. In this case we define $\langle \mathcal{F}, \mathcal{F}' \rangle = 0$.

Consider two partial measured foliations $\mathcal{F}_1, \mathcal{F}_2$. We say that $\mathcal{F}_1, \mathcal{F}_2$ are *weakly transverse* if we can write $\mathcal{F}_i = \mathcal{F}'_i \cup \mathcal{F}''_i$, where each of $\mathcal{F}'_i, \mathcal{F}''_i$ is a union of foliation components of \mathcal{F}_i , such that $\mathcal{F}'_1, \mathcal{F}'_2$ are topologically equivalent and $\mathcal{F}''_1, \mathcal{F}''_2$ are transverse and have disjoint boundary singularities.

Suppose that $\mathcal{F}, \mathcal{F}'$ are weakly transverse. A *compression disc* for $\mathcal{F}, \mathcal{F}'$ is a component of $\text{Cl}(S - (\text{Supp}(\mathcal{F}) \cup \text{Supp}(\mathcal{F}')))$ which is a nonpunctured disc whose boundary is the concatenation of one nonsingular boundary leaf segment of \mathcal{F} and one nonsingular boundary leaf segment of \mathcal{F}' . We say that $\mathcal{F}, \mathcal{F}'$ *intersect efficiently* if they are weakly transverse and have no compression discs. For example, the horizontal and vertical measured foliations of a quadratic differential intersect efficiently.

Every pair $[\mathcal{F}_1], [\mathcal{F}_2] \in \mathcal{MF}$ has efficiently intersecting representatives. The quickest way to see this is using the measured geodesic laminations λ_1, λ_2 representing $[\mathcal{F}_1], [\mathcal{F}_2]$ with respect to a particular hyperbolic structure. Then we can write $\lambda_i = \lambda'_i \cup \lambda''_i$ where λ'_1, λ'_2 have the same underlying topological lamination, and where λ''_1, λ''_2 are transverse; weakly transverse representatives $\mathcal{F}_1, \mathcal{F}_2$ are obtained by appropriately collapsing λ_1, λ_2 . This argument shows that $\mathcal{F}_1, \mathcal{F}_2$ can always be chosen to be canonical models.

Given two classes $[\mathcal{F}_1], [\mathcal{F}_2] \in \mathcal{MF}$ the intersection number $\langle [\mathcal{F}_1], [\mathcal{F}_2] \rangle$ is defined to be the infimum of $\langle \mathcal{F}'_1, \mathcal{F}'_2 \rangle$ over all weakly transverse representatives $\mathcal{F}'_1, \mathcal{F}'_2$. The infimum is realized if and only if $\mathcal{F}'_1, \mathcal{F}'_2$ intersect efficiently. To see briefly why this is true, if there are compression discs then by pulling leaves across them the intersection number is reduced. Conversely, if there are no compression discs then for any other pair of representatives $\mathcal{F}''_1, \mathcal{F}''_2$ which intersect efficiently, one can slice along separatrices to make everything canonical and then construct an ambient

isotopy taking $\mathcal{F}_1'', \mathcal{F}_2''$ jointly to $\mathcal{F}_1', \mathcal{F}_2'$, thereby showing that $\langle \mathcal{F}_1', \mathcal{F}_2' \rangle = \langle \mathcal{F}_1'', \mathcal{F}_2'' \rangle$. The converse can also be proved easily using measured geodesic laminations.

The arguments above show that $[\mathcal{F}_1], [\mathcal{F}_2] \in \mathcal{MF}$ have intersection number zero if and only if there are representatives $\mathcal{F}_1, \mathcal{F}_2$, such that each component of \mathcal{F}_1 is either topologically equivalent to some component of \mathcal{F}_2 or is disjoint from \mathcal{F}_2 , and vice versa.

Joint filling of a pair of measured foliations is defined by using the following:

Lemma 2.10.1 (and definition of joint filling). *Suppose $\mathcal{F}_1, \mathcal{F}_2$ are efficiently intersecting partial measured foliations on S . The following are equivalent*

- (1) *For every $[\mathcal{F}] \in \mathcal{M}$, either $\langle [\mathcal{F}_1], [\mathcal{F}] \rangle \neq 0$ or $\langle [\mathcal{F}_2], [\mathcal{F}] \rangle \neq 0$.*
- (2) *$\mathcal{F}_1, \mathcal{F}_2$ are not topologically equivalent and arational; and for every $[c] \in \mathcal{C}$, either $\langle [\mathcal{F}_1], [c] \rangle \neq 0$ or $\langle [\mathcal{F}_2], [c] \rangle \neq 0$.*
- (3) *No foliation component of \mathcal{F}_1 is topologically equivalent to a foliation component of \mathcal{F}_2 ; and each connected component of $\text{Cl}(S - (\text{Supp}(\mathcal{F}_1) \cup \text{Supp}(\mathcal{F}_2)))$ is a nonpunctured or once-punctured disc.*

If any of these are satisfied then we say that $\mathcal{F}_1, \mathcal{F}_2$ jointly fill.

For example, the horizontal and vertical measured foliations of any quadratic differential jointly fill. A converse to this was proved by Gardiner and Masur [GM91]: if $\mathcal{F}_1, \mathcal{F}_2$ are a jointly filling pair of measured foliations then there is a quadratic differential q such that $\mathcal{F}_1 = \mathcal{F}^h(q)$ and $\mathcal{F}_2 = \mathcal{F}^v(q)$.

Also, it is not hard to see that any two canonical models $\mathcal{F}, \mathcal{F}'$ can be isotoped so that they intersect efficiently.

Joint filling of a pair of points $[\mathcal{F}_1], [\mathcal{F}_2]$ in \mathcal{MF} is defined by the first alternative in the lemma. The equivalence of the first and third alternatives shows how to check joint filling by inspecting any efficiently intersecting representatives $\mathcal{F}_1, \mathcal{F}_2$.

Proof of Lemma 2.10.1. 1 implies 2, because if 2 is false, then one can disprove by either taking $\mathcal{F} = \mathcal{F}_1$ in the first case, or taking $\mathcal{F} = c$ in the second case.

2 implies 3, because if 3 is false then one can disprove 2 as follows. Suppose first that $\mathcal{F}_1, \mathcal{F}_2$ have topologically equivalent foliation components $\mathcal{F}_1', \mathcal{F}_2'$; if $\mathcal{F}_1, \mathcal{F}_2$ are arational then $\mathcal{F}_i = \mathcal{F}_i'$ and that immediately disproves 2; whereas if $\mathcal{F}_1, \mathcal{F}_2$ are not arational then there is a component c of $\partial \text{Supp}(\mathcal{F}_1) = \partial \text{Supp}(\mathcal{F}_2)$ which is an essential curve, and c is used to disprove 2. Suppose next that $\text{Cl}(S - (\text{Supp}(\mathcal{F}_1) \cup \text{Supp}(\mathcal{F}_2)))$ has a component C which is not a nonpunctured or once-punctured disc; it follows that C contains an essential curve c which is used to disprove 2.

It remains to prove that 3 implies 1. For this purpose we fix a complete, finite area hyperbolic structure on S and use the fact that each element of \mathcal{MF} is represented uniquely by a measured geodesic lamination on S (see [CB88]). Let $\lambda_1, \lambda_2, \mu$ be measured geodesic laminations representing $[\mathcal{F}_1], [\mathcal{F}_2], [\mathcal{F}]$, respectively; it follows that $\langle [\mathcal{F}_i], [\mathcal{F}] \rangle = \langle \lambda_i, \mu \rangle$. From 3 it follows that λ_1, λ_2 are transverse and each component of $S - (\lambda_1 \cup \lambda_2)$ is a nonpunctured or once punctured disc. Letting μ be any other measured geodesic lamination, it follows that μ has transverse intersection points with either λ_1 or λ_2 , which implies that $\langle \lambda_1, \mu \rangle \neq 0$ or $\langle \lambda_2, \mu \rangle \neq 0$. \diamond

2.11 Thurston's classification of mapping classes

In this section we review the classification, and we outline the proof for later reference, but for details the reader is referred to [Thu88], [FLP⁺79], or [CB88].

A homeomorphism $\phi: S \rightarrow S$ is *pseudo-Anosov* if there exists a transverse pair of measured foliations $\mathcal{F}_s, \mathcal{F}_u$ and a number $\lambda > 1$ such that $\phi(\mathcal{F}_u) = \lambda \cdot \mathcal{F}_u$ and $\phi(\mathcal{F}_s) = \lambda^{-1} \cdot \mathcal{F}_s$. The measured foliation \mathcal{F}_u is called the *unstable foliation* for ϕ , and note that its leaves are expanded by a factor of λ when measured by \mathcal{F}_s . The measured foliation \mathcal{F}_s is the *stable foliation* for ϕ , and its leaves are contracted by a factor of λ when measured by \mathcal{F}_u . Because the foliations $\mathcal{F}_s, \mathcal{F}_u$ have a self-map stretching or contracting every leaf by λ , they have no saddle connections, and so they have the same singularity type, which is defined to be the singularity type of the pseudo-Anosov homeomorphism ϕ .

An element of $\mathcal{MCG}(S)$ is pseudo-Anosov if it is represented by a pseudo-Anosov homeomorphism.

A homeomorphism $\phi: S \rightarrow S$, and its mapping class, are called *reducible* if there exists a system of pairwise disjoint, pairwise nonisotopic essential closed curves $\mathcal{C} = c_1 \cup \dots \cup c_n$ such that $\phi(\mathcal{C})$ is isotopic to \mathcal{C} . Such a curve system \mathcal{C} is called a *reducing system* for ϕ .

Theorem 2.11.1 ([Thu88]). *Every mapping class is either finite order, reducible, or pseudo-Anosov.*

Outline of the proof. The first step is Theorem 2.7.1 which says that the compactified Teichmüller space $\overline{\mathcal{T}} = \mathcal{T} \cup \mathcal{PMF}$ is homeomorphic to a ball. The mapping class group \mathcal{MCG} acts on $\overline{\mathcal{T}}$ by homeomorphisms, and so by the Brouwer fixed point theorem each $\Phi \in \mathcal{MCG}$ has a fixed point in $\overline{\mathcal{T}}$. If there is a fixed point in \mathcal{T} itself then Φ has finite order, because the action of \mathcal{MCG} on \mathcal{T} is properly discontinuous.

Suppose then that there is a fixed point $\mathcal{P}[\mathcal{F}] \in \mathcal{PMF}$, and so there exists a representative measured foliation \mathcal{F} , a number $\lambda > 0$, and a representative ϕ of Φ such that $\phi(\mathcal{F})$ is equivalent to $\lambda\mathcal{F}$ up to Whitehead equivalence and isotopy.

If \mathcal{F} is not arational, then replacing \mathcal{F} by any canonical model, we may assume that $\phi(\mathcal{F}) = \lambda\mathcal{F}$. It follows that the union of the boundary components of the component foliations of \mathcal{F} , with duplicates deleted in each isotopy class, forms a reducing system for ϕ .

The remaining case, where \mathcal{F} is arational, breaks down according to the following result, which can be found for example in [FLP⁺79]:

Lemma 2.11.2. *Suppose that \mathcal{F} is an arational measured foliation and Φ is a mapping class such that $\Phi[\mathcal{F}] = \lambda[\mathcal{F}]$ in \mathcal{MF} . If $\lambda = 1$ then Φ has finite order. If $\lambda \neq 1$ then Φ is pseudo-Anosov, and moreover if each saddle connection of \mathcal{F} is collapsed to a point then there exists a pseudo-Anosov homeomorphism ϕ representing Φ such that the following holds: if $\lambda > 1$ then \mathcal{F} is the unstable foliation of ϕ ; whereas if $\lambda < 1$ then \mathcal{F} is the stable foliation of ϕ .*

◇

3 Train tracks

In this section we review the theory of train tracks, due to W. Thurston. The lore of train tracks was absorbed and digested by many people beginning with Thurston's lectures in the late 1970's. Bits of this lore are scattered throughout the literature. Some of it is expounded in [Thu87], [CB] (especially volume 2), [CB88], [Pen92], and [Bon99], and we refer to these sources in the many instances when we omit detailed proofs.

3.1 Pretracks

A *pretrack* on S is a nonempty, smooth, closed 1-complex $\tau \subset S$ with the following property: for each vertex v of τ there is unique tangent line $L \subset T_v(S)$ such that for some neighborhood U of v the intersection $\tau \cap U$ is a union of smooth open arcs in S , all of which are tangent to L at v . A point $v \in \tau$ at which τ is not a 1-manifold is called a *switch* of τ , and the set of switches is denoted $\text{Sw}(\tau)$. A component of $\tau - \text{Sw}(\tau)$ is called a *branch*. A branch can be a smooth circular component of τ , or it can be an open arc with two ends, each end located at some switch. A noncircular branch will sometimes be confused with its closure in τ ; this confusion should cause no trouble, because distinct branches have distinct closures. Each noncircular branch b has two *ends* in the sense of Freudenthal, and each end of b is *located* at a particular switch s of τ , meaning that s is the accumulation point of the end. Given a switch v with tangent line L , the set of branch ends located at v is partitioned by the two directions of L into two nonempty subsets, called the two *sides* at v . Because a switch is a non-manifold point, at least one of its sides has two or more branch ends. A switch v is *semigeneric* if one of its sides has exactly one branch end; this side is called the *one-ended side* of v , and the other is called the *multi-ended side* of v . A semigeneric switch v is *generic* if the multi-ended side has exactly two branch ends. We say that τ is (semi)generic if all of its switches are (semi)generic.

A *track with terminals* τ is like a pretrack except that we drop the condition that each vertex v has a neighborhood U which is a union of smooth open arcs, and instead we require that v has a neighborhood U which is either a union of smooth open arcs or a union of smooth half-open arcs each with its endpoint at v . Any vertex of the second kind which is not of the first kind is called a *terminal* of τ .

Given a pretrack $\tau \subset S$, a *train path* in τ is a smooth path in S whose domain is a connected 1-manifold, whose image is contained in τ , and whose derivative is everywhere nonzero. A *closed train path* is one whose domain is a circle. The *combinatorial length* of a train path α in τ is the number of components of $\text{domain}(\alpha) - \alpha^{-1}(\Sigma)$ where Σ is the union of the set $\text{Sw}(\tau)$ and one point in each

circular component of τ . A train path is *bi-infinite* if its domain is diffeomorphic to the real line and the restriction to each half-line has infinite combinatorial length.

A map $f: \tau \rightarrow \tau'$ between two pretracks is an *abstract carrying map* if, for every train path α in τ , $f \circ \alpha$ is a train path in τ' . Note the distinction between an abstract carrying map and an “ordinary” carrying map, to be defined later; roughly speaking, given two pretracks $\tau, \tau' \subset S$, an ordinary carrying map $\tau \rightarrow \tau'$ must be approximable by smooth embeddings $\tau \hookrightarrow S$ which are isotopic to the inclusion $\tau \subset S$, whereas an abstract carrying map has no such restriction. If $f: \tau \rightarrow \tau'$ is a homeomorphism, and if both f and its inverse are abstract carrying maps, then f is called an *abstract diffeomorphism*. The abstract diffeomorphism type of a pretrack is completely determined by its underlying structure as a 1-complex together with the side partition of the set of branch ends at each switch.

When S has finite type, we shall define a train track on S to be a pretrack $\tau \subset S$ whose complementary components all have negative index. That is, we may regard the metric completion of $S - \tau$ as a surface with cusps on its boundary, and we define the index of a “surface with cusps on its boundary” as the Euler characteristic minus $\frac{1}{2}$ times the number of boundary cusps. Requiring negative index has the effect of ruling out a small list of diffeomorphism types: nonpunctured and once punctured discs with no cusps; nonpunctured discs with one or two cusps; and nonpunctured annuli with no cusps. We wish to make these notions very precise, using some concepts which will prove useful in several contexts, so we shall make a detour to discuss surfaces-with-corners.

A *surface-with-corners* is an oriented surface-with-boundary F equipped with an atlas of charts mapping to *local models*, certain closed subsets of \mathbf{R}^2 , with C^1 overlap maps. Each point of F is categorized into one of several types, according to local models for a neighborhood basis: *interior* points; *regular boundary* points; *regular corners*; *reflex corners*; *regular cusps*; and *reflex cusps*. Given a point $x \in F$ and a chart $U \ni x$, the set U will map to an open subset of the local model corresponding to x , and we describe each local model under the assumption that x maps to the origin of \mathbf{R}^2 . For an interior point the local model is: \mathbf{R}^2 . For a regular boundary point: any closed half-plane whose boundary passes through the origin. For a regular corner: any closed quadrant. For a reflex corner: the closure of the complement of any closed quadrant. For a regular cusp: any set of the form $\{(x, y) \in \mathbf{R}^2 \mid x \geq 0, f(x) \leq y \leq g(x)\}$, where $f, g: [0, \infty) \rightarrow \mathbf{R}$ are two C^1 functions which vanish to first order at $x = 0$ and which satisfy $f(x) < g(x)$ for all $x > 0$. For a reflex cusp: the closure of the complement of a cusp local model. A *generalized corner* is either a regular corner, reflex corner, regular cusp, or reflex cusp. A *convex corner* is either a regular corner or a regular cusp. A surface-with-corners is *smooth* if it has no generalized corners.

One common way that reflex corners and reflex cusps occur is when S is a

smooth surface and $F \subset S$ is a convex subsurface-with-corners. The subsurface-with-corners $\text{Cl}(S - F)$ then has a reflex corner for every regular corner of F , and a reflex cusp for every regular cusp of F .

In most contexts we will drop the word “regular” and refer to simply a “corner” or a “cusp”. In certain contexts where no confusion can arise, we will sometimes use the term “cusp” to refer to either a regular cusp or a reflex cusp.

Suppose that F is any C^1 surface and $K \subset F$ is a C^1 subcomplex with no isolated points. The subsurface $F - K$ has a natural *completion* denoted $\mathcal{C}(F - K)$, which is a surface-with-corners with interior $F - K$, such that the inclusion $F - K \hookrightarrow F$ extends to a C^1 map $(\mathcal{C}(F - K), \partial\mathcal{C}(F - K)) \rightarrow (\text{Cl}(F - K), \text{Fr}(K))$. The completion may be defined by choosing any C^0 Riemannian metric on F , restricting to $F - K$, and then taking the metric completion. The map $(\mathcal{C}(F - K), \partial\mathcal{C}(F - K)) \rightarrow (\text{Cl}(F - K), \text{Fr}(K))$ is called the *overlay map*. The completion and the overlay map are well-defined up to C^1 diffeomorphism, independent of the choice of metric. We remark that $\mathcal{C}(F - K)$ can exhibit any of the four types of generalized corners, depending on how edges of K meet at vertices.

We will be particularly interested in two special cases which exhibit only convex corners. If τ is a pretrack in S then, besides interior points and regular boundary points, the surface-with-corners $\mathcal{C}(S - \tau)$ can have regular cusps, but cannot have any of the three other generalized corners; we call this a *surface with cusps*. If τ, τ' is a transverse pair of pretracks in S then $\mathcal{C}(S - (\tau \cup \tau'))$ can have regular corners or cusps; nonconvex corners do not occur.

A surface-with-corners has *finite type* if it is C^1 diffeomorphic to a compact surface-with-corners minus a finite set of interior points called *punctures*. If C is a surface-with-corners of finite type, we define the *Euler index* $\iota(C)$ as follows. In the special case where C has only convex corners the formula is

$$\iota(C) = \chi(C) - \frac{1}{2}\#(\text{cusps of } C) - \frac{1}{4}\#(\text{corners of } C)$$

where $\chi(C)$ is the ordinary Euler characteristic. More generally the formula is

$$\begin{aligned} \iota(C) = \chi(C) - \frac{1}{2}(\#(\text{cusps}) - \#(\text{reflex cusps})) \\ - \frac{1}{4}(\#(\text{regular corners}) - \#(\text{reflex corners})) \end{aligned}$$

The Euler index is additive under gluing, which means: if S is a surface-with-cusps obtained by C^1 gluing along edges of surfaces-with-cusps S_1, \dots, S_n , then

$$\iota(S) = \iota(S_1) + \dots + \iota(S_n)$$

The Euler index can also be defined using obstruction theory, a viewpoint which is useful in proving additivity of the Euler index. Define the *cross bundle* of C

to be the fiber bundle whose fiber over each point x is the set of *crosses* in the tangent plane to x , a cross consisting of an unordered transverse pair of unoriented 1-dimensional subspaces; the two lines in a cross are called the *slats* of the cross. The cross bundle is bundle homotopy equivalent to an oriented circle bundle over C ; its 4-fold cover is the oriented frame bundle of C , which itself is bundle homotopy equivalent to the tangent circle bundle of C . Over the boundary of C and in a neighborhood of the punctures, there is a section of the cross bundle well-defined up to proper homotopy, as follows. The cross is tangent to the boundary, meaning that at each boundary point one of the two slats is tangent to the boundary, and at a corner point the two slats are tangent to the two respective sides of the boundary. Near a puncture the cross is *polar*, meaning that in some coordinate disc centered on the puncture the two slats of the cross are tangent to the level curves of polar coordinates, one slat tangent to circles and the other to rays. Using orientability of C , the obstruction to extending this partial section over all of C takes values in \mathbf{Z} , and the Euler index $\iota(C)$ equals $\frac{1}{4}$ times the obstruction. The fraction $\frac{1}{4}$ correlates with the 4-fold covering by the oriented frame bundle; for example, if C has no generalized corners then $\iota(C) = \chi(C)$.

3.2 Train tracks

Consider now our fixed finite type surface (S, P) . Given any pretrack τ on S , each component of $\mathcal{C}(S - \tau)$ is a surface-with-cusps.

A *train track* on S is a pretrack $\tau \subset S$ such that each component of $\mathcal{C}(S - \tau)$ has negative Euler index. This rules out components of $\mathcal{C}(S - \tau)$ the following types: *nullgons*, which are smooth discs with no punctures, having Euler index 1; *monogons*, nonpunctured discs with 1 cusp, having Euler index $\frac{1}{2}$; *peripheral smooth discs*, which are smooth discs with 1 puncture, having Euler index 0; *bigons*, nonpunctured discs with 2 cusps, having Euler index 0; and smooth annuli, having Euler index 0. All other types of components are allowed. Of course, the sum of the Euler indices of components of $\mathcal{C}(S - \tau)$ equals $\chi(S)$. Some examples of train tracks are shown in Figure 6.

There are two important equivalence relations on the set of train tracks on S . The first is *isotopy* which means ambient isotopy in the surface S : two train tracks τ, τ' are isotopic if there exists $f \in \text{Homeo}_0(S)$ such that $f(\tau) = \tau'$. The second relation is *combinatorial equivalence* or *topological equivalence*, meaning that there exists $f \in \text{Homeo}_+(S)$ such that $f(\tau) = \tau'$. The phrase “combinatorial equivalence” refers to the fact that, when the surface $\mathcal{C}(S - \tau)$ is reglued along its boundary to yield the surface S , and when the same is done to $\mathcal{C}(S - \tau')$, then the gluing patterns are identical if and only if there exists $f \in \text{Homeo}_+(S)$ such that $f(\tau) = \tau'$. A *combinatorial invariant* of a train track τ is a property of τ which is

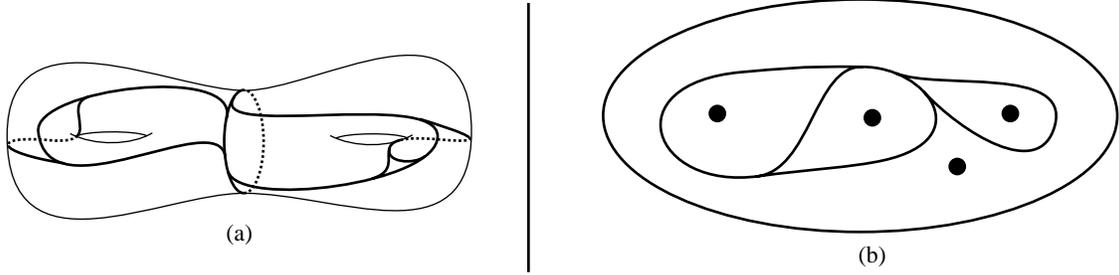


Figure 6: Examples of train tracks. In (a), $\mathcal{C}(S - \tau)$ consists of two cusped triangles and one cusped 4-gon. In (b), $\mathcal{C}(S - \tau)$ consists of four punctured monogons.

invariant under combinatorial equivalence. Many of the properties of train tracks that we investigate are combinatorial invariants; at first we will try to mention this explicitly, but eventually we will forget to mention it in the many cases where it is true.

One measure of complexity of a train track is its number of cusps, which is a combinatorial invariant. The next lemma shows how this is bounded in terms of the topology of S . A train track τ is *complete* if each component of $\mathcal{C}(S - \tau)$ has the smallest possible negative Euler index, namely, $-\frac{1}{2}$; a component has Euler index $-\frac{1}{2}$ if and only if it is a nonpunctured disc with three cusps or a once-punctured disc with one cusp. Completeness is a combinatorial invariant.

Lemma 3.2.1 (Cusp count). *If S has p punctures then any train track $\tau \subset S$ has at most $6|\chi(S)| - 2p$ cusps, and the maximum is achieved if and only if τ is complete.*

Proof. Suppose that some component of $\mathcal{C}(S - \tau)$ has Euler index $< -\frac{1}{2}$. We shall add certain branches τ to complete it. When a branch b is added to a component C of $\mathcal{C}(S - \tau)$, either b is nonseparating in which case $\iota(\mathcal{C}(C - b)) = \iota(C) < 0$, or b is separating and we are careful that both components have negative Euler index, and so the result is still a train track. A new simple closed curve component added to τ leaves the number of cusps unchanged, whereas adding a branch with two ends on τ adds two cusps. At the end we'll count cusps of a complete train track.

In any component of genus ≥ 1 , add to τ a nonseparating simple closed. Next, in any component with two or more boundary components, add a nonseparating branch with ends on two different boundary components. Next, if any component has two or more punctures then add a separating branch which has punctures on both sides, taking care not to create a component which is a once-punctured disc with no cusps. At this stage each component of $\mathcal{C}(S - \tau)$ is a disc with ≤ 1

puncture. If any component is a once-punctured disc with ≥ 2 sides, then add a branch encircling the puncture and creating a once-punctured disc with one cusp. At this stage each once-punctured component has one cusp. If any component is a nonpunctured disc with four or more sides then add a branch with ends on two distinct sides, taking care not to create a bigon by arranging that the switch orientations at the two ends both agree with the boundary orientation of the disc. The train track is now complete.

If τ is complete, each component of $\mathcal{C}(S - \tau)$ has Euler index $-\frac{1}{2}$. Let A be the number of nonpunctured components with three cusps and B the number of once-punctured disc components with one cusp. We have

$$\begin{aligned} |\chi(S)| &= \frac{1}{2} \#(\text{components of } \mathcal{C}(S - \tau)) = \frac{A}{2} + \frac{B}{2} \\ &= \frac{A}{2} + \frac{p}{2} \\ A &= 2|\chi(S)| - p \\ \#(\text{cusps}) &= 3A + B = 3A + p \\ &= 6|\chi(S)| - 2p \end{aligned}$$

◇

If τ is a train track on S and $\sigma \subset \tau$ is a pretrack, then σ is a train track, called a *subtrack* of τ . This follows because each component C of $\mathcal{C}(S - \sigma)$ is obtained by gluing components of $\mathcal{C}(S - \tau)$, whose Euler indices add up to $\iota(C)$, and so $\iota(C) < 0$.

As a simple exercise, a 3-punctured sphere contains no train track; this can be proved by adding up the Euler indices of possible components of $\mathcal{C}(S - \tau)$ to get $\chi(S) = -1$, and deriving a contradiction.

The support of a train track. Given a train track $\tau \subset S$, the *support* of τ is a subsurface $\text{Supp}(\tau)$ which is the union of a closed regular neighborhood $\nu(\tau)$ and any components of $S - \nu(\tau)$ that are nonpunctured or once-punctured discs. The support of τ is unique up to isotopy.

A subsurface $F \subset S$ is *essential* if:

- (1) F is finite type surface-with-boundary,
- (2) F is a closed subset of S , and
- (3) ∂F is an essential curve family in S .

Conditions (1) and (2) hold for F if and only if they hold for $\text{Cl}(S - F)$. Assuming (1) and (2) hold, condition (3) is equivalent to the statement that no component of

F or of $\text{Cl}(S - F)$ is a nonpunctured or once-punctured disc, and so F is essential if and only if $\text{Cl}(S - F)$ is essential. Note that each component of an essential subsurface is essential, and a disjoint union of essential subsurfaces is essential.

We record here a simple but important property:

Proposition 3.2.2. *For any train track $\tau \subset S$, the subsurface $\text{Supp}(\tau)$ is essential.*

Proof. Properties (1) and (2) for $\text{Supp}(\tau)$ are evident. Consider a component c of $\partial \text{Supp}(\tau)$, and suppose that $c = \partial D$ for some nonpunctured once-punctured disc D . If $D \cap \tau = \emptyset$ then D is a component of $\text{Cl}(S - \nu(\tau))$, contradicting the definition of $\text{Supp}(\tau)$.

Suppose then that $D \cap \tau \neq \emptyset$, so $D \cap \tau$ is a subtrack τ' of τ and $\tau' \cap \partial D \subset \tau \cap c = \emptyset$. The component of $\mathcal{C}(D - \tau')$ that contains c is a nonpunctured annulus, possibly with cusps, and hence has Euler index ≤ 0 . Since $\iota(D) \geq 0$ it follows that some component C of $\mathcal{C}(D - \tau')$ disjoint from c has Euler index ≥ 0 , but C is also a component of $\mathcal{C}(S - \tau)$, contradicting the definition of a train track. \diamond

Corollary 3.2.3. *If τ is a train track on S then for each component C of $S - \tau$, the inclusion induced homomorphism $\pi_1 C \rightarrow \pi_1 S$ is injective.*

Proof. If C is a nonpunctured or once-punctured disc this is obvious. Otherwise, C is isotopic to the interior of some component of $\text{Cl}(S - \text{Supp}(\tau))$, and hence, applying Proposition 3.2.2, it suffices to quote the well known fact that if F is an essential subsurface of S then the induced homomorphism $\pi_1 F \rightarrow \pi_1 S$ is injective. \diamond

Complexity of subsurfaces. As Lemma 3.2.1 indicates the number of cusps is an important measure of the complexity of a train track τ . Another important combinatorial invariant is the topological type of the support of τ . We shall need a somewhat crude measure of the complexity of the support, defined as follows.

Given an orientable, finite type surface-with-boundary F , define $g = g(F)$ to be the genus, $p = p(F)$ to be the number of punctures of $\text{int}(F)$ (that is, the total number of punctures and boundary components of F), and $x = x(F) = |\chi(F)| = 2g - 2 + p$. Define the complexity of F to be $d(F) = 3x - p = 6g - 6 + 2p$. Ruling out $g = 0, p = 0, 1, 2$ and $g = 1, p = 0$ it follows that $d(F)$ is a non-negative integer equal to the dimension of the Teichmüller space of $\text{int}(F)$, the thrice punctured sphere being the only example where $d(F) = 0$.

The following simple and well known fact is the basis of many induction arguments in surface topology:

Lemma 3.2.4. *If F', F are essential subsurfaces of S with $F' \subset F$ then $d(F') \leq d(F)$ with equality if and only if F' is isotopic to F .*

Proof. To prove the claim, first we may isotope F' so that for any component $c \subset \partial F$, if c is isotopic to a component $c' \subset \partial F'$ then $c = c'$, and otherwise $c \subset \text{int}(F')$. It follows that $F'' = \text{Cl}(F - F')$ is an essential subsurface and no component of F'' is a nonpunctured annulus, and so $d(F'') \geq 0$. Consider the essential curve system $C = \partial F' = \partial F'' \subset F$, and note that F' is isotopic to F if and only if $F'' \neq \emptyset$ if and only if $C \neq \emptyset$. Since C is a disjoint union of circles we have $\chi(F') + \chi(F'') = \chi(F)$, and so, in the case $C \neq \emptyset$, we have

$$\begin{aligned} x(F') + x(F'') &= x(F) \\ p(F') + p(F'') &= p(F) + 2 \cdot \#C \\ d(F') + d(F'') &= d(F) - 2 \cdot \#C \\ &< d(F) \\ d(F') &< d(F) \end{aligned}$$

◇

Types of branches. Let τ be a pretrack. Given a branch b and an end e of b located at a switch s of τ , we say that e is an *inflow* of b if e is the unique branch end on its side of s , and otherwise e is an *outflow*. If b has two inflows then it is a *sink branch*, if b has two outflows then we say it is a *source branch*, and if it has one inflow and one outflow then it is a *transition branch*.¹ By convention we sometimes think of a nongeneric switch as a “degenerate sink branch”. Every transition branch has a *flow orientation*, pointing from the inflow to the outflow.

When τ is semigeneric then this classification of branches can be restated as follows. Given a semigeneric switch s , the *switch orientation* of s is the orientation on the tangent line of s pointing from the multi-ended side to the one-ended side; it can be visualized by imagining s as an arrowhead. A branch b is a sink branch if the switch orientation at each end of b points into b , a source branch if both switch orientations point out of b , and a transition branch if one points inward and one points outward.

The terminology of a “source branch” evokes a flow pouring outward at each end; for a sink branch the flow pours inward, and for a transition branch the flow pours in one end and out the other. Sink branches are always embedded. Source branches need not be embedded, for their two ends might lie on the multi-ended side of the same switch. Transition branches need not be embedded, for their two ends might lie on opposite sides of the same switch. In a recurrent pretrack, however,

¹Source, sink, and transition branches are called *small*, *large*, and *mixed* or *half-large* branches, respectively, in [Pen92].

transition branches are always embedded; see the proof of Proposition 3.12.1 for a more general statement.

Given a generic pretrack τ , we make a further classification of source branches b of τ . For each branch b of τ , a *side* of b is formally defined to be a smooth closed arc $\tilde{b} \subset \partial\mathcal{C}(S - \tau)$ such that the overlay map restricts to a bijection between $\text{int}(\tilde{b})$ and $\text{int}(b)$. Each branch of τ has exactly two sides. The overlay map induces a bijection between the ends of b and the ends of each side \tilde{b} . Given an end η of a source branch b , the cusp v at η is *located* on one of the two sides \tilde{b} of b , characterized as the unique side so that the end $\tilde{\eta}$ of \tilde{b} corresponding to η accumulates at the cusp of $\mathcal{C}(S - \tau)$ corresponding to v . Furthermore, drawing \tilde{b} as a vertical arc in the plane with $\tilde{\eta}$ at the lower endpoint of \tilde{b} and v pointing downward, the arc \tilde{b} is either on the Left or the Right side of the cusp v . We say that b is an LL *source branch* if each end of b is on the Left side of its incident cusp, b is an RR *source branch* if each end is on the Right side of the incident cusp, and b is an LR *source branch* if one end is on the Left and the other is on the Right side of its incident cusp. Note that if b is an LR source branch then b is the overlay image of some side β of some nonsmooth component c of $\partial\mathcal{C}(S - \tau)$, and moreover the endpoints of b are identified to the same switch if and only if c has exactly one cusp. On the other hand, if b is an LL or RR source branch then b is *not* the overlay image of any side of any component of $\partial\mathcal{C}(S - \tau)$.

General structure of a finite pretrack. Consider a finite pretrack τ .

A smooth closed curve γ embedded in τ is called a *sink loop* if it is a union of transition branches; equivalently, γ has an orientation, also called the *flow orientation*, such that every switch on γ is a semigeneric switch at which the switch orientation agrees with the flow orientation.

Let t be the closure of a component of

$$\tau - \cup\{\text{source branches}\} - \cup\{\text{sink branches}\} - \cup\{\text{sink loops}\}$$

where the union of sink branches includes degenerate sink branches, that is, nongeneric switches. Then t is a union of transition branches, and t is a directed tree with a particular vertex called the *outflow* such that the flow orientation on each branch of t points towards the outflow. We refer to these trees as the *transition trees* of τ ; they are pairwise disjoint, except that two transition trees may have the same outflow located at a nongeneric switch of τ . When τ has no sink loops then every transition branch lies in a transition tree, and their union is called the *transition forest* of τ .

Let ξ be a sink branch, nongeneric switch, or a sink loop. The *basin* of ξ is the union of ξ with any transition trees that intersect ξ nontrivially. There is a

deformation retraction from the basin of ξ onto ξ , taking each transition tree to its outflow. The basin is either contractible, when ξ is a sink branch or nongeneric switch, or is homotopy equivalent to a circle, when ξ is a sink loop.

The entire pretrack τ is the disjoint union of its basins, together with source branches acting as bridges among the basins. A source branch can have its two ends on different basins, or on the same one; indeed, the two ends can be located at the same switch, on the same side of that switch, or on opposite sides of a nongeneric switch.

3.3 Train paths in the universal cover

The results in this section are adapted from [CB], with extensions to punctured surfaces.

Consider a train track τ . We'll assume there is a fixed hyperbolic structure on S , complete and of finite area, so that \tilde{S} is isometric to the hyperbolic plane \mathbf{H}^2 . The total lift of τ to \tilde{S} is a pretrack $\tilde{\tau}$.

Each component C of $\mathcal{C}(\tilde{S} - \tilde{\tau})$ is the universal cover of a component of $\mathcal{C}(S - \tau)$, by Corollary 3.2.3, and so either C is compact and projects homeomorphically to a component of $\mathcal{C}(S - \tau)$, or the stabilizer of C under the deck transformation group is infinite. In particular, if C is compact then $\iota(C) < 0$.

The negative index components of $\mathcal{C}(\tilde{S} - \tilde{\tau})$ give enough semblance of negative curvature so that train paths in $\tilde{\tau}$ have many of the same large scale properties of geodesics in a negatively curved metric space, as the next several propositions show.

Proposition 3.3.1. *Let τ be a train track.*

- (1) *Every train path in $\tilde{\tau}$ is embedded.*
- (2) *The intersection of any two train paths in $\tilde{\tau}$ is connected.*

Proof. Suppose there exists a nonembedded train path in $\tilde{\tau}$. Letting $\gamma: [a, b] \rightarrow \tilde{\tau}$ be a minimal such path, it follows that $\gamma(a) = \gamma(b)$ and γ is otherwise one-to-one. By the Schoenflies Theorem γ bounds a closed disc D . The boundary of D is smooth except for a possible singularity at the point $\gamma(a) = \gamma(b)$, where D has either a cusp, a smooth point, or a reflex cusp. It follows that $\iota(D) = \frac{1}{2}$, 1, or $\frac{3}{2}$. But this is impossible: D is a union of compact components of $\mathcal{C}(\tilde{S} - \tilde{\tau})$, each with negative Euler index, whose sum is $\iota(D)$.

Suppose that there exist two train paths with disconnected intersection. The train paths themselves are embedded, and so it follows that there are subpaths α, β with the same endpoints and disjoint interiors. The simple closed curve $\alpha \cup \beta$ bounds a disc D , again by the Schoenflies Theorem. The boundary of D is smooth

in $\text{int}(\alpha) \cup \text{int}(\beta)$, and at each of the common boundary points D has either a smooth point, or a cusp, or a reflex cusp. From this it follows that the minimum value of $\iota(D)$ is zero, occurring when both points are cusps; from there $\iota(D)$ can take any value in increasing increments of $\frac{1}{2}$ up to 2, occurring when there are two reflex cusps. This leads to a contradiction as above. \diamond

Note that the results of Proposition 3.3.1 apply not just to a finite train path $[a, b] \mapsto \tilde{\tau}$, but to a *train ray* which is (the image of) a proper train path $[0, \infty) \mapsto \tilde{\tau}$, and to a *train line* which is (the image of) a proper train path $(-\infty, +\infty) \mapsto \tilde{\tau}$.

In a metric space, the *Hausdorff distance* between two sets A, B is the infimum of $r \in [0, \infty]$ such that $A \subset N_r(B)$ and $B \subset N_r(A)$.

Proposition 3.3.2. *Let τ be a train track.*

- (1) *If ρ, ρ' are two train rays in $\tilde{\tau}$ which have finite Hausdorff distance in \tilde{S} , then ρ, ρ' eventually coincide.*
- (2) *If ρ, ρ' are two train lines in $\tilde{\tau}$ which have finite Hausdorff distance in \tilde{S} , then ρ, ρ' coincide.*

Proof. We may assume that each component of $\mathcal{C}(S - \tau)$ is a nonpunctured or once punctured disc, for if this is not true then we can inductively add new branches to τ to get a larger and larger sequence of train tracks until it becomes true. With this assumption, it follows that $\tilde{\tau}$ is “uniformly connected” in \tilde{S} , meaning that for all $r \geq 0$ there exists $s \geq 0$ such that for any $x, y \in \tilde{\tau}$, if $d(x, y) \leq r$ then there is a path γ in $\tilde{\tau}$ from x to y of combinatorial length $\leq s$, where the combinatorial length of γ is defined as the number of branches of $\tilde{\tau}$ traversed by the path. Note that γ need not be a train path; it can make illegal turns at the switches of $\tilde{\tau}$.

To prove (1), assuming that ρ, ρ' do not eventually coincide, it follows from Proposition 3.3.1 that their intersection, if nonempty, is a finite train path. Let r be the Hausdorff distance between ρ and ρ' , and let s be as above, and so each point of ρ is connected to a point of ρ' by a path in $\tilde{\tau}$ of combinatorial length $\leq s$. Choose a sequence of points $x_i \in \rho$ going out the end, and for each of them choose a path γ_i from x_i to a point $x'_i \in \rho'$ such that γ_i has combinatorial length $\leq s$. We may assume that γ_i is a simple path. Replacing γ_i by a subpath and moving the points x_i, x'_i a uniformly bounded amount, if necessary, we may assume that γ_i intersects ρ only in x_i and ρ' only in x'_i . Fix i . If j is sufficiently large then we obtain a simple closed curve $\alpha_j = \gamma_i^{-1} * \sigma_i * \gamma_j * \sigma'_j^{-1}$ where σ_j, σ'_j are subsegments of ρ, ρ' , respectively. By the Schoenflies theorem, α_j bounds a disc D_j . As j goes to ∞ , the lengths of σ_j, σ'_j go to ∞ . Each component of $\mathcal{C}(\tilde{S} - \tilde{\tau})$ in D_j is compact, and it follows that the number of such components, the “area” of D_j , goes to ∞ . By summing Euler indices, it follows that $\iota(D_j)$ goes to $-\infty$, because each compact

component of $\mathcal{C}(\tilde{S} - \tilde{\tau})$ has index $\leq -\frac{1}{2}$. However, the number of singularities on $\partial D_j = \alpha_j$ is $\leq 2(s+1)$, and it follows that $\iota(D) \geq -(s+1)$, a contradiction.

To prove (2), each of the two ends of ρ, ρ' eventually coincide, and the intersection of ρ, ρ' is connected, and so $\rho = \rho'$. \diamond

The following result shows that the analogy between train paths and geodesics is quite precise, in the large scale: with respect to any complete, finite area hyperbolic structure on S , train paths in the universal cover \mathbf{H}^2 are uniformly quasigeodesic. In the case where S is compact and τ is transversely recurrent, a stronger conclusion due to Thurston says that one can pick the hyperbolic structure on S so that train paths of τ have arbitrarily small geodesic curvature; this argument can be found in [Pen92], Theorem 1.4.4, communicated via notes of Nat Kuhn.

Given metric spaces X, Y and numbers $\lambda \geq 1, \epsilon \geq 0$, a map $f: X \rightarrow Y$ is a (λ, ϵ) -*quasigeodesic embedding* if

$$\frac{1}{\lambda}d(x, y) - \epsilon \leq d(f(x), f(y)) \leq \lambda d(x, y) + \epsilon, \quad \text{for all } x, y \in X$$

If in addition for each $y \in Y$ there exists $x \in X$ such that $d(f(x), y) \leq \epsilon$, then we say that f is a (λ, ϵ) -*quasi-isometry* between X and Y .

A (λ, ϵ) -*quasigeodesic* in a metric space X is a (λ, ϵ) -quasi-isometric embedding $p: I \rightarrow X$, where I is a closed, connected subset of \mathbf{R} . We call p a *quasigeodesic segment, ray, or line* depending on whether I is a finite interval, half-line, or the whole line.

In the hyperbolic plane \mathbf{H}^2 , every quasigeodesic ray converges to a unique point in $\partial\mathbf{H}^2$, the two ends of a quasigeodesic line converge to distinct points in $\partial\mathbf{H}^2$, and for each $\lambda \geq 1, \epsilon \geq 0$ there exists $r \geq 0$ such that each (λ, ϵ) -quasigeodesic segment, ray, or line has Hausdorff distance $\leq r$ from the unique geodesic with the same ends (finite or infinite).

Proposition 3.3.3. *Let τ be a train track in S . There exists $\lambda \geq 1, \epsilon \geq 0$ such that every train path in $\tilde{\tau}$ is a (λ, ϵ) quasigeodesic in $\tilde{S} \approx \mathbf{H}^2$.*

Combining this with Proposition 3.3.2 we obtain:

Corollary 3.3.4. *The end of every train ray converges to a unique infinite endpoint in $\partial\mathbf{H}^2$. Two train rays converging to the same infinite endpoint on $\partial\mathbf{H}^2$ eventually coincide. The two ends of every train line converge to distinct infinite endpoints. Two train lines with the same infinite endpoints coincide.* \diamond

Proof of Proposition 3.3.3. First we give the proof when S is compact. As above, by adding branches we may assume that each component of $\mathcal{C}(S - \tau)$ is a nonpunctured

disc. Each component of $\mathcal{C}(\tilde{S} - \tilde{\tau})$ is therefore diffeomorphic to a component of $\mathcal{C}(S - \tau)$ that it covers, and so $\tilde{\tau}$ is a train track in \tilde{S} .

Claim: The inclusion map from $\tilde{\tau}$, with its path metric, into \mathbf{H}^2 , is a quasi-isometry. This is an incidence of the basic principle of geometric group theory, which says that if X, Y are proper, geodesic metric spaces on which a group G acts properly discontinuously and cocompactly, then any G -equivariant map $X \mapsto Y$ is a quasi-isometry. There is also a direct proof, which uses only the following facts about the train track $\tilde{\tau}$: there are only finitely many isometry types of components of $\mathcal{C}(\tilde{S} - \tilde{\tau})$; and the decomposition of \tilde{S} into these components is uniformly locally finite. Here is the proof, which will be useful in what follows.

Evidently $d_{\mathbf{H}^2}(x, y) \leq d_{\tilde{\tau}}(x, y)$ for $x, y \in \tilde{\tau}$. Also, letting D be the maximum diameter of a component of $\mathcal{C}(\tilde{S} - \tilde{\tau})$, for every $y \in \mathbf{H}^2$ there exists $x \in \tilde{\tau}$ such that $d_{\mathbf{H}^2}(x, y) \leq D$. The only nontrivial thing to check is the inequality

$$\frac{1}{K}d_{\tilde{\tau}}(x, y) - C \leq d_{\mathbf{H}^2}(x, y)$$

with constants $K \geq 1$, $C \geq 0$ independent of $x, y \in \tilde{\tau}$. Let γ be the \mathbf{H}^2 geodesic between x and y . There is a constant μ such that any two points in $\tilde{\tau}$ whose distance in \mathbf{H}^2 is at most $2D$ are connected by a path in $\tilde{\tau}$ of length at most μ . There exists a sequence of points $x = x_0, x_1, \dots, x_N, x_{N+1} = y \in \gamma \cap \tilde{\tau}$, monotonic on γ , such that $D < d(x_{i-1}, x_i) \leq 2D$ for $i = 1, \dots, N$, and $d(x_N, x_{N+1}) \leq D$. It follows that

$$\begin{aligned} d_{\mathbf{H}^2}(x, y) &\geq \sum_{i=1}^N d_{\mathbf{H}^2}(x_{i-1}, x_i) \\ &\geq \frac{D}{\mu} \sum_{i=1}^N d_{\tilde{\tau}}(x_{i-1}, x_i) \\ &\geq \frac{D}{\mu} \sum_{i=1}^{N+1} d_{\tilde{\tau}}(x_{i-1}, x_i) - \frac{D}{\mu} \mu \\ &\geq \frac{D}{\mu} d_{\tilde{\tau}}(x, y) - D \end{aligned}$$

completing the proof of the claim.

Supposing the proposition is not true, for any $I > 0$ there exists a finite train path γ with endpoints x, y such that $\text{Length}(\gamma) \geq I$ and

$$I \cdot d(x, y) \leq \text{Length}(\gamma).$$

It follows that x, y are connected by a path δ in $\tilde{\tau}$, not necessarily a train path, such that $\text{Length}(\delta) \leq Kd(x, y) + C$, for some constants $K \geq 1$, $C \geq 0$ independent of

γ ; just take δ to be a geodesic in the metric space $\tilde{\tau}$. We have

$$\begin{aligned} \text{Length}(\delta) &\leq \frac{K}{I} \text{Length}(\gamma) + C \\ &\leq \frac{K+C}{I} \text{Length}(\gamma) \end{aligned} \tag{3.1}$$

We may assume that δ is an embedded path. For sufficiently large I we shall obtain a contradiction, using an index argument.

We reduce to the case that γ and δ have disjoint interiors, still maintaining 3.1. The path δ may be decomposed into subpaths as

$$\delta = \delta_1 * \cdots * \delta_L$$

in such a way that each subdivision point $\delta_l \cap \delta_{l+1}$ lies in γ , and each δ_l either lies on γ or has interior disjoint from γ . Associated to δ_l is a subpath of γ denoted γ_l , such that $\partial\delta_l = \partial\gamma_l$, and either $\gamma_l = \delta_l$ or γ_l, δ_l have disjoint interiors. Note that $\gamma = \cup_{l=1}^L \gamma_l$; this follows because each point of γ not lying on δ lies on some minimal subsegment γ' whose endpoints do lie in δ , and δ contains some minimal subsegment δ' whose endpoints lie in opposite components of $\gamma - \text{int}(\gamma')$, implying that $\delta' = \delta_k$ and $\gamma' \subset \gamma_k$ for some $k = 1, \dots, L$. It follows that

$$\begin{aligned} \sum_{l=1}^L \text{Length}(\delta_l) &= \text{Length}(\delta) \leq \frac{K+C}{I} \text{Length}(\gamma) \\ &\leq \sum_{l=1}^L \frac{K+C}{I} \text{Length}(\gamma_l) \end{aligned}$$

and so there exists $l \in \{1, \dots, L\}$ such that

$$\text{Length}(\delta_l) \leq \frac{K+C}{I} \text{Length}(\gamma_l)$$

If we are in the case that $\delta_l = \gamma_l$ then for sufficiently large I we obtain a contradiction, completing the reduction.

Assuming then that 3.1 holds with γ and δ having disjoint interiors, their union bounds a closed disc D . The ‘‘combinatorial area’’ of D , namely the number of components of $\mathcal{C}(\tilde{S} - \tilde{\tau})$ contained in D , is $\geq A \text{Length}(\gamma)$ for some constant $A > 0$ independent of γ ; this follows because each component of $\mathcal{C}(\tilde{S} - \tilde{\tau})$ has a uniformly bounded number of branches on its boundary. Each component of $\mathcal{C}(\tilde{S} - \tilde{\tau})$ having Euler index $\leq -\frac{1}{2}$, it follows that

$$\iota(D) \leq -\frac{A}{2} \text{Length}(\gamma)$$

On the other hand, each singularity on ∂D is a switch lying on δ , and so the number of such singularities is $\leq B \text{Length}(\delta)$ for some constant $B > 0$ independent of γ . In calculating the Euler index of D , each of these singularities contributes an amount $\geq -\frac{1}{2}$, and we therefore have

$$\iota(D) \geq 1 - \frac{B}{2} \text{Length}(\delta) \geq 1 - \frac{B(K+C)}{2I} \text{Length}(\gamma)$$

and so

$$\frac{A}{2} \leq -1 + \frac{B(K+C)}{2I}$$

The constants A, B, C, K being independent of γ , for sufficiently large I this is a contradiction. This completes the proof under the assumption that S is closed.

When S has punctures, then we must import ideas of relative hyperbolicity. The problem is that the inclusion $\tilde{\tau} \hookrightarrow \mathbf{H}^2$ is not a quasi-isometry. However, we will add branches to enlarge $\tilde{\tau}$, so that the inclusion becomes a quasi-isometry. First, by adding branches to τ we may assume that each component of $\mathcal{C}(S - \tau)$ is a nonpunctured or once-punctured disc. Denote $\sigma^0 = \tilde{\tau}$; this will be the first of an increasing sequence $\sigma^0 \subset \sigma^1 \subset \sigma^2 \subset \dots$ of pretracks on \tilde{S} , whose union $\sigma^\infty = \cup \sigma^i$ will be a train track. The sequence σ^n will be constructed by adding branches in each noncompact component of $\mathcal{C}(\tilde{S} - \sigma^0)$. The train track σ^∞ will have the key properties needed to insure that the inclusion map $\sigma^\infty \hookrightarrow \mathbf{H}^2$ is a quasi-isometry, namely:

- $\mathcal{C}(\tilde{S} - \sigma^\infty)$ has finitely many isometry types of components;
- These components are uniformly locally finite.

Once σ^∞ is constructed and these key properties are observed, all the above arguments go through, and it follows that train paths in σ^∞ are uniformly quasigeodesic, and so the same is true for train paths in the subtrack $\tilde{\tau}$ of σ^∞ .

Let $P \subset \partial \mathbf{H}^2$ be the set of parabolic points for the action of $\pi_1 S$. There is a one-to-one correspondence between points $p \in P$ and noncompact components N_p of $\mathcal{C}(\tilde{S} - \sigma^0)$, N_p being stabilized by the subgroup of $\pi_1 S$ that stabilizes p . For each p we describe the construction of the sequence σ^n in N_p . Here is a detailed description of the construction, accompanied by Figure 7 which should give the idea clearly.

Fix $p \in P$. Choose conformal coordinates $\alpha \times \beta: \mathbf{H}^2 \rightarrow \mathbf{R} \times \mathbf{R}$, by taking a conformal isomorphism between \mathbf{H}^2 and the upper half plane taking p to the point at infinity, and letting $\alpha = x$ and $\beta = \log(y)$. Note that each $\alpha_s = \alpha^{-1}(s)$

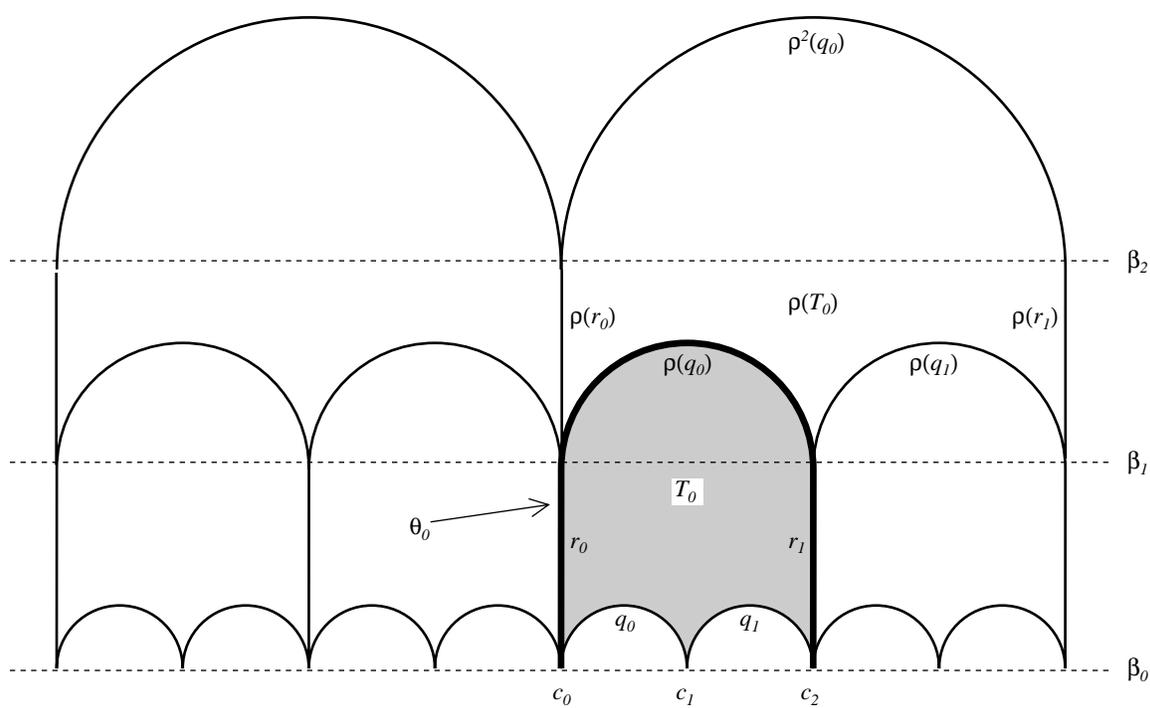


Figure 7: The pretracks $\sigma^0 \subset \sigma^1 \subset \sigma^2 \subset \dots$, and their union, the train track σ^∞ .

is a geodesic with one infinite endpoint at p , and the t coordinate gives an arc length parameterization of α_s ; each $\beta_t = \beta^{-1}(t)$ is a horocycle attached to p , and the s coordinate multiplied by e^{-t} gives an arc length parameterization of β_t ; and $(\alpha \times \beta)(s \times t)$ approaches p as $t \rightarrow +\infty$. By isotoping τ , adding a constant to the function β , and multiplying the function α by a constant, if necessary, we may make the following additional assumptions:

- The cusps of N_p are the points $c_i \in \beta_0$ with $(\alpha \times \beta)(c_i) = i \times 0$, for $i \in \mathbf{Z}$.
- The train path $q_i \subset \partial N_p$ between c_i and c_{i+1} has length $< 2 \log(2)$, and so $\beta \mid \partial N_p$ takes values < 2 .
- Each train path q_i intersects the geodesics α_s transversely except at the cusps.
- Letting D_p be the subgroup of $\text{Isom}(\mathbf{H}^2)$ stabilizing the point set $\{c_i \mid i \in \mathbf{Z}\}$, and noting that D_p fixes p and acts as the infinite dihedral group on the horocycle β_0 , we may assume that ∂N_p is stabilized by the action of D_p .

Let $\rho: \mathbf{H}^2 \rightarrow \mathbf{H}^2$ be the hyperbolic isometry which translates along the line α_0 a distance of 2 towards the point p . Note that $\rho(\partial N_p)$ is disjoint from N_p , and the cusp $\rho(c_i) \in \rho(\partial N_p)$ lies on α_{2i} . Let r_i be the geodesic segment from $\rho(c_i)$ to c_{2i} . Now we define

$$\begin{aligned} \sigma^1 &= \sigma^0 \cup \rho(\partial N_p) \cup \bigcup_{-\infty}^{+\infty} r_i \\ &= \sigma^0 \cup \bigcup_{-\infty}^{+\infty} \theta_i \end{aligned}$$

where

$$\theta_i = r_i \cup \rho(q_i) \cup r_{i+1}$$

is a train path in σ^1 . Note that σ^1 is indeed a pretrack, and the compact components of $\mathcal{C}(\tilde{S} - \sigma^1)$ consist of the compact components of $\mathcal{C}(\tilde{S} - \sigma^0)$ together with a sequence of trigons T_i , $i \in \mathbf{Z}$, whose three sides are θ_i , q_{2i} , and q_{2i+1} . Note that the trigons T_{i-1} and T_i intersect along the branch r_i .

We now define σ^n by induction:

$$\sigma^n = \sigma^{n-1} \cup \bigcup_{-\infty}^{+\infty} \rho^{n-1}(\theta_i)$$

and it is straightforward to see that $\sigma^\infty = \bigcup_0^\infty \sigma^n$ satisfies the desired key properties. \diamond

3.4 Ties of a train track

Consider a pretrack τ in S . A *tie bundle* over τ , denoted $\nu \rightarrow \tau$ or $\nu(\tau) \rightarrow \tau$, is an analogue of a normal bundle of a submanifold, embedded into the ambient manifold using the Tubular Neighborhood Existence Theorem of differential topology (e.g. Theorem 5.2 of [Hir76]). When τ is semigeneric the tie bundle is well defined up to isotopy, depending only on the isotopy class of τ itself, just as in the Tubular Neighborhood Uniqueness Theorem (Theorem 5.3 of [Hir76]). When τ is not semigeneric then a tie bundle is not well defined, but depends on finitely many choices for its construction.

A *tie bundle over τ* is a 2-complex ν equipped with a map $f: (S, \nu, S - \nu) \rightarrow (S, \tau, S - \tau)$, satisfying the following conditions:

- $f: S \rightarrow S$ is homotopic to the identity.
- $f: S - \nu \rightarrow S - \tau$ is a homeomorphism.
- For each $x \in \tau$, $f^{-1}(x)$ is a closed arc, called the *tie* over x .
- For each branch b of τ , setting $\mathring{R}_b = f^{-1}(\text{int}(b))$, the restriction $f: \mathring{R}_b \rightarrow b$ is a fiber bundle map.
- For each switch s of τ there is a neighborhood U of s in τ such that the restriction $f|_{f^{-1}(U)}$ has the structure described below (and see Figure 8).

For a precise description of the local structure of a tie bundle over a neighborhood U of a switch s , suppose that s has m branch ends on one side and n branch ends on the other side. Consider the strip $\mathbf{R} \times [-1, +1]$, foliated by vertical arcs $x \times [-1, +1]$. Choose two monotonically increasing sequences

$$y_1^- < \cdots < y_{m-1}^-, \quad y_1^+ < \cdots < y_{n-1}^+ \in (-1, +1)$$

When s is semigeneric, m or n equals 1 and the corresponding sequence is empty, simplifying the construction. Choose C^∞ functions

$$g_1^-, h_1^-, \dots, g_{m-1}^-, h_{m-1}^-: (-\infty, 0] \rightarrow (-1, +1)$$

such that $g_i^-(0) = h_i^-(0) = y_i^-$ for all i and all derivatives vanish at 0, and such that

$$g_1^-(x) < h_1^-(x) < \cdots < g_{m-1}^-(x) < h_{m-1}^-(x), \quad x \in (-\infty, 0)$$

Choose

$$g_1^+, h_1^+, \dots, g_{n-1}^+, h_{n-1}^+: [0, +\infty) \rightarrow (-1, +1)$$

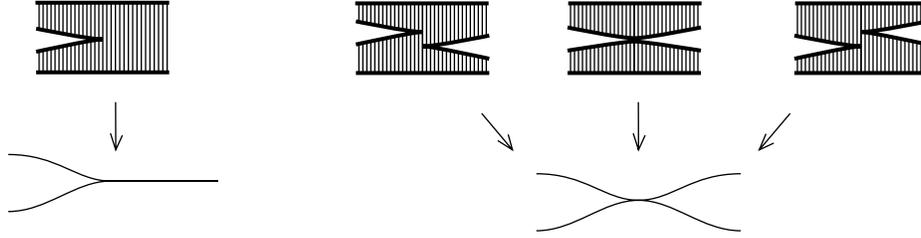


Figure 8: Local models of a tie bundle near a switch s . When s is semigeneric then the local model is unique, but when s is not semigeneric there are finitely many local models up to isotopy, depending on how the cusps on either side of the tie over s are interleaved along that tie. For example, when there are two branch ends on each side of s , and hence one cusp on each side of the tie, then up to isotopy there are three different tie bundles over a neighborhood of s : either of the two cusps can be higher, or their positions can coincide.

similarly with respect to y_j^+ . From the strip $\mathbf{R} \times [-1, +1]$ remove any point (x, y) such that $x < 0$ and $g_i^-(x) < y < h_i^-(x)$ for some i , or such that $x > 0$ and $g_j^+(x) < y < h_j^+(x)$ for some j . The resulting subset of $\mathbf{R} \times [-1, +1]$, decomposed into the components in which the intervals $x \times [-1, +1]$ intersect the subset, is an example of a local model for the tie bundle over U . Figure 8 shows the cases $m = 2, n = 1$ and $m = 2, n = 2$. Once the sequences y_i^-, y_j^+ are picked, the local model is well-defined up to isotopy, independent of the choice of the functions $g_i^-, h_i^-, g_j^+, h_j^+$. However, the choices of the sequences y_i^-, y_j^+ definitely affect the isotopy class of the tie bundle over U . When s is not semigeneric, in which case neither sequence y_i^-, y_j^+ is empty, the isotopy classes of the local model for the tie bundle over U are in one-to-one correspondence with different ways to interleave the two sequences y_i^-, y_j^+ . Since there are finitely many choices for interleaving these two sequences, that leads to finitely many choices for the local model over U . For example if any y_i^- equals any y_j^+ then ν has a pinch at the corresponding point of $f^{-1}(s)$; this shows that in general ν is a pinched subsurface. But as long as there are no coincidences of the form $y_i^- = y_j^+$, it follows from the construction that each of the points $0 \times y_i^-, 0 \times y_j^+$ is a reflex cusp of ν .

To summarize, if τ is semigeneric then its tie bundle ν is a surface with reflex cusps, well defined up to isotopy, which we can take to be a regular neighborhood of τ . The ties forms a boundary transverse foliation of ν whose only singularities are 2-pronged transverse boundary singularities corresponding to the reflex cusps. When τ is semigeneric we are therefore justified in speaking of *the* tie bundle over τ . In general, if we do not assume τ is semigeneric, then there is a finite collection of

tie bundles over τ up to isotopy.

For any tie bundle $\nu \rightarrow \tau$, there is a natural bijection between the reflex cusps of ν and the cusps of τ . To describe this bijection precisely, suppose that we represent a cusp of τ by an “illegal turn” $\alpha * \beta$, where α, β are short train paths, the concatenation $\alpha * \beta$ is not smoothable at the concatenation point, and the path $\alpha * \beta$ pulls back via the overlay map to a path on the boundary of $\mathcal{C}(S - \tau)$; in this case there is a unique lift of $\alpha * \beta$ to ν that lies on the boundary of ν , and the nonsmooth point of this lift corresponds to a reflex cusp of ν .

In the context of a tie bundle ν , we will often refer to reflex cusps of ν simply as cusps of ν .

3.5 Carrying maps between train tracks

Consider train tracks τ, τ' on S . A *carrying map* from τ' to τ is a smooth map $f: (S, \tau') \rightarrow (S, \tau)$ whose restriction to each tangent line of τ' is an isomorphism onto a tangent line of τ , such that there exists a smooth homotopy $H: S \times [0, 1] \rightarrow S$ with $H(x, 0) = x$, $H(x, 1) = f(x)$, and $H \mid S \times t$ a diffeomorphism for each $t < 1$.

We say that τ *carries* τ' , denoted $\tau \succcurlyeq \tau'$, if there exists a carrying map from $f: (S, \tau') \rightarrow (S, \tau)$. Often we will abbreviate the notation and write simply $f: \tau' \rightarrow \tau$ for this carrying map, always assuming that f is ambiently defined and satisfies the above definition.

The following proposition shows that the carrying relation between semigeneric train tracks behaves well with respect to tie bundles:

Proposition 3.5.1. *Let τ, τ' be semigeneric train tracks and let $\nu(\tau), \nu(\tau')$ be tie bundles. The following are equivalent:*

- (1) τ carries τ' .
- (2) The inclusion $\tau' \subset S$ is smoothly isotopic to an inclusion $\tau' \hookrightarrow \text{int}(\nu(\tau))$ whose image is transverse to the ties of $\nu(\tau)$.
- (3) The inclusion $\nu(\tau') \subset S$ is smoothly isotopic to an inclusion $\nu(\tau') \hookrightarrow \text{int}(\nu(\tau))$ such that each tie of $\nu(\tau')$ is a subtie of $\nu(\tau)$.

We extend the notion of carrying from train tracks to tie bundles, by writing $\nu(\tau) \succcurlyeq \nu(\tau')$ if condition (3) holds.

Either of the maps $\tau' \hookrightarrow \nu(\tau)$ or $\nu(\tau') \hookrightarrow \nu(\tau)$ is called a *carrying injection*; the same terminology is used later for measured foliations carried by τ .

When τ is not semigeneric then properties (1), (2), (3) are not equivalent; we will discuss this situation after the proof.

Proof. Since τ is semigeneric, one may construct the tie bundle $\nu(\tau)$ so that $\tau \subset \nu(\tau)$ and τ is transverse to the ties of $\nu(\tau)$ (this is where the proof breaks down when τ is not semigeneric), and similarly for τ' .

To prove $1 \implies 2$, if H is a homotopy from the identity to $f: (S, \tau') \rightarrow (S, \tau)$ as in the definition of a carrying map, then the f_t satisfies the conclusions of 2 for t sufficiently close to 1. The implication $2 \implies 3$ is obvious. To prove $3 \implies 1$, start with the inclusion $\tau' \subset \nu(\tau') \hookrightarrow \nu(\tau)$, and then follow by a homotopy of the inclusion map $\nu(\tau)$ which isotopically shrinks ties until, at the end of the homotopy, one obtains the tie bundle map $\nu(\tau') \rightarrow \nu(\tau)$. \diamond

Note that if τ is not semigeneric then property (2) always fails for $\tau' = \tau$, despite the obvious fact that τ carries itself. Also, if $\tau' = \tau$ then property (3) can hold for some choices of tie bundles and fail for others. It is possible to get a weaker equivalence which works for nonsemigeneric train tracks. For example: $\tau \succcurlyeq \tau'$ if and only if for each tie bundle $\nu(\tau')$ that has no pinch points there exists a tie bundle $\nu(\tau)$ such that $\nu(\tau') \succcurlyeq \nu(\tau)$.

In the proof above, a direct proof that (2) \implies (1) can be seen as follows. Suppose that $\tau' \hookrightarrow \text{int}(\nu(\tau))$ is a carrying injection. For each fiber I of $\nu(\tau)$, let I' be the smallest subinterval of I containing $\tau' \cap I$. Then by collapsing each I' to a point, the image of τ' under this collapsing is isotopic to a subtrack of τ , and hence after this isotopy is performed we obtain a carrying map $f: \tau' \rightarrow \tau$. The intervals I' are called the *fibers* of the carrying map $f: \tau' \rightarrow \tau$.

Given that $\tau \succcurlyeq \tau'$, although carrying maps from τ' to τ need not be unique, the following proposition shows to what extent uniqueness holds:

Proposition 3.5.2 (Homotopy uniqueness of carrying maps). *Consider two train tracks τ, τ' . Given two carrying maps $f_0, f_1: \tau' \rightarrow \tau$, there exists a smooth homotopy of carrying maps from f_0 to f_1 . In particular, $\text{image}(f_0) = \text{image}(f_1)$. Assuming τ is generic, given a tie bundle $\nu \rightarrow \tau$, and given two carrying injections $g_0, g_1: \tau' \rightarrow \nu$, there exists a smooth isotopy of carrying injections from g_0 to g_1 .*

Remark. Even more is true: the space of carrying maps $\tau' \rightarrow \tau$ is contractible, as is the space of carrying injections $\tau' \rightarrow \nu$. We omit the proof.

Corollary 3.5.3. *Given two train tracks τ, τ' such that $\tau \succcurlyeq \tau'$, and given a subtrack $\sigma' \subset \tau'$, there exists a unique subtrack $\sigma \subset \tau$ such that $\sigma \stackrel{F}{\succcurlyeq} \sigma'$.*

Proof. Consider the carrying map $\tau' \rightarrow \tau$. The composition $\sigma' \hookrightarrow \tau' \rightarrow \tau$ is a carrying map, and by applying Proposition 3.5.2 it follows that any other carrying map $\sigma' \rightarrow \tau$ has the same image. \diamond

Proof of Proposition 3.5.2. We shall prove only that f_0, f_1 are homotopic through carrying maps; the proof that g_0, g_1 are isotopic through carrying maps can then be easily derived.

Choose a complete, finite area, hyperbolic metric on S . The universal cover \tilde{S} is identified with the hyperbolic plane \mathbf{H}^2 , with isometric deck transformation group $\pi_1 S$.

Consider each carrying map $f_i: (S, \tau') \rightarrow (S, \tau)$, $i = 0, 1$, and let $H_i: S \times [0, 1] \rightarrow S$ be the homotopy from f_i to the identity on S occurring in the definition of a carrying map. We may adjust the homotopy H_i so that in some neighborhood of each puncture it is stationary. It follows that for each $x \in S$ the length of its track $H_i(x \times [0, 1])$ under the homotopy is uniformly bounded. Let $\tilde{f}_i: \mathbf{H}^2 \rightarrow \mathbf{H}^2$ be the unique lift of f_i whose extension to $\partial\mathbf{H}^2$ is the identity. The homotopy H_i lifts to a homotopy \tilde{H}_i from \tilde{f}_i to the identity, such that the length of the track of each point under \tilde{H}_i is uniformly bounded.

Let $\tilde{\tau}, \tilde{\tau}'$ be the total lifts of τ, τ' in \tilde{S} , and so $\tilde{f}_i, i = 0, 1$, may be regarded as a carrying map $\tilde{\tau}' \rightarrow \tilde{\tau}$.

Consider any train line ℓ in $\tilde{\tau}'$, and the image train lines $\tilde{f}_0(\ell)$ and $\tilde{f}_1(\ell)$ in $\tilde{\tau}$. Since \tilde{f}_0 is homotopic to \tilde{f}_1 by a homotopy whose tracks have bounded length, it follows that the Hausdorff distance between $\tilde{f}_0(\ell)$ and $\tilde{f}_1(\ell)$ is finite. Applying Proposition 3.3.2 it follows that $\tilde{f}_0(\ell) = \tilde{f}_1(\ell)$.

Given a branch b of $\tilde{\tau}'$, choose a train line ℓ through b , and let $\mu = \tilde{f}_0(\ell) = \tilde{f}_1(\ell)$. It follows that $\tilde{f}_0(b), \tilde{f}_1(b)$ are both contained in μ . Let $\mu(b)$ be the smallest closed, connected subset of μ containing both $\tilde{f}_0(b)$ and $\tilde{f}_1(b)$. We can now define a homotopy from the map $\tilde{f}_0|_b$ to the map $\tilde{f}_1|_b$ which, at each moment, maps b into $\mu(b)$ by a map which has derivative = 1 near each endpoint of b . This can be done independently of the train line ℓ , because if we choose any other train line ℓ' through b , and let $\mu' = \tilde{f}_0(\ell') = \tilde{f}_1(\ell')$, then $\mu \cap \mu'$ is connected by Proposition 3.3.1, and it follows that $\mu(b) = \mu'(b)$. This can moreover be done equivariantly with respect to the action of $\pi_1 S$.

The condition that the homotopy from $\tilde{f}_0|_b$ to $\tilde{f}_1|_b$ have derivative = 1 near the endpoints of b implies that as b varies over all branches of $\tilde{\tau}'$, the homotopies piece together to give a smooth homotopy $\tilde{\tau}' \times [0, 1] \rightarrow \tilde{\tau}$ from the map $\tilde{f}_0|_{\tilde{\tau}'}$ to the map $\tilde{f}_1|_{\tilde{\tau}'}$. Moreover this homotopy is $\pi_1 S$ equivariant. It is now a simple matter to extend this to a $\pi_1 S$ equivariant homotopy $\tilde{H}: \tilde{S} \times [0, 1] \rightarrow \tilde{S}$ through carrying maps, which descends to a homotopy $H: S \times [0, 1] \rightarrow S$ that proves the proposition. \diamond

3.6 Measured foliations carried on train tracks

Consider a train track τ , a tie bundle $\nu(\tau) \rightarrow \tau$, and a partial measured foliation \mathcal{F} . A *carrying injection* from \mathcal{F} to $\nu(\tau)$ is a homeomorphism $f: S \rightarrow S$ isotopic to the identity such that $f(\mathcal{F}) \subset \nu(\tau)$ with all leaves of $f(\mathcal{F})$ transverse to the ties of $\nu(\tau)$. We will often denote a carrying injection as $\mathcal{F} \hookrightarrow \nu(\tau)$ and identify \mathcal{F} with its image. Note that when τ is not semigeneric the partial measured foliation \mathcal{F} may have a carrying injection into some tie bundles over τ but not over others.

A partial measured foliation \mathcal{F} is *carried by* a train track τ if there exists a tie bundle $\nu(\tau) \rightarrow \tau$ and a carrying injection $\mathcal{F} \hookrightarrow \nu(\tau)$; in this case the composition $\mathcal{F} \hookrightarrow \nu(\tau) \rightarrow \tau$ is called a *carrying map* from \mathcal{F} to τ . Note that the image of every leaf segment of \mathcal{F} under a carrying map $\mathcal{F} \rightarrow \tau$ is a train path in τ .

A class $[\mathcal{F}] \in \mathcal{MF}$ is carried by a train track τ if some representative \mathcal{F} is carried by τ . We shall occasionally abuse terminology and say that a partial measured foliation \mathcal{F} is carried by τ if the class of \mathcal{F} is carried by τ , equivalently, if \mathcal{F} is equivalent to some \mathcal{F}' which has a carrying injection into some tie bundle over τ .

Given a partial measured foliation \mathcal{F} , a carrying map $\mathcal{F} \rightarrow \tau$ is said to be *full* if and only if it is surjective. For a carrying injection $\mathcal{F} \hookrightarrow \nu(\tau)$, we distinguish between a *full carrying injection* which means that the image intersects every tie, and a *carrying bijection* which is a surjective carrying injection. A carrying bijection will be denoted $\mathcal{F} \hookrightarrow \nu(\tau)$. A carrying injection $\mathcal{F} \hookrightarrow \nu(\tau)$ is said to be *canonical* if \mathcal{F} is a canonical model and the inclusion is surjective; in this case, the singular set of \mathcal{F} equals the set of reflex cusps of $\nu(\tau)$, and \mathcal{F} has no proper saddle connections. If there exists a canonical carrying injection $\mathcal{F} \hookrightarrow \nu(\tau)$ then we say that \mathcal{F} is canonically carried on τ ; this implies that the singularity types of \mathcal{F} and of τ are identical.

Existence of carrying. The next proposition shows that a train track always exists carrying a given class in \mathcal{MF} .

Proposition 3.6.1. *For every measured foliation \mathcal{F} there exists an equivalent partial measured foliation \mathcal{F}' and a train track τ carrying \mathcal{F}' . In fact, if \mathcal{F}' is a canonical model then there exists a train track τ that canonically carries \mathcal{F}' .*

Proof. We'll give two somewhat similar proofs for this proposition; for another proof see Corollary 1.7.6 of [Pen92].

Let \mathcal{F}' be a canonical model for \mathcal{F} . For each component \mathcal{F}'_i we produce a train track $\tau_i \subset \text{Supp}(\mathcal{F}'_i)$ canonically carrying \mathcal{F}'_i , and it follows that $\tau = \cup_i \tau_i$ canonically carries \mathcal{F}' . If \mathcal{F}'_i is annular then τ_i is a simple closed curve isotopic to any leaf of \mathcal{F}'_i .

Thus we have reduced to the case of an arational measured foliation \mathcal{F} . We may assume that \mathcal{F} has no saddle connections by collapsing them all. Choose a transverse arc β disjoint from the singularities. Let Σ be a union of finite separatrices of \mathcal{F}

with the property that $\Sigma \cap \text{int}(\beta) = \partial\Sigma$, and that Σ is maximal with respect to this property. This completely determines Σ , and Σ can be constructed explicitly by starting off along each separatrix of \mathcal{F} until that separatrix first hits $\text{int}(\beta)$, which must happen because of arationality of \mathcal{F} . Let x, y be the outermost points on β which lie in Σ . It may or may not happen that $\partial\beta = \{x, y\}$; if not, truncate β so that it does happen. Let \mathcal{F}' be the partial measured foliation obtained from \mathcal{F} by slicing open along Σ , and note that \mathcal{F}' is a canonical model for \mathcal{F} . Note that β pulls back to a transverse arc for \mathcal{F}' whose endpoints lie on $\partial\text{Supp}(\mathcal{F}')$ and which contains each cusp of \mathcal{F}' . Each component of $\text{Supp}(\mathcal{F}') - \beta$ is a rectangle minus its vertical sides, and it can be completed to a compact rectangle by adding segments of β ; these two vertical sides may overlap on β , and so the compact rectangle may be only immersed in $\text{Supp}(\mathcal{F}')$, not embedded. We thus obtain a finite collection of rectangles R_1, \dots, R_k . Choose a vertical foliation of each rectangle, thus defining a boundary transverse “vertical foliation” of $\text{Supp}(\mathcal{F}')$ by arcs transverse to \mathcal{F}' , with a 2-pronged boundary transverse singularity at each singularity of \mathcal{F}' . The desired train track τ canonically carrying \mathcal{F}' is obtained by collapsing to a point each vertical leaf of \mathcal{F}' , and so the surface $\text{Supp}(\mathcal{F}')$ equipped with its vertical foliation may be regarded as a tie bundle ν over τ with a canonical carrying injection $\mathcal{F}' \hookrightarrow \nu$.

Notice in the above construction of τ that the image of β under the carrying map $\mathcal{F}' \rightarrow \tau$ is a nongeneric switch of τ , and indeed it is the only switch of τ . Also, the image of each rectangle R_i is a source branch of τ , and these are all of the branches of τ . If the unique switch of τ is slightly uncombed then τ becomes a “one-sink train track” in the sense of Section 9.4.

Here is another proof, at least of the first sentence, as illustrated in Figure 9. We will use the Gardiner–Masur Theorem 2.5.1 which says that every measured foliation \mathcal{F} can be altered up to Whitehead equivalence so that \mathcal{F} is the horizontal foliation of a quadratic differential q . Let $\mathcal{F}^v(q)$ denote the vertical foliation of q . Define a *horizontal separatrix system* or *h-separatrix system* of q to be a finite union of separatrices of \mathcal{F} , denoted ξ . We say that ξ is *v-filling* if every closed leaf and every half-infinite leaf of $\mathcal{F}^v(q)$ intersects ξ . For example, in each arational component of \mathcal{F} choose an infinite separatrix and then choose a sufficiently long finite subseparatrix, and take the union of those separatrices with all the saddle connections of \mathcal{F} to obtain ξ ; if the finite separatrices in arational components are sufficiently long then, by density of the infinite separatrices, it follows that ξ is v-filling. Applying Fact 2.6.2, let $\mathcal{F}(\xi)$ to be the partial measured foliation obtained from \mathcal{F} by slicing along ξ , with partial fulfillment map $r: (S, \mathcal{F}(q, \sigma)) \rightarrow (S, \mathcal{F}(\xi))$. Letting $\nu = \text{Supp}(\mathcal{F}(q, \xi))$, by pulling back the leaves of $\mathcal{F}^v(q)$ under r we obtain a boundary transverse foliation of ν by closed arcs, called the *vertical foliation*, whose singularities are all 2-pronged boundary transverse singularities, coinciding with the

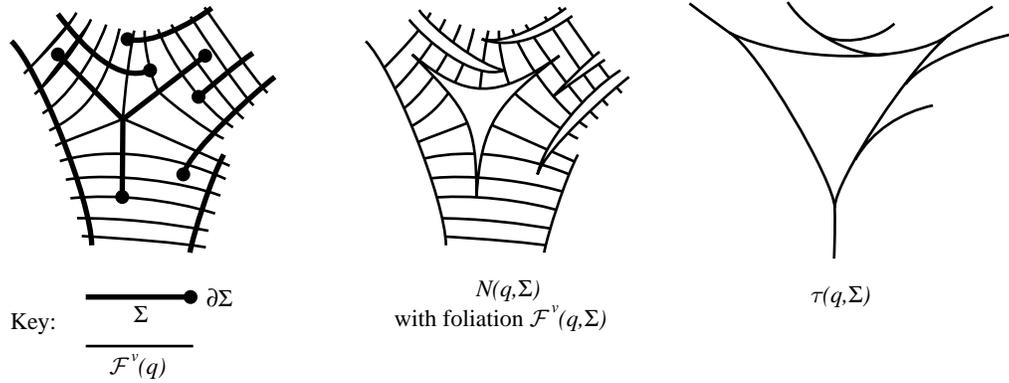


Figure 9: The train track associated to a quadratic differential q and a ν -filling h-separatrix system Σ .

singularities of $\mathcal{F}(\xi)$. Now collapse each leaf of the vertical foliation to a point to get the desired train track τ , so ν is a tie bundle τ and $\mathcal{F}(q, \xi)$ has a carrying bijection in ν .

◇

Uniqueness of carrying maps. For any class in \mathcal{MF} , there are infinitely many train tracks carrying that class that are pairwise distinct up to isotopy, and even up to comb equivalence. Fixing the train track τ , we are interested in uniqueness properties of the carrying map to τ , or of the carrying injection to a tie bundle over τ . For example, having defined the statement “ \mathcal{F} is carried by τ ” to mean that “ \mathcal{F} is carried by some tie bundle $\nu(\tau)$ ”, there are several questions: is the tie bundle uniquely determined? If \mathcal{F} has two different carrying maps to the same tie bundle, how are those maps related? If \mathcal{F} has two different carrying maps to the same train track, how are the images related? How are carrying maps of non-isotopic but equivalent measured foliations related? The next few results will answer these questions.

Consider a partial measured foliation \mathcal{F} , a train track τ , and a carrying map $f: \mathcal{F} \rightarrow \tau$. Each leaf ℓ of \mathcal{F} determines, by mapping via f , a train path on τ , either a closed train path when ℓ is compact, or a bi-infinite train path when ℓ is not compact. The following proposition says that this set of train paths depends only on the equivalence class of \mathcal{F} .

Proposition 3.6.2 (Uniqueness of train paths). *Given equivalent partial measured foliations $\mathcal{F}, \mathcal{F}'$, a train track τ , and carrying maps $f: \mathcal{F} \rightarrow \tau, f': \mathcal{F}' \rightarrow \tau$, the*

set of train paths determined by f is the same as the set of train paths determined by f' .

Proof. We argue up in the universal cover \tilde{S} . The bijection between leaves of \mathcal{F} and leaves of \mathcal{F}' discussed in Section 2.6 lifts to a bijection between leaves of $\tilde{\mathcal{F}}$ and leaves of $\tilde{\mathcal{F}}'$. Moreover, it is clear that corresponding leaves ℓ of $\tilde{\mathcal{F}}$ and ℓ' of $\tilde{\mathcal{F}}'$ have finite Hausdorff distance, because under a Whitehead move or a partial fulfillment, a leaf is moved a finite distance to its corresponding leaf. The maps f and f' also move points a finite distance. Thus, $\tilde{f}(\ell)$ and $\tilde{f}'(\ell')$ are bi-infinite train paths in $\tilde{\tau}$ at finite Hausdorff distance. Applying Proposition 3.3.2 it follows that $\tilde{f}(\ell) = \tilde{f}'(\ell')$, that is, the set of train paths in $\tilde{\tau}$ determined by \tilde{f} is the same as the set of train paths determined by \tilde{f}' . The same statement downstairs is an immediate consequence. \diamond

The next corollary says that the image in a train track τ of a class in \mathcal{MF} carried on τ is well defined, independent of the choice of a carrying map $\mathcal{F} \rightarrow \tau$ where \mathcal{F} represents the given class. This corollary will be used repeatedly throughout the paper.

Corollary 3.6.3 (Uniqueness of carrying image). *Under the same hypotheses as Proposition 3.6.2, the images of f, f' are the same.*

Proof. The image of f is equal to the union of the train paths determined by f , and similarly for f' . Now apply Proposition 3.6.2. \diamond

The image of a partial measured foliation \mathcal{F} carried by τ is obviously a subtrack of τ , and it is called the *carrying image* of \mathcal{F} in τ . In particular, we can define \mathcal{F} , or its class in \mathcal{MF} , to be *fully carried* by τ if some carrying map is surjective, and this is well-defined independent of the carrying map.

The next corollary says that, for a class in \mathcal{MF} fully carried on a train track τ , the tie bundle ν over τ that surjectively carries that class is well-defined.

Corollary 3.6.4 (Uniqueness of carrying tie bundle). *Let τ be a train track, ν, ν' two tie bundles of τ , and $\mathcal{F}, \mathcal{F}'$ equivalent partial measured foliations (possibly with pinch points), and suppose that there exist carrying bijections $\mathcal{F} \hookrightarrow \nu, \mathcal{F}' \hookrightarrow \nu'$. Then ν, ν' are isotopic tie bundles.*

Proof. Consider a nongeneric switch s of τ , choose a transverse orientation for τ at s , and list the branch ends on the two sides of s , in increasing order, as e_1, \dots, e_M and f_1, \dots, f_N . Let t, t' be the ties in ν, ν' lying over s , with orientations determined by the choice of transverse orientation on s . Consider a cusp (e_i, e_{i+1}) on one side of s and a cusp (f_j, f_{j+1}) on the other side of s . These two cusps of τ lift, via the tie

bundle map $\nu \rightarrow \tau$, to cusps of ν whose relative position depends solely on the train paths in τ determined by \mathcal{F} . For example, the lift of (e_i, e_{i+1}) is greater than the lift of (f_j, f_{j+1}) in t if and only if some leaf segment of \mathcal{F} induces a train path of τ of the form $e_m * f_n$ where $1 \leq m \leq i$ and $j+1 \leq n \leq N$. The same is true in ν' —the relative position of the lifts to ν' of the cusps (e_i, e_{i+1}) and (f_j, f_{j+1}) depend in the same way on the train paths in τ determined by \mathcal{F}' . But Proposition 3.6.2 says that $\mathcal{F}, \mathcal{F}'$ determine the same train paths in τ , and hence cusps in τ lift with the same positions in ν, ν' , implying that ν, ν' are isotopic. \diamond

Proposition 3.6.2 and its corollaries settle uniqueness questions regarding the target of a carrying map, and now we turn to the uniqueness of the carrying map itself.

Proposition 3.6.5. *Consider a train track τ , a class $[\mathcal{F}] \in \mathcal{MF}$ carried on τ , and a tie bundle ν over τ .*

- (1) *If \mathcal{F} is a representative of $[\mathcal{F}]$ and if $f_0, f_1: \mathcal{F} \hookrightarrow \nu$ are carrying injections, then f_0, f_1 are isotopic through carrying injections $f_t: \mathcal{F} \hookrightarrow \nu$, $t \in [0, 1]$.*
- (2) *Furthermore, if both f_0, f_1 are carrying bijections, then the isotopy can be chosen so that f_t is surjective; in fact, replacing \mathcal{F}_i by $f_i(\mathcal{F}_i)$, so that $\text{Supp}(\mathcal{F}_i) = \nu$, there is an isotopy from \mathcal{F}_0 to \mathcal{F}_1 that is the identity on $S - \text{int}(\nu)$ and that preserves each tie of ν . As a consequence, if f_0, f_1 are canonical carrying injections then the isotopy can be chosen so that f_t is canonical.*
- (3) *If $\mathcal{F}_0, \mathcal{F}_1$ are representatives of $[\mathcal{F}]$ (possibly with pinch points), and if $f_0: \mathcal{F}_0 \hookrightarrow \nu$ and $f_1: \mathcal{F}_1 \hookrightarrow \nu$ are carrying bijections, then $\mathcal{F}_0, \mathcal{F}_1$ are isotopic (and so the conclusions of (2) hold as well).*

For a simple proof of (1) when \mathcal{F} is an essential closed curve, see Proposition 3.7.3 below.

Proof of Proposition 3.6.5 (1,2). This proof is similar to Proposition 3.5.2. We work up in the universal cover \tilde{S} , and by equipping S with a complete, finite area hyperbolic structure we may identify $\tilde{S} \approx \mathbf{H}^2$. We have a lifted train track $\tilde{\tau}$, tie bundle $\tilde{\nu}$, and partial measured foliation $\tilde{\mathcal{F}}$ in \tilde{S} . Choose a tie bundle map $\pi: (S, \nu) \rightarrow (S, \tau)$ which is the identity in a neighborhood of each puncture, and choose a lift $\tilde{\pi}$ that is the identity on the circle at infinity and so moves points a uniformly bounded amount. The maps f_0, f_1 can be chosen to be the identity near each puncture, and choosing the lifted maps \tilde{f}_0, \tilde{f}_1 to be the identity in the circle at infinity, it follows that \tilde{f}_0, \tilde{f}_1 move points a uniformly bounded amount. All isotopies that we construct will be equivariant, so that they descend to isotopies on S .

It is convenient to replace $\tilde{\mathcal{F}}$ by $\tilde{f}_0(\tilde{\mathcal{F}})$, so that \tilde{f}_0 becomes the identity map Id. Note that for each leaf ℓ of \mathcal{F} the bi-infinite train paths $\tilde{\pi}(\ell)$, $\tilde{\pi}(\tilde{f}_1(\ell))$ in $\tilde{\tau}$ are identical, by Proposition 3.3.2.

We will define an equivariant isotopy between Id and \tilde{f}_1 , which will descend to the desired isotopy in S . The isotopy is carried out in four steps; it will be convenient after each step to replace $\tilde{\mathcal{F}}$ by its isotoped image.

For the proof of (2), it will usually be clear that if f_0, f_1 are surjective then each f_t is surjective; in a couple of the steps we will supply details.

Step 1: Proper saddle connections. Let α be a proper saddle connection of $\tilde{\mathcal{F}}$. We isotope $\tilde{\mathcal{F}}$ equivariantly, with the effect that after the isotopy we have $\tilde{\pi}(\alpha) = \tilde{\pi}\tilde{f}_1(\alpha)$.

There are exactly two leaves λ, μ of $\tilde{\mathcal{F}}$ containing α , and we have $\lambda \cap \mu = \alpha$. Note that we have a train path $\gamma = \tilde{\pi}(\lambda) \cap \tilde{\pi}(\mu) = \tilde{\pi}f_1(\lambda) \cap \tilde{\pi}f_1(\mu)$. The path γ contains both of the train paths $\tilde{\pi}(\alpha)$ and $\tilde{\pi}\tilde{f}_1(\alpha)$. Let β be the smallest subsegment of γ containing both $\tilde{\pi}(\alpha)$ and $\tilde{\pi}\tilde{f}_1(\alpha)$, let $\beta_\lambda = \lambda \cap \pi^{-1}(\beta)$, and let $\beta_\mu = \mu \cap \pi^{-1}(\beta)$. Let N_α be the union of all subsegments of ties of $\nu(\tilde{\tau})$ with one endpoint on β_λ and one endpoint on β_μ ; some of these segments degenerate to a point on α . Now define an isotopy of $\tilde{\mathcal{F}}$ which is supported in a small neighborhood of N_α , after which $\tilde{\pi}\alpha = \tilde{\pi}\tilde{f}_1(\alpha)$. This can be done equivariantly, by choosing the sets N_α to be equivariant and pairwise disjoint as α varies over all proper saddle connections.

In the case of carrying bijections, note that $\gamma = \tilde{\pi}(\alpha) = \tilde{\pi}\tilde{f}_1(\alpha) = \beta$, and so the isotopy can be taken to move points along ties and to be the identity on $\partial\tilde{\nu}$, and so each moment of the isotopy defines a carrying bijection from $\mathcal{F} \hookrightarrow \nu$.

Step 2: Singularities. Consider a singularity s of $\tilde{\mathcal{F}}$, which must be a tangential 3-pronged boundary singularity. We isotope $\tilde{\mathcal{F}}$ equivariantly, with the effect that after the isotopy the points s and $\tilde{f}_1(s)$ coincide with a certain reflex cusp of $\tilde{\nu}$.

The procedure is similar to step 1: there are exactly two leaves λ, μ of $\tilde{\mathcal{F}}$ containing s . Consider $\alpha = \lambda \cap \mu$. It is possible that α is a proper saddle connection with endpoint s , in which case we are done by step 1. The other possibility is that α is an infinite separatrix based at s , in which case $\gamma = \tilde{\pi}(\lambda) \cap \tilde{\pi}(\mu) = \tilde{\pi}f_1(\lambda) \cap \tilde{\pi}f_1(\mu)$ is an infinite train path ending at a cusp of $\tilde{\tau}$ that lifts to a certain cusp c of $\tilde{\nu}$. We can now proceed as in step 1 to define the required isotopy. In the case of carrying bijections, we have $s = \tilde{f}_1(s) = c$ and the isotopy can be chosen to move points along ties.

Step 3: Boundary saddle connections. With Steps 1 and 2 completed, it follows automatically that for each boundary saddle connection α of \mathcal{F} , we have $\tilde{\pi}\alpha = \tilde{\pi}f_1(\alpha)$, and so we can isotop points along ties to take α to $f_1(\alpha)$.

Step 4: Taking leaves to leaves. For each fiber I of $\nu(\tilde{\tau})$, define an isotopy of I with the effect that for each leaf ℓ of \mathcal{F} that intersects I , the point $\ell \cap I$ isotopes to the point $\tilde{f}_1(\ell) \cap I$. We may do this so that each component of $I - \mathcal{F}$ maps affinely to a component of $I - \tilde{f}_1(\mathcal{F})$. The existence and continuity of this isotopy follows from Steps 1 and 2. As constructed, this isotopy is equivariant. After the isotopy, the effect is that for each leaf ℓ of \mathcal{F} we have $\ell = \tilde{f}_1(\ell)$.

Step 5: Fixing up each leaf. After steps 1–4, For each leaf ℓ of \mathcal{F} , the map \tilde{f}_1 takes ℓ to itself. Define a straight-line isotopy of ℓ , isotoping each point $x \in \ell$ to its image $\tilde{f}_1(x) \in \ell$. These isotopies piece together continuously and equivariantly. \diamond

Proof of Proposition 3.6.5 (3). Let $\mathcal{F}_0, \mathcal{F}_1$ be equivalent partial measured foliations with support ν and transverse to the ties. The key point is to show that the proper saddle connections of $\mathcal{F}_0, \mathcal{F}_1$ agree in the appropriate sense. This will be accomplished by slicing along proper saddle connections to obtain canonical models, and then studying how the canonical models sit inside ν in order to identify the saddle connections of $\mathcal{F}_0, \mathcal{F}_1$.

Let \mathcal{F}'_i be the canonical model obtained from \mathcal{F}_i by slicing open along all proper saddle connections, and so we have carrying injections $\mathcal{F}'_0, \mathcal{F}'_1 \hookrightarrow \nu$. Each boundary singularity s of \mathcal{F}'_i came from a boundary singularity of \mathcal{F}_i which was not an endpoint of a proper saddle connection, and so s was unmoved by the slicing operations, implying that s coincides with a reflex cusp of ν .

By uniqueness of canonical models, the partial measured foliations $\mathcal{F}'_0, \mathcal{F}'_1$ are isotopic, and so we can apply steps 1–5 used in the proof of part (1). Note that in step 1 there is nothing to do. In step 2, since each singularity already coincides with a reflex cusp of ν , there is still nothing to do. Thus, we see that the singular sets of $\mathcal{F}'_0, \mathcal{F}'_1$ are identical, coinciding with a certain set of reflex cusps of ν .

Before continuing, do a small isotopy of \mathcal{F}'_0 and \mathcal{F}'_1 through carrying injections that pulls their images into the interior of ν , and so that they still have the same singularity sets.

Now continue with step 3, isotoping \mathcal{F}'_0 and \mathcal{F}'_1 so that their boundary saddle connections are identical; clearly this can be done without moving the singularity sets. Once this is done, it follows that $\text{Supp}(\mathcal{F}'_0) = \text{Supp}(\mathcal{F}'_1)$, a subsurface contained in the interior of ν which we shall denote σ .

Consider the surface $\text{Cl}(\nu - \sigma)$, whose boundary is the disjoint union of $\partial\nu$ and $\partial\sigma$, with a reflex cusp of $\text{Cl}(\nu - \sigma)$ at each reflex cusp of ν , and a cusp of $\text{Cl}(\nu - \sigma)$ at each reflex cusp of σ . The surface $\text{Cl}(\nu - \sigma)$ is foliated by tie segments of ν , with singularities at the cusps and reflex cusps. If we cut $\text{Cl}(\nu - \sigma)$ along the tie segments passing through the reflex cusps of $\text{Cl}(\nu - \sigma)$, the closures of the complementary components are of three types, *boundary rectangles*, *cuspid triangles*, and *fake saddle connections*. A *boundary rectangle* of $\text{Cl}(\nu - \sigma)$ is a rectangle foliated by ties with one horizontal side coinciding with a side of ν , and the other horizontal side lying on $\partial\sigma$. A *cuspid triangle* is a triangle with one vertex at a cusp s of $\text{Cl}(\nu - \sigma)$, so that the side opposite s is a tie segment whose interior contains a reflex cusp c of ν , the two sides adjacent to s lying on $\partial\sigma$, and the cuspid triangle is foliated by tie segments with a singularity at s . A *fake saddle connection* is a rectangle foliated by tie segments, each of whose vertical sides has interior containing a reflex cusp of ν , and whose horizontal sides each lie on $\partial\sigma$.

To summarize, starting from $\mathcal{F}_0, \mathcal{F}_1$, we have sliced open along proper saddle connections to obtain canonical models, pulled the canonical models away from the boundary into the interior of ν , and then isotoped along ties so that the canonical models have the same support σ . By tracing through these steps it is clear that each proper saddle connection of \mathcal{F}_i corresponds to a fake saddle connection of $\text{Cl}(\nu - \sigma)$. To be precise, if in each fake saddle connection we construct a horizontal path that connects the two reflex cusps, then the union of these paths over all fake saddle connections can be isotoped along ties to coincide with the proper saddle connections of \mathcal{F}_0 , and can also be isotoped to coincide with the proper saddle connections of \mathcal{F}_1 . It follows that we can isotop \mathcal{F}_0 along ties so that its proper saddle connections are taken to the proper saddle connections of \mathcal{F}_1 .

Once this is accomplished, we can then carry out step 4 of (1) to produce an isotopy along ties that takes \mathcal{F}_0 to \mathcal{F}_1 . \diamond

3.7 Invariant weights on train tracks

By definition, carrying of a partial measured foliation \mathcal{F} by a train track τ is a property of the equivalence class of \mathcal{F} , and so it is well-defined to say that the class $[\mathcal{F}] \in \mathcal{MF}$ is carried by τ ; to be precise, this means that there exists a representative \mathcal{F} of $[\mathcal{F}]$, a tie bundle ν over τ , and a carrying injection $\mathcal{F} \hookrightarrow \nu$. Let $\mathcal{MF}(\tau)$ denote the subset of \mathcal{MF} consisting of all classes of partial measured foliations carried by τ . Let $\mathcal{PMF}(\tau) \subset \mathcal{PMF}$ denote the projective image of $\mathcal{MF}(\tau)$.

Given a train track τ we now review the parameterization of $\mathcal{MF}(\tau)$ using “invariant weights” on τ .² There is a corresponding projective parameterization of

²Invariant weights are called “transverse measures” in [Pen92] and elsewhere, though the usage is inconsistent. They are called “edge weight systems” in [Bon99].

$\mathcal{PMF}(\tau)$.

Let $B(\tau)$ be the set of branches and $\text{Sw}(\tau)$ the set of switches of τ . A *weight function* on τ is a function $w: B(\tau) \rightarrow [0, \infty)$, i.e. an element $w \in [0, \infty)^{B(\tau)}$. For each switch $s \in \text{Sw}(\tau)$, label the two sides of s arbitrarily as $+$ and $-$, and let b_1^+, \dots, b_k^+ and b_1^-, \dots, b_l^- be the lists of branches with ends on the corresponding sides of s ; the number of times a branch occurs in one of these lists equals the number of ends of that branch on the corresponding side of s . The *switch condition* of s is the equation

$$w(b_1^+) + \dots + w(b_k^+) = w(b_1^-) + \dots + w(b_l^-)$$

If a weight function $w \in [0, \infty)^{B(\tau)}$ satisfies the switch conditions for all $s \in \text{Sw}(\tau)$ then w is called an *invariant weight* on τ . Alternatively, an invariant weight is a function $w: \tau \rightarrow [0, \infty)$ which is constant on the interior of each branch, such that for each switch s and for each side of s , if b_1, \dots, b_k are the branches on that side, listed once for each branch end on that side, then $w(s) = w(b_1) + \dots + w(b_k)$.

The set of invariant weights on τ is denoted $W(\tau) \subset [0, \infty)^{B(\tau)}$. Since $W(\tau)$ is defined by a finite set of linear equations, it follows that $W(\tau)$ is a closed, finite sided, polyhedral subcone of the cone $[0, \infty)^{B(\tau)}$. The image of $W(\tau)$ under the projectivization map $[0, \infty)^{B(\tau)} \rightarrow \mathcal{P}[0, \infty)^{B(\tau)}$ is denoted $\mathcal{PW}(\tau)$; this is a closed, finite-sided subpolyhedron of the simplex $\mathcal{P}[0, \infty)^{B(\tau)}$. The face lattice of $W(\tau)$ and of $\mathcal{PW}(\tau)$ is explained below in Proposition 3.11.3.

For example, given a partial measured foliation \mathcal{F} , a tie bundle $\nu(\tau)$, and a carrying injection $\mathcal{F} \subset \nu(\tau)$, there is an induced invariant weight $w_{\mathcal{F}} \in W(\tau)$ whose value at any $x \in \tau$ is the \mathcal{F} -transverse measure of the tie over x . Note that \mathcal{F} is fully carried if and only if $w_{\mathcal{F}}$ has positive values.

Now we define a map $W(\tau) \rightarrow \mathcal{MF}$.

Consider $w \in W(\tau)$ and suppose at first that $w(b) > 0$ for all branches b of τ . Choose a tie bundle $p: \nu(\tau) \rightarrow \tau$, and for each b choose a horizontal measured foliation on $R_b = \text{Cl}(\overset{\circ}{R}_b)$ transverse to the tie bundle so that each tie in $\overset{\circ}{R}_b$ has transverse measure $w(b)$, and so that for each switch s , the Borel measures on either side of the tie $p^{-1}(s)$ match up along $p^{-1}(s)$ to give a positive Borel measure of total mass $w(s)$. The weights $w(b)$ together with the matching condition determines the structure of the tie bundle over a neighborhood of each switch. The horizontal measured foliations in the rectangles R_b piece together to define a horizontal partial measured foliation \mathcal{F}_w with support $\nu(\tau)$, whose singular set is identical to the singular set of the tie bundle of $\nu(\tau)$ (note that $\nu(\tau)$ and \mathcal{F}_w may have pinch points). In the unpinched case, the proof that \mathcal{F}_w is a partial measured foliation follows from the observation that $\mathcal{C}(S - \tau)$ and $\mathcal{C}(S - \text{Supp}(\mathcal{F}_w))$ have identical structures as cusped surfaces, cusps of $\mathcal{C}(S - \tau)$ corresponding to 3-pronged tangential boundary

singularities of \mathcal{F}_w ; the fact that τ is a bigon track then allows one to verify the requirements on components of $\mathcal{C}(S - \text{Supp}(\mathcal{F}_w))$ in order for \mathcal{F}_w to be a partial measured foliation. For example, a k -punctured disc component of $\mathcal{C}(S - \tau)$ with n -cusps on the boundary leads to a k -punctured disc component of $\mathcal{C}(S - \mathcal{F}_w)$ along whose boundary the foliation \mathcal{F}_w has exactly n 3-pronged boundary singularities.

Consider next the general case where w is not necessarily positive on all branches of τ . The union of branches on which w is positive is a subtrack of τ . The construction of \mathcal{F}_w can therefore be carried out in a tie bundle over this subtrack.

The construction of the partial measured foliation \mathcal{F}_w is unique up to isotopy, depending only on w . We therefore have defined a map $W(\tau) \rightarrow \mathcal{MF}$.

Theorem 3.7.1. *For any train track τ , the map $W(\tau) \rightarrow \mathcal{MF}$ is an embedding with image $\mathcal{MF}(\tau)$, and the induced map $\mathcal{PW}(\tau) \rightarrow \mathcal{PMF}$ is an embedding with image $\mathcal{PMF}(\tau)$.* \diamond

Proof. We prove injectivity, referring the reader to [Pen92] Theorem 2.7.4, and [Bon99] for a proof of continuity.

To prove the theorem it suffices to show that if $\mathcal{F}_0, \mathcal{F}_1$ are two partial measured foliations with carrying injections $\mathcal{F}_i \subset \nu(\tau)$, if $w_i \in W(\tau)$ is the invariant weight induced by \mathcal{F}_i , and if $\mathcal{F}_0, \mathcal{F}_1$ are Whitehead equivalent, then $w_0 = w_1$. By slicing $\mathcal{F}_0, \mathcal{F}_1$ along their proper saddle connections, we may assume that they are canonical models, and so they are isotopic.

If τ is semigeneric, then $\nu(\tau)$ with its tie bundle is uniquely defined, and so we may apply Proposition 3.6.5 to conclude that there is an isotopy from \mathcal{F}_0 to \mathcal{F}_1 through carrying injections, implying that $w_0 = w_1$.

If τ is not semigeneric, we may easily perturb it to be semigeneric as described later, without altering $W(\tau)$ or the map $W(\tau) \rightarrow \mathcal{MF}$. \diamond

Corollary 3.7.2. *For any train track τ , $\mathcal{PMF}(\tau)$ is a compact subset of \mathcal{PMF} .* \diamond

Remark on the proofs of Proposition 3.6.5 and Theorem 3.7.1. In the case of an essential simple closed curve inducing an integer invariant measure, there are particularly simple proofs using an index argument, which do not require the arguments in the universal cover. We record this here:

Proposition 3.7.3. *For any train track τ , if c, c' are essential simple closed curves with carrying injections $c, c' \subset \nu(\tau)$, and if c, c' are isotopic, then they are isotopic through carrying injections. It follows that if c, c' induce weights $w, w' \in W(\tau)$ then $w = w'$.*

Proof. Perturb the carrying injections $c, c' \hookrightarrow \nu(\tau)$ so that c, c' intersect transversely. We proceed by induction on the cardinality of $c \cap c'$.

In the case when $c \cap c' = \emptyset$, it follows that $c \cup c'$ bounds an annulus A . When $c \cap c' \neq \emptyset$ then there is a *bigon*, a nonpunctured disc d whose boundary is a union of two arcs $\alpha \subset c, \alpha' \subset c'$ so that $\alpha \cap \alpha' = \partial\alpha = \partial\alpha'$; by making the intersection points of c and c' tangential, we may regard the bigon as a disc with two cusps.

Claim: in either case, the annulus A or the bigon d is contained in $\nu(\tau)$, with ties crossing from one side to the other. For the annulus case this immediately implies $w = w'$. For the bigon case this implies that c can be isotoped across the bigon, reducing the cardinality of $c \cap c'$, preserving the induced weight w . By induction we are done.

The proof of the claim in both cases is an index argument. The Euler indices of A and d are both zero, and their boundaries have neighborhoods contained in $\nu(\tau)$ with the tie bundle transverse to the boundaries. It follows that the sum of the indices of the components of $\text{Cl}(A - \nu(\tau))$ or of $\text{Cl}(d - \nu(\tau))$ equals zero. However, each of these components is a component of the closure of $S - \nu(\tau)$, and the latter all have negative Euler index, a contradiction. \diamond

3.8 Naturality of carrying

The carrying relation between train tracks is natural with respect to carrying of measured foliations:

Proposition 3.8.1. *Given a carrying relation $\tau \gg \tau'$ between train tracks, we have $\mathcal{MF}(\tau') \subset \mathcal{MF}(\tau)$ and $\mathcal{PMF}(\tau') \subset \mathcal{PMF}(\tau)$. There is a natural commutative diagram*

$$\begin{array}{ccc} W(\tau') & \longrightarrow & W(\tau) \\ \downarrow & & \downarrow \\ \mathcal{MF}(\tau') & \xrightarrow{\subset} & \mathcal{MF}(\tau) \end{array}$$

where the map $W(\tau') \rightarrow W(\tau)$ is well-defined independent of the choice of a carrying map.

Proof. Given a partial measured foliation \mathcal{F} , composing a carrying injection $\mathcal{F} \subset \nu(\tau')$ with the inclusion $\nu(\tau') \subset \nu(\tau)$ of Proposition 3.5.1 defines a carrying injection $\mathcal{F} \subset \nu(\tau)$, showing that $\mathcal{MF}(\tau') \subset \mathcal{MF}(\tau)$. Pushing forward invariant weights under a carrying map $\tau' \rightarrow \tau$ defines the map $W(\tau') \rightarrow W(\tau)$. Commutativity of the diagram is easily checked. The fact that $W(\tau') \rightarrow W(\tau)$ is well-defined follows from injectivity of the maps $W(\tau) \rightarrow \mathcal{MF}(\tau)$, $W(\tau) \rightarrow \mathcal{MF}(\tau')$. \diamond

3.9 Filling properties of train tracks

A train track τ is *filling* if each component of $S - \tau$ is a nonpunctured or once-punctured disc. If τ is a filling train track then the *singularity type* of τ is the pair of sequences $(i_3, i_4, \dots; p_1, p_2, \dots)$ where i_n (resp. p_n) is the number of nonpunctured (resp. punctured) components of $\mathcal{C}(S - \tau)$ with n cusps. These concepts are combinatorial invariants of τ .

Note that if \mathcal{F} is a partial measured foliation canonically carried on a filling train track τ , then \mathcal{F} is arational and the singularity types of \mathcal{F} and of τ are identical.

However, a filling, recurrent train track might not canonically carry anything, in fact it might not carry anything arational at all, belying the terminology “filling”. Here is an example.

Suppose that \mathcal{F} is a canonical partial arational measured foliation with two components $\mathcal{F}_1, \mathcal{F}_2$, each an orientable foliation, separated from each other by a single component of $\mathcal{C}(S - \mathcal{F})$ which is a nonpunctured annulus. Choose a train track τ' which canonically carries \mathcal{F} , and so τ' has components τ'_1, τ'_2 canonically carrying $\mathcal{F}_1, \mathcal{F}_2$ respectively, and there is a unique component of $\mathcal{C}(S - \tau')$, an annulus, separating τ'_1 from τ'_2 . The train track τ' is necessarily orientable. Now choose short smooth train paths α, α' on opposite components of ∂A , and pinch these together to create a weakly filling train track τ . In the terminology to be introduced below, τ' is obtained from τ by a central splitting, and this splitting is forced. In practical terms this means that τ and τ' carry the exact same set of measured foliations, that is, $\mathcal{MF}(\tau) = \mathcal{MF}(\tau')$, as is easy to see. In particular, τ does not carry any arational measured foliation.

See the remarks after Proposition 3.13.2 for a further discussion of filling properties.

3.10 Diagonal extension of train tracks

A train track σ on S is *complete* if each component of $\mathcal{C}(S - \sigma)$ is a trigon or once-punctured monogon. If σ is filling but not complete then we wish to define several extensions of σ .

Let P be an n -cusped polygon with $k \leq 1$ punctures whose Euler index $\iota(P) = 1 - \frac{n}{2} - k$ is negative. A *diagonal* in P is a properly embedded smooth arc whose two endpoints meet two distinct cusps tangentially. An *awl* in P is a smooth 1-complex consisting of the union of the boundary of a once-punctured monogon in the interior of P together with an arc, so that the monogon boundary and the arc meet at a generic train track switch, and the other endpoint of the arc meets a cusp of P tangentially. A diagonal that is homotopic rel endpoints to a side of P is said to be *side parallel*, and if P is a once-punctured monogon then any awl in P is also said

to be *side parallel*; in either case, the diagonal or the awl cuts P into two regions one of which is a nonpunctured bigon. Two diagonals with disjoint interiors and the same endpoints which bound a nonpunctured bigon are said to be *parallel*. A *diagonal system* in P is a union of diagonals and awls which are disjoint except at cusps of P , so that the system cuts P into regions of negative Euler index; in other words, there are no side parallel diagonals or awls, and no two diagonals are parallel. If moreover each region has index $-\frac{1}{2}$, i.e. it is a nonpunctured trigon or once punctured monogon, then the diagonal system is *complete*. The number of diagonals and awls in a diagonal system in P is $\leq n + 2k - 3$, with equality if and only if it is complete. Of course, if $\iota(P) = -\frac{1}{2}$ then the only complete diagonal system is empty. More generally, the number of complete diagonal systems in a nonpunctured n -gon is given by the Catalan number a_{n-2} where $a_k = \binom{2k}{k} - \binom{2k}{k-1}$ ([Knu73] page 531), and for a once-punctured n -gon the number of completions is na_{n-1} . In particular if $\iota(P) < -\frac{1}{2}$, that is if P is not already a trigon or a punctured monogon, then a complete diagonal system is not unique.

Consider a filling train track σ on S . A *diagonal extension* of σ is any train track ρ containing σ such that for each component P of $\mathcal{C}(S - \sigma)$, $\rho \cap P$ is a diagonal system in P .

Let $\mathcal{MF}^\ell(\tau) \subset \mathcal{MF}$ be the union of $\mathcal{MF}(\sigma)$ over all diagonal extensions of τ , and let $\mathcal{PMF}^\ell(\tau)$ be the projective image in \mathcal{PMF} , equal to the union of $\mathcal{PMF}(\sigma)$ over all diagonal extensions of τ .³ A partial measured foliation or essential simple closed curve is *legal* with respect to τ if its class is contained in $\mathcal{MF}^\ell(\tau)$, i.e. if it is carried by some diagonal extension of τ .

We noted earlier that for a train track τ , the map $W(\tau) \rightarrow \mathcal{MF}$ is an embedding, whose image we have denoted $\mathcal{MF}(\tau)$. This can be extended to obtain a description of $\mathcal{MF}^\ell(\tau)$ as follows.

Lemma 3.10.1. *Let τ be a filling train track and let \mathcal{F} be a partial measured foliation or essential simple closed curve which is legal with respect to τ . Then there is a unique minimal diagonal extension of τ which carries \mathcal{F} . To be precise, suppose that $\mathcal{F}_1, \mathcal{F}_2$ are equivalent to \mathcal{F} , and for $i = 1, 2$ suppose that \mathcal{F}_i is carried on a diagonal extension σ_i of τ so that the image of the carrying map $\mathcal{F}_i \rightarrow \sigma_i$ contains $\sigma_i - \tau$. Then σ_1 is isotopic to σ_2 rel τ .*

As a corollary, we can describe the structure of $\mathcal{PMF}^\ell(\tau)$:

Corollary 3.10.2. *Suppose τ is a filling train track, and let Σ be the lattice of diagonal extensions of τ under inclusion. Then $\mathcal{PMF}^\ell(\tau)$ is obtained from the disjoint union of the $\mathcal{PMF}(\sigma)$, $\sigma \in \Sigma$, by identifying $\mathcal{PMF}(\sigma)$ with its image in $\mathcal{PMF}(\sigma')$ whenever $\sigma \subset \sigma'$. \diamond*

³ $\mathcal{MF}^\ell(\tau)$ is denoted $PE(\tau)$ in [MM99].

We shall prove Lemma 3.10.1 in detail only on the special case that $\mathcal{F}_1, \mathcal{F}_2$ are simple closed curves; the general proof will be sketched briefly.

Proof of Lemma 3.10.1 for simple closed curves $c_i = \mathcal{F}_i$. Assuming that σ_1, σ_2 are not isotopic rel σ , we prove that c_1, c_2 are not isotopic. For any component P of $\mathcal{C}(S - \sigma)$, let σ_i^P denote the set of diagonals and awls added in P to form σ_i . We may assume that σ_1^P, σ_2^P intersect transversely and efficiently in P , meaning that each component of $\mathcal{C}(P - (\sigma_1^P \cup \sigma_2^P))$ has nonpositive index; moreover, we can arrange that the interior of the monogon in a awl σ_1^P is disjoint from σ_2^P except possibly for a single arc on the boundary of the monogon of a awl of σ_2^P , and similarly with indices 1, 2 reversed. Each component of index 0 is either a disc with two corners and one cusp or a disc with four corners and no cusps; the corners are located at transverse intersection points of σ_1^P, σ_2^P .

Case 1: σ_1 and σ_2 have no transverse intersections. It follows that $\sigma_1 \cup \sigma_2$ is a train track carrying both c_1 and c_2 . Since $\sigma_1 - \sigma \neq \sigma_2 - \sigma$, it follows that c_1 and c_2 are fully carried on different subtracks of $\sigma_1 \cup \sigma_2$. Applying Proposition 3.7.3 it follows that c_1, c_2 are not isotopic.

Case 2: σ_1 and σ_2 have a transverse intersection. Choose a tie bundle $\nu(\sigma)$. Extend $\nu(\sigma)$ to tie bundles $\nu(\sigma_1)$ and $\nu(\sigma_2)$, so that in a component P of $\mathcal{C}(S - \sigma)$, each transverse intersection point of σ_1^P, σ_2^P corresponds to a component of $\nu(\sigma_1) \cap \nu(\sigma_2)$ which is a square on which the tie bundles intersect like the horizontal and vertical foliations of the square $[0, 1] \times [0, 1]$; we shall refer to these as the *crossing squares* of $\nu(\sigma_1), \nu(\sigma_2)$. By the hypothesis of Case 2 there is at least one crossing square on the surface. Choose carrying injections $c_i \subset \nu(\sigma_i)$, $i = 1, 2$, and perturb them so that the simple closed curves c_1, c_2 intersect transversely; there is at least one point of $c_1 \cap c_2$ in each crossing square.

Now we mimic the proof of Proposition 3.7.3. Suppose that c_1, c_2 have a bigon (α_1, α_2) bounding a nonpunctured disc d , with $d \cap c_i = \alpha_i \subset \partial d$.

We claim that $d \subset \nu(\sigma)$, and that the restriction of the tie bundle of $\nu(\sigma)$ to d consists of arcs each crossing d from α_1 to α_2 .

Once this claim is established, then we can pull c_1 across d along ties, reducing the cardinality of $c_1 \cap c_2$, preserving the fact that we have carrying injections $c_i \hookrightarrow \nu(\sigma_i)$, and preserving the image of the carrying map $c_i \rightarrow \sigma_i$. We may then continue pulling c_1 across bigons one by one, and by induction we are reduced to the case that c_1, c_2 have no bigons. But $c_1 \cap c_2$ is still nonempty, having a transverse intersection point in each crossing square, and so c_1, c_2 are not isotopic.

Now we prove the claim. The bigon d has two corners. If either of these corners is in $\nu(\sigma)$, smooth that corner so that the strands of c_1, c_2 are tangent to each other

at that point, keeping c_1, c_2 transverse to the tie of $\nu(\sigma)$ at that point; this converts that corner of d into either a cusp or a smooth boundary point. The bigon d may now be regarded as a disc with smooth boundary except at ≤ 2 special boundary points, each a cusp or a corner; it follows that $\iota(d) \geq 0$, with equality if and only if d has two cusps. Consider the following decomposition of d into surfaces-with-corners:

$$\begin{aligned} d &= \underbrace{(d \cap (\nu(\sigma_1) \cup \nu(\sigma_2)))}_{d_1} \cup \underbrace{\text{Cl}(d - (\nu(\sigma_1) \cup \nu(\sigma_2)))}_{d_2} \\ &= \quad \quad \quad \cup \quad \quad \quad \end{aligned}$$

The Euler index of d_1 is zero, because the cross bundle over d_1 has a section with the property that for each point $x \in \partial d_1$, each tangent line of ∂d_1 at x is parallel to one slat of the cross over x . In each crossing square the two slats of the cross are parallel to the two fibrations of the square; outside of each crossing square one slat of the cross is transverse to the tie bundle. Each component of d_2 has nonpositive Euler index, because each corresponds to a component of $S - (\sigma_1 \cup \sigma_2)$ preserving diffeomorphism type. It follows that $\iota(d) \leq 0$.

Combining what we know about $\iota(d)$, it follows that $\iota(d) = 0$, that each component of d_2 has index 0, and that d has two cusps and no corners. In particular, since d has no corners, d is disjoint from the crossing squares. However, as mentioned above, the components of $S - (\sigma_1 \cup \sigma_2)$ of index 0 are each incident to a transverse intersection point of $\sigma_1 \cap \sigma_2$, and so each component of d_2 of index 0 is incident to a corner of an crossing square. If d_2 were nonempty then d would not be disjoint from the crossing squares, and so we must have $d_2 = \emptyset$. This implies $d \subset \nu(\sigma_1) \cup \nu(\sigma_2)$. Also, for each $i = 1, 2$, d must be disjoint from $\nu(\sigma_i) - \nu(\sigma)$, because otherwise d would intersect a crossing square. This shows that $d \subset \nu(\sigma)$. \diamond

Here is a brief sketch of the proof when $\mathcal{F}_i, i = 1, 2$ are partial arational measured foliations; the general case is obtained by combining this with the annular case, proved above. We have a carrying injection $\mathcal{F}_i \subset \nu(\sigma_i)$. We can perturb $\mathcal{F}_1, \mathcal{F}_2$ to intersect transversely. It still makes sense to talk about bigons of $\mathcal{F}_1, \mathcal{F}_2$, and to pull \mathcal{F}_1 across a bigon to simplify the intersection, although one might first have to slice \mathcal{F}_1 and \mathcal{F}_2 along the boundary of the bigon. If $\mathcal{F}_1, \mathcal{F}_2$ are *not* topologically equivalent, meaning that their underlying topological foliations are not equivalent, then one can pull across finitely many bigons until $\mathcal{F}_1, \mathcal{F}_2$ intersect efficiently, and in this case the proof goes through as above. If $\mathcal{F}_1, \mathcal{F}_2$ are topologically equivalent, then there is an infinite sequence of bigon pulls, so that the intersection number limits to zero; but we get a positive lower bound to the intersection number by considering the weight functions that $\mathcal{F}_1, \mathcal{F}_2$ induce on σ_1, σ_2 , and taking the product at a transverse intersection point.

3.11 Recurrence

Let τ be a pretrack. A train path $\gamma: [a, b] \rightarrow \tau$ is *closed* if $\gamma(a) = \gamma(b)$ and gluing a to b induces a smooth map from the circle to τ . We say that τ is *recurrent* if for each branch of τ there exists a closed train path passing over that branch. Recurrence is a combinatorial invariant of train tracks.

Here are several equivalent formulations of recurrence:

Theorem 3.11.1. *Given a pretrack τ , the following are equivalent:*

- (1) *There exists a positive invariant weight on τ .*
- (2) *There exists a surjective closed train path in τ .*
- (3) *There exists a positive integral invariant weight on τ .*
- (4) *τ is recurrent (that is, for every branch b there exists a closed train path of τ passing over b).*
- (5) *For every branch b there exists an integral invariant weight w_b on τ such that $w_b(b) > 0$.*

Sketch of proof. (see [Pen92] §1.3 for more details) It is evident that $2 \iff 3$, and $4 \iff 5$. It is evident that $3 \implies 1$; for the converse use rational approximation and clear the denominators. It is evident that $3 \implies 5$; for the converse take the sum of invariant weights w_b over all branches b . \diamond

When τ is recurrent, the dimension of $W(\tau)$ is easily computed:

Proposition 3.11.2. *If τ is a connected, recurrent pretrack then*

$$\dim(W(\tau)) = \begin{cases} \#\text{branches}(\tau) - \#\text{switches}(\tau) & \text{if } \tau \text{ is not orientable} \\ \#\text{branches}(\tau) - \#\text{switches}(\tau) + 1 & \text{if } \tau \text{ is orientable} \end{cases}$$

The proof in the nonorientable case is taken from Lemma 2.1.1 of [Pen92], and the orientable case is similar.

Proof. The switch conditions define a natural linear map $\sigma: \mathbf{R}^{B(\tau)} \rightarrow \mathbf{R}^{\text{Sw}(\tau)}$, and the kernel of this map intersected with $[0, \infty)^{B(\tau)}$ is equal to $W(\tau)$. Assuming τ is recurrent, $W(\tau) \cap (0, \infty)^{B(\tau)}$ is nonempty, and it follows that $\dim(W(\tau))$ equals $\#\text{branches}(\tau) - \text{rk}(\text{image}(\sigma))$. We must therefore compute $\text{rk}(\text{image}(\sigma))$, which equals the maximum number of linearly independent switch conditions.

Suppose first that τ is oriented, so τ may be regarded as a directed graph. We must prove $\text{rk}(\text{image}(\sigma)) = \#\text{switches}(\tau) - 1$. There is one evident linear

relation among switch conditions: the sum of the switch conditions equals zero, because in this sum each branch occurs once with a $+1$ coefficient and once with a -1 coefficient. A basic fact of graph theory says that orientability together with recurrence and connectedness of τ implies that τ is *strongly connected*, which means that if b, b' are any two branches of τ , oriented by restricting the ambient orientation, then there exists a train path (directed path) of the form $b = b_1 * \cdots * b_n = b'$ (see e.g. [Pen92] Proposition 1.3.7). Consider two switches $s \neq s'$. Let b be a branch with tail s , let b' be a branch with head s' , and let $b = b_1 * \cdots * b_n = b'$ be a train path as above. The counting measure of this train path satisfies all switch conditions except those at s and s' , showing that the switch conditions of s, s' are not contained in the span of the remaining switch conditions, thereby proving that $\text{rk}(\text{image}(\sigma)) = \#\text{switches}(\tau) - 1$.

If τ is nonorientable and recurrent, then the orientation double cover of τ is connected, recurrent, and orientable, and therefore strongly connected. It follows that if b, b' are any two branches of τ with arbitrarily assigned orientations, then there is a train path $b = b_1 * \cdots * b_n = b'$. We must prove that the set of switch conditions is linearly independent. Given a switch s , let b be an oriented branch with tail s , and let b' be the same branch but with opposite orientation. Let $b = b_1 * \cdots * b_n = b'$ be a train path as above. The counting measure of this train path satisfies each switch condition except that at s , showing that the switch condition of s is not a linear combination of the remaining switch conditions, and thereby proving $\text{rk}(\text{image}(\sigma)) = \#\text{switches}(\tau)$. \diamond

The following proposition uses recurrent subtracks of τ to explain the face structure of $W(\tau)$ and $\mathcal{PW}(\tau)$:

Proposition 3.11.3. *Let τ be a pretrack.*

- (1) *The face lattice of $W(\tau)$ is isomorphic to the inclusion lattice of recurrent subtracks of τ . A face F of $W(\tau)$ corresponds to a recurrent subtrack σ of τ if and only if σ is the union of those branches of τ which have positive coordinate for every invariant weight in the interior of F .*
- (2) *The face lattice of $\mathcal{PW}(\tau)$ is isomorphic to the sublattice of nonempty recurrent subtracks of τ .*
- (3) *If τ is semigeneric then vertices of the face lattice of $\mathcal{PW}(\tau)$ correspond to subtracks of τ which are either simple closed curves or dumbbell tracks (see Figure 10). If τ is not semigeneric then we must also allow “degenerate dumbbells” where the central branch of the dumbbell is reduced to a nonsemigeneric switch.*

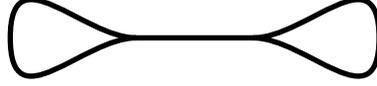


Figure 10: A dumbbell.

Proof. (1) is proved as follows (see also [Pen92] page 117). Given a face F , note that F is the intersection of $W(\tau)$ with some face of the simplicial cone $[0, \infty)^{B(\tau)}$. Letting σ_F be the union of branches of τ having positive weight in the interior of F , if w is in the interior of F then it follows that σ_F is the support of w , showing that σ_F is a recurrent pretrack. Inclusion of faces clearly corresponds to inclusion of recurrent subtracks. If σ is a recurrent subtrack of τ , let F_σ be the set of invariant weights with positive coordinate on each branch of σ and zero coordinate elsewhere in τ . Recurrence of σ implies that F_σ is nonempty. Clearly F_σ is the intersection of $W(\tau)$ with a face of the simplicial cone $[0, \infty)^{B(\tau)}$, and so F_σ is a face of $W(\tau)$.

(2) follows immediately from (1).

To prove (3), let σ be a minimal, nonempty, recurrent subtrack of τ , that is, $\mathcal{PW}(\sigma)$ is a vertex of $\mathcal{PW}(\tau)$. Assuming τ is semigeneric, we must prove that σ is either a simple closed curve or a dumbbell.

First, every closed train path on σ surjects onto σ , for suppose there exists a nonsurjective closed train path γ with invariant weight $w = w_\gamma$. Since σ is recurrent there also exists a positive invariant weight w' . But the invariant weights $w, w' \in \mathbf{R}^{B(\sigma)}$ have distinct supports, and therefore they are linearly independent, from which it follows that $\mathcal{PW}(\sigma)$ is not a point, a contradiction.

Now let $\gamma: S^1 \rightarrow \sigma$ be a closed train path of minimal length, where length is the number of branch traversals of γ . Note that γ traverses each branch at most once in each direction; if not then γ has the form $abcb$, for some train paths a, c , with b traversed in the same direction each time, and then ab defines a shorter closed train path, contradiction.

If γ traverses each branch at most once in either direction then, since τ is semigeneric, it also follows that γ hits each switch at most once, and so σ is a simple closed curve and we are done.

If γ traverses some branches twice, albeit in opposite directions, it follows that γ can be written as

$$\gamma = a_1 * e_1 * \cdots * a_{2n} * e_{2n}$$

where each a_i is a maximal train path in σ over which γ is one-to-one, each e_i is a

maximal train path over which γ is two-to-one, and there is a transposition of the index set $\{1, \dots, 2n\}$ denoted $i \leftrightarrow \bar{i}$ such that $e_{\bar{i}}$ is the same as e_i traversed in the opposite direction, that is, $e_{\bar{i}} = \bar{e}_i$. Here again we use that τ is semigeneric to write γ in this form.

If $n = 1$ then σ is a dumbbell.

If $n \geq 2$ we shall derive a contradiction. Choosing two pairs of indices i, \bar{i}, j, \bar{j} as above and writing $e = e_i, e' = e_j$, then γ has one of two forms:

$$\gamma = aebe'c\bar{e}d\bar{e}' \quad \text{or} \quad aeb\bar{e}ce'd\bar{e}'$$

If γ has the first form $aebe'c\bar{e}d\bar{e}'$ then $ae\bar{c}\bar{e}'$ defines a shorter closed train path, contradicting minimality. If γ has the second form $aeb\bar{e}ce'd\bar{e}'$, then $aeb\bar{e}\bar{a}e'd\bar{e}'$ is a closed train path in σ disjoint from c , contradicting that all train paths in σ are surjective.

The proof of (3) when τ is not semigeneric reduces to the semigeneric case by a combing argument. \diamond

Here are several other equivalent formulations of recurrence, which make it clear that recurrence is a (quickly) decidable property of a pretrack:

Proposition 3.11.4. *Given a pretrack τ , let $\tilde{\tau} \rightarrow \tau$ be the orientation double cover. The following are equivalent:*

- (1) τ is recurrent.
- (2) $\tilde{\tau}$ is recurrent.
- (3) For every branch b of $\tilde{\tau}$ there exists a simple closed train path in $\tilde{\tau}$ passing over b .
- (4) For every branch b of τ there exists a closed train path passing over b at most once in each direction.
- (5) For every branch b of τ there exists a subtrack of τ containing b which is either a simple closed curve or a dumbbell.

Property 3, for example, shows that recurrence is decidable in linear time, for there is an algorithm which decides in linear time as a function of the number of branches whether the directed graph $\tilde{\tau}$ has simple directed loops passing over every branch.

Sketch of proof. For $1 \implies 2$ lift a positive invariant weight from τ to $\tilde{\tau}$, and for $2 \implies 1$ push forward a positive invariant weight from $\tilde{\tau}$ to τ . For $2 \implies 3$, starting from any closed train path passing over b , use cut and paste to find a simpler one, and continue by induction. It is evident that $2 \longleftarrow 3 \iff 4 \longleftarrow 5$.

It remains to prove that recurrence of τ implies 5. As a consequence of Proposition 3.11.3(1) and (2), combined with recurrence of τ , for each branch b of τ there exists a recurrent subtrack σ of τ which contains b , such that $\mathcal{PW}(\sigma)$ is a single point. By Proposition 3.11.3(3) it follows that σ is a simple closed curve or a dumbbell. \diamond

3.12 Slide moves and comb equivalence

There is a natural equivalence relation among semigeneric train tracks: τ, τ' are *carrying equivalent* if each is carried by the other. By Propositions 3.8.1 this implies that $\mathcal{MF}(\tau) = \mathcal{MF}(\tau')$.

The existence of a “slide move”⁴ between two generic train tracks (Figure 11) implies carrying equivalence, as does a more general “combing move” between two semigeneric train tracks (Figure 12). We introduce these moves here, and the equivalence relation that they generate, called “comb equivalence”. Then in Proposition 3.12.2 we will prove that comb equivalence is the same as carrying equivalence, for recurrent train tracks.

Slide moves. Let τ be a generic pretrack and let b be an embedded transition branch. A *slide* of τ along b results in a generic pretrack τ' , as depicted in Figure 11; this move is denoted $\tau \asymp^b \tau'$. Letting the ends of b be located at distinct switches p_-, p_+ , so that the switch orientation at p_- points into b and at p_+ points out of b , and orienting b to agree with these two switch orientations, the effect of the slide move is to exchange the switches at p_- and p_+ , sliding the switch at p_- forward and the switch at p_+ backward until these switches pass each other, resulting in the pretrack τ' . The diffeomorphism types of completed complementary components is unchanged by a slide, and in particular the list of their Euler indices is unchanged; it follows that if τ is a bigon track or train track then the same is true of τ' .

To be precise, choose a disc neighborhood D of b intersecting τ transversely in four points r, s, t, u , so that $D \cap \tau$ consists of arcs $\overline{rp_-}$, $\overline{sp_-}$ on the two-ended side of p_- , the branch $b = \overline{p_-p_+}$ on the one-ended side of p_- and the two ended side of p_+ , the branch $\overline{tp_+}$ on the two-ended side of p_+ , and the branch $\overline{p_+u}$ on the one-ended side of p_+ . The effect of the slide is to remove the arcs $\overline{rp_-}$ and $\overline{tp_+}$, and to smoothly insert arcs $\overline{rp_+}$ and $\overline{tp_-}$, so that on the two-ended side of p_- we have

⁴Slide moves are called *shifts* in [Pen92].

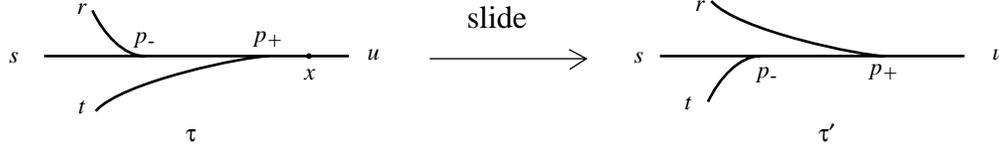


Figure 11: A slide.

the arcs $\overline{tp_-}$ and $\overline{sp_-}$ and on the one-ended side we have $\overline{p_-p_+}$, and so that on the two-ended side of p_+ we have the arcs $\overline{p_-p_+}$ and $\overline{rp_+}$ and on the one-ended side we have $\overline{p_+u}$.

Note that sliding is a symmetric relation among generic pretracks: if $\tau \rightarrow \tau'$ is a slide then so is $\tau' \rightarrow \tau$. It follows that the reflexive, transitive closure of the slide relation is an equivalence relation called *slide equivalence*. If one pretrack in a slide equivalence class is a bigon track, resp. train track, then they all are.

If $\tau \succcurlyeq \tau'$ then $\tau \succcurlyeq \tau'$. The carrying map $\tau' \mapsto \tau$ is defined as follows (see figure 11). Letting $x \in \tau$ be a point in the interior of $\overline{p_+u}$, we have:

$$\begin{aligned}
 \tau' &\mapsto \tau \\
 \overline{tp_-} &\mapsto \overline{tp_+} \\
 \overline{p_-p_+} &\mapsto \overline{p_+x} \\
 \overline{p_+u} &\mapsto \overline{xu} \\
 \overline{sp_-} &\mapsto \overline{sp_-} \cup \overline{p_-p_+} \\
 \overline{rp_+} &\mapsto \overline{rp_-} \cup \overline{p_-p_+} \cup \overline{p_+x}
 \end{aligned}$$

where x is some point in the interior of $\overline{p_+u}$.

By symmetry of sliding, if τ, τ' are slide equivalent it follows that they are carrying equivalent, and so $\mathcal{MF}(\tau) = \mathcal{MF}(\tau')$ and $\mathcal{MF}^\perp(\tau) = \mathcal{MF}^\perp(\tau')$ (see also [Pen92] Proposition 2.2.2).

Comb equivalence. Slide moves themselves do not have a natural generalization to semigeneric train tracks. However, note that given a slide $\tau \succcurlyeq^b \tau'$ with inverse slide $\tau' \succcurlyeq^{b'} \tau$, there is a semigeneric train track τ'' which is obtained from τ by collapsing b and is also obtained by τ' by collapsing b' . Such collapses are examples of *comb* moves on semigeneric train tracks, which we now study.

Suppose that τ is a semigeneric pretrack and b is an embedded transition branch. We define *combing* along b , resulting in a semigeneric pretrack τ' , as follows (see Figure 12). The endpoints of b are switches s_- , s_+ , with the switch orientation at s_- pointing into b and at s_+ pointing out of b . On the multi-ended side of s_-

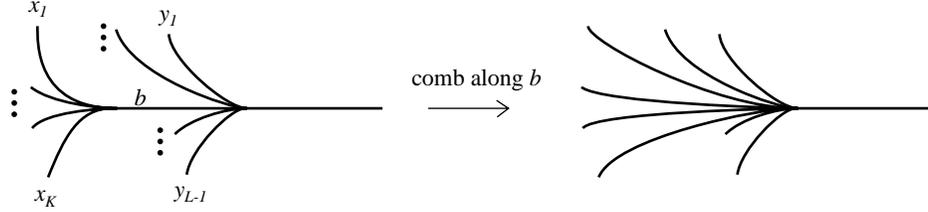


Figure 12: Combing along a transition branch.

there are $K \geq 2$ ends represented by segments $\overline{x_k s_-}$ for $k = 1, \dots, K$, and the one-ended side of s_+ is represented by $\overline{s_- s_+}$. On the multi-ended side of s_+ there are $L \geq 2$ ends represented by segments $\overline{y_l s_+}$, $l = 1, \dots, L-1$ and by $\overline{s_- s_+}$, and on the one-ended side of s_+ there is a segment $\overline{s_+ z}$. The combed pretrack τ' is defined by collapsing b to a single switch s , whose multi-ended side has $K + L - 1$ ends represented by segments $\overline{x_k s}$, $k = 1, \dots, K$ and $\overline{y_l s}$, $l = 1, \dots, L-1$, and whose one-ended side is represented by $\overline{s z}$. Note that the collapse map extends to a map $f: (S, \tau) \rightarrow (S, \tau')$ which is one-to-one except for the identification of b to a point, and such that f is homotopic to the identity on S .

Comb equivalence among semigeneric pretracks is the equivalence relation generated by combing along embedded transition branches. We use the notation $\tau \approx \tau'$ to denote comb equivalence of τ and τ' .

Note that there is a carrying map $\tau' \rightarrow \tau$, mapping each arc $\overline{y_l s}$ to $\overline{y_l s_-} \cup \overline{s_- s_+}$. There is also a carrying map $\tau \rightarrow \tau'$, mapping each arc $\overline{x_k s}$ to $\overline{x_k s_-} \cup \overline{s_- s_+}$. Comb equivalence therefore implies carrying equivalence, and so if τ, τ' are comb equivalent then $\mathcal{MF}(\tau) = \mathcal{MF}(\tau')$ and $\mathcal{MF}^\perp(\tau) = \mathcal{MF}^\perp(\tau')$.

Every semigeneric pretrack τ is comb equivalent to a generic pretrack. In fact, near any semigeneric, nongeneric switch s , τ can be *uncombed* to produce a family of generic switches. If s has $k \geq 3$ ends on the multi-ended side, then uncombing produces a tree of $k-2$ transition branches and $k-1$ generic switches. By uncombing each semigeneric switch, the result is called a *generic uncombing* of τ .

We now have two equivalence relations among generic pretracks: slide equivalence, which is the reflexive, transitive closure of slide moves; and comb equivalence, which is the restriction to generic pretracks of the equivalence relation among semigeneric pretracks generated by comb moves. These two notions give the exact same relation among generic pretracks. We saw earlier that a slide move $\tau \approx^b \tau'$ can be realized by combing τ along b and then uncombing to get τ' . In the other direction, if τ_0 is a semigeneric pretrack, if we then comb along some embedded transition branch to obtain τ_1 , and if τ'_i is a generic uncombing of τ_i for $i = 1, 2$, then a

sequence of slide moves from τ'_0 to τ'_1 is easily constructed.

Among recurrent train tracks, comb equivalence can be characterized nicely as follows:

Proposition 3.12.1. *If τ is a semigeneric, recurrent pretrack, then up to isotopy there is a unique semigeneric pretrack τ' comb equivalent to τ such that τ' has no transition branches.*

Any semigeneric pretrack without transition branches is said to be *completely combed*, and the train track τ' is called the *complete combing* of τ . As the proof shows, τ' is obtained from τ by collapsing each transition branch of τ to a point.

Proof. We may assume that τ is connected, and we may also assume that τ is not a smooth closed curve.

Since τ is recurrent, it has no sink loops, and so the union of transition branches forms a forest, the transition forest. By collapsing each component of the transition forest of τ to a point, the result is a pretrack τ' without transition branches. Moreover, it is clear that when the transition branches of τ are inductively collapsed in any order, the result is isotopic to τ' . We refer to τ' as the *complete combing* of τ . In carrying out this induction we make use of the fact that each transition branch in a recurrent pretrack is embedded, so that the combing move is defined; and also recurrence is preserved under a combing move, so we can continue to collapse an arbitrary transition branch.

Suppose τ_0, τ_1 are semigeneric, recurrent pretracks such that τ_1 is obtained from τ_0 by combing along a transition branch b . Then the complete combings of τ_0, τ_1 are isotopic. To see why, we may start the complete combing of τ_0 by first combing along b , producing τ_1 , and then continuing with any complete combing of τ_1 . The resulting pretrack is the complete combing of both τ_0 and τ_1 .

It follows that all semigeneric pretracks in the comb equivalence class of τ have the same complete combing up to isotopy. In particular, if τ'_0, τ'_1 are two semigeneric pretracks in the comb equivalence class of τ each of which has no transition branches, then they are each their own complete combings, and so they are isotopic to each other. \diamond

Comb equivalence and carrying equivalence. We now prove that these two equivalence relations are identical.

Proposition 3.12.2. *Two recurrent, semigeneric train tracks τ, τ' are comb equivalent if and only if they are carrying equivalent.*

This proposition can fail if τ, τ' are not semigeneric. For instance, if τ has one semigeneric switch s and if τ' is obtained from τ by completely uncombing the switch s , then $\tau' \succ \tau$ but $\tau \not\prec \tau'$.

Proof. Comb equivalence obviously implies carrying equivalence. Assume then that τ, τ' are carrying equivalent. We may assume that τ, τ' are each completely combed, and with this assumption we shall prove that τ, τ' are isotopic. Choose carrying maps $\tau \xrightarrow{f} \tau' \xrightarrow{g} \tau$. Consider the composed carrying map $\tau \xrightarrow{h=gf} \tau$.

We claim that h has a very simple structure: h shrinks each sink branch and stretches each source branch, moving each switch forward into the adjacent sink branch. More precisely:

- Each sink branch b is mapped into itself preserving orientation.
- Each source branch b' is mapped over itself, by an injection whose image $h(b')$ lies in the union of b' with the two incident sink branches, so that the map $h^{-1}(b') \mapsto b'$ preserves orientation.

Once this claim is established, the proposition is proved as follows.

First, it follows from the claim that the map $h: \tau \rightarrow \tau$ is almost one-to-one, in the following sense. For each switch s , with incident sink edge b , and incident source edges b_1, \dots, b_k , there are points $x_i \in b_i$ and $y \in b$ such that each of $\overline{x_i s}$ is mapped by h onto \overline{sy} ; and except for these identifications, h is one-to-one. This implies in turn that $f: \tau \rightarrow \tau'$ is almost one-to-one in a similar sense: for each switch s , with the above notation, there are points $x'_i \in \overline{x_i s}$ and $y' \in \overline{sy}$ such that each of $\overline{x'_i s}$ is mapped by f onto $\overline{sy'}$; and except for these identifications, f is one-to-one. Second, the carrying map $f: \tau \rightarrow \tau'$ is surjective, for by Proposition 3.5.2 f is smoothly homotopic to any other carrying map from τ to τ' , and surjectivity is preserved under smooth homotopy of carrying maps; from the definition of comb equivalence one obtains the existence of one surjective carrying map. From this description of f , it follows directly that τ' is isotopic to τ .

It remains to verify the claim. Consider the lifted train track $\tilde{\tau}$ in the universal covering surface \tilde{S} , and let $\partial\tilde{S}$ be the circle at infinity. There is a unique lifted carrying map $\tilde{h}: (S, \tilde{\tau}) \rightarrow (S, \tilde{\tau})$ which extends to the identity on $\partial\tilde{S}$, and by adjusting h to be the identity on a neighborhood of each puncture we may assume that \tilde{h} moves each point a uniformly bounded amount. For each sink branch b of $\tilde{\tau}$ let $t(b)$ be the union of train lines containing b . Note that $t(b)$ is itself a train track. The intersection of any two train lines is connected, by Proposition 3.3.1. It follows that $t(b)$ is a tree, with a unique sink branch b and all other branches being transition branches, and so $t(b)$ is a kind of “infinite basin” of b . By Proposition 3.3.2, \tilde{h} maps each bi-infinite train path in $\tilde{\tau}$ diffeomorphically to itself, preserving orientation. This implies that $\tilde{h}(t(b)) = t(b)$, and since \tilde{h} restricts to a carrying map from $t(b)$ to itself, this implies in turn that \tilde{h} maps b into itself, preserving orientation.

For each source branch b' of $\tilde{\tau}$, let b_1, b_2 be the two sink branches incident to b' , and so the endpoint $s_i = b' \cap b_i$ is mapped by \tilde{h} to a point $x_i \in b_i$. We now

have two train paths in $\tilde{\tau}$ connecting x_1 to x_2 , namely $\tilde{h}(b')$ and $\overline{x_1 s_1} * \overline{s_1 s_2} * \overline{s_2 x_2} = \overline{x_1 s_1} * b' * \overline{s_2 x_2}$. These two train paths are identical, and so \tilde{h} maps b' over itself, preserving orientation.

These conclusions for the map \tilde{h} imply the desired conclusions for h , proving the claim. \diamond

The proof of Proposition 3.12.2 contains in it a technical result which will be useful later. Given semigeneric pretracks τ, τ' , a *switch fold* from τ to τ' is a surjective carrying map $f: \tau \rightarrow \tau'$ which satisfies the following properties:

- For each switch s of τ there exists a switch s' of τ' such that s, s' have the same number of ends on the multi-ended side, and such that one of the following two possibilities occurs:
 - $f(s) = s'$; or
 - on the multi-ended side of s the ends are represented segments $\overline{x_i s}$, $i = 1, \dots, n$, and the one-ended side of s' is represented by a segment $\overline{s' y'}$, such that f takes s to y' , each x_i to s' , and each of the segments $\overline{x_i s}$ to the segment $\overline{s' y'}$.
- Other than the identifications just described near certain switches, f is one-to-one.

It is clear that the existence of a switch fold from τ to τ' implies that τ, τ' are isotopic. The technical result we need is the following:

Lemma 3.12.3. *If τ is a completely combed train track, then every carrying map from τ to itself is a switch fold.* \diamond

3.13 Splitting of train tracks

First we define elementary splitting, an operation on generic train tracks. Then we broaden this to define a splitting in the most general terms. Finally, we indicate an intermediate generalization which is much easier to work with, called a “wide splitting”.

Elementary splittings. Let τ be a generic pretrack. Given a sink branch b of τ we shall define the three *elementary splittings of τ along b* , the *Right*, *Left*, and *Central*⁵ elementary splittings along b , denoted $\tau \succ \tau_R$, $\tau \succ \tau_L$, $\tau \succ \tau_C$. Left and Right splittings are jointly referred to as *parity splittings*, although to confuse

⁵Central splittings are called *collisions* in [Pen92].

the issue the noun *parity* will refer to one of the three terms “Left”, “Right”, and “Central”, or their one letter abbreviations “L”, “R”, and “C”. Throughout the paper we use the variable D to stand for a parity, and when $D \in \{L, R\}$ then we use \bar{D} for the complement of D in the set $\{L, R\}$. In the splitting notation $\tau \succ \tau'$, if we need to emphasize the splitting arc b or the parity L, R, or C, then they can be added to the notation, e.g. $\tau \succ_R^b \tau'$. We note that the determination of Left or Right parity depends on an orientation of the surface, or at least on a choice of an orientation in a neighborhood of the sink branch b . Parity splittings can be defined on a nonoriented surface, but there is no natural way to assign Left and Right parity without first choosing a local orientation.

Elementary splittings are depicted in Figure 13, which show an oriented disc neighborhood D of b , the pretrack τ before splitting, and the pretracks τ_R, τ_L, τ_C after splitting; outside of this neighborhood the pretrack is unaffected by the splitting. Note that if τ is a train track then so are τ_R, τ_C, τ_L . This is immediate for parity splittings because the list of surface-with-corner diffeomorphism types of completed complementary components is preserved. For a central splitting this follows after a moment’s thought, by noting that the central splitting has one of two effects: two components are combined into one, adding Euler indices; or one component is changed in its diffeomorphism type, by adding 1 to its genus, but its Euler index is unaffected; the diffeomorphism types and Euler indices of all other components are unaffected.

For precise definitions, let p, q be the endpoints of b , and suppose that D is chosen so that τ intersects ∂D transversely in four points r, s, t, u in counterclockwise order around ∂D , so that $\tau \cap \partial D$ consists of the segments $\overline{rp}, \overline{sp}, \overline{tq}, \overline{uq}$, and $b = \overline{pq}$.

The pretrack τ_C obtained by central elementary splitting along b is defined by removing these five segments and smoothly inserting disjoint segments $\overline{ru}, \overline{st}$ in D . We define a carrying map $f: (S, \tau_C) \rightarrow (S, \tau)$, equal to the identity outside of D , whose effect on D is as follows. Choose subdivisions $\overline{ru} = \overline{rp_1} * \overline{p_1q_1} * \overline{q_1u}$ and $\overline{st} = \overline{sp_2} * \overline{p_2q_2} * \overline{q_2t}$. Choose a rectangle $r \subset D$ whose horizontal sides $\overline{p_1q_1}$ and $\overline{p_2q_2}$ equal its intersection with τ_C , and foliate r by vertical arcs. The map f collapses each vertical arc of r to a point and is otherwise one-to-one, and it has the following effect on various arcs in D : $\overline{rp_1} \mapsto \overline{rp}$, $\overline{sp_2} \mapsto \overline{sp}$, $\overline{tq_2} \mapsto \overline{tq}$, $\overline{uq_1} \mapsto \overline{uq}$, and $\overline{p_1q_i} \mapsto \overline{pq}$ for $i = 1, 2$. The rectangle r is called the *collapsing rectangle* of the carrying map f .

The pretracks τ_R, τ_L obtained by Left and Right elementary splitting along b are defined as follows. Define τ_R by adding to τ_C a branch $b_R = \overline{p_1q_2}$, with a generic switch at p_1 whose one-ended side is tangent to rp_1 , and a generic switch at q_2 whose one-ended side is tangent to tq_2 . Define τ_L by adding to τ_C a branch $b_L = \overline{q_1p_2}$, with a generic switch at q_1 whose one-ended side is tangent to uq_1 , and a generic switch at p_2 whose one-ended side is tangent to sp_2 . For each $D \in \{L, R\}$ the branch b_D is called the *post-splitting branch* of the splitting $\tau \succ \tau_D$, and b_D is a source branch of

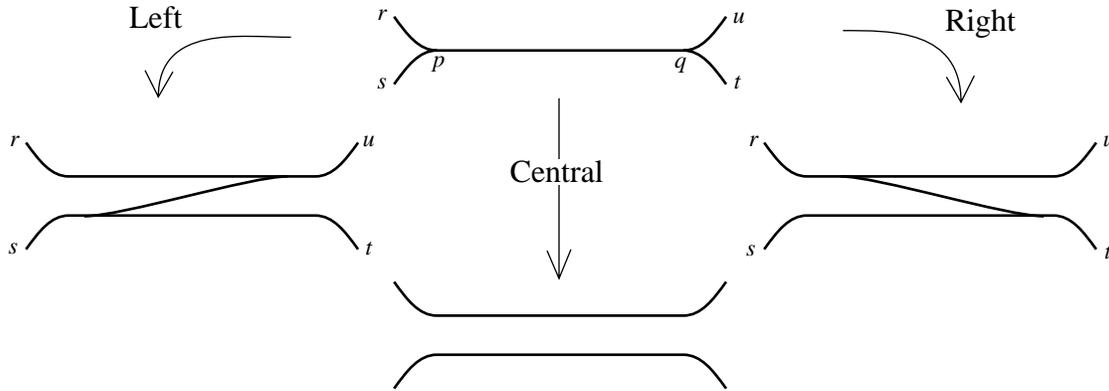


Figure 13: Elementary splittings. The meaning of Left and Right is better appreciated by rotating the page 90° , from which vantage one sees that under a Left splitting each cusp moves to the Left, and under a Right splitting each cusp moves to the Right.

type DD in the train track τ_D . Note that in the notation above, the postsplitting branches b_D can be chosen to have transverse intersection with the vertical fibration of the collapsing rectangle r , from which it follows that the map f defined above can be regarded not just as a carrying map from τ_C to τ , but also as a carrying map from τ_D to τ .

Splitting is a wholly asymmetric relation among train tracks: if $\tau \rightarrow \tau'$ is a splitting then $\tau' \rightarrow \tau$ is *not* a splitting. For central splittings this follows from consideration of the completed complementary components. For parity splittings we shall give a proof of a more general fact below in Proposition 3.14.1: given any sequence of train tracks $\tau_0, \tau_1, \dots, \tau_n$ such that for each $i = 1, \dots, n$ either $\tau_{i-1} \succcurlyeq \tau_i$ is a slide move or $\tau_{i-1} \succ \tau_i$ is a splitting, the two train tracks τ_0, τ_n are comb equivalent if and only if each $\tau_{i-1} \succcurlyeq \tau_i$ is a slide move.

From the discussion above, given an elementary splitting $\tau \succ \tau'$ we have $\tau \succcurlyeq \tau'$, and so by Propositions 3.8.1 it follows that $\mathcal{MF}(\tau) \supset \mathcal{MF}(\tau')$. One can be more precise, as follows.

Consider a train track τ and a sink branch b , and use the labels p, q, r, s, t, u be as in Figure 13. The train path \overline{su} is called a *Left crossing* of b , and the train path \overline{rt} is a *Right crossing* of b . With these notions we easily have:

Fact 3.13.1 (Splitting inequalities). *Let $\tau \succ \tau_L, \tau \succ \tau_R, \tau \succ \tau_C$ be the Left, Right, and Central splittings along b . If $w \in W(\tau)$ then, letting $\mathcal{F} = \mathcal{F}_w \in \mathcal{MF}(\tau)$ with carrying map $\mathcal{F} \rightarrow \tau$, we have:*

- $\mathcal{F} \in \mathcal{MF}(\tau_C) \iff$ no leaf segment of \mathcal{F} maps to a Left or Right crossing of $b \iff w(\overline{r\bar{p}}) = w(\overline{u\bar{q}}) \iff w(\overline{s\bar{p}}) = w(\overline{q\bar{t}})$.
- $\mathcal{F} \in \mathcal{MF}(\tau_L) \iff$ no leaf segment of \mathcal{F} maps to a Right crossing of $b \iff w(\overline{r\bar{p}}) \leq w(\overline{u\bar{q}}) \iff w(\overline{s\bar{p}}) \geq w(\overline{q\bar{t}})$. If these occur then: \mathcal{F} is fully carried by $\tau_L \iff$ the inequalities are strict.
- $\mathcal{F} \in \mathcal{MF}(\tau_R) \iff$ no leaf segment of \mathcal{F} maps to a Left crossing of $b \iff w(\overline{r\bar{p}}) \geq w(\overline{u\bar{q}}) \iff w(\overline{s\bar{p}}) \leq w(\overline{q\bar{t}})$. If these occur then: \mathcal{F} is fully carried by $\tau_R \iff$ the inequalities are strict.

◇

Next we have:

Proposition 3.13.2 (Splitting lattice). *With the notation above, we have a diagram of inclusions:*

$$\begin{array}{ccc}
 & \mathcal{MF}(\tau) = \mathcal{MF}(\tau_L) \cup \mathcal{MF}(\tau_R) & \\
 \nearrow & & \nwarrow \\
 \mathcal{MF}(\tau_L) & & \mathcal{MF}(\tau_R) \\
 \nwarrow & & \nearrow \\
 & \mathcal{MF}(\tau_C) = \mathcal{MF}(\tau_L) \cap \mathcal{MF}(\tau_R) &
 \end{array}$$

If τ is recurrent, then either all three of τ_L, τ_C, τ_R are recurrent or exactly one of them is recurrent (this result is from [Pen92] Lemma 2.1.3). Moreover:

- If all of τ_L, τ_C, τ_R are recurrent then all inclusions above are proper.
- If only τ_C is recurrent then all inclusions are equalities.
- If only τ_L is recurrent then $\mathcal{MF}(\tau_L) = \mathcal{MF}(\tau)$ and $\mathcal{MF}(\tau_C) = \mathcal{MF}(\tau_R)$ and the other two inclusions are proper.
- If only τ_R is recurrent then $\mathcal{MF}(\tau_R) = \mathcal{MF}(\tau)$ and $\mathcal{MF}(\tau_C) = \mathcal{MF}(\tau_L)$ and the other two inclusions are proper.

Remarks. When τ' is the unique one of τ_L, τ_R, τ_C that is recurrent, the splitting $\tau \succ \tau'$ is said to be *forced*. See [Pen92] figure 2.1.4 for examples of forced splittings of arbitrary parity L, R, C.

The case of a forced Central splitting $\tau \succ \tau_C$ is particularly interesting, because in that case any measured foliation carried by τ is forced to have more saddle

connections than may be apparent. In particular, the example of Section 3.9 gives a filling train track τ and a forced Central splitting $\tau \succ \tau_C$ such that τ_C is not filling, and hence *no* measured foliation carried by τ is arational, belying the apparent fact that τ “fills” the surface.

This example motivates a definition of a *strongly filling* train track τ , which means that there is no sequence of forced splittings going from τ to a nonfilling train track. It might be interesting to investigate the relationship between strong filling and other properties of τ such as: τ carries some arational measured foliation; the set of vertex curves of τ fills S ; the set of all essential closed curves carried by τ fills S .

Proof. The diagram of inclusions, and the equations shown in that diagram, are an immediate consequence of Fact 3.13.1.

The fact that either all three or exactly one of τ_L, τ_C, τ_R is recurrent is proved in [Pen92] Lemma 2.1.3, and that proof shows that the four cases depend on which of the two sets $\mathcal{MF}(\tau_L) - \mathcal{MF}(\tau_C)$, $\mathcal{MF}(\tau_R) - \mathcal{MF}(\tau_C)$ is empty. If both are empty then τ_C is the only recurrent one and all inclusions are equalities. If only $\mathcal{MF}(\tau_L) - \mathcal{MF}(\tau_C)$ is empty then τ_R is the only recurrent one and only the inclusions $\mathcal{MF}(\tau_L) \hookrightarrow \mathcal{MF}(\tau)$, $\mathcal{MF}(\tau_C) \hookrightarrow \mathcal{MF}(\tau_R)$ are proper. The case where only $\mathcal{MF}(\tau_R) - \mathcal{MF}(\tau_C)$ is empty is similar. If none are empty then all three of τ_L, τ_C, τ_R are recurrent. \diamond

General splittings. The notion of an elementary splitting has the following built in weakness: given slide equivalent generic train tracks τ, τ' , the set of train tracks obtainable from τ by an elementary splitting need not correspond bijectively, up to isotopy nor even up to slide equivalence, with the set of train tracks obtainable from τ' by elementary splitting. For example, if $\tau \succ \tau'$ is a slide move along a transition branch $b_0 \subset \tau$, and if there is a sink branch b_1 sharing an endpoint with b_0 , then none of the splittings of τ along b_1 are slide equivalent to any splittings of τ' .

In order to remedy this situation (somewhat) we introduce the following:

Definition (General definition of a splitting). Given train tracks τ, τ' , we say that $\tau \succ \tau'$ is a *splitting* if there exist generic train tracks τ_1, τ'_1 comb equivalent to τ, τ' , respectively, such that $\tau_1 \succ \tau'_1$ is an elementary splitting. Moreover, two splittings $\tau \succ \sigma$, $\tau' \succ \sigma'$ are said to be *comb equivalent* if τ, σ are comb equivalent and τ', σ' are comb equivalent.

Thus, every general splitting is comb equivalent to an elementary splitting.

The general definition of splittings has its own problems, in particular lack of constructibility: given a train track τ , how do I construct a representative of each comb equivalence class of splittings $\tau \succ \sigma$? There is an obvious but tedious answer

to this question: enumerate the train tracks obtained from τ by comb equivalence, then do all elementary splittings on those train tracks, and finally weed out all but one representative of each comb equivalence class. A better answer is given as follows.

Wide splittings. We introduce a particular class of general splittings called “wide splittings” which have most of the advantages and few of the disadvantages of both elementary splittings and general splittings. In particular, if τ, σ are comb equivalent then the train tracks obtained from τ by wide splittings correspond bijectively up to comb equivalence with those obtained from σ by wide splitting (see Proposition 3.13.3).

Given a semigeneric pretrack τ , a *splitting arc* is an embedded arc α in S with the following properties:

- (1) $\alpha \cap \tau$ is an embedded train path in τ , consisting of one sink branch, no source branches, and a finite (possibly empty) union of transition branches.
- (2) $\alpha \cap \tau \subset \text{int}(\alpha)$.
- (3) For each of the two components η of $\alpha - \tau$ there exists a component P of $\mathcal{C}(S - \tau)$ such that $\mathcal{C}(\eta)$ has one endpoint at a cusp of P ; we say that α *enters* that cusp at the end η .

The important equivalence relation among splitting arcs of τ is *relative isotopy*, meaning ambient isotopy relative to τ .

If τ is generic then a splitting arc α is uniquely determined up to relative isotopy just by the train path $\alpha \cap \tau$. If τ is not generic then this is not true: one can have splitting arcs α, α' so that the train paths $\alpha \cap \tau$ and $\alpha' \cap \tau$ are identical, but some switch s at an endpoint of this train path is nongeneric, and α, α' enter different cusps at s . In general a splitting arc α is determined, up to relative isotopy, by $\alpha \cap \tau$ and by the cusps into which α enters at its two ends, but in fact α is determined just by the two cusps. To see why, each splitting arc α contains a unique sink branch b , and $\alpha \cap \tau$ is contained in the basin of b ; since the basin of b is a tree then the path $\alpha \cap \tau$ is determined by its endpoints. This argument gives a way to enumerate splitting arcs containing b up to relative isotopy: there are $k_1 \cdot k_2$ of them, where the two half basins of b contain k_1, k_2 cusps, respectively.

If τ, σ are comb equivalent semigeneric train tracks then there is a natural bijection between the splitting arcs of τ and those of σ up to relative isotopy. One way to see this is that a comb equivalence induces a natural bijection of cusps, preserving the property of a pair of cusps that the two members of the pair lie on opposite halves of the basin of some sink branch. Another way to see this is by induction

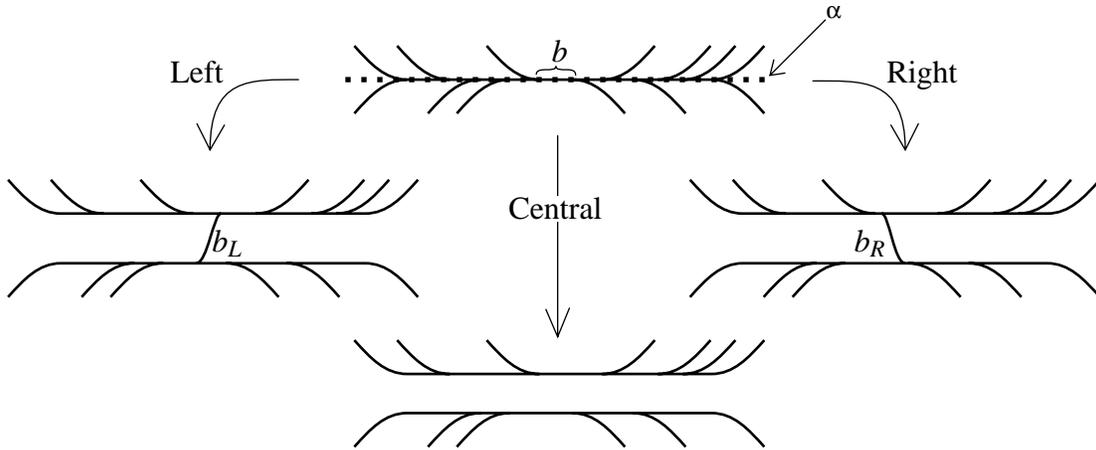


Figure 14: Wide splitting of a semigeranic train track along a splitting arc α .

on the number of comb moves from τ to σ : if σ is obtained from τ by collapsing a transition branch $b \subset \tau$, then for each splitting arc α of τ , the image of α under the quotient map which collapses b to a point is a splitting arc β of σ , and this induces the desired bijection.

We shall define left, right, and central wide splittings of a semigeranic pretrack τ along a splitting arc α . These are depicted in Figure 14. These wide splittings are denoted $\tau \succ_L^\alpha \tau_L$, $\tau \succ_C^\alpha \tau_C$, $\tau \succ_R^\alpha \tau_R$. When the splitting arc is explicitly or implicitly determined, we will often drop the adjective “wide” and simply say, for example, “the result of a left splitting of τ (along α) is τ_L ”.

We’ll give two equivalent definitions of wide splittings, the first similar to the definition of an elementary splitting, and the second which shows that a wide splitting is an instance of a general splitting by explicitly factoring it into comb moves and an elementary splitting.

We start with the central splitting $\tau \succ \tau_C$. The idea is to insert a knife and split the surface open along α , pushing one part of τ upward and the other part downward. To be precise, let D be a disc with two cusps embedded in S in such a way that α is properly embedded in D with its endpoints at the two cusps of D . Foliate D so that each leaf is an arc between the two sides of D with leaves degenerating to points at each of the two cusps, and so that α is transverse to the foliation intersecting each leaf in exactly one point. Choose a map $f: S \rightarrow S$ homotopic to the identity, so that f collapses each leaf of D to the point where that leaf intersects α , f is otherwise one-to-one, and the restriction of f to $S - \text{int}(D)$ is a smooth submersion. Defining $\tau_C = f^{-1}(\tau) - \text{int}(D)$, it follows easily that τ_C is a

train track, and we have defined the central splitting $\tau \succ \tau_C$ along α . Note that the map f is a carrying map from τ_C to τ , called a *pinch map* associated to the central splitting.

For later purposes we note that a central splitting $\tau \succ \tau_C$ can be defined along any path α that satisfies items (2,3) in the definition of a splitting arc, and the following weaker version of (1):

- (1') $\alpha \cap \tau$ is an embedded train path in τ with endpoints at switches, and at each endpoint the switch orientation points into $\alpha \cap \tau$.

If α satisfies (1'), (2), and (3) then we say that α is a *C-splitting arc*, and a central splitting $\tau \succ \tau_C$ along α is defined as above; note that in this situation Left and Right splittings will *not* be defined in general. Also, if α is a C-splitting arc then the number of sink branches in α is one more than the number of source branches, and so α is an ordinary splitting arc if and only if it contains exactly one sink branch.

Next we define the Left and Right wide splittings $\tau \succ \tau_L$, $\tau \succ \tau_R$ of τ along a splitting arc α . Continuing the above notation, note that $\tau_C \cap \partial D$ consists of two embedded train paths $\alpha_-, \alpha_+ \subset \tau_C$, each taken diffeomorphically to $\alpha \cap \tau$ by the pinch map f . We smoothly identify D with a subset of \mathbf{R}^2 so that each leaf of D is vertical in \mathbf{R}^2 , so that α_- and α_+ are horizontal, and so that orientation is preserved. Under this identification, $f^{-1}(\alpha \cap \tau)$ is a rectangle $R = [x_-, x_+] \times [y_-, y_+]$ where the notation is chosen so that $\alpha_- = [x_-, x_+] \times y_-$ and $\alpha_+ = [x_-, x_+] \times y_+$. By definition of a splitting arc, α contains a unique sink branch b . Let b^* be a subarc contained in the interior of b . Let $b_\pm^* = \alpha_\pm \cap f^{-1}(b^*)$, so $b_\pm^* = [x_\pm^*, x_\pm^*] \times y_\pm$ for some subinterval $[x_\pm^*, x_\pm^*] \subset (x_-, x_+)$. Let b_L be a smooth arc properly embedded in D , meeting α_- tangentially at the point $x_-^* \times y_-$, meeting α_+ tangentially at $x_+^* \times y_+$, and having positive slope everywhere in the rectangle R . Let b_R be similarly defined, meeting α_- at $x_+^* \times y_-$ and α_+ at $x_-^* \times y_+$, having negative slope. Define $\tau_L = \tau_C \cup b_L$ and $\tau_R = \tau_C \cup b_R$. Note that f is a carrying map from τ_L to τ , and from τ_R to τ , in this context called a *fold map* associated to a parity splitting. The branches b_L, b_R of τ_L, τ_R respectively are called the *postsplitting branches*. Note that τ_L, τ_R are evidently train tracks.

This completes the first definition of wide splitting.

Here is a second and equivalent definition of a wide splitting in the restricted case that τ is generic. Let $\alpha \cap \tau = a_1 * \cdots * a_m * b * c_n * \cdots * c_1$, written as a concatenation of m transition branches a_i , a sink branch b , and n transition branches c_j . By doing $m + n$ successive slide moves $\tau = \tau_0 \succ \cdots \succ \tau_{m+n}$ along the transition branches of α , starting from the outermost transition branches a_1, c_1 and working inward, it follows that there is an elementary splitting $\tau_{m+n} \succ \tau'$, and this τ' agrees with the

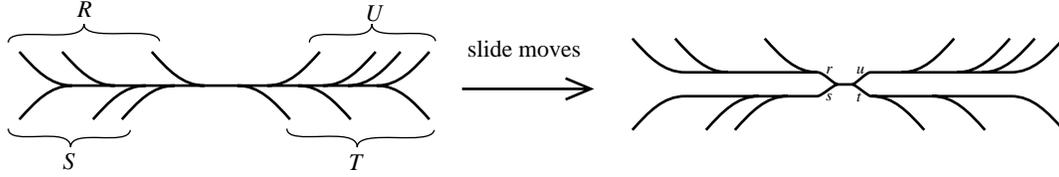


Figure 15: A wide splitting on a generic train track can be factored as an elementary splitting following a sequence of slide moves. The slide moves depicted here, applied to the top train track from Figure 14, result in a train track with sink branch depicted here, and following this by elementary splittings along the sink branch produces the same train tracks as the wide splittings depicted in Figure 14.

result of wide splitting of τ along α for any of the parities L, R, C (see Figure 15). This shows directly that $\tau \succ \tau'$ is a general splitting.

For a wide splitting $\tau \succ \tau'$ on a nongeneric train track τ , it is not hard to produce generic perturbations τ_1 of τ and τ'_1 of τ' so that $\tau_1 \succ \tau'_1$ is a wide splitting of the generic train track τ_1 , and having just shown that $\tau_1 \succ \tau'_1$ is a general splitting it follows that $\tau \succ \tau'$ is as well.

Note that if τ is generic and $\alpha \cap \tau = b$ is an sink branch, then the definition of wide splitting along α coincides with the definition of elementary splitting along b .

The following results exhibit the nice properties of wide splittings:

Proposition 3.13.3. *Given comb equivalent, semigeneric train tracks τ, σ , a splitting arc α of τ , the corresponding splitting arc β of σ , and $d \in \{L, R, C\}$, the result of a d -splitting of τ along α is comb equivalent to the result of a d -splitting of σ along β .*

Proof. By induction it suffices to prove this for a single comb move collapsing a transition branch b of τ to a point yielding σ . Choosing a regular neighborhood D of b we may assume that $\tau - D = \sigma - D$. If $\alpha \cap b = \emptyset$ then the proposition is obvious with $\beta = \alpha$. If $b \subset \alpha$ then the result of collapsing b produces the desired splitting arc β of σ and the proposition is again obvious. \diamond

Proposition 3.13.4 (Wide Splitting Proposition). *Given semigeneric train tracks τ, τ' the following are equivalent:*

- (1) *There exists a wide splitting $\tau \succ \tau'_1$ so that τ', τ'_1 are comb equivalent.*
- (2) *There exists a wide splitting $\tau_2 \succ \tau'_2$ so that τ, τ_2 are comb equivalent and τ', τ'_2 are comb equivalent.*

(3) For any semigeneric τ_3 comb equivalent to τ there exists a wide splitting $\tau_3 \succ \tau'_3$ so that τ', τ'_3 are comb equivalent.

(4) $\tau \succ \tau'$ is a general splitting.

Proof of Proposition 3.13.4. Obviously (3) \implies (1) \implies (2) \implies (4). Also, (4) \implies (2), because an elementary splitting is an instance of a wide splitting.

We prove that (2) \implies (3). Assuming (2), if τ_3 is comb equivalent to τ then it is comb equivalent to τ_2 . If α is the splitting arc of τ_2 along which the d -splitting $\tau_2 \succ \tau'_2$ is defined, $d \in \{L, R, C\}$, then by Proposition 3.13.3 there is a splitting arc β of τ_3 such that if $\tau_3 \succ \tau'_3$ is a d -splitting along β then τ'_2, τ'_3 are comb equivalent. \diamond

Corollary 3.13.5. *Given a semigeneric train track τ and a splitting arc α , the Splitting Lattice 3.13.2 holds for the three splittings of τ along α .*

Proof. There exists a generic train track τ' comb equivalent to τ , and an elementary splitting branch b of τ' , such that α corresponds to b under the comb equivalence $\tau \approx \tau'$. It follows from Proposition 3.13.4 that the three splittings $\tau \succ \tau_L, \tau \succ \tau_C, \tau \succ \tau_R$ of τ along α are comb equivalent to the three splittings $\tau' \succ \tau'_L, \tau' \succ \tau'_C, \tau' \succ \tau'_R$. The result now follows by applying Proposition 3.13.3 to τ' and b . \diamond

By induction using Proposition 3.13.4 we obtain the following result, which tells us that wide splittings can be used exclusively to study splitting sequences, if one so desires:

Corollary 3.13.6. *Given a sequence of splittings $\tau_0 \succ \tau_1 \succ \dots$ and a semigeneric train track τ'_0 comb equivalent to τ_0 there exists a sequence of wide splittings $\tau'_0 \succ \tau'_1 \succ \dots$ such that τ_i, τ'_i are comb equivalent for each i .* \diamond

3.14 Factoring a carrying map into splittings

Consider a carrying relation $\tau \succ \tau'$. Recall that it is a *full carrying*, denoted $\tau \overset{F}{\succ} \tau'$, if the carrying map $\tau' \rightarrow \tau$ is surjective. Also, we say that $\tau \succ \tau'$ is a *homotopic carrying*, denoted $\tau \overset{H}{\succ} \tau'$, if the carrying map $\tau' \rightarrow \tau$ is a homotopy equivalence. Because the carrying map $\tau' \rightarrow \tau$ is uniquely determined up to homotopy through carrying maps (Proposition 3.5.2), it follows that both of these properties are well-defined independent of the choice of carrying map.

The following proposition shows that full carrying is the transitive closure of comb equivalence and splitting, and homotopic carrying is the transitive closure of comb equivalence and parity splitting. A version of the following proposition is found in [Pen92], Theorem 2.4.1.

Proposition 3.14.1. *If τ, τ' are train tracks, then $\tau \xrightarrow{F} \tau'$ if and only if one of the following two mutually exclusive alternatives happens:*

- (1) τ is comb equivalent to τ' ; or
- (2) there exists a finite sequence of splittings

$$\tau = \tau_0 \succ \tau_1 \succ \cdots \succ \tau_n = \tau'$$

Furthermore, $\tau \xrightarrow{H} \tau'$ if and only if either (1) holds or

- (2') there exists a finite sequence of Left or Right splittings

$$\tau = \tau_0 \succ \tau_1 \succ \cdots \succ \tau_n = \tau'$$

In particular, $\tau \xrightarrow{H} \tau'$ implies $\tau \xrightarrow{F} \tau'$.

As a corollary we get a simple criterion for checking when a full carrying is a homotopy carrying. Given any train track τ on S let $i(\tau)$ be the number of cusps of $\mathcal{C}(S - \tau)$, and so by Lemma 3.2.1 we have $i(\tau) \leq 6|\chi(S)| - 2p$ where p is the number of punctures of S . Note that $i(\tau) = 2|\chi(\tau)|$, because for each switch s , if the valence equals v_s then the number of cusps at s is $v_s - 2$, and so

$$\begin{aligned} i(\tau) &= \sum_s (v_s - 2) = \sum_s v_s - (2 \cdot \#\text{vertices}(\tau)) \\ &= 2 \cdot \#\text{edges}(\tau) - 2 \cdot \#\text{vertices}(\tau) \\ &= 2|\chi(\tau)| \end{aligned}$$

Corollary 3.14.2. *If $\tau \xrightarrow{F} \tau'$ then $i(\tau) \geq i(\tau')$, with equality if and only if $\tau \xrightarrow{H} \tau'$. Moreover, if equality does not hold then $i(\tau) \geq i(\tau') + 2$.*

Proof. If τ, τ' are comb equivalent then this is clear. Otherwise, choose a splitting sequence $\tau = \tau_0 \succ \cdots \succ \tau_n = \tau'$. By Proposition 3.14.1, $\tau \xrightarrow{H} \tau'$ if and only if no Central splitting occurs in this splitting sequence. On the other hand, a parity splitting preserves i whereas a Central splitting decreases it by 2. \diamond

Proof of Proposition 3.14.1. By Proposition 3.12.2, (1) implies both $\tau \xrightarrow{F} \tau'$ and $\tau \xrightarrow{H} \tau'$. Given an elementary splitting $\tau \succ \tau'$, clearly $\tau \xrightarrow{F} \tau'$, and if the splitting has parity L or R then $\tau \xrightarrow{H} \tau'$. Given a general splitting $\tau \succ \tau'$, because it factors into comb equivalences and an elementary splitting, it follows

by transitivity that $\tau \overset{\text{F}}{\succ} \tau'$, and if the parity is L or R then $\tau \overset{\text{H}}{\succ} \tau'$. A further application of transitivity shows that (2) implies $\tau \overset{\text{F}}{\succ} \tau'$ and (2') implies $\tau \overset{\text{H}}{\succ} \tau'$.

Suppose that $\tau \overset{\text{F}}{\succ} \tau'$. By generic uncombing, we may assume that τ, τ' are generic. Apply Proposition 3.5.1 to obtain an inclusion of tie bundles $\nu(\tau') \subset \nu(\tau)$, and by isotopy we may assume that $\nu(\tau') \subset \text{int}(\nu(\tau))$. Let $p: \nu(\tau) \rightarrow \tau$ be the fibration map. Consider the surface $\text{Cl}(\nu(\tau) - \nu(\tau'))$, with tie bundle obtained by taking components of intersection with ties of $\nu(\tau)$. Let M be the subsurface of $\text{Cl}(\nu(\tau) - \nu(\tau'))$ whose ties have both endpoints in $\partial\nu(\tau')$. The tie bundle of M is actually a nonsingular foliation of M by arcs, with leaves degenerating to a point at each cusp of $\partial\nu(\tau')$ contained in ∂M . Note that each component of ∂M contains two “leaves” each of which is either a degenerate leaf at a cusp of $\nu(\tau')$ or a nondegenerate leaf containing a cusp of $\nu(\tau)$. Moreover, since $\nu(\tau')$ has neither bigons nor smooth annuli, it follows that each component of M has at least one nondegenerate boundary leaf. Let $M \rightarrow J$ be the quotient map crushing each leaf to a point, so J is a compact 1-manifold. Choose a section $J \rightarrow M$ such that each endpoint of J is a cusp of either $\nu(\tau)$ or $\nu(\tau')$, and so each component of J has at least one endpoint at a singularity of $\nu(\tau)$, corresponding to a nondegenerate boundary leaf of the corresponding component of M . Let J_1, \dots, J_K be the components. Each path $p \mid J_k$ is a train path in τ , with at least one endpoint at a switch s_k of τ such that the switch orientation at s_k points into $p \mid J_k$ at s_k ; such endpoints of $p \mid J_k$ correspond to endpoints of J_k in a nondegenerate boundary leaf of M . Now start splitting along these train paths, starting from the endpoint s_k . The first move is an elementary split or a slide as appropriate, depending on whether the first branch b along J_k is a sink branch or transition branch. If b is a transition branch then slide along b . If b is a sink branch, with endpoint s opposite s_k , then either J_k ends at s and we should do an elementary central splitting, or J_k turns to the Left or Right at s and we should do an elementary splitting of the corresponding parity. Let $\tau = \tau_0 \rightarrow \tau_1$ be this operation, and note that we get an inclusion $\nu(\tau') \subset \nu(\tau_1) \subset \nu(\tau)$, which allows us to continue elementary splits and slides along J by induction. Note that if the first move along J_k is a slide, then one of two alternatives happens: either J_k consists entirely of transition branches and so J_k is embedded by recurrence of τ , thereby generating a slide equivalence; or J_k contains some sink branch and the result, up to and including the first elementary splitting, is a general splitting. Continuing by induction, at the end of this sequence the resulting train track is τ' (without assuming full carrying, the best we can say is that τ' is a subtrack of the resulting train track at the end of the sequence).

Suppose now that $\tau \overset{\text{H}}{\succ} \tau'$. We claim that $\tau \overset{\text{F}}{\succ} \tau'$. If not, the carrying map $\tau' \rightarrow \tau$ factors through inclusion of a proper subtrack $\sigma \subset \tau$. It follows that the induced isomorphism $H_1(\tau') \rightarrow H_1(\tau)$ factors through the inclusion induced

homomorphism $H_1(\sigma) \rightarrow H_1(\tau)$. But since τ has no valence 1 vertices, the image of $H_1(\sigma) \rightarrow H_1(\tau)$ is a proper subset of $H_1(\tau)$, contradiction.

Since $\tau \xrightarrow{F} \tau'$, the above argument shows that either (1) or (2) holds. Assuming (2) holds, to prove (2') we must show that none of the splittings $\tau_i \succ \tau_{i+1}$ is central. If $\tau_i \succ \tau_{i+1}$ is a central splitting then one of two alternatives occurs: either τ_{i+1} has one more component than τ_i ; or the carrying map $\tau_{i+1} \rightarrow \tau_i$ induces a bijection of components, it restricts to a homeomorphism on all but one of the components, and the rank of the remaining component of τ_{i+1} is one less than the rank of its image in τ_i . On the other hand a parity splitting, being a homotopy equivalence, induces a bijection of components so that corresponding components have the same rank. This immediately implies that if there exists one or more central splittings in the splitting sequence then $\tau \xrightarrow{H} \tau'$ is *not* a homotopy equivalence.

To prove that (1) and (2) are mutually exclusive, suppose on the contrary that τ, τ' are comb equivalent, but that we also have a sequence of at least one splitting from τ to τ' . We may assume that τ, τ' are completely combed, and hence they are isotopic, by Proposition 3.12.1. We may also assume that there is a sequence of carrying relations $\tau = \tau_0 \succ \tau_1 \succ \cdots \succ \tau_k = \tau'$ so that each pair $\tau_{i-1} \succ \tau_i$ is either a comb equivalence or an elementary splitting between generic train tracks. For each $i = 1, \dots, k$ choose a carrying map $\tau_i \rightarrow \tau_{i-1}$; in particular, if $\tau_{i-1} \succ \tau_i$ is an elementary splitting, choose the carrying map constructed in the definition of an elementary splitting. By composition we get a carrying map $\tau' \rightarrow \tau$. By Lemma 3.12.3 this carrying map is a switch fold. However, it is clearly impossible to factor a switch fold as a composition $\tau' = \tau_k \rightarrow \cdots \rightarrow \tau_i \rightarrow \tau_{i-1} \rightarrow \cdots \rightarrow \tau_0 = \tau$ where $\tau_i \rightarrow \tau_{i-1}$ is the standard carrying map of an elementary splitting. This contradiction shows that (1) and (2) are mutually exclusive. \diamond

3.15 Combinatorial types of train tracks.

The combinatorial type of τ is determined by the pattern of gluing of the components of $\mathcal{C}(S - \tau)$. Two train tracks are combinatorially equivalent if their gluing patterns are the same. We make these notions precise as follows.

Consider a train track τ that fills S . Each component of $\mathcal{C}(S - \tau)$ is a non-punctured or once-punctured cusped polygon of negative Euler index. Consider the overlay map $m: \mathcal{C}(S - \tau) \rightarrow S$. The set $m^{-1}(\text{cusps}(\tau))$ is called the *vertex set* of $\mathcal{C}(S - \tau)$, it includes the cusps of $\mathcal{C}(S - \tau)$, and it includes points which are not cusps called *smooth vertices* of $\mathcal{C}(S - \tau)$. The vertex set of $\mathcal{C}(S - \tau)$ subdivides $\partial\mathcal{C}(S - \tau)$ into finitely many arcs, called the *branch preimages* of $\mathcal{C}(S - \tau)$. The overlay map $\mathcal{C}(S - \tau) \rightarrow S$ identifies branch preimages in pairs to produce the branches of τ , and it also identifies vertices of $\mathcal{C}(S - \tau)$ to produce the switches of τ , in such a way

that exactly two smooth vertices are identified to each switch of τ . The *gluing data* for τ consists of the cusped surface $\mathcal{C}(S - \tau)$ and the pairing of branch preimages.

Two filling train tracks τ, τ' are said to be *combinatorially equivalent*, or to have the same *combinatorial type*, if they satisfy one of the equivalent conditions in the following result:

Proposition 3.15.1 (Combinatorial equivalence of filling train tracks). *Given two filling train tracks τ, τ' , the following are equivalent:*

- (1) *There exists a bijection between the set of branch preimages of $\mathcal{C}(S - \tau)$ and the set of branch preimages of $\mathcal{C}(S - \tau')$, which respects cyclic ordering, cusps, smooth vertices, and gluing data.*
- (2) *There exists a diffeomorphism $f: \mathcal{C}(S - \tau) \rightarrow (S - \tau')$ that respects the vertices and gluing data.*
- (3) *There exists a diffeomorphism $F: (S, \tau) \rightarrow (S, \tau')$.*

Proof. The proofs of (1) \iff (2) \iff (3) are all immediate. \diamond

The advantage of definition (1) is that it is clearly finitistic: the combinatorial type of τ is specified by writing down a finite amount of data, consisting of the cycles of branch preimages, the smooth and cusp vertices, and the branch pre-image pairing. As a corollary there are finitely many combinatorial types of train tracks on a surface S of given topological type, and one can easily describe an algorithm to enumerate them. We will have more to say about this enumeration, at least with regard to the train tracks that occur in one cusp expansions, later on.

A diffeomorphism $F: (S, \tau) \rightarrow (S, \tau')$ is said to be a *combinatorial equivalence* between τ and τ' . Note that F induces a bijection between $\text{cusps}(\tau)$ and $\text{cusps}(\tau')$, and the following lemma tells us that this bijection depends only on the isotopy class of F .

Lemma 3.15.2 (Naturality of cusps). *Given filling train tracks τ, τ' and combinatorial equivalences $F, F': (S, \tau) \rightarrow (S, \tau')$, the following are equivalent:*

- (1) *F, F' are isotopic.*
- (2) *F, F' induce the same bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$.*
- (3) *There exists $v \in \text{cusps}(\tau)$ such that $F(v) = F'(v)$ in $\text{cusps}(\tau')$.*

When we augment a train track with extra data then we can define combinatorial equivalence of the augmented train track. For example, define a *one cusp train track* to be an ordered pair (τ, v) where τ is a train track and $v \in \text{cusps}(\tau)$. Two one cusp

train tracks (τ, v) , (τ', v') are combinatorial equivalent if there exists $F \in \text{Homeo}_+$ such that $F(\tau) = \tau'$ and $F(v) = v'$.

Given a filling train track $\tau \subset S$, the *stabilizer* of τ is the subgroup $\text{Stab}(\tau) \subset \mathcal{MCG}$ that is represented by $F \in \text{Homeo}_+$ such that $F(\tau) = \tau$. The stabilizer $\text{Stab}(\tau, v)$ of a one cusp train track is similarly defined. Given two filling train tracks τ, τ' , define $C(\tau, \tau')$ to be the set of elements in \mathcal{MCG} that are represented by combinatorial equivalences from τ to τ' . The following result is an immediate consequence of the naturality of cusps 3.15.2.

Corollary 3.15.3 (Rigidity of cusps).

- If τ is a filling train track on S then $\text{Stab}(\tau)$ is a finite subgroup. Moreover, if (τ, v) is a filling, one cusp train track then $\text{Stab}(\tau, v)$ is the trivial subgroup of \mathcal{MCG} .
- If τ, τ' are two filling train tracks on S then the set of mapping classes taking τ to τ' (up to isotopy) is a finite subset; in fact, if it is not empty then it is a left coset of $\text{Stab}(\tau)$ and a right coset of $\text{Stab}(\tau')$. Moreover, given filling, one cusp train tracks (τ, v) , (τ', v') , there is at most one mapping class taking (τ, v) to (τ', v') (up to isotopy).

◇

Proof of Lemma 3.15.2. To prove (1) \implies (2), by replacing τ with $F(\tau)$ we may assume that F, F' are isotopic to the identity and both are carrying maps from τ to τ' . Now mimic the proof of Proposition 3.5.2, showing that two ambient diffeomorphisms $(S, \tau) \rightarrow (S, \tau')$ that are each isotopic to the identity are isotopic to each other through diffeomorphisms $(S, \tau) \rightarrow (S, \tau')$. The bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$ is evidently constant throughout this isotopy.

The implication (2) \implies (3) is obvious.

To prove (3) \implies (1), starting from v , add one branch at a time to produce a larger and larger sequence connected subgraphs of τ . Each time a branch is added, one proves inductively that F and F' take that branch to the same branch of τ' , using the fact that F and F' both respect cyclic ordering of branch ends around each switch. It follows that F and F' take each switch of τ to the same switch of τ' , and they take each branch of τ to the same branch of τ' . By considering the boundary 1-cycle of each component of $\mathcal{C}(S - \tau)$, it follows that F and F' take each component of $\mathcal{C}(S - \tau)$ to the same component of $\mathcal{C}(S - \tau')$. It follows that by isotoping along branches and extending we can construct an ambient isotopy of F after which $F|_{\tau}$ and $F'|_{\tau}$ are identical. Because each component of $\mathcal{C}(S - \tau)$ and of $\mathcal{C}(S - \tau')$ is a nonpunctured or once-punctured disc, and because F, F' preserve punctures, we can follow this by another ambient isotopy after which $F = F'$. ◇

4 Train track expansions of measured foliations

In this section we formalize the concept of train track expansions, and explain their construction; this much is done very quickly in Section 4.1. In Section 4.2 we then show that a measured foliation is annular if and only if its fully carrying train track expansions are finite.

For future applications we shall require a unified description of *all* train track expansions of a given measured foliation \mathcal{F} that are based at a given train track τ that fully carries \mathcal{F} . This is carried out in Sections 4.3 and 4.4. The main idea is that, after choosing \mathcal{F} within its equivalence class so that it has a carrying bijection onto a tie bundle ν over τ , an expansion of \mathcal{F} based at τ corresponds to an expanding family of finite separatrices of \mathcal{F} . This idea is used in Sections 4.3 to give another construction of train track expansions. Then in Section 4.4 we prove Proposition 4.4.1, which says that this new construction given in Section 4.3 is universal: every expansion of \mathcal{F} based at τ arises by the new construction. Proposition 4.4.1 will be a key technical tool used in several later results.

4.1 Definition and existence of train track expansions

A *splitting sequence* means any finite or infinite sequence of splittings of the form $\tau_0 \succ \tau_1 \succ \dots$. A splitting sequence is *complete* if it is either infinite or it is finite and the last train track τ_n is an *essential curve family*, that is, a disjoint union of essential simple closed curves.

The number of central splittings in any splitting sequence on S is bounded by $3|\chi(S)| - p$, for by Lemma 3.2.1 any train track has at most $6|\chi(S)| - 2p$ cusps, and a parity splitting preserves the number of cusps while a central splitting reduces the number by 2. By the same reasoning, a finite, complete splitting sequence starting at a train track τ will have exactly $i(\tau)/2$ central splittings, where $i(\tau)$ is the number of cusps of τ .

Given a measured foliation \mathcal{F} , a complete splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ is a *train track expansion* of \mathcal{F} if

$$[\mathcal{F}] \in \bigcap_{n \geq 0} \mathcal{MF}(\tau_n)$$

In other words, for each n there is a partial measured foliation equivalent to \mathcal{F} which is carried by τ_n .

As mentioned in the introduction, train track expansions are easily constructed:

Theorem 4.1.1. *For any measured foliation \mathcal{F} and any train track τ there exists a train track expansion $\tau = \tau_0 \succ \tau_1 \succ \dots$ of \mathcal{F} such that each τ_i fully carries \mathcal{F} .*

Proof. By induction it suffices to show that if τ fully carries \mathcal{F} , and if α is a splitting arc of τ , then there exists a splitting $\tau \succ \tau'$ along α such that τ' fully carries \mathcal{F} . This is an immediate consequence of the Splitting Lattice 3.13.2 in the case of an elementary splitting of a generic train track, or Corollary 3.13.5 for wide splittings. \diamond

By combining Theorem 4.1.1 with Proposition 3.6.1 to obtain τ , we have:

Corollary 4.1.2 (Existence of expansions). *Every measured foliation \mathcal{F} has a train track expansion.* \diamond

4.2 Expansions of annular measured foliations

Measured foliations with finite train track expansions are precisely those whose foliation components are all annular, in analogy with the fact that a real number has a finite continued fraction expansion if and only if it is rational.

Lemma 4.2.1. *If the measured foliation \mathcal{F} has a finite train track expansion, ending at some τ_N which is a curve family, then \mathcal{F} is annular, and the core of each annulus component of \mathcal{F} is isotopic to some component of τ_N . Conversely, if the measured foliation \mathcal{F} is annular, and if $\tau_0 \succ \tau_1 \succ \dots$ is any splitting sequence such that each τ_i fully carries \mathcal{F} , then the sequence is finite.*

Proof. The first sentence is obvious from the definitions. The second sentence is well known; here is a detailed proof.

For the moment we make no assumption that \mathcal{F} is annular; that assumption will be brought in at the very last moment.

Consider a train track τ fully carrying a measured foliation \mathcal{F} . Let $w(b; \mathcal{F})$ be the weight induced on a branch $b \subset \tau$ by \mathcal{F} . Enumerate the sink branches of τ as b_1, b_2, \dots, b_k so that $w(b_1; \mathcal{F}) \geq w(b_2; \mathcal{F}) \geq \dots \geq w(b_k; \mathcal{F})$; note that k is bounded by a constant depending only on the topology of S . Define $\omega(\tau, \mathcal{F})$ to be this decreasing sequence of positive numbers. The sequence $\omega(\tau, \mathcal{F})$ measures, in some sense, how good of an approximation τ is of \mathcal{F} : the smaller the values of $\omega(\tau, \mathcal{F})$, the better the approximation. Order the set of nonincreasing positive sequences so that $(w_i \mid i = 1, \dots, k) > (w'_i \mid i = 1, \dots, k')$ if either $k > k'$ and $w_i = w'_i$ for all $i = 1, \dots, k'$, or there exists an integer j with $1 \leq j \leq \min\{k, k'\}$ such that $w_i = w'_i$ for $i < j$ and $w_j > w'_j$.

Noting the obvious fact that $\omega(\tau, \mathcal{F})$ is unchanged by a comb equivalence, we now prove:

Lemma 4.2.2. *Given a splitting $\tau \succ \tau'$ and a measured foliation \mathcal{F} fully carried by τ and by τ' , we have*

$$\omega(\tau, \mathcal{F}) > \omega(\tau', \mathcal{F})$$

Given a splitting $\tau \succ \tau'$ along a sink branch b , where τ, τ' each fully carry \mathcal{F} , by changing τ, τ' within their comb equivalence classes we may assume that $\tau \succ \tau'$ is an elementary splitting along a sink branch b . We use the notation of Figure 13, and we apply the Splitting Inequalities 3.13.1. We claim that the sequence $\omega(\tau', \mathcal{F})$ is obtained from $\omega(\tau, \mathcal{F})$ by deleting the entry $w(b; \mathcal{F}) = w(\overline{pq}, \mathcal{F})$ and inserting up to two smaller entries. We will prove this by carefully describing the inserted entries in $\omega(\tau', \mathcal{F})$, assuming that $\tau \succ \tau'$ is a Left splitting; Right splittings are handled by the exact same argument after reversing orientation, and Central splittings are handled by a similar argument.

Referring to Figure 13, it is possible that the points s, u lie on the same source branch b_s of τ , in which case s, u lie on the same sink branch b'_s of τ' , and in this case we insert one entry into $\omega(\tau', \mathcal{F})$, namely

$$w(b'_s; \mathcal{F}) = w(b_s; \mathcal{F}) = w(\overline{pq}, \mathcal{F}) - w(\overline{rp}, \mathcal{F}) < w(\overline{pq}, \mathcal{F}) = w(b, \mathcal{F})$$

where the strictness of the inequality follows from the fact that τ fully carries \mathcal{F} and so $w(\overline{rp}, \mathcal{F}) > 0$.

Assume now that s, u do not lie on the same source branch of τ . It follows that s lies on a sink branch b'_s of τ' if and only if s lies on a transition branch b_s of τ , in which case we insert an entry $w(b'_s; \mathcal{F}) < w(b; \mathcal{F})$ by the same inequalities above. Similarly, u lies on a sink branch b'_u of τ' if and only if u lies on a transition branch b_u of τ , in which case we insert the entry $w(b'_u; \mathcal{F}) < w(b; \mathcal{F})$.

This completes the proof that $\omega(\tau; \mathcal{F})$ is strictly decreased by a splitting, when τ fully carries \mathcal{F} .

Suppose now that \mathcal{F} is an annular measured foliation and $\tau_0 \succ \tau_1 \succ \dots$ is any expansion of \mathcal{F} where each τ_i fully carries \mathcal{F} . Since carrying depends only on the underlying unmeasured foliation of \mathcal{F} , we may change the measure on \mathcal{F} so that the total transverse measure of each annular foliation component of \mathcal{F} is equal to 1, and $\tau_0 \succ \tau_1 \succ \dots$ is still an expansion of \mathcal{F} . It follows that $\omega(\tau, \mathcal{F})$ has integer entries, and since $\omega(\tau, \mathcal{F})$ has length at most K , where K is a constant depending only on the topology of S , it follows that $\omega(\tau, \mathcal{F})$ takes values in a well-ordered set. It follows that the number of splittings in any splitting sequence fully carrying \mathcal{F} is finite. \diamond

4.3 A unified construction of train track expansions

Despite the simplicity of the proof of Theorem 4.1.1, we wish to give a different viewpoint on the construction of train track expansions. This will produce a unified picture of *all* train track expansions of a measured foliation \mathcal{F} based at a given train track τ that fully carries \mathcal{F} , as we shall prove in Proposition 4.4.1. The latter

result will be useful in several situations later, where we analyze the properties of an arbitrary train track expansion of \mathcal{F} .

Suppose that \mathcal{F} is a measured foliation fully carried on a train track τ . In order to construct an expansion beginning with τ it suffices to find an expansion beginning with any comb equivalent train track, and so we may assume that τ is generic.

Let ν denote a tie bundle of τ . Replacing \mathcal{F} by an equivalent partial measured foliation, we may assume that \mathcal{F} has a carrying bijection into ν . Thus, we shall regard $\nu = \text{Supp}(\mathcal{F})$ as a foliated tie bundle, whose horizontal foliation is \mathcal{F} , and whose ties form the *vertical foliation* of ν denoted $\mathcal{F}^v(\nu)$. The singularities of \mathcal{F} are all 3-pronged boundary singularities, coinciding with the set of cusps of ν .

Choose once and for all a positive transverse Borel measure on $\mathcal{F}^v(\nu)$. Restricting this measure to each leaf segment of \mathcal{F} makes the leaf segment isometric to an interval in the real line.

The train track of a separatrix family. To review the setup, \mathcal{F} is a partial measured foliation with a carrying bijection $\mathcal{F} \hookrightarrow \nu$, for some tie bundle $\nu \rightarrow \tau$ over a train track τ . A *separatrix family* of \mathcal{F} is a union of proper, finite separatrices of \mathcal{F} , one for each proper separatrix germ; we allow the case of a degenerate proper separatrix consisting of just a boundary singularity. Because we have fixed a carrying bijection $\mathcal{F} \hookrightarrow \nu$, all singularities of \mathcal{F} are 3-pronged boundary singularities, and it follows that for each separatrix family ξ , if we write ξ as the disjoint union of its components $\xi = \sqcup_i \xi_i$, then each ξ_i is either a nondegenerate proper separatrix, or a degenerate proper separatrix, or a proper saddle connection of \mathcal{F} . Let $\partial\xi$ be the set of all points which are either degenerate separatrices of ξ or nonsingular endpoints of nondegenerate separatrices of ξ . Denote $\text{int } \xi = \xi - \partial\xi$.

To each separatrix family ξ we associate a tie bundle $\nu(\xi)$ over a train track $\sigma(\xi)$, and a partial measured foliation $\mathcal{F}(\xi)$ with support $\nu(\xi)$ and transverse to the ties, so that \mathcal{F} is equivalent to $\mathcal{F}(\xi)$. The tie bundle $\nu(\xi)$ with its horizontal foliation $\mathcal{F}(\xi)$ is obtained simply by slicing along ξ ; fact 2.6.2 shows that $\mathcal{F}(\xi)$ is equivalent to \mathcal{F} , with a partial fulfillment map $r: (S, \mathcal{F}(\xi)) \rightarrow (S, \mathcal{F})$. By pulling back the tie foliation $\mathcal{F}^v(\nu)$ via the map r we obtain a tie foliation of $\nu(\xi)$ which we denote $\mathcal{F}^v(\nu(\xi))$. Again, each cusp of $\nu(\xi)$ is a 3-pronged boundary singularity of $\mathcal{F}(\xi)$ and a 3-pronged boundary transverse singularity of $\mathcal{F}^v(\nu(\xi))$, and the foliations $\mathcal{F}(\xi)$ and $\mathcal{F}^v(\nu(\xi))$ are otherwise nonsingular and transverse. Now define the train track $\sigma(\xi)$ by collapsing to a point each tie of $\mathcal{F}^v(\nu(\xi))$. The foliated tie bundle $\nu(\xi)$ thus serves a dual role: as the support of $\mathcal{F}(\xi)$; and as a tie bundle over $\sigma(\xi)$. Since \mathcal{F} and $\mathcal{F}(\xi)$ are equivalent in \mathcal{MF} it follows that $\sigma(\xi)$ fully carries \mathcal{F} .

Notice that $\sigma(\xi)$ depends not only on ξ but also on the tie foliation $\mathcal{F}^v(\nu)$.

Here is another description of the tie bundle $\nu(\xi)$ and the train track $\sigma(\xi)$. Given

a separatrix family ξ of \mathcal{F} , a *shunt* of ξ is a subsegment α of a tie of ν such that $\text{int}(\alpha) \cap \xi$ is a nonempty subset of $\partial\xi$, and α is maximal with respect to this property. It follows that $\partial\alpha \subset \partial\nu \cup (\xi - \partial\xi)$, and so $\alpha \cap \partial\xi \subset \text{int}(\alpha)$. Each shunt α of ξ contains at least one point of $\partial\xi$, and the *degeneracy* of α , denoted $d(\alpha)$, is the cardinality of $\alpha \cap \partial\xi$. We say that α is *generic* if $d(\alpha) = 1$. We say that ξ is generic if it has no degenerate separatrices and each shunt is generic. A shunt α is *semigeneric* if ξ approaches each point of $\alpha \cap \partial\xi$ from the same side of α . The separatrix family ξ is *semigeneric* if it has no degenerate separatrices and each shunt is semigeneric. In particular, generic implies semigeneric.

Let A be the union of the shunts of ξ . Now let us momentarily forget all about the tie foliation $\mathcal{F}^v(\nu)$ *except for* the set A , a union of tie segments. Although we mentioned above that $\sigma(\xi)$ depends on $\mathcal{F}^v(\nu)$, we wish to show now that $\sigma(\xi)$ actually depends only on ξ and A , meaning that $\sigma(\xi)$ can be reconstructed entirely from the surface ν , the set of segments ξ , and the set of segments A , without regard to forgotten structure of $\mathcal{F}^v(\nu)$. First slice open ν along ξ as above to produce a subsurface $\nu(\xi)$ equipped with a map $r: \nu(\xi) \rightarrow \nu$, and note that $\partial\nu(\xi) = r^{-1}(\partial\nu \cup \xi)$. Although we have mostly forgotten about the vertical foliation, we have retained A ; let $A' = r^{-1}(A)$. Notice that each component R of $\nu(\xi) - (\partial\nu(\xi) \cup A')$ can be given the structure of the interior of a rectangle, with horizontal sides on $\partial\nu(\xi)$ and vertical sides on A' ; to be more precise, the completion $\mathcal{C}(R)$ is a closed disc whose boundary subdivides into four segments, two opposite segments called the *horizontal sides* mapping to $\partial\nu(\xi)$, and the other two segments called the *vertical sides* mapping to A' . We may therefore construct a vertical foliation of the closure \overline{R} , each leaf having one endpoint on one horizontal side. This vertical foliation is well-defined up to an isotopy that fixes each point of $\overline{R} \cap A'$ pointwise. Taking the union over all R defines a new vertical foliation of $\nu(\xi)$. One can now define a train track by collapsing each leaf of the new vertical foliation to a point. But the new vertical foliation of $\nu(\xi)$ is isotopic rel A' to the forgotten vertical foliation obtained by pulling back of $\mathcal{F}^v(\nu)$, and it follows that the new train track is isotopic to $\sigma(\xi)$.

From either construction of $\sigma(\xi)$ it follows easily that $\sigma(\xi)$ is (semi)generic if and only if ξ is (semi)generic.

The space of separatrix families. Let Ξ (an upper case ξ) be the space of separatrix families of \mathcal{F} , topologized using the Hausdorff topology on compact subsets of S . For each finite separatrix ξ_i , let $\text{Length}(\xi_i)$ be its total length measured with respect to the transverse measure on $\mathcal{F}^v(\nu)$. For each $\xi \in \Xi$ with components $\xi = \sqcup_i \xi_i$, define $\text{Length}(\xi) = \sum_i \text{Length}(\xi_i)$. The function $\text{Length}: \Xi \rightarrow [0, \infty)$ is continuous and proper. The set of generic separatrix families forms an open dense subset of Ξ .

Canonical models. Recall that \mathcal{F} has a carrying bijection $\mathcal{F} \hookrightarrow \nu$ into the tie bundle ν over a train track τ that fully carries \mathcal{F} .

There is a parameterization of Ξ which is most well-behaved when \mathcal{F} is canonically carried by τ , that is, when \mathcal{F} is a canonical model. In this case no component of a separatrix family is a saddle connection. Thus, letting I be the set of germs of proper separatrices of \mathcal{F} , each $\xi \in \Xi$ has a canonical decomposition $\xi = \sqcup_{i \in I} \xi_i$ where ξ_i is the (possibly degenerate) component of ξ representing the germ i . We obtain a homeomorphism $\mathcal{L}: \Xi \rightarrow [0, \infty)^I$ whose coordinate maps \mathcal{L}_i are

$$\mathcal{L}_i(\xi) = \text{Length}(\xi_i)$$

The decomposition of a separatrix family $\xi = \sqcup_i \xi_i$ induces a stratification of Ξ labelled by subsets of I : the stratum Ξ_J corresponding to a subset $J \subset I$ is the set of $\xi \in \Xi$ such that ξ_i is nondegenerate if and only if $i \in J$. The homeomorphism \mathcal{L} takes the stratum Ξ_J to the open J -coordinate subspace $\{x \in [0, \infty)^I \mid x_i \neq 0 \iff i \in J\}$.

The map which associates to each canonical model \mathcal{F} its space of separatrix families Ξ is functorial, in the sense that if $\mathcal{F}, \mathcal{F}'$ are isotopic canonical models and if $f: S \rightarrow S$ is isotopic to the identity taking \mathcal{F} to \mathcal{F}' , then there is an induced homeomorphism $f_*: \Xi \rightarrow \Xi'$ between the spaces of separatrix families, and all the usual axioms of a functor are satisfied. Moreover, f_* preserves strata in the sense that if I, I' are the proper separatrix germs of $\mathcal{F}, \mathcal{F}'$ respectively, and if $f_*: I \rightarrow I'$ is the induced bijection, then for each subset $J \subset I$ we have $f_*(\Xi_J) = \Xi'_{f_*(J)}$. Even more, the induced bijection $f_*: I \rightarrow I'$ is independent of the choice of f , meaning that any other isotopy from \mathcal{F} to \mathcal{F}' induces the same bijection $I \rightarrow I'$; this is an immediate consequence of Proposition 2.8.1.

The general case. In the general case, where we do not assume that \mathcal{F} is a canonical model, the space of separatrix families Ξ is defined and topologized as above, but the parameterization of Ξ is a little more complicated. Let I be the set of separatrix germs of \mathcal{F} which do not lie on saddle connections, let J be the set of saddle connections of \mathcal{F} , and for each $j \in J$ let I_j be the pair of separatrix germs associated to the two ends of j . For each $j \in J$ and each $\xi \in \Xi$, *either* ξ contains the whole saddle connection j of length $\text{Length}(j)$, *or* ξ contains disjoint separatrices, one for each end of j , whose lengths sum to less than $\text{Length}(j)$. This suggests the following: for each $j \in J$, start with the ordered pairs in $[0, \infty)^{I_j} \approx [0, \infty) \times [0, \infty)$ whose two coordinates sum to at most $\text{Length}(j)$, and pass to a quotient space by identifying to a single point all those ordered pairs whose two coordinates sum to exactly $\text{Length}(j)$; let \mathbf{R}_j denote this quotient space. For convenience we shall identify the special point of the quotient space \mathbf{R}_j with the real number $\text{Length}(j)$,

so a point in \mathbf{R}_j is *either* an ordered pair of positive numbers whose sum is less than $\text{Length}(j)$ or the single number $\text{Length}(j)$. By taking lengths in the canonical decomposition of $\xi \in \Xi$ we obtain a homeomorphism \mathcal{L} between Ξ and the space

$$([0, \infty)^I \times (\times_{j \in J} \mathbf{R}_j))$$

Partial ordering on Ξ . Another useful structure on the space of separatrix families Ξ is the partial ordering by inclusion, which agrees with coordinatewise partial ordering in the following sense. When \mathcal{F} has no saddle connections then $\xi \subseteq \xi'$ if and only if $\mathcal{L}(\xi)(i) \leq \mathcal{L}(\xi')(i)$ for each $i \in I$. When \mathcal{F} has saddle connections a similar statement is true, with the proviso that for $j \in J$ the partial ordering on the j -coordinate space \mathbf{R}_j has the special point $\text{Length}(j)$ as a global maximum, and otherwise the elements \mathbf{R}_j which are ordered pairs of real numbers are ordered coordinatewise.

When \mathcal{F} has only annular foliation components then there are no infinite separatrices and so $\Xi \approx \times_{j \in J} \mathbf{R}_j$. It follows that Ξ has a unique maximum point, namely the union of saddle connections; this is also the unique point where the length function $\text{Length}(\xi)$ achieves its maximum.

When \mathcal{F} has at least one partial arational foliation component then Ξ has no maximum point and the length function $\text{Length}(\xi)$ has no upper bound. The space Ξ is noncompact and has an end, called the *positive end*, with a neighborhood basis in the end compactification given by sets of the form $\text{Length}^{-1}(L, +\infty)$.

Expanding separatrix families. An *expanding separatrix family* is a continuous map from a closed, connected subset $J \subset \mathbf{R}$ into Ξ , denoted $t \rightarrow \xi(t)$ for $t \in J$, such that J is a strictly increasing map with respect to the partial ordering on Ξ . We say that $t \rightarrow \xi(t)$ is *complete* if it has no upper bound in Ξ , equivalently, one of the following happens: $\cup_{t \in J} \xi(t)$ contains an infinite separatrix in some arational component of \mathcal{F} ; or \mathcal{F} is annular and $\cup_{t \in J} \xi(t)$ is the unique maximum point of Ξ , equal to the union of all saddle connections of \mathcal{F} .

We study various events that occur in expanding separatrix families: a “central splitting event”, a “parity splitting event”, and a “slide event”. These events, visualized in Figure 16, exactly parallel the elementary splits and slides among generic train tracks.

A parameter value $t \in J$ is a *nonevent* if there exists $\epsilon > 0$ such that the separatrix families $\xi(t-\epsilon, t+\epsilon)$ are generic and their saddle connections are constant. It follows that the train tracks $\sigma(\xi(t-\epsilon, t+\epsilon))$ are constant up to isotopy.

A parameter value $t \in J$ is a *central splitting event* if, as $s \in J$ increases past t , a new saddle connection appears in $\xi(t)$, and nothing else happens. That is, there exists $\epsilon > 0$ and a saddle connection β of \mathcal{F} such that: the families $\xi(t-\epsilon, t+\epsilon)$ are

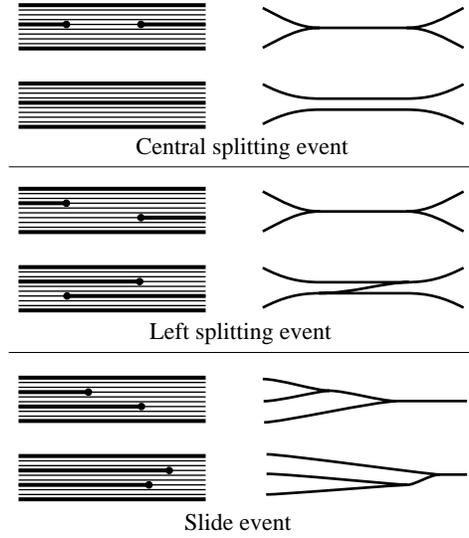


Figure 16: Events.

generic; the saddle connections for $\xi(t - \epsilon, t)$ are constant; and the saddle connections for $\xi[t, t + \epsilon)$ are constant, consisting of β together with the saddle connections for $\xi(t - \epsilon, t)$. Under these circumstances, the train tracks $\sigma(\xi(t - \epsilon, t))$ are constant up to isotopy, equal to τ say; the train tracks $\sigma(\xi[t, t + \epsilon))$ are constant up to isotopy, equal to τ' say; and there is an elementary central splitting $\tau \succ \tau'$.

To define a *parity splitting event*, we first define a special kind of shunt α , called a *parity splitting shunt*, which means that $d(\alpha) = 2$ and ξ approaches the two points of $\alpha \cap \partial\xi$ from opposite sides of α (see Figure 17). A parameter value $t \in J$ is a *parity splitting event* if, as s passes t , a parity splitting shunt α appears momentarily, and nothing else happens. As $s \in J$ passes t , the separatrices whose endpoints lie on α pass by each other in opposite directions like ships passing in the night. To be precise, there exists $\epsilon > 0$ such that: the shunts of $\xi(t)$ are all generic except for one parity splitting shunt α ; the families $\xi(t - \epsilon, t)$ and $\xi(t, t + \epsilon)$ are generic; and the saddle connections of $\xi(t - \epsilon, t + \epsilon)$ are constant. The train tracks $\sigma(\xi(t - \epsilon, t))$ are constant up to isotopy, equal to τ say; the train tracks $\sigma(\xi(t, t + \epsilon))$ are constant up to isotopy, equal to τ' say; and we have an elementary Left or Right splitting $\tau \succ \tau'$. The parity of $\tau \succ \tau'$ depends on whether the two ships pass each other on their left sides or their right sides.

To define a *slide event*, we first define a *slide shunt* to be a shunt α such that $d(\alpha) = 2$ and α is not a parity splitting shunt, equivalently, ξ approaches the two points of $\alpha \cap \partial\xi$ from the same side of α (see Figure 17). A parameter value $t \in J$ is

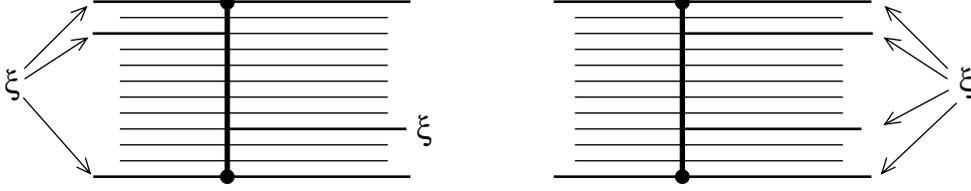


Figure 17: On the left: a parity splitting shunt. On the right: a slide shunt.

a *slide event* if, as s passes t , a slide shunt α comes into existence momentarily, and the separatrices whose endpoints lie on α pass each other like one ship overtaking another as they go along parallel courses in the same direction. To be precise, there exists $\epsilon > 0$ with the following properties: the shunts of $\xi(t)$ are all generic except for one slide shunt; the separatrix families $\xi(t - \epsilon, t)$ and $\xi(t, t + \epsilon)$ are generic; the saddle connections of $\xi(t - \epsilon, t + \epsilon)$ are constant; and the following property holds:

- Let i, i' be the separatrix germs for the two separatrices $\xi_i(t), \xi_{i'}(t)$ of $\xi(t)$ having endpoints $x, x' \in \partial\alpha$. Then the sign of $\text{Length}(\xi_i(s)) - \text{Length}(\xi_{i'}(s))$ changes as s passes through t from positive to negative or from negative to positive.

The train tracks $\sigma(\xi(t - \epsilon, t))$ are constant up to isotopy, equal to τ say, the train tracks $\sigma(\xi(t, t + \epsilon))$ are constant up to isotopy, equal to τ' say, and $\tau \succ \tau'$ is a slide.

We now define an expanding separatrix family $\xi(t), t \in J$, to be *generic* if every parameter value $t \in J$ is either a nonevent, a central splitting event, a parity splitting event, or a slide event. It follows that the events occur at a discrete set of times, and so the events subdivide the interval J into a concatenation of subintervals $J_i \subset \mathbf{R}$, for i in some subinterval $I \subset \mathbf{Z}$, such that if $i, i + 1 \in I$ then $\text{Max } J_i = \text{Min}(J_{i+1})$. For each $i \in I$ the generic train tracks $\sigma(\xi(\text{int } J_i))$ are in the same isotopy class, represented by τ_i say. Each $\tau_i \rightarrow \tau_{i+1}$ is an elementary splitting or a slide, depending on the type of the event at the parameter value $\text{Max } J_i = \text{Min } J_{i+1}$. We refer to $(\tau_i)_{i \in I}$ as the *elementary move sequence* associated to the family $\xi(t)$.

It is evident that any complete expanding separatrix family $t \rightarrow \xi(t), t \in J$, can be perturbed to be generic, where the perturbation is constant on any open subset of J where the family is already generic.

Here is the central technical point in the construction of a train track expansion:

Claim 4.3.1. *If $t \rightarrow \xi(t), t \in J$, is a complete, generic expanding separatrix family, and if \mathcal{F} has an arational component \mathcal{F}' , then the family has infinitely many splitting events.*

For the proof, we may choose \mathcal{F}' so that $\cup_{t \in J} \xi(t)$ contains an infinite separatrix of \mathcal{F}' with germ i . For each $t \in J$ let $x(t)$ be the point of $\partial \xi$ lying on $\xi_i(t)$. Consider a nonevent $s \in J$; it suffices to prove the existence of a splitting event $t > s$ in J . Corresponding to the point $x(t)$ is some switch of the train track $\sigma(\xi(t))$ denoted $y(t)$. At time s , the switch $y(s)$ is in the basin $A(b)$ of some sink branch b of the train track $\sigma(\xi(s))$. As $t \geq s$ increases, and $\xi_i(t)$ gets longer without bound, the switches in opposite half-basins of $A(b)$ approach each other, perhaps undergoing some slide events, until at some time $t \geq s$ the sink branch collapses to a point; this time t is a splitting event, proving the claim.

Constructing train track expansions using separatrix families. To complete our second construction, choose a complete, generic, expanding separatrix family $t \rightarrow \xi(t)$, $t \in J$. Let $t_0 < t_1 < \dots$ be the ordered list of events, starting with $t_0 = \inf(J)$. Choose $s_i \in (t_i, t_{i+1})$, and define $\tau_i = \sigma(\xi(s_i))$, and so each τ_i fully carries \mathcal{F} . Since $\tau = \sigma(\xi(t_0))$ is generic it follows that $\tau = \tau_0$.

Now let $i_1 < i_2 < \dots$ be the ordered list of all indices such that t_{i_n} is a splitting event. It follows that $\tau_{i_n} \succ \tau_{i_{n+1}}$ is a general splitting for each n , because it factors into slide moves and a single elementary splitting.

When \mathcal{F} has an arational component the splitting sequence $\tau_{i_0} \succ \tau_{i_1} \succ \dots$ is complete, by Claim 4.3.1. When \mathcal{F} has only annular components then the last event $t_{I_N} = t_I = \sup(J)$ is a central splitting event and $\sigma(\xi(t_I))$ is obtained by splitting \mathcal{F} along the union of all saddle connections, so $\sigma(\xi(t_I))$ is a union of simple closed curves; by our conventions, the last train track in the sequence is $\tau_{i_{N-1}} = \sigma(\xi(t_{i_{N-1}}))$, but the splitting $\tau_{i_{N-1}} \succ \sigma(\xi(t_I))$ factors into slide moves and a single central splitting, and so the splitting sequence $\tau_{i_0} \succ \dots \succ \tau_{i_{N-1}} \succ \sigma(\xi(t_I))$ is complete.

This completes our second proof of Theorem 4.1.1.

4.4 Universality of the construction.

The second proof of Theorem 4.1.1 offered in Section 4.3 shows that from a generic expanding separatrix family we obtain a train track expansion, and moreover the construction clearly produces an expansion for which each train track fully carries the measured foliation. For later purposes we need a converse which says that every train track expansion fully carrying the measured foliation arises, as described in the above construction, from a generic expanding separatrix family.

To formulate this converse, we need some terminology.

In the proof above, when $t \rightarrow \xi(t)$, $j \in J$ is a complete expanding separatrix family, with ordered list of events was denoted $\inf(J) = t_0 < t_1 < \dots$, we chose $s_i \in (t_i, t_{i+1})$, and we obtained a sequence of train tracks $\tau_i = \sigma(\xi(s_i))$. Each

carrying relation $\tau_i \succ \tau_{i+1}$ is either a slide move or an elementary splitting. The train track sequence τ_0, τ_1, \dots will be called the *elementary move sequence* of the family $\xi(t)$.

An *elementary factorization* of a splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ is a sequence

$$\tau'_{0,0} \succ \dots \succ \tau'_{0,n_0} = \tau'_{1,0} \succ \dots \succ \tau'_{1,n_1} = \tau'_{2,0} \succ \dots$$

where each τ'_{ij} is generic, τ_i and $\tau'_{i,0}$ are comb equivalent for each $i \geq 0$, and each subsequence $\tau'_{i,0} \succ \dots \succ \tau'_{i,n_i}$ consists of one elementary splitting and $n_i - 1$ slide moves.

The second construction of train track expansions gives one direction of the following result:

Proposition 4.4.1 (Slicing and Expansion). *If $\tau_0 \succ \tau_1 \succ \dots$ is a complete splitting sequence and \mathcal{F} is a measured foliation, then the following are equivalent:*

- $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion of \mathcal{F} with the property that each τ_i fully carries \mathcal{F} .
- There exists a generic train track τ fully carrying \mathcal{F} , and a foliated tie bundle ν over τ whose horizontal foliation is equivalent to \mathcal{F} in \mathcal{MF} , and there exists a complete, generic, expanding separatrix family $t \rightarrow \xi(t)$, $t \in J$, in ν whose elementary move sequence is an elementary factorization of the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$.

Proof of Proposition 4.4.1. For the yet unproved direction of this proposition, suppose that $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion of \mathcal{F} where each τ_i fully carries \mathcal{F} . Replacing each τ_i by a comb equivalent generic train track, if necessary, we may assume that each τ_i is generic. Then by factoring and reindexing we may assume that each carrying relation $\tau_i \succ \tau_{i+1}$ is either a slide move or an elementary splitting.

Replacing \mathcal{F} by an equivalent partial measured foliation, choose a foliated tie bundle ν over τ_0 whose horizontal foliation is \mathcal{F} . Choose a positive transverse Borel measure on $\mathcal{F}^v(\nu)$ so that we can measure length along leaves of \mathcal{F} .

We start with the completely degenerate separatrix system $\xi(0)$ consisting of the set of singularities of \mathcal{F} , so $\sigma(\xi(0)) = \tau_0$.

Now we proceed by induction. Suppose that we have constructed a generic expanding separatrix family $\xi[0, t_i]$ with $\xi(t_i)$ generic, whose elementary move sequence is τ_0, \dots, τ_i with $\tau_i = \sigma(\xi(t_i))$. Thus, by slicing ν along $\xi(t_i)$ we obtain a foliated tie bundle $\nu(\xi(t_i))$ over τ_i whose horizontal foliation is equivalent in \mathcal{MF} to \mathcal{F} . Let $r: \nu(\xi(t_i)) \rightarrow \nu$ be the partial fulfillment map.

The carrying relation $\tau_i \succ \tau_{i+1}$ is either a slide move along a transition branch b of τ_i , or an elementary splitting along a sink branch b of τ_i . Let x, y be the endpoints

of b ; if b is a transition branch choose x to be the tail end and y to be the head end. Let ξ, η be the cusps of $\nu(\xi(t_i))$ projecting to x, y respectively. Let $R(b)$ be the inverse image of b in $\nu(\xi(t_i))$, so ξ is contained in the interior of a vertical side of $R(b)$. Choose a horizontal leaf segment α with one end on ξ and extending all the way across $R(b)$ and through the opposite vertical side of $R(b)$ by a tiny amount, unless α runs into η on the opposite boundary component in which case it stops there. Slicing $\nu(\xi(t))$ along α produces a foliated tie bundle over a train track τ' , with horizontal foliation equivalent to \mathcal{F} .

We claim that τ' is isotopic to τ_{i+1} . This is obvious in the case where b is a transition branch, for in that case τ' is evidently obtained by a slide move along b . When b is a sink branch, it follows from the construction that $\tau_i \succ \tau'$ is a splitting along b and τ' fully carries \mathcal{F} . However, of the three splittings of τ_i along b —Left, Right, and Central—only one of them produces a train track fully carrying \mathcal{F} , and by hypothesis that one is τ_{i+1} , proving the claim.

The result of slicing $\nu(\xi(t))$ along α is identical to the result of slicing ν along $\xi(t_i) \cup r(\alpha)$. Thus we can extend $\xi[0, t_i]$ to $\xi[0, t_{i+1}]$ with $\xi(t_{i+1}) = \xi(t_i) \cup r(\alpha)$, where $t_{i+1} = t + \text{Length}(\alpha)$, completing the induction.

We have defined the required expanding separatrix family $t \rightarrow \xi(t)$, $t \in J$, but it remains prove completeness.

Let Ξ be the space of separatrix systems in ν , partially ordered by inclusion as discussed in the proof of Theorem 4.1.1. Since the expanding separatrix family $\xi(t)$ is continuous and increasing in Ξ , as t increases it follows that either Ξ has no maximum point and $\xi(t)$ diverges to the positive end of Ξ in which case $\xi(t)$ is complete, or $\xi(t)$ approaches a supremum $\xi(\infty) \in \Xi$ and we must prove that $\xi(\infty)$ equals the maximum point of Ξ .

The only bad case is that $\xi(t)$ converges to $\xi(\infty) \in \Xi$ but $\xi(\infty)$ is not equal to a maximum point of Ξ (this incorporates the case that Ξ has no maximum point and yet $\xi(t)$ does not diverge to the positive end). In this case we shall derive a contradiction.

Since $\xi(\infty)$ is not a maximum of Ξ , the train track $\sigma(\xi(\infty))$ is not an essential curve family. The expanding separatrix family $\xi(t)$, $t \in J$ is either compact with $\xi(\sup(J)) = \xi(\infty)$, or it extends continuously to a compact family achieving its maximum at $\xi(\infty)$. In either case, compactness implies that there are only finitely many splitting events, and so the original splitting sequence is finite. But the last train track in the splitting sequence is either $\xi(\infty)$ or something which splits to $\xi(\infty)$, and in particular it is not an essential curve family. This contradicts the hypothesis that the original splitting sequence is complete. \diamond

5 Convergence of train track expansions

In this section we prove Theorem 5.1.1, which generalizes the fact that the continued fraction expansion of an irrational number converges to that number. The main tools are the Slicing and Expansion Proposition 4.4.1 and the Ascoli-Arzelà Theorem.

5.1 Convergence to an arational measured foliation

Given a measured foliation \mathcal{F} , let $\mathcal{MF}(\mathcal{F})$ denote the space of transverse measures on the underlying singular topological foliation of \mathcal{F} , with the weak* topology, and let $\mathcal{PMF}(\mathcal{F})$ denote the projectivization of $\mathcal{MF}(\mathcal{F})$. There are maps $\mathcal{MF}(\mathcal{F}) \hookrightarrow \mathcal{MF}$ and $\mathcal{PMF}(\mathcal{F}) \hookrightarrow \mathcal{PMF}$ commuting with projectivization, and these maps are continuous and injective: letting τ be any train track carrying \mathcal{F} with carrying injection $\mathcal{F} \subset \nu(\tau)$, the natural map $\mathcal{MF}(\mathcal{F}) \rightarrow W(\tau)$ is evidently continuous and injective; composing with the continuous injection $W(\tau) \rightarrow \mathcal{MF}$ finishes the proof. By general results of measurable dynamics, $\mathcal{PMF}(\mathcal{F})$ is a compact space with the structure of a Choquet simplex, whose extreme points are the ergodic transverse measures on \mathcal{F} , and by finite dimensionality of \mathcal{PMF} it follows that $\mathcal{PMF}(\mathcal{F})$ is a finite dimensional simplex. As remarked in the introduction, the set of measured foliations on S whose projective class lies in $\mathcal{PMF}(\mathcal{F})$ is exactly the unmeasured equivalence class of \mathcal{F} , where $\mathcal{F}, \mathcal{F}'$ are unmeasured equivalent if their equivalence classes have representatives with the same underlying unmeasured singular foliation.

When \mathcal{F} is the stable foliation of a pseudo-Anosov homeomorphism ϕ then \mathcal{F} is uniquely ergodic, because under iteration of ϕ each point of \mathcal{PMF} except the unstable foliation approaches \mathcal{F} , and yet ϕ preserves the set $\mathcal{PMF}(\mathcal{F})$. However, in general an arational measured foliation need not be uniquely ergodic. Examples in the context of interval exchange maps were given in [KN76], and the translation into measured foliation language is now standard (see also [Mas92]).

If \mathcal{F} is a measured foliation and if $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion of \mathcal{F} then evidently the topological foliation underlying \mathcal{F} is carried by each τ_i . In other words, train track expansions are sensitive only to the unmeasured equivalence class of a measured foliation. It follows that $\mathcal{MF}(\mathcal{F}) \subset \cap_i \mathcal{MF}(\tau_i)$ and similarly for \mathcal{PMF} .

To summarize the discussion, the best one can realistically hope for is that $\mathcal{MF}(\mathcal{F}) = \cap_i \mathcal{MF}(\tau_i)$. Our next theorem proves this when \mathcal{F} is arational:

Theorem 5.1.1 (Expansion Convergence Theorem). *Let \mathcal{F} be an arational measured foliation, and let $\tau_0 \succ \tau_1 \succ \dots$ be a train track expansion of \mathcal{F} . Then*

$$\mathcal{MF}(\mathcal{F}) = \cap_i \mathcal{MF}(\tau_i) \quad \text{and} \quad \mathcal{PMF}(\mathcal{F}) = \cap_i \mathcal{PMF}(\tau_i).$$

In the special case that τ_0 canonically carries⁶ \mathcal{F} , expansion convergence is proved in [PP87], Theorem 3.1, whose hypotheses are described at the beginning of Section 2.1 of that paper. Another case of expansion convergence is given in [Pen92] Theorem 3.3.2.

5.2 Proof of the expansion convergence theorem

Suppose that the splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ is an expansion of measured foliations \mathcal{F} and \mathcal{G} , and \mathcal{F} is arational. We shall show that \mathcal{F} and \mathcal{G} are equivalent in \mathcal{MF} to partial measured foliations which have identical underlying topological singular foliations.

We first carry out the proof with an following additional assumption:

Special Case 1: Suppose that each τ_i fully carries both \mathcal{F} and \mathcal{G} . Applying Proposition 4.4.1, we may assume that each τ_i is generic, and we have the following conclusion. Let $\nu = \nu(\tau_0)$ be a foliated tie bundle over τ_0 , whose horizontal foliation $\mathcal{F}^h(\nu)$ is equivalent in \mathcal{MF} to \mathcal{F} , and whose vertical foliation $\mathcal{F}^v(\nu)$ consists of the ties over τ_0 . Then there is a complete, generic, expanding separatrix family $\xi(t)$ in ν whose elementary move sequence is an elementary factorization of the splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$. Altering each τ_i within its comb equivalence class, we thus obtain parameter values $t_0 < t_1 < \cdots$, so that if ν is sliced along $\xi_i = \xi(t_i)$ the result is a foliated tie bundle ν_i over τ_i , whose horizontal foliation $\mathcal{F}^h(\nu_i)$ is again equivalent in \mathcal{MF} to \mathcal{F} . Recall that $\partial\xi_i$ is the set of nonsingular endpoints of ξ_i . We have partial fulfillment maps $r_i: (S, \mathcal{F}^h(\nu_i), \mathcal{F}^v(\nu_i)) \rightarrow (S, \mathcal{F}^h(\nu), \mathcal{F}^v(\nu))$ as well as partial fulfillment maps $r_i^j: (S, \mathcal{F}^h(\nu_j), \mathcal{F}^v(\nu_j)) \rightarrow (S, \mathcal{F}^h(\nu_i), \mathcal{F}^v(\nu_i))$ with the property that $r_i \circ r_i^j = r_j$ whenever $j \geq i$. We also have a tie collapse $f_i: (S, \nu_i) \rightarrow (S, \tau_i)$. Dropping the subscript 0 we write $r^j = r_0^j: (S, \mathcal{F}^h(\nu_j), \mathcal{F}^v(\nu_j)) \rightarrow (S, \mathcal{F}^h(\nu), \mathcal{F}^v(\nu))$. Letting $|dx|$ denote the transverse measure on $\mathcal{F}^v(\nu)$ and $|dy|$ the transverse measure on $\mathcal{F}^h(\nu)$, we obtain a metric $dx^2 + dy^2$ on ν , which is a Euclidean metric with geodesic boundary, except at the singular points where the metric has a cusp.

The inverse image under f_i of the switch set of τ_i is a finite set of ties in ν_i which subdivide ν_i into rectangles, in one-to-one correspondence with the branches of τ_i ; we call this the *rectangle decomposition* of ν_i .

By pushing forward the rectangle decomposition of ν_i we obtain a collection of Euclidean rectangles in ν denoted \mathcal{R}_i . The union of horizontal sides of \mathcal{R}_i , denoted $\partial_h \mathcal{R}_i$, is $\xi_i \cup \partial\nu$. The union of vertical sides of \mathcal{R}_i , denoted $\partial_v \mathcal{R}_i$, is the set of

⁶The statement “ τ canonically carries \mathcal{F} ” is equivalent to the statement in [PP87] that “ τ is suited to \mathcal{F} ”.

maximal vertical segments in ν whose interiors are disjoint from ξ_i but which have nonempty intersection with $\partial\xi_i$; each endpoint of each segment of $\partial\nu\mathcal{R}_i$ is contained in $\xi_i - \partial\xi_i$. The rectangles of \mathcal{R}_i are in one-to-one correspondence with the branches of τ_i ; given a branch $b \in \text{Br}(\tau_i)$ let $R_b \in \mathcal{R}_i$ be the corresponding rectangle. The components of $\partial\nu\mathcal{R}_i$ are in one-to-one correspondence with the switches of τ_i ; given a switch $s \in \text{Sw}(\tau_i)$ let I_s be the corresponding component of $\partial\nu\mathcal{R}_i$. Let $\ell_h(R), \ell_v(R)$ be the horizontal and vertical side lengths of a rectangle $R \in \mathcal{R}_i$.

By arationality of \mathcal{F} , the set $\bigcup_i \xi_i$ is dense in ν . It follows that

$$\max_{R \in \mathcal{R}_i} \ell_v(R) \rightarrow 0 \quad \text{as } i \rightarrow \infty$$

Since the measured foliation \mathcal{G} is carried by τ_i , there exists an invariant weight $w_i \in W(\tau_i)$ corresponding to \mathcal{G} . Consider a vertical segment I_s , $s \in \text{Sw}(\tau_i)$. On one side of I_s there is a rectangle $R_{b_1} \in \mathcal{R}_i$, $b_1 \in \text{Br}(\tau)$, such that I_s is the entire vertical side of R_{b_1} . On the other side of I_s there are two rectangles $R_{b_2}, R_{b_3} \in \mathcal{R}_i$, $b_2, b_3 \in \text{Br}(\tau)$, such that I_s is a concatenation of the form $I_s = I_{s_2} * I_{s_3}$, where I_{s_j} is a vertical side of R_{b_j} . Define a positive Borel measure ν_i on I_s , whose restriction to I_{s_j} is given by

$$\nu \mid I_{s_j} = \frac{w_i(b_j)}{\int_{I_{s_j}} |dy|} |dy|$$

We thus have $\nu_i(I_s) = w_i(b_1)$ for $j = 2, 3$, and $\nu_i(I_{s_j}) = w_i(b_j)$.

For each $b \in \text{Br}(\tau)$ we may now extend ν_i to a measured foliation on the rectangle R_b as follows. On each component of $\partial\nu R_b$ we have constructed a positive Borel measure of total mass $w_i(b)$, which induces a measure preserving homeomorphism f_b between the two components of $\partial\nu R_b$. In the $|dy|$ metric the homeomorphism f_b is piecewise linear. Now construct a foliation of R_b each of whose leaves is a Euclidean geodesic whose endpoints are a pair of points on opposite components of $\partial\nu R_b$ that correspond under the homeomorphism f_b , and extend ν_i to a transverse measure.

We have constructed a partial measured foliation \mathcal{G}_i and a carrying injection $\mathcal{G}_i \subset \nu$. From this construction it is clear that $[\mathcal{G}_i] = [\mathcal{G}]$. For each $i \geq 0$ we construct a homeomorphism $\phi_i: (\nu, \mathcal{G}_0) \rightarrow (\nu, \mathcal{G}_i)$ as follows. Given a vertical fiber I of ν we have $\nu_0(I) = \nu_i(I)$, and the restriction of ϕ_i to I is the unique orientation preserving homeomorphism such that $(\phi_i)_*(\nu_0) = \nu_i$.

For purposes of applying the Ascoli-Arzelà Theorem we must verify that the sequence ϕ_i is equicontinuous. This follows by applying two observations to each vertical fiber I of ν : as $i \rightarrow \infty$ the set $I \cap \xi_i$ determines a subdivision of I whose mesh goes to 0; and for all $j \geq i$ the maps ϕ_j are identical on the set $I \cap \xi_i$.

Applying Ascoli-Arzelà let $\phi_\infty: \nu \rightarrow \nu$ be the limit map. It is a continuous surjection preserving the tie foliation of ν . We show that ϕ_∞ is a homeomorphism,

taking the underlying singular foliation of \mathcal{G} to the underlying singular foliation of \mathcal{F} .

What is already clear from the construction is that each leaf segment of \mathcal{G} is mapped by ϕ_∞ to a horizontal path in ν , that is, a leaf segment of $\mathcal{F}^h(\nu)$. By irrationality of \mathcal{F} , this already implies that \mathcal{G} has no closed leaves. To prove that ϕ_∞ takes \mathcal{G} to \mathcal{F} it remains only to show that ϕ_∞ is one-to-one. Consider a vertical fiber of ν , and arguing by contradiction suppose that there are two points $p \neq q$ in that vertical fiber such that $\phi_\infty(p) = \phi_\infty(q)$. Letting I be the subsegment of the vertical fiber with endpoints p, q , it follows that ϕ_∞ is constant on I . Let $z = \phi_\infty(I)$. Each bi-infinite leaf of \mathcal{G} passing through I maps to a bi-infinite horizontal path in ν passing through z . Since \mathcal{G} has no closed leaves, there are infinitely many distinct bi-infinite leaves of \mathcal{G} passing through I , and these project to infinitely many distinct train paths in τ_0 . It follows that there are infinitely many distinct bi-infinite horizontal paths in ν passing through z , but this is clearly absurd.

This completes the proof of Theorem 5.1.1 under the assumption that each τ_i fully carries both \mathcal{F} and \mathcal{G} .

Special Case 2: Removing the assumption that each τ_i fully carries \mathcal{G} . For each i let τ'_i be the subtrack of τ_i which fully carries \mathcal{G} . We need the following:

Fact 5.2.1. *Given a splitting $\tau \succ \sigma$, a partial measured foliation \mathcal{G} carried by τ and σ , if $\tau' \subset \tau$, $\sigma' \subset \sigma$ are the subtracks fully carrying \mathcal{G} , then $\tau' \succ \sigma'$ is either a splitting or a comb equivalence.*

Proof. It suffices to consider the case where $\tau \succ \sigma$ is a wide splitting of τ along some splitting arc α . Consider Figure 14, whose top portion shows a neighborhood of α . Divide the branch ends incident to α into four sets: upper-left, lower-left, upper-right, or lower-right. Consider which of these four sets has nonempty intersection with τ' . If τ' has nonempty intersection with no more than one of upper-left and lower-left, or if τ' has nonempty intersection with no more than one of upper-right and lower-right, then $\tau' \succ \sigma'$ is a comb equivalence. Otherwise, $\tau' \succ \sigma'$ is a splitting. \diamond

A splitting sequence can have only finitely many central splittings, as shown in Section 4.1. By truncation we may assume that neither of the sequences $\tau_0 \succ \tau_1 \succ \dots$ or $\tau'_0 \succ \tau'_1 \succ \dots$ has any central splittings. It follows that for each splitting $\tau_i \succ \tau_{i+1}$, the carrying relation $\tau'_i \succ \tau'_{i+1}$ is either a parity splitting or a comb equivalence.

We now carry out the construction of the sequence of partial measured foliations \mathcal{G}_i and carrying injections $\mathcal{G}_i \subset \nu$, as above, noting that as a consequence of

the discussion in the previous paragraph, the subsurface $\text{Supp}(\mathcal{G}_i)$ is independent of i ; denote this subsurface $\nu^{\mathcal{G}}$. We can now define the sequence of maps $\phi_i: \nu^{\mathcal{G}} \rightarrow \nu^{\mathcal{G}}$ taking \mathcal{G}_0 to \mathcal{G}_i and apply Ascoli-Arzela as before, obtaining in the limit a homeomorphism $\phi_\infty: \nu^{\mathcal{G}} \rightarrow \nu^{\mathcal{G}}$. Arguing as before, ϕ_∞ takes each bi-infinite leaf of \mathcal{G} to a bi-infinite horizontal path in ν , but each of the latter is dense in ν since \mathcal{F} is arational, and it follows that $\text{image}(\phi_\infty) = \nu$ and so $\nu = \nu^{\mathcal{G}}$. The argument is now completed as before.

The general case: removing the assumption that each τ_i fully carries \mathcal{F} . Let σ_i be the subtrack of τ_i which fully carries \mathcal{F} and let τ'_i be the subtrack which fully carries \mathcal{G} . Truncate so that none of the splitting sequences (τ_i) , (σ_i) , (τ'_i) has central splittings. We obtain a sequence of carrying maps $\sigma_0 \succcurlyeq \sigma_1 \succcurlyeq \cdots$ consisting of comb equivalences and splittings, such that σ_i carries the arational measured foliation \mathcal{F} , and hence each σ_i fills.

We claim that infinitely many of the carryings $\sigma_i \succcurlyeq \sigma_{i+1}$ are splittings, in fact, there is a constant A depending only on the topology of S such that for each i the finite sequence $\sigma_i \succcurlyeq \sigma_{i+1} \succcurlyeq \cdots \succcurlyeq \sigma_{i+A}$ contains at least one splitting. The claim follows because σ_i is a filling subtrack of τ_i , and there is a bound on the length of a splitting sequence starting at τ_i which does not split the subtrack σ_i .

By ignoring the comb equivalences in the sequence $\sigma_0 \succcurlyeq \sigma_1 \succcurlyeq \cdots$ we may therefore regard it as a complete splitting sequence, which is therefore a train track expansion of \mathcal{F} . As in the proof under the special assumption above, we may apply Proposition 4.4.1 to the tie bundle $\nu = \nu(\sigma_0)$ as above, equipped with a Euclidean metric on ν whose horizontal measured foliation is equivalent to \mathcal{F} , and the conclusion is that there is a complete, generic, expanding separatrix family $\xi(t)$ in ν whose elementary move sequence is an elementary factorization of the expansion $\sigma_0 \succcurlyeq \sigma_1 \succcurlyeq \cdots$. As above we obtain rectangle decompositions \mathcal{R}_i of ν , with horizontal boundaries $\partial_h \mathcal{R}_i = \xi_i \cup \partial\nu$ and vertical boundaries $\partial_v \mathcal{R}_i$, so that the vertical widths of the rectangles approach zero.

Since \mathcal{G} itself is not carried by σ_0 , we must be a little careful at this stage.

Consider a tie bundle $\nu(\tau_0)$, chosen so that we have a commutative diagram

$$\begin{array}{ccc} \nu = \nu(\sigma_0) & \xrightarrow{\subset} & \nu(\tau_0) \\ \pi' \downarrow & & \downarrow \pi \\ \sigma_0 & \xrightarrow{\subset} & \tau_0 \end{array}$$

where π, π' are tie collapsing maps. Note that for each vertical fiber I of $\nu(\tau_0)$, the intersection $I \cap \nu$ either is empty or is an entire vertical fiber of $\nu(\sigma_0)$, depending on whether $\pi(I) \in \sigma_0$.

The map $\pi: (S, \nu(\tau_0)) \rightarrow (S, \tau_0)$ may be factored

$$(S, \nu(\tau_0)) \xrightarrow{\psi} (S, \hat{\nu}) = (S, \nu(\sigma_0) \cup \hat{\tau}) \xrightarrow{\pi'} (S, \tau_0)$$

as follows. The map π' is the tie bundle of $\nu = \nu(\sigma_0)$, extended to S by a homeomorphism of $S - \nu = S - \nu(\sigma_0)$ to $S - \sigma_0$. The set $\hat{\tau}$ is the closure of $\pi'^{-1}(\text{Cl}(\tau_0 - \sigma_0))$; note that $\text{Cl}(\tau_0 - \sigma_0)$ is a train track with terminals, having one stop for each point $s \in \sigma_0$ which is a switch of τ_0 but not a switch of σ_0 . Note that π' restricts to a homeomorphism between $\hat{\tau}$ and $\text{Cl}(\tau_0 - \sigma_0)$. The map ψ has the following effect on each vertical fiber I of $\nu(\tau_0)$: if $I \cap \nu \neq \emptyset$ then ψ takes I homeomorphically to $I \cap \nu$; otherwise $\psi(I)$ is the point of $\hat{\tau}$ which is taken to $\pi(I)$ by π' . What we have done is to partially collapse $\nu(\tau_0)$, in effect gluing the train track with terminals $\hat{\tau}$ to the tie bundle $\nu(\sigma_0)$. The gluing is done smoothly, so that a smooth path in S which is a concatenation of paths transverse to vertical fibers of $\nu(\sigma_0)$ and train paths in $\hat{\tau}$ projects to a train path of τ_0 . Each “vertical fiber” of $\hat{\nu}$ is either a vertical fiber of ν or a point of $\hat{\tau}$.

We also have a tie bundle $\nu(\tau'_0)$ over τ'_0 . Choose a weight function w on τ_0 corresponding to \mathcal{G} , nonzero precisely on the branches of τ'_0 , and use this to construct a partial measured foliation supported on $\nu(\tau'_0)$ equivalent to \mathcal{G} in \mathcal{MF} , still denoted \mathcal{G} . Now we proceed as earlier, constructing Borel measures ν_i on vertical fibers of $\hat{\nu}$, by using \mathcal{G} -heights of the rectangles of \mathcal{R}_i for the vertical fibers of ν , and using \mathcal{G} -atomic measures for the points of $\hat{\tau}$. Each transverse measure ν_i may be used to construct a continuous map $\phi_i: \mathcal{G} \rightarrow \hat{\nu}$ in the same manner as before, except that bands of leaves will be squeezed together in $\hat{\tau}$, so by all appearances the map ϕ_i may not be one-to-one. Nevertheless, equicontinuity still follows by the same argument and so we obtain a limit map ϕ_∞ , still perhaps not one-to-one. The image of a leaf segment of \mathcal{G} is now a horizontal path in $\hat{\nu}$, which means a bi-infinite smooth concatenation of horizontal paths in ν and train paths in $\hat{\tau}$. The same argument as before therefore shows that if ϕ_∞ is not one-to-one then infinitely many distinct leaves of \mathcal{G} map to infinitely many distinct bi-infinite horizontal paths through a single point of $\hat{\nu}$. However, since σ_0 fills the surface, it follows that each point of $\hat{\nu}$ lies on only finitely many bi-infinite horizontal paths: there are only finitely many maximal smooth paths in $\hat{\nu}$ which are disjoint from the interior of ν , and once a horizontal path enters the interior of ν at a boundary singularity, it never again leaves the interior, by a-rationality of \mathcal{F} . Thus we obtain the same contradiction as before, thereby proving that ϕ_∞ is one-to-one. It follows that image of ϕ_∞ is contained in ν , for if some point $x \in \hat{\tau} = \hat{\nu} - \nu$ were in the image, then infinitely many distinct leaves of \mathcal{G} would have image passing through x . It therefore follows as before that ϕ_∞ takes \mathcal{G} to \mathcal{F} .

This completes the proof of the Expansion Convergence Theorem.

Remark. The Expansion Convergence Theorem can fail without assuming \mathcal{F} is arational. For instance, suppose that \mathcal{F} has two arational foliation components. Then we can choose a generic expanding separatrix family which becomes dense in one foliation component but not in the other. The corresponding train track expansion thus converges in one of the arational foliation components but not the other. This example accords with the remark in the introduction that our definition of train track expansions works well for arational measured foliations but not otherwise.

Remark. The Expansion Convergence Theorem brings up the following question: if the infinite splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ is not an expansion of an arational measured foliation, what exactly is it an expansion of? In other words, what is the set $\cap_{i=0}^{\infty} \mathcal{MF}(\tau_i)$? We will answer this question later in Theorem 8.5.1. The following generalization of the Expansion Convergence theorem provides some clues, and will play an important role in the proof of Theorem 8.5.1.

5.3 Convergence to a partial arational measured foliation

For later purposes we shall need a generalization of the Expansion Convergence Theorem which describes expansions of partial arational measured foliations.

Suppose that \mathcal{F}, \mathcal{G} are partial measured foliations, \mathcal{F} is partial arational, and $\tau_0 \succ \tau_1 \succ \cdots$ is an expansion of each. We would like to say how \mathcal{F}, \mathcal{G} are related. We cannot conclude in general that $\mathcal{G} \in \mathcal{MF}(\mathcal{F})$ —perhaps \mathcal{F}, \mathcal{G} are distinct foliation components of some larger measured foliation of which $\tau_0 \succ \tau_1 \succ \cdots$ is an expansion. A weaker but still useful conclusion would be that $\langle [\mathcal{F}], [\mathcal{G}] \rangle = 0$, but even this may not hold: perhaps \mathcal{F}, \mathcal{G} are both carried on a subtrack $\sigma_0 \subset \tau_0$ which is never split during the entire splitting sequence, in which case it follows that $\mathcal{MF}(\sigma_0) \subset \cap_i \mathcal{MF}(\tau_i)$, and it might be possible to choose \mathcal{F}, \mathcal{G} whose classes lie in $\mathcal{MF}(\sigma_0)$ so that $\langle [\mathcal{F}], [\mathcal{G}] \rangle \neq 0$. Thus, we need an additional hypothesis which rules out the existence of such a subtrack. Let $\sigma_i \subset \tau_i$ be the subtrack fully carrying \mathcal{F} , and so we obtain a sequence of carryings $\sigma_0 \succ \sigma_1 \succ \cdots$ each of which is either a comb equivalence or a splitting. If this sequence contains infinitely many splittings then we say that \mathcal{F} is *infinitely split* by the expansion $\tau_0 \succ \tau_1 \succ \cdots$.

Theorem 5.3.1. *If \mathcal{F} is a partial arational measured foliation, and if $\tau_0 \succ \tau_1 \succ \cdots$ is an expansion of \mathcal{F} which splits \mathcal{F} infinitely, then for every $\mathcal{G} \in \cap_n \mathcal{MF}(\tau_n)$ we have $\langle [\mathcal{F}], [\mathcal{G}] \rangle = 0$.*

The proof of the theorem will give very explicit information about how the carrying maps $\mathcal{F} \rightarrow \tau_i, \mathcal{G} \rightarrow \tau_i$ are related. This information will be useful later, so we will describe it in detail in Proposition 5.3.2.

Proof. By working on one foliation component of \mathcal{G} at a time, it suffices to assume that \mathcal{G} has a unique foliation component.

We follow the proof of the general case in the Expansion Convergence Theorem: let $\sigma_0 \subset \tau_0$ be the subtrack fully carrying \mathcal{F} , let ν be a tie bundle over σ_0 , and by altering \mathcal{F} within its equivalence class we assume that we have a carrying bijection $\mathcal{F} \xrightarrow{c} \nu$, so vertical segments of ν are ties and horizontal segments of ν are leaf segments of \mathcal{F} . The sequence of carrying maps $\sigma_0 \succ \sigma_1 \succ \cdots$ consists of comb equivalences and infinitely many splittings, and so we can collapse the comb equivalences to obtain a complete splitting sequence which is therefore an expansion of \mathcal{F} , each train track of which fully carries \mathcal{F} . Letting $\hat{\nu}$ be obtained from ν by smoothly attaching the train track with terminals $\hat{\tau}_0 = \text{Cl}(\tau_0 - \sigma_0)$, by the same proof as above we alter \mathcal{G} within its equivalence class and obtain a map $\phi_\infty: \mathcal{G} \rightarrow \hat{\nu}$ so that each leaf of \mathcal{G} maps to a horizontal path in $\hat{\nu}$, which must be a smooth concatenation of train paths in $\hat{\tau}_0$ and leaf segments of \mathcal{F} . Note that we no longer have available the argument that each point of $\hat{\nu}$ lies on only finitely many bi-infinite horizontal paths.

Let Σ be the union of saddle connections of \mathcal{F} . We consider two cases, depending on whether $\phi_\infty(\mathcal{G}) \cap \text{Supp}(\mathcal{F})$ is a subset of Σ .

Case 1: $\phi_\infty(\mathcal{G}) \cap \text{Supp}(\mathcal{F}) \subset \Sigma$. Slicing \mathcal{F} open along Σ to obtain a canonical model \mathcal{F}' , it follows that \mathcal{G} can be isotoped into the complement of $\text{Supp}(\mathcal{F}')$, proving that $\langle [\mathcal{F}], [\mathcal{G}] \rangle = 0$.

Case 2: $\phi_\infty(\mathcal{G}) \cap \text{Supp}(\mathcal{F}) \not\subset \Sigma$. It follows that some half-infinite leaf segment ℓ of \mathcal{G} is mapped by ϕ_∞ to a half-infinite leaf segment of \mathcal{F} . This implies that \mathcal{G} is not annular, and so \mathcal{G} is partial arational. Also, since ℓ is dense in $\text{Supp}(\mathcal{G})$ and $\phi_\infty(\ell)$ is dense in $\nu = \text{Supp}(\mathcal{F})$ it also follows that $\phi_\infty(\mathcal{G}) = \nu$.

It remains to show that ϕ_∞ is one-to-one. Unfortunately, unlike all cases of Theorem 5.1.1, it need not be true that each point of ν lies on only finitely many bi-infinite horizontal paths, and so we cannot reach the same contradiction as before that a finite cardinal equals an infinite cardinal. The problem is that by concatenating saddle connections we may be able to obtain infinitely many bi-infinite horizontal paths in ν . However, each point $x \in \nu - \Sigma$ lies on at most countable many horizontal paths in ν , projecting to at most countably many bi-infinite train paths in σ_0 . If ϕ_∞ is not one-to-one on \mathcal{G} then there exists a transversal α of \mathcal{G} such that $\phi_\infty(\alpha)$ is a single point $x \in \nu$. The set of bi-infinite leaves of \mathcal{G} passing through points of α is uncountable, since \mathcal{G} is partial arational, and no two of them determine the same bi-infinite train path of σ_0 . Thus, we obtain a similar contradiction, that a countable cardinal equals an uncountable cardinal. \diamond

Remark. Unlike in the case where \mathcal{F} is arational, when \mathcal{F} is only partial arational it is possible that $\phi_\infty(\mathcal{G}) \subset \text{Supp}(\mathcal{F})$ and yet $\langle [\mathcal{F}], [\mathcal{G}] \rangle = 0$. To see how this can happen, consider a carrying bijection $\mathcal{F}_i \xleftrightarrow{\nu} \nu(\sigma_i)$ as in the above proof, where \mathcal{F}_i is equivalent to \mathcal{F} . We may assume that for sufficiently large i , say $i \geq I$, the partial measured foliations \mathcal{F}_i are stable, meaning that no proper saddle connections are sliced for $i \geq I$, and hence all the \mathcal{F}_i are isotopic for $i \geq I$. Let Σ_i be the union of all saddle connections of \mathcal{F}_i , proper and not proper. The set Σ_i may be interpreted as a train track with terminals, and as such it may have a true train track as a subset; let σ_i be the maximal such train track. The stability assumption implies that the isotopy type of σ_i is constant for $i \geq I$. Any measured foliation carried by \mathcal{G} has the property that $\tau_0 \succ \tau_1 \succ \dots$ is an expansion of \mathcal{G} , and also that $\langle [\mathcal{F}], [\mathcal{G}] \rangle$, and moreover if the above argument is applied to \mathcal{G} then $\phi_\infty(\mathcal{G}) \subset \Sigma_i$.

The proof of Theorem 5.3.1 gives explicit information about how the carrying maps $\mathcal{F} \rightarrow \tau_i$ and $\mathcal{G} \rightarrow \tau_i$ are related, which we will store in the next proposition for later consumption.

Consider a train track τ carrying a partial arational measured foliation \mathcal{F} . We assume that we have a carrying injection $\mathcal{F} \hookrightarrow \nu(\tau)$. Let $\sigma \subset \tau$ be the subtrack fully carrying \mathcal{F} , and we may assume that we have a carrying bijection $\mathcal{F} \xleftrightarrow{\nu} \nu(\sigma)$, and a commutative diagram

$$\begin{array}{ccc} \nu(\sigma) & \xrightarrow{\subset} & \nu(\tau) \\ \pi' \downarrow & & \downarrow \pi \\ \sigma & \xrightarrow{\subset} & \tau \end{array}$$

where the vertical arrows are tie bundle projections.

Now slice \mathcal{F} along its proper saddle connections, to obtain a canonical model \mathcal{F}' with (possibly nonsurjective) carrying injection $\mathcal{F}' \hookrightarrow \nu(\sigma)$. By restricting the tie bundle to $\text{Supp}(\mathcal{F}')$ we obtain a tie bundle $\nu(\sigma')$ of a train track σ' that canonically carries \mathcal{F}' . The injection $\mathcal{F}' \hookrightarrow \nu(\sigma)$ is far from unique, however, and this implies that σ' is not unique. But by choosing the carrying injection $\mathcal{F}' \hookrightarrow \nu(\tau)$ carefully we get a natural choice of σ' . Namely, start with a carrying bijection $\mathcal{F} \xleftrightarrow{\nu} \nu(\sigma)$, and then slice along the proper saddle connections of \mathcal{F} to get \mathcal{F}' , taking care that $\text{Supp}(\mathcal{F}')$ contains a neighborhood in $\nu(\sigma)$ of any cusp of $\nu(\sigma)$ which is *not* the endpoint of a proper saddle connection of \mathcal{F} . We refer to this construction by saying that σ' is obtained from τ by *canonical splitting compatible with a canonical model of \mathcal{F}* .

Next, choose a carrying injection $\mathcal{F}' \hookrightarrow \nu(\tau)$ whose image is contained in the interior of $\nu(\tau)$. By restricting the tie bundle to $\nu(\tau) - \text{int}(\text{Supp}(\mathcal{F}'))$ we obtain a tie bundle $\nu(\tau')$ of a train track with terminals τ' . In this case τ' is uniquely determined

up to isotopy, independent of the carrying injection \mathcal{F}' ; this is a consequence of Proposition 3.6.5. We say that τ' is obtained from τ by *cutting out* a canonical model of \mathcal{F} . Note that τ' is *never* a train track: it must have one terminal for each boundary singularity of \mathcal{F}' . Thus, the maximal pretrack τ'' contained in τ' is a proper subset of τ' . But at least τ'' is a true train track on S ; this follows because each component of the $\mathcal{C}(S - \tau'')$ is a union of components of $\mathcal{C}(S - \tau)$ and hence has negative index. (See Section 8.3 for a full treatment of “cutting out one train track from another”).

Note that the two processes we have described—cutting out a canonical model, and canonical splitting compatible with a canonical model—yield what are in some sense complementary tracks σ', τ'' , in the sense that each component of $\mathcal{C}(S - (\sigma' \cup \tau''))$ that touches both σ' and τ'' is a topological annulus with one boundary component on σ' and the other on τ'' .

In case (1) of the proof of Theorem 5.3.1, it is clear that if \mathcal{F}' is a canonical model obtained by slicing \mathcal{F} along the union of proper saddle connections Σ , then one can choose \mathcal{G} in its equivalence class so that there are carrying injections $\mathcal{F}' \hookrightarrow \nu(\tau_0)$, $\mathcal{G} \hookrightarrow \nu(\tau_0)$ which have disjoint images. It follows that \mathcal{G} is carried by the track with terminals τ'_0 that is obtained from τ_0 by cutting out \mathcal{F}' , and so \mathcal{G} is carried by the maximal train track τ''_0 contained in τ'_0 . The same argument holds for each $i \geq 0$; just apply the whole argument to the truncated splitting sequence to $\tau_i \succ \tau_{i+1} \succ \dots$.

We record this as follows. Given pairwise disjoint essential subsurfaces $H_1, \dots, H_K \subset S$ and sets of measured foliations $F_1, \dots, F_K \subset \mathcal{MF}(S)$ such that each $\mathcal{F} \in F_k$ has support contained in H_k , the *join* $F_1 * \dots * F_K$ is defined to be the subset of \mathcal{MF} represented by all measured foliations of the form $\mathcal{F}_1 \cup \dots \cup \mathcal{F}_K$ where each \mathcal{F}_k is either empty or is an element of F_k . For example, if a train track τ is a union of pairwise disjoint subtracks τ_1, \dots, τ_n then $\mathcal{MF}(\tau) = \mathcal{MF}(\tau_1) * \dots * \mathcal{MF}(\tau_n)$.

Proposition 5.3.2. *Consider partial measured foliations \mathcal{F}, \mathcal{G} such that \mathcal{F} is partial arational and \mathcal{G} has one foliation component, and let $\tau_0 \succ \tau_1 \succ \dots$ be a train track expansion of both \mathcal{F} and \mathcal{G} which splits \mathcal{F} infinitely. For each i let σ'_i be obtained from τ_i by canonical splitting compatible with a canonical model of \mathcal{F} , let τ'_i be obtained from τ_i by cutting out a canonical model of \mathcal{F} , and let τ''_i be the maximal train track in τ'_i . Then either $\mathcal{G} \in \mathcal{MF}(\tau''_i)$, or $\mathcal{G} \in \mathcal{MF}(\mathcal{F})$. It follows that*

$$\cap_i \mathcal{MF}(\tau_i) = \mathcal{MF}(\mathcal{F}) * \cap_i \mathcal{MF}(\tau''_i)$$

◇

6 Arational Expansion Theorem

TO DO:

- Introduce one cusp train tracks. Introduce cusp jumps. Explain how this breaks the finite order symmetry. Put an extra restriction that there is only one cusp jump, and that at the very end.
- Describe graphs $\tilde{\Gamma}$, $\tilde{\Gamma}^d$, \mathcal{MCG} action, Γ , Γ^d .
- Give the FDA point of view. Arationality, Canonical arationality, can both be checked by an FDA.
- Describe examples of Keane from [Kea77], and prove their arationality.
 - These examples are *not* one sink expansions: he switches the transverse orientation of the sink at the end of each block.
 - The block he describes does kill all the branches needed for a sink expansion (see my diagrams on the left edge of p. 190 of the xerox of [Kea77]).
 - So, do the same principles of arationality work when you do that switcheroo for a one sink expansion?

The Arational Expansion Theorem 6.3.2 will give a criterion which is necessary and sufficient for a given splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ to be a train track expansion of some arational measured foliation. This criterion, called the “iterated rational killing criterion”, is couched entirely in terms of the combinatorial structure of the splitting sequence. What it says, intuitively, is that for any $i \geq 0$ there exists $j > i$ such that any structure in τ_i which could account for a nonarational measured foliation must be killed in τ_j .

Since the set of measured foliations $\cap_i \mathcal{MF}(\tau_i)$ does not change when we truncate a splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$, we may assume that all splittings are parity splittings, and hence all carrying relations $\tau_i \succ \tau_j$ are homotopic carryings; we will make this assumption henceforth in our discussion of the iterated killing criterion.

Recall from Proposition 3.14.1 that a homotopic carrying relation $\tau \succ \tau'$ is either a slide equivalence or it factors as a sequence of parity splittings. Given a homotopic carrying relation $\tau \succ \tau'$, we will define what it means for this pair to satisfy the rational killing criterion—intuitively, any structure in τ which could account for a nonarational measured foliation is killed in τ' . The rational killing criterion is stable in the sense that if $\tau \succ \tau' \succ \tau'' \succ \tau'''$ then the rational killing criterion for $\tau' \succ \tau''$ implies the rational killing criterion for $\tau \succ \tau'''$.

Given a parity splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$, we will define the iterated rational killing criterion by requiring that for each m there exists $n > m$ such that $\tau_m \succ \tau_n$ satisfies the rational killing criterion.

As explained in the introduction, in the case that S is a punctured torus the rational killing criterion for a finite splitting sequence $\tau_i \succ \cdots \succ \tau_j$ can be easily stated: each of the two simple curves in τ_i is killed in τ_j ; equivalently, each of the parities R and L occurs at least once in the splitting sequence $\tau_i \succ \cdots \succ \tau_j$.

To formulate the rational killing criterion in general for a carrying relation $\tau \succ \tau'$, our task is to identify combinatorial structures in τ which should be killed in τ' . In this task we are guided by the following intuition: the rational killing criterion means that one cannot easily detect the presence of a nonarational measured foliation—various nonfilling train tracks associated to τ , which carry only nonarational measured foliations, are “killed” before they can get to τ' . On the other hand, the number of objects to be killed should be kept low, for purposes of computational efficiency—on the torus, we need only keep track of two objects at a time, namely the two simple closed curves in τ_i .

The rational killing criterion for $\tau \succ \tau'$ is a conjunction of two criteria, the “subtrack killing criterion”, and the “splitting cycle killing criterion”. We describe these two criteria briefly here; full discussion is given in Sections 6.1 and 6.2 below.

To describe the subtrack killing criterion for a carrying pair $\tau \succ \tau'$, consider a subtrack $\sigma \subset \tau$. We say that σ *survives* in τ' if there exists a subtrack $\sigma' \subset \tau'$ such that σ homotopically carries σ' ; this σ' , if it exists, is unique, and is called the *descendant* of σ in τ' . If σ does not survive in τ' then we say that σ is *killed* in τ' . We say that $\tau \succ \tau'$ satisfies the *subtrack killing criterion* if every nonfilling subtrack of τ is killed in τ' . Failure of the subtrack killing criterion gives easily detectable nonarational measured foliations carried on both τ and τ' , namely, anything carried on σ' .

For now we shall roughly describe the splitting cycle killing criterion in the special case of a nonpunctured surface; the full definition incorporating punctured surfaces is given in Section 6.2. Given a filling train track σ a *splitting cycle* of σ is a smooth, simple closed curve γ which is a union of C-splitting arcs of σ and arcs properly embedded in $\mathcal{C}(S - \sigma)$ with endpoints at cusps, with a few other technical restrictions on γ that are not so important for now. Consider a filling subtrack $\sigma \subset \tau$ which survives in τ' with descendant $\sigma' \subset \tau'$. We say that a splitting cycle γ of σ *survives* in σ' if there exists a splitting cycle γ' of σ' such that γ, γ' are isotopic in S . This γ' , if it exists, is unique up to isotopy rel σ' , and is called the *descendant* of γ in σ' . If γ does not survive in σ' then we say that γ is *killed* in τ' . We say that $\tau \succ \tau'$ satisfies the *splitting cycle killing criterion* if for each filling subtrack $\sigma \subset \tau$ which survives with descendant $\sigma' \subset \tau'$, each splitting cycle of σ is killed in σ' . Failure of the splitting cycle killing criterion gives easily detectable

nonarational measured foliations carried on both τ and τ' : let ρ be the nonfilling train track obtained from σ by splitting along γ , and let ρ' be the nonfilling train track obtained from σ' by splitting along γ' . Note that $\tau \succcurlyeq \rho$ and $\tau' \succcurlyeq \rho'$, and also $\rho \succcurlyeq \rho'$, hence anything carried by ρ' is a nonarational measured foliation carried by both τ and τ' .

Here, then, is a restatement of the rational killing criterion for a homotopic carrying relation $\tau \overset{\text{H}}{\succcurlyeq} \tau'$, at least in the nonpunctured case: for each pair of subtracks $\sigma \subset \tau$, $\sigma' \subset \tau'$, if $\sigma \overset{\text{H}}{\succcurlyeq} \sigma'$ then both σ, σ' fill, and each splitting cycle of σ is killed in σ' .

The arationality theorem, or rather its proof, will show that if the splitting sequence $\tau_0 \succ \tau_1 \succ \tau_2 \succ \cdots$ is an expansion of a nonarational measured foliation \mathcal{F} then for some m and all $n \geq m$ the rational killing criterion fails for $\tau_m \overset{\text{H}}{\succcurlyeq} \tau_n$, because one can use \mathcal{F} to detect either a nonfilling subtrack or a splitting cycle of a filling subtrack of τ_m which survives in τ_n . Conversely, if there exists m such that for all $n \geq m$ the pair $\tau_m \overset{\text{H}}{\succcurlyeq} \tau_n$ fails the rational killing criterion, then one can piece together enough evidence—either nonfilling subtracks of the τ_i or splittings cycles of filling subtracks of the τ_i —to pinpoint a nonarational measured foliation of which $\tau_0 \succ \tau_1 \succ \tau_2 \succ \cdots$ is an expansion.

Sections 6.1 and 6.2 contain a complete discussion of survival of subtracks and splitting cycles. It follows that the rational killing criterion for a carrying pair $\tau \overset{\text{H}}{\succcurlyeq} \tau'$ depends in a very simple way on the combinatorial structure of this pair—indeed, it can be checked very quickly by the simplest type of computing machine, in the sense that for a given train track τ the set consisting of splitting sequences $\tau = \tau_0 \succ \cdots \succ \tau_n = \tau'$ such that $\tau \overset{\text{H}}{\succcurlyeq} \tau'$ satisfies the rational killing criterion forms a regular language in some sense. This fact will allow us to conclude that the language of train track expansions of arational measured foliations has nice language theoretic properties (Section ??).

The Arational Expansion Theorem 6.3.2 is stated in full in Section 6.3. Section 6.4 contains an application to criteria for canonical expansions of measured foliations. The proof is covered in Sections 6.5–6.8.

Section 9 contains further applications, namely recipes for construction of pseudo-Anosov homeomorphisms. In later sections this information, combined with stable equivalence, will be used to classify pseudo-Anosov mapping classes up to conjugacy.

6.1 Survival of subtracks

Recall from Section 3.14 that a carrying relation $\sigma \overset{\text{H}}{\succcurlyeq} \sigma'$ is said to be homotopic, denoted $\sigma \overset{\text{H}}{\succcurlyeq} \sigma'$, if the carrying map $\sigma' \rightarrow \sigma$ is a homotopy equivalence. Recall also that $\sigma \overset{\text{H}}{\succcurlyeq} \sigma'$ implies $\sigma \overset{\text{F}}{\succcurlyeq} \sigma'$; moreover, assuming that $\sigma \overset{\text{F}}{\succcurlyeq} \sigma'$, it follows that $\sigma \overset{\text{H}}{\succcurlyeq} \sigma'$

where is this section?

if and only if $i(\sigma) = i(\sigma')$ which occurs if and only if σ, σ' are comb equivalent or there exists a sequence of parity splittings from σ to σ' .

Consider a carrying relation $\tau \succcurlyeq \tau'$ and subtrack $\sigma \subset \tau$. Given a subtrack $\sigma' \subset \tau'$, Corollary 3.5.3 shows that $\sigma \overset{\text{H}}{\succcurlyeq} \sigma'$ if and only if the carrying map $\tau' \rightarrow \tau$ takes σ' onto σ by a homotopy equivalence. If such a σ' exists then we say that σ *survives* in τ' , that σ' is a *descendant* of σ in τ' , and that σ is the *ancestor* of σ' in τ . If σ does not survive in τ then σ is *killed* in τ .

The characterization of survival for comb equivalent train tracks is particularly simple:

Proposition 6.1.1. *Given comb equivalent train tracks $\tau \succcurlyeq \tau'$, each subtrack of τ survives in τ' , and the descent relation is a bijection between the set of subtracks of τ and the set of subtracks of τ' .*

Proof. For a single comb move this is obvious, and the general case follows by induction on the number of comb moves between τ and τ' . \diamond

Next we characterize survival for an elementary parity splitting $\tau \succ \tau'$ of parity $d \in \{\text{L}, \text{R}\}$ along a sink branch $b \subset \tau$. Let b have regular neighborhood N and let $\{r, s, t, u\} = \tau \cap \partial N$ as depicted in Figure 13, so in particular r, s are contained in one component of $(\tau \cap N) - \text{int}(b)$, t, u are contained in the other component, and with respect to the (counterclockwise) boundary orientation on ∂N the points are listed in circular order as r, s, t, u . Given a subtrack $\sigma \subset \tau$, we say that σ contains a *Left crossing of b* if $\{s, u\} \subset \sigma$, and σ contains a *Right crossing of b* if $\{r, t\} \subset \sigma$. These notions are independent of the choice of N and the labelling of the points of $\tau \cap \partial N$.

Proposition 6.1.2. *Let $\tau \succ \tau'$ be an elementary splitting of parity $d \in \{\text{L}, \text{R}\}$ along a sink branch b , let \bar{d} be the complement of d in $\{\text{L}, \text{R}\}$, let N be a neighborhood of b as depicted in Figure 13, and let σ be a subtrack of τ . The following are equivalent:*

- *The subtrack σ survives in τ' .*
- *There exists a subtrack $\sigma' \subset \tau'$ such that $\sigma \overset{\text{F}}{\succcurlyeq} \sigma'$.*
- *Either σ does not contain a \bar{d} -crossing of b or σ contains a d -crossing of b .*

Furthermore, if σ survives, then its descendant $\sigma' \subset \tau'$ is the unique maximal subtrack of τ' such that $\sigma \overset{\text{F}}{\succcurlyeq} \sigma'$; equivalently, σ' is the unique subtrack such that $\sigma' - N = \sigma - N$, and σ' contains the post-splitting arc of τ' if and only if σ contains a d -crossing of b . Moreover, σ is isotopic to σ' unless $\sigma \cap N = N$, in which case $\sigma \succ \sigma'$ is a d -splitting along b .

Finally, descent is a bijection between a subset of the set of subtracks of τ and a subset of the set of subtracks of τ' .

Proof. Assume that $\tau \succ \tau'$ has parity $d = L$, and so we must prove that σ is killed if and only if $\{r, t\} \subset \sigma$ and $\{s, u\} \not\subset \sigma$.

If $\{r, t\} \subset \sigma$ and $\{s, u\} \not\subset \sigma$ then $\sigma \cap \partial N = \{r, t\}$ or $\{r, s, t\}$ or $\{r, t, u\}$, and in each case one can easily see that there does not even exist a subtrack $\sigma' \subset \tau'$ such that $\sigma' \cap \partial N = \sigma \cap \partial N$. It follows that there is no subtrack σ' such that $\sigma \stackrel{F}{\succ} \sigma'$, and so there is no σ' such that $\sigma \stackrel{H}{\succ} \sigma'$.

Conversely, suppose that either $\{r, t\} \not\subset \sigma$ or $\{s, u\} \subset \sigma$, and so $\sigma \cap \{r, s, t, u\} = \emptyset$ or $\{r, u\}$ or $\{s, t\}$ or $\{s, u\}$ or $\{r, s, u\}$ or $\{s, t, u\}$ or $\{r, s, t, u\}$. In each case there clearly does exist a subtrack $\sigma' \subset \tau'$ such that $\sigma' \cap \partial N = \sigma \cap \partial N$, and moreover one can choose σ' so that $\sigma' - N = \sigma - N$, which guarantees that $\sigma \stackrel{F}{\succ} \sigma'$. In addition, σ' is uniquely determined if we require that σ be maximal with respect to these conditions, or equivalently, that σ' contains the post-splitting branch if and only if $\{s, u\} \subset \sigma \cap \partial N$. With this choice of σ' , it is easy to check case by case that $\sigma \stackrel{H}{\succ} \sigma'$, and that σ is isotopic to σ' unless $\{r, s, t, u\} \subset \sigma$ in which case $\sigma \succ \sigma'$ is a Left splitting.

For the final conclusion, assuming σ' is a descendant of σ , it suffices to notice from the above discussion that σ determines σ' ; conversely, σ' determines σ by the requirement that $\sigma \stackrel{F}{\succ} \sigma'$. \diamond

Remark. Observe in the above proof that if σ survives and $\sigma \stackrel{F}{\succ} \sigma'$ but σ' is *not* the descendant of σ , then the only way this can happen is if $\{r, s, t, u\} \subset \sigma$ and σ' does not contain the post-splitting arc; this implies that $\sigma \succ \sigma'$ is a central splitting.

Since any splitting relation $\tau \succ \tau'$ can be factored as a comb equivalence followed by an elementary splitting followed by a comb equivalence, by combining Propositions 6.1.1 and 6.1.2 we obtain a description of survival for subtracks of τ . Nevertheless, in light of Proposition 3.13.6 which says that any splitting sequence can be replaced by a wide splitting sequence, it is convenient to get an even more explicit description of survival for a wide splitting, paralleling the description for an elementary splitting.

Let τ be a generic train track and let α be a splitting arc of τ , containing a sink branch b of τ . Let B_α denote the set of oriented branches of τ terminating at a switch lying in α . Each element of B_α lies on one of the two sides of α . Also, each element of B_α terminates in one of the two components of $\alpha - \text{int}(b)$. Together we obtain a partition of B_α into four sets, where $b, b' \in B_\alpha$ are in the same partition element if they lie on the same side of α and touch the same component of $\alpha - \text{int}(b)$. Each partition element is nonempty, and we denote the partition $B_\alpha = R \cup S \cup T \cup U$, where R, S touch one component of $\alpha - \text{int}(b)$ and T, U touch the other component, where R, U lie on one side of α and S, T lie on the other side, and where R, S, T, U

are listed in circular order with respect to the orientation on a neighborhood of α . In analogy with the four points r, s, t, u in Figure 13, one can easily label the sets R, S, T, U in Figure 14. Given a subtrack $\sigma \subset \tau$, we say that σ contains a Left crossing of α if σ contains a branch of S and a branch of U , and σ contains a Right crossing of α if σ contains a branch of R and a branch of T .

Proposition 6.1.3. *Given a wide splitting of generic train tracks $\tau \succ \tau'$ of parity $d \in \{L, R\}$ along a splitting arc α , a subtrack $\sigma \subset \tau$ survives the splitting if and only if σ contains a d crossing or σ does not contain a \bar{d} crossing of α . If σ survives then its descendant σ' is the unique subtrack of τ' such that $\sigma \succ_{\text{F}} \sigma'$ and such that σ' contains the postsplitting branch of τ' if and only if σ contains a d -crossing of α .*

Proof. Consider the slide equivalence $\tau \succ \tau''$ depicted in Figure 15 with sink branch $b'' \subset \tau''$, which has the property that $\tau'' \succ \tau'$ is an elementary splitting of parity d along b'' , and the postsplitting branch of τ' is the same for the wide splitting $\tau \succ \tau'$ and the elementary splitting $\tau'' \succ \tau'$. Applying Proposition 6.1.1, σ survives in τ'' with descendant $\sigma'' \subset \tau''$, and in the notation of Figure 15, σ contains the branch R if and only if σ'' contains r , and similarly for the pairs S, s and T, t and U, u . It follows that σ contains a d crossing of b if and only if σ'' contains a d crossing of b'' , and similarly for \bar{d} crossings. Applying Proposition 6.1.1 it follows that σ survives in τ' if and only if σ contains a d crossing or σ does not contain a \bar{d} crossing. \diamond

The following lemma gives a description of survival and descent for any homotopic carrying relation factored as a sequence of parity splittings. By combining it with Lemmas 6.1.2 or 6.1.3, we obtain an efficient computational criterion for verifying survival of a subtrack.

The lemma is stated more generally in terms of any carrying relation factored as a sequence of carrying relations:

Lemma 6.1.4. *Given a sequence of carrying relations $\tau_0 \succ \tau_1 \succ \cdots \succ \tau_n$, a subtrack $\sigma_0 \subset \tau_0$ survives in τ_n with descendant $\sigma_n \subset \tau_n$ if and only if there exist subtracks $\sigma_i \subset \tau_i$ for $i = 1, \dots, n-1$ such that for each $i = 1, \dots, n$ the subtrack σ_{i-1} survives the splitting $\tau_{i-1} \succ \tau_i$ with descendant σ_i . For a parity splitting sequence it follows, by combining with Proposition 6.1.2 above, that the descent relation is a bijection between a subset of the set of subtracks of τ_0 and a subset of the set of subtracks of τ_n .*

Remark There exist splitting sequences $\tau_0 \succ \cdots \succ \tau_n$ such that no proper subtrack of τ_0 survives in τ_n . For instance, on a once punctured torus this holds whenever the sequence contains splittings of both parities L and R. Examples where the bijection is empty also exist on a general finite type surface S : any train track τ_0

has only finitely many proper subtracks, and they can be killed off one by one in some carefully chosen parity splitting sequence $\tau_0 \succ \tau_1 \succ \cdots \succ \tau_n$; furthermore, we can then extend the splitting sequence beyond τ_n in an arbitrary manner, say $\tau_n \succ \cdots \succ \tau_m$, and the result is that no subtrack of τ_0 survives to τ_m . This shows that in some sense the property that no subtrack survives a splitting sequence is a generic property.

Proof of Lemma 6.1.4. If the subtracks σ_i exist as described then σ_0 survives with descendant σ_n , because the composition of homotopy equivalent carrying maps $\sigma_n \rightarrow \sigma_{n-1} \rightarrow \cdots \rightarrow \sigma_0$ is a homotopy equivalent carrying map.

To prove the converse, suppose then that $\sigma_0 \subset \tau_0$ survives with descendant $\sigma_n \subset \tau_n$. Working backwards by induction starting with σ_n , let $\sigma_{i-1} \subset \tau_{i-1}$ be the image of σ_i under the carrying map $\tau_i \rightarrow \tau_{i-1}$. It follows that $\sigma_{i-1} \overset{F}{\succ} \sigma_i$ and so by Corollary 3.14.2 we have $i(\sigma_0) \geq i(\sigma_1) \geq \cdots \geq i(\sigma_n)$. Since $\sigma_0 \overset{H}{\succ} \sigma_n$ it follows again by Corollary 3.14.2 that $i(\sigma_0) = i(\sigma_n)$, implying that $i(\sigma_{i-1}) = i(\sigma_i)$ for all $i = 1, \dots, n$, and so again by Corollary 3.14.2 each of the carrying relations $\sigma_{i-1} \overset{H}{\succ} \sigma_i$ is homotopic. \diamond

6.2 Survival of splitting cycles

Recall that a C-splitting arc of a semigeneric train track τ is an embedded arc $\alpha \subset S$ such that $\alpha \cap \tau$ is an embedded train path contained in the interior of α , each end of $\alpha \cap \tau$ is a switch pointing into the train path $\alpha \cap \tau$, and each of the two components of $\mathcal{C}(\alpha - \tau)$ has an endpoint on some cusp of some component of $\mathcal{C}(S - \tau)$. If τ is generic then a C-splitting arc is completely determined by its intersection with τ , up to isotopy rel τ . We remind the reader that a C-splitting arc α can contain any number of sink branches, and that a C-splitting arc α is an (unmodified) splitting arc if and only if α contains a unique sink branch.

Consider a filling, semigeneric train track σ (we don't bother defining splitting cycles for nonfilling train tracks). A *splitting cycle* of σ is either a simple closed curve or a proper line with ends at distinct punctures of P , denoted γ , that satisfies the following properties:

- For each component a of $\gamma \cap \sigma$ there is a subsegment $\alpha \subset \gamma$ such that α is a C-splitting arc and $\alpha \cap \sigma = a$.
- For each component Q of $\mathcal{C}(S - \sigma)$, the set $\gamma \cap Q$ is one of the following:
 - the empty set, or
 - a compact arc with endpoints at distinct cusps of Q , or

- a half-open arc, with finite endpoint at a cusp of Q , and with infinite end escaping to the puncture of Q .

The equivalence relation we use for splitting cycles of σ is *relative isotopy*, meaning isotopy relative to σ . Note that if σ is generic then two splitting cycles γ, γ' of σ are relatively isotopic if and only if $\gamma \cap \sigma = \gamma' \cap \sigma$. This equivalence can fail when σ is nongeneric, but in general σ is determined, up to relative isotopy, by the relative isotopy classes of the \mathcal{C} -splitting arcs that σ contains. This shows that σ has finitely many splitting cycles up to relative isotopy, and gives an easy method for enumerating them.

It is possible that two splitting cycles which are not relatively isotopic are nonetheless ambiently isotopic. For instance suppose that the splitting cycle γ contains a \mathcal{C} -splitting arc α , and that some component Q of $\mathcal{C}(S - \sigma)$ has a side β which embeds in α via the overlay map. Then γ can be isotoped so that the segment β is replaced with a diagonal of Q parallel to the side β , and the result of this isotopy is a splitting cycle γ' which is not relatively isotopic to γ .

Consider a carrying pair $\sigma \rightsquigarrow \sigma'$ of filling train tracks. Given a splitting cycle γ of σ , we say that γ *survives* in σ' if there exists a splitting cycle γ' of σ' and a carrying map $f: (S, \sigma') \rightarrow (S, \sigma)$ such that:

- (1) γ' is transverse to the fibers of f ,
- (2) the map $f|_{\gamma'}$ is injective, and
- (3) $f(\gamma') = \gamma$.

In this case γ' is a *descendant* of γ and γ is an *ancestor* of γ' . Note that properties (1), (2) alone imply that $f(\gamma')$ is a splitting cycle of σ , and so γ' has an ancestor in σ ; (3) identifies the ancestor with γ . If γ does not survive in σ' then we say γ is *killed* in σ' .

As in the last section, we can get a more precise, combinatorial characterization of survival of splitting cycles for various special kinds of carrying relations, starting with comb equivalence.

Proposition 6.2.1. *If $\sigma \rightsquigarrow \sigma'$ is a comb equivalence of filling train tracks, then every central splitting of σ survives in σ' , and the descent relation induces a bijection between all relative isotopy classes of splitting cycles of σ and all those of σ' .*

Proof. Fix carrying maps $f: (S, \sigma) \rightarrow (S, \sigma')$ and $\bar{f}: (S, \sigma') \rightarrow (S, \sigma)$. For any splitting cycle γ of σ , once γ is isotoped rel σ to be transverse to the fibers of the carrying map f , it follows that $f(\gamma)$ is a splitting cycle of σ' with a well-defined isotopy class rel σ' . By symmetry of comb equivalence a similar statement holds

in the other direction. For any splitting cycle γ of σ the splitting cycle $\bar{f} \circ f(\gamma)$ is isotopic to γ rel σ , and similarly in the other direction. \diamond

Next we consider survival of splitting cycles under an elementary splitting of parity $d \in \{L, R\}$. Recall that a subtrack is killed if and only if it contains a \bar{d} crossing of the sink branch but not a d crossing. With a splitting cycle γ the idea is similar, but somewhat complicated by the fact that $\gamma \cap \sigma$ may have endpoints at one or two of the endpoints of the sink branch.

Let $\sigma_0 \succ \sigma_1$ be a splitting of parity $d \in \{L, R\}$ along a sink branch $b \subset \sigma_0$, and as usual N is a regular neighborhood of b , with $\{r, s, t, u\} = \sigma_0 \cap \partial N$ and with $\partial b = \{p, q\}$, so that p, r, s are in one component of $(\sigma_0 \cap N) - \text{int}(b)$ and q, t, u are in the other, and r, s, t, u are in positive (counterclockwise) circular order on ∂N ; see Figures 13 and 18). Given a C-splitting arc γ_0 of σ_0 , we say that γ_0 contains a *partial Left crossing of b* if $\gamma_0 \cap \sigma_0 \cap N \subset \overline{su}$, which occurs if and only if $\gamma_0 \cap \sigma_0 \cap N = \overline{su}$ or \overline{pu} or \overline{sq} or \overline{pq} . Similarly, γ_0 contains a *partial Right crossing of b* if $\gamma_0 \cap \sigma_0 \cap N \subset \overline{rt}$ which occurs if and only if $\gamma_0 \cap \sigma_0 \cap N = \overline{rt}$ or \overline{rq} or \overline{pt} or \overline{pq} . Note in particular that $\gamma_0 \cap \sigma_0 \cap N = \overline{pq}$ if and only if γ_0 contains both a partial Left and partial Right crossing.

Proposition 6.2.2. *Consider an elementary splitting $\sigma_0 \succ \sigma_1$, with σ_0, σ_1 generic and filling, of parity d along a sink branch $b \subset \sigma_0$. A splitting cycle γ_0 of σ_0 survives the splitting if and only if γ_0 does not contain a partial \bar{d} -crossing. If γ_0 does survive then its descendant γ_1 is determined up to isotopy by the requirement that $\gamma_0 - N = \gamma_1 - N$. Descent is a bijection (up to relative isotopy) between a subset of the splitting cycles of σ and a subset of the splitting cycles of σ_1 .*

When $d = L$ the proposition asserts that the following two (obviously equivalent) statements are true:

- (1) γ_0 survives $\iff \gamma_0 \cap \sigma_0 \cap N = \emptyset$ or \overline{ru} or \overline{st} or \overline{su} or \overline{pu} or \overline{sq}
 \iff one of cases (a,b,c,d,e,f) holds in Figure 18.
- (2) γ_0 is killed $\iff \gamma_0 \cap \sigma_0 \cap N = \overline{pq}$ or \overline{rt} or \overline{pt} or \overline{rq}
 \iff one of cases (g,h,i,j) holds in Figure 18.

And, when $d = R$, the following two statements are true:

- (3) γ_0 survives $\iff \gamma_0 \cap \sigma_0 \cap N = \emptyset$ or \overline{ru} or \overline{st} or \overline{rt} or \overline{pt} or \overline{rq}
 \iff one of cases (a,b,c,h,i,j) holds in Figure 18.
- (4) γ_0 is killed $\iff \gamma_0 \cap \sigma_0 \cap N = \overline{su}$ or \overline{pu} or \overline{sq} or \overline{pq}
 \iff one of cases (d,e,f,g) holds in Figure 18.

There is a similar result for wide splittings which we leave for the reader to formulate.

Proof. First we claim that the ancestor relation is well defined on relative isotopy classes. As part of this claim, if we fix a carrying map $f: (S, \sigma_1) \rightarrow (S, \sigma_0)$, and if γ_1, γ'_1 are two relatively isotopic splitting cycles each satisfying properties (1), (2) in the definition of survival, then clearly γ_1, γ'_1 intersect the same fibers of f , implying that $f(\gamma_1) \cap \sigma_0 = f(\gamma'_1) \cap \sigma_0$ and so $f(\gamma_1), f(\gamma'_1)$ are relatively isotopic splitting cycles of σ_0 .

To prove the claim, consider a splitting cycle γ_1 of σ_1 and a carrying map $f: (S, \sigma_1) \rightarrow (S, \sigma_0)$ so that (1), (2) in the definition of survival are satisfied, and hence $f(\gamma_1)$ is a splitting cycle of σ_0 . Consider any other carrying map $f': (S, \sigma_1) \rightarrow (S, \sigma_0)$. By the previous paragraph it suffices to show that one can isotope γ_1 relative to σ_1 so that properties (1), (2) are satisfied for f' . Applying Proposition 3.5.2, choose a homotopy of carrying maps $F_t: (S, \sigma_1) \rightarrow (S, \sigma_0)$ such that $F_0 = f, F_1 = f'$. As t varies, the fibers of F_t vary continuously, and we can find an isotopy γ_{1t} of γ_1 so that γ_{1t} is transverse to the fibers of F_t . If for some value of t the map $f_t \mid \gamma_{1t}$ is one-to-one then the same is clearly true for nearby values of t , and moreover $f_t(\gamma_{1t}) \cap \sigma_0$ is constant for nearby values of t . But as t approaches a limit point t' , the constant set $f_t(\gamma_{1t}) \cap \sigma_0$ approaches a limit which equals the set $f_{t'}(\gamma_{1t'}) \cap \sigma_0$, and hence $f_{t'} \mid \gamma_{1t'}$ is also one-to-one. This shows that properties (1), (2) are open and closed with respect to the parameter t , and since they are satisfied for $t = 0$ they must be satisfied for $t = 1$, proving the claim.

We have shown that the ancestor relation is a well-defined function from a subset of the relative isotopy classes of splitting cycles of σ_1 to a subset of the relative isotopy classes of splitting cycles of σ_0 . We must identify these two subsets.

From the argument above, we may choose the carrying map $f: (S, \sigma_1) \rightarrow (S, \sigma_0)$ arbitrarily; choose f to be the fold map associated to the splitting.

Up to changing orientation of S if necessary, we may assume that $d = L$. Let b' be the postsplitting branch of σ_1 , with endpoints p', q' denoted so that r, u, p' are in one component of $(\sigma_1 \cap N) - \text{int}(b')$ and s, t, q' are in the other component.

Consider a splitting cycle γ_1 of σ_1 transverse to the fibers of f . Note that $\gamma_1 \cap \sigma_1 \cap N$ has at most two components, because each component must contain one of the points p', q' .

We shall show that if $\gamma_1 \cap \sigma_1 \cap N$ has two components then it has no ancestor, whereas if $\gamma_1 \cap \sigma_1 \cap N$ has at most one component then it has an ancestor, and the ancestor is one of the cases (a,b,c,d,e,f) of Figure 18.

Consider first the case that $\gamma_1 \cap \sigma_1 \cap N$ has at most one component. It follows that $\gamma_1 \cap \sigma_1 \cap N$ is one of the six sets \emptyset or \overline{ru} or \overline{st} or \overline{su} or $\overline{q'_L u}$ or $\overline{sp'_L}$. In each of these cases clearly $f \mid \gamma_1$ is injective, $\gamma_0 = f(\gamma_1)$ is a splitting cycle of σ , and

$\gamma_0 \cap \sigma_0 \cap N$ is, respectively, one of the six sets in (1) above, as listed in the following table and depicted in the indicated cases of Figure 18:

	(a)	(b)	(c)	(d)	(e)	(f)
$\gamma_0 \cap \sigma_0 \cap N$	\emptyset	\overline{ru}	\overline{st}	\overline{su}	\overline{pu}	\overline{sq}
$\gamma_1 \cap \sigma_1 \cap N$	\emptyset	\overline{ru}	\overline{st}	\overline{su}	$\overline{q'_L u}$	$\overline{sp'_L}$

Moreover, from this description it clearly follows that the ancestor function is injective on relative isotopy classes, for the table clearly shows that if $\gamma_0 = f(\gamma_1)$ is the ancestor of γ_1 then $\gamma_0 \cap \sigma_0$ determines $\gamma_1 \cap \sigma_1$.

Noting that cases (a,b,c,d,e,f) of Figure 18 are precisely the cases of parity $d = L$ in which it is claimed that γ_0 survives, it remains to prove that γ_1 has no ancestor in the case that $\gamma_1 \cap \sigma_1 \cap N$ has two components. One of these components must be \overline{ru} or $\overline{q'_L u}$ and the other must be $\overline{sp'_L}$ or \overline{st} . For each $x \in [p, q]$ the fiber $f^{-1}(x)$ contains points $x_0 \in [rq'_L]$, $x_1 \in [p'_L q'_L]$, $x_2 \in [p'_L, t]$, and also contains an arc I having endpoints x_0, x_2 and whose interior contains x_1 ; each of the subarcs $[x_0, x_1], [x_1, x_2] \subset I$ contains a point of γ' different from x_1 . This implies that $f \upharpoonright \gamma_1$ is not injective, and hence γ_1 has no ancestor. \diamond

Finally, as with subtracks we have a characterization of survival for splitting cycles in the context of a sequence of carrying maps:

Lemma 6.2.3. *Given a sequence of carrying relations $\sigma_0 \succ \dots \succ \sigma_n$ of filling train tracks, a splitting cycle $\gamma \subset \sigma_0$ survives in σ_n with descendant γ' if and only if there exist splitting cycles $\gamma = \gamma_0, \gamma_1, \dots, \gamma_n = \gamma'$ of $\sigma_0, \sigma_1, \dots, \sigma_n$, respectively, such that for each $i = 1, \dots, n$ the splitting cycle γ_{i-1} survives the carrying $\sigma_{i-1} \succ \sigma_i$ with descendant γ_i . For a parity splitting sequence it follows, by combining with Proposition 6.2.2 above, that descent is a bijection between a subset of the set of splitting cycles of σ_0 and a subset of the set of splitting cycles of σ_n .*

Proof. The “if” direction is clear. For the other direction, choose carrying maps $f_i: (S, \sigma_i) \rightarrow (S, \sigma_{i-1})$, $i = 1, \dots, n$, and for $0 \leq i < j \leq n$ define a carrying map $f_{ji}: (S, \sigma_j) \rightarrow (S, \sigma_i)$ by $f_{ji} = f_{i+1} \circ \dots \circ f_j$. Choose $\gamma' = \gamma_n$ in its relative isotopy class so that it is tranverse to the fibers of f_{0n} and $f_{0n} \upharpoonright \gamma_n$ is one-to-one with image γ . It follows that γ_n is tranverse to the fibers of $f_{n-1,n}$ and $f_{n-1,n} \upharpoonright \gamma_n$ is one-to-one, so its image γ_{n-1} is a splitting cycle of σ_{n-1} , and moreover γ_{n-1} is tranverse to the fibers of f_{n-1} and $f_{n-1} \upharpoonright \gamma_{n-1}$ is one-to-one. An induction argument completes the proof. \diamond

6.3 Statement of the Arational Expansion Theorem

We shall first discuss the *subtrack killing criterion* and then the *splitting cycle killing criterion*. Then, putting them together, we obtain the *rational killing criterion*

which is the key to formulating the Arational Expansion Theorem.

The subtrack killing criterion. A carrying pair $\tau \succcurlyeq \tau'$ is said to satisfy the *subtrack killing criterion* if every nonfilling subtrack of τ is killed in τ' .

Note that the subtrack killing criterion is stable under extension in the sense that if $\tau \succcurlyeq \tau' \succcurlyeq \tau'' \succcurlyeq \tau'''$ and if $\tau' \succcurlyeq \tau''$ satisfies the subtrack killing criterion then $\tau \succcurlyeq \tau'''$ also satisfies the criterion: for if $\sigma \subset \tau$ survives with descendant $\sigma''' \subset \tau'''$ then by Lemma 6.1.4 it follows that σ also survives with descendants $\sigma' \subset \tau'$ and $\sigma'' \subset \tau''$, and that σ' survives with descendant σ'' .

The main utility of the subtrack killing criterion, and of other “killing criteria” which are stable under extension, comes from iteration. Given a property \mathcal{A} of carrying pairs which is stable under extension, an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ satisfies the *iterated property* \mathcal{A} if for every $m \geq 0$ there exists $n \geq m$ such that $\tau_m \succcurlyeq \tau_n$ satisfies property \mathcal{A} . This is equivalent to saying that property \mathcal{A} is satisfied starting from τ_0 to some τ_{m_1} , and also from τ_{m_1} to some τ_{m_2} , and also from τ_{m_2} to some τ_{m_3} , and so on. The proof of equivalence is where “stable under extension” is needed.

It is sometimes useful to observe that two different criteria for a carrying pair produce the same iterated criterion for a splitting sequence. Such is the case for the subtrack killing criterion. Define the *strong subtrack filling criterion* for a carrying pair $\tau \succ \tau'$ to say that for each nonfilling subtrack $\sigma' \subset \tau'$, its carrying image $\sigma \subset \tau$ is filling. This criterion is evidently stable under extension.

Lemma 6.3.1. *For each splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, the iterated subtrack filling criterion is equivalent to the iterated strong subtrack filling criterion. To be precise:*

- (1) *If $\rho \succcurlyeq \rho'$ satisfies the strong subtrack filling criterion then it satisfies the subtrack filling criterion.*
- (2) *There exists a constant K depending only on the topology of S such that for any carrying sequence $\rho_0 \succcurlyeq \rho_1 \succcurlyeq \dots \succcurlyeq \rho_K$, if each $\rho_{k-1} \succcurlyeq \rho_k$ satisfies the subtrack filling criterion, $k = 1, \dots, K$, then $\rho_0 \succcurlyeq \rho_K$ satisfies the strong subtrack filling criterion.*

Proof. Suppose $\rho \succ \rho'$ satisfies the strong subtrack filling criterion, and let $\sigma \subset \rho$ be a nonfilling subtrack. If σ survives with descendant $\sigma' \subset \rho'$ then σ' does not fill, and so the carrying image of σ' in ρ does fill, but that image is σ , contradiction.

For the converse, suppose we have a carrying sequence $\rho_0 \succcurlyeq \rho_1 \succcurlyeq \dots \succcurlyeq \rho_K$ such that $\rho_{k-1} \succcurlyeq \rho_k$ satisfies the subtrack killing criterion for each $k = 1, \dots, K$, and suppose that there exists a nonfilling subtrack $\sigma' \subset \rho_K$ whose carrying image $\sigma \subset \rho_0$

does not fill. We must bound K . Starting from $\sigma_K = \sigma'$ and working backwards inductively, let $\sigma_{k-1} \subset \rho_{k-1}$ be the carrying image of $\sigma_k \subset \rho_k$; we end up at the subtrack $\sigma_0 = \sigma' \subset \rho_0$. Clearly we have $\sigma_0 \xrightarrow{F} \sigma_1 \xrightarrow{F} \cdots \xrightarrow{F} \sigma_K$, and so by Corollary 3.14.2 we have $i(\sigma_0) \geq i(\sigma_1) \geq \cdots \geq i(\sigma_K)$. Since σ_0 does not fill, none of the σ_k fill, and so by applying the subtrack filling criterion K times it follows that $\sigma_{k-1} \subset \rho_{k-1}$ does not survive in ρ_k . Thus, the full carrying relation $\sigma_{k-1} \xrightarrow{F} \sigma_k$ is *not* a homotopic carrying relation, and by applying Corollary 3.14.2 again we get $i(\sigma_0) > i(\sigma_1) > \cdots > i(\sigma_K) \geq 0$, and in fact $i(\sigma_{k-1}) \geq i(\sigma_k) + 2$. But $i(\sigma_0) \leq 6|\chi(S)| - 2p$ by Lemma 3.2.1, and so $K \leq 3|\chi(S)| - p$, where p is the number of punctures of S . \diamond

The strong subtrack killing criterion seems sometimes to be easier to verify: special knowledge about a carrying pair $\tau \succ \tau'$ may make it evident that every nonfilling subtrack of τ' has filling carrying image in τ .

On the other hand, as the statement of the lemma shows, the subtrack killing criterion seems to be verifiable on much shorter segments of a splitting sequence than the strong subtrack killing criterion. This can be useful in computer verification of arationality, such as the type of computations that go into checking when a mapping class is pseudo-Anosov given an invariant train track. We shall discuss this issue more later.

The splitting cycle killing criterion. A carrying pair $\tau \succ \tau'$ satisfies the *splitting cycle killing criterion* if for each filling subtrack $\sigma \subset \tau$ which survives with descendant $\sigma' \subset \tau'$, no splitting cycle of σ survives in σ' . This criterion is stable under extension, for suppose that $\tau \succ \tau' \succ \tau'' \succ \tau'''$ and $\tau' \succ \tau''$ satisfies the criterion. Let $\sigma \subset \tau$ survive with descendant $\sigma''' \subset \tau'''$, and suppose that γ is a splitting cycle of σ which survives with descendant a splitting cycle γ''' of σ''' . Applying Proposition 6.1.4 it follows that σ has descendants $\sigma' \subset \tau'$, $\sigma'' \subset \tau''$ and that σ' survives with descendant σ'' . Applying Proposition 6.2.3 it follows that there are splitting cycles γ' of σ' and γ'' of σ'' which are descendants of γ so that γ'' is a descendant of γ' .

Statement of the theorem. A carrying pair $\tau \succ \tau'$ satisfies the *rational killing criterion* if it satisfies both the subtrack killing criterion and the splitting cycle killing criterion; equivalently, for every subtrack σ of τ , if σ survives in τ' with descendant σ' then σ, σ' fill, and no splitting cycle of σ survives in σ' . The rational killing criterion is stable under extension, being the conjunction of two criteria which are stable under extension. Given a finite splitting sequence $\tau = \tau_0 \succ \cdots \succ \tau_n = \tau'$,

by applying the results of Sections 6.1 and 6.2 we obtain a very efficient decision procedure for the rational killing criterion.

Note that an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \tau_2 \succ \cdots$ satisfies the rational killing criterion if and only if it satisfies the subtrack killing criterion and the splitting cycle killing criterion. This is a useful comment because experience shows that a given splitting cycle is easier to kill than a given subtrack, and so it is easier to check the splitting cycle killing criterion. In other words, one might be able to verify the splitting cycle killing criterion after fewer splittings than the subtrack killing criterion.

Theorem 6.3.2 (Arational Expansion Theorem). *Let \mathcal{F} be a measured foliation and $\tau_0 \succ \tau_1 \succ \tau_2 \succ \cdots$ a train track expansion of \mathcal{F} . Then \mathcal{F} is arational if and only if the expansion $\tau_0 \succ \tau_1 \succ \tau_2 \succ \cdots$ satisfies the iterated rational killing criterion.*

By combining this with the Expansion Convergence Theorem 5.1.1, we obtain:

Corollary 6.3.3. *For any splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$, exactly one of the following happens: either $\cap \mathcal{MF}(\tau_i) = \mathcal{MF}(\mathcal{F})$ for some arational measured foliation \mathcal{F} ; or the sequence fails the iterated rational killing criterion. \diamond*

When the iterated rational killing criterion fails, the structures of $\tau_0 \succ \tau_1 \succ \cdots$ and of the set $\cap_i \mathcal{MF}(\tau_i)$ are explained in Section 8.

The proof of Theorem 6.3.2 will be given in Sections 6.5 and 6.6. For now we turn to an application of the theorem.

6.4 Application: Canonical train track expansions

Given a splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ which is an expansion of an arational measured foliation \mathcal{F} , we say that it is a *canonical expansion* if each τ_i canonically carries \mathcal{F} . It follows that the singularity type of each τ_i matches the singularity type of \mathcal{F} , which implies in turn that there are no central splittings in the sequence, and hence each carrying $\tau_i \succ \tau_j$, $i < j$, is a homotopic carrying.

Just as the arational expansion theorem gives a combinatorial method for detecting when a splitting sequence is an arational expansion, as an application we shall develop a combinatorial method for detecting when a splitting sequence is a canonical expansion.

Consider a homotopic carrying $\tau \xrightarrow{H} \tau'$ where τ, τ' fill. Given a C-splitting arc α of τ , we say that α *survives* in τ' if for any carrying map $f: \tau' \rightarrow \tau$ there exists a C-splitting arc α' of τ' such that α' is transverse to the fibers of f , the map $f|_{\alpha'}$ is injective, and $\alpha = f(\alpha')$; in this case we say that α' is a *descendant* of α in τ' . If α does not survive in τ' then we say that α is *killed* in τ' .

The methods of Section 6.2 can be applied to give a combinatorial characterization of survival and descent of C-splitting arcs, as follows:

Proposition 6.4.1. *Consider a homotopic carrying $\tau \xrightarrow{H} \tau'$.*

- *If τ, τ' are comb equivalent then every C-splitting arc of τ survives in τ' and vice versa, and descent is a bijection up to relative isotopy.*
- *If $\tau \succ \tau'$ is an elementary splitting, of parity d along a sink branch $b \subset \tau$, and if N is a neighborhood of b such that $N - \tau = N - \tau'$, then a C-splitting arc α of τ survives the splitting if and only if α does not contain a partial \bar{d} crossing of b , and if so then α has a descendant α' characterized up to relative isotopy by the requirement that, for some neighborhood N of b , we have $N - \alpha = N - \alpha'$.*
- *For any parity splitting sequence $\tau = \tau_0 \succ \cdots \succ \tau_n = \tau'$, a C-splitting arc α of τ_0 survives with descendant α' in τ_n if and only if there exist C-splitting arcs α_i of τ_i with $\alpha = \alpha_0$, $\alpha' = \alpha_n$, such that each α_{i-1} survives the splitting $\tau_{i-1} \succ \tau_i$ with descendant α_i .*

In all these situations, descent is a bijection up to relative isotopy between a subset of the C-splitting arcs of τ and those of τ' .

Proof. One can prove these statements by the same methods used in Section 6.2 to characterize survival and descent of central splitting *cycles*; alternatively, one can actually reduce the present statements to the results of Section 6.2 by introducing extra punctures as in the proof of Theorem 6.4.2 below. \diamond

We say that the pair $\tau \xrightarrow{H} \tau'$ satisfies the *canonical killing criterion* if no proper subtrack of τ survives in τ' , and no C-splitting arc of τ survives in τ' . This criterion is stable under extension. A splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ satisfies the *iterated canonical killing criterion* if for each $m \geq 0$ there exists $n > m$ such that $\tau_m \xrightarrow{H} \tau_n$ satisfies the canonical killing criterion. Note that this is equivalent to conjunction of the iterated proper subtrack killing criterion and the iterated C-splitting arc killing criterion.

Theorem 6.4.2 (Canonical expansions). *A splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ is a canonical expansion of some arational measured foliation if and only if it satisfies the iterated canonical killing criterion.*

Proof. We know that each $\tau_i \succ \tau_{i+1}$ is a parity splitting, and that each τ_i fills. In particular, there is at most one puncture in each component of $\mathcal{C}(S - \tau_i)$. Note that each carrying map $f_i: \tau_i \rightarrow \tau_{i-1}$, $i \geq 1$, induces a bijection between components of $\mathcal{C}(S - \tau_i)$ and of $\mathcal{C}(S - \tau_{i-1})$. By isotoping each train track τ_i , we may assume by

induction that for each nonpunctured component C of $\mathcal{C}(S - \tau_0)$ there exists a point $p_C \in \text{int}(C)$ such that $f_i(p_C) = p_C$, and p_C lies in the interior of the component of $\mathcal{C}(S - \tau_i)$ corresponding to C . Let S' be the surface obtained from S by removing the point p_C from C , for each nonpunctured component C of $\mathcal{C}(S - \tau_0)$. It follows that $\tau_0 \succ \tau_1 \succ \dots$ is a splitting sequence on S' . Moreover, the intersection $\cap_i \mathcal{MF}(\tau_i)$ is well-defined whether we interpret it on the surface S or on the surface S' .

Choose \mathcal{F} so that $[\mathcal{F}] \in \cap_i \mathcal{MF}(\tau_i)$, and so that \mathcal{F} is a partial measured foliation with a carrying injection $\mathcal{F} \hookrightarrow \nu(\tau_0)$ for some tie bundle $\nu(\tau_0)$. Thus, we may interpret \mathcal{F} as a partial measured foliation on the surface S or on S' . Note that by the Expansion Convergence Theorem 5.1.1, or rather by its proof, the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ is a canonical expansion of *some* arational measured foliation if and only if it is a canonical expansion of \mathcal{F} .

First we reduce to the case that \mathcal{F} is fully carried on each τ_i . If not, then there exists I such that for all $j \geq I$ the carrying map $\mathcal{F} \rightarrow \tau_j$ is not full, and so its image is a proper subtrack σ_j . We have a full carrying sequence $\sigma_I \xrightarrow{\mathcal{F}} \sigma_{I+1} \xrightarrow{\mathcal{F}} \dots$, and by Proposition 3.14.1, together with the fact the the number of central splittings in any splitting sequence is bounded, the carrying $\sigma_i \xrightarrow{\mathcal{F}} \sigma_{i+1}$ is homotopic for sufficiently large i , say $i \geq I$. It follows that σ_I is a proper subtrack of τ_I that survives in each τ_i , $i \geq I$, contradicting the iterated canonical killing criterion.

Assuming now that \mathcal{F} is fully carried on each τ_i , we may choose \mathcal{F} so that the carrying injection $\mathcal{F} \hookrightarrow \nu(\tau_0)$ is surjective. Since all splittings $\tau_i \succ \tau_{i+1}$ are parity splittings, it follows that \mathcal{F} is a canonical model on S if and only if each τ_i canonically carries \mathcal{F} , if and only if τ_0 canonically carries \mathcal{F} .

To prove Theorem 6.4.2, we shall prove that the equivalence of the following chain of statements:

- (1) $\tau_0 \succ \tau_1 \succ \dots$ satisfies the iterated canonical killing criterion on S .
- (2) $\tau_0 \succ \tau_1 \succ \dots$ satisfies the iterated rational killing criterion on S' .
- (3) \mathcal{F} is arational on S' .
- (4) \mathcal{F} is a canonical model on S .

The equivalence of (2) and (3) is just the Arational Expansion Theorem 6.3.2.

Proof that (2) implies (1). Assuming that (1) fails, there exists I such that for all $j \geq I$ the pair $\tau_I \not\xrightarrow{\mathcal{F}} \tau_j$ fails the canonical killing criterion, and so τ_I has either a proper subtrack or a C-splitting arc which survives into τ_j . Since there are only finitely many subtracks and finitely many C-splitting arcs, it follows that *either* there is a proper subtrack of τ_I which survives into τ_j for all $j \geq I$, *or* there is a C-splitting arc of τ_I which survives into τ_j for all $j \geq I$.

Case 1: Suppose that there exists a proper subtrack $\sigma \subset \tau_I$ which survives in each τ_j . Since σ is proper, there exists a branch b of τ_I with $b \not\subset \sigma$. The surface S' contains an essential simple closed curve c such that $c \subset S' - (\tau_I - b) \subset S' - \sigma$: if both sides of b are incident to the same component C of $S' - \tau_I$ then c is contained in $C \cup b$ intersecting b once; whereas if the two sides of b are incident to different components C_1, C_2 of $S' - \tau_I$, then c bounds a twice punctured disc contained in $C_1 \cup C_2 \cup b$. It follows that σ does not fill S' , and so the iterated rational killing criterion fails for $\tau_0 \succ \tau_1 \succ \dots$ in S' .

Case 2: Suppose that there exists a C-splitting arc α of τ_I which survives into each τ_j . We may assume that S' is not a once-punctured torus, because in this case τ_I has a unique C-splitting arc, which must be split in τ_{I+1} , contradicting survival of α .

If the two endpoints of α lie in separate components C_1, C_2 of $S' - \tau_I$, containing punctures p_1, p_2 respectively, then we can stretch α out to the punctures p_1, p_2 to produce a splitting cycle of τ_I on S' which survives into each τ_j .

On the other hand, suppose the two endpoints of α lie in the same component C of $S' - \tau_I$. Since the two ends of α enter different cusps of C it follows that C is a once-punctured k -gon for some $k \geq 2$. Form a closed curve $\gamma = \alpha \cup \beta$ in S' where β connects the endpoints of α in C . It follows that γ is a splitting cycle of τ_I , and since α survives in each τ_j it follows that γ survives in each τ_j .

Having produced a splitting cycle of τ_I which survives in τ_j , it follows again that the iterated rational killing criterion fails for $\tau_0 \succ \tau_1 \succ \dots$ on the surface S' .

Proof that (1) implies (2) Assuming that (2) fails, there exists I such that for all $j \geq I$ the pair $\tau_I \not\stackrel{H}{\succ} \tau_j$ fails the rational killing criterion in S' . We may assume that no proper subtrack of τ_I survives in τ_j , and so there exists a splitting cycle γ_I of τ_I which survives with a descendant γ_j in τ_j . Let $f: \tau_j \rightarrow \tau_I$ be a carrying map taking γ_j diffeomorphically to γ_I . Let α be any C-splitting arc contained in γ_I . It follows that α survives in τ_j , with descendant $f^{-1}(\alpha)$, and so $\tau_I \not\stackrel{H}{\succ} \tau_j$ fails the canonical killing criterion in the surface S .

Proof that (3) implies (4) Because we have a carrying injection $\mathcal{F} \subset \nu(\tau_I)$, every saddle connection of \mathcal{F} is either proper or boundary.

If \mathcal{F} is not a canonical model then \mathcal{F} has a proper saddle connection α . If the endpoints of α are in the same component C of $S' - \text{Supp}(\mathcal{F})$ then there is an essential simple closed curve c obtained by connecting the endpoints of α by an arc properly embedded in C . Clearly $\langle \mathcal{F}, c \rangle = 0$ and so \mathcal{F} is not arational on S' . If the endpoints of α are in distinct components C_1, C_2 of $S' - \text{Supp}(\mathcal{F})$, containing

punctures p_1, p_2 respectively, then there is an essential simple closed curve c which is homotopic to a closed curve obtained by travelling along α in one direction, looping around p_1 in C_1 , then travelling along α in the other direction, and looping around p_2 in C_2 to close up. Again $\langle \mathcal{F}, c \rangle = 0$ and so \mathcal{F} is not arational in S' .

Proof that (4) implies (3) If \mathcal{F} is a canonical model on S then \mathcal{F} has no proper saddle connections, and it immediately follows that \mathcal{F} is arational in S' . \diamond

6.5 Necessity of the iterated rational killing criterion

Suppose that $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion of an arational measured foliation \mathcal{F} . By the expansion convergence theorem 5.1.1, it follows that everything in $\cap \mathcal{PMF}(\tau_i) = \mathcal{PMF}(\mathcal{F})$ is arational. Arguing by contradiction, we shall assume that $\tau_0 \succ \tau_1 \succ \dots$ fails the iterated rational killing criterion, and prove that $\cap \mathcal{PMF}(\tau_i)$ contains a measured foliation which is not arational.

Failure of the iterated rational killing criterion says that there exists M such that for each $n > M$ there exist a subtrack $\sigma_M^n \subset \tau_M$ which survives with descendant $\sigma_n \subset \tau_n$, and either σ_M^n does not fill, or it does fill and some splitting cycle of σ_M^n survives in σ_n .

Since τ_M has only finitely many subtracks, by passing to a subsequence we may assume that the subtrack σ_M^n is independent of n , and we denote it σ_M . This subtrack survives in τ_n for arbitrarily large n , and its descendants are uniquely determined by σ_M . It follows that there exists a sequence of homotopic carryings $\sigma_M \succ \sigma_{M+1} \succ \dots$ such that σ_m is a subtrack of τ_m for each $m \geq M$. Such a sequence is called a *line of descent* for the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, and by Proposition 6.1.2 it follows that each carrying $\sigma_m \succ \sigma_{m+1}$ is either a comb equivalence or a parity splitting.

Case 1: σ_M does not fill. We have $\mathcal{PMF}(\sigma_i) \subset \mathcal{PMF}(\tau_i)$ and so $\cap_i \mathcal{PMF}(\sigma_i) \subset \cap_i \mathcal{PMF}(\tau_i)$. But $\cap_i \mathcal{PMF}(\sigma_i)$ is a nonempty set of nonarational measured foliations contained in $\cap_i \mathcal{PMF}(\tau_i)$, a contradiction.

Case 2: σ_M fills. Since σ_M has only finitely many splitting cycles, by passing to a subsequence we may assume that there is a splitting cycle γ_M of σ_M which survives infinitely, with descendant γ_n in Σ_n .

Let ρ_n be the train track obtained from σ_n by simultaneous central splitting along each of the c-splitting arcs in the splitting cycle γ_n . Note that ρ_n does not fill, because ρ_n is disjoint from γ_n , and either γ_n is an essential closed curve or a proper line with ends at distinct punctures.

We claim that $\rho_M \succ \rho_{M+1} \succ \dots$. Noting that $\tau_n \succ \sigma_n \succ \rho_n$ it follows that we have an inclusion $\mathcal{PMF}(\rho_n) \subset \mathcal{PMF}(\tau_n)$ for all n , and so $\cap \mathcal{PMF}(\rho_i) \subset \cap \mathcal{PMF}(\tau_i)$. From the claim it follows that the set $\cap \mathcal{PMF}(\rho_i)$ is a nested intersection, therefore nonempty. Since ρ_i does not fill, the set $\cap \mathcal{PMF}(\rho_i)$ consists entirely of nonarational measured foliations, each contained in $\cap \mathcal{PMF}(\tau_i)$, a contradiction.

It remains to prove the claim, which can be restated in the following general form: given a parity splitting $\sigma \succ \sigma'$ of filling train tracks, and given a splitting cycle γ of σ which survives with descendant γ' in σ' , if we split σ along γ to produce ρ , and if we split σ' along γ' to produce ρ' , then $\rho \succ \rho'$. Applying Propositions 3.13.4 and 6.2.1, we may assume that $\sigma \succ \sigma'$ is an elementary splitting of parity $d \in \{L, R\}$ along a sink branch b . We adopt the notation of Proposition 6.2.2: in particular the carrying map $f: (S, \sigma') \rightarrow (S, \sigma)$. We also have carrying maps $g: (S, \rho) \rightarrow (S, \sigma)$, $g': (S, \rho') \rightarrow (S, \sigma')$ for the multiple central splittings.

Our task is to solve a mapping problem

$$\begin{array}{ccc}
 \sigma & \xleftarrow{f} & \sigma' \\
 \uparrow g & & \uparrow g' \\
 \rho & \xleftarrow{F} & \rho'
 \end{array}$$

where $F: (S, \rho') \rightarrow (S, \rho)$ is a carrying map. Up to change of parity and notation, it suffices to consider cases (a,c,d,e) in Proposition 6.2.2 and Figure 18, with $d = L$. Case (a) is the easiest, and cases (c,d,e) are depicted in Figure 19. Here is a formal description of F .

Given two partitions d, d' of a set X we write $d < d'$ if for each $A \in d$ there exists $A' \in d'$ such that $A \subset A'$, and we say d is a *subpartition* of d' . Given a surjective map $m: X \rightarrow Y$, denote the point inverse image partition of X as $m^\# = \{m^{-1}(y) \mid y \in Y\}$. More generally, given a partition Δ of Y , denote the pullback of Δ by $m^\#(\Delta) = \{m^{-1}(D) \mid D \in \Delta\}$, and note that $m^\# < m^\#(\Delta)$. Given a composition of surjective maps $X \xrightarrow{m} Y \xrightarrow{n} Z$ we have $(n \circ m)^\# = m^\#(n^\#)$. Given a surjective map $m: X \rightarrow Y$ and a partition Δ of X such that $m^\# < \Delta$, the pushforward of Δ is a partition of Y denoted $m_\#(\Delta) = \{m(D) \mid D \in \Delta\}$.

Let $r_{\sigma'}$ be the collapsing rectangle for the carrying map $f: \sigma' \rightarrow \sigma$ with post-splitting branch b' crossing $r_{\sigma'}$ diagonally from corner p' to corner q' . Let r_ρ be the disjoint union of collapsing rectangles for $g: (S, \rho) \rightarrow (S, \sigma)$. Let $r_{\rho'}$ be the disjoint union of collapsing rectangles for $g': (S, \rho') \rightarrow (S, \sigma')$. Note that the nontrivial elements of $f^\#$ are precisely the vertical segments of $r_{\sigma'}$, and similarly for $g^\#$ and $g'^\#$. Noting that the map F must satisfy $F^\# < F^\#(g^\#) = (g \circ F)^\# = (f \circ g')^\# = g'^\#(f^\#)$, this suggests that we may define $F: S \rightarrow S$ by describing $F^\#$ as a subpartition of $g'^\#(f^\#)$. To check that this works, one must carry out the verification that F

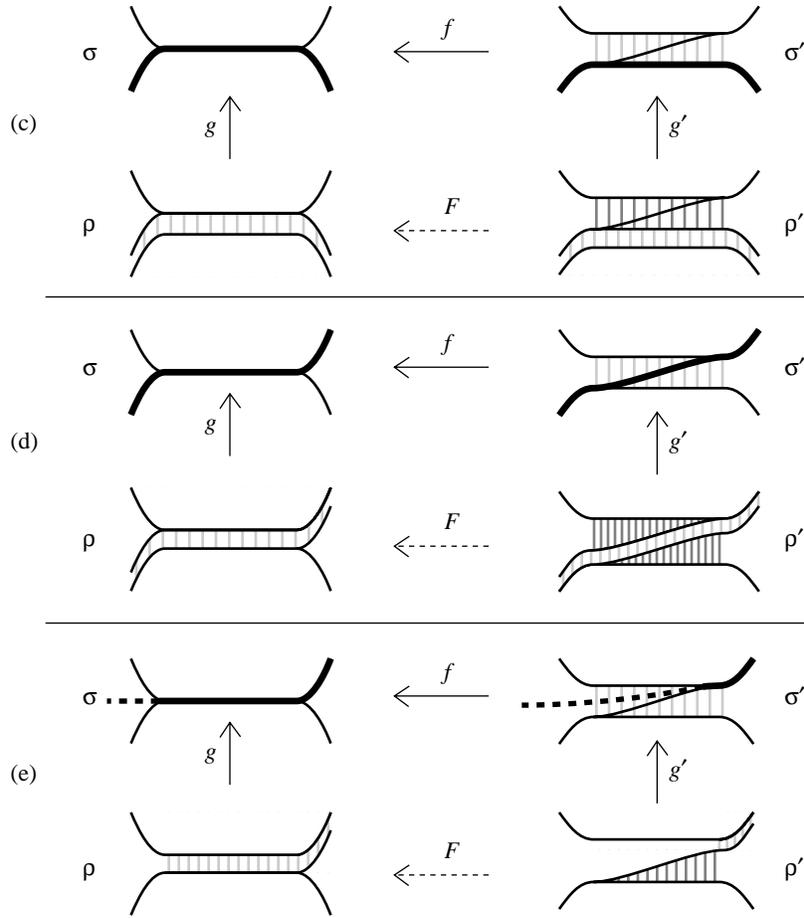


Figure 19: The three cases (c,d,e) correspond to similar cases in Figure 18. In case (c) there is an elementary splitting $\rho \succ \rho'$, and in cases (d,e) there is an isotopy between ρ, ρ' . Collapse of the darker vertical segments defines $F: (S, \rho') \rightarrow (S, \rho)$; collapse of lighter vertical segments defines the other carrying maps.

is homotopic to the identity and that $F_{\#}(g^{\#}(f^{\#})) = g^{\#}$, tasks which are left to the reader; we shall merely describe the nontrivial elements of $F^{\#}$, referring the reader to Figure 19 to understand the description of $F^{\#}$ and to aid in the necessary verifications in cases (c,d,e).

Case (a): $\gamma \cap D = \emptyset$. Letting $\hat{r} = g'^{-1}(r_{\sigma'})$, it follows that g' restricts to a bijection between \hat{r} and $r_{\sigma'}$. Moreover, $\hat{r} \cap r_{\rho'} = \emptyset$. The partition $g'^{\#}(f^{\#})$ is the set of vertical leaves of \hat{r} and of $r_{\rho'}$. We define the nontrivial elements of $F^{\#}$ to be the vertical leaves of \hat{r} . This makes F the carrying map of an elementary splitting $\rho \succ \rho'$.

Case (c): $\gamma \cap D = \overline{st}$. In this case $r_{\sigma'} \cap \gamma'$ is a horizontal side of $r_{\sigma'}$. It follows that the set $\hat{r} = \text{Cl}(g'^{-1}(r_{\sigma'}) - r_{\rho'})$ is a rectangle, mapping bijectively to $r_{\sigma'}$ under g' . Again define the nontrivial elements of $F^{\#}$ to be the vertical leaves of \hat{r} . This again makes F the carrying map of an elementary splitting $\rho \succ \rho'$.

Case (d): $\gamma \cap D = \overline{su}$. In this case $r_{\sigma'} \cap \gamma' = b'$. It follows that the set $\text{Cl}(g'^{-1}(r_{\sigma'}) - r_{\rho'})$ is union of two triangles, each with one cusp and two corners, each vertically foliated. The leaves of these foliations are defined to be the nontrivial elements of $F^{\#}$. This makes F a switch fold from ρ' to ρ , proving that ρ, ρ' are isotopic.

Case (e): $\gamma \cap D = \overline{pu}$. In this case $\alpha = r_{\sigma'} \cap \gamma'$ is an arc with one end at q' , intersecting the opposite vertical side of $r_{\sigma'}$ in an interior point of that side, crossing each vertical fiber once. Note that the branch b' divides $r_{\sigma'}$ into two triangles, each having one cusp and two corners; one of these triangles contains α ; the other, denoted t , is disjoint from α except at the corner q' . Let $\hat{t} = \text{Cl}(g'^{-1}(t) - g'^{-1}(q'))$, foliated by pulling back the vertical leaves of $r_{\sigma'}$ intersected with t , and define the nontrivial elements of $F^{\#}$ to be the vertical leaves of \hat{t} . It follows that F is a switch fold from ρ' to ρ , and again ρ, ρ' are isotopic.

This completes the proof that the iterated rational killing criterion is necessary for $\tau_0 \succ \tau_1 \succ \tau_2 \succ \dots$ to be a train track expansion of an arational measured foliation.

6.6 Sufficiency of the iterated rational killing criterion

Let \mathcal{F} be a nonarational measured foliation, and let $\tau_0 \succ \tau_1 \succ \dots$ be a train track expansion of \mathcal{F} . We shall prove that this splitting sequence fails the iterated rational killing criterion.

First we reduce to the following special case:

Full Carrying Hypothesis Each τ_i fills S and fully carries \mathcal{F} .

To carry out this reduction, let $\tau'_i \subset \tau_i$ be the subtrack fully carrying \mathcal{F} . Then we have $\tau'_i \overset{\text{F}}{\succ} \tau'_{i+1}$ for all $i \geq 0$, and this carrying is either a comb equivalence or a splitting. Applying Corollary 3.14.2, for sufficiently large i , say $i \geq I$, we have $\tau'_i \overset{\text{H}}{\succ} \tau'_{i+1}$, and so the subtrack τ'_I survives infinitely. If τ'_I does not fill then we immediately contradict the iterated rational killing criterion for the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$. If τ'_I fills then so do all τ'_i with $i \geq I$. If only finitely many of the carryings $\tau'_I \overset{\text{H}}{\succ} \tau'_{I+1} \overset{\text{H}}{\succ} \dots$ are splittings then for sufficiently large i , say $i \geq I' \geq I$, the carrying $\tau'_i \overset{\text{H}}{\succ} \tau'_{i+1}$ is a comb equivalence; but in this case it follows that every splitting cycle of $\tau'_{I'}$ survives infinitely, also contradicting the iterated rational killing criterion. Infinitely many of the carryings $\tau'_{I'} \overset{\text{H}}{\succ} \tau'_{I'+1} \overset{\text{H}}{\succ} \dots$ are therefore splittings, and by ignoring the comb equivalences in this sequence we may regard it as a train track expansion of \mathcal{F} each track of which fully carries \mathcal{F} . This expansion clearly inherits the iterated rational killing criterion from the original expansion $\tau_0 \succ \tau_1 \succ \dots$, completing the reduction.

The Full Carrying Hypothesis allows us to apply Proposition 4.4.1. Replacing each τ_i by some comb equivalent generic train track, and borrowing notation from Sections 4.3–4.4, we therefore have the following objects:

- a foliated tie bundle $\mu: \nu_i \rightarrow \tau_i$, whose horizontal foliation \mathcal{F}_i^h is a partial measured foliation equivalent to \mathcal{F} , and with vertical foliation denoted \mathcal{F}_i^v , where the map μ_i plays the dual role of a carrying map $\mu_i: \mathcal{F}_i^h \rightarrow \tau_i$.
- a generic expanding separatrix family $\xi(t)$ of ν_0 , which we may assume is parameterized by length;
- a sequence $t_0 < t_1 < \dots$ such that, setting $\xi_i = \xi(t_i)$, we have $\tau_i = \tau_0(\xi(t_i))$, and moreover the elementary move sequence of $\xi[t_i, t_{i+1}]$ is a sequence of slide moves and a single elementary splitting from τ_i to τ_{i+1} ;
- an identification of the foliated tie bundle ν_i with $\nu(\xi(t_i))$, and of the partial measured foliation \mathcal{F}_i^h with $\mathcal{F}(\xi(t_i))$;
- a partial fulfillment map $r_i: (S, \mathcal{F}_i^h, \mathcal{F}_i^v) \rightarrow (S, \mathcal{F}_0^h, \mathcal{F}_0^v)$

Replacing \mathcal{F} within its equivalence class we may assume $\mathcal{F} = \mathcal{F}_0^h$. For the most part we will drop the “horizontal” superscript h , for example writing \mathcal{F}_i instead of \mathcal{F}_i^h , but we will retain the more rarely used “vertical” superscript v .

Remark. Because the sequence $\tau_0 \succ \tau_1 \succ \dots$ has no central splittings it follows that the partial measured foliations \mathcal{F}_i are all isotopic to each other. We shall nevertheless continue to distinguish them from each other by the notation \mathcal{F}_i .

Since \mathcal{F} has at least one arational foliation component, by Proposition 4.4.1 it follows that the expanding separatrix family $\xi(t)$ is complete, meaning that $\cup_t \xi(t)$ contains at least one infinite separatrix. Let \mathcal{F}' be the closure of the union of infinite separatrices contained in $\cup_t \xi(t)$, equivalently, \mathcal{F}' is the union of components of \mathcal{F} that are infinitely split by the sequence $\tau_0 \succ \tau_1 \succ \dots$.

Now we reduce to a further special case:

Infinitely split hypothesis $\mathcal{F} = \mathcal{F}'$, that is, every component of \mathcal{F} is infinitely split by the sequence $\tau_0 \succ \tau_1 \succ \dots$.

To carry out the reduction, for each j let σ_j denote the subtrack of τ_j which fully carries \mathcal{F}' , and so we have a sequence of full carryings $\sigma_0 \succ \sigma_1 \succ \dots$ such that each $\sigma_i \succ \sigma_{i+1}$ is either a comb equivalence or a splitting. By Corollary 3.14.2, all but finitely many of these carryings are comb equivalences or parity splittings, and so for some M we have a line of descent $\sigma_M \overset{\text{H}}{\succ} \sigma_{M+1} \overset{\text{H}}{\succ} \dots$. By truncating the original splitting sequence we may assume that $M = 0$. If σ_0 fails to fill S then we contradict the iterated rational killing criterion for the sequence $\tau_0 \succ \tau_1 \succ \dots$. We may therefore assume that σ_0 fills S . We know that \mathcal{F}' , a subfoliation of the nonarational foliation \mathcal{F} , is not arational. It suffices therefore to prove the theorem for the expansion $\sigma_0 \succ \sigma_1 \succ \dots$ of \mathcal{F}' , because nonarationality of \mathcal{F}' will imply that $\sigma_0 \succ \sigma_1 \succ \dots$ fails the iterated rational killing criterion, which immediately implies that $\tau_0 \succ \tau_1 \succ \dots$ also fails the criterion. This finishes the reduction.

At this point we have reduced to the case that \mathcal{F} is fully carried on each τ_i and that each component of \mathcal{F} is infinitely split, that is, $\cup_t \xi(t)$ is dense in \mathcal{F} . Notice that in carrying out this reduction we have used the iterated subtrack killing criterion several times. We have not yet used the iterated splitting cycle killing criterion.

The idea of the proof is to use nonarationality of \mathcal{F} to exhibit a splitting cycle of some τ_i which survives infinitely. Consider the carrying bijection $\mathcal{F}_i \xrightarrow{\nu} \nu(\tau_i)$ and the set Σ_i consisting of the union of all proper saddle connections of \mathcal{F}_i . If we knew that the carrying map $\nu(\tau_i) \rightarrow \tau_i$ were injective on Σ_i , then we could find the desired splitting cycle. For example, if there exists some proper saddle connection of \mathcal{F}_i whose endpoints are cusps of the same component of $\text{Cl}(S - \nu(\tau_i))$ then by connecting up those cusps by a path in the complement of $\nu(\tau_i)$ we obtain an essential closed curve γ , and if the saddle connection injects into τ_i then the image of γ under the carrying map $(S, \nu(\tau_i)) \rightarrow (S, \tau_i)$ is a splitting cycle of τ_i which survives forever. More generally, nonarationality of \mathcal{F} implies the existence of an essential closed curve or proper line γ that intersects $\nu(\tau_i)$ in a collection of proper

saddle connections, and if the union of these saddle connections injects in τ_i then the image of γ in τ_i is a splitting cycle which survives forever.

There is, of course, no reason a priori to expect the set of proper saddle connections of \mathcal{F}_i to inject in $\nu(\tau_i)$, but the Stability Lemma 6.7.1 will show that this in fact is true for sufficiently large i , using the fact that $\cup_i \xi(t_i)$ is dense in \mathcal{F} . The idea of the Stability Lemma is simple: if i is sufficiently large so that $\xi(t_i)$ is sufficiently dense in $\mathcal{F} = \mathcal{F}_0$, it follows that any two distinct points on $\Sigma = \Sigma_0$ which lie on the same tie of $\nu(\tau_0)$ are separated by a point of $\xi(t_i)$, which implies that Σ_i embeds in τ_i .

6.7 The Stability Lemma.

We shall formulate the Stability Lemma for a general collection of leaf paths of a certain type. Once the lemma is proved, we will apply it to the collection consisting of all proper saddle connections.

Consider a partial measured foliation \mathcal{G} . A *leaf path* of \mathcal{G} is a continuous path $\theta: [a, b] \rightarrow \text{Supp}(\mathcal{G})$ which can be decomposed as a concatenation of leaf segments of \mathcal{G} . A *legal leaf path* of \mathcal{G} consists of a leaf path θ equipped with a transverse orientation $v_\theta(x)$ defined at the point θx and varying continuously with $x \in [a, b]$, such that θ is perturbable in the direction v_θ to nearby leaf paths. More formally, we require that the map $\theta: [a, b] \rightarrow \text{Supp}(\mathcal{G})$ extends to a map $\Theta: [a, b] \times [0, \epsilon] \rightarrow \text{Supp}(\mathcal{G})$, $\epsilon > 0$, such that for each $t \in [0, \epsilon]$ the map $\Theta \mid [a, b] \times t$ is a leaf path, and for each $x \in [a, b]$ the map $\Theta \mid x \times [0, \epsilon]$ is an embedded segment transverse to \mathcal{G} that represents the transverse orientation $v_\theta(x)$ at the point $p = \theta x = \Theta(x, 0)$. Note that θ is legal if and only if points sufficiently close to θ on the side of v_θ lie in $\text{Supp}(\mathcal{G})$, and if $s = \theta x$ is a singularity of \mathcal{G} , then the restriction of θ to some interval $(x - \eta, x + \eta)$, $\eta > 0$, consists of two separatrices bounding a sector of s , such that v_θ points into that sector. A *legal system* of \mathcal{G} is a map $\theta: I_\theta \rightarrow \text{Supp}(\mathcal{G})$ such that I_θ is a disjoint union of finitely many compact arcs, the restriction of θ to each component of I_θ is a legal leaf path, and the following embedding condition holds: given $x \neq y \in I_\theta$, either $\theta x \neq \theta y$, or $\theta x = \theta y$ and $v_\theta(x) \neq v_\theta(y)$.

Consider now a measured foliation \mathcal{F} and an expansion $\tau_0 \succ \tau_1 \succ \dots$, such that \mathcal{F} is fully carried by each τ_i and each component of \mathcal{F} is infinitely split by $\tau_0 \succ \tau_1 \succ \dots$. We adopt the notation described in Section 6.6. Given a legal system $\theta: I_\theta \rightarrow \text{Supp}(\mathcal{F})$ of \mathcal{F} , when \mathcal{F} is sliced along the separatrix family ξ_j the legal system θ lifts uniquely to a legal system $\theta_j: I_\theta \rightarrow \text{Supp}(\mathcal{F}_j)$. To put it another way, θ pulls back via the partial fulfillment map $r_j: \mathcal{F}_j \rightarrow \mathcal{F}$ to θ_j . Under the carrying map $\mu_j: \mathcal{F}_j \rightarrow \tau_j$, the system θ_j projects to a system of train paths $\mu_j \circ \theta_j: I_\theta \rightarrow \tau_j$.

What we at first might like the Stability Lemma to say is that for any legal system θ of \mathcal{F} , for all sufficiently large j the system of train paths $\mu_j \circ \theta: I_\theta \rightarrow \tau_j$ is injective, but this does not take into account the possibility that two points of I_θ may map to the same point of \mathcal{F} with opposite transverse orientations. Taking this into account, we can now state the Stability Lemma:

Lemma 6.7.1 (Stability Lemma). *Let $\tau_0 \succ \tau_1 \succ \dots$ be a train track expansion of a measured foliation \mathcal{F} that infinitely splits each component of \mathcal{F} , so that \mathcal{F} is fully carried on each τ_i ; we adopt the notation of Section 6.6. For each legal system θ of \mathcal{F} there exists J such that for all $j \geq J$, the following hold:*

- (1) *the system of train paths $\mu_j \circ \theta_j: I_\theta \rightarrow \tau_j$ is at most two-to-one;*
- (2) *for each $x \neq y \in I_\theta$, if $\mu_j \circ \theta_j(x) = \mu_j \circ \theta_j(y)$ then $\theta_j x = \theta_j y$ in $\text{Supp}(\mathcal{F}_j)$. It follows, by definition of a legal leaf system, that the transverse orientations $v_{\theta_j}(x), v_{\theta_j}(y)$ point in opposite transverse directions at the point $\theta_j x = \theta_j y$.*

We refer to the statements (1)–(2) by saying that θ stabilizes in τ_j .

Corollary 6.7.2. *Under the same hypotheses as the Stability Lemma, there exists $J \geq 0$ such that for all $j \geq J$, the set of proper saddle connections of \mathcal{F}_j injects in τ_j under the carrying map $\mu_j: \nu(\tau_j) \rightarrow \tau_j$.*

Proof. Let Σ_i be the set of proper saddle connections of \mathcal{F}_i . Let $\theta: I_\theta \rightarrow \text{Supp}(\mathcal{F}_0)$ be a legal system for \mathcal{F}_0 consisting of a continuous choice of transverse orientation on Σ_0 , and so $x \neq y \in I_\theta$ implies $\theta(x) \neq \theta(y)$. The lifted legal system $\theta_j: I_\theta \rightarrow \text{Supp}(\mathcal{F}_j)$ therefore has the property that $x \neq y \in I_\theta$ implies $\theta_j(x) \neq \theta_j(y)$. Applying the Stability Lemma there exists $J \geq 0$ such that for each $j \geq J$ the system of train paths $\mu_j \circ \theta_j: I_\theta \rightarrow \tau_j$ is injective. But $\Sigma_j \subset \text{image}(\theta_j)$, and so $\mu_j \upharpoonright \Sigma_j$ is injective. \diamond

Proof of the Stability Lemma 6.7.1. If the conclusion is true for $j = J$ then it is clearly true for all $j \geq J$. So we need only find one J for which the conclusion holds.

Claim: without loss of generality, we may assume $\theta: I_\theta \rightarrow \text{Supp}(\mathcal{F})$ has the following property:

- (A) If $x \in I_\theta$ and if $\theta x \notin \partial \text{Supp}(\mathcal{F})$ then there exists $y \neq x \in I_\theta$ such that $\theta x = \theta y$. It follows, from the definition of a legal system, that $v_\theta(x), v_\theta(y)$ point in opposite directions.

In this property, recalling that all singularities of \mathcal{F} are on $\partial \text{Supp}(\mathcal{F})$, it follows that θx is a regular point, and so there are exactly two directions transverse to \mathcal{F} at

x . If property (A) does not already hold, then we can augment θ by throwing in an additional finite number of legal leaf paths so that it does hold. If the augmentation of θ stabilizes in τ_J then clearly θ itself also stabilizes in τ_J , proving the claim.

The legal system $\theta: I_\theta \rightarrow \text{Supp}(\mathcal{F})$ may be extended to a map $\Theta: I_\theta \times [0, \epsilon] \rightarrow \text{Supp}(\mathcal{F})$, $\epsilon > 0$, so that for each $t \in [0, \epsilon]$ the restriction of Θ to each component of $I_\theta \times t$ is a leaf path, and for each $x \in I_\theta$ the restriction of Θ to $x \times [0, \epsilon]$ is an embedded segment transverse to \mathcal{F} which, at the point $\theta x = \Theta(x, 0)$, represents the transverse orientation $v_\theta(x)$. Moreover, the embedding condition in the definition of a legal system guarantees that, by replacing the interval $[0, \epsilon]$ with a smaller interval if necessary, the restriction of Θ to the set $I_\theta \times (0, \epsilon]$ is an embedding with image disjoint from $\text{image}(\theta) = \text{image}(\Theta \mid (I_\theta \times 0))$.

Choose J as follows: since $\bigcup_j \xi_j$ is dense in $\text{Supp}(\mathcal{F})$, there exists J such that for each $x \in I_\theta$ the open transverse segment $\Theta(x \times (0, \epsilon))$ has nonempty intersection with ξ_J .

We prove that θ stabilizes in τ_J .

Suppose by contradiction that conclusion 1 does not hold, so there exist three distinct points $x, y, z \in I_\theta$ such that $\mu_J(\theta_J x) = \mu_J(\theta_J y) = \mu_J(\theta_J z)$ in τ_J . It follows that $\theta_J x, \theta_J y, \theta_J z$ all lie on a single tie of ν_J . Each point of that tie has at most two transverse directions, so at least two of $\theta_J x, \theta_J y, \theta_J z$ are distinct. If all three are distinct then we may permute the notation so that $\theta_J y$ lies between $\theta_J x$ and $\theta_J z$, and so that $v_{\theta_J}(y)$ points into the interior of the segment $\overline{\theta_J x, \theta_J y}$ in the tie. Similarly, if two of $\theta_J x, \theta_J y, \theta_J z$ are the same then we permute the notation so that $\theta_J y = \theta_J z$ and so that $v_{\theta_J}(y)$ points into the interior of the segment $\overline{\theta_J x, \theta_J y}$. Finally, we may assume that the interior of the segment $\overline{\theta_J x, \theta_J y}$ is disjoint from $\text{image}(\theta_J)$, for if not then we may replace $\theta_J x$ by the point of $\text{image}(\theta_J)$ lying on the segment which is closest to $\theta_J y$. Letting $\delta = \overline{\theta_J x, \theta_J y}$ and $\eta = y$, the pair δ, η satisfies the following property:

- (*) δ is a subsegment of a tie of ν_J ; $\eta \in I_\theta$; $\theta_J(\eta) \in \partial\delta \subset \text{image}(\theta_J)$; $\text{int}(\delta) \cap \text{image}(\theta_J) = \emptyset$; and the transverse direction $v_{\theta_J}(\eta)$ points into δ at the point $\theta_J(\eta)$.

But the existence of such a pair δ, η leads to a contradiction. To see why, the image $r_J(\delta) \subset \nu_0$ satisfies the following properties: $r_J(\delta)$ is a leaf segment of \mathcal{F}_0^v ; $\theta(\eta) \in \partial r_J(\delta) \subset \text{image}(\theta)$; $\text{int}(r_J(\delta)) \cap \text{image}(\theta) = \emptyset$; and the transverse direction $v_\theta(\eta)$ points into $r_J(\delta)$ at the point $\theta(\eta)$. Now consider the segment $\Theta(\eta \times [0, \epsilon])$, which starts at the point $\theta(\eta)$ and follows the leaf of \mathcal{F}_0^v in the direction of $v_\theta(\eta)$, stopping before reaching another point of $\text{image}(\theta)$. It follows that $\Theta(\eta \times (0, \epsilon)) \subset \text{int}(r_J(\delta))$. But also, $\xi_J \cap \text{int}(r_J(\delta)) = \emptyset$, because $r_J(\delta)$ lifts to the connected vertical segment δ of ν_J . We therefore have $\Theta(\eta \times (0, \epsilon)) \cap \xi_J = \emptyset$, contradicting the choice of J .

Now we prove conclusion 2. Consider two points $x \neq y \in I_\theta$ with $\mu_J(\theta_J x) = \mu_J(\theta_J y)$, and suppose that $\theta_J x \neq \theta_J y$. Let β be the segment of the tie of ν_J with endpoints $\theta_J x, \theta_J y$. By assumption (A) above, by changing x if necessary we may assume that $v_{\theta_J}(x)$ points into β . Then, taking $\delta = \beta$ and $\eta = x$, the pair δ, η satisfies (*), and we obtain a contradiction as before.

◇

6.8 Sufficiency of the iterated rational killing criterion: conclusion.

Continuing with the notation of Section 6.6, the corollary to the Stability Lemma 6.7.1 says that for sufficiently large J , each $j \geq J$ satisfies the property that the set Σ_j of proper saddle connections of \mathcal{F}_j embeds in τ_j under the carrying map $\mu_j: \nu(\tau_j) \rightarrow \tau_j$. It follows immediately that the image of Σ_j is a pairwise disjoint family of \mathcal{C} -splitting arcs of τ_j .

Let

$$G_j = \text{Cl}(S - \text{Supp}(\mathcal{F}_j)) \cup (\text{proper saddle connections of } \mathcal{F}_j)$$

Let Γ_j denote the graph obtained from G_j by collapsing to a point each component of $\text{Cl}(S - \text{Supp}(\mathcal{F}_j))$, where the vertices of Γ_j correspond bijectively with the components of $\text{Cl}(S - \text{Supp}(\mathcal{F}_j))$, and the edges correspond bijectively with the proper saddle connections of \mathcal{F}_j . Each component of $\text{Cl}(S - \text{Supp}(\mathcal{F}_j))$ has zero or one puncture, and the corresponding vertex of Γ_j is labelled as a nonpuncture or a puncture vertex, respectively. The partial fulfillment map $\text{Supp}(\mathcal{F}_j) \rightarrow \text{Supp}(\mathcal{F}_{j-1})$ induces a label preserving isomorphism $\Gamma_j \rightarrow \Gamma_{j-1}$.

Define a *cycle* in the graph Γ_j to be either an immersed closed curve or an immersed arc whose endpoints are located at distinct puncture vertices. Since \mathcal{F}_j is not arational, it follows that cycles exist. Each cycle of minimal length is embedded. Let c_J be a cycle of minimal length in Γ_J , and so c_J is either an embedded arc whose interior contains no puncture vertex, or c_J is an embedded circle containing at most one puncture vertex. There exist up to isotopy either one or two splitting cycles of τ_J which project to c_J under the quotient map $G_J \rightarrow \Gamma_J$. In the case when c_J is a circle containing exactly one puncture vertex, corresponding to a certain component of $\mathcal{C}(S - F_J)$ containing a puncture p , then there are up to isotopy two splitting cycles projecting to c_J , and they differ by going around p on opposite sides. In all other cases there is up to isotopy a unique splitting cycle projecting to c_J . For each $j \geq J$ we may inductively choose the cycle c_j in Γ_j so that it corresponds to c_{j-1} under the isomorphism $\Gamma_j \rightarrow \Gamma_{j-1}$. The splitting cycles of τ_j corresponding to c_j are in bijective correspondence with the splitting cycles of τ_{j-1} corresponding to c_{j-1} , in such a way that corresponding splitting cycles are related by descent. Thus, the choice of corresponding splitting cycles for c_J, c_{J+1}, \dots produces a sequence splitting

cycles $\gamma_J, \gamma_{J+1}, \dots$ for $\tau_J, \tau_{J+1}, \dots$, respectively, such that γ_J survives forever with descendant γ_j in τ_j .

This shows that $\tau_0 \succ \tau_1 \succ \dots$ fails the iterated splitting cycle killing criterion, completing the proof of the Arational Expansion Theorem.

Certain additional details emerge from the proof of sufficiency of the iterated rational killing criterion, which can be used to produce a criterion that, while a priori stronger than the iterated rational killing criterion, is in fact equivalent.

For instance, if γ_j is the sequence of splitting cycles produced at the end of the proof, then γ_j has the following property: if γ_j is a circle then there is at most one punctured component Q of $\mathcal{C}(S - \tau_j)$ that γ_j meets; and if γ_j is a proper line then there are exactly two punctured components of $\mathcal{C}(S - \tau_j)$ that γ_j meets.

One can also use the Stability Lemma to prove that for sufficiently large j there is an infinitely surviving splitting cycle γ_j of τ_j such that for each component α of $\gamma_j \cap \tau_j$, the \mathbb{C} -splitting arc α contains at most two sink branches of τ_j .

We may therefore take the two properties of γ_j described in the previous two paragraphs, add each of these properties to the definition of a splitting cycle, and thereby get a stronger statement of the iterated splitting cycle killing criterion and hence of the iterated rational killing criterion, and yet this stronger statement is equivalent to the original statement. The advantage of this is that, with stronger conditions on splitting cycles, there are fewer splitting cycles that need to be considered, for example in the pseudo-Anosov algorithms of Section 9.

7 Stable equivalence of expansion complexes

The goal of this section is to establish dictionary entries between stable equivalence phenomena for continued fraction expansions of irrational numbers and for train track expansions of arational measured foliations. As explained in the introduction, given an arational measured foliation \mathcal{F} and a train track τ canonically carrying \mathcal{F} , the lack of naturality of a train track expansion of \mathcal{F} based at τ means that a naive generalization of stable equivalence on the torus will not work. Instead we give two versions of stable equivalence which get around this lack of naturality in different ways.

Section 7.1 describes stable equivalence of “one cusp train track expansions”, a natural finite set of train track expansions of \mathcal{F} based at τ . The advantage of one cusp expansions are their finiteness, the disadvantage is that they are very far from encompassing all expansions of \mathcal{F} based at τ .

In Section 7.2 we package all train track expansions of \mathcal{F} based at τ into a single object, the *expansion complex* of \mathcal{F} based at τ , and we describe stable equivalence of expansion complexes. Expansion complexes are based on the results of Section 4. The expansion complex of \mathcal{F} based at τ would not be worth much if it were just the mere *set* of train track expansions of \mathcal{F} based at τ , but instead it has a rich structure which explains how the different train track expansions of \mathcal{F} based at τ are related to each other, and in particular it clarifies the special role played by one cusp expansions. In Theorem 7.2.2 we describe the structure of the expansion complex. Roughly speaking, the expansion complex of \mathcal{F} based at τ is a higher dimensional version of a train track expansion, where the dimension N is equal to the number of cusps of τ , or equivalently the number of infinite separatrices of \mathcal{F} . Instead of requiring the N separatrices to grow in synchrony as we did in Proposition 4.4.1 where train track expansions were characterized in terms of “expanding separatrix families”, we allow them to grow independently, filling out a parameter space $\Xi = [0, \infty)^N$, the “space of separatrix families” as it was dubbed in Section 4.3. The space Ξ is decomposed into cells labelled by the isotopy type of a train track associated to each cell, these cells form a polyhedral cell decomposition, and the face relation among cells can be described by an elementary combinatorial relation between train tracks. We thus obtained a labelled cell structure on Ξ , which is the expansion complex of \mathcal{F} based at τ . The one cusp expansions emerge from this structure in a very natural way: there are simply the expansions obtained by focussing on each coordinate axis of $\Xi = [0, \infty)^N$.

The main result of Section is Theorem 7.2.3, which is the marked version of stable equivalence for expansion complexes. As a consequence to the Stable Equivalence Theorem, Corollary 7.2.6 explains how to use expansion complexes to classify arational measured foliations on S up to the relation of unmeasured equivalence (see

Section 5.1). The proof of this corollary uses the Expansion Convergence Theorem 5.1.1.

Section 7.3 reduces the proof of Theorem 7.2.3 to the structure theorem for expansion complexes, Theorem 7.2.2, and the latter is proved in Section 7.4.

The unmarked version stable equivalence for expansion complexes, which classifies measured foliations up to topological equivalence is stated and proved in Section 7.7, Theorem 7.7.1.

7.1 Stable equivalence of one cusp train track expansions

TO DO:

- Check that these notions are not already defined elsewhere.

In this section we describe one cusp train track expansions, and we state the marked and unmarked versions of stable equivalence for one cusp expansions. We shall also show how the unmarked version follows from the marked version. The proof of the marked version will be given later, as a corollary to stable equivalence of expansion complexes.

Given a generic train track τ and a cusp v of τ , the ordered pair (τ, v) is called a *one cusp train track*. Let s be the switch at which v is located. Consider the branch b of τ which is incident to s on the one-ended side of s , and so s is an inward pointing switch of b . The branch b must therefore be either a sink branch or a transition branch, and in the latter case the ends of b are at distinct switches, since τ is recurrent. If b is a transition branch let $\tau \succcurlyeq \tau'$ be the slide move along b , let $v' \in \text{cusps}(\tau')$ be the cusp corresponding to v via the homotopic carrying $\tau \xrightarrow{H} \tau'$, and we say that the relation $(\tau, v) \succcurlyeq (\tau', v')$ is a *one cusp slide*; notice that this is *not* a symmetric relation, unlike ordinary slides. If on the other hand b is a sink branch, choose $d \in \{L, R\}$, let $\tau \succ \tau'$ be an elementary splitting of parity d along b , let $v' \in \text{cusps}(\tau')$ be the cusp corresponding to v via the homotopic carrying $\tau \xrightarrow{H} \tau'$, and we say that $(\tau, v) \succ (\tau', v')$ is a *one cusp elementary splitting* of parity d .

Next we describe how wide splittings arise naturally in the one cusp setting. Given a one cusp train track (τ, v) , consider a train path which starts at s and moves away in the direction of the switch orientation. As this train path continues it may encounter other switches, and as long as these switches continue to point in the same direction, the train path continues uniquely. Since τ is recurrent, the train path must eventually encounter a switch $s' \neq s$ that points in the opposite direction. Let α be the unique embedded train path thus described, starting from s that stops at s' , and note that α is the shortest splitting arc of τ that has s as one of its endpoints; all interior switches of α point in the same direction as the switch

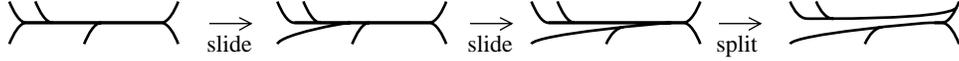


Figure 20: If $(\tau, v) \succ (\tau', v')$ is a one cusp splitting along α , and if α contains $k + 1$ switches including its endpoints, $k \geq 1$, then $(\tau, v) \succ (\tau', v')$ factors into $k - 1$ one cusp slides followed by a one cusp elementary splitting.

s. Choose $d \in \{L, R\}$ and let $\tau \succ \tau'$ be the wide splitting of parity d along α . Let $v' \in \text{cusps}(\tau')$ be the cusp corresponding to v via the homotopic carrying $\tau \xrightarrow{H} \tau'$. We say that $(\tau, v) \succ (\tau', v')$ is a *one cusp splitting* of parity d . Note that the parity d and the isotopy class of the pair (τ, v) completely determine the isotopy class of (τ', v') .

Each one cusp splitting $(\tau, v) \succ (\tau', v')$ factors uniquely (up to isotopy) into zero or more one cusp slides followed by a one cusp elementary splitting (see Figure 20). To see this, let α be the splitting arc along which $(\tau, v) \succ (\tau', v')$ is defined. If α is a sink branch then $(\tau, v) \succ (\tau', v')$ is an elementary one cusp splitting. Otherwise, let $s = s_0, s_1, \dots, s_k = s', k \geq 2$, be the switches in $\text{int}(\alpha)$ in order. Then we have a factorization

$$(\tau, v) = (\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots \succ (\tau_{k-1}, v_{k-1}) \succ (\tau_k, v_k) = (\tau', v')$$

into $k - 1$ one cusp slides followed by a one cusp elementary splitting.

The following is just a special case of the splitting inequalities 3.13.1, where the case of a central splitting is precluded by the assumption that τ canonically carries \mathcal{F} :

Fact 7.1.1. *If a train track τ canonically carries an arational measured foliation \mathcal{F} , and if $v \in \text{cusps}(\tau)$, then there exists a unique $d \in \{L, R\}$ such that the one cusp splitting $(\tau, v) \succ (\tau', v')$ of parity d results in a train track τ' that canonically carries \mathcal{F} .*

A *one cusp splitting sequence* based at τ_0 is splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ such that, for some sequence of cusps $v_i \in \text{cusps}(\tau_i)$, $i \geq 0$, each of the splittings $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ is a one cusp splitting. The *elementary factorization* of a one cusp splitting sequence is the unique sequence obtained by factoring each one cusp splitting into one cusp slides followed by a one cusp elementary splitting.

Lemma 7.1.2. *If S is not the punctured torus, and if $\tau_0 \succ \tau_1 \succ \dots$ is a one cusp splitting sequence that satisfies the subtrack killing criterion, then there is a unique sequence $v_i \in \text{cusps}(\tau_i)$, $i \geq 0$, such that each $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ is a one cusp splitting. In particular, uniqueness of the sequence v_i holds if $\tau_0 \succ \tau_1 \succ \dots$ is an expansion of an arational measured foliation.*

Sketch of proof. In fact what is true is that if there is a distinct sequence $v'_i \in \text{cusps}(\tau_i)$ such that each $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ is a one cusp splitting, then there is a surviving subtrack $\sigma_0 \subset \tau_0$ whose support is either an annulus or a one-holed torus.

Since the sequences $(v_i)_{i \geq 0}$ and $(v'_i)_{i \geq 0}$ are distinct, and since both of these sequences respects all of the bijections $\text{cusps}(\tau_i) \leftrightarrow \text{cusps}(\tau_{i+1})$, it follows that $v_i \neq v'_i$ for all $i \geq 0$. Letting α_i be the splitting arc of τ_i along which $\tau_i \succ \tau_{i+1}$ is defined, it follows that v_i, v'_i are located at opposite ends of α_i , pointing into the interior of α_i . Let $\beta_i \subset \tau_i$ be the post-splitting arc of $\tau_i \succ \tau_{i+1}$. By examining Figure 14 (see also Figure 13) it follows that $\beta_i \cup \alpha_i$ is an embedded closed curve in τ_i denoted γ_i . If the parity of $\tau_i \succ \tau_{i+1}$ is constant then the curves γ_i are all isotopic, finishing the proof. Suppose that the parity is not constant, and by truncating we may assume that $\tau_0 \succ \tau_1 \succ \tau_2$ have opposite parity. In this case one can check that γ_1 and γ_2 have intersection number 1, and the smallest subtrack σ_1 of τ_1 that carries both γ_1 and γ_2 has support a one-holed torus. Moreover, σ_1 survives the splitting $\tau_1 \succ \tau - 2$ with descendant $\sigma_2 \subset \tau_2$. An inductive argument shows that σ_1 survives forever. \diamond

Consider an arational measured foliation \mathcal{F} , which we take to be a canonical model. Let τ_0 be a train track that canonically carries \mathcal{F} . A *one cusp train track expansion* of \mathcal{F} based at τ_0 is a one cusp splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ which is an expansion of \mathcal{F} based at τ_0 ; we say that v_0 is the *starting cusp* for this expansion. As a consequence of Fact 7.1.1, for any $v_0 \in \text{cusps}(\tau_0)$ there exists a unique one cusp expansion of \mathcal{F} based at τ_0 with starting cusp v_0 , and hence unique cusps $v_i \in \text{cusps}(\tau_i)$ so that each $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ is a one cusp splitting. For this reason we shall interchangeably write $\tau_0 \succ \tau_1 \succ \dots$ or $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ to represent a one cusp expansion.

We have the following obvious naturality property for one cusp expansions:

Lemma 7.1.3. *Let τ be a generic train track canonically carrying an arational measured foliation \mathcal{F} , and choose $v \in \text{cusps}(\tau)$. Any homeomorphism ϕ takes the one cusp expansion of \mathcal{F} based at τ with initial cusp v to the one cusp expansion of $\phi(\mathcal{F})$ based at $\phi(\tau)$ with initial cusp $\phi(v)$.* \diamond

Notice as a consequence of Lemma 7.1.2, if S is not the once-punctured torus, any two one cusp expansions of \mathcal{F} , based at the same train track τ , and with distinct starting cusps, are distinct.

Consider now an arational measured foliation \mathcal{F} canonically carried on a train track τ_0 . Fix a tie bundle ν_0 over τ_0 with vertical measured foliation \mathcal{F}_0 , and we assume that we have a carrying bijection $\mathcal{F} \hookrightarrow \nu_0$. Recall that Proposition 4.4.1 shows how to associate, to each complete, expanding separatrix family of \mathcal{F} , a train

track expansion of \mathcal{F} based at τ_0 . By following through the construction in the proof of Proposition 4.4.1, we easily obtain:

Lemma 7.1.4. *With the above notation, given $v \in \text{cusps}(\tau_0)$, let ℓ be the infinite separatrix of \mathcal{F} corresponding to v , and let ξ_t , $t \geq 0$, be the growing separatrix family such that ξ_t is the initial segment of ℓ of \mathcal{F}_v transverse measure t . Then the elementary move sequence of the growing separatrix family ξ_t , $t \geq 0$, coincides with the elementary factorization of the one cusp expansion of \mathcal{F} with starting cusp v . \diamond*

As a consequence of this lemma, starting from the elementary move sequence $\tau_0 \succ \tau_1 \succ \dots$ associated to ξ_t , $t \geq 0$, one can recover the one cusp expansion of \mathcal{F} based at τ_0 with starting cusp v , in the following manner. Let $v_i \in \text{cusps}(\tau_i)$ corresponds to $v \in \text{cusps}(\tau_0)$ via the homotopic carrying $\tau_0 \xrightarrow{\text{H}} \tau_i$. Define a subsequence $0 = n_0 < n_1 < \dots$ inductively so that $\tau_{n_{i-1}} \succ \tau_{n_{i-1}} + 1 \succ \dots \succ \tau_{n_i}$ consists of zero or more one cusp slides followed by a one cusp elementary splitting. It follows that $\tau_{n_0} \succ \tau_{n_1} \succ \dots$ is the one cusp expansion of \mathcal{F} based at τ_0 with starting cusp v .

Now we are almost ready to state the main result of this section. Given a canonical, arational measured foliation \mathcal{F} and two train tracks τ, τ' that canonically carry \mathcal{F} there are bijections $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\mathcal{F}) \leftrightarrow \text{cusps}(\tau')$. Composition yields a bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$ that depends naturally on the triple $(\tau, \mathcal{F}, \tau')$, called the *cusp bijection via separatrices of \mathcal{F}* . Notice we are *not* saying that the bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$ depends naturally on just the pair (τ, τ') alone.

Theorem 7.1.5 (One cusp stable equivalence, marked version). *Given an arational measured foliation \mathcal{F} and two train tracks τ, τ' canonically carrying \mathcal{F} , if $v \in \text{cusps}(\tau)$ and $v' \in \text{cusps}(\tau')$ correspond via separatrices of \mathcal{F} , then the one cusp expansion of \mathcal{F} based at τ with starting cusp v is stably equivalent to the one cusp expansion of \mathcal{F} based at τ' with starting cusp v' .*

The proof will be given as a corollary to the main Stable Equivalence Theorem 7.2.3. For now we describe several corollaries, including the unmarked version of one cusp stable equivalence.

By applying Lemma 7.1.4 we obtain an explicit description of the bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$: it is the composition

$$\text{cusps}(\tau) \leftrightarrow s(\mathcal{F}) \leftrightarrow \text{cusps}(\tau')$$

where $s(\mathcal{F})$ is the set of infinite separatrices of \mathcal{F} , and the bijections between $s(\mathcal{F})$ and the sets $\text{cusps}(\tau)$, $\text{cusps}(\tau')$ are induced by carrying bijections from \mathcal{F} to tie bundles ν, ν' over τ, τ' , respectively.

Our first corollary gives a classification of arational measured foliations up to unmeasured equivalence, in terms of stable equivalence classes of one cusp expansions:

Corollary 7.1.6. *Given arational measured foliations $\mathcal{F}, \mathcal{F}'$ canonically carried by train tracks τ, τ' , the following are equivalent:*

- (1) $\mathcal{PMF}(\mathcal{F}) = \mathcal{PMF}(\mathcal{F}')$.
- (2) *There exist cusps $v \in \text{cusps}(\tau)$, $v' \in \text{cusps}(\tau')$ such that the one cusp expansion of \mathcal{F} based at τ with starting cusp v is stably equivalent to the one cusp expansion of \mathcal{F}' based at τ' with starting cusp v' .*
- (3) *There exists a bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$ such that if $v \in \text{cusps}(\tau)$ and $v' \in \text{cusps}(\tau')$ correspond, then the one cusp expansion of \mathcal{F} based at τ with starting cusp v is stably equivalent to the one cusp expansion of \mathcal{F}' based at τ' with starting cusp v' .*

Proof. Obviously (3) \implies (2). The implication (2) \implies (1) is an immediate consequence of the Expansion Convergence Theorem 5.1.1. To prove (1) \implies (3), note that unmeasured equivalence of $\mathcal{F}, \mathcal{F}'$ implies that each expansion of \mathcal{F} is also an expansion of \mathcal{F}' and vice versa, and so (3) follows by an application of Theorem 7.1.5. \diamond

Corollary 7.1.6 gives a complete classification of arational measured foliations up to topological equivalence. However, there are two disadvantages of this result.

First, given a one cusp splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, how do we tell whether it is a one cusp expansion of an arational measured foliation? Answer: this happens if and only if $\tau_0 \succ \tau_1 \succ \dots$ satisfies the canonical killing criterion. In Section 11.1 we will show that for a one cusp splitting sequence, the canonical killing criterion has a particularly nice form.

Second, given two one cusp splitting sequences $\tau_0 \succ \tau_1 \succ \dots$ and $\tau'_0 \succ \tau'_1 \succ \dots$, how do we tell whether they are one expansions of the same arational measured foliation? This question is quite subtle, and we must wait for the general Stable Equivalence Theorem 7.2.3 for an answer; see the remarks after the statement of that theorem.

Unmarked stable equivalence of one cusp expansions. Recall from Section 3.15 that two one cusp train tracks (τ, v) , (τ', v') are combinatorially equivalent if there exists $\phi \in \text{Homeo}_+$ such that $\phi(\tau, v) = (\tau', v')$.

Consider a pair of one cusp splitting sequences $\tau_0 \succ \tau_1 \succ \dots$ and $\tau'_0 \succ \tau'_1 \succ \dots$. Choose cusp sequences $v_i \in \text{cusps}(\tau_i)$, $v'_i \in \text{cusps}(\tau'_i)$ so that $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ and $(\tau'_i, v'_i) \succ (\tau'_{i+1}, v'_{i+1})$ are one cusp splittings. These two splitting sequences are said to be *unmarked stably equivalent* if there exist integers $a, b \geq 0$ such that for all $i \geq 0$ the following hold:

- The one cusp train tracks $(\tau_{a+i}, v_{a+i}), (\tau'_{b+i}, v'_{b+i})$ are combinatorially equivalent.
- The splittings $\tau_{a+i} \succ \tau_{a+i+1}$ and $\tau'_{b+i} \succ \tau'_{b+i+1}$ have the same parity.

Theorem 7.1.7 (One cusp stable equivalence, unmarked version). *Given two arational measured foliations $\mathcal{F}, \mathcal{F}'$, the following are equivalent:*

- (1) *There exists a mapping class ϕ such that $\mathcal{PMF}(\phi(\mathcal{F})) = \mathcal{PMF}(\mathcal{F}')$.*
- (2) *For any train tracks τ, τ' that canonically carry $\mathcal{F}, \mathcal{F}'$, there exists cusps v, v' of τ, τ' , respectively, such that the one cusp expansion of \mathcal{F} based at τ with starting cusp v is unmarked stably equivalent to the one cusp expansion of \mathcal{F}' based at τ' with starting cusp v' .*
- (3) *For any train tracks τ, τ' that canonically carry $\mathcal{F}, \mathcal{F}'$, respectively, there is a bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$, such that if $v \in \text{cusps}(\tau)$ and $v' \in \text{cusps}(\tau')$ correspond, then the one cusp expansion of \mathcal{F} based at τ with starting cusp v is unmarked stably equivalent to the one cusp expansion of \mathcal{F}' based at τ' with starting cusp v' .*

Proof of Theorem 7.1.7, given Theorem 7.1.5. The implication (3) \implies (2) is obvious.

Next we prove (1) \implies (3). Let $\mathcal{F}'' = \phi(\mathcal{F})$, and so $\mathcal{PMF}(\mathcal{F}'') = \mathcal{PMF}(\mathcal{F}')$. Let $\tau'' = \phi(\tau)$. Applying Corollary 7.1.6, it follows that there is a bijection $\text{cusps}(\tau'') \leftrightarrow \text{cusps}(\tau')$, so that if $v'' \in \text{cusps}(\tau'')$ corresponds to $v' \in \text{cusps}(\tau')$ then the one cusp expansion of \mathcal{F}'' based at τ'' with starting cusp v'' is unmarked stably equivalent to the one cusp expansion of \mathcal{F}' based at τ' with starting cusp v' . Next note that the action of a mapping class on a one-cusp splitting sequence preserves the combinatorial type of each train track and also preserves each parity, and so induces an unmarked stable equivalence. It follows that under the bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau'')$ induced by ϕ , if $v \in \text{cusps}(\tau)$ corresponds to $v'' \in \text{cusps}(\tau'')$ then the expansion of \mathcal{F} based at τ with starting cusp v is unmarked stably equivalent to the expansion of \mathcal{F}'' based at τ'' with starting cusp v'' . We thus obtain by composition a bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau'') \leftrightarrow \text{cusps}(\tau')$ that satisfies (3).

Finally we prove (2) \implies (1). We have one cusp expansions

$$\begin{aligned} \tau &= \tau_0 \succ \tau_1 \succ \cdots \\ \tau' &= \tau'_0 \succ \tau'_1 \succ \cdots \end{aligned}$$

of $\mathcal{F}, \mathcal{F}'$, respectively, with starting cusps $v_0 \in \text{cusps}(\tau_0)$ and $v'_0 \in \text{cusps}(\tau'_0)$. We have unique cusp sequences $v_i \in \text{cusps}(\tau_i), v'_i \in \text{cusps}(\tau'_i)$ so that $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$

and $(\tau'_i, v'_i) \succ (\tau'_{i+1}, v'_{i+1})$ are one cusp splittings. Choose integers $a, b \geq 0$, so that for each $i \geq 0$ there is a combinatorial equivalence $\theta_i: (\tau_{a+i}, v_{a+i}) \rightarrow (\tau'_{b+i}, v'_{b+i})$, and so that the parities of the splittings $\tau_{a+i} \succ \tau_{a+i+1}$ and $\tau'_{b+i} \succ \tau'_{b+i+1}$ are identical; let $d_i \in \{L, R\}$ be this parity.

We claim all of the combinatorial equivalences θ_i for $i \geq 0$ are in the same mapping class Φ . To prove this, notice that splittings are natural, and so the train track $\theta_i(\tau_{a+i+1})$ is obtained from $\tau'_{b+i} = \theta_i(\tau_{a+i})$ by a splitting of parity d_i . This implies that $\theta_i(\tau_{a+i+1}) = \tau'_{b+i+1} = \theta_{i+1}(\tau_{a+i+1})$. Moreover, the cusp bijections induced by splittings are natural, so the bijection $\text{cusps}(\tau_{a+i}) \leftrightarrow \text{cusps}(\tau_{a+i+1})$ is taken by θ_i to the bijection $\text{cusps}(\tau'_{b+i}) \leftrightarrow \text{cusps}(\tau'_{b+i+1})$, implying that $\theta_i(v_{a+i+1}) = v'_{b+i+1}$. The homeomorphism $\theta_{i+1}^{-1}\theta_i$ therefore stabilizes the isotopy class of the one cusp train track $(\tau_{a+i+1}, v_{a+i+1})$. Applying Corollary 3.15.3 it follows that θ_{i+1} is isotopic to θ_i .

Let Θ be the mapping class of each θ_i , $i \geq 0$. Letting $\tau''_i = \Theta(\tau_i)$, $v''_i = \Theta(v_i)$, it follows that $(\tau''_0, v''_0) \succ (\tau''_1, v''_1) \succ \dots$ is a one cusp expansion of $\Theta(\mathcal{F})$ which is stably equivalent to $(\tau'_0, v'_0) \succ (\tau'_1, v'_1) \succ \dots$. Applying Theorem 7.1.5, we have $\mathcal{PMF}(\Theta(\mathcal{F})) = \mathcal{PMF}(\mathcal{F}')$. \diamond

7.2 Expansion complexes of arational measured foliations

In this subsection we describe expansion complexes and state two results, Theorem 7.2.2 which describes some of their structure, and Theorem 7.2.3 which describes (marked) stable equivalence. The proofs of these theorems come in the following subsections. We shall delay until Section 7.7 the statement of unmarked stable equivalence of expansion complexes, which is somewhat more technical to state than the unmarked stable equivalence of one cusp expansions.

Consider an arational measured foliation \mathcal{F} and fix a train track τ_0 that carries \mathcal{F} . Assuming that τ_0 is generic and that τ_0 canonically carries \mathcal{F} , we shall define the expansion complex of \mathcal{F} based at τ_0 . It is possible to drop the assumption that τ_0 is generic, but this would complicate the proofs without changing the statements in any significant way. It is also possible to drop the assumption that τ_0 carries \mathcal{F} canonically, but in this case the statements would be significantly more complicated and the expansion complex less well behaved; we will briefly describe some of the features of noncanonical expansion complexes in Section 7.6.

Since τ_0 is generic there is an isotopically unique tie bundle ν_0 over τ_0 , and each singular tie of ν_0 is generic. Let \mathcal{F}_ν denote the foliation of ν_0 by ties. Since τ_0 canonically carries \mathcal{F} , we may choose \mathcal{F} in its equivalence class to be a canonical model with a carrying bijection $\mathcal{F} \xrightarrow{c} \nu_0$. Let \mathcal{I} be the set of proper separatrix germs of \mathcal{F} , in one-to-one correspondence with the cusps of ν_0 . The components of any separatrix family ξ of \mathcal{F} are indexed by \mathcal{I} , where ξ_i is the component representing

the germ $i \in \mathcal{I}$. Recall that a component ξ_i is *degenerate* if ξ_i reduces to a cusp of ν_0 , and otherwise ξ_i is *nondegenerate*.

The underlying topological space of the expansion complex based at τ_0 is the space $\Xi = \Xi(\mathcal{F}, \tau_0)$ of separatrix families of \mathcal{F} , equipped with a natural stratification labelled by subsets $\mathcal{J} \subset \mathcal{I}$, where

$$\Xi_{\mathcal{J}} = \{\xi \in \Xi \mid \forall i \in \mathcal{I}, \xi_i \text{ is nondegenerate} \implies i \in \mathcal{J}\}$$

In other words, $\Xi_{\mathcal{J}}$ is the closure of the set of $\xi \in \Xi$ such that ξ_i is nondegenerate if and only if $i \in \mathcal{J}$. Choosing a positive transverse Borel measure on \mathcal{F}_v , since \mathcal{F} has no proper saddle connections we obtain a homeomorphism $\Xi \approx [0, \infty)^{\mathcal{I}}$, taking the stratum $\Xi_{\mathcal{J}}$ to the coordinate subspace

$$[0, \infty)^{\mathcal{J}} = \{x \in [0, \infty)^{\mathcal{I}} \mid x_i > 0 \implies i \in \mathcal{J}\}$$

Recall also that if \mathcal{F}' is isotopic to \mathcal{F} , with separatrix germs \mathcal{I}' and space of separatrix families Ξ' , then there is a natural bijection $\mathcal{I} \leftrightarrow \mathcal{I}'$, inducing a bijection of subsets $\mathcal{J} \leftrightarrow \mathcal{J}'$, such that any isotopy $f: \mathcal{F} \rightarrow \mathcal{F}'$ induces a homeomorphism $\Xi \rightarrow \Xi'$ taking $\Xi_{\mathcal{J}}$ to $\Xi'_{\mathcal{J}'}$.

Recall that to each $\xi \in \Xi$ there is associated a tie bundle $\nu(\xi)$ over a train track $\sigma(\xi)$: the tie bundle $\nu(\xi)$ is obtained by slicing ν_0 along ξ and pulling back ties of ν_0 to get ties of $\nu(\xi)$; and the train track $\sigma(\xi)$ is obtained by collapsing ties of $\nu(\xi)$ to get $\sigma(\xi)$. There are natural bijections

$$\mathcal{I} \leftrightarrow \{\text{components of } \xi\} \leftrightarrow \{\text{cusps of } \nu(\xi)\} \leftrightarrow \{\text{cusps of } \sigma(\xi)\}$$

The second bijection is defined by the property that each cusp of $\nu(\xi)$ is taken by the partial fulfillment map $\nu(\xi) \rightarrow \nu_0$ to the boundary point of the corresponding component of ξ . The third bijection is lifting of cusps, explained at the end of Section 3.4. Given $i \in \mathcal{I}$ we sometimes use c_i to denote the corresponding cusp of $\sigma(\xi)$.

A cusp of $\sigma(\xi)$ is said to be degenerate or nondegenerate as it corresponds with a component of ξ the same type. The ordered pair $(\sigma(\xi), D)$, consisting of the train track $\sigma(\xi)$ and the set D of degenerate cusps, is called the *decorated train track* associated to ξ .

Abstracting the above concepts, a *decorated train track* on S is an ordered pair (σ, D) consisting of a train track σ and a subset $D \subset \text{cusps}(\sigma)$ called the set of *degenerate cusps* (whose complement $\text{cusps}(\sigma) - D$ is called the set of *nondegenerate cusps*), such that at each switch s of σ there is at most one degenerate cusp. If there is a degenerate cusp located at the switch s then we say that s is a *degenerate switch*, otherwise s is a *nondegenerate switch*. Often we will suppress the notation D , using σ to refer to a decorated train track (σ, D) .

The cell structure of the expansion complex. Define an equivalence relation on Ξ where two separatrix families $\xi, \xi' \in \Xi$ are equivalent if the decorated train tracks $\sigma(\xi), \sigma(\xi')$ are isotopic. Each equivalence class is labelled with the isotopy class of a decorated train track σ , isotopic to $\sigma(\xi)$ for each ξ in the equivalence class. Distinct equivalence classes have distinct labels up to isotopy. We define the *expansion complex* of \mathcal{F} based at τ_0 to be the stratified space Ξ , decomposed into its equivalence classes, each equivalence class labelled with the appropriate isotopy class of decorated train track. The equivalence class of Ξ labelled by the isotopy class of the decorated train track σ is denoted $\overset{\circ}{c}(\sigma)$; as we shall see, this is an open polyhedral cell. We have the following obvious naturality property for expansion complexes:

Lemma 7.2.1. *Let τ be a generic train track canonically carrying a canonical, arational measured foliation \mathcal{F} . Let Ξ be the expansion complex of \mathcal{F} based at τ and \mathcal{I} the set of infinite separatrix germs of \mathcal{F} . Fix a homeomorphism f , let Ξ' be the expansion complex of $f(\mathcal{F})$ based at $f(\tau)$, and let \mathcal{I}' be the set of separatrix germs. Then f induces a bijection $f: \mathcal{I} \rightarrow \mathcal{I}'$ and a homeomorphism $f: \Xi \rightarrow \Xi'$ such that $f(\Xi_{\mathcal{J}}) = \Xi'_{f(\mathcal{J})}$ for each $\mathcal{J} \subset \mathcal{I}$, and such that for each decomposition element $\overset{\circ}{c}(\sigma)$ of Ξ we have $f(\overset{\circ}{c}(\sigma)) = \overset{\circ}{c}(f(\sigma))$. \diamond*

Define a *linear cell decomposition* of Ξ to be a CW-decomposition of $\Xi \approx [0, \infty)^{\mathcal{I}}$ such that the closure of each open k -cell is a finite sided, convex polyhedron of dimension k called a *closed k -cell*, the boundary of each closed k -cell is a union of lower dimensional closed cells, and each stratum $\Xi_{\mathcal{J}}$ is a subcomplex. It follows that the intersection of any two closed cells is a union of closed cells, possibly empty. Also, if c, d are closed cells of X and $\text{int}(d) \cap \partial c \neq \emptyset$ then $d \subset \partial c$. Given two closed cells c, c' of Ξ , we define c to be a *face* of c' if $c \subset c'$.

Theorem 7.2.2 (Cell Structure Theorem). *Given an arational measured foliation \mathcal{F} canonically carried by a generic train track τ_0 , if $\Xi = \Xi(\mathcal{F}, \tau_0)$ is the expansion complex of \mathcal{F} based at τ_0 then the equivalence classes of Ξ are open cells of a linear cell decomposition of Ξ . Letting $c(\sigma)$ denote the closed cell of Ξ whose interior is labelled by a decorated train track σ , we have:*

- (1) *The dimension of $c(\sigma)$ equals the number of nondegenerate switches of σ .*
- (2) *The smallest stratum $\Xi_{\mathcal{J}}$ containing $c(\sigma)$ is labelled by the subset $\mathcal{J} \subset \mathcal{I}$ corresponding to the set of nondegenerate cusps of σ .*
- (3) *There is a natural relation defined on pairs of decorated train tracks which is independent of \mathcal{F} and of τ_0 , called the face relation and denoted $\sigma \sqsubset \sigma'$, which*

has the property that for any cells $c(\sigma)$, $c(\sigma')$ of Ξ , the cell $c(\sigma)$ is a face of the cell $c(\sigma')$ if and only if $\sigma \sqsubset \sigma'$.

This theorem will be proved in Sections 7.4 and 7.5.

In order to understand and motivate the Cell Structure Theorem, we describe some features of Ξ .

Coordinate axes and one cusp expansions. A *coordinate axis* of Ξ is a 1-dimensional stratum $\Xi_{\mathcal{J}}$, where $\mathcal{J} = \{i\}$ is a singleton. Note that $\Xi_{\mathcal{J}}$ is an expanding separatrix family ξ_t , $t \in [0, \infty)$. The 0-cells on $\Xi_{\mathcal{J}}$ are precisely the events in the family ξ_t . Letting $c(\sigma_0), c(\sigma_1), c(\sigma_2), \dots$ be the sequence of 1-cells on $\Xi_{\mathcal{J}}$, by Lemma 7.1.4 it follows that $\sigma_0 \succ \sigma_1 \succ \sigma_2$ is a one cusp splitting sequence, in fact it is the one cusp expansion of \mathcal{F} based at τ_0 whose starting cusp corresponds to the separatrix germ i . We thus obtain a natural bijection between the coordinate axes of Ξ and the one cusp expansions of \mathcal{F} based at τ_0 .

Some faces. As an example of (3) we describe the face relation between cells of codimensions 0 and 1 in Ξ . As a consequence, we describe how generic train track expansions of \mathcal{F} based at τ_0 arise from generic growing separatrix families that are in general position with respect to the cell structure of Ξ .

A closed cell $c(\sigma)$ is top dimensional if and only if each separatrix family $\xi \in \text{int}(c(\sigma))$ is generic, which occurs if and only if the train track σ is generic and has no degenerate cusps.

There are two types of codimension-1 cells $c(\sigma)$. The first type occurs when σ is generic and has exactly one degenerate cusp; equivalently, for each $\xi \in c(\sigma)$, all components of ξ are nondegenerate except for one degenerate component ξ_i , all shunts of ξ are generic except for the shunt t_i passing through $\partial\xi_i$, and $t_i \cap \partial\xi = \partial\xi_i$. In this case $c(\sigma)$ is a top dimensional cell of the codimension-1 stratum $\Xi_{\mathcal{J}}$ where $\mathcal{J} = \mathcal{I} - \{i\}$. The cell $c(\sigma)$ is a face of a single top-dimensional cell of Ξ , obtained by changing the degenerate cusp of σ to a nondegenerate cusp, or equivalently, allowing the degenerate separatrix ξ_i to grow into a very short, nondegenerate separatrix.

The second type of codimension-1 cell $c(\sigma)$ occurs when the train track σ is generic except for a single valence four switch s , and σ has no degenerate cusps. There are two subtypes, corresponding to the two types of elementary moves between generic train tracks, namely slide moves and elementary splittings. In the “slide move” subtype, s has three branch ends on one side and one branch end on the other side; the cell $c(\sigma)$ is a face of two top-dimensional cells, labelled by the two generic train tracks obtained by completely combing the switch s in two distinct ways up to isotopy, and hence these two train tracks differ by a slide move. In the “elementary splitting” subtype, s has two branch ends on each side, and so the

switch s may be regarded as a degenerate sink branch; the cell $c(\sigma)$ is a face of two top-dimensional cells, one labelled by the train track τ obtained by uncombing the switch s creating a nondegenerate sink branch b , and the other labelled by a train track τ' which is obtained from τ by either a Left or Right elementary splitting on b . The distinction between a Left and a Right splitting is not determined from the train track σ itself, although it is determined by the ordering of the two cusps on the unique nongeneric singular tie of the tie bundle $\nu(\xi)$, for any $\xi \in \text{int}(c(\sigma))$.

Notice that a codimension-1 face of the “elementary splitting” type has a natural transverse orientation, from τ to τ' where $\tau \succ \tau'$ is the elementary move; whereas a codimension-1 face of the “slide move” type has no natural transverse orientation.

From the description of the face relation between codimension-0 and codimension-1 cells of $\Xi = \Xi(\mathcal{F}, \tau_0)$, one can see that an expansion of \mathcal{F} of the form $\tau_0 \succ \tau_1 \succ \tau_2 \succ \dots$ consisting entirely of slide moves and elementary splittings corresponds to a growing separatrix family in Ξ that starts in the unique top dimensional cell incident to τ_0 (obtained by labelling all the cusps of τ_0 as nondegenerate), and that is in general position with respect to the cell structure, such that each time the path passes through a codimension-1 face of the elementary splitting type it does so in the positive transverse direction.

Stable equivalence. We now define stable equivalence of expansion complexes. Note that the topological space $[0, \infty)^{\mathcal{I}}$ has one end, in the sense of Freudenthal: a locally compact space X has one end if each compact subset is contained in another compact subset whose complement is connected, or equivalently, each compact subset has a unique unbounded complementary component. A *neighborhood of ∞* in X is any subset containing the unbounded complementary component of some compact subset.

Consider an arational measured foliation \mathcal{F} , and two generic train tracks τ, τ' canonically carrying \mathcal{F} . Let $\Xi = \Xi(\mathcal{F}, \tau)$, $\Xi' = \Xi(\mathcal{F}, \tau')$ be the expansion complexes of \mathcal{F} based at τ, τ' , respectively. We say that Ξ, Ξ' are *stably equivalent* if there exists a label preserving isomorphism $\psi: \hat{\Xi} \rightarrow \hat{\Xi}'$ between subcomplexes $\hat{\Xi}, \hat{\Xi}'$ of Ξ, Ξ' respectively, such that $\hat{\Xi}$ is a neighborhood of ∞ in Ξ , and $\hat{\Xi}'$ is a neighborhood of ∞ in Ξ' . Since labels determine strata according to the Cell Structure Theorem 7.2.2 (2), it follows that the labelling of strata by subsets of the index set \mathcal{I} is natural with respect to stable equivalence, meaning that ψ restricts to an isomorphism between $\Xi_{\mathcal{J}} \cap \hat{\Xi}$ and $\Xi'_{\mathcal{J}} \cap \hat{\Xi}'$, for every subset $\mathcal{J} \subset \mathcal{I}$.

Theorem 7.2.3 (Stable Equivalence Theorem). *For each arational measured foliation \mathcal{F} , any two expansion complexes of \mathcal{F} are stably equivalent.*

Remark. This theorem depends on the convention that an expansion complex of \mathcal{F} is based at a train track which canonically carries \mathcal{F} . A more general class of expansion complexes based at train tracks which carry noncanonically will be discussed later, and for such expansion complexes the Stable Equivalence Theorem fails in general.

Corollaries. Before launching into the proofs of the above theorems, here are some corollaries. The first corollary describes an asymptotic property of the carrying relation among train tracks labelling cells of Ξ .

Corollary 7.2.4. *If Ξ is an expansion complex of an arational measured foliation \mathcal{F} , then for each closed cell $c(\sigma)$ of Ξ there is a neighborhood of ∞ such that for all closed cells $c(\sigma')$ contained in that neighborhood we have $\sigma \succ \sigma'$.*

Proof. Let $\sigma = \sigma(\xi)$ for $\xi \in \Xi$. We may assume that σ is generic, for otherwise we may perturb ξ to obtain a generic uncombing of σ , which carries the exact same σ' as are carried by the original σ . Note that σ canonically carries \mathcal{F} , so the expansion complex Ξ' of \mathcal{F} based at σ is defined. Evidently σ carries every train track labelling a cell of Ξ' , and since Ξ, Ξ' are stably equivalent it follows that σ carries every train track labelling cells in some neighborhood of ∞ in Ξ . \diamond

The following corollary describes a kind of weak stable equivalence property that is satisfied by canonical train track expansions:

Corollary 7.2.5. *For each arational measured foliation \mathcal{F} , if $\tau_0 \succ \tau_1 \succ \cdots$ and $\tau'_0 \succ \tau'_1 \succ \cdots$ are two canonical expansions of \mathcal{F} , then for each $i \geq 0$ there exists $j \geq 0$ such that $\tau_i \succ \tau'_j$. Symmetrically, for each $j \geq 0$ there exists $i \geq 0$ such that $\tau'_j \succ \tau_i$.*

It follows from this corollary that one can construct a train track expansion of \mathcal{F} that starts from $\tau_0 = \tau_{i_0}$, then splits to some τ'_{j_0} , then splits to some τ_{i_1} , then to some τ'_{j_1} , then to some τ_{i_2} , and so on, alternating infinitely between the two given expansions.

Proof of Corollary 7.2.5. We may assume τ_0, τ'_0 are generic. Let Ξ, Ξ' be the expansion complexes of \mathcal{F} based at τ_0, τ'_0 , respectively. By Proposition 4.4.1, each train track τ_i labels some cell of Ξ , and as $i \rightarrow \infty$ the cells labelled by τ_i approach the end of Ξ ; a similar statement holds for Ξ' . Fixing i , we apply Corollary 7.2.4 to conclude that τ_i carries all train tracks labelling cells of Ξ in some neighborhood of ∞ . By stable equivalence of Ξ and Ξ' , this includes the train tracks τ'_j for all sufficiently large j . \diamond

In Section 5.1 we defined two arational measured foliations $\mathcal{F}, \mathcal{F}'$ to be unmeasured equivalent if their equivalence classes have representatives with the same underlying singular foliation. The next corollary gives a combinatorial classification of unmeasured equivalence classes of arational measured foliations:

Corollary 7.2.6. *Given arational measured foliations $\mathcal{F}, \mathcal{F}'$ on S , the following are equivalent:*

- (1) $\mathcal{F}, \mathcal{F}'$ are unmeasured equivalent, that is, $\mathcal{PMF}(\mathcal{F}) = \mathcal{PMF}(\mathcal{F}')$.
- (2) There exists an expansion complex Ξ for \mathcal{F} and an expansion complex Ξ' for \mathcal{F}' such that Ξ, Ξ' are isomorphic.
- (3) Every expansion complex of \mathcal{F} is stably equivalent to every expansion complex of \mathcal{F}' .

Remark. In Section 7.1, after the statement of Corollary 7.1.6, we left open the following question: given two one cusp splitting sequences $\tau_0 \succ \tau_1 \succ \cdots$ and $\tau'_0 \succ \tau'_1 \succ \cdots$ satisfying the canonical killing criterion, how do we tell whether these are expansions of the same arational measured foliation? The best answer that we can give for now is to say this happens if and only if there is an expansion complex Ξ such that the two given sequences are stably equivalent, respectively, to the one cusp expansions associated to two of the axes of Ξ . This is still an unsatisfying answer, because we do not have a handle on how to decide the existence of Ξ . A hint at a better answer comes by realizing that any two axes Ξ_i, Ξ_j of Ξ form the boundary of a 2-stratum $\Xi_{\{i,j\}}$ of Ξ , although this immediately begs the question of how to decide the existence of this 2-stratum. We shall return to this issue briefly when we discuss Two Cusp Splitting Circuits in Section ??.

Proof. The implication (2) \implies (3) follows immediately from the Stable Equivalence Theorem 7.2.3.

To prove (1) \implies (2), suppose that $\mathcal{F}, \mathcal{F}'$ are unmeasured equivalent. Let Ξ be an expansion complex of \mathcal{F} based at a train track τ canonically carrying \mathcal{F} . Choose $\mathcal{F}, \mathcal{F}'$ to be canonical models for their equivalence classes. It follows that the underlying singular foliations of $\mathcal{F}, \mathcal{F}'$ are isotopic, and so replacing \mathcal{F}' up to isotopy if necessary, we may assume that $\mathcal{F}, \mathcal{F}'$ have the *same* underlying singular foliation, and in particular we may identify each of them with the same horizontal foliation of the tie bundle ν over τ . Let Ξ' be the expansion complex of \mathcal{F}' based at τ , and so Ξ and Ξ' have the same underlying set of separatrix families. The key observation is that in the definition of an expansion complex, the construction of the cells and their labels is entirely independent of the transverse measure on

the horizontal foliation of ν : given a separatrix family ξ of \mathcal{F} , the decorated train track $\sigma(\xi)$ is unaffected by a change of the transverse measure on \mathcal{F} , an immediate consequence of the definition of $\sigma(\xi)$ given in Section 4.3. It follows that Ξ and Ξ' are isomorphic.

To prove (3) \implies (1), let Ξ be an expansion complex of \mathcal{F} based at τ and Ξ' an expansion complex of \mathcal{F}' based at τ' , where τ, τ' canonically carry \mathcal{F} . By hypothesis of (3), Ξ and Ξ' are stably equivalent. Consider a train track expansion $\tau = \tau_0 \succ \tau_1 \succ \dots$ of \mathcal{F} . Applying Proposition 4.4.1 and the definition of an expansion complex, it follows that each τ_i is the label of a cell $c(\tau_i)$ of Ξ . Noting that the cells $c(\tau_i)$ are distinct, they must escape all compact subsets of Ξ as $i \rightarrow \infty$. From stable equivalence of Ξ and Ξ' it follows that there exists $I \geq 0$ such that for each $i \geq I$ the train track τ_i labels a cell $c'(\tau_i)$ of Ξ' . In particular, τ_i carries \mathcal{F}' for $i \geq I$. This implies that $\mathcal{F}' \in \cap_i \mathcal{PMF}(\tau_i)$. Since also $\mathcal{F} \in \cap_i \mathcal{PMF}(\tau_i)$, we may apply the Expansion Convergence Theorem 5.1.1 to conclude that $\mathcal{F}, \mathcal{F}'$ are unmeasured equivalent. \diamond

Finally, we show how One Cusp Stable Equivalence is a corollary of ordinary Stable Equivalence:

Proof of One Cusp Stable Equivalence, marked version, 7.1.5. In one sentence, the proof follows from an obvious correspondence between one cusp expansions and axes of the expansion complex. Here are all the words.

Let τ, τ' be generic train tracks each canonically carrying an arational measured foliation \mathcal{F} , which we take to be a canonical model. Choose tie bundles ν, ν' over τ, τ' and carrying bijections $\phi, \phi': \mathcal{F} \rightarrow \nu, \nu'$, respectively, so ν has horizontal foliation $\phi(\mathcal{F})$ with expansion complex Ξ , and ν' has horizontal foliation $\phi'(\mathcal{F})$ with expansion complex Ξ' . Let $s(\cdot)$ denote the set of infinite separatrices of a canonical model. The desired bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$ is defined as the composition

$$\text{cusps}(\tau) \leftrightarrow \text{cusps}(\nu) \leftrightarrow s(\phi(\mathcal{F})) \leftrightarrow s(\mathcal{F}) \leftrightarrow s(\phi'(\mathcal{F})) \leftrightarrow \text{cusps}(\nu') \leftrightarrow \text{cusps}(\tau')$$

The index set \mathcal{I} , formally equal to the set of germs of infinite separatrices of \mathcal{F} , which is in turn identified with $s(\mathcal{F})$, is used to index all of the sets in this string of bijections.

Applying Theorem 7.2.3, we obtain a label preserving isomorphism $\psi: \hat{\Xi} \rightarrow \hat{\Xi}'$ where $\hat{\Xi} \subset \Xi$, $\hat{\Xi}' \subset \Xi'$ are subcomplex neighborhoods of infinity. Moreover, as remarked just before the statement of Theorem 7.2.3, $\psi(\hat{\Xi}_{\mathcal{J}}) = \hat{\Xi}'_{\mathcal{J}}$ for each subset $\mathcal{J} \subset \mathcal{I}$.

Fix an index $i \in \mathcal{I}$, and consider the separatrix $s_i \in s(\mathcal{F})$ and the corresponding cusps $v_i \in \text{cusps}(\tau)$, $v'_i \in \text{cusps}(\tau')$. There is a growing separatrix family of $\phi(\mathcal{F})$ determined by choosing finite initial segments of the separatrix $\phi(s_i) \in s(\mathcal{F}')$, and

this corresponds to the axis Ξ_i of Ξ , and moreover the one cusp expansion of \mathcal{F} based at τ with initial cusp v_i is exactly the expansion obtained from the axis Ξ_i ; this is the content of Lemma 7.1.4. Similarly, the growing separatrix family consisting of finite initial segments of $\phi(\mathcal{F}')$ corresponds to the axis Ξ'_i of Ξ , and the expansion obtained from this axis agrees up to comb equivalence with the one cusp expansion of \mathcal{F} based at τ' with initial cusp v'_i . Letting \mathcal{J} be the singleton set $\{i\}$, we have $\psi(\hat{\Xi}_i) = \hat{\Xi}'_i$, and so the one cusp expansions of \mathcal{F} based at τ, τ' with initial cusps v_i, v'_i are stably equivalent, finishing the proof. \diamond

7.3 Proving the Stable Equivalence Theorem

In this section we reduce the Stable Equivalence Theorem 7.2.3 to the Cell Structure Theorem 7.2.2. The latter will be proved in the following sections.

Consider an expansion complex Ξ of an arational measured foliation \mathcal{F} based at a train track τ . Consider a closed cell $c(\sigma)$ of Ξ . Although we know a priori that tie bundles over σ may not be unique up to isotopy, nevertheless we do have uniqueness in the context of Ξ , as the next statement shows:

Lemma 7.3.1 (Tie bundle uniqueness in expansion complexes). *For all $\xi, \xi' \in \Xi$, the decorated tie bundles $\nu(\xi), \nu(\xi')$ are isotopic if and only if the decorated train tracks $\sigma(\xi), \sigma(\xi')$ are isotopic.*

Proof. For undecorated tie bundles and train tracks, this follows from Corollary 3.6.4 using the fact that \mathcal{F} has carrying bijections into both of the tie bundles $\nu(\xi), \nu(\xi')$. The correspondence of the sets of nondegenerate cusps follows from the fact that the nondegenerate cusps of $\nu(\xi)$ are the pullbacks of the nondegenerate cusps of $\sigma(\xi)$, and similarly for $\nu(\xi')$ and $\sigma(\xi')$. \diamond

As a consequence of this lemma, if σ is a decorated train track labelling some cell $c(\sigma) \subset \Xi$, we can unambiguously write $\nu(\sigma)$ for the tie bundle labelling the same cell, well-defined by the property that $\nu(\sigma) = \nu(\xi)$ for $\xi \in \text{int}(c(\sigma))$. We can also denote the cell $c(\sigma)$ in the unique form $c(\nu)$.

The heart of the Stable Equivalence Theorem is the following statement:

- (*) Let \mathcal{F} be an arational measured foliation, τ_0, τ'_0 two generic train tracks that canonically carry \mathcal{F} , and Ξ, Ξ' the expansion complexes of \mathcal{F} based at τ_0, τ'_0 . For each cell $c(\sigma)$ sufficiently close to infinity in Ξ there exists a cell $c'(\sigma')$ of Ξ' such that σ, σ' are isotopic as decorated tie bundles.

First we shall prove this statement, and later we shall see how it combines with the Cell Structure Theorem to prove the Stable Equivalence Theorem.

Let ν_0, ν'_0 be the tie bundles over τ_0, τ'_0 for which there is a carrying bijection of canonical models $\mathcal{F}_0, \mathcal{F}'_0$ for \mathcal{F} . By Proposition 2.8.1, canonical models for \mathcal{F} are unique up to isotopy, and so by replacing \mathcal{F} in its equivalence class and by isotoping \mathcal{F}_0 we may assume that $\mathcal{F}_0 = \mathcal{F}'_0$. This implies that the tie bundles ν_0, ν'_0 are supported on the same subsurface of S , denoted N_0 , and that the horizontal foliations of ν_0, ν'_0 are the same foliation, namely \mathcal{F} . There are two distinct tie bundle maps $N_0 \rightarrow \tau_0, N_0 \rightarrow \tau'_0$ with distinct tie foliations $\mathcal{F}_v, \mathcal{F}'_v$ respectively.

Let Ξ be the expansion complex of \mathcal{F} based at τ_0 , and Ξ' the expansion complex of \mathcal{F} based at τ'_0 . Ignoring the labelled cell structure, the underlying stratified sets of Ξ and Ξ' are identical: an element of this set is a separatrix family ξ of \mathcal{F} . The foliation \mathcal{F}_v is used to determine a decorated tie bundle associated to each $\xi \in \Xi$ which we will denote $\nu(\xi; \mathcal{F}_v)$, with quotient train track $\sigma(\xi; \mathcal{F}_v)$. The foliation \mathcal{F}'_v determines a decorated tie bundle associated to each $\xi' \in \Xi'$ denoted $\nu(\xi'; \mathcal{F}'_v)$, with quotient train track $\sigma(\xi'; \mathcal{F}'_v)$. The fact that the train tracks τ_0, τ'_0 may be very different implies that the tie foliations $\mathcal{F}_v, \mathcal{F}'_v$ of N_0 may also be very different, implying in turn that $\sigma(\xi; \mathcal{F}_v)$ and $\sigma(\xi; \mathcal{F}'_v)$ may be very different for the same separatrix family ξ of \mathcal{F} .

The key idea for proving (*) is that if $\xi \in \Xi$ is a sufficiently long separatrix family, then it cuts the leaves of the foliations \mathcal{F}_v and \mathcal{F}'_v into very short segments, but any two foliations of N_0 transverse to \mathcal{F} look very much alike when focussing only on sufficiently short segments, and so we would expect that $\sigma(\xi; \mathcal{F}_v)$ and $\sigma(\xi; \mathcal{F}'_v)$ may not be so different after all.

To prove (*), by Lemma 7.3.1 it suffices to show that for each $\xi \in \Xi$ sufficiently close to infinity there exists $\xi' \in \Xi'$ such that the decorated tie bundles $\nu(\xi; \mathcal{F}_v)$, $\nu(\xi'; \mathcal{F}'_v)$ are isotopic. The separatrix family ξ' may not be exactly the same as ξ , but it will be only a small perturbation.

We remarked earlier that in order to construct the undecorated tie bundle $\nu(\xi; \mathcal{F}_v)$, in addition to N_0 and ξ one needs only the \mathcal{F}_v -shunts of ξ ; one does not need the rest of the information of the foliation \mathcal{F}_v . One can also construct the non-degenerate cusps of $\nu(\xi; \mathcal{F}_v)$ without further information, because they correspond to the nondegenerate components of ξ .

Thus, to prove (*) it suffices to show that:

- (**) For each $\xi \in \Xi$ sufficiently close to infinity there exists an isotopy of N_0 preserving \mathcal{F} , taking each \mathcal{F}_v shunt of ξ to a leaf segment of \mathcal{F}'_v , and taking ξ to some separatrix family ξ' .

This proves (*) because it follows that the images of the \mathcal{F}_v shunts of ξ under the isotopy are precisely the \mathcal{F}'_v shunts of ξ' , and the cusp labelling is preserved for free.

We shall think of \mathcal{F} as the horizontal direction and \mathcal{F}'_v as the “true” vertical direction in N_0 , and so the first part of (**) can be reworded to say that for each $\xi \in$

Ξ sufficiently close to infinity there is a horizontal isotopy of N_0 which straightens out each \mathcal{F}_v shunt of ξ until it is truly vertical.

Define a Riemannian metric on N_0 of the form $ds^2 = |dx|^2 + |dy|^2$, where $|dy|$ denotes the transverse measure on \mathcal{F} and $|dx|$ is any fixed transverse measure on \mathcal{F}'_v ; all distance measurements are taken with respect to this Riemannian metric. It is a Euclidean metric with geodesic boundary, except at the reflex cusp singularities. Define a *rectangle* in N_0 to be any embedding $I \times J \hookrightarrow N_0$, where I, J are compact arcs, such that $s \times J$ is a leaf segment of \mathcal{F}'_v and $I \times t$ is a leaf segment of \mathcal{F} for each $s \in I, t \in J$.

The outline of the proof of (***) is as follows. First we choose ξ sufficiently long so that the union of its \mathcal{F}_v shunts can be isolated into a collection of pairwise disjoint rectangles. Then we shall define a horizontal isotopy of N_0 supported in these rectangles which straightens each \mathcal{F}_v shunt of ξ to be truly vertical.

As a preliminary step, there exists a closed neighborhood R_s of each reflex cusp s of N_0 , and there exists a horizontal isotopy of N_0 , such that after the isotopy the foliation \mathcal{F}_v restricted to each R_s is truly vertical. This will be useful in several places.

The *interior* of a subset X of N_0 , denoted $\text{int}(X)$, means the largest open subset of N_0 containing X .

Lemma 7.3.2 (Rectangle Lemma). *There exists a constant $\delta > 0$ such that if A is a subset of N_0 and $\text{diam}(A) < \delta$ then either $A \subset \text{int}(R_s)$ for some s or $A \subset \text{int}(R)$ for some rectangle R such that $\text{diam}(R) < 3 \text{diam}(A)$.*

Proof. Choose a finite collection of rectangles $R_i, i = 1, \dots, I$, such that the interiors of the sets R_1, \dots, R_n , together with the interiors of the sets R_s for reflex cusps s of N_0 , form an open cover of the compact metric space N_0 . Let δ be a Lebesgue number for this covering. If $\text{diam}(A) < \delta$ then either $A \subset \text{int}(R_s)$ for some cusp s and we are done, or $A \subset \text{int}(R_i)$ for some $i = 1, \dots, I$. In the latter case let R' be the smallest rectangle such that $A \subset R' \subset R_i$. A simple argument in Euclidean geometry shows that $\text{diam}(R') \leq 2\sqrt{2} \text{diam}(A) < 3 \text{diam}(A)$. By choosing the rectangle $R \subset \text{int}(R_i)$ slightly larger than R' we obtain $A \subset \text{int}(R)$, and we still have the inequality $\text{diam}(R) < 3 \text{diam}(A)$. \diamond

We now fix some sufficiently small constant ϵ , whose size will be determined as the argument proceeds. By choosing ξ sufficiently long it follows that each \mathcal{F}_v shunt of ξ has diameter less than ϵ . Let \mathcal{A} be the set of \mathcal{F}_v shunts of ξ , and so \mathcal{A} has cardinality at most equal to $i(\tau_0)$, the number of cusps of N_0 .

Our first requirement on ϵ is that $\epsilon < \delta$, and so for each $\alpha \in \mathcal{A}$ we can apply the Rectangle Lemma, reaching one of two possible conclusions. In the first case, we have $\alpha \subset \text{int}(R_s)$ for some cusp s , but α is a leaf segment of \mathcal{F}_v and so it follows

that α is truly vertical; in this case we set $R(\alpha) = \alpha$ by convention, a “vertically degenerate rectangle”. In the second case, we have $\alpha \subset \text{int}(R(\alpha))$ for some rectangle $R(\alpha)$ with

$$\text{diam}(R(\alpha)) < 3 \text{diam}(\alpha) < 3\epsilon = \epsilon_1$$

Note that this inequality is satisfied in the first case as well.

Let $\mathcal{R}_1 = \{R(\alpha) \mid \alpha \in A\}$. Each $R \in \mathcal{R}_1$ is a rectangle of diameter less than ϵ_1 , and \mathcal{R}_1 has cardinality $\#\mathcal{R}_1 \leq i(\tau_0)$. For each shunt $\alpha \in \mathcal{A}$, either α is truly vertical and is an element of \mathcal{R}_1 , or α is contained in the interior of some element of \mathcal{R}_1 .

We want the set of rectangles \mathcal{R}_1 to be pairwise disjoint, so that we can carry out horizontal isotopies supported in them. But if they are not pairwise disjoint, then we shall start coalescing them and applying the Rectangle Lemma inductively, producing smaller and smaller collections $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3, \dots$ of larger and larger rectangles, continuing until we obtain a pairwise disjoint collection \mathcal{R}_K . Since the cardinality of \mathcal{R}_1 is at most $i(\tau_0)$, the process must stop with $K \leq i(\tau_0)$. Each application of the Rectangle Lemma will require a certain inequality, which will be justified by a sequence of stricter and stricter upper bounds on ϵ .

Assuming \mathcal{R}_1 is not pairwise disjoint, we carry out the next step of the induction to produce \mathcal{R}_2 . Each component C of the set $\cup \mathcal{R}_1$ has diameter less than

$$\epsilon_1 \cdot \#\mathcal{R}_1 \leq 3i(\tau_0)\epsilon$$

Our second requirement on ϵ is that $3i(\tau_0)\epsilon < \delta$, allowing us to apply the Rectangle Lemma to the components of $\cup \mathcal{R}_1$. If C is a component of $\cup \mathcal{R}_1$ then there are two possible conclusions. In the first case, we have $C \subset R_s$ for some cusp s , and then C is a union of degenerate rectangles of the form $R(\alpha) = \alpha$ for truly vertical shunts α , but any two of these are pairwise disjoint, so in fact $C = \alpha$ for some truly vertical shunt α , and we obtain a vertically degenerate rectangle $R(C) = \alpha$. In the second case, we have $C \subset \text{int}(R(C))$ for some rectangle $R(C)$ with

$$\text{diam}(R(C)) < 3 \text{diam}(C) < 3^2 i(\tau_0)\epsilon = \epsilon_2$$

The same inequality is evidently satisfied in the first case.

Let \mathcal{R}_2 be the set of rectangles $R(C)$ for components C of $\cup \mathcal{R}_1$, each with diameter less than ϵ_2 , and with cardinality $\#\mathcal{R}_2 < \#\mathcal{R}_1 \leq i(\tau_0)$.

We continue inductively as long as the collection \mathcal{R}_k is not pairwise disjoint, using a sequence of stricter and stricter conditions on ϵ of the form $3^k i(\tau_0)^k \epsilon < \delta$ in order to apply the Rectangle Lemma. Eventually we obtain a collection of pairwise disjoint rectangles \mathcal{R}_K with $K \leq i(\tau_0)$, such that for each shunt $\alpha \in A$, either α is strictly vertical and is an element of \mathcal{R}_K , or α is contained in the interior of an element of \mathcal{R}_K .

We have achieved our intermediate goal of isolating the shunts in a collection \mathcal{R}_K of pairwise disjoint rectangles.

Consider now any rectangle $R \in \mathcal{R}_K$. Assuming R is not vertically degenerate, we must construct a horizontal isotopy supported on R which will straighten out the set of shunts $\mathcal{A}_R = \{\alpha \in \mathcal{A} \mid \alpha \subset \text{int}(R)\}$ so that each of them is truly vertical. Choose an isometric parameterization $R \approx [0, a] \times [0, b]$, and let $\pi: R \rightarrow [0, b]$ be projection.

Define a relation \triangleleft on any finite, pairwise disjoint collection of transverse arcs of the horizontal foliation of R , where $\beta \triangleleft \beta'$ if $\beta \neq \beta'$ and there exists $(x, y) \in \beta$, $(x', y') \in \beta'$ such that $x < x'$. Note that it is impossible to have $\beta \triangleleft \beta'$ and $\beta' \triangleleft \beta$, because β, β' are disjoint.

We claim that for any finite, pairwise disjoint collection of transverse arcs in R , the ordering \triangleleft has a minimum element, and we prove this by induction on the size of the collection on which \triangleleft is defined. The key point we must demonstrate is that the transitive closure of \triangleleft , which we denote $<$, is a partial order: if $\alpha < \alpha'$ then $\alpha' \not< \alpha$. To prove this, choose a chain $\alpha = \alpha_1 \triangleleft \alpha_2 \triangleleft \cdots \triangleleft \alpha_n = \alpha'$. Using α_1, α_2 we can define a transverse arc β such that $\pi(\beta) = \pi(\alpha_1 \cup \alpha_2)$, $\alpha_1 \triangleleft \beta \triangleleft \alpha_2$, and none of $\alpha_3, \dots, \alpha_n$ come between α_1 and β nor between β and α_2 . If $\pi(\alpha_1) \supset \pi(\alpha_2)$ then choose β to be slightly to the right of α_1 , and if $\pi(\alpha_1) \subset \pi(\alpha_2)$ then choose β to be slightly to the left of α_2 . Otherwise, let β start slightly to the right of the endpoint of α_1 whose projection is disjoint from the projection of α_2 , then β travels close to and on the right side of α_1 until it comes between α_1 and α_2 , then β moves over to be close to and on the left side of α_2 , and β then continues in this manner ending slightly to the left of the endpoint of α_2 whose projection is disjoint from the projection of α_1 . Clearly we have $\beta \triangleleft \alpha_3 \triangleleft \cdots \triangleleft \alpha_n$, and so by induction we have $\alpha_n \not< \beta$, which implies $\alpha_n \not< \alpha_1$, proving the claim.

Let α be a \triangleleft minimal element of \mathcal{A}_R . Then we can easily construct a horizontal isotopy supported in $R \approx [0, a] \times [0, b]$ after which the following holds: α is truly vertical, of the form $x \times [c, d]$, and for each $\alpha' \in \mathcal{A}_R - \{\alpha\}$ and each $(x', y') \in \alpha'$ we have $x < x'$. We may now focus on the subrectangle $[x', a] \times [0, b]$ of R , which contains one fewer shunt, and continue by induction.

To summarize, we have proved that, after a preliminary horizontal isotopy of N_0 that makes \mathcal{F}_v truly vertical in a neighborhood R_s of each cusp $s \in N_0$, there exists a horizontal isotopy of N_0 supported in a pairwise disjoint collection of at most $i(\tau_0)$ rectangles, at the end of which each \mathcal{F}_v shunt of ξ becomes truly vertical. As mentioned earlier, if ξ' is the image of ξ under this isotopy, then the isotopy takes $\nu(\xi; \mathcal{F}_v)$ to $\nu(\xi'; \mathcal{F}'_v)$. We must also prove that the isotopy preserves decorations, but the isotopy obviously induces a bijection between nondegenerate components of ξ and of ξ' .

This completes the proof of (**) and hence of (*).

Now we prove the Stable Equivalence Theorem by applying (*) together with the Cell Structure Theorem 7.2.2. Given a cell complex X let $\mathcal{C}(X)$ denote the set of cells of X .

Let \mathcal{F} , τ_0 , τ'_0 , $\Xi = \Xi(\mathcal{F}, \tau_0)$, $\Xi' = \Xi(\mathcal{F}, \tau'_0)$ be as in (*). Define a relation Ψ between $\mathcal{C}(\Xi)$ and $\mathcal{C}(\Xi')$ where $c(\sigma) \Psi c(\sigma')$ if and only if σ, σ' are isotopic as decorated train tracks. Evidently the relation Ψ is a bijection between a subset of $\mathcal{C}(\Xi)$ and a subset of $\mathcal{C}(\Xi')$. By applying (*), it follows that there is a subcomplex $\hat{\Xi} \subset \Xi$ which is a neighborhood of infinity such that $\mathcal{C}(\hat{\Xi})$ is contained in the projection of the relation Ψ to the first factor. By applying Theorem 7.2.2 (3), the set of cells of Ξ' that correspond, via \mathcal{P} , to cells of $\hat{\Xi}$ forms a subcomplex $\hat{\Xi}' \subset \Xi'$. We thus have a bijection $\Psi: \mathcal{C}(\hat{\Xi}) \rightarrow \mathcal{C}(\hat{\Xi}')$.

We claim that $\hat{\Xi}'$ is a neighborhood of infinity in Ξ' . To see why, one knows, by applying (*) with the roles of Ξ, Ξ' reversed, that there is a subcomplex $\hat{\Xi}'_0 \subset \Xi'$ that is a neighborhood of infinity in Ξ' and whose cells are contained in the image of the projection of Ψ to the second factor. Since $\hat{\Xi}$ is a neighborhood of infinity in Ξ , it follows that after deleting a finite number of cells in $\hat{\Xi}'_0$ we may assume that each cell in $\hat{\Xi}'_0$ is of the form $\Psi(c)$ for some $c \in \mathcal{C}(\hat{\Xi})$. After this deletion, $\hat{\Xi}'_0$ is still a neighborhood of infinity in Ξ' . But it now follows that $\hat{\Xi}'_0 \subset \hat{\Xi}'$ and hence $\hat{\Xi}'$ is a neighborhood of infinity in Ξ' .

By Theorem 7.2.2 (3), the map $\Psi: \mathcal{C}(\hat{\Xi}) \rightarrow \mathcal{C}(\hat{\Xi}')$ respects the face relation of cells. By using the rest of the Cell Structure Theorem 7.2.2 and working inductively up through the skeleta, we can promote Ψ to a label preserving, cellular isomorphism $\psi: \hat{\Xi} \rightarrow \hat{\Xi}'$, as follows. First, restricting Ψ to the 0-cells defines a label preserving cellular isomorphism of 0-skeleta $\psi^{(0)}: \hat{\Xi}^{(0)} \rightarrow \hat{\Xi}'^{(0)}$. Assuming that we have a label preserving cellular isomorphism of i -skeleta $\psi^{(i)}: \hat{\Xi}^{(i)} \rightarrow \hat{\Xi}'^{(i)}$, consider an $i+1$ cell $c(\sigma) \in \mathcal{C}(\hat{\Xi})$, and so $\Psi(c(\sigma)) = c(\sigma') \in \mathcal{C}(\hat{\Xi}')$ where σ, σ' are isotopic as decorated train tracks. Applying Theorem 7.2.2 (1), $c(\sigma')$ also has dimension $i+1$, and so $\partial c(\sigma) \subset \Xi^{(i)}$ and $\partial c(\sigma') \subset \Xi'^{(i)}$. By Theorem 7.2.2 (3), Ψ restricts to a bijection between $\mathcal{C}(\partial c(\sigma))$ and $\mathcal{C}(\partial c(\sigma'))$. Applying the induction hypothesis, $\psi^{(i)}$ restricts to a label preserving cellular isomorphism from $\partial c(\sigma)$ to $\partial c(\sigma')$. Since Ξ, Ξ' are linear cell complexes by Theorem 7.2.2, the subcomplexes $\partial c(\sigma)$ and $\partial c(\sigma')$ are i -spheres, and the homeomorphism $\psi^{(i)}$ between these spheres extends to a homeomorphism between the $i+1$ balls $c(\sigma)$ and $c(\sigma')$. Piecing these homeomorphisms together, one for each $(i+1)$ -cell $c(\sigma)$ of $\hat{\Xi}$, we obtain a label preserving cellular isomorphism $\psi^{(i+1)}: \hat{\Xi}^{(i+1)} \rightarrow \hat{\Xi}'^{(i+1)}$. Continuing by induction on i , we obtain a label preserving cellular isomorphism $\psi: \hat{\Xi} \rightarrow \hat{\Xi}'$.

This proves the Stable Equivalence Theorem, modulo the Cell Structure Theorem 7.2.2.

7.4 Cells of the expansion complex

In this and the next section we prove the Cell Structure Theorem 7.2.2.

Fix a generic train track τ with tie bundle $\nu_0 \rightarrow \tau$ and tie foliation \mathcal{F}_ν , and fix a canonical arational measured foliation \mathcal{F} with a carrying bijection $\mathcal{F} \hookrightarrow \nu_0$. Let Ξ be the expansion complex of \mathcal{F} based at τ . Let \mathcal{I} be the set of proper separatrix germs of \mathcal{F} , so each $\xi \in \Xi$ has components $\xi = \sqcup_{i \in \mathcal{I}} \xi_i$ with ξ_i a representative of the germ i , possibly degenerating to a point. Fixing a transverse measure on \mathcal{F}_ν we obtain a parameterization $\Xi \approx [0, \infty)^{\mathcal{I}}$.

We will use the symbols such as x, ξ to represent points in Ξ . Generally a Roman letter x will be a variable point in Ξ , in which case we abuse notation by using x_i to denote a real variable representing the i^{th} coordinate of x , that is, the length of the i^{th} component of x (rather than the i^{th} component itself). Generally a Greek letter ξ will be a constant point in Ξ , in which case the lengths of its components $\xi = \sqcup_i \xi_i$ will be represented by constant real numbers $\ell_i = \ell_i(\xi) = \text{Length}(\xi_i)$.

Cells of Ξ . The equivalence class of $\xi \in \Xi$ is the set

$$\overset{\circ}{c}(\xi) = \{x \in \Xi \mid \sigma(\xi), \sigma(x) \text{ are isotopic as decorated train tracks.}\}$$

Our primary goal is to show that $\overset{\circ}{c}(\xi)$ is the interior of a polyhedral cell. We accomplish this by writing down a system $X_\xi = E_\xi \cup I_\xi$ consisting of 1st degree equations E_ξ and strict inequalities I_ξ in the variables $x_i, i \in \mathcal{I}$, and showing that the solution set of the system X_ξ equals $\overset{\circ}{c}(\xi)$.

The system X_ξ is produced in two stages. First we define a *homotopic carrying injection* $f_\xi: \sigma(\xi) \rightarrow \nu_0$, which means a carrying injection that is also a homotopy equivalence; the map f_ξ will be well-defined up to isotopy along ties. We will also construct f_ξ so that it is *decoration preserving* in the appropriate sense. Then, given any decorated train track σ contained in ν_0 such that the inclusion map $\sigma \hookrightarrow \nu_0$ is a decoration preserving homotopic carrying injection, we shall define a system $X_\sigma = E_\sigma \cup I_\sigma$ consisting of equations E_σ and strict inequalities I_σ in the variables $x_i, i \in \mathcal{I}$. From the construction it will be evident that X_σ is well-defined up to isotopy of σ along ties of ν_0 . Thus, for any $\xi \in \Xi$, the system $X_\xi = X_{f_\xi(\sigma(\xi))}$ is well defined.

We shall prove that $\overset{\circ}{c}(\xi)$ is contained in the solution set of X_ξ as follows. Not only will X_σ be well-defined up to isotopy of σ along ties of ν_0 , but it will be well defined up to isotopy of σ through decoration respecting carrying injections into ν_0 . Also, it will be obvious from the construction that $x = \xi$ is in the solution set of X_ξ . Given any $\eta \in \overset{\circ}{c}(\xi)$, the decorated train tracks $\sigma(\xi), \sigma(\eta)$ are isotopic, and by applying Lemma 3.5.2, it follows that $f_\xi(\sigma(\xi))$ and $f_\eta(\sigma(\eta))$ are isotopic through

carrying injections, and hence the systems X_ξ and X_η are identical. In particular they have the same solution set, and since $x = \eta$ is in the solution set of X_η , it is also in the solution set of X_ξ .

The other inclusion, showing that the solution set of X_ξ is contained in $\overset{\circ}{c}(\xi)$, will follow from a convexity argument.

The carrying injection $f_\xi: \sigma(\xi) \rightarrow \nu_0$. Recall that $\sigma(\xi)$ is constructed by slicing ν_0 along ξ to obtain a tie bundle $\nu(\xi)$, and then collapsing the ties of $\nu(\xi)$ to obtain $\sigma(\xi)$. The carrying injection f_ξ , depicted in Figure 21, is described as follows.

There is a natural bijection between shunts of ξ , singular ties of $\nu(\xi)$, and switches of $\sigma(\xi)$; let $t_s(\xi)$ be the shunt corresponding to a switch $s \in \sigma(\xi)$. For each switch $s \in \sigma(\xi)$ define its image $f_\xi(s)$ to be a point of $t_s(\xi)$, with the following restriction. If s is nondegenerate then $f_\xi(s)$ may be any point in $t_s(\xi)$. Supposing s is degenerate, let v be the degenerate cusp of $\sigma(\xi)$ located at s , and let c be the corresponding cusp of ν_0 . It follows that $t_s(\xi)$ is a subtie of the tie of ν_0 through c , and moreover c is contained in the interior of $t_s(\xi)$. In this case we define $f_\xi(s) = c$; in making this definition, we are using the fact that $\sigma(\xi)$ has at most one degenerate cusp at each switch, in order for f_ξ to be well-defined on degenerate switches.

There is a natural bijection between maximal subrectangles of ν_0 whose interiors are disjoint from $\xi \cup \{\text{shunts of } \xi\}$, maximal subrectangles of $\nu(\sigma)$ whose interiors are disjoint from the singular ties of $\nu(\sigma)$, and branches of $\sigma(\xi)$; let $R_b(\xi)$ be the rectangle in ν_0 corresponding to a branch b of $\sigma(\xi)$. Consider a branch $b \subset \sigma(\xi)$, with one end on a switch s and the other end on a switch s' . One vertical side of the rectangle $R_b(\xi)$ is a subsegment $\partial_s R_b(\xi)$ of $t_s(\xi)$, and the other vertical side is a subsegment $\partial_{s'} R_b(\xi)$ of $t_{s'}(\xi)$. It is possible that $f_\xi(s)$ does not lie in $\partial_s R_b(\xi)$, and $f_\xi(s')$ may not lie in $\partial_{s'} R_b(\xi)$. We now define $f_\xi(b)$ to be a segment transverse to \mathcal{F}_v with endpoints $f_\xi(s)$, $f_\xi(s')$, which is isotopic staying transverse to \mathcal{F}_v to a horizontal leaf of $R_b(\xi)$; to be explicit, starting from $f_\xi(s)$, the path $f_\xi(b)$ travels transverse to \mathcal{F}_v staying close to $t_s(\xi)$ until it enters the interior of $R_b(\xi)$, then $f_\xi(b)$ travels along a leaf of \mathcal{F} until it approaches $t_{s'}(\xi)$, and then $f_\xi(b)$ travels transverse to \mathcal{F}_v staying close to $t_{s'}(\xi)$ until it ends at $f_\xi(s')$. If one of the switches, say s , is degenerate, then there seems to be a possible obstruction to this definition of $f_\xi(b)$: the shunt $t_s(\xi)$ contains a reflex cusp r of ν_0 , and if $R_b(\xi)$ is on the wrong side of r then it seems possible that as the path $f_\xi(b)$ travels transverse to \mathcal{F}_v staying close to $t_s(\xi)$, it may be forced to cross the boundary of ν_0 into the complement $S - \nu_0$. However, this possibility is precluded by our requirement that $f_\xi(s) = r$, which guarantees that $R_b(\xi)$ is on the correct side of r .

Although the construction of the map $f_\xi: \sigma(\xi) \rightarrow \nu_0$ depends on choices, it is

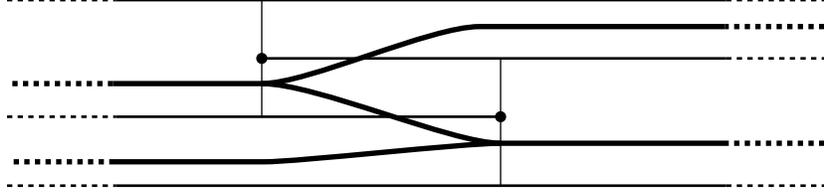


Figure 21: In the tie bundle ν_0 , the separatrix family ξ determines a train track $\sigma(\xi)$ which has a carrying injection into ν_0 .

clearly well defined up to isotopy along leaves of the foliation \mathcal{F}_v . Notice that f_ξ is a homotopy equivalence: in fact, the composition $\sigma(\xi) \rightarrow \nu_0 \rightarrow \tau_0$ is homotopic to the natural homotopic carrying map of $\sigma(\xi)$ onto τ_0 , and the tie bundle projection $\nu_0 \rightarrow \tau_0$ is a homotopy equivalence.

The system X_σ . Consider now a decorated train track σ . We assume that $\sigma \subset \nu_0$ and that the inclusion map is a *homotopic carrying injection*, that is, σ is transverse to the ties of ν_0 , and the inclusion $\sigma \hookrightarrow \nu_0$ is a homotopy equivalence. There is a natural bijection between cusps of ν_0 and cusps of σ ; for each $i \in \mathcal{I}$ let c_i denote the cusp of ν_0 at which the separatrix germ i is based, and let v_i be the corresponding cusp of σ .

Abstracting the above discussion, we say that the inclusion $\sigma \hookrightarrow \nu_0$ *respects decoration* if, for each degenerate cusp v_i of σ located at a switch s , we have $s = c_i$. Thus, for any $\xi \in \Xi$ the inclusion of $f_\xi(\sigma(\xi))$ into ν_0 respects decoration.

Given a decoration respecting, homotopic carrying injection $\sigma \hookrightarrow \nu_0$, we shall define a system X_σ of equations E_σ and strict inequalities I_σ in the variables x_i , $i \in \mathcal{I}$. We will also check that the form of each equation and inequality in this system depends on σ only up to isotopy through decoration respecting carrying injections.

The bijection $c_i \leftrightarrow v_i$ between cusps(ν_0) and cusps(σ) is characterized as follows (see Figure 22). Suppose first that v_i is a degenerate cusp, located at a switch s ; because $\sigma \hookrightarrow \nu_0$ respects decorations, the switch s coincides with a cusp of ν_0 , and that cusp must be c_i . Suppose now v_i is a nondegenerate cusp, and it follows that there exists a unique triangle T_i in ν_0 , with corner vertices p, q and cusp vertex v_i , such that the side \overline{pq} of T_i is the maximal tie segment of ν_0 containing c_i and with interior disjoint from σ , each side $\overline{pv_i}, \overline{qv_i}$ of T_i is a train path in σ , the interior of T_i is disjoint from σ , and T_i is foliated by segments of \mathcal{F}_v each connecting a point of $\overline{pv_i}$ to a point of $\overline{qv_i}$, with one degenerate leaf of the foliation at the cusp v_i . To be formally correct, T_i lives not in ν_0 but in the surface-with-cusps $\mathcal{C}(\nu_0 - \sigma)$, but

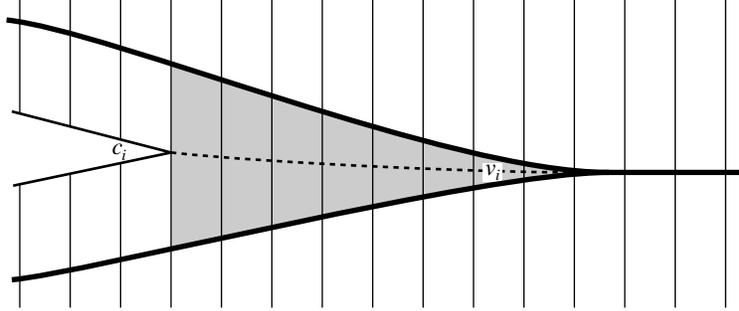


Figure 22: A cusp triangle $T_i \subset \nu_0$, with one cusp and two corners, associated to a nondegenerate cusp v_i of σ and the associated cusp c_i of ν_0 . Also shown is a section of T_i (dotted line). The triangle T_i is foliated by tie segments, with one degenerate leaf at the cusp vertex v_i , with vertical side on the tie through c_i , with train paths in σ as horizontal sides, and with interior disjoint from σ .

the interior of T_i embeds in ν_0 via the overlay map $\mathcal{C}(\nu_0 - \sigma) \rightarrow \nu_0$, and so we shall abuse terminology by referring to T_i as living in ν_0 . The triangle T_i is called the *cusp triangle* of σ associated to the cusp v_i .

We extend the definition of $l_i(\sigma)$ to the case where v_i is a degenerate cusp of σ , by taking $l_i(\sigma) = 0$.

For each branch b of σ the \mathcal{F}_v transverse measure of σ is denoted $l_b(\sigma)$.

Notice that when $\sigma = f_\xi(\sigma(\xi))$ for $\xi \in \Xi$, then clearly we have $l_i(\sigma) = l_i(\xi)$, and $l_b(\sigma) = l_b(\xi)$. From these identities it will follow that $x = \xi$ is a solution of each equation and inequality in the system X_ξ .

Before proceeding we need a construction. Consider a continuous path $\gamma: [a, b] \rightarrow \nu_0$. At each point of ν_0 there are two transverse orientations of \mathcal{F}_v . The transverse orientations at each point of ν_0 fit together to form a two-to-one covering space over ν_0 , the transverse orientation bundle of ν_0 . A *transverse orientation* of γ is a choice of transverse orientation at each point $\gamma(t)$ which varies continuously in $t \in [a, b]$ —in other words, a continuous lift of γ to the transverse orientation bundle. Each path γ has exactly two transverse orientations, and a choice of transverse orientation for a single point $t \in [a, b]$ extends to a unique transverse orientation on γ . A *transversely oriented* path in ν_0 consists of a path together with a choice of transverse orientation. Given a transversely oriented path $\gamma: [a, b] \rightarrow \nu_0$, we define the *signed intersection number* $\langle \gamma, \mathcal{F}_v \rangle$ as follows: for each $t \in [a, b]$ the choice of transverse orientation at $\gamma(t)$ extends uniquely to a transverse orientation of \mathcal{F}_v in a small neighborhood U of t , which determines in turn a unique closed, singular 1-cocycle on U with the property that the absolute value of this 1-cocycle equals

the given transverse measure on \mathcal{F}_v , and the 1-cocycle evaluates to a positive number on any vector based at $\gamma(t)$ that points in the same direction as the specified transverse orientation. By pulling back these locally defined 1-forms via γ we obtain a well-defined 1-cocycle on $[a, b]$, whose integral is equal to $\langle \gamma, \mathcal{F}_v \rangle$. From this definition it is clear that $\langle \gamma, \mathcal{F}_v \rangle$ is independent of γ up to a homotopy for which the endpoints $\gamma(a), \gamma(b)$ are fixed and for which the orientation varies continuously in the homotopy parameter.

The system of equations E_σ . The equations in E_σ are of two types.

First, for each degenerate cusp v_i of σ , there is a *degeneracy equation* $x_i = 0$. Given $\xi \in \Xi$, the fact that the inclusion $f_\xi: \sigma(\xi) \hookrightarrow \nu_0$ respects decorations obviously implies that $x = \xi$ is in the solution set of $x_i = 0$.

Second, for each pair $i, j \in \mathcal{I}$ such that the cusps v_i, v_j of σ are located at the same switch s of σ , there is an equation $E_{ij}(\sigma)$ which gives numerical expression to the fact that the two separatrices ξ_i, ξ_j have boundary points on the same shunt of ξ .

In discussing the equation $E_{ij}(\sigma)$, we shall consider the path $\rho_{ij}(\sigma)$ which is a concatenation of two pieces: a path in the cusp triangle T_i from c_i to s , concatenated at s with a path in T_j from s to c_j (if one of v_i, v_j is degenerate then one of the two pieces of $\rho_{ij}(\sigma)$ will degenerate to a point). We choose a transverse orientation on $\rho_{ij}(\sigma)$ so that at the initial point c_i the orientation points in the same direction as the cusp c_i . Since each of the two pieces of $\rho_{ij}(\sigma)$ is well-defined up to homotopy rel endpoints, and since the initial orientation is well-defined, we obtain the following fact:

Fact 7.4.1. *As σ varies by isotopy through decoration respecting carrying injections, the cusps v_i, v_j continue to be located at the same switch, and the path $\rho_{ij}(\sigma)$ varies by a homotopy that preserves endpoints and orientation. It follows that the signed intersection number $\langle \rho_{ij}(\sigma), \mathcal{F}_v \rangle$ is constant as σ varies.*

Note that, if desired, the path $\rho_{ij}(\sigma)$ can be chosen to be transverse to \mathcal{F}_v , *except* that in the case that v_i, v_j are on the same side of the switch s , the path $\rho_{ij}(\sigma)$ cannot be transverse at s . However, because of the way we defined signed intersection number, transversality properties of $\rho_{ij}(\sigma)$ are unnecessary in the arguments that follow, and we shall need this freedom in one particular place.

There are two cases in the definition of $E_{ij}(\sigma)$, depending on whether v_i, v_j are on the same or opposite sides of s ; see Figure 23.

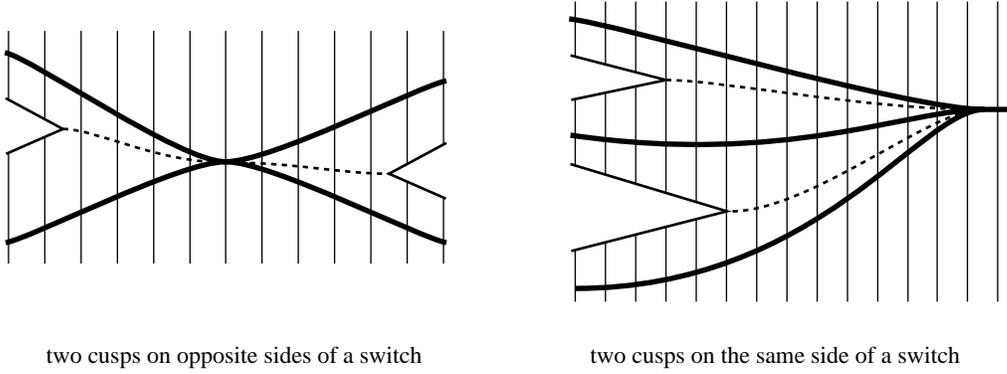


Figure 23: If the cusps v_i, v_j of σ are located at the same switch of σ , then there is an equation $E_{ij}(\sigma)$ of the form $x_i \pm x_j = \langle \rho, \nu_0 \rangle$, where ρ is a path from c_i to c_j consisting of a path in T_i from c_i to s , concatenated with a path in T_j from s to c_j . The \pm sign is $+$ if v_i, v_j are on opposite sides of s , and $-$ otherwise.

The equation $E_{ij}(\sigma)$, Case 1: If v_i, v_j are on opposite sides of the switch s then the equation is

$$\begin{aligned} x_i + x_j &= \ell_i(\sigma) + \ell_j(\sigma) & E_{ij}(\sigma), \text{ Case 1} \\ &= \langle \rho_{ij}(\sigma), \mathcal{F}_v \rangle \end{aligned}$$

It is clear that $x = \xi$ is a solution of $E_{ij}(\sigma)$. As σ varies by isotopy through carrying injections, Fact 7.4.1 implies that the equation $E_{ij}(\sigma)$ is invariant.

The equation $E_{ij}(\sigma)$, Case 2: If v_i, v_j are on the same side of the switch s then the equation is

$$\begin{aligned} x_i - x_j &= \ell_i(\sigma) - \ell_j(\sigma) & E_{ij}(\sigma), \text{ Case 2} \\ &= \langle \rho_{ij}(\sigma), \nu_0 \rangle \end{aligned}$$

Again it is clear that $x = \xi$ is a solution of $E_{ij}(\sigma)$. Again, as σ varies through carrying injections, Fact 7.4.1 says that $E_{ij}(\sigma)$ is invariant.

The inequalities I_σ . The system of strict inequalities I_σ consists of two types of inequalities.

First, for each nondegenerate cusp v_i of σ there is a *nondegeneracy inequality* $x_i > 0$. The fact that the inclusion $f_\xi: \sigma(\xi) \hookrightarrow \nu_0$ respects decorations obviously implies that $x = \xi$ satisfies $x_i > 0$.

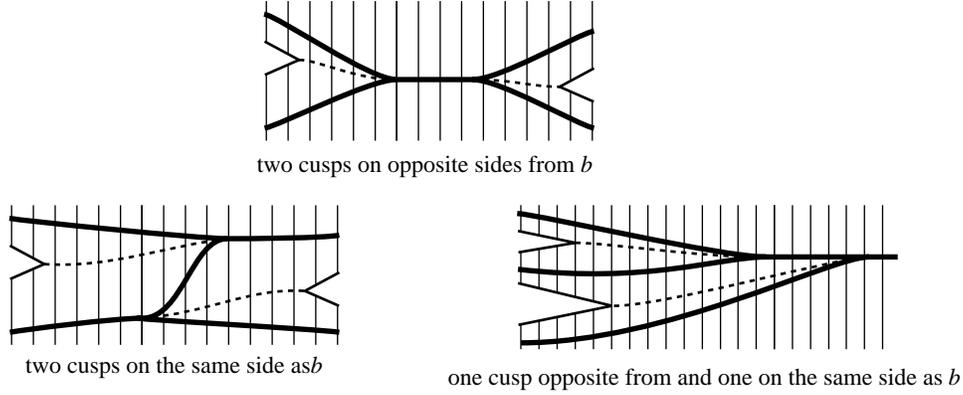


Figure 24: If the cusps v_i, v_j of σ are located at opposite endpoints s, s' of a branch b , then there is an inequality $I_{ibj}(\sigma)$ of the form $x_i \pm x_j <> \langle \rho_{ibj}(\sigma), \mathcal{F}_v \rangle$, where the right hand side is signed intersection number with \mathcal{F}_v of the path $\rho_{ibj}(\sigma)$ from c_i to c_j consisting of a section of T_i concatenated with b concatenated with a section of T_j . The \pm sign is $+$ if the cusps v_i, v_j are both on opposite sides from b , or both on the same side as b , at the switches s, s' respectively; the sign is $-$ otherwise. The inequality is $<$ if v_i is on the opposite side from b at s , and $>$ otherwise.

Second, for each $i, j \in \mathcal{I}$ and each branch b of σ with distinct endpoints $s \neq s'$ such that v_i, v_j are located at s, s' respectively, there is an inequality $I_{ibj}(\sigma)$ which is a numerical expression of the fact that the boundary points of the separatrices ξ_i, ξ_j lie on the shunts over the endpoints s, s' of b .

In discussing the inequality $I_{ibj}(\sigma)$ we shall need the path $\rho_{ibj}(\sigma)$ which is a concatenation of three pieces: (i) a path T_i from c_i to s , (ii) the branch b from s to s' , and (iii) a path in T_j from s' to c_j . Either of the two outer pieces may degenerate to points, depending on which of v_i, v_j are degenerate cusps of σ . We orient $\rho_{ibj}(\sigma)$ so that at its initial point c_i it points in the same direction as the cusp c_i . We have:

Fact 7.4.2. *As σ varies by isotopy through decoration respecting carrying injections, the cusps v_i, v_j continue to be located at opposite ends of the same branch b , and the path $\rho_{ibj}(\sigma)$ varies by a homotopy that preserves endpoints and orientation. It follows that the signed intersection number $\langle \rho_{ibj}(\sigma), \mathcal{F}_v \rangle$ is constant as σ varies.*

As with the paths $\rho_{ij}(\sigma)$, each of the paths $\rho_{ibj}(\sigma)$ may be chosen to be transverse to \mathcal{F}_v , except possibly at the points s, s' , depending on whether the cusps v_i, v_j are on the same or opposite side of b at s, s' , respectively.

The description of the inequality $I_{ibj}(\sigma)$ breaks into three cases, depending on how many of v_i, v_j are on the same side as b at s, s' , respectively; see Figure 24.

The inequality $I_{ibj}(\sigma)$, Case 1: If v_i, v_j are both on opposite sides from b at the switches s, s' , respectively, then the inequality is

$$\begin{aligned} x_i + x_j &< \ell_i(\sigma) + \ell_b(\sigma) + \ell_j(\sigma) && I_{ibj}(\sigma), \text{ Case 1} \\ &= \langle \rho_{ibj}(\sigma), \mathcal{F}_v \rangle \end{aligned}$$

This inequality holds when $x = \xi$ because then the left hand side is equal to $\ell_i(\sigma) + \ell_j(\sigma)$. Invariance of this inequality under isotopy of σ through carrying injections follows from Fact 7.4.2.

The inequality $I_{ibj}(\sigma)$, Case 2: Suppose one of v_i, v_j is on the opposite side of b at its respective switch s, s' , and the other is on the same side as b . Changing notation if necessary we assume that v_i is on the opposite side as b and v_j is on the same side. In this case the inequality is

$$\begin{aligned} x_i - x_j &< \ell_i(\sigma) + \ell_b(\sigma) - \ell_j(\sigma) && I_{ibj}(\sigma), \text{ Case 2} \\ &= \langle \rho_{ibj}(\sigma), \mathcal{F}_v \rangle \end{aligned}$$

This holds when $x = \xi$ because then the left hand side equals $\ell_i(\sigma) - \ell_j(\sigma)$. Invariance under isotopy of σ through carrying injections follows from Fact 7.4.2.

The inequality $I_{ibj}(\sigma)$, Case 3: Suppose that both of the cusps v_i, v_j are on the same side as b at their respective switches. In this case the inequality is

$$\begin{aligned} x_i + x_j &> \ell_i(\sigma) - \ell_b(\sigma) + \ell_j(\sigma) && I_{ibj}(\sigma), \text{ Case 3} \\ &= \langle \rho_{ibj}(\sigma), \mathcal{F}_v \rangle \end{aligned}$$

When $x = \xi$, the left hand side equals $\ell_i(\sigma) + \ell_j(\sigma)$, and so the inequality holds. Invariance under isotopy of σ through carrying injections follows from Fact 7.4.2.

This completes the definition of the system of equations $X_\sigma = E_\sigma \cup I_\sigma$, and the proof of invariance of X_σ under decoration respecting carrying isotopy of σ . As noted earlier, for each $\xi \in \Xi$ we have therefore defined the system $X_\xi = X_{\sigma(\xi)}$ and we have proved that $\overset{\circ}{c}(\xi)$ is contained in the solution set of X_ξ .

Polyhedral structure of $c(\xi)$. One direction is left in the proof that $\overset{\circ}{c}(\xi)$ equals the solution set of X_ξ : we must prove that the solution set is contained in $\overset{\circ}{c}(\xi)$. The key to the proof is the fact that X_ξ is convex, being the intersection of affine subspaces and affine half-spaces. Suppose that $\eta \in \Xi$ is another point in the solution set of the system X_ξ . By convexity the entire segment $\overline{\xi\eta}$ is contained in

the solution set. Letting x vary along this segment from ξ to η , we shall show that the isotopy class of the decorated train track $\sigma(x)$ is constant. We identify $\sigma(x)$ with its image $f_x(\sigma(x))$ in ν_0 under the decoration respecting carrying injection f_x .

As x varies in $\xi\eta$, first we note that there are only two ways that the *undecorated* isotopy class of $\sigma(x)$ could change. One way is if, for some nongeneric switch s , two cusps v_i, v_j located at s split apart so that they are located at different switches; but this occurrence is prevented by the requirement that the equation $E_{ij}(\sigma(\xi))$ holds true, which forces v_i, v_j to stay located at the same switch. The other way that $\sigma(x)$ could change is that, for some branch b of $\sigma(x)$ with ends at distinct points, as x varies b shrinks to a point, so that cusps v_i, v_j which were previously located at distinct ends of b are now located at the same switch; but this is prevented by the requirement that the strict inequality $I_{ibj}(\sigma(\xi))$ holds true, which forces v_i, v_j to stay located at opposite ends of b , and hence b cannot shrink to a point. This shows that the undecorated isotopy class of $\sigma(x)$ is invariant.

Next we note that as x varies in $\overline{\xi\eta}$, the set of coordinates $i \in \mathcal{I}$ for which $x_i = 0$ obviously remains invariant, because for all $i \in \mathcal{I}$ we have $\xi_i = 0$ if and only if $\eta_i = 0$ if and only if the degeneracy equation $x_i = 0$ is in the system X_ξ . But this implies that the subset of \mathcal{I} for which the cusp v_i of $\sigma(x)$ is degenerate is a constant subset, and hence the *decorated* isotopy class of $\sigma(x)$ is invariant. Thus, for all $x \in \overline{\xi\eta}$, and in particular for $x = \eta$, the decorated train track $\sigma(x)$ is isotopic to $\sigma(\xi)$.

We have proved that $\overset{\circ}{c}(\xi)$ is equal to the solution set of the family of equations and inequalities X_ξ , and so $\overset{\circ}{c}(\xi)$ is the interior of a finite sided, convex polyhedron $c(\xi)$. To prove that $c(\xi)$ is a convex polyhedral cell it suffices to prove that $\overset{\circ}{c}(\xi)$ is bounded in Ξ . Arguing by contradiction, if $\overset{\circ}{c}(\xi)$ were unbounded then there would be an unbounded, straight ray $\rho: [0, \infty) \rightarrow \Xi \approx [0, \infty)^{\mathcal{I}}$ starting from $\rho(0) = \xi$, so that ρ is completely contained in $\overset{\circ}{c}(\xi)$. If ρ were a generic ray then we could apply Claim 4.3.1 to conclude that the set of train tracks labelling the equivalence classes intersected by ρ contain the train tracks in some infinite splitting sequence, namely, some train track expansion of \mathcal{F} , contradicting that this set of decorated train tracks consists of the singleton $\{\sigma(\xi)\}$. However, the proof of Claim 4.3.1 works in the nongeneric case as well. One need only notice that starting from $x = \xi$, as x moves along ρ towards infinity, at least one component of the expanding separatrix family x gets longer and longer without bound, and hence at some moment $x = \rho(s)$ there must be two components of x whose boundary points lie on the same shunt of x , whereas the two corresponding components of ξ do not have endpoints on the same shunt of ξ ; at that moment the decorated train track $\sigma(x)$ is isotopically distinct from $\sigma(\xi)$.

We have now proved that each equivalence class $\overset{\circ}{c}(\xi)$ is the interior of a closed, convex polyhedral cell $c(\xi)$. Item (1) of Theorem 7.2.2 follows by observing that

the dimension of $c(\xi)$ is equal to the number of switches s of $\sigma(\xi)$ such that the shunt $t_s(\xi)$ which contains $f_\xi(\sigma(\xi))$ does not contain any cusp of ν_0 , which equals the number of switches of $\sigma(\xi)$ at which no degenerate cusp of $\sigma(\xi)$ is located. Item (2) follows by observing that the stratum $\Xi_{\mathcal{J}}$ containing $c(\xi)$ is determined by $\mathcal{J} = \{i \in \mathcal{I} \mid \xi_i \text{ is nondegenerate}\}$, but this set corresponds to the set of nondegenerate cusps of the decorated train track $\sigma(\xi)$. Item (3) will be proved below.

7.5 The face relation in the expansion complex.

To complete the proof of the Cell Structure Theorem 7.2.2: we must show that the equivalence classes $\overset{\circ}{c}(\xi)$ of Ξ form the open cells of a linear cell decomposition, or in other words we must show that the boundary of the closed cell $c(\xi)$ is a union of closed cells of lower dimensions; and we must show that the face relation among closed cells satisfies item (3) of Theorem 7.2.2. Both of these tasks require a study of the face relation among train tracks, with and without decorations.

The face relation on train tracks. As motivation we start by reconsidering several familiar constructions in the context of face inclusions.

Consider the examples of the face relation which were discussed after the statement of Theorem 7.2.2. Let σ be a generic train track and b a branch. We say that b is *collapsible* if b has endpoints at distinct switches and b is not the overlay image of a side of any component of $\mathcal{C}(S - \sigma)$. To be precise, every sink branch is collapsible, a transition branch is collapsible if and only if its ends are located at distinct switches, and a source branch is collapsible if and only if it is of type LL or RR. If b is a sink branch or a transition branch with ends at distinct switches, collapsing b results in a train track σ' which satisfies the face relation $\sigma' \sqsubset \sigma$; we have described this previously by saying that σ' is obtained by combing σ along b . If b is a DD source branch with $D \in \{L, R\}$, collapsing b results in a train track σ' which satisfies the face relation $\sigma' \sqsubset \sigma$, and in this case $\sigma' \succ \sigma$ is a splitting of parity D : the result of collapsing b is to produce a nongeneric switch in σ' with two branch ends on each side, and then if one does a generic uncombing of σ' at this switch, one can follow that with an elementary splitting of parity D to recover σ . The difference between the current discussion and the discussion given after the statement of Theorem 7.2.2 is that currently the train tracks are undecorated.

One can generalize the above examples by taking σ to be generic, and taking F to be a pairwise disjoint collection of collapsible branches. By collapsing each component of F we obtain a train track σ' which gives another example of the face relation $\sigma' \sqsubset \sigma$. The codimension of this example is equal to the cardinality of F .

One more class of familiar examples occurs when F is a neighborhood in σ of the basin of a sink branch, or a disjoint union of such basins. In fact, in the case where

F is a neighborhood of the union of all basins, then the result of collapsing σ along F is exactly the same as what we have previously called the complete combing of σ .

Collapsing a forest. We shall define a face relation $\sigma' \sqsubset \sigma$ on undecorated train tracks by choosing a certain “collapsible forest” F in a train track σ , and letting σ' be obtained from σ by collapsing each component of F to a point. We now turn to a discussion of collapsible forests.

Given a train track σ and a subcomplex $F \subset \sigma$ which is a disjoint union of nondegenerate trees, and assume that no point of the frontier of F relative to σ is a switch of σ ; we call F a *forest* in σ , or a *tree* if F has one component. Define the *heart* of a forest F to be the union of all branches of σ that are contained in F . A tree is *heartless* if its heart is empty, equivalently, it contains no more than one switch and so is either contained in the interior of a branch or is a regular neighborhood of a single switch. A tree that is not heartless deformation retracts onto its heart.

A forest F in σ is said to be *collapsible* if the following is true:

- (A) Letting $\{F_n\}$ be the set of components of F , there is a pairwise disjoint collection of smoothly embedded rectangles $\{R_n\}$, called *frames* for the components $\{F_n\}$, such that the following properties hold. For each n there is a smooth parameterization $f: I \times J \rightarrow R_n$, $I = J = [0, 1]$, such that denoting $\partial_h R_n = f(I \times \partial J)$ and $\partial_v R_n = f(\partial I \times J)$, and defining the vertical direction at a point $f(x, y) \in R$ to be the tangent line to $f(x \times J)$, we have
- $F_n = \sigma \cap R_n$.
 - σ is disjoint from $\partial_h R_n$ and is transverse to $\partial_v R_n$, and hence $\text{Fr}(F_n) = F_n \cap \partial R_n = F_n \cap \partial_v R_n$.
 - For any point $x \in F_n$ the tangent line $T_x \sigma$ is not vertical in R_n .

Up to isotopy there is a unique train track σ' such that $\sigma - \cup R_n = \sigma' - \cup R_n$ and such that for each n , $\sigma' \cap R_n$ is a collapsible tree with frame R_n and with a single switch in R_n (see Figure 25); we say that σ' is obtained from σ by *collapsing* F . There is a smooth map $q: \sigma \rightarrow \sigma'$, called a *forest collapse*, satisfying the following properties: q is homotopic to the inclusion $\sigma \hookrightarrow S$, each component of the heart of F collapses to a switch of σ' , q is otherwise one-to-one, and q is an immersion on the complement of the heart. A forest collapse $q: \sigma \rightarrow \sigma'$ is uniquely defined up to postcomposition by a diffeomorphism $\sigma' \rightarrow \sigma'$ smoothly isotopic to the identity map; the proof of uniqueness is similar to the proof of Lemma 3.5.2. A forest collapse $q: \sigma \rightarrow \sigma'$ therefore sets up a well-defined bijection $\text{cusps}(\sigma) \leftrightarrow \text{cusps}(\sigma')$.

Notice that a nontrivial forest collapse $q: \sigma \rightarrow \sigma'$ is *not* a carrying map, because it is not an immersion—for each point x in the heart of the forest, the tangent line

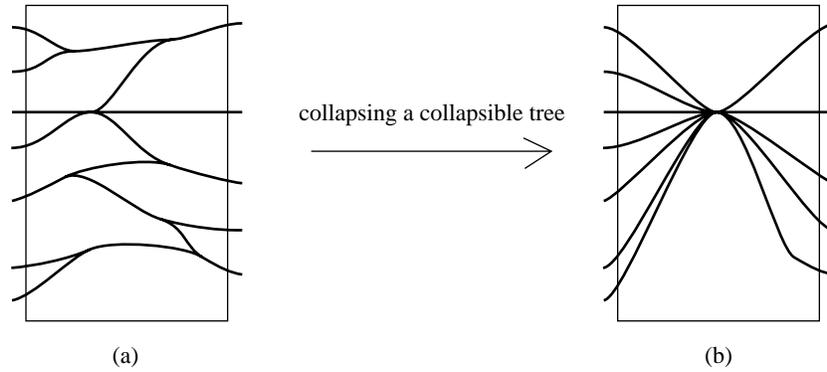


Figure 25: (a) shows a collapsible tree with rectangular frame in a train track σ . Up to isotopy there is a unique way to collapse the tree (b).

of σ at x is mapped trivially by the derivative of q . In particular, given a train path γ in σ , the image path $q \circ \gamma$ is not necessarily a train path in σ' . However, the only reason that $q \circ \gamma$ fails to be a train path is that, while γ is traveling through the forest, $q \circ \gamma$ slows down to a dead stop for a while before starting up again. The important point is that $q \circ \gamma$ does *not* make any illegal turns. In particular, there is a C^0 perturbation of $q \circ \gamma$ that is a train path: where $q \circ \gamma$ came to a dead stop before the perturbation, after the perturbation it can slow down to a crawl for a while before speeding up again, restoring the property of being an immersion.

We now define the face relation $\sigma' \sqsubset \sigma$ on (undecorated) train tracks to mean that σ' is, up to isotopy, obtained by collapsing some collapsible forest of σ . It is clear that this relation is natural in the sense that if $f, f': S \rightarrow S$ are isotopic homeomorphisms, we have $\sigma' \sqsubset \sigma$ if and only if $f'(\sigma') \sqsubset f(\sigma)$.

As a final remark on the undecorated face relation, we noted earlier that if $\sigma' \sqsubset \sigma$ then a forest collapse map $q: \sigma \rightarrow \sigma'$ is not a carrying map, but it is still interesting to ask whether σ' carries σ . The answer is no, in general. For a counterexample, if b is a sink branch or a collapsible source branch of σ , and if σ' is obtained by collapsing b , then σ' does not carry σ ; notice, however, that in each of these cases the switch s' that is the image of b can be completely uncombed producing a train track that does carry σ . In general, although the relation $\sigma' \sqsubset \sigma$ need not imply that σ carries σ' , it does imply that σ' is comb equivalent to a train track that carries σ : for each nongeneric switch s of σ' that is the image of the heart of some component of the collapsing forest, if s is slightly uncombed to produce a sink branch connecting two semigeneric switches, the resulting train track σ'' is comb equivalent to σ' and σ'' homotopically carries σ . It is easy to see that the composed bijection $\text{cusps}(\sigma') \leftrightarrow \text{cusps}(\sigma'') \leftrightarrow \text{cusps}(\sigma)$, the first induced by comb equivalence

and the second by carrying, agrees with the bijection induced by the forest collapse $q: \sigma \rightarrow \sigma'$.

Another description of collapsible forests. There is an equivalent description of collapsible forests which makes them easier to recognize in some contexts; we will need this description below. We claim that a forest $F \subset \sigma$ is collapsible if and only if it satisfies the following condition:

(B) F does not contain the overlay image of any side of any component of $\mathcal{C}(S - \sigma)$.

This description makes it clear, for example, that if b is a branch of a generic train track σ and if F is a regular neighborhood of b in σ then F is a collapsible forest if and only if the branch b is collapsible in the motivational sense discussed earlier.

The equivalence of (A) and (B) is proved as follows.

Proof of (A) \implies (B). Assuming (A), without loss of generality we may take F to have a single component with frame R . For each component D of $\mathcal{C}(R - F)$, note that D is a disc-with-corners having one of two types: either D has two corners and one cusp, and in this case $D \cap F$ is a union of two train paths meeting at the cusp, coinciding with two of the three sides of D ; or D has four corners and no cusps, in which case $D \cap F$ is a smooth train path coinciding with one of the four sides of D . In either case, D does not contain two cusps, and so D does not contain the overlay image of any side of any component of $\mathcal{C}(S - \sigma)$. This proves (B). \diamond

Proof of (B) \implies (A). Assuming (B), again without loss of generality we may assume that F has a single component. We must construct the frame R . To start with we take R to be a smoothly embedded disc in S whose boundary is transverse to σ such that $F = \sigma \cap R$ is properly embedded in R . We must alter the smooth structure on R , introducing four corners so as to make R a rectangle, and so that R becomes a frame for F . Consider a component D of $\mathcal{C}(R - F)$, and note that $\partial D = \alpha \cup \beta$ where α is a subsegment of ∂R , β is a topological arc embedded in F , and $\alpha \cap \beta = \partial\alpha = \partial\beta$. Note that D has exactly two corners, at the two points $\partial\alpha = \partial\beta$, and it may also have cusps in the interior of β . If there are two or more cusps on β then the segment between two consecutive cusps on β is the overlay image of a side of $\mathcal{C}(S - \sigma)$, contradicting (B). It follows that there is either zero or one cusp on β . There are now two cases to consider. In the first case D has two corners and one cusp, and so D has Euler index zero, and we can impose a transverse pair of line fields on D , one parallel to the two segments of β and one parallel to α . If D has no cusp then alter the smooth structure of R by introducing two corners in the interior of α , subdividing α into three segments, and converting D into a disc with four corners; after this alteration, D has Euler index zero, and we

can impose a transverse pair of line fields on D , one of which is parallel to β and to the middle segment of α , and the other of which is parallel to the two end segments of α . Employing additivity of Euler index it follows that the index of R equals zero, and so R is a disc with four corners, that is, a rectangle. Each component of $\mathcal{C}(R - F)$ has a line field that is transverse to F , and we can adjust that line field so that it is actually perpendicular to F , with respect to some fixed Riemannian metric on S . The two line fields defined on the components of $\mathcal{C}(R - F)$ now fit together continuously to give a transverse pair of line fields on R , one of which is parallel to F and to one pair of opposite sides of R , and the other of which is transverse to F and parallel to the other pair of opposite sides of R . These two line fields are not necessarily uniquely integrable; indeed, assuming F has at least one switch the line field parallel to F is definitely not uniquely integrable. However, by a C^0 perturbation we may make the two line fields uniquely integrable, after which they fit together to define a product structure on R . The line field which, before perturbation, was parallel to F is no longer parallel. However, the line field which was transverse to F is still transverse, and so F is nowhere vertical on R . This establishes (A). \diamond

The face relation on decorated train tracks. We revert to the formal, ordered pair notation for decorated train tracks: consider two decorated train tracks (σ, D) , (σ', D') , with $D \subset \text{cusps}(\sigma)$ and $D' \subset \text{cusps}(\sigma')$. We define $(\sigma', D') \sqsubset (\sigma, D)$ if $\sigma' \sqsubset \sigma$ and if the subset $D \subset \text{cusps}(\sigma)$ is contained in the image of the subset $D' \subset \text{cusps}(\sigma')$ under the bijection $\text{cusps}(\sigma') \leftrightarrow \text{cusps}(\sigma)$ defined above. For example, if $\sigma = \sigma'$ and if $D \subset D'$ then we have $(\sigma', D') \sqsubset (\sigma, D)$.

As in the undecorated case, the face relation among decorated train tracks is natural: if $f, f': S \rightarrow S$ are isotopic homeomorphisms, then $(\sigma, D) \sqsubset (\sigma', D')$ if and only if $f(\sigma, D) \sqsubset f'(\sigma', D')$.

The linear cell decomposition of Ξ . We are now ready to complete the proof of Theorem 7.2.2, by showing that the cells of Ξ form a linear cell decomposition, and proving property (3). Both of these tasks are accomplished by proving that for any two distinct closed cells $c(\sigma) \neq c(\sigma')$ of Ξ the following statements are equivalent:

- (i) $c(\sigma') \subset c(\sigma)$
- (ii) $\overset{\circ}{c}(\sigma') \subset \partial c(\sigma)$
- (iii) $\overset{\circ}{c}(\sigma') \cap \partial c(\sigma) \neq \emptyset$
- (iv) $\sigma' \sqsubset \sigma$.

Assuming for the nonce that these are equivalent we proceed as follows. Define the n -skeleton of Ξ to be the union of the closed n -dimensional cells. We must prove that the boundary of a closed n -cell $c(\sigma)$ is a union of closed cells of dimension $\leq n - 1$. Since $\partial c(\sigma)$ an $n - 1$ dimensional topological sphere, it suffices to prove that $\partial c(\sigma)$ is a union of cells, but this follows immediately from the equivalence of (ii) and (iii). Item (3) of Theorem 7.2.2 follows from the equivalence of (i) and (iv), together with the fact observed earlier that the relation \sqsubset is natural.

Now we prove the equivalence of (i)–(iv). To prove (i) \implies (ii), since $c(\sigma') \neq c(\sigma)$, and since interiors of cells form a decomposition of Ξ , it follows that $\overset{\circ}{c}(\sigma') \cap \overset{\circ}{c}(\sigma) = \emptyset$, and as assuming (i) it follows that $\overset{\circ}{c}(\sigma') \subset \partial c(\sigma)$. Obviously (ii) \implies (iii).

Proof of (iii) \implies (iv). Choose $\xi \in \overset{\circ}{c}(\sigma') \cap \partial c(\sigma)$ and $\eta \in \overset{\circ}{c}(\sigma)$, and so $\sigma' = \sigma(\xi)$ and $\sigma = \sigma(\eta)$. Let $[\xi, \eta]$ denote the straight line segment in Ξ from ξ to η , and let $(\xi, \eta) = [\xi, \eta] - \{\xi\}$. By convexity we have $(\xi, \eta) \subset \overset{\circ}{c}(\sigma(\eta))$. We may choose the train tracks $\sigma(\xi), \sigma(\eta)$ in their isotopy classes so that the inclusion maps $\sigma(\xi) \hookrightarrow \nu_0, \sigma(\eta) \hookrightarrow \nu_0$ are decoration respecting, homotopic carrying injections. For $x \in (\xi, \eta]$ we may identify the train track $\sigma(x)$ with its image in ν_0 under the decoration respecting, homotopic carrying injection $f_x: \sigma(x) \rightarrow \nu_0$.

Although not strictly necessary for the proof, we may assume that the train tracks $\sigma(x)$ vary continuously in the parameter $x \in (\xi, \eta]$, and as x approaches ξ we wish to describe how the decorated train tracks $\sigma(x)$ degenerate to $\sigma(\xi)$. We shall show that degeneration can be described as the collapse of a certain forest and the degeneration of certain cusps. The degeneration is controlled by studying the behavior of the system X_x as $x \rightarrow \xi$. As x varies in $(\xi, \eta]$, the decorated isotopy class of the train tracks $\sigma(x)$ is constant, and so the form of the equations and inequalities in the system X_x is constant. However, the difference between the two sides of an inequality in X_x can vary for $x \in (\xi, \eta]$, and this variation will determine the degeneration.

We shall identify the set of branches of $\sigma(\eta)$ with the set of branches of $\sigma(x)$ for any $x \in (\xi, \eta]$; let this set be denoted B , and we use its elements as indexes of the inequalities $I_{ibj}(\sigma(\eta))$.

Each equation in E_η being a closed condition, these equations all hold for $x \in [\xi, \eta]$. So, as x approaches ξ in $(\xi, \eta]$, the only thing that can happen to the system X_x is that one or more of the strict inequalities in I_η degenerate to equations at $x = \xi$. Let B' be the set of branches $b \in B$ such that, for some $i, j \in \mathcal{I}$, the strict inequality $I_{ibj}(\sigma(x))$ degenerates to an equation as $x \rightarrow \xi$, that is, the difference of the two sides of $I_{ibj}(\sigma(x))$ approaches zero. Let F be a regular neighborhood in $\sigma(\eta)$ of $\bigcup B'$, and so F deformation retracts to $\bigcup B'$. Note that $\bigcup B'$ is the union of all complete branches of $\sigma(\eta)$ that are contained in F . We shall show

that F is a collapsible forest and that, ignoring decorations, $\sigma(\xi)$ is obtained from $\sigma(\eta)$ by collapsing F . Having done this, consideration of decorations will show that $\sigma(\xi) \sqsubset \sigma(\eta)$.

First note that for each $b \in B'$ we have

$$\lim_{x \rightarrow \xi} \int_{f_x(b)} \mathcal{F}_v = 0 \quad (**)$$

To prove the claim, consider the inequality $I_{ibj}(\sigma(\eta))$, in any of cases 1, 2, or 3 as described earlier. For each $x \in \overset{\circ}{c}(\eta)$, this inequality can also be written in the form $I_{ibj}(\sigma(x))$, from which it follows that the difference between the two sides of the inequality is equal in absolute value to $\ell_b(\sigma(x)) = \int_{f_x(b)} \mathcal{F}_v$, which therefore goes to zero as $x \rightarrow \xi$.

First we prove that F is a forest of $\sigma(\eta)$. By construction of F , no point in the frontier of F relative to $\sigma(\eta)$ is a switch. If F is not a forest then there exists a simple closed curve γ in $\bigcup B'$. Letting $\gamma(x) = f_x(\gamma) \subset \sigma(x)$, from $(*)$ it follows that $\lim_{x \rightarrow \xi} \int_{\gamma(x)} \mathcal{F}_v = 0$. But this is impossible, for the following reasons. The curve γ is homotopically nontrivial in $\sigma(\eta)$, and so the curves $\gamma(x)$ are homotopically nontrivial in ν_0 . But clearly there is a positive lower bound for the value of $\int_c \mathcal{F}_v$ over all homotopically nontrivial simple closed curves c in ν_0 , a contradiction which proves that F is a forest.

Next we prove that F is collapsible, and for this we will use the characterization (B) of collapsibility, by proving that $\bigcup B'$ does not contain the overlay image of any side of $\mathcal{C}(S - \sigma(\eta))$. Suppose otherwise, that there is a side of $\mathcal{C}(S - \sigma(\eta))$ whose overlay image is contained in $\bigcup B'$. Letting s be the corresponding side of $\mathcal{C}(S - \sigma(x))$, it follows that

$$\lim_{x \rightarrow \xi} \int_s \mathcal{F}_v = 0$$

We shall derive a contradiction by proving that there is a positive lower bound for $\int_s \mathcal{F}_v$ where s is any side of $\mathcal{C}(S - \sigma)$ and σ is any train track which includes into ν_0 by a homotopic carrying injection. Since the carrying injection $\sigma \hookrightarrow \nu_0$ is a homotopy equivalence, it follows that there is a bijection between the sides of $\mathcal{C}(S - \sigma)$ and the sides of ν_0 , so that if s is a side of $\mathcal{C}(S - \sigma)$ and if s' is the corresponding side of ν_0 , then s' is homotopic along leaves of \mathcal{F}_v to some subsegment of s ; in particular, $\int_s \mathcal{F}_v \geq \int_{s'} \mathcal{F}_v$. But ν_0 has only finitely many sides, so the desired positive lower bound is the minimum over all sides of ν_0 of the \mathcal{F}_v transverse measure of that side.

This completes the proof that F is a collapsible forest.

Now we must check that the (undecorated) train track obtained from $\sigma(\eta)$ by collapsing F is isotopic to $\sigma(\xi)$. For this we turn directly to an examination of each $x \in [\xi, \eta]$ as a separatrix family in ν_0 , the corresponding set of shunts $t_s(x)$

in bijective correspondence with the switches s of $\sigma(x)$, and the corresponding set of rectangles $R_b(x)$ in bijective correspondence with the branches b of $\sigma(x)$. In particular, as $x \rightarrow \xi$ in $[\xi, \eta]$, we study how the shunts of x limit to the shunts of ξ , and we study how the rectangles $R_b(x)$ behave for $b \in B'$.

For each shunt $t_s(\xi)$ choose a thin rectangle $R_s(\xi)$ in ν_0 which contains $t_s(\xi)$ as a vertical fiber distinct from the two vertical sides, so the interior of each horizontal side of $R_s(\xi)$ contains one of the endpoints of $t_s(\xi)$. Choose these rectangles so thin that they have pairwise disjoint interiors and vertical sides; there may exist points $p \in \xi$ such that two horizontal sides of rectangles contain p , one approaching from each side of ξ . Let $R'_s(\xi)$ be a rectangle obtained from $R_s(\xi)$ by shaving off a thin rectangle containing each horizontal side of $R_s(\xi)$, so each horizontal fiber of $R'_s(\xi)$ is also a horizontal fiber of $R_s(\xi)$. For $x, y \in \Xi$ let $|x - y|$ denote the sup norm $\sup_{i \in \mathcal{I}} |\text{Length}(x_i) - \text{Length}(y_i)|$. For $i \in I$ let $\Delta_i(x, y)$ be the symmetric difference of the finite separatrices x_i and y_i , and so $|x - y| = \sup \text{Length}(\Delta_i(x, y))$. By choosing a number $\epsilon > 0$ sufficiently small, it follows that if $x \in [\xi, \eta]$ and if $|\xi - x| < \epsilon$ then the following hold:

$$(1) \text{Length}(\Delta_i) < \epsilon \text{ for all } i \in \mathcal{I},$$

$$(2) \bigcup_{i \in \mathcal{I}} \Delta_i \subset \bigcup_{s \in \text{Sw}(\sigma(\xi))} R_s(\xi)$$

$$(3) \text{ Given } b \in B, \text{ we have } R_b(x) \subset \bigcup_{s \in \text{Sw}(\sigma(\xi))} R_s(\xi) \text{ if and only if } b \in B'.$$

Consider each train track $\sigma(x)$, whose injection into ν_0 is identified with the map $f_x: \sigma(x) \rightarrow \nu_0$; we may assume that $\sigma(x)$ is chosen so that its intersection with each rectangle $R_s(\xi)$ is contained in $R'_s(\xi) - \partial_h R'_s(\xi)$. It follows that $\sigma(x) \cap R'_s(\xi)$ is a collapsible tree with frame $R'_s(\xi)$, for each $s \in \text{Sw}(\sigma(\xi))$, and so

$$F'_x = \bigcup_{s \in \text{Sw}(\sigma(\xi))} \sigma(x) \cap R'_s(\xi)$$

is a collapsible forest. Moreover, it is clear from the construction that $\sigma(\xi)$ is the train track obtained from $\sigma(x)$ by collapsing F'_x . Certain components of F'_x are trivial in that they do not contain any branches of $\sigma(x)$, and are just regular neighborhoods of switches of $\sigma(x)$; letting F_x be obtained from F'_x by removing the trivial components, it still follows that $\sigma(\xi)$ is obtained from $\sigma(x)$ by collapsing F_x . From item (3) above it follows that for each $b \in B$, we have $f_x(b) \subset F'_x$ if and only if $b \in B'$, which occurs if and only if $b \in F$. Thus, under the identification of $\sigma(x)$ with $\sigma(\eta)$, both F_x and F are regular neighborhoods of $\bigcup B'$, and so by uniqueness of regular neighborhoods it follows that $\sigma(\xi)$ is obtained from $\sigma(\eta)$ by collapsing F .

Finally, we note that the degeneration equations $x_i = 0$ that are in E_η hold for all $x \in (\xi, \eta]$, and since each of these equations is a closed condition it also holds for $x = \xi$. This implies that the set of degenerate cusps of $\sigma(\eta)$ corresponds to a subset of the degenerate cusps of $\sigma(\xi)$, under the bijection $\text{cusps}(\sigma(\eta)) \leftrightarrow \text{cusps}(\sigma(\xi))$ induced by the forest collapse $\sigma(\eta) \rightarrow \sigma(\xi)$. This shows that $\sigma(\xi) \sqsubset \sigma(\eta)$, completing the proof of (iii) \implies (iv).

Proof of (iv) \implies (i). Assuming $\sigma' \sqsubset \sigma$ we must show $\overset{\circ}{c}(\sigma')$, the solution set of $X_{\sigma'}$, is contained in the closure of the solution set of X_σ . Fix a forest F of σ whose collapse produces σ' . We may assume that both σ and σ' are included into ν_0 by decoration preserving, homotopic carrying maps.

Consider a degeneracy equation $x_i = 0$ in E_σ . There is one such equation associated to each degenerate cusp v_i of σ , and since $\sigma' \sqsubset \sigma$ it follows that the corresponding cusp v'_i of σ' is also degenerate, and hence $x_i = 0$ holds in $\overset{\circ}{c}(\sigma')$. Consider next a nondegeneracy inequality $x_i > 0$ in I_σ ; evidently $x_i \geq 0$ holds throughout Ξ , and in particular in $\overset{\circ}{c}(\sigma')$.

It suffices therefore to consider the equations $E_{ij}(\sigma)$ and strict inequalities $I_{ibj}(\sigma)$ in X_σ . We shall show that to each equation $E_{ij}(\sigma)$ in X_σ there is an identical equation $E_{ij}(\sigma')$ in $X_{\sigma'}$. We shall also show that to each strict inequality $I_{ibj}(\sigma)$ there is either an identical inequality $I_{ibj}(\sigma')$ in $X_{\sigma'}$ or there is an equation $E_{ij}(\sigma')$ in $X_{\sigma'}$ which is obtained from $I_{ibj}(\sigma')$ by changing the strict inequality sign to an equation sign; the distinction between the two cases is determined by whether $b \subset F$.

Because we have already taken care of the degeneracy equations and inequalities, and because the forms of the remaining equations and inequalities are invariant under isotopy through carrying injections, even those which do not respect decorations, we may assume that σ, σ' are contained in the interior of ν_0 ; this will simplify the following discussion.

First we alter the choice of σ in a particular way. For each switch s of σ' , choose a small rectangle R_s whose interior contains s ; by choosing each R_s sufficiently small it follows that the rectangles R_s are pairwise disjoint. From the definition of the relation $\sigma' \sqsubset \sigma$, it follows that we may choose σ in its isotopy class so that

$$\sigma - \bigcup_{s \in \text{Sw}(\sigma')} R_s = \sigma' - \bigcup_{s \in \text{Sw}(\sigma')} R_s$$

and so that for each $s \in \text{Sw}(\sigma')$, either $\sigma' \cap R_s = \sigma \cap R_s$, or R_s is a frame of a component of the collapsing forest F . From this description it follows that the collapse map $q: \sigma \rightarrow \sigma'$ can be extended to a map $Q: \nu_0 \rightarrow \nu_0$ that is the identity outside of $\bigcup R_s$, that is one-to-one except for the identifications made by q , and

that has the property that for each cusp v_i of σ and corresponding cusp v'_i of σ' , the map Q takes the cusp triangle T_i of σ associated to v_i to the cusp triangle T'_i of σ' associated to v'_i . In particular, any path in T_i from c_i to v_i is taken by Q to a path in T'_i from c_i to v'_i ; this will help in understanding how Q acts on the paths $\rho_{\dots}(\sigma)$.

Consider now an equation $E_{ij}(\sigma)$ in X_σ , corresponding to a pair of cusps v_i, v_j of σ that are located at the same switch s of σ . The equation has the form

$$x_i \pm x_j = \langle \rho_{ij}(\sigma), \mathcal{F}_v \rangle$$

Note that v'_i, v'_j are also located at a single switch of σ' , and so the system $X_{\sigma'}$ has an equation of the form

$$x_i \pm x_j = \langle \rho_{ij}(\sigma'), \mathcal{F}_v \rangle$$

The \pm signs of these two equations are identical, because of the observation that v_i, v_j are on the same side of s if and only if v'_i, v'_j are on the same side of s' . Also, the map Q obviously takes $\rho_{ij}(\sigma)$ to a path which is homotopic rel endpoints to $\rho_{ij}(\sigma')$, because Q takes T_i to T'_i and T_j to T'_j , and since $\rho_{ij}(\sigma), \rho_{ij}(\sigma')$ have the same initial orientation at the cusp c_i it follows that the right hand sides of the above two equations are identical. The equations $E_{ij}(\sigma), E_{ij}(\sigma')$ are therefore identical, and so the solution set $\overset{\circ}{c}(\sigma')$ of the system $X_{\sigma'}$ is contained in the solution set of $E_{ij}(\sigma)$.

Consider next an inequality $I_{ibj}(\sigma)$ in X_σ , corresponding to a pair v_i, v_j of cusps of σ located at opposite endpoints of a branch b of σ . The inequality has the form

$$x_i \pm x_j <> \langle \rho_{ibj}(\sigma), \mathcal{F}_v \rangle \qquad I_{ibj}(\sigma)$$

where the choice of sign \pm and of inequality $<>$ depend on whether v_i, v_j are located on the same side as b at their respective switches. We break into two cases, depending on whether b is contained in the forest F .

If b is not contained in the forest F , then the image of b under the forest collapse q is a branch of σ' also denoted b , and there is an inequality $I_{ibj}(\sigma')$ in $X_{\sigma'}$ of the form

$$x_i \pm x_j <> \langle \rho_{ibj}(\sigma'), \mathcal{F}_v \rangle \qquad I_{ibj}(\sigma')$$

Since b is not collapsed, the cusp v_i is on the same side as b at its switch if and only if the cusp v'_i is on the same side as b , and similarly for v_j, v'_j , and it follows that the \pm sign and the $<>$ sign are chosen identically in $I_{ibj}(\sigma)$ and $I_{ibj}(\sigma')$. Also, the map Q evidently takes $\rho_{ibj}(\sigma)$ to a path that is homotopic rel endpoints to $\rho_{ibj}(\sigma')$, and so the right hand sides of the above two equations are identical. It follows that the solution set $\overset{\circ}{c}(\sigma')$ of $X_{\sigma'}$ is contained in the solution set of the inequality $I_{ibj}(\sigma)$.

Suppose now that b is contained in the forest F , and so the image of b under q is a switch s' of σ' . The cusps v'_i, v'_j of σ' are therefore both located at s' , and so there is an equation $E_{ij}(\sigma')$ in $X_{\sigma'}$, of the form

$$x_i \pm x_j = \langle \rho_{ij}(\sigma'), \mathcal{F}_v \rangle \quad E_{ij}(\sigma')$$

If v_i, v_j are both on the opposite side from b at their switches, or if v_i, v_j are both contained on the same side from b at their switches, it follows that v'_i, v'_j are on opposite sides of the switch s' , and in this case the \pm sign is chosen to be a $+$ in both the inequality $I_{ij}(\sigma)$ and the equation $E_{ij}(\sigma')$. Also, the image under Q of the path $\rho_{ij}(\sigma)$ is obviously homotopic rel endpoints to the path $\rho_{ij}(\sigma')$, and both paths have the same initial orientation at c_i , so the constants on the right hand sides of $I_{ij}(\sigma)$ and $E_{ij}(\sigma')$ are identical. In other words, $E_{ij}(\sigma')$ is the equation obtained from $I_{ij}(\sigma)$ by changing the inequality sign into an equality. This shows that the solution set of $\overset{\circ}{c}(\sigma')$ is contained in the *closure* of the solution set of the inequality $I_{ij}(\sigma)$.

This completes the proof that $\overset{\circ}{c}(\sigma') \subset \partial c(\sigma)$, which is the last step in the proof of Theorem 7.2.2.

7.6 Remarks on noncanonical expansion complexes.

Suppose that \mathcal{F} is an arational measured foliation and τ is a generic train track which carries \mathcal{F} noncanonically, with tie bundle ν . Replacing \mathcal{F} within its equivalence class we may assume that we have a carrying bijection $\mathcal{F} \xrightarrow{\sim} \nu$, but now \mathcal{F} is not a canonical model, instead it has at least one proper saddle connection. We defined and parameterized the space Ξ of separatrix families of \mathcal{F} in Section 4.3. Associated to each element $\xi \in \Xi$ there is still a decorated tie bundle $\nu(\xi)$, and we can use isotopy classes of decorated tie bundles to define an equivalence relation on Ξ as before. The cell structure theorem remains true: equivalence classes are the interiors of cells of a cell decomposition on Ξ .

But the stable equivalence theorem fails in this context: if Ξ, Ξ' are two generalized expansion complexes of the same arational element of \mathcal{MF} then Ξ, Ξ' need not be stably equivalent. Namely, suppose that Ξ is obtained from $\mathcal{F} \xrightarrow{\sim} \nu$ and Ξ' is obtained from $\mathcal{F}' \xrightarrow{\sim} \nu'$, and assume that $\mathcal{F}, \mathcal{F}'$ are equivalent. A little thought shows that stable equivalence of Ξ, Ξ' implies that $\mathcal{F}, \mathcal{F}'$ are isotopic, but two noncanonical measured foliations which are equivalent need not be isotopic. What one can prove is a converse: if $\mathcal{F}, \mathcal{F}'$ are isotopic then Ξ, Ξ' are stably equivalent. One can also prove that if \mathcal{F}' is obtained from \mathcal{F} by slicing along some of the proper saddle connections of \mathcal{F} then Ξ' is stably equivalent to a subcomplex of Ξ that, although it is not a neighborhood of infinity in Ξ , nevertheless intersects every neighborhood of infinity.

7.7 Unmarked stable equivalence of expansion complexes

We have previously defined (marked) stable equivalence of two expansion complexes Ξ, Ξ' to mean the existence of a cellular isomorphism $\psi: \hat{\Xi} \rightarrow \hat{\Xi}'$ between subcomplexes that are neighborhoods of infinity, such that ψ preserves the isotopy class of the decorated train track labelling each cell. We shall define *unmarked* stable equivalence in a similar manner except that ψ is only required to preserve the combinatorial types of train tracks. This turns out not to be quite the correct statement—we must also require that ψ is natural with respect to the face relation among train tracks, and in order to do this we must require not just the existence of combinatorial equivalences between labels of corresponding cells $c(\sigma), c'(\sigma') = \psi(c(\sigma))$, but we must also require the existence of a system of combinatorial equivalences that fits together locally in a nice way.

Here is the precise statement.

Consider the expansion complex Ξ of an arational measured foliation \mathcal{F} based at a generic train track τ that canonically carries \mathcal{F} . Consider two cells Ξ , one a face of the other: $c(\sigma_1) \subset c(\sigma_2)$. Applying the Cell Structure Theorem 7.2.2, there is a map $q_{\sigma_2\sigma_1}: (S, \sigma_2) \rightarrow (S, \sigma_1)$ which collapses some collapsible forest of σ_2 and is otherwise injective, and which is homotopic to the identity on S .

Given two filling train tracks σ, σ' on S , define an *unmarked forest collapse* to be a map $g: \sigma \rightarrow \sigma'$ which factors as the composition of a forest collapse $q: \sigma \rightarrow \sigma''$ postcomposed with an orientation preserving diffeomorphism of S taking σ'' to σ' ; in other words, we simply drop the requirement that g is homotopic to inclusion.

Consider now expansion complexes Ξ, Ξ' of arational measured foliations $\mathcal{F}, \mathcal{F}'$ based at generic train tracks τ, τ' , respectively. An *unmarked stable equivalence* between Ξ and Ξ' consists of the following data:

- A cellular isomorphism $F: \hat{\Xi} \rightarrow \hat{\Xi}'$ between subcomplexes $\hat{\Xi} \subset \Xi, \hat{\Xi}' \subset \Xi'$, each a neighborhood of infinity.
- For each corresponding pair of cells $c(\sigma) \subset \hat{\Xi}$, $c'(\sigma') = F(c(\sigma)) \subset \hat{\Xi}'$, a decoration preserving combinatorial equivalence $f_{\sigma\sigma'}: (S, \sigma) \rightarrow (S', \sigma')$

subject to the following condition: for any pair of cells in $\hat{\Xi}$ one of which is a face of the other, $c(\sigma_1) \subset c(\sigma_2)$, letting the corresponding cells of $\hat{\Xi}'$ be $c'(\sigma'_1) = F(c(\sigma_1))$, $c'(\sigma'_2) = F(c(\sigma_2))$, the following diagram commutes up to homotopy through un-

marked forest collapse maps:

$$\begin{array}{ccc}
 \sigma_2 & \xrightarrow{f_{\sigma_2\sigma'_2}} & \sigma'_2 \\
 q_{\sigma_2\sigma_1} \downarrow & & \downarrow q_{\sigma'_2\sigma'_1} \\
 \sigma_1 & \xrightarrow{f_{\sigma_1\sigma'_1}} & \sigma'_1
 \end{array} \quad (*)$$

We describe the latter requirement by saying that the system of stable equivalences f respects face relations. We say that Ξ, Ξ' are *unmarked stably equivalent* if an unmarked stable equivalence between them exists.

Remarks. It may seem that the choice of a system of combinatorial equivalences $f_{\sigma\sigma'}$ is somewhat heavy handed. However, up to isotopy this is a uniformly finite choice per cell, in that the set of isotopy classes of combinatorial equivalences from σ to σ' is a coset of a finite subgroup of \mathcal{MCG} , as shown in Corollary 3.15.3. Moreover, it should be reasonably clear that most (combinatorial types of) filling train tracks have trivial stabilizer, and so in most cases $f_{\sigma\sigma'}$ is unique up to isotopy. To formulate a rigorous statement along these lines, it should be true that in the finite set of combinatorial types of filling train tracks on S , the percentage with trivial stabilizer approaches 100% as the topological type of S goes to infinity.

The requirement that the system of combinatorial equivalences $f_{\sigma\sigma'}$ respects face relations is designed to guarantee that in fact the maps $f_{\sigma\sigma'}$ are all in the same mapping class—this observation is the key step in the proof of Theorem 7.7.1. Without this requirement, one runs up against the following phenomenon. One can construct a filling train track σ and two distinct forest collapses $f: \sigma \rightarrow \sigma'$, $g: \sigma \rightarrow \sigma''$, such that σ', σ'' are not isotopic but they *are* combinatorially equivalent—in fact, one can take σ to be a generic, recurrent, filling train track on the punctured torus, and σ', σ'' the train tracks obtained by collapsing the two source branches of σ . The requirement of respecting face relations is thus necessary in order to “tie together” the combinatorial equivalences between cells. In general, even if by some coincidence each decorated train track labelling a cell of $\hat{\Xi}$ had trivial stabilizer in \mathcal{MCG} , one would still need additional information tying together the system of combinatorial equivalences. For example, in the unmarked stable equivalence theorem for one cusp expansions, all stabilizers are trivial by Corollary 3.15.3, but we still need the requirement that parity be respected—in the one cusp setting, this requirement is essentially the same as respecting face relations. Although we do not have an example at hand, it seems likely that there are examples of one cusp expansions where successive train tracks are combinatorially equivalent but

the expansions are *not* unmarked stably equivalent; there may also be examples of this phenomenon among expansion complexes.

Theorem 7.7.1 (Unmarked Stable Equivalence Theorem). *If $\mathcal{F}, \mathcal{F}'$ are arational measured foliations canonically carried on generic train tracks τ, τ' with expansion complexes Ξ, Ξ' , respectively, then the following are equivalent:*

- (1) *The measured foliations $\mathcal{F}, \mathcal{F}'$ are topologically equivalent, that is, there exists $\Phi \in \mathcal{MCG}$ such that $\Phi(\mathcal{PMF}(\mathcal{F})) = \mathcal{PMF}(\mathcal{F}')$.*
- (2) *The expansion complexes Ξ, Ξ' are unmarked stably equivalent.*

Proof. One direction is clear: if (1) is true, then the (marked) Stable Equivalence Theorem 7.2.3 implies that $\Phi(\Xi)$ and Ξ' are stably equivalent, and so Ξ, Ξ' are unmarked stably equivalent, using Φ as a face respecting combinatorial equivalence between cell labels, proving (2).

For the other direction, assume (2) is true. Each of $\hat{\Xi}, \hat{\Xi}'$ has a unique non-compact component, and that component is still a neighborhood of infinity, and so replacing each with the noncompact component we may assume that $\hat{\Xi}, \hat{\Xi}'$ are connected.

The central claim is that as $c(\sigma)$ varies over all cells of $\hat{\Xi}$, letting $c'(\sigma') = F(c(\sigma))$, the homeomorphisms $f_{\sigma\sigma'}$ all lie in the same mapping class. Since the poset of faces is connected, to prove this claim it suffices to show that for any face relation $c(\sigma_1) \subset c(\sigma_2)$ in $\hat{\Xi}$, setting $c(\sigma'_1) = F(c(\sigma_1))$ and $c(\sigma'_2) = F(c(\sigma_2))$, the two homeomorphisms $f_{\sigma_1\sigma'_1}$ and $f_{\sigma_2\sigma'_2}$ are in the same mapping class. In the commutative diagram (*), each of the maps may be extended to an ambient map $S \rightarrow S$ in a unique way up to proper homotopy, and the diagram of the ambient extensions commutes up to proper homotopy. We use proper homotopy here because of the presence of punctures. The ambient extensions of the two maps $q_{\sigma_2\sigma_1}, q_{\sigma'_2\sigma'_1}$ are each properly homotopic to the identity, and it follows that $f_{\sigma_1\sigma'_1}$ and $f_{\sigma_2\sigma'_2}$ are properly homotopic. But it is well known that proper homotopy implies isotopy for homeomorphisms between finite type surfaces.

Let Φ be the common mapping class of the maps $f_{\sigma\sigma'}$. It follows that $\Phi(\Xi)$ is the expansion complex of $\Phi(\mathcal{F})$ based at $\Phi(\tau)$. Moreover, clearly $\Phi(\Xi)$ is *marked* stably equivalent to Ξ' . Applying Corollary 7.2.6, we have $\Phi(\mathcal{PMF}(\mathcal{F})) = \mathcal{PMF}(\Phi(\mathcal{F})) = \mathcal{PMF}(\mathcal{F}')$. \diamond

8 The structure of infinite splitting sequences

The Arational Expansion Theorem 6.3.2, and in particular the proof in Section 6.5 of the necessity of the iterated rational killing criterion, brings up a question: what, exactly, is carried by an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, when it does not carry an arational measured foliation? That is, given a complete splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, what can you say about $\cap_i \mathcal{MF}(\tau_i)$ in general?

One motivation to tackle this question comes from a gap in our dictionary between continued fraction expansions and train track expansions: for a given continued fraction, finite or infinite, we know what it is an expansion of; but for a given complete splitting sequence, so far we only know what it is an expansion of in the case where the sequence satisfies the iterated rational killing criterion.

To answer this question we must understand the structure of infinite splitting sequences, in a manner analogous to Thurston's classification of mapping classes. In Theorems 8.2.1 and 8.5.1 below, we will show that an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ exhibits a blending of two extreme kinds of behavior. These two behaviors can be distinguished by considering $\mathcal{F} \in \cap_i \mathcal{MF}(\tau_i)$ and asking whether \mathcal{F}_i is infinitely split by the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$.

Theorem 8.2.1 describes a special case which is analogous to a mapping class that has no pseudo-Anosov pieces. Given a train track τ we shall describe what it means for τ to be "twistable" along a given closed curve c , and if so then we shall describe a *twist splitting* $\tau \succ \tau'$ along c . Examples of twist splittings occur in the description of Penner's recipe [Pen92]. An infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ is said to be a "twist splitting sequence" if there is an essential curve system C such that each $\tau_i \succ \tau_{i+1}$ is a twist splitting along some component of C , and each component of C occurs infinitely often in this manner. If $\tau_0 \succ \tau_1 \succ \dots$ is a twist splitting sequence then Theorem 8.2.1 will show that there exists a train track μ , comb equivalent to a subtrack of each τ_i , such that $\cap_i \mathcal{MF}(\tau_i) = \mathcal{MF}(\mu)$, and moreover no measured foliation is infinitely split by the sequence. Conversely, if nothing is infinitely split by the sequence $\tau_0 \succ \tau_1 \succ \dots$ then it is eventually a twist splitting sequence. The proof of Theorem 8.2.1 is quite different from anything else in this work, involving more train track combinatorics and less measured foliations.

Theorem 8.5.1 and Corollary 8.5 give the structure of a general splitting sequence $\tau_0 \succ \tau_1 \succ \dots$. Theorem 8.5.1 says that $\cap_i \mathcal{MF}(\tau_i)$ is the join of sets of the form $\mathcal{MF}(\mathcal{F}_1), \dots, \mathcal{MF}(\mathcal{F}_J), \mathcal{MF}(\mu)$, where $\mathcal{F}_1, \dots, \mathcal{F}_J$ is a collection of partial arational measured foliations with pairwise disjoint supports, each of which is infinitely split by the sequence $\tau_0 \succ \tau_1 \succ \dots$, and where μ is a train track on S with support disjoint from the supports of $\mathcal{F}_1, \dots, \mathcal{F}_J$ such that $\mathcal{MF}(\mu)$ consists of all measured foliations in $\cap_i \mathcal{MF}(\tau_i)$ that are not infinitely split by the sequence. The foliations $\mathcal{F}_1, \dots, \mathcal{F}_J$ have canonical expansions which can be read off from the given

sequence $\tau_0 \succ \tau_1 \succ \dots$. Also, the subsurface $\mathcal{C}(S - \cup_j \text{Supp}(\mathcal{F}_j))$ is the support of a twist splitting sequence which can be read off from the given sequence, and from this twist splitting sequence one reads off μ . The foliations $\mathcal{F}_1, \dots, \mathcal{F}_J$ are analogous to the pseudo-Anosov pieces in the Thurston decomposition; the complementary twist splitting sequence is analogous to the finite order pieces and the Dehn twists along the bounding annuli.

Theorems 8.2.1 and 8.5.1 will be used later in constructing efficient algorithms for reading off the Thurston decomposition of a mapping class in terms of an invariant train track.

Section 8.1 gives examples and definitions of twist splittings. Theorem 8.2.1, characterizing twist splitting sequences, is stated and proved in Section 8.2, one of the longest sections of this work. Theorem 8.5.1 is stated and proved in Section 8.5.

8.1 Twist splittings

Theorem 8.2.1 will characterize those infinite splitting sequences $\tau_0 \succ \tau_1 \succ \dots$ with the property that nothing in $\cap_i \mathcal{MF}(\tau_i)$ is infinitely split.

We start with a class of examples—for an explicit example in this class see Figure 26. Construct a generic train track τ on S by starting with an essential closed curve c and attaching source branches to c so that c becomes a subtrack of τ . Choose a parity $d \in \{L, R\}$, and make sure that as each branch end approaches c it turns in the direction d as it comes in tangentially. Also, make sure that there is at least one branch end on each side of c , and so c contains at least one sink branch; let $a, b \geq 1$ be the number of branch ends on the two sides of c . Let T_c be the positive Dehn twist around c if $d = R$, and a negative Dehn twist around c if $d = L$. It follows that there is a sequence of splittings of parity d , of the form $\tau = \tau_0 \succ \dots \succ \tau_n = T_c(\tau)$, where $n = a \cdot b$.

Now extend the finite splitting sequence $\tau_0 \succ \dots \succ \tau_n$ to an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ which is T_c -periodic, meaning that $T_c(\tau_i) = \tau_{i+n}$ for all $i \geq 0$. We claim that $\cap_i \mathcal{MF}(\tau_i) = \mathcal{MF}(c)$. The inclusion $\mathcal{MF}(c) \subset \cap_i \mathcal{MF}(\tau_i)$ is clear. For the opposite inclusion, consider a measured foliation \mathcal{F} carried by each τ_i , inducing a weight function μ_i on τ_i . By isotoping c off of itself in one direction we obtain an essential closed curve c' , and in the other direction we obtain c'' . Because all branch ends turn in the same direction as the approach c , both of the curves c', c'' intersect τ_0 efficiently, and it follows that c', c'' intersect \mathcal{F} efficiently. The intersection number $\langle \mathcal{F}, c \rangle = \langle \mathcal{F}, c' \rangle = \langle \mathcal{F}, c'' \rangle$ may therefore be computed by summing the μ_0 weights over all branch ends on the c' side of c , and also summing over all branch ends on the c'' side of c . But if any one of these weights, say x , is positive, then as the n splittings of parity d are performed, going from τ_0 to $\tau_n = T_c(\tau_0)$, each branch weight of μ_n along c is less than the corresponding branch

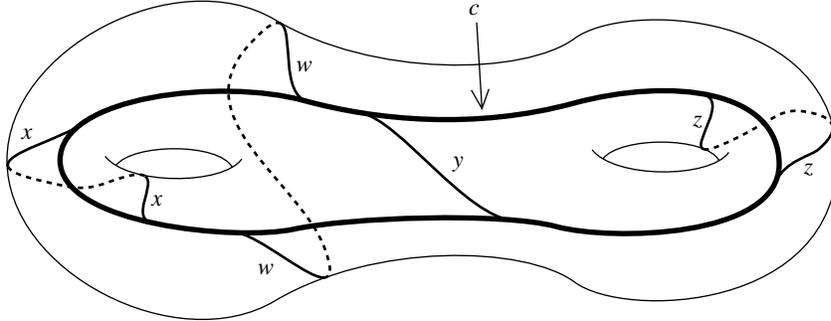


Figure 26: A twist splitting on a train track τ , obtained from an essential closed curve c by attaching branches w, x, y, z . Each branch end turns to the Left as it approaches c , with 4 branch ends approaching from the inside and 4 from the outside. There is a sequence of $4 \cdot 4 = 16$ Left splittings going from τ to the train track $T_c(\tau)$, where T_c is the negative Dehn twist around c . If \mathcal{F} is any measured foliation carried by τ then $\langle \mathcal{F}, c \rangle = x + 2y + z = x + 2w + z$, and so $\langle \mathcal{F}, c \rangle = 0$ if and only if \mathcal{F} is an annular foliation whose leaves are isotopic to c .

weight on μ_0 by an amount at least equal to x . Moreover, the branch weight x itself is unchanged. Thus, as the Dehn twist T_c is iterated, branch weights along c continue to be decremented by at least the amount x for each twist. This contradicts the Archimedean principle when x is positive, and it follows that $\langle \mathcal{F}, c \rangle = 0$ and μ_0 equals zero for each branch not on c . This implies that $[\mathcal{F}] \in \mathcal{MF}(c)$ as required.

We shall describe a generalization of the above example, called a *twist splitting sequence*. Theorem 8.2.1 will say that a splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ has the property that no element of $\bigcap_i \mathcal{MF}(\tau_i)$ is infinitely split if and only if the sequence can be truncated so that what remains after truncation is a twist splitting sequence. Moreover, in this case $\bigcap_i \mathcal{MF}(\tau_i)$ is very easily described: it is equal to $\mathcal{MF}(\rho)$ for some subtrack ρ of a train track in the truncated sequence.

Twist splittings along an essential curve. Given a train track τ and a closed curve c , we say that c is *simply carried* by τ if the inclusion map $c \hookrightarrow S$ is isotopic to a smooth embedding with image in τ . A weaker property is that c is *semisimply carried* by τ , which means that there exists a local embedding $f: c \times [0, 1] \rightarrow S$ such that $\tau_c = f(c \times 1)$ is the subtrack of τ fully carrying c , the restricted map $f|_{c \times [0, 1)}$ is an embedding with image disjoint from $f(c \times 1)$, and $f(c \times t)$ is isotopic to c for $t \in [0, 1)$. Denote $A_c = f(c \times [0, 1])$, called a *collar* of c ; note that

$A_c - \tau_c$ is a half-open annulus in S , foliated by circles isotopic to c . By a regular homotopy of the map f we may assume that $f^{-1}(\tau)$ is a pretrack containing $c \times 1$, and that each branch of $f^{-1}(\tau) - c$ is transverse to the circle direction of $c \times [0, 1]$, with one endpoint at a switch on $c \times 1$ and the other endpoint transverse on $c \times 0$.

For example, if $\sigma \subset \tau$ is a dumbbell subtrack then there is an essential closed curve c that is semisimply carried by τ with carrying image σ , but there is of course no closed curve that is simply carried by τ with image σ . In general, if c is semisimply but not simply carried by τ then the carrying map $c \rightarrow \tau_c$ is one-to-one over source branches of τ_c , two-to-one over sink branches of τ_c , and τ_c has no transition branches.

Let c be semisimply carried by τ , and $f: c \times [0, 1] \rightarrow A_c$ a parameterization of a collar as above, and so $f^{-1}(\tau)$ is a pretrack in $c \times [0, 1]$ containing $c \times 1$ and transverse to the circle foliation. We say that c is a *twist curve* for τ with *twist collar* A_c if the following conditions hold:

- $f^{-1}(\tau) - (c \times 1) \neq \emptyset$, equivalently, $f^{-1}(\tau)$ contains at least one switch on $c \times 1$.
- There exists an orientation on c , called the *twist orientation*, which agrees with the switch orientation of each switch of $f^{-1}(\tau)$ contained in $c \times 1$.
- The subtrack τ_c contains at least one sink branch of τ .

Using the twist orientation on c we get a product orientation on $c \times [0, 1]$ and we can ask whether the local diffeomorphism $f: c \times [0, 1] \rightarrow S$ preserves orientation. If it does then we say that c is a *Right twist curve*, and otherwise c is a *Left twist curve*. The parity $d \in \{L, R\}$ of the twist curve c is equivalently characterized by saying that as a branch approaches τ_c through A_c it turns in the direction d .

Note that if a twist curve $c = \tau_c$ is simply carried on τ then, given a two-sided collar neighborhood of c , one or both halves of this collar may be twist collars; if both halves are twist collars then they must impart opposite twist orientations on c , in order for τ_c to contain a sink branch, as in the earlier example. Unlike the examples in Figure 26, it is *not* necessary in general for c to be simply carried, and even if it is it is *not* necessary for τ_c to have twist collars on both sides.

We shall define certain homotopic carryings $\tau \succcurlyeq \tau'$ called *twist carryings* along a twist curve of τ . Let c be a d twist curve of τ , $d \in \{L, R\}$. First, if $\tau \cong \tau'$ is any comb equivalence then c is still a d twist curve for τ' , and $\tau \cong \tau'$ is an example of twisting along c . Consider now a generic train track τ and a sink branch $b \subset \tau$ such that $b \subset \tau_c$ and such that the carrying map $c \rightarrow \tau$ contains a d crossing of b . In this situation the elementary splitting $\tau \succ \tau'$ of parity d along b is another example of a twist carrying along c , called an *elementary twist splitting* along c ; in this situation, it again follows that c is a d twist curve of τ' . A general twist splitting $\tau \succ \tau'$

along c is defined, as usual, by the existence of generic train tracks τ_1, τ'_1 comb equivalent to τ, τ' respectively such that τ_1, τ'_1 is an elementary twist splitting along c . General twist carryings are defined by iteration: $\tau \succ \tau'$ is a twist carrying along c if it is a comb equivalence or it factors as a sequence of twist splittings along c , $\tau = \tau_0 \succ \tau_1 \succ \cdots \succ \tau_n = \tau'$.

Here is another description of twist carryings. Consider a train track c with twist curve c and twist collar $f: c \times [0, 1] \rightarrow A_c$. Let $q: \mathbf{R} \times [0, 1] \rightarrow c \times [0, 1]$ be a universal covering map which takes the standard orientation on \mathbf{R} to the twist orientation on c . Let $F = f \circ q: \mathbf{R} \times [0, 1] \rightarrow A_c$. A *twist map* $h: c \times [0, 1] \rightarrow c \times [0, 1]$ relative to $c \times 0$ is defined to be a diffeomorphism whose lift to the universal cover has the form

$$H(x, t) = (\exp(2\pi i \tilde{H}(\theta, t)), t)$$

where the smooth function $\tilde{H}: \mathbf{R} \times [0, 1] \rightarrow \mathbf{R}$ satisfies: $\tilde{H}(\theta, t) = \theta$ for t in some neighborhood of 0; $\tilde{H}(\theta + 1, t) = \tilde{H}(\theta, t) + 1$; $\frac{\partial}{\partial x} \tilde{H} > 0$; and $\frac{\partial}{\partial t} \tilde{H} \geq 0$. Let τ' be obtained from τ by removing $\tau \cap A = f(f^{-1}(\tau)) = F(F^{-1}(\tau))$ from τ and inserting $F(H(F^{-1}(\tau)))$. Let τ_1, τ'_1 be any train tracks comb equivalent to τ, τ' respectively. Then $\tau_1 \succ \tau'_1$ is a twist carrying along c , and moreover every twist carrying along c has this form up to isotopy.

Twist splittings along an essential curve family. Let C be an essential, pairwise nonisotopic curve family in S . If τ is a train track such that each $c \in C$ is a twist curve of τ , then we say that C is a *twist family* for τ , and we denote $\tau_C = \cup_{c \in C} \tau_c$. One can choose simultaneous twist collars $f_c: c \times [0, 1] \rightarrow A_c \subset S$, for all $c \in C$, to have the following property:

Disjointness property $(A_c - \tau_c) \cap (A_{c'} - \tau_{c'}) = \emptyset$ if $c \neq c' \in C$.

To see why this is possible, assume that one twist collar A_c is chosen, and consider $c' \neq c \in C$. The carrying map $f_{c'}: c' \rightarrow \tau$ is defined, and we want to extend this to a twist collar $f_{c'}: c' \times [0, 1] \rightarrow S$ with the above disjointness property. This is possible if we can show that $f_{c'}(c')$ cannot approach τ_c from within $A_c - \tau_c$. If it does then there are two cases. If $f_{c'}: c' \rightarrow \tau$ pulls back continuously via f_c to a map $c' \rightarrow c \times [0, 1]$ then it follows that $c = c'$, a contradiction. If $f_{c'}: c' \rightarrow \tau$ does not pull back continuously to $c \times [0, 1]$ then, as c' approaches c from within $A_c - \tau_c$, it follows along c for a while but must eventually cross to the other side of c , since it does not have a continuous pullback; this implies that c and c' have nonzero intersection number, also a contradiction.

Given a twist system C for τ , the fact that simultaneous twist collars exist with the disjointness property implies that if $\tau \succ \tau'$ is a comb equivalence or a twist splitting along some component of C then C is also a twist family for τ' . Note

moreover that τ'_C is isotopic to τ_C , so we may assume that $\tau'_C = \tau_C$ as subsets of S . We may also assume that we have a family of simultaneous twist collars for τ which is also a family of simultaneous twist collars for τ' . Henceforth we will make use of these assumptions without further mention.

We define general twist carryings along C by iteration: assuming C is a twist family for τ , we define $\tau \succcurlyeq \tau'$ to be a twist carrying along C if it is a comb equivalence or it factors as a sequence $\tau = \tau_0 \succcurlyeq \tau_1 \succcurlyeq \cdots \succcurlyeq \tau_n = \tau'$ of twist splittings along components of C .

Given an essential curve family C we say that an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ is *infinitely twisted along C* if each splitting $\tau_i \succ \tau_{i+1}$ is a twist splitting along some component of C , and each component of C occurs infinitely often as the curve along which $\tau_i \succ \tau_{i+1}$ is split.

Normal form for twist splitting sequences. An infinite splitting sequence which is infinitely twisted along an essential curve family can be put into a nice normal form, up to comb equivalence.

A train track τ with a twist system C is said to be in *normal form rel C* if each switch of τ lying on τ_C is a generic switch of τ , and each branch of τ intersecting $A_c - \tau_c$ is a source branch, for each $c \in C$. To put τ in normal form within its comb equivalence class: first uncomb to make τ generic; then inductively collapse transition branches that intersect $A_c - \tau_c$ for some c ; and then, given a nongeneric, semigeneric switch s of τ such that $s \in \tau_C$, carefully uncomb s to produce generic switches connected by a train path of transition branches that are each contained in τ_C .

Consider now any twist splitting $\tau \succ \tau'$ along $c \in C$, and by comb equivalence we may assume that τ, τ' are generic and that $\tau \succ \tau'$ is a splitting along a sink branch b of τ contained in τ_c . Extend b slightly into the cusps at each end of b to give a splitting arc α along which $\tau \succ \tau'$ is split, with $\alpha \cap \tau = b$. Converting τ into a normal form τ_1 by applying the process above, let α_1 be a splitting arc of τ_1 corresponding to α under the comb equivalence $\tau \approx \tau_1$, and from the description of the normal form process it follows easily that $\alpha_1 \cap \tau_1 \subset \tau_C$. In terms of the parameterization $f_c: c \times [0, 1] \rightarrow A_c$, we can pull back $\alpha_1 \cap \tau_1$ continuously to an arc in $c \times 1$ denoted $\hat{\alpha}$. Define a twist map $h: c \times [0, 1] \rightarrow c \times [0, 1]$ which moves $\hat{\alpha}$ off of itself and is constant outside of a small neighborhood of $\hat{\alpha}$. As described earlier, this twist map determines a twist carrying $\tau_1 \succ \tau'_1$, and clearly τ'_1 is comb equivalent to τ' . Moreover, the branches outside of τ_C are unaltered except for where they attach to τ_C , and by construction of h it follows that τ_C contains no nongeneric switches of τ'_1 , and thus τ'_1 is in normal form rel C . We say that the twist splitting $\tau_1 \succ \tau'_1$ is in normal form rel C .

By induction using the construction above it follows that if $\tau_0 \succ \tau_1 \succ \dots$ is infinitely twisted along some essential curve family C , then one can put the entire splitting sequence in normal form rel C , as follows. First change τ_0 in its comb equivalence class to be in normal form rel C . Assuming inductively that the initial subsequence $\tau_0 \succ \dots \succ \tau_i$ is in normal form rel C , by the previous paragraph we may alter τ_{i+1} within its comb equivalence class so that it, and the splitting $\tau_i \succ \tau_{i+1}$, are both in normal form rel C .

A digression: Relations between twist carryings and Dehn twists. Consider a twist carrying $\tau \succ \tau'$ along a twist curve c of τ , defined in terms of a rotation map $h: c \times [0, 1] \rightarrow c \times [0, 1]$. There is an interesting analytic relation between the rotation number $r(h)$ and the number of twist splittings it takes to factor $\tau \succ \tau'$. Recall the definition of $r(h)$: using the lift $H(x, t) = (\exp(2\pi i \tilde{H}(\theta, t)), t)$ of h as defined earlier, letting $g(\theta) = \tilde{h}(\theta, 1)$, we have

$$r(h) = \lim_{n \rightarrow \infty} \frac{g^n(\theta)}{n}$$

which is independent of $\theta \in \mathbf{R}$.

Suppose now that $\tau \succ \tau'$ factors as twist splittings $\tau = \tau_0 \succ \dots \succ \tau_N$ in normal form rel c . Then there exist constants $K \geq 1$, $A \geq 0$ depending only on τ such that

$$r(h) - A \leq N \leq Kr(h) + A$$

The lower bound follows because each Dehn twist requires at least one twist splitting. The constant K for the upper bound is $K = a \cdot b$, where a is the number of branch ends in $A_c - \tau_c$, and where b is the number of points on $c \times 1$ which map to switches of τ pointing in the opposing direction. The coarseness constant A is needed because the rotation number of a circle map is only an asymptotic measure of how the map rotates the circle.

By using the rotation number inequality above one can prove the following: letting T_c be the Dehn twist along c , positive if $d = \mathbf{R}$ and negative if $d = \mathbf{L}$, then there exists a constant A such that for any twist carrying $\tau \succ \tau'$ along c there exists an integer $n \geq 0$ such that one can get from $T_c^n(\tau)$ to τ' by a sequence of at most A twist splittings.

8.2 Characterization of twist splitting sequences

A *twist splitting sequence* is an infinite splitting sequence which is infinitely twisted along some essential curve family.

Theorem 8.2.1. *For any infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, the following are equivalent:*

- (1) No element of $\cap_i \mathcal{MF}(\tau_i)$ is infinitely split.
- (2) There exists I such that the truncated sequence $\tau_I \succ \tau_{I+1} \succ \dots$ is a twist splitting sequence.

In this case there is a train track μ , comb equivalent to a subtrack of τ_i for each $i \geq I$, such that $\cap_i \mathcal{MF}(\tau_i) = \mathcal{MF}(\mu)$.

For the rest of this section we will prove (2) \implies (1). The converse direction will be proved in Sections 8.3—8.4.

Proof of (2) \implies (1). The following lemma shows how to define the train track μ :

Lemma 8.2.2. *Consider a train track τ with a twist system C , and with simultaneous twist collars A_c for $c \in C$. Let μ be the largest recurrent subtrack of τ that is disjoint from $\cup_{c \in C} (A_c - \tau_c)$. Then μ is the unique recurrent subtrack of τ with the property that for all $[\mathcal{F}] \in \mathcal{MF}(\tau)$, we have $[\mathcal{F}] \in \mathcal{MF}(\mu)$ if and only if $\langle \mathcal{F}, c \rangle = 0$ for all $c \in C$.*

Note that $\mu \neq \emptyset$ because, by the disjointness property for simultaneous twist collars, each subtrack τ_c is contained in μ . In the example of Figure 26, $\mu = \tau_c = c$.

Proof. We may choose each $c \in C$ in its isotopy class to coincide with the image of $c \times \frac{1}{2}$ under the map $c \times [0, 1] \rightarrow A_c$. With this choice, the curve c and the train track τ intersect efficiently, because all branch ends intersecting c turn in the same direction as they approach τ_c . In particular, for any \mathcal{F} carried by τ inducing a weight function $w \in W(\tau)$, we can compute the intersection number $\langle \mathcal{F}, c \rangle$ by summing $w(x)$ over $x \in c \cap \tau$. It follows that $\langle \mathcal{F}, c \rangle = 0$ if and only if $w(b) = 0$ for each branch b intersecting $A_c - \tau_c$, and therefore $\langle \mathcal{F}, c \rangle = 0$ for all $c \in C$ if and only if \mathcal{F} is carried by μ . \diamond

Proof of (2) \implies (1). Assume that $\tau_0 \succ \tau_1 \succ \dots$ is infinitely split along the essential curve family C , and is in normal form rel C . As mentioned earlier we may assume that $\tau_C = \cup_{c \in C} \tau_c$ is a subtrack of each τ_i , and that there is a simultaneous family of twist collars A_c , $c \in C$, for each τ_i . Let μ_i be the largest subtrack of τ_i disjoint from each $A_c - \tau_c$. Note that μ_i survives the splitting $\tau_i \succ \tau_{i+1}$ with descendant μ_{i+1} , and that $\mu_i \approx \mu_{i+1}$ is a comb equivalence; indeed, since the splittings are in normal form rel C it follows that all the μ_i are isotopic to each other; let μ be a train track in this isotopy class, identified with μ_0 . We therefore have $\mathcal{MF}(\mu) \subset \cap_i \mathcal{MF}(\tau_i)$ and moreover nothing in $\mathcal{MF}(\mu)$ is infinitely split. To prove (2) \implies (1) it therefore suffices to show that $\cap_i \mathcal{MF}(\tau_i) \subset \mathcal{MF}(\mu)$. By applying the above lemma it suffices to show that if $\mathcal{F} \in \cap_i \mathcal{MF}(\tau_i)$ then $\langle \mathcal{F}, c \rangle = 0$ for all $c \in C$.

Let $w_i: \tau_i \rightarrow [0, \infty)$ be the invariant weight induced by \mathcal{F} . Using the intersection number formula for $\langle \mathcal{F}, c \rangle$ from the proof of the above lemma, if b is any branch of τ_0 intersecting $A_c - \tau_c$ it suffices to prove that $w_0(b) = 0$.

Suppose then that $c \in C$ is fixed and that there is a branch b of τ_0 intersecting $A_c - \tau_c$ such that $r = w_0(b) > 0$. We shall derive a contradiction, similar to the contradiction derived in the earlier example, by showing that the values of w_i on points of τ_c are infinitely decremented by amounts of at least r , which contradicts the Archimedean principle.

Consider the parameterization $f_c: c \times [0, 1] \rightarrow S$ of the twist collar A_c . A point $x \in c \times 1$ is called an *opposing switch* of τ_i if $f_c(x)$ is a switch of τ_i whose switch orientation pulls back via the derivative of f_c to an orientation of the tangent line of x in $c \times 1$ which points in the direction opposite to the twist orientation on $c \times 1$. Enumerate the branch ends of $f_c^{-1}(\tau_0)$ in $A_c - \tau_c$ in cyclic order, starting with b , as $b = b_1, \dots, b_M$; equivalently, we may think of these as the branches of $f_c^{-1}(\tau_0) - (c \times 1)$, and by normality they end at points in $c \times 1$ which are distinct from each other and from the opposing switches. In the twist splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, we know that infinitely many of the twist splittings occur along c ; enumerate these as $\tau_{i_j} \succ \tau_{i_{j+1}}$, $1 \leq j < \infty$, and let $h_j: c \times [0, 1] \rightarrow c \times [0, 1]$ be a twist map which determines the twist splitting $\tau_{i_j} \succ \tau_{i_{j+1}}$. For $m = 1, \dots, M$ let $b_m^0 = b_m$ and by induction on $j \geq 0$ note that $b_m^j = h_{j-1}(b_m^{j-1})$, are precisely the branches of $f_c^{-1}(\tau_{i_j}) - (c \times 1)$. Thus, h_j has the effect of taking some cyclic subinterval of the branches b_1^j, \dots, b_M^j and moving each of them past one or more opposing switches of τ_{i_j} , leaving the rest of the branches b_m^j unmoved.

Notice that $w_{i_j}(b_m^j)$ is constant independent of j .

Notice also that for each $x \in \tau_C$ the weights $w_i(x)$ form a nonincreasing sequence: $w_{i+1}(x) - w_i(x)$ is equal to the sum of the weights of the branch ends which are moved past x under the twist splitting $\tau_i \succ \tau_{i+1}$.

However, we claim that for any point $x \in \tau_c$ there are infinitely many values of j for which the branch end b_1^j is moved past x , and since $w_{i_j}(b_1^j) = w_0(b) = r$ is positive, this claim contradicts the Archimedean principle.

To prove the claim, suppose by contradiction that for sufficiently large j , say $j \geq J$, the branch b_1^j is not moved past x . It follows that no branch ends are moved past $f_c^{-1}(x)$ by h_j for any $j \geq J$, and hence h_j only moves points forward within the set $(c \times 1) - f_c^{-1}(x)$ which consists of one or two open subintervals of the circle $c \times 1$. This implies that each time a branch end is moved past an opposing switch in the set $(c \times 1) - f_c^{-1}(x)$ it cannot move past that switch again. Since there are only finitely many branch ends and finitely many opposing switches, eventually no branch end moves past an opposing switch in $c \times 1$. This contradicts that the sequence $\tau_0 \succ \tau_1 \succ \dots$ is infinitely split along the twist curve c , completing the

proof of (2) \implies (1) in Theorem 8.2.1.

8.3 Cutting one train track out of another

In this subsection we develop some concepts that will be needed in the proof that (1) \implies (2) in Theorem 8.2.1, as well as in the proof of Theorem 8.5.1.

Consider an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$. Recall that a line of descent consists of a sequence of subtracks $\sigma_i \subset \tau_i$ defined for all sufficiently large i , say $i \geq I$, such that each σ_i survives the splitting $\tau_i \succ \tau_{i+1}$ with descendant σ_{i+1} . Each line of descent is a carrying sequence $\sigma_I \succ \sigma_{I+1} \succ \cdots$ such that each $\sigma_i \succ \sigma_{i+1}$ is either a comb equivalence or a parity splitting. Up to truncation there are only finitely many lines of descent, and so we can and often will truncate the splitting sequence so that every line of descent starts in τ_0 .

For any $\mathcal{F} \in \cap_i \mathcal{MF}(\tau_i)$, letting $\sigma_i \subset \tau_i$ be the subtrack that fully supports \mathcal{F} , we obtain a full carrying sequence $\sigma_0 \succ \sigma_1 \succ \cdots$. This need not be a line of descent because not every carrying $\sigma_i \succ \sigma_{i+1}$ need be homotopic, however only finitely many of them are not homotopic and hence for some $I \geq 0$ we obtain a line of descent $\sigma_I \succ \sigma_{I+1} \succ \cdots$, and by truncation we usually assume that $I = 0$.

The strategy for proving (1) \implies (2) in Theorem 8.2.1 will be to examine the structure of the lines of descent associated to elements of $\cap_i \mathcal{MF}(\tau_i)$. More specifically, assuming that $\tau_0 \succ \tau_1 \succ \cdots$ satisfies (1), we shall look for a certain line of descent which can be cut out of $\tau_0 \succ \tau_1 \succ \cdots$, to produce a simpler infinite splitting sequence that still satisfies (1). Applying induction, this simpler sequence will satisfy (2), and then we put the pieces back together to conclude (2) for the sequence $\tau_0 \succ \tau_1 \succ \cdots$.

The “cutting out” operation referred to above will be defined whenever one train track τ carries another σ . The result of cutting σ out of τ will not be a train track but a more general object called a “track with terminals”.

Tracks with terminals. A *track with terminals* in S is a smooth, compact 1-complex $\pi \subset S$ satisfying all the axioms of a pretrack except that for each $x \in \pi$, we do not require the existence of a smooth open arc in π containing x , but we do require the existence of a smooth *compact* arc in π that contains x , possibly as an endpoint. A point which does not lie on any smooth open arc in π is called a *terminal*. Given a terminal $s \in \pi$, on one side of the tangent line at s there exists at least one incident branch end; the number of such branch ends is called the *valence* of s . A valence 1 terminal is called a *stop* [Pen92]. Each branch of π incident to two stops is an entire component of π diffeomorphic to a closed interval in the real line.

If π is a track with terminals, there is a subtrack with terminals $\pi' \subset \pi$ obtained from π by *pulling out the stops*, which means that we remove from π each stop and the interior of each branch incident to a stop, repeating such removals inductively until there are no stops, and finally removing any isolated points that remain. Put another way, π' is the largest track with terminals in π that has no stops.

Many of the definitions concerning train tracks extend in a straightforward way to tracks with terminals, such as: tie bundles; comb moves and comb equivalence; carrying; elementary splittings, splitting arcs, and wide splittings; splitting sequences, etc. We shall refer as needed to extensions of basic results from the category of train tracks to the category of tracks-with-terminals; the proofs of these extended results are left to the reader. One slight change is that for purposes of defining sink branches, source branches, and transition branches of a track-with-terminals π , for any branch $b \subset \pi$ with an endpoint at a terminal s we should regard s as an outflow of b . Thus, if b has an inflow at its opposite end then b is a transition branch and we can comb π along b .

Consider a track with terminals π . The union of any two pretracks that are contained in π is a pretrack contained in π , and hence there is a maximal pretrack in π . In the contexts that we will be studying, these pretracks will always be train tracks, which can be easily checked as needed, and hence we will abuse terminology and refer to the maximal pretrack of π as the *maximal train track* of π .

A track with terminals π is *transient* if either of the following equivalent conditions hold: the maximal train track in π is empty; there are no closed train paths in π ; there is an upper bound to the combinatorial length of each train path in π .

Lemma 8.3.1. *Given a transient track with terminals π , there is an upper bound to the length of any splitting sequence starting with π .*

Proof. The proof requires studying maximal train paths in π . Each maximal train path has both ends at terminals. There is only a finite number of maximal train paths, because if there were infinitely many then they would be indefinitely long, and so there would be one which crosses some branch twice in the same direction, giving a closed train path.

Let $M(\pi)$ be the number of maximal train paths. The number $M(\pi)$ is invariant under comb equivalence of transient tracks-with-terminals.

Given a splitting $\pi \succ \pi'$ where π is transient, note that π' is also transient. Moreover, the fold map $\pi' \rightarrow \pi$ takes the set of maximal train paths of π' injectively to the set of maximal train paths of π . We claim that the image of this map misses at least one maximal train path in π . For instance, if $\pi \succ \pi'$ is a parity splitting then we may assume by comb equivalence that $\pi \succ \pi'$ is an elementary splitting of parity $d \in \{L, R\}$ along a sink branch $b \subset \pi$, in which case any \bar{d} crossing of b

extends to a maximal train path in π which is not the image under folding of any maximal train path in π' . The case of a central splitting is similar. This proves that $M(\pi') < M(\pi)$, and it follows by induction that any splitting sequence starting with π has at most $M(\pi)$ splittings. \diamond

We will need to strengthen the above lemma. Given a track with terminals π , a smoothly embedded circle $\gamma \subset \pi$ is called a *peripheral circle* if γ is a boundary circle of a compact annulus $A \subset S$ such that $A \cap \pi = \gamma$. We say that π is *transient rel periphery* if its maximal recurrent train track is a disjoint union of peripheral circles.

Lemma 8.3.2. *If the track with terminals π is transient rel periphery, then there is a bound to the length of any splitting sequence starting with π .*

Proof. Unlike in the transient case, we cannot prove this lemma by counting all maximal train paths, because there are infinitely many of them: a train path can cycle around a peripheral curve any number of times.

Instead of counting arbitrary train paths, we consider train paths in π which are *embeddable* meaning that there exists an embedded lift to the interior of the I -bundle over π . A train path which enters a peripheral curve γ , goes more than one time around γ , and then leaves γ is not embeddable. However, a train path α which enters γ and then spirals infinitely around γ is embeddable; in fact, if an embeddable α goes more than once around γ then it is constrained to stay in γ in at least one direction.

Any embeddable train path can be extended to a maximal embeddable train path. To see how, suppose α is an embeddable train path with at least one endpoint that is not at a terminal, and we regard α as lifted to the I -bundle over π . Starting from any endpoint of α that is not at a terminal, start extending α in the I -bundle; there is no obstruction to continuing the extension indefinitely while staying embedded in the I -bundle, unless the path arrives at the tie over a terminal. Once one endpoint has been extended either infinitely or to a terminal, the other endpoint, if it exists, can then be extended.

For any branch $b \subset \pi$ which is not in a peripheral circle, a train path cannot traverse b more than once in each direction, for otherwise one easily constructs a recurrent subtrack which is not a peripheral circle. Similarly, for any peripheral circle γ , if a train path $\alpha: I \rightarrow \pi$ traverses two branches of γ in the same direction then all branches traversed by α between these two are also in γ , and hence $\alpha^{-1}(\gamma)$ contains at most two components.

The previous paragraph, together with the constraint on embeddable train paths going more than once around a peripheral curve, implies that there are only finitely many maximal embeddable train paths α up to reparameterization of the domain of

the train path, because α is determined up to finite ambiguity by which peripheral curves it intersects and which nonperipheral branches it traverses.

Let $N(\pi)$ be the number of maximal embeddable train paths. Note that $N(\pi)$ is invariant under comb equivalence. For any splitting $\pi \succ \pi'$, the track with terminals π' is also transient rel its periphery Γ' , and the fold map $\pi' \rightarrow \pi$ maps Γ' diffeomorphically onto Γ . Now the proof proceeds exactly as in Lemma 8.3.1, showing that $N(\pi') < N(\pi)$ by mapping the set of maximal embeddable train paths in π' injectively but not surjectively to the set of maximal embeddable train paths in π . By induction it follows that any splitting sequence starting with π has at most $N(\pi)$ splittings. \diamond

Note: When π is transient the bound $N(\pi)$ is better than the bound $M(\pi)$ in Lemma 8.3.1. However, $N(\pi)$ is still probably a pretty bad bound: a single splitting most likely wipes out lots of maximal embeddable train paths.

Cutting one train track out of another. Here is our main construction of tracks with terminals.

Let τ, σ be train tracks on S such that $\tau \succ \sigma$. We shall define a track with terminals $\rho \subset S - \sigma$ which is obtained from τ by “cutting out” σ , denoted $\rho = \tau \setminus \sigma$.

Choose tie bundles $\nu(\sigma), \nu(\tau)$ and a carrying injection $\nu(\sigma) \subset \text{int}(\nu(\tau))$. Define $\nu(\rho'') = \nu(\tau) - \text{int}(\nu(\sigma))$, and decompose $\nu(\rho'')$ into ties, each a maximal subtie of $\nu(\tau)$ which is disjoint from $\text{int}(\nu(\sigma))$. Collapse each tie of $\nu(\rho'')$ to a point to obtain a track with terminals ρ'' having one terminal for each cusp of σ . Note that each terminal of ρ'' is a stop at the endpoint of a branch $b \subset \rho''$ whose opposite endpoint is a switch with one branch end on the side of b and two or more branch ends on the opposite side. Let ρ' be obtained from ρ'' by pulling all the stops, so each component of ρ'' deformation retracts to a component of ρ' , and all components of ρ' occur in this way. Note that $\nu(\rho'') \cap \nu(\sigma) = \partial\nu(\sigma)$. Each component γ of $\partial\nu(\sigma)$ is thus contained in $\nu(\rho'')$, and the projection $\gamma \mapsto \rho''$ is locally injective except where it folds around a stop. When the stops are pulled then the map $\gamma \rightarrow \rho''$ homotopes to a local injection $\gamma \rightarrow \rho'$, called a *peripheral curve* of ρ' . Given a component C of ρ' , we say that C is *trivial* if it is the embedded image of some peripheral curve; equivalently, there exists a component c of $\partial\nu(\tau)$ which projects diffeomorphically onto C under the projection $\nu(\rho'') \mapsto \rho'$. Now define $\rho = \tau \setminus \sigma$ by removing from ρ' all trivial components. Each nontrivial peripheral curve of ρ' has image contained in ρ and is called a *peripheral curve* of ρ . The set of peripheral curves of ρ is denoted Γ_ρ .

For example, if $\tau = \sigma$, or more generally if we have a homotopic carrying $\tau \stackrel{\text{H}}{\succ} \sigma$, then every component of ρ' is a trivial peripheral curve and so $\rho = \tau \setminus \sigma = \emptyset$. This

indicates that what we are mainly interested in is cutting out train tracks σ which are carried by τ but not homotopically carried, usually proper subtracks or smooth closed curves.

Note that as a consequence of Proposition 3.5.2, if τ carries σ then the carrying injection $\nu(\sigma) \hookrightarrow \nu(\tau)$ is well defined up to homotopy through carrying injections. It follows that the tie bundle $\nu(\rho'')$, and the set of curves $\nu(\rho'') \cap \nu(\sigma)$, are well defined independent of the choice of the carrying injection $\nu(\sigma) \hookrightarrow \nu(\tau)$, showing that ρ and Γ_ρ are also well defined.

In the special case that σ is a subtrack of τ , each peripheral curve of ρ' is embedded in ρ' , and hence the same is true of ρ . This is true because for each tie t of $\nu(\rho'')$, either $t \cap \partial\nu(\sigma)$ is at most one point, or t projects to a point of $\rho'' - \rho'$, or t projects to a terminal of ρ' .

When σ is a subtrack of τ then there is an alternate description of $\tau \setminus \sigma$ which may be enlightening. Consider the surface-with-cusps $\mathcal{C}(S - \sigma)$, with overlay map $q: \mathcal{C}(S - \sigma) \rightarrow S$. Note that $\rho' = q^{-1}(\tau)$ is a track with terminals in $\mathcal{C}(S - \sigma)$ containing $\partial\mathcal{C}(S - \sigma)$ and with one terminal at each cusp of $\mathcal{C}(S - \sigma)$; no terminal of ρ' is a cusp. The peripheral curves of ρ' are precisely the components of $\partial\mathcal{C}(S - \sigma)$, making clear that each peripheral curve of ρ' is embedded in ρ' . The track with terminals ρ is defined by first removing from ρ' any peripheral curve which is an entire component of ρ' , and then isotoping the rest into $\text{int}(\mathcal{C}(S - \sigma)) = S - \sigma$.

Tracks with moats. When $\tau \gg \sigma$ and $\rho = \tau \setminus \sigma$, the union $\omega = \sigma \cup \rho$ is a track with terminals in S which has a particular structure that we now describe. A *track with moats* in S is a track with terminals $\omega \subset S$ such that each component of $\mathcal{C}(S - \omega)$ is either a surface with cusps of negative Euler index, or an annulus μ called a *moat*, such that the following hold:

- For each moat μ there is an integer $n \geq 0$ such that one component of $\partial\mu$ has n cusps and the other component has n reflex cusps, and hence the Euler index of μ equals zero. If $n = 0$ we say μ is a *smooth moat*.
- For each moat μ there is a distinguished component of $\partial\mu$ called the *peripheral curve*, such that when μ is not smooth the peripheral curve is the component of $\partial\mu$ with reflex cusps (the peripheral curves need not be embedded in general).
- There is a foliation of μ by arcs connecting the two components of $\partial\mu$.
- There is a smooth map $\phi: S \rightarrow S$ homotopic to the identity which collapses (the overlay image of) each leaf of each moat to a point, and is otherwise one-to-one and locally diffeomorphic. This map ϕ is called *draining the moats*.

Let Γ be a set consisting of one peripheral curve for each moat. We index moats by their peripheral curves, letting μ_γ be the moat with peripheral curve γ , and denoting $t_\gamma = \partial\mu_\gamma - \gamma$. We will occasionally confuse μ_γ with its image in S ; there is no loss in doing this because the image of μ_γ determines μ_γ unambiguously.

When $\tau \succ \sigma$ and $\rho = \tau \setminus \sigma$ then $\omega = \sigma \cup \rho$ is a track with moats whose peripheral curves are precisely the peripheral curves of ρ as in the definition of “cutting out”. This can be seen by noticing that in the tie bundle $\nu(\tau)$, we can isotope $\nu(\rho'') = \nu(\tau) - \text{int}(\nu(\sigma))$ until $\nu(\rho'')$ is disjoint from $\nu(\sigma)$, after which each component of $\partial\nu(\sigma)$ includes as one boundary circle of an annulus component of $\text{Cl}(S - (\nu(\sigma) \cup \nu(\rho'')))$. Some of these annuli are adjacent to trivial components of $\nu(\rho'')$, but the rest of these annuli give rise to the moats of ω .

We need a moderate generalization of “cutting out”, as follows. Consider a track with terminals π in S , and suppose that π has no stops. Given a subtrack σ carried by π , we choose in the obvious way a carrying injection $\nu(\sigma) \hookrightarrow \text{int}(\nu(\pi))$, and from this point on the definition of the track with terminals $\rho = \pi \setminus \sigma$ and of the peripheral curves of ρ are word for word as in the case when π is a train track. If π has the additional structure of a track with moats then it follows easily that $\sigma \cup \rho$ is a track with moats, such that each moat of π is still a moat of $\sigma \cup \rho$, with additional moats coming from nontrivial components of $\partial\nu(\sigma)$.

Lemma 8.3.3. *Given a track with terminals π , if σ is the maximal train track of π , and if $\rho = \pi \setminus \sigma$, then the maximal train track of ρ is precisely the smooth periphery of ρ , that is, the set of smooth peripheral curves.*

Proof. Any bi-infinite train path γ in ρ must map, via draining the moat $\sigma \cup \rho \rightarrow \pi$, into σ , because σ is the maximal train track in π . This is only possible if γ is carried by the smooth periphery of ρ , that being the only part of ρ whose image under the map $\sigma \cup \rho \rightarrow \pi$ is contained in σ . It follows that the maximal train track of ρ is precisely the smooth periphery. \diamond

Given a splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ and an infinitely surviving subtrack $\sigma_0 \subset \tau_0$ with line of descent $\sigma_0 \succ \sigma_1 \succ \dots$, for each i we will be interested in the track with terminals $\pi_i = \tau_i \setminus \sigma_i$. The following lemma shows that, in the context of tracks with terminals, we obtain a carrying sequence $\pi_0 \succ \pi_1 \succ \dots$ such that each $\pi_i \succ \pi_{i+1}$ is either a comb equivalence or a splitting.

Lemma 8.3.4. *Given a parity splitting $\tau \succ \tau'$ and a subtrack $\sigma \subset \tau$ that survives with descendant σ' , if $\pi = \tau \setminus \sigma$ and $\pi' = \tau' \setminus \sigma'$ then $\pi \succ \pi'$ and this carrying is either a comb equivalence or a splitting.*

Proof. Consider first the case that τ, τ' are generic and $\tau \succ \tau'$ is an elementary splitting of parity $d \in \{\text{L}, \text{R}\}$ along a sink branch $b \subset \tau$. For concreteness we assume

$d = L$. Let a neighborhood $N(b)$ be denoted as in Figure 13, with $\partial N(b) \cap \tau = \{r, s, t, u\}$, and given a subset $P \subset \{r, s, t, u\}$ let \overline{P} denote the smallest connected subset of $N(b) \cap \tau$ containing P . The set $\sigma \cap N(b)$ is called the *interaction* of σ with the splitting. Since σ survives a Left splitting, we have the following restrictions on the interaction: either $\overline{su} \subset \sigma$ or $\overline{rt} \not\subset \sigma$. This implies that $\sigma \cap N(b)$ fits into one of the cases named in the following table:

Name of Interaction	$\sigma \cap N(b)$	$\pi \rightsquigarrow \pi'$
Empty	\emptyset	Splitting
Tangential	\overline{ru} or \overline{st}	Splitting
Diagonal	\overline{su}	Isotopy
Y	\overline{rsu} or \overline{stu}	Isotopy
X	\overline{rstu}	Isotopy

The table shows the nature of the carrying $\pi \rightsquigarrow \pi'$ in each case, either a splitting or an isotopy, and we refer to Figure 27 for verification.

Consider next the case of a comb equivalence $\tau \approx \tau'$ in which τ' obtained from τ by combing a transition branch $b \subset \tau$. In this case it is easy to see that $\pi \approx \pi'$ is a comb equivalence, in fact either a comb move along a transition branch or an isotopy; an analysis as in Figure 27 may easily be carried out in this case.

For a general splitting $\tau \succ \tau'$ the proof follows by induction, factoring it into comb moves and an elementary splitting. \diamond

8.4 Proof of (1) \implies (2) in Theorem 8.2.1

Consider an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ that does not infinitely split anything in $\cap_i \mathcal{MF}(\tau_i)$.

As the proof goes along, certain statements will only be true for sufficiently large values of the index i , and as we accumulate more such statements the lower bound $i \geq I$ for the truth of all these statements will increase. Instead of keeping track of the increasing value of the lower bound I , we will simply truncate the original splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ so that the statements become true for all i .

Note that $\tau_0 \succ \tau_1 \succ \dots$ has only finitely many lines of descent up to truncation: this follows from the fact that each τ_i has only finitely many subtracks, and that two lines of descent are equivalent up to truncation if and only if they coincide in any τ_i where they are both defined. By truncating $\tau_0 \succ \tau_1 \succ \dots$ we may therefore assume that each line of descent starts in τ_0 , and hence has the form $\sigma_0 \rightsquigarrow \sigma_1 \rightsquigarrow \dots$ for $\sigma_i \subset \tau_i$. Furthermore, each line of descent either has infinitely many splittings or finitely many, and by further truncation we can assume that each line of descent which has finitely many splittings actually consists entirely of comb equivalences, thereby having the form $\sigma_0 \approx \sigma_1 \approx \dots$.

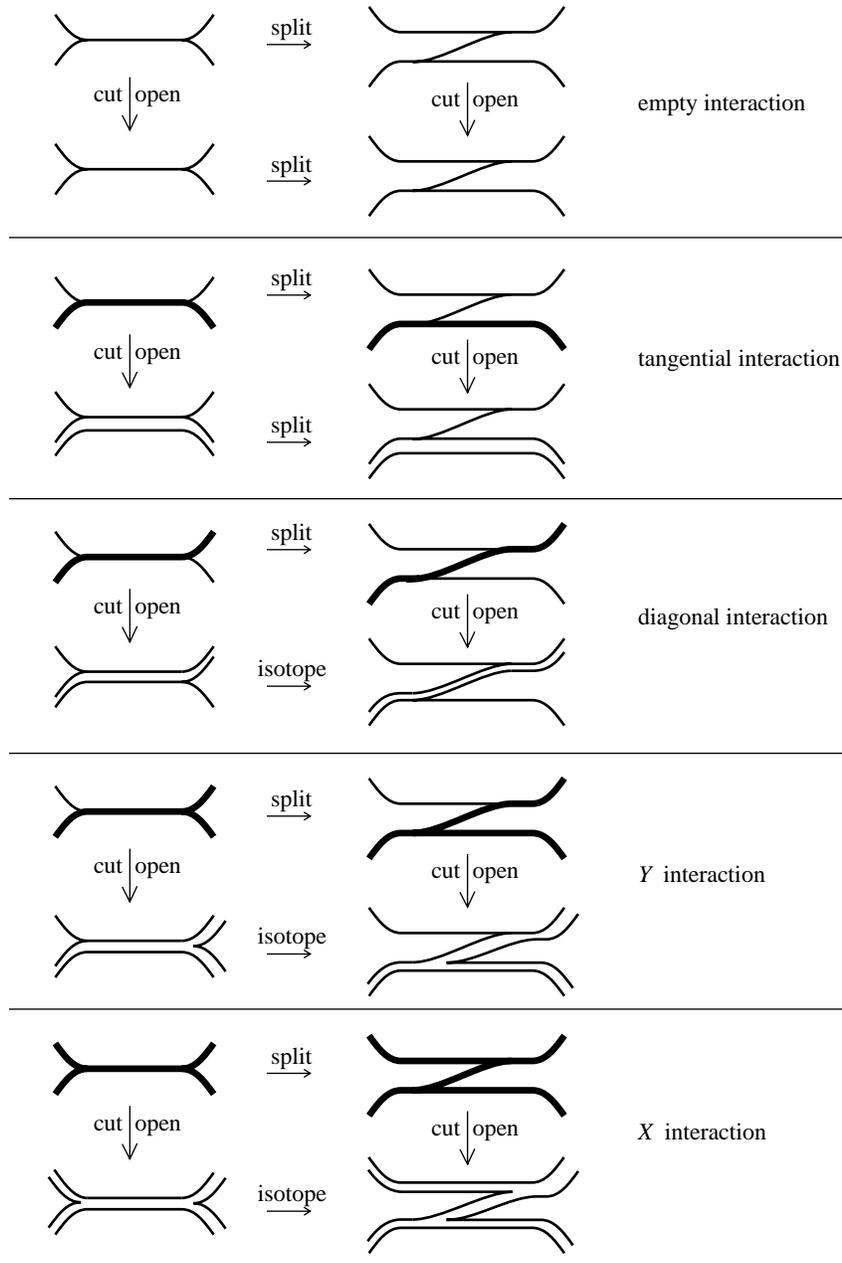


Figure 27: Interactions of a surviving subtrack $\sigma \subset \tau$ with an elementary Left splitting $\tau \succ \tau'$ with descendant $\sigma' \subset \tau'$. The subtracks σ, σ' are shown in bold. The carrying $\pi \succ \pi'$ is shown just below the splitting in each case.

Given any splitting $\tau \succ \tau'$ and a subtrack $\sigma \subset \tau$ that survives with descendant $\sigma' \subset \tau'$, it is evident that σ is a smooth peripheral circle of τ if and only if σ' is a smooth peripheral circle of τ' . Furthermore, each smooth peripheral circle of τ does indeed survive in τ' . It follows that for any line of descent $\sigma_0 \succ \sigma_1 \succ \cdots$, if there exists i such that σ_i is a smooth peripheral circle of τ_i then this is true for all i . Such lines of descent will be called *peripheral*, and we want to avoid them, because for any train track τ and subtrack σ which is a disjoint union of smooth peripheral circles, it is evident that cutting τ along σ yields τ back again. In fact we shall prove the existence of a line of descent which is nonperipheral in a strong sense.

Given a splitting $\tau \succ \tau'$ and a subtrack $\sigma \subset \tau$ that survives with descendant $\sigma' \subset \tau'$, and given a smooth peripheral circle $\gamma \subset \tau$ with descendant $\gamma' \subset \tau'$, it is evident that $\gamma \subset \sigma$ if and only if $\gamma' \subset \sigma'$. It follows that for any line of descent $\sigma_0 \succ \sigma_1 \succ \cdots$, if some σ_i contains a peripheral circle then they all do. A subtrack (or its line of descent) which contains no peripheral circle is called *aperipheral*.

Step 1: There is an aperipheral, noninfinitely split line of descent.
We begin with:

Claim 8.4.1. *If the train track τ is not transient rel periphery, then there exists an aperipheral recurrent subtrack of τ .*

To prove the claim, let $\sigma \subset \tau$ be a recurrent subtrack that is not contained in the union of peripheral circles. By an obvious induction it suffices to show that if σ contains some peripheral circle γ of τ then there exists a recurrent subtrack $\sigma' \subset \sigma$ such that $\gamma \not\subset \sigma'$ but $\sigma' - \gamma = \sigma - \gamma$. Let \mathcal{F} be a partial measured foliation fully carried on σ , with a surjective inclusion of \mathcal{F} into the tie bundle $\nu(\sigma)$. Since γ lifts to a smooth boundary component of $\nu(\sigma)$ it follows that \mathcal{F} has an annular component \mathcal{F}' carried by γ . Let $w \in W(\sigma)$ be the invariant weight induced by \mathcal{F} , and let $r > 0$ be the smallest value of $w(b)$ for branches $b \in \gamma$; choose $b \subset \gamma$ so that $w(b) = r$. The transverse weight of \mathcal{F}' is clearly equal to r . It follows that if \mathcal{F}'' is the union of all foliation components of \mathcal{F} except \mathcal{F}' , then the maximal subtrack $\sigma' \subset \sigma$ carrying \mathcal{F}'' is a recurrent train track not containing b . Since \mathcal{F}' is carried by γ it follows that $\sigma' - \gamma = \sigma - \gamma$, proving the claim.

The proof of the claim clearly shows that if \mathcal{F} is carried by a train track τ , and if $\sigma \subset \tau$ is the subtrack fully carrying \mathcal{F} , then σ contains a peripheral curve γ of τ if and only if \mathcal{F} has an annular component with core isotopic to γ .

Fix $I \geq 0$ and consider the train track τ_I . Since τ_I is the start of an infinite splitting sequence, by applying Lemma 8.3.2 it follows that τ_I is not transient rel periphery. This implies that τ_I has an aperipheral recurrent subtrack; choose a measured foliation \mathcal{F} fully carried by this subtrack. Clearly we have $[\mathcal{F}] \in \cap_{i=1}^I \mathcal{MF}(\tau_i)$. Moreover, the subtrack of τ_0 fully carrying \mathcal{F} is aperipheral, because τ_0 and τ_I have

the same peripheral circles up to isotopy, and none of them is the core of an annular component of \mathcal{F} .

To summarize, for each $I \geq 0$ there exists a recurrent, aperipheral subtrack $\sigma \subset \tau_0$ such that $\mathcal{MF}(\sigma) \cap (\bigcap_{i=1}^I \mathcal{MF}(\tau_i)) \neq \emptyset$. Since τ_0 has only finitely many subtracks, it follows that there is an aperipheral recurrent subtrack $\sigma \subset \tau_0$ such that $\mathcal{MF}(\sigma) \cap (\bigcap_{i=1}^{\infty} \mathcal{MF}(\tau_i)) \neq \emptyset$. Choosing \mathcal{F} in this set, it follows that the line of descent fully carrying \mathcal{F} is aperipheral and is not infinitely split, finishing Step 1.

Step 2: Cutting open the splitting sequence. Let $\sigma_0 \approx \sigma_1 \approx \dots$ be a non-peripheral line of descent, consisting entirely of comb equivalences, for the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$. We may assume that each σ_i is connected. Cut σ_i out of τ_i to obtain $\pi'_i = \tau_i \setminus \sigma_i$. Applying Lemma 8.3.4 we obtain a carrying sequence $\pi'_0 \succ \pi'_1 \succ \dots$ where each $\pi'_i \succ \pi'_{i+1}$ is a comb equivalence or a splitting. Let $\rho_i \subset \pi'_i$ be the maximal subtrack. Clearly we have a full carrying $\rho_i \succ \rho_{i+1}$ which is either a comb equivalence or a splitting, possibly a central splitting, but the sequence $\rho_0 \succ \rho_1 \succ \dots$ can have at most finitely many central splittings and so by truncation we may assume that they are all comb equivalences or parity splittings. Cut ρ_i out of π'_i to obtain $\pi_i = \pi'_i \setminus \rho_i$. Applying Lemma 8.3.4 and truncating we may again assume that the sequence $\pi_0 \succ \pi_1 \succ \dots$ consists of comb equivalences and parity splittings. By Lemma 8.3.3 each π_i is transient rel a pairwise disjoint collection of smooth peripheral curves, and hence by Lemma 8.3.2 the sequence $\pi_0 \succ \pi_1 \succ \dots$ contains only finitely many splittings. After another truncation we have a sequence of comb equivalences $\pi_0 \approx \pi_1 \approx \dots$. Since the train track ρ_i is carried by τ_i it follows that $\bigcap_i \mathcal{MF}(\rho_i) \subset \bigcap_i \mathcal{MF}(\tau_i)$. Given $\mathcal{F} \in \bigcap_i \mathcal{MF}(\rho_i)$, since \mathcal{F} is not infinitely split by $\tau_0 \succ \tau_1 \succ \dots$ it follows that \mathcal{F} is not infinitely split by $\rho_0 \succ \rho_1 \succ \dots$.

If $\rho_0 \succ \rho_1 \succ \dots$ has infinitely many splittings then, by ignoring the comb equivalences in the sequence, we may regard it as an infinite splitting sequence, and hence the sequence $\rho_0 \succ \rho_1 \succ \dots$ satisfies (1). In this case we wish to apply induction to conclude that $\rho_0 \succ \rho_1 \succ \dots$ satisfies (2), but in order to do so we need to know that ρ_i is less complex than τ_i in some well ordered sense. This is where it is crucial that σ_i is a nonperipheral subtrack of τ_i . For example, if σ_i were a disjoint union of peripheral circles of τ_i then $\rho_i = \tau_i \setminus \sigma_i$ would be isotopic to τ_i , giving us no foundation for an inductive proof.

Recall that if F is a finite type surface-with-boundary, of genus g and with a total of p punctures and boundary components, then the complexity of F is defined to be the dimension of the Teichmüller space of the interior of F , namely $d(F) = 6g - 6 + 2p$. Note that $d(\text{Supp}(\rho_i))$ and $d(\text{Supp}(\tau_i))$ are each constant independent of i . We want to apply Lemma 3.2.4 to obtain the inequality

$$d(\text{Supp}(\rho_i)) < d(\text{Supp}(\tau_i))$$

To start with $\text{Supp}(\sigma_i)$ is an essential subsurface of $\text{Supp}(\tau_i)$. Now it is not true that $\pi'_i \subset S - \text{Supp}(\sigma_i)$, because a nonpunctured or once punctured disc component A of $S - \sigma_i$ need not be disjoint from π'_i . However, if the set $A \cap \pi'_i$ is nonempty then it is a transient component of π'_i and so is disjoint from ρ_i . This implies that $\text{Supp}(\rho_i) \subset \text{Supp}(\tau_i) - \text{Supp}(\sigma_i) \subset \text{Supp}(\tau_i)$, and so by Lemma 3.2.4 we at least have the nonstrict inequality $d(\text{Supp}(\rho_i)) \leq d(\text{Supp}(\tau_i))$. Now we use that σ_i is not peripheral, which implies that no component of $\text{Supp}(\tau_i) - \text{Supp}(\sigma_i)$ is isotopic to $\text{Supp}(\tau_i)$, and hence $\text{Supp}(\rho_i)$ is not isotopic to $\text{Supp}(\tau_i)$, and applying Lemma 3.2.4 it follows that the inequality is strict.

By induction on complexity of subsurfaces it follows that the sequence $\rho_0 \succ \rho_1 \succ \dots$ either has finitely many splittings or it has infinitely many splittings and, regarding it as an infinite splitting sequence by ignoring comb equivalences in the sequence, the sequence satisfies (2). Applying induction we may conclude that, after further truncation, either we have an essential curve family $C' \neq \emptyset$ along which the sequence $\rho_0 \succ \rho_1 \succ \dots$ is infinitely twisted, or we have a sequence of comb equivalences $\rho_0 \approx \rho_1 \approx \dots$ and in this case we set $C' = \emptyset$.

We may regard C' as a train track carried on each ρ_i , so we may cut C' out of ρ_i to obtain $\omega_i = \rho_i \setminus C'$.

Let us now review what has happened. Cutting σ_i out of τ_i we obtain the track-with-terminals $\pi'_i = \tau_i \setminus \sigma_i$ and the track with moats $\sigma_i \cup \pi'_i$, with set of moats denoted M_1 that is independent of i . Drainage of the moats in M_1 allows reconstruction of τ_i from $\sigma_i \cup \pi'_i$. Next, cutting the maximal subtrack ρ_i out of π'_i we obtain the track with terminals $\pi_i = \pi'_i \setminus \rho_i$, and a track with moats $\sigma_i \cup \rho_i \cup \pi_i$. Each moat in M_1 is still a moat of $\sigma_i \cup \rho_i \cup \pi_i$, and together with a new collection of moats M_2 arising from cutting ρ_i out of π'_i , the set of all moats of $\sigma_i \cup \rho_i \cup \pi_i$ is just $M_1 \cup M_2$, again independent of i . Drainage of the moats $M_1 \cup M_2$ allows τ_i to be reconstructed from $\sigma_i \cup \rho_i \cup \pi_i$. Finally, cutting the essential curve system C' out of ρ_i we obtain $\omega_i = \rho_i \setminus C'$, and a track with moats denoted

$$\hat{\tau}_i = \sigma_i \cup \pi_i \cup \omega_i$$

Again each moat of $M_1 \cup M_2$ is a moat of $\hat{\tau}_i$, and together with a new collection of moats M_3 , one for each curve in C' , the set of all moats of $\hat{\tau}_i$ is $M = M_1 \cup M_2 \cup M_3$, again independent of i . Drainage of the moats in M allows τ_i to be reconstructed from $\hat{\tau}_i$. The key feature of this construction is that we have a sequence of comb equivalences $\hat{\tau}_0 \approx \hat{\tau}_1 \approx \dots$.

Let $q_i: \hat{\tau}_i \rightarrow \tau_i$ be the map which drains the moats M .

Step 3: Moat analysis. We analyze how the moats M of $\hat{\tau}_i$ interact with the splittings in the sequence $\tau_0 \succ \tau_1 \succ \dots$, stating the key point in Lemma 8.4.2 and

proving (1) \implies (2) using this lemma. The proof of the lemma is given in steps 4 and 5.

Given a moat $\mu \in M$ let γ_μ be its core curve.

We defined earlier what it means for the splitting $\tau_i \succ \tau_{i+1}$ to be a twist splitting with respect to some twist curve. We generalize this by saying that a moat μ *interacts* with the splitting $\tau_i \succ \tau_{i+1}$ if there is a smooth train path $\beta \subset \partial\mu$ such that $h_i(\beta)$ is a d -crossing of the sink branch along which $\tau_i \succ \tau_{i+1}$ is split, where d is the parity of the splitting; equivalently, there is a smooth train path $\beta \subset \partial\mu$ such that $h_{i+1}(\beta)$ contains the post-splitting arc of τ_{i+1} . Note that the each moat in M_3 , obtained from C' , interacts with infinitely many of the splittings $\tau_0 \succ \tau_1 \succ \dots$. The goal of moat analysis is to determine exactly which moats in the whole collection M interact with infinitely many of the splittings.

Note that it is possible for as many as two moats to interact with a splitting $\tau_i \succ \tau_{i+1}$.

Classify the moats of M into four types:

- Nonsmooth moats.
- Smooth moats which are not twistable, meaning that each component of the boundary of the moat contains a sink branch of $\hat{\tau}_i$, a property which is independent of i . Note that a smooth moat μ is not twistable if and only if γ_μ is not a twist curve of τ_i .
- Smooth twistable moats μ which are only finitely split, meaning that they interact with only finitely many splittings in the sequence $\tau_0 \succ \tau_1 \succ \dots$. Equivalently, γ_μ is a twist curve but only finitely many of the splittings $\tau_0 \succ \tau_1 \succ \dots$ are twist splittings along γ_μ .
- Smooth twistable moats μ which are infinitely split, meaning that infinitely many of the splittings $\tau_0 \succ \tau_1 \succ \dots$ are twist splittings for γ_μ ; obviously this includes all moats with core curve in C' .

The meanings of “smooth” and “twistable” are defined with respect to any $\hat{\tau}_i$, and they are independent of i .

Here is the key point:

Lemma 8.4.2. *Any nonsmooth moat and any smooth nontwistable moat interacts with only finitely many splittings.*

Using the lemma we finish the proof of (1) \implies (2) as follows.

Define C to be the core curves of the smooth, twistable moats which interact with infinitely many splittings of $\tau_0 \succ \tau_1 \succ \dots$. Note that $C' \subset C$. Evidently each

curve $c \in C$ is a twist curve for each τ_i , and infinitely many splittings are twist splittings along c .

By applying the lemma, we can truncate the splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ so that any nonsmooth or smooth nontwistable moat interacts with none of the splittings, and we can truncate further so that any smooth, twistable, finitely split moat interacts with none of the splittings. Thus, the only moats μ which interact with $\tau_0 \succ \tau_1 \succ \cdots$ are those which are smooth, twistable, and infinitely twisted, implying that each splitting $\tau_i \succ \tau_{i+1}$ is a twist splitting along some curve in C . This proves (2).

To start the proof of Lemma 8.4.2, consider a moat $\mu \in M$ with core curve γ_μ . Each boundary circle of the annulus $\mathcal{C}(\mu)$ maps to a track with terminals in $\hat{\tau}_i$; denote these two subsets $\partial_i^\pm \mu \subset \hat{\tau}_i$. If μ is smooth then $\partial_i^\pm \mu$ are subtracks of $\hat{\tau}_i$.

Consider a C-splitting arc α of τ_i ; and recall that each end of α emerges into some cusp of τ_i . We say that α *crosses the moat* μ if α can be written as a concatenation of the form $\alpha = \alpha_- * \alpha_0 * \alpha_+$ so that the following properties hold:

- (1) For $\epsilon \in \{-, +\}$ the path α_ϵ lifts smoothly to a path $\hat{\alpha}_\epsilon$ whose intersection with $\hat{\tau}_i$ is a union of transition branches ending at the first point where $\hat{\alpha}_\epsilon$ intersects $\partial_i^\epsilon \mu$.
- (2) For each $\epsilon \in \{-, +\}$ the path α_0 lifts smoothly to a path $\hat{\alpha}_0^\epsilon$ in $\partial_i^\epsilon \mu$, in such a way that $\alpha_- * \alpha_0$ lifts smoothly to $\hat{\alpha}_- * \hat{\alpha}_0^-$ and $\alpha_0 * \alpha_+$ lifts smoothly to $\hat{\alpha}_0^+ * \alpha_+$.

Note that $\hat{\alpha}_-, \hat{\alpha}_+$ hit the periphery of $\hat{\tau}_i$ only at their inner endpoints, for if they hit the periphery of $\hat{\tau}_i$ any earlier then they could never leave afterwards since they consist solely of transition paths, but that means that they hit $\partial \mu$ earlier than their inner endpoint, a contradiction.

Step 4: Proof of Lemma 8.4.2 for a smooth, nontwistable moat μ . Since μ is nontwistable, the two sides $\partial^\epsilon \mu$, $\epsilon \in \{-, +\}$, each contain a sink branch of $\hat{\tau}_i$. Let \mathcal{A}_i be the set of relative isotopy classes of C-splitting arcs α which cross μ such that at least one of the lifts $\hat{\alpha}_0^-, \hat{\alpha}_0^+ \subset \hat{\tau}_i$ does *not* contain a sink branch of $\hat{\tau}_i$. Let $N_i = |\mathcal{A}_i|$.

We claim that $N_{i-1} \geq N_i$ with strict inequality if the moat μ interacts with the splitting $\tau_{i-1} \succ \tau_i$; this clearly implies that μ interacts with only finitely many splittings, finishing Step 4. In proving the claim, note that the number N_i is an invariant of the comb equivalence class of τ_i , and so we are free to assume that $\tau_{i-1} \succ \tau_i$ is an elementary splitting of some parity $d \in \{L, R\}$.

We use the usual technique of defining an injective map $\mathcal{A}_i \rightarrow \mathcal{A}_{i-1}$ which is nonsurjective if μ interacts with the splitting $\tau_{i-1} \succ \tau_i$. Given $\alpha \in \mathcal{A}_i$, to show that α corresponds to some element of \mathcal{A}_{i-1} the key is showing that the carrying map $\tau_i \rightarrow \tau_{i-1}$ is injective on α . Consider the concatenation $\alpha = \alpha_- * \alpha_0 * \alpha_+$ and lifts $\hat{\alpha}_-, \hat{\alpha}_0^-, \hat{\alpha}_0^+, \hat{\alpha}_+$.

Arguing by contradiction, suppose that α does not inject in τ_{i-1} , which implies that α contains each of the switches at the ends of the postsplitting branch of τ_i ; see Figure 28. There exists a moat μ' which interacts with the splitting $\tau_{i-1} \succ \tau_i$, and hence the postsplitting branch b of τ_i is contained in the carrying image of the curve $\gamma_{\mu'}$. However, the only peripheral points of $\hat{\tau}_i$ contained in $\hat{\alpha}_-$ or $\hat{\alpha}_+$ are their inner endpoints, and so these endpoints lie on $\partial_i^- \mu', \partial_i^+ \mu'$, respectively, implying that $\mu' = \mu$. Thus, the postsplitting branch b lifts to $b^\epsilon \subset \partial_i^\epsilon \mu$ for each $\epsilon \in \{-, +\}$. Moreover, the two endpoints of the b lift to the inner endpoints of $\hat{\alpha}_-, \hat{\alpha}_+$, respectively. Notice that α_0 must go around all of the carrying map $\gamma_\mu \rightarrow \tau_i$, from one endpoint of b to the other, without traversing b itself. It follows that for each $\epsilon \in \{-, +\}$ the path $\hat{\alpha}_0^\epsilon$ must go all the way around $\partial_i^\epsilon \mu$ except for b^ϵ . But since $\partial_i^\epsilon \mu$ contains a sink branch of $\hat{\tau}_i$ it follows that $\hat{\alpha}_0^\epsilon$ contains a sink branch, contradicting the fact that $\alpha \in \mathcal{A}_i$.

We have shown that α injects in τ_{i-1} . The concatenation $\alpha = \alpha_- * \alpha_0 * \alpha_+$ used to conclude $\alpha \in \mathcal{A}_i$ works, with small adjustments, to show that the image of α under the carrying map $\tau_i \rightarrow \tau_{i-1}$ is an element of the set \mathcal{A}_{i-1} . This defines the map $\mathcal{A}_i \rightarrow \mathcal{A}_{i-1}$. The map is obviously injective.

Assuming μ interacts with the splitting $\tau_{i-1} \succ \tau_i$ we prove that $\mathcal{A}_i \rightarrow \mathcal{A}_{i-1}$ is nonsurjective by exhibiting an element of \mathcal{A}_{i-1} which is not in the image. Let b' be the sink branch along which $\tau_{i-1} \succ \tau_i$ is split, and since μ interacts with the splitting it follows that the carrying map of $\gamma_\mu \rightarrow \tau_{i-1}$ traverses a d -crossing of b' . This implies that if α' is the c -splitting arc intersecting τ_{i-1} in b' then α' crosses μ . Obviously neither of the lifts $\hat{\alpha}'_0^+, \hat{\alpha}'_0^-$ contains a sink branch of $\hat{\tau}_{i-1}$, since the two lifts of b' to $\hat{\tau}_{i-1}$ are *not* sink branches of $\hat{\tau}_{i-1}$. This shows that $\alpha' \in \mathcal{A}_{i-1}$, but clearly α' is not in the image of \mathcal{A}_i .

This finishes Step 4.

Remark: The proof of step 4 will fail if μ is a twistable moat of $\hat{\tau}_i$, as it should in order for infinite twisting to be possible. The reason the proof fails is that the carrying map $\tau_i \rightarrow \tau_{i-1}$ need not be injective on α if one of $\partial^- \mu, \partial^+ \mu$ consists solely of transition branches of $\hat{\tau}_i$, in which case no map from \mathcal{A}_i to \mathcal{A}_{i-1} is defined.

Step 5: Proof of Lemma 8.4.2 for nonsmooth moats. We will sketch this proof briefly; it is similar to but simpler than Step 4.

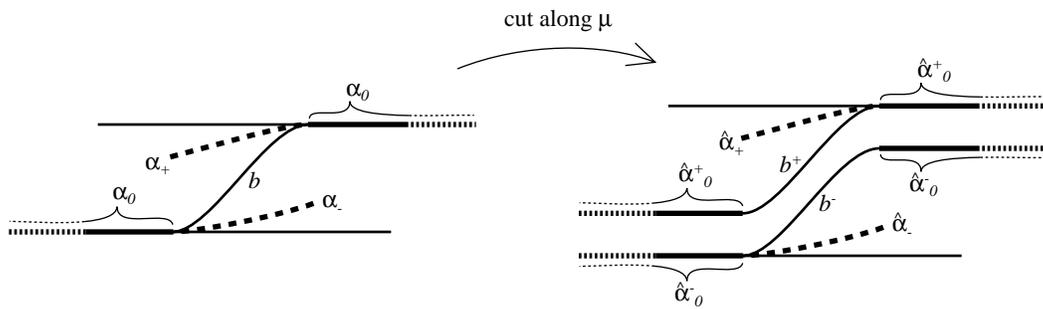


Figure 28: If the c-splitting arc α of τ_i crosses the smooth nontwistable moat μ , and if α contains the endpoints of the postsplitting branch b of $\tau_{i-1} \succ \tau_i$, then α_0 goes all the way around the carrying image of the closed curve γ_μ except for the branch b , implying that each of $\hat{\alpha}_0^+$, $\hat{\alpha}_0^-$ must contain a sink branch of $\hat{\tau}_i$. This figure shows the case that α^-, α^+ each intersect τ_i in a single point, which in the present situation must be an endpoint of b ; the cases where α_- or α_+ intersect τ_i nontrivially are similarly depicted.

Fix a nonsmooth moat μ of $\hat{\tau}_i$. For a C-splitting arc α of τ_i , the definition of α crossing μ is exactly the same. The set \mathcal{A}_i is defined to be the set of all C-splitting arcs of τ_i which cross μ (the issue of sink branches of $\hat{\tau}_i$ is irrelevant here). Let $N_i = |\mathcal{A}_i|$.

Again we prove $N_{i-1} \geq N_i$ with strict inequality whenever $\tau_{i-1} \succ \tau_i$ intersects with the moat μ , and again the technique is to produce an injective map $\mathcal{A}_i \rightarrow \mathcal{A}_{i-1}$ which is not surjective whenever $\tau_{i-1} \succ \tau_i$ intersects with the moat μ . Let $\alpha = \alpha_- * \alpha_0 * \alpha_+$ be in \mathcal{A}_i , with lifts $\hat{\alpha}_-, \hat{\alpha}_0^+, \hat{\alpha}_0^-, \hat{\alpha}_+$. Again the key is showing that the carrying map $\tau_i \rightarrow \tau_{i-1}$ is injective on α . Again, the only way this could fail is if the splitting involves μ , implying that the postsplitting branch b of τ_i is contained in $h_i(s)$ for some side s of μ , one endpoint of b lying in α_- , and the other endpoint in α_+ . But this is impossible because the train path α_0 must traverse the train path $h_i(s)$ from one endpoint of b to the other, staying disjoint from the interior of b (see Figure 28), but b disconnects its two endpoints in the train path $h_i(s)$ — here we are using the fact that s is an arc, rather than a circle as in Step 4. The map $\mathcal{A}_i \rightarrow \mathcal{A}_{i-1}$ is thus defined. Injectivity is obvious as before, and when μ is involved with the splitting $\tau_{i-1} \succ \tau_i$ then failure of surjectivity is proved as before.

This completes step 5, and the proof of (1) \implies (2) in Theorem 8.2.1. \diamond

8.5 General splitting sequences.

The following theorem gives the structure of a general splitting sequence.

Theorem 8.5.1. *Given an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, there exists a train track μ and a set of partial arational measured foliations $\{\mathcal{F}_1, \dots, \mathcal{F}_J\}$, $J \geq 0$, such that the supports of $\mu, \mathcal{F}_1, \dots, \mathcal{F}_J$ are pairwise disjoint, and such that*

$$\cap_i \mathcal{MF}(\tau_i) = \mathcal{MF}(\mathcal{F}_1) * \dots * \mathcal{MF}(\mathcal{F}_J) * \mathcal{MF}(\mu)$$

Moreover, an element of $\cap_i \mathcal{MF}(\tau_i)$ is infinitely split if and only if it has a component in $\mathcal{MF}(\mathcal{F}_j)$ for some $j = 1, \dots, J$.

Corollary 8.5.2. *Given an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ with $\mathcal{F}_1, \dots, \mathcal{F}_J, \mu$ as above, the following description of $\tau_0 \succ \tau_1 \succ \dots$ holds. After truncation there exists a line of descent $\sigma_0 \succ \sigma_1 \succ \dots$, and for each i there exists a pairwise disjoint collection A_i of C-splitting arcs of σ_i , so that A_i survives $\sigma_i \succ \sigma_{i+1}$ with descendant A_{i+1} , and so that if ρ_i is the train track obtained from σ_i by C-splitting along the arcs of A_i , then:*

- *The components of ρ_i can be listed as ρ_{ij} , $j = 1, \dots, J$, so that for each $j = 1, \dots, J$ the sequence $\rho_{0j} \succ \rho_{1j} \succ \dots$ consists of comb equivalences and infinitely many parity splittings and, ignoring its comb equivalences, it is a canonical expansion of \mathcal{F}_j on the surface $\text{Supp}(\mathcal{F}_j)$.*

- Letting ω_i be the maximal train track of the track with terminals $\tau_i \setminus \rho_i$, we obtain a sequence of comb equivalences and parity splittings $\omega_0 \succ \omega_1 \succ \dots$ which satisfies one of two possibilities: either it consists entirely of comb equivalences; or, ignoring its comb equivalences, it is an infinite twist splitting sequence. Moreover, there is a train track μ' comb equivalent to a subtrack of each of the ω_i such that $\cap_i \mathcal{MF}(\omega_i) = \mathcal{MF}(\mu') = \mathcal{MF}(\mu)$.

In these results, either the train track μ or the set $\{\mathcal{F}_j\}$ may be empty. If the latter is empty then the results reduce to Theorem 8.2.1.

Proof of Theorem 8.5.1. The proof of the theorem will be obtained by piecing together the conclusions of the Arational Expansion Theorem 6.3.2 and Theorem 8.2.1, using Section 5.3 as a bridge.

The first case is that $\tau_0 \succ \tau_1 \succ \dots$ does not infinitely split anything in $\cap_i \mathcal{MF}(\tau_i)$, in which case the theorem follows immediately from Theorem 8.2.1.

Consider now the case that $\tau_0 \succ \tau_1 \succ \dots$ infinitely splits some measured foliation which we denote \mathcal{F}_1 . By truncation we may assume there are no central splittings, and so the number of cusps of τ_i is constant independent of i . We will proceed by induction on this number.

Since \mathcal{F}_1 is infinitely split, it follows that some foliation component of \mathcal{F}_1 is infinitely split, and replacing \mathcal{F}_1 by that component we may assume that \mathcal{F}_1 has a single, infinitely split component. By Lemma 4.2.1, \mathcal{F}_1 is partial arational. Let $\pi_i \subset \tau_i$ be the carrying image of \mathcal{F}_1 , and by truncation we obtain a sequence $\pi_0 \succ \pi_1 \succ \dots$ consisting of comb equivalences and infinitely many parity splittings. Since there are no central splittings, we may choose \mathcal{F}_1 in its equivalence class so that for each i there is a carrying bijection $\mathcal{F}_1 \xleftrightarrow{\nu} \nu(\pi_i)$. Let Σ be the set of proper saddle connections of \mathcal{F}_1 . Applying the corollary of the Stability Lemma 6.7.1 to \mathcal{F}_1 , by truncation we may assume that each tie bundle $\nu(\pi_i) \rightarrow \pi_i$ is injective on Σ , taking each component of Σ to a C-splitting arc of π_i ; let B_i be this collection of C-splitting arcs, and so B_i survives $\pi_i \succ \pi_{i+1}$ with descendant B_{i+1} . Let ρ_{1i} be obtained from π_i by central splittings along the arcs of B_i . We may choose a carrying injection $\nu(\rho_{1i}) \subset \text{int } \nu(\pi_i)$, and note that $\nu(\rho_{1i})$ is the carrying image of a canonical model for \mathcal{F}_1 . Let τ'_i be the maximal train track in the track with terminals $\tau_i \setminus \rho_{1i}$, and by truncation we obtain a sequence of comb equivalences and parity splittings $\tau'_0 \succ \tau'_1 \succ \dots$. Applying Proposition 5.3.2 it follows that

$$\cap_i \mathcal{MF}(\tau_i) = \mathcal{MF}(\mathcal{F}_1) * \cap_i \mathcal{MF}(\tau'_i)$$

If the sequence $\tau'_0 \succ \tau'_1 \succ \dots$ has only finitely many parity splittings then we may truncate so that it consists entirely of comb equivalences, and the theorem follows.

If this sequence has infinitely many parity splittings and no element of $\cap_i \mathcal{MF}(\tau'_i)$ is infinitely split then we apply Theorem 8.2.1 and the theorem follows.

Consider then the remaining case that the sequence $\tau'_0 \succ \tau'_1 \succ \dots$ has infinitely many splittings and some element of $\cap_i \mathcal{MF}(\tau'_i)$ is infinitely split. The number of cusps of τ'_i is a constant independent of i , and moreover it is strictly less than the number of cusps of τ_i . By induction, the sequence satisfies the conclusions of the theorem, which we state as follows. After truncation, there exists a train track μ , and measured foliations $\mathcal{F}_2, \dots, \mathcal{F}_J$ for some $J \geq 2$, such that

$$\cap_i \mathcal{MF}(\tau'_i) = \mathcal{MF}(\mathcal{F}_2) * \dots * \mathcal{MF}(\mathcal{F}_J) * \mathcal{MF}(\mu)$$

and such that an element of $\cap_i \mathcal{MF}(\tau'_i)$ is infinitely split by the sequence $\tau'_0 \succ \tau'_1 \succ \dots$ if and only if it has a foliation component in one of $\mathcal{MF}(\mathcal{F}_2), \dots, \mathcal{MF}(\mathcal{F}_J)$. Together with the previous equation, gives us the desired expression for $\cap_i \mathcal{MF}(\tau_i)$, and it follows that an element of $\cap_i \mathcal{MF}(\tau_i)$ is infinitely split by the sequence $\tau_0 \succ \tau_1 \succ \dots$ if and only if it has a foliation component in one of $\mathcal{MF}(\mathcal{F}_1), \mathcal{MF}(\mathcal{F}_2), \dots, \mathcal{MF}(\mathcal{F}_J)$. \diamond

Proof of Corollary 8.5.2. Let σ_i be the subtrack of τ_i that fully carries the measured foliation $\mathcal{F} = \mathcal{F}_1 \cup \dots \cup \mathcal{F}_J$, and truncate so that the sequence $\sigma_0 \succ \sigma_1 \succ \dots$ consists of comb equivalences and infinitely many parity splittings. Choose \mathcal{F} in its equivalence class so that there is a carrying bijection $\mathcal{F} \xrightarrow{\nu} \nu(\sigma_0)$, and let Σ be the union of proper saddle connections of \mathcal{F} . Since $\sigma_0 \succ \sigma_1 \succ \dots$ has no central splittings it follows that there is a carrying bijection $\mathcal{F} \xrightarrow{\nu} \nu(\sigma_i)$. By Theorem 8.5.1, each of $\mathcal{F}_1, \dots, \mathcal{F}_J$ is infinitely split, and so by the corollary to the Stability Lemma 6.7.1 we may truncate so that the image of Σ under the carrying map $\mathcal{F} \rightarrow \sigma_i$ is a pairwise disjoint collection of saddle connections A_i . Let ρ_i be obtained from σ_i by central splitting along all the arcs of A_i . It follows that ρ_i has components $\rho_{i1}, \dots, \rho_{iJ}$ so that ρ_{ij} canonically carries \mathcal{F}_j and so that $\rho_{0j} \succ \rho_{1j} \succ \dots$ consists of comb equivalences and infinitely many parity splittings, and is a canonical expansion of \mathcal{F}_j .

Let ω_i be the maximal train track of the track with terminals $\tau_i \setminus \rho_i$. Applying Proposition 5.3.2 J times in succession it follows that for any $\mathcal{G} \in \cap_i \mathcal{MF}(\tau_i)$, if \mathcal{G} has no component in $\mathcal{MF}(\mathcal{F}_1), \dots, \mathcal{MF}(\mathcal{F}_J)$, equivalently, if $\mathcal{G} \in \mathcal{MF}(\mu)$, then $\mathcal{G} \in \cap_i \mathcal{MF}(\omega_i)$. Conversely, if $\mathcal{G} \in \cap_i \mathcal{MF}(\omega_i)$ then \mathcal{G} is not infinitely split by the sequence $\tau_0 \succ \tau_1 \succ \dots$ and so $\mathcal{G} \in \mathcal{MF}(\mu)$. This shows that $\cap_i \mathcal{MF}(\omega_i) = \mathcal{MF}(\mu)$. It is clear from this that nothing in $\cap_i \mathcal{MF}(\omega_i)$ is infinitely split by the sequence $\omega_0 \succ \omega_1 \succ \dots$, and so by Theorem 8.2.1 it also follows that there is a train track μ' comb equivalent to a subtrack of each ω_i such that $\cap_i \mathcal{MF}(\omega_i) = \mathcal{MF}(\mu')$. \diamond

9 Constructing pseudo-Anosov homeomorphisms

TO DO:

- Given the FDA point of view.
- Recall Γ , Γ^d from earlier section, Section ??.
- Define splitting circuits, canonical splitting circuits, etc.
- Corollary: pseudo-Anosov splitting circuits form a regular language, in both noncanonical and canonical versions. For each singularity type, canonical loops of that type form a regular language.
- In each construction, emphasize the regular language point of view.
- In the one-sink section, explain the connection with interval exchange maps, in the case of oriented train tracks. Explain Kerckhoff's contribution here.

Many constructions of pseudo-Anosov mapping classes proceed by first producing the mapping class ϕ together with an “invariant train track” for ϕ , and then applying some method which uses the invariant train track to verify the pseudo-Anosov property for ϕ . An exception to this is Thurston's original construction in [Thu88], and similar constructions of Teichmüller curves [Vee89], [McM02a], which are far outside the scope of this work.

In this section we will apply the results of Sections 6 and 8 to give proofs of several constructions of pseudo-Anosov mapping classes, all deriving from the same general method involving invariant train tracks.

Section 9.1 will describe a general method which, starting from a homeomorphism ϕ and an invariant train track τ , decides whether the mapping class of ϕ is pseudo-Anosov. Such methods were commonly discussed in the early 1980's, but the first careful description of a general method is due to Casson [CB], which has been a useful method in the literature for checking when mapping classes are pseudo-Anosov, e.g. [Bau92]. In Theorem 9.1.1 we describe Casson's method as well as several other equivalent methods which offer some computational improvements. In Theorem 9.1.2 we will also describe several equivalent methods which, when ϕ is pseudo-Anosov, allow one to decide whether the singularity type of ϕ matches the singularity type of τ .

Section 9.1 also contains two general algorithms derived from Theorems 9.1.1 and 9.1.2, one using an invariant train track to decide whether a mapping class is pseudo-Anosov, the other which decides equality of singularity types of a pseudo-Anosov mapping class and an invariant train track. These algorithms ought to be

efficiently programmable on the computer. Similar algorithms in [CB] require some linear programming, but we avoid this by using splitting cycles and C -splitting arcs.

In the following sections, we illustrate the methods of Theorems 9.1.1 and 9.1.2 by using them to justify some constructions of pseudo-Anosov homeomorphisms from the literature, including: a recipe due to R. Penner [Pen88], in Section 9.3; and in Section 9.4 a construction due to this author [Mos86], which generates some power of every pseudo-Anosov mapping class. These constructions all share a certain flavor, reminiscent of the iterated killing criteria of Section 6. In each of these constructions, special classes of invariant train tracks are used to find special algorithms which are much more efficient than the general algorithm, although they ultimately derive from the general algorithm.

9.1 Invariant train tracks

Given a homeomorphism $\phi: S \rightarrow S$, an *invariant train track* of ϕ is a train track τ satisfying $\tau \succ \phi(\tau)$. In this case the action of ϕ , and the identification of $\mathcal{MF}(\tau)$ with $W(\tau)$, induces a linear map $\phi: W(\tau) \rightarrow W(\tau)$.

In Theorem 9.1.1 we give several equivalent criteria which employ a ϕ -invariant train track τ to decide whether the mapping class of ϕ is pseudo-Anosov. One such criterion, due to Casson [CB], is included as item (2) in the theorem. Other criteria offer some computational improvements. The proof of equivalence of these criteria depend on the results of Sections 6 and 8, although as remarked after the proof, if one is willing to weaken Casson's criterion then one can get a proof of equivalence using only the results of Section 6.

When ϕ is pseudo-Anosov, Theorem 9.1.1 also describes how to obtain the unstable foliation \mathcal{F} of ϕ from the action of ϕ on $\mathcal{PMF}(\tau)$, namely, $\bigcap_i \mathcal{PMF}(\phi^i(\tau)) = \{\mathcal{P}[\mathcal{F}]\}$.

We will also prove Theorem 9.1.2 which, in the case that ϕ is pseudo-Anosov, gives several equivalent criteria for deciding whether a ϕ -invariant train track τ canonically carries the unstable foliation for ϕ ; if this happens then we say that τ is a *canonical invariant train track*.

In order to state these theorems we must describe the relation between invariant train tracks, periodic splitting sequences, and associated measured foliations.

Given a mapping class ϕ , an infinite splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ is said to be ϕ -*periodic* if there exists $N > 0$ such that $\phi(\tau_i)$ is isotopic to τ_{i+N} for all $i \geq 0$. The number N is the *period* of the sequence. Notice that the number of cusps $i(\tau_n)$ is a constant independent of n ; this follows from ϕ -periodicity together with the fact that $i(\tau_n)$ is nonincreasing (Corollary 3.14.2).

There is a very close connection between ϕ -invariant train tracks and ϕ -periodic splitting sequences. If $\tau_0 \succ \tau_1 \succ \dots$ is ϕ -periodic then clearly τ_0 is ϕ -invariant.

Conversely, if τ is ϕ -periodic then there exists a ϕ -periodic splitting sequence $\tau = \tau_0 \succ \cdots \succ \tau_N = \phi(\tau_0) \succ \cdots$. To see why the converse is true, first note that the carrying $\tau \succ \phi(\tau)$ is full, for otherwise there is a proper subtrack $\sigma \subset \tau$ that fully carries $\phi(\tau)$, and so $i(\tau) > i(\sigma) \geq i(\phi(\tau)) = i(\tau)$, a contradiction; recall that $i(\cdot)$ denotes the number of cusps, and the second inequality is a consequence of Corollary 3.14.2. Knowing that $\tau \not\stackrel{F}{\succ} \phi(\tau)$, Corollary 3.14.2 and the fact that $i(\tau) = i(\phi(\tau))$ implies that $\tau \stackrel{H}{\succ} \phi(\tau)$, and the existence of a finite splitting sequence $\tau = \tau_0 \succ \cdots \succ \tau_N = \phi(\tau)$ to $\phi(\tau)$ now follows from Proposition 3.14.1. This finite sequence evidently extends in a unique way (up to isotopy) to an infinite ϕ -periodic splitting sequence.

Theorem 9.1.1. *Given a homeomorphism $\phi: S \rightarrow S$, an invariant train track τ for ϕ , and a ϕ -periodic splitting sequence $\tau = \tau_0 \succ \tau_1 \succ \cdots \succ \tau_N = \phi(\tau_0) \succ \cdots$, the following are equivalent:*

- (1) *The mapping class of ϕ is pseudo-Anosov.*
- (2) [CB] *Every ϕ -invariant subtrack $\sigma \subset \tau_0$ fills S and satisfies the following:*
 - C_σ : *For every diagonal extension ρ of σ which is also an invariant train track of ϕ , the induced map $\phi: W(\rho) \rightarrow W(\rho)$ has no fixed points.*
- (3) *Every ϕ -invariant subtrack $\sigma \subset \tau_0$ fills S and satisfies the following:*
 - D_σ : *The carrying $\sigma \succ \phi(\sigma)$ is not a comb equivalence, and there does not exist a sequence of splitting cycles $\gamma_1, \dots, \gamma_N$ of σ such that $\phi(\gamma_i)$ is isotopic to γ_{i+1} for each $i \in \mathbf{Z}/N$.*
- (4) *For each subtrack $\sigma \subset \tau_0$, if σ is nonfilling then it is eventually killed, and if σ is not eventually killed then each splitting cycle of σ is eventually killed.*
- (5) *The splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ satisfies the iterated rational killing criterion.*

Furthermore, if these properties hold, and if \mathcal{F} is an unstable measured foliation for ϕ , then

$$\bigcap_{j \geq 0} \mathcal{PMF}(\phi^j(\tau_0)) = \bigcap_{i \geq 0} \mathcal{PMF}(\tau_i) = \{\mathcal{P}[\mathcal{F}]\}$$

Assuming that ϕ is pseudo-Anosov, we can also get a list of equivalent conditions that characterize when the splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ is a canonical expansion of the unstable foliation \mathcal{F} .

Theorem 9.1.2. *Given a pseudo-Anosov homeomorphism $\phi: S \rightarrow S$ with unstable measured foliation \mathcal{F} , and given an invariant train track τ for ϕ and a ϕ -periodic splitting sequence $\tau = \tau_0 \succ \tau_1 \succ \cdots \succ \tau_N = \phi(\tau_0) \succ \cdots$, the following are equivalent:*

- (1) *The singularity types of \mathcal{F} and of τ are identical.*
- (2) *τ canonically carries \mathcal{F} .*
- (3) *$\tau_0 \succ \tau_1 \succ \cdots$ is a canonical expansion of \mathcal{F} .*
- (4) *$\tau_0 \succ \tau_1 \succ \cdots$ satisfies the iterated canonical killing criterion.*
- (5) *Each proper subtrack of τ_0 is eventually killed, and each C-splitting arc of τ_0 is eventually killed.*
- (6) *There are no proper ϕ -invariant subtracks of τ_0 , and there does not exist a sequence of C-splitting arcs $\alpha_1, \dots, \alpha_N$ of τ_0 such that for each $i \in \mathbf{Z}/N$ the splitting arc α_i survives the carrying $\tau \succ \phi(\tau)$ with descendant $\phi(\alpha_{i-1})$.*

When the conditions of Theorem 9.1.2 hold then we say that τ is a *canonical invariant train track* for ϕ .

As we shall see in the following sections, item 5 is very useful for checking that ϕ is pseudo-Anosov with the correct singularity type.

Before turning to the proofs of Theorems 9.1.1 and 9.1.2, we first show how they can be used to develop some algorithms for deciding the pseudo-Anosov property and the canonical invariant train track property.

Algorithms. As pointed out in [CB] the equivalence of (1) and (2) in Theorem 9.1.1 provides the basis for an algorithm which, given a mapping class ϕ and an invariant train track τ for ϕ , decides whether ϕ is pseudo-Anosov. The most laborious part of this procedure is enumerating diagonal extensions ρ which are also invariant train tracks and deciding whether any of the linear maps $\phi: W(\rho) \rightarrow W(\rho)$ have a fixed point.

Item (3) in Theorem 9.1.1 reduces the labor by an algorithm employing splitting cycles. Item (6) in Theorem 9.1.2 uses a similar algorithm to decide whether τ is a canonical invariant train track, thereby deriving the singularity type of ϕ . Here is a description of these algorithms.

First we need an important subroutine, namely, an enumeration of the ϕ -invariant subtracks of τ . Choose a splitting sequence $\tau = \tau_0 \succ \tau_1 \succ \cdots \succ \tau_n = \phi(\tau)$. For each subtrack σ , use Propositions 6.1.3 and Lemma 6.1.4 to decide whether σ survives in τ_n ; if it does survive let its sequence of descendants be $\sigma = \sigma_0 \succ \sigma_1 \succ \cdots \succ \sigma_n =$

$\sigma' \subset \tau_n$. Now compare σ' to $\phi(\sigma)$, and we claim that σ is a ϕ invariant subtrack if and only if $\sigma' = \phi(\sigma)$. One direction of this claim is clear; for the other direction, if σ is ϕ -invariant then the fact that $\sigma \succ \phi(\sigma)$ and the fact that σ and $\phi(\sigma)$ have the same number of cusps implies, by Corollary 3.14.2, that σ survives the carrying $\tau \succ \phi(\tau)$ with descendant $\phi(\sigma)$.

Another useful subroutine checks, for each ϕ -invariant subtrack σ , whether the carrying $\sigma \succ \phi(\sigma)$ is a comb equivalence. To decide this quickly, note that each carrying $\sigma_i \succ \sigma_{i+1}$ is either a comb equivalence or a splitting, and the distinction between the two is easily decidable; then $\sigma \succ \phi(\sigma)$ is a comb equivalence if and only if the entire sequence $\sigma = \sigma_0 \succ \cdots \succ \sigma_n = \phi(\sigma)$ consists of comb equivalences.

A pseudo-Anosov algorithm. Here is the algorithm that carries out property (3) of Theorem 9.1.1 and thus decides whether ϕ is pseudo-Anosov:

- (1) Search for invariant subtracks σ of τ which do not fill S . If one is found, stop: ϕ is not pseudo-Anosov. Otherwise continue.
- (2) Search for invariant subtracks σ such that $\sigma \succ \phi(\sigma)$ is a comb equivalence. If one is found, stop: ϕ is not pseudo-Anosov. Otherwise continue.
- (3) For each invariant subtrack σ — which at this stage is guaranteed to fill S and not to be comb equivalent to $\phi(\sigma)$ — do the following:
 - (a) Enumerate the set of splitting cycles Γ of σ .
 - (b) Compute the action of ϕ on the set Γ as follows: for each $\gamma \in \Gamma$, decide whether γ survives the carrying $\sigma \succ \phi(\sigma)$ — using, say, Lemma 6.2.3 together with a splitting sequence from σ to $\phi(\sigma)$. If γ survives the carrying $\sigma \succ \phi(\sigma)$, with descendant γ' a splitting cycle of $\phi(\sigma)$, then set $\phi(\gamma') = \gamma$.
 - (c) Decide whether the action of ϕ on Γ has a periodic point. If one is found, stop: ϕ is not pseudo-Anosov. Otherwise continue with the next σ .
- (4) If the algorithm reaches this step, ϕ is pseudo-Anosov.

An algorithm for canonical invariant train tracks. Here is the algorithm which carries out item (6) of Theorem 9.1.2, thereby computing the singularity type of a pseudo-Anosov homeomorphism ϕ :

- (1) Search for proper invariant subtracks of τ . If one is found, stop: τ is not a canonical invariant train track for ϕ .

- (2) If $\tau \succ \phi(\tau)$ is a comb equivalence, stop: τ is not a canonical invariant train track for ϕ .
- (3) Enumerate the set \mathcal{A} of splitting arcs of τ .
- (4) Compute the action of ϕ on \mathcal{A} as follows: for each $\alpha \in \mathcal{A}$, decide whether α survives the carrying $\tau \succ \phi(\tau)$, and if it does, with descendant α' a splitting arc of $\phi(\tau)$, set $\phi(\alpha') = \alpha$.
- (5) Decide whether the action of ϕ on \mathcal{A} has a periodic orbit. If one is found stop: τ is not a canonical invariant train track for ϕ .
- (6) If the algorithm reaches this step, τ is a canonical invariant train track for ϕ .

Now we turn to:

Proof of Theorem 9.1.1. Let ϕ be a mapping class with invariant train track τ and with ϕ -periodic splitting sequence $\tau = \tau_0 \succ \cdots \succ \tau_N = \phi(\tau_0) \succ \cdots$.

The last sentence of the proof, regarding the unstable foliation, will be proved below as part of the proof that (5) \implies (1).

For any subtrack $\sigma \subset \tau$, we note that if σ is ϕ -invariant then σ survives forever with line of descent $\sigma = \sigma_0 \succ \sigma_1 \succ \cdots$, for subtracks $\sigma_i \subset \tau_i$.

One fact we will use more than once is that *every* line of descent for $\tau_0 \succ \tau_1 \succ \cdots$ is a truncation of a line of descent that starts in τ_0 . To see why, consider a line of descent $\sigma_I \succ \sigma_{I+1} \succ \cdots$ starting in some τ_I . We claim that the sequence $\sigma_I \succ \sigma_{I+1} \succ \cdots$ is periodic with respect to some power ϕ^M and with period MN , where N is the period of $\tau_0 \succ \tau_1 \succ \cdots$. To see why this is true, the train tracks $\phi^{-m}(\sigma_{I+mN})$ are all subtracks of τ_I , and so they must eventually repeat, say, $\phi^{-m}(\sigma_{I+mN}) = \phi^{-m'}(\sigma_{I+m'N})$ with $m < m'$. Setting $M = m' - m$ the claim follows by ϕ periodicity of the sequence $\tau_0 \succ \tau_1 \succ \cdots$. Now we can extend the line of descent $\sigma_I \succ \sigma_{I+1} \succ \cdots$ back to σ_0 using the periodicity formula $\phi^M(\sigma_i) = \sigma_{i+MN}$, and the result is evidently a line of descent starting in τ_0 .

Proof of (4) \iff (5). Note that (5) obviously implies (4).

For the converse, item (4) says that lines of descent *that start in* τ_0 satisfy the requirements of the iterated rational killing criterion, whereas item (5) says the same thing for *all* lines of descent. Thus, we need only note that every line of descent is a truncation of a line of descent that does indeed start in τ_0 , as just proved above.

Proof of (5) \implies (1). The map $\phi: \mathcal{PMF}(\tau_0) \rightarrow \mathcal{PMF}(\tau_N)$ composed with the inclusion $\mathcal{PMF}(\tau_N) \hookrightarrow \mathcal{PMF}(\tau_0)$ defines a map $\phi: \mathcal{PMF}(\tau_0) \rightarrow \mathcal{PMF}(\tau_0)$ which, by the Brouwer fixed point theorem, has a fixed point $\mathcal{P}[\mathcal{F}]$.

Since $\phi(\mathcal{PMF}(\tau_i)) = \mathcal{PMF}(\tau_{i+N})$ it follows that $\mathcal{P}[\mathcal{F}] \in \bigcap_i \mathcal{PMF}(\tau_i)$, that is, $\tau_0 \succ \tau_1 \succ \dots$ is a train track expansion of \mathcal{F} . Since $\tau_0 \succ \tau_1 \succ \dots$ satisfies the iterated rational killing criterion it follows that \mathcal{F} is arational.

Choose coordinates for $W(\tau_0)$, so that the self-map $\phi: \mathcal{MF}(\tau_0) \rightarrow \mathcal{MF}(\tau_0)$ is expressed as a linear map $W(\tau_0) \rightarrow W(\tau_0)$ given by a square integer matrix M . It follows that \mathcal{F} is expressed as a non-negative eigenvector of M , with eigenvalue $\lambda > 0$, and since M is an integer matrix we clearly have $\lambda \geq 1$.

Suppose that $\lambda = 1$. Let $\sigma_i \subset \tau_i$ be the subtrack fully carrying \mathcal{F} , and let w_i be the invariant weight on σ_i induced by \mathcal{F} , so the carrying induced injection $W(\sigma_{i+N}) \hookrightarrow W(\sigma_i)$ takes w_{i+N} to w_i . The line of descent $\sigma_0 \succ \sigma_1 \succ \dots$ is periodic for ϕ , and moreover the map $\phi: W(\sigma_i) \rightarrow W(\sigma_{i+N})$ takes w_i to w_{i+N} , because of the assumption that $\lambda = 1$. It follows that none of the entries of w_i approach zero as $i \rightarrow \infty$, and so by Lemma 4.2.2 the carrying $\sigma_i \succ \sigma_{i+1}$ is a comb equivalence for sufficiently large i . But by ϕ periodicity it follows that the entire sequence consists of comb equivalences $\sigma_0 \approx \sigma_1 \approx \dots$. This implies that *every* splitting cycle of σ_0 survives infinitely, contradicting the iterated rational killing criterion for $\tau_0 \succ \tau_1 \succ \dots$.

We conclude therefore that $\lambda > 1$. Now we can quote [FLP⁺79] for the theorem that if ϕ is a homeomorphism, \mathcal{F} is a measured foliation, $\lambda > 1$, and $\phi[\mathcal{F}] = \lambda[\mathcal{F}]$, then ϕ is isotopic to a pseudo-Anosov homeomorphism whose unstable foliation is equivalent to \mathcal{F} . This can be seen directly by showing that the action of ϕ^{-1} on $W^\perp(\sigma_i)$ has an eigenvector of eigenvalue λ , representing a measured foliation \mathcal{F}' , such that $\mathcal{F}, \mathcal{F}'$ are transverse and so that ϕ is isotopic to a pseudo-Anosov homeomorphism with stable foliation \mathcal{F}' and unstable foliation \mathcal{F} .

From Theorem 5.1.1 we have that

$$\bigcap_j \mathcal{PMF}(\phi^j(\tau)) = \bigcap_i \mathcal{PMF}(\tau_i) = \mathcal{PMF}(\mathcal{F})$$

The fact that \mathcal{F} is the unstable foliation of a pseudo-Anosov homeomorphism ϕ is well known to imply that \mathcal{F} is uniquely ergodic — briefly, the action of ϕ on the Choquet simplex $\mathcal{PMF}(\mathcal{F})$ would have to preserve the set of extremal points of the simplex, giving a finite periodic set for ϕ . However, ϕ acts on \mathcal{PMF} with source-sink dynamics, the stable foliation being a repelling fixed point and the unstable foliation being an attracting fixed point, and so $\mathcal{PMF}(\mathcal{F})$ must be a zero-dimensional simplex, that is, $\mathcal{PMF}(\mathcal{F}) = \{\mathcal{P}[\mathcal{F}]\}$. This proves the last sentence of Theorem 9.1.1.

Proof of (1) \implies (2). Suppose that ϕ is pseudo-Anosov. If σ is any ϕ -invariant train track then σ must fill S , because otherwise the boundary of the support of σ is a ϕ -invariant essential curve system. Also, ϕ has no fixed points in \mathcal{MF} , and so ϕ has no fixed point in $W(\rho)$ for any invariant train track ρ .

Proof of (2) \implies (3). Consider a ϕ -invariant subtrack $\sigma \subset \tau$ with line of descent $\sigma = \sigma_0 \succ \sigma_1 \succ \dots$. Assuming that σ fills S and satisfies criterion C_σ , we prove that σ satisfies criterion D_σ . Suppose then that D_σ fails.

One way D_σ can fail is if the carrying $\sigma \succ \phi(\sigma)$ is a comb equivalence, and in this case some power of ϕ acts as the identity on $W(\sigma)$. Since $\mathcal{PW}(\sigma)$ is a compact disc, the action of ϕ on $\mathcal{PW}(\sigma)$ has a fixed point, and so the action of ϕ on $W(\sigma)$ has an eigenvector with positive eigenvalue λ . But since some power of ϕ is the identity on $W(\sigma)$, it follows that $\lambda = 1$, that is, ϕ has a fixed point in $W(\sigma)$, contradicting (2).

The other way D_σ can fail is if there exists a sequence of splitting cycles $\gamma_1, \dots, \gamma_N$ of σ such that $\phi(\gamma_i)$ is isotopic to γ_{i+1} for each $i \in \mathbf{Z}/N$.

We consider the case that each of the splitting cycles $\gamma_1, \dots, \gamma_N$ is a circle; the case of a proper line will be sketched later. Let $\Gamma = \{\gamma_1, \dots, \gamma_N\}$. Choose ϕ in its isotopy class, and choose the elements of Γ in their isotopy classes rel σ , so that Γ is in general position in $S - \sigma$, and so that Γ is ϕ -equivariant, meaning that for each $\gamma \in \Gamma$ the closed curve $\phi(\gamma)$ is an element of Γ , not just isotopic to an element of Γ . For each transverse intersection point x of Γ in $S - \sigma$, choose a surgery direction on x , replacing the two crossing strands $\overline{ab}, \overline{cd}$ near x by two disjoint strands with the same set of four endpoints; there are two ways to make this choice, either $\overline{ac}, \overline{bd}$, or $\overline{ad}, \overline{bc}$, and the choice surgery directions of each x should be made in a ϕ -equivariant manner. This has the effect of replacing the curve system Γ by a new ϕ -equivariant curve system without transverse intersection points in $S - \sigma$; the inessential components of this curve system are all contained in the complement of σ , and so by deleting them we obtain a ϕ -equivariant essential curve system Γ' , such that each component of $\Gamma' - \sigma$ is a diagonal of σ . This implies that there is a ϕ -invariant diagonal extension ρ of σ such that each element of Γ' is carried by ρ . Moreover, each $\gamma' \in \Gamma'$ induces an invariant weight $w_{\gamma'} \in W(\rho)$, and $\sum_{\gamma' \in \Gamma'} w_{\gamma'}$ is a fixed point for ϕ in $W(\rho)$, again contradicting (2).

In the case that each $\gamma_1, \dots, \gamma_N$ is a proper line, the elements of the set Γ are not the proper lines $\gamma_1, \dots, \gamma_N$ themselves, but instead closed curves homotopic to the boundaries of regular neighborhoods of $\gamma_1, \dots, \gamma_N$. The rest of the proof goes through.

Having proved that (4) \iff (5) \implies (1) \implies (2) \implies (3), the following will finish the proof of the theorem:

Proof of (3) \implies (5). This proof uses the results on the structure of general splitting sequences from Section 8. We break into two cases, depending on whether or not some element of $\cap_i \mathcal{MF}(\tau_i)$ is infinitely split by the sequence $\tau_0 \succ \tau_1 \succ \dots$.

Case 1: No element of $\cap_i \mathcal{MF}(\tau_i)$ is infinitely split. Applying Theorem 8.2.1, the sequence $\tau_0 \succ \tau_1 \succ \dots$ is eventually an infinite twist sequence, and it then follows by ϕ -periodicity that it is an infinite twist sequence starting from τ_0 . Let C be the essential curve family along which $\tau_0 \succ \tau_1 \succ \dots$ is infinitely twisted. Let $\sigma_i \subset \tau_i$ be the carrying image of C in τ_i , and so $\sigma_0 \succ \sigma_1 \succ \dots$ is a line of descent. Noting that C is also the curve system along which the truncated sequence $\tau_N = \phi(\tau_0) \succ \tau_{N+1} = \phi(\tau_1) \succ \dots$ is infinitely twisted, it follows by naturality that $\phi(C) = C$, and so $\phi(\sigma_i) = \sigma_{i+N}$. However, since C is a family of twist curves for τ_0 , for each $c \in C$ the twist annulus A_c has interior disjoint from σ_0 , and so σ_0 does not fill S . This contradicts item (3), and so Case 1 cannot occur.

Case 2: Some element of $\cap_i \mathcal{MF}(\tau_i)$ is infinitely split. By applying Corollary 8.5.2, we obtain the following description of the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$.

For some $J \geq 1$ we have a set of partial arational measured foliations $\mathcal{F}^1, \dots, \mathcal{F}^J$ with pairwise disjoint supports H^1, \dots, H^J . We have a line of descent $\sigma_0 \succ \sigma_1 \succ \dots$ and a collection of pairwise disjoint C-splitting arcs A_i of σ_i , such that A_i survives $\sigma_i \succ \sigma_{i+1}$ with descendant A_{i+1} . Letting ρ_i be the train track obtained from σ_i by simultaneously splitting along the arcs in A_i , we have a carrying sequence $\rho_0 \succ \rho_1 \succ \dots$. The components of each train track ρ_i can be listed as $\rho_i^1, \dots, \rho_i^J$, such that for each $j = 1, \dots, J$ we have a sequence $\rho_0^j \succ \rho_1^j \succ \dots$ consisting entirely of comb equivalences and infinitely many splittings, so that $\text{Supp}(\rho_i^j)$ is isotopic to H^j , and so that $\rho_0^j \succ \rho_1^j \succ \dots$ is a canonical expansion of \mathcal{F}^j . Finally, a partial arational measured foliation in $\cap_i \mathcal{MF}(\tau_i)$ is infinitely split if and only if it is contained in one of $\mathcal{MF}(\mathcal{F}^1), \dots, \mathcal{MF}(\mathcal{F}^J)$.

We claim that the line of descent $\sigma_0 \succ \sigma_1 \succ \dots$ is ϕ -periodic. To see why, first note that the action of ϕ on \mathcal{MF} must preserve the subset $\{\mathcal{MF}(\mathcal{F}^1), \dots, \mathcal{MF}(\mathcal{F}^J)\}$, because this is the set of classes of partial arational elements of $\cap_i \mathcal{MF}(\tau_i)$ that are infinitely split by the sequence $\tau_0 \succ \tau_1 \succ \dots$, and it is also the set of classes of partial arational elements of $\phi(\cap_i \mathcal{MF}(\tau_i)) = \cap_i \mathcal{MF}(\phi(\tau_i))$ that are infinitely split by the sequence $\phi(\tau_0) = \tau_N \succ \phi(\tau_1) = \tau_{N+1} \succ \dots$. Thus, the carrying image of $\mathcal{F}^1 \cup \dots \cup \mathcal{F}^J$ in τ_i , which is just σ_i , is taken by ϕ to the carrying image in τ_{i+N} , which is just σ_{i+N} , proving the claim.

From item (3) it follows that σ_0 fills S and that no splitting cycle of σ_0 survives infinitely. But this immediately implies that ρ_0 is connected, that is, $J = 1$, and that ρ_0 fills S : if either of these failed, then from the set of splitting arcs A_0 we

could extract a subset which lie on an infinitely surviving splitting cycle of σ_0 . Thus, $\mathcal{F} = \mathcal{F}^1$ is arational, and since $\tau_0 \succ \tau_1 \succ \dots$ is an expansion of \mathcal{F}^1 it follows from Theorem 6.3.2 that $\tau_0 \succ \tau_1 \succ \dots$ satisfies the iterated rational killing criterion. \diamond

Remark on the proof. In the above proof, we used the results of Section 8 to get the final implication (3) \implies (5). If we allowed ourselves to weaken items (2) and (3) in the theorem, then we could avoid using the results of Section 8. The point is that instead of quantifying over every ϕ -invariant subtrack of τ_0 in items (2) and (3), we could instead quantify over every sequence of subtracks $\sigma^1, \dots, \sigma^M$ of τ_0 such that $\phi(\sigma^m) \not\asymp \sigma^{m+1}$, where m varies cyclically over \mathbf{Z}/M ; we'll call this a ϕ -periodic subtrack orbit. For instance, in place of item (3) we would have:

(3') For every ϕ -periodic subtrack orbit $\sigma^1, \dots, \sigma^M$ of τ_0 , each σ^m fills S and there does not exist a sequence $\gamma_1, \dots, \gamma_{MK}$, such that γ_{m+Mk} is a splitting cycle of σ^m for each $k = 1, \dots, K$ and each $m = 1, \dots, M$, and such that $\phi(\gamma_i)$ is isotopic to γ_{i+1} for each $i \in \mathbf{Z}/MK$.

One can similarly restate item (2) to get (2') by quantifying over ϕ -periodic subtrack orbits rather than over ϕ -invariant subtracks. The point is that (3) obviously implies that $\tau_0 \succ \tau_1 \succ \dots$ satisfies the iterated rational killing criterion, *without* quoting the results of Section 8. The implications (1) \implies (2') and (2') \implies (3') are proved exactly as before.

Now we turn to:

Proof of Theorem 9.1.2. We can make use of Theorem 9.1.1 in several places to make the proof shorter.

Proof of (1) \iff (2). The direction (2) \implies (1) is obvious. For the other direction, suppose τ and \mathcal{F} have the same singularity type. Let \mathcal{F}' be a partial measured foliation equivalent to \mathcal{F} which has a carrying injection into a tie bundle $\nu(\tau)$. Slicing along all proper saddle connections of \mathcal{F}' we obtain a canonical model \mathcal{F}'' of \mathcal{F} carried on $\nu(\tau)$. If τ does not canonically carry \mathcal{F} then it follows that either \mathcal{F}' is not fully carried by τ , or \mathcal{F}' is fully carried and there is at least one proper saddle connection; in either case, it follows that the number of boundary singularities of \mathcal{F}'' is strictly less than the number of cusps of τ , and so the singularity type of τ is not equal to the singularity type of \mathcal{F}'' (which equals the singularity type of \mathcal{F}).

Proof of (2) \implies (3). We know from Theorem 9.1.1 that $\tau_0 \succ \tau_1 \succ \dots$ is an expansion of \mathcal{F} , and we also know from item (2) that $\tau = \tau_0$ canonically carries \mathcal{F} . We need only observe that canonical carrying of \mathcal{F} follows for each τ_i , $i \geq 0$. To see

why this is true, suppose by induction that τ_i canonically carries \mathcal{F} . Since $\tau_i \succ \tau_{i+1}$ is a parity splitting and since τ_{i+1} carries \mathcal{F} it follows that τ_{i+1} fully carries \mathcal{F} . If τ_{i+1} did not canonically carry \mathcal{F} then there would exist a partial measured foliation \mathcal{F}' equivalent to \mathcal{F} and a carrying bijection $\mathcal{F}' \xrightarrow{\nu} \nu(\tau_{i+1})$ such that \mathcal{F}' has at least one proper saddle connection, but composing with a surjection $\nu(\tau_{i+1}) \rightarrow \nu(\tau_i)$ which takes each tie into a tie and is injective on the interior we obtain a carrying bijection $\mathcal{F}' \xrightarrow{\nu} \nu(\tau_i)$, contradicting that τ_i canonically carries \mathcal{F} .

Proof of (3) \implies (4) \implies (5) \implies (6). The implication (3) \implies (4) is an immediate consequence of Theorem 6.4.2, (4) \implies (5) is obvious, and (5) \implies (6) follows from the fact that every ϕ -invariant subtrack of τ_0 survives forever, and every C-splitting arc of τ_0 contained in a sequence as in (6) survives forever.

Proof of (6) \implies (2). From Theorem 9.1.1 we know that $\cap_i \mathcal{MF}(\tau_i) = \{\mathcal{F}\}$. Letting $\sigma_i \subset \tau_i$ be the subtrack that fully carries \mathcal{F} , since the map $\phi: \mathcal{PMF}(\tau_0) \rightarrow \mathcal{PMF}(\tau_N)$ takes $\mathcal{P}[\mathcal{F}]$ to $\mathcal{P}[\mathcal{F}]$ it follows that ϕ takes σ_0 to σ_N , and so σ_0 is an invariant subtrack. By item (6) we must have $\sigma_0 = \tau_0$, that is, \mathcal{F} must be fully carried by τ_0 . Choose \mathcal{F}' equivalent to \mathcal{F} so that there is a carrying bijection $\mathcal{F}' \xrightarrow{\nu} \nu(\tau_0)$. By the corollary to the Stability Lemma 6.7.1, and by ϕ -periodicity, any proper saddle connection of \mathcal{F} yields a C-splitting arc of τ_0 which survives infinitely, which is ruled out by item (6). It follows that \mathcal{F}' is a canonical model for \mathcal{F} , and hence τ_0 canonically carries \mathcal{F} . ◇

9.2 Automata for canonical splitting circuits (in progress).

The results of the previous section can be re-interpreted in terms of finite automata and regular languages. In particular, given a splitting sequence $\tau_0 \succ \cdots \succ \tau_n = \Phi(\tau_0) \succ \cdots$ which is periodic with respect to some mapping class Φ , the period loop $\tau_0 \succ \cdots \succ \tau_n = \Phi(\tau_0)$ can be interpreted as a cyclic word in a certain language; this word is called a “splitting circuit”. We shall study the sublanguage of all “canonical splitting circuits”, meaning those splitting circuits that arise from periodic, canonical expansions of unstable foliations of pseudo-Anosov mapping classes. Our main theorem is that the set of canonical splitting circuits forms a regular language. The tools that we describe here will be applicable in many other situations where we study various killing criteria on splitting circuits, and we show in several such situations that the class of such circuits is a regular language.

The idea behind these results is that the killing criteria that we use to study splitting sequences all have a finitistic character that makes them amenable to study using finite automata.

Automata and regular languages. Let \mathcal{A} be a finite set, called an *alphabet*. Let \mathcal{A}^* be the set of finite sequences of elements of \mathcal{A} , called *words* over \mathcal{A} . A *language* over \mathcal{A} is a subset of \mathcal{A}^* .

A *finite automaton* over \mathcal{A} is a finite directed graph G whose vertices $S(G)$ are called *states* and whose directed edges $T(G)$ are called *transitions*, together with a labelling of each transition by some element of \mathcal{A} , and with a specified subset of the states called the *start states* and another subset called the *stop states*. An *accepted path* in G is a finite, directed path going from a start state to a stop state. Let $L(G)$ be the set of all words $(a_1, \dots, a_n) \in \mathcal{A}^*$ such that, for some accepted path $e_1 * \dots * e_n$ in G , the label of e_i is a_i for all $i = 1, \dots, n$. A *language over G* is a subset of $L(G)$. A language over \mathcal{A} is said to be *regular* if it is the same as the language over some finite automaton defined over \mathcal{A} . In \mathcal{A}^* , the union of any two regular languages is regular, the intersection of any two regular languages is regular, and the complement of any regular language is regular; see [HU79].

Often one requires a finite automaton G to be *deterministic* meaning that for each state s of G and each letter a of the alphabet there is at most one transition leaving s and labelled with a . This property will hold for all of our examples and so it is harmless to assume it, nevertheless it does not alter any of the basic language theory, because of the result that the language accepted by a nondeterministic automaton will also be accepted by some deterministic automaton [HU79].

Suppose now that we fix a finite directed graph G which is *self labelled*, meaning that the alphabet \mathcal{A} is the set of transitions of G and every transition is labelled with itself. We assume that every state is both a start and a stop state. An *automaton over G* consists of a finite automaton G' over the set of transitions of G , with the property that there exists a directed map $f: G' \rightarrow G$ taking states to states and transitions to transitions, such that each transition e' of G' is labelled by the transition $f(e')$ of G ; the map f is uniquely determined. The language $L(G')$ can be described as the sublanguage of $L(G)$ represented by directed paths in G that have a continuous lift which is an accepted path in G' . It follows that a sublanguage of $L(G)$ is regular if and only if it is the full language $L(G')$ of some automaton G' over G .

Here are some important examples.

Loops. Given a finite automaton, the *language of loops in G* , denoted $\Lambda(G)$, is the sublanguage of $L(G)$ labelled by accepted paths which start and stop at the same state. The language $\Lambda(G)$ is regular, because we can construct an automaton $G' = G \times S(G)$ over G , whose start and stop states consist of the diagonal in $S(G') = S(G) \times S(G)$, and clearly $L(G') \cap L(G)$ is the language of loops in G .

Killing automata. Suppose that G' is an automaton over G . We interpret G' as a *killing automaton* over G by considering the *sublanguage of G killed by G'* , denoted $L(G, G')$, to be the set of all words labelling directed paths p in G such that p does *not* lift to an accepted path in G' . This language is a regular language because it equals $(\mathcal{A}^* - L(G')) \cap L(G)$.

Loops and killing automata. Combining the previous two examples, given an automaton G' over G define the *language of loops killed by G'* , denoted $\Lambda(G; G')$, to be the language of all loops Λ in G such that no iterate Λ^n lifts to G' ; here Λ^n is the loop obtained by going n times around Λ . The language $\Lambda(G; G')$ is regular, as we now show. First, letting m be the maximum number of vertices of G' over a vertex of G , it follows that a loop Λ of G is in the language $\Lambda(G; G')$ if and only if none of the loops $\Lambda, \Lambda^2, \dots, \Lambda^m$ lift to G' . By using complements and intersections it therefore suffices to prove for each fixed integer $k \geq 1$ that the language of all loops Λ in G such that Λ^k lifts to G' is a regular language. Define an automaton G'^k over G as follows. For each state s of G , the states of G'^k over s consist of all ordered k -tuples of states of G' over s . Given two states $(s_1, \dots, s_k), (s'_1, \dots, s'_k)$ of G'^k , there is a transition $(s_1, \dots, s_k) \rightarrow (s'_1, \dots, s'_k)$ over an edge e of G if and only if there are transitions $s_1 \rightarrow s'_1, \dots, s_k \rightarrow s'_k$ each over e . It follows that Λ^k lifts to G' if and only if Λ lifts to a loop of G'^k , and so the language of loops Λ whose k th iterate lifts to G' coincides with the regular language $\Lambda(G) \cap \Lambda(G'^k)$.

We are also interested in infinitary languages over a finite alphabet \mathcal{A} , which are sublanguages of the set \mathcal{A}^∞ of all infinite sequences (a_1, a_2, \dots) of elements in \mathcal{A} . Consider a finite automaton G over \mathcal{A} ; we ignore start and stop states. There is an infinitary language denoted $L^\infty(G)$ consisting of all elements of \mathcal{A}^∞ which label infinite directed paths in G . Note that there is a natural bijection between periodic elements of $L^\infty(G)$ with a specified period and the language of loops $\Lambda(G)$: an element (a_1, a_2, \dots) of period n corresponds to a loop $(a_1, \dots, a_n) \in L(G)$.

Given a finite automaton $f: G' \rightarrow G$, a path $e_1 * e_2 * \dots$ satisfies the *iterated killing criterion* associated to G' if for each $m \geq 0$ there exists $n \geq m$ such that the subpath $e_m * \dots * e_n$ does not lift to G' . The language of all words over G labelling paths that satisfy the iterated killing criterion associated to G' is denoted $L^\infty(G; G')$. Note that the period bijection described above restricts to a bijection between the periodic elements of $L^\infty(G; G')$ with specified period and elements of $\Lambda(G; G')$.

The graph of train tracks. Define the *graph of train tracks* on S , denoted $\hat{\Gamma} = \hat{\Gamma}(S)$ to be the directed graph whose vertices are isotopy classes of filling train tracks on S , with a directed edge between (the isotopy classes of) train tracks τ, τ'

if and only if $\tau \succ \tau'$ is a splitting of parity L or R. A splitting sequence on S , based at τ_0 , with no C-splittings, and with all train tracks filling, is up to isotopy the same thing as an infinite directed path in $\hat{\Gamma}$ starting at (the isotopy class of) τ_0 .

Define the *canonical killing graph* over $\hat{\Gamma}$ to be the directed graph $\hat{\Gamma}^\#$ equipped with a directed map $\hat{\Gamma}^\# \rightarrow \hat{\Gamma}$ defined as follows. Given a vertex τ of $\hat{\Gamma}$, the vertices of $\hat{\Gamma}^\#$ lying over τ consist of all ordered pairs (τ, x) where x is either a proper, nonempty subgraph of τ or a C-splitting arc of τ . The directed edges of $\hat{\Gamma}^\#$ lying over a given directed edge $\tau \rightarrow \tau'$ of $\hat{\Gamma}$ consist of all edges of the form $(\tau, x) \rightarrow (\tau', x')$ where x survives the splitting $\tau \succ \tau'$ with descendant x' . The directed map $\hat{\Gamma}^\#$ can therefore be thought of as the forgetful map, which forgets x .

In this language, a splitting sequence on S corresponds to a directed path $\tau_0 \succ \tau_1 \succ \dots$ in $\hat{\Gamma}$, and the sequence satisfies the canonical killing criterion if each vertex τ_i on the path is the initial vertex of a finite subpath $\tau_i \succ \dots \succ \tau_j$ that does not lift to $\hat{\Gamma}^\#$.

The mapping class group acts on $\hat{\Gamma}$ and $\hat{\Gamma}^\#$ respecting the map $\hat{\Gamma}^\# \rightarrow \hat{\Gamma}$. We would like to pass to quotient graphs under this action, but since the action of \mathcal{MCG} on $\hat{\Gamma}$ can have nontrivial stabilizers in \mathcal{MCG} , the quotient map $\hat{\Gamma} \rightarrow \hat{\Gamma}/\mathcal{MCG}$ need not satisfy unique path lifting. The stabilizers are always finite subgroups, however, and if we are willing to augment a train track with a cusp then the stabilizer becomes trivial, by Corollary 3.15.3. In order to get free actions we therefore proceed as follows.

Define a directed graph $\tilde{\Gamma}$ whose vertices are the isotopy types of filling, one cusp train tracks (τ, v) . For every splitting $\tau \succ \tau'$ of parity L or R between filling train tracks, and for every cusp $v \in \text{cusps}(\tau)$, there is a unique $v' \in \text{cusps}(\tau')$ that corresponds to v under the bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$ induced by the splitting, and we define a *splitting edge* in $\tilde{\Gamma}$ from the isotopy type of (τ, v) to the isotopy type of (τ', v') . Also, for every filling train track and any two cusps $v \neq v' \in \text{cusps}(\tau)$, we define a directed *change of cusp edge* in $\tilde{\Gamma}$ from (τ, v) to (τ, v') . The mapping class group evidently acts on $\tilde{\Gamma}$, and by Corollary 3.15.3 the action is free, with quotient graph Γ and a covering map $\tilde{\Gamma} \rightarrow \Gamma$ whose deck transformation group is \mathcal{MCG} . Note that the vertices of Γ are precisely the combinatorial types of filling, one cusp train tracks. The directed edges of Γ can be thought of as the combinatorial types of splitting edges and of change of cusp edges. In particular, it follows that Γ is a finite directed graph. We will use square brackets $[\cdot]$ to denote combinatorial types, so the vertex of Γ corresponding to a vertex (τ, v) of $\tilde{\Gamma}$ is denoted $[\tau, v]$.

Every directed loop in Γ determines a conjugacy class in \mathcal{MCG} . This is just a special case of the general fact that if a group G acts freely on a space \tilde{X} with quotient $X = \tilde{X}/G$, then every continuous loop $p: [0, 1] \rightarrow X$, $p(0) = p(1)$, determines a well-defined conjugacy class in G , as follows. Choose a lift $\tilde{p}(0)$ of the base point $p(0)$, extend to a continuous lift $\tilde{p}: [0, 1] \rightarrow \tilde{X}$, and since $\tilde{p}(0)$ and $\tilde{p}(1)$ are in

the same orbit under the action of G there exists $g \in G$ such that $g \cdot \tilde{p}(0) = \tilde{p}(1)$. The conjugacy class of g is well-defined because for any other choice of a lift $\tilde{p}'(0)$ determining a lift $\tilde{p}': [0, 1] \rightarrow \tilde{X}$, when we take $g' \in G$ so that $g' \cdot \tilde{p}'(0) = \tilde{p}'(1)$, and when we take $h \in G$ so that $h \cdot \tilde{p}(0) = \tilde{p}'(0)$, then we clearly have $h \cdot \tilde{p}(1) = \tilde{p}'(1)$ and so $h \cdot g \cdot \tilde{p}(0) = g' \cdot h \cdot \tilde{p}(0)$ implying that $hg = g'h$.

Define a *splitting circuit* to be a directed loop in Γ .

Proposition 9.2.1. *Given a splitting circuit $[\tau_0, v_0] \rightarrow [\tau_1, v_1] \rightarrow \dots \rightarrow [\tau_n, v_n]$, if we choose a lift (τ_0, v_0) , and if $\Phi \in \mathcal{MCG}$ is the mapping class corresponding to this lift, then τ_0 is an invariant train track for Φ . If we choose another lift (τ'_0, v'_0) , and if $\Phi' \in \mathcal{MCG}$ corresponds to this lift, then τ_0 is a canonical invariant train track for Φ if and only if τ'_0 is a canonical invariant train track for Φ' .*

Proof. The path in $\tilde{\Gamma}$ that lifts the given splitting circuit is a path of the form $(\tau_0, v_0) \rightarrow (\tau_1, v_1) \rightarrow \dots \rightarrow (\tau_n, v_n) = \Phi(\tau_0, v_0)$ where for each $i = 1, \dots, n$ either $\tau_{i-1} \succ \tau_i$ is a splitting or $\tau_{i-1} = \tau_i$ and the move $(\tau_{i-1}, v_{i-1}) \rightarrow (\tau_i, v_i)$ is a change of cusp. In other words we have a carrying sequence $\tau_0 \succcurlyeq \tau_1 \succcurlyeq \dots \tau_n = \Phi(\tau_0)$ where each carrying $\tau_{i-1} \succcurlyeq \tau_i$ is either a splitting or an isotopy, implying that τ_0 is an invariant train track for Φ .

Suppose that τ_0 is a canonical invariant train track for Φ , and so Φ is a pseudo-Anosov mapping class with unstable foliation \mathcal{F}^u satisfying $\mathcal{PMF}(\mathcal{F}^u) = \bigcap_{i=0}^{\infty} \Phi^i(\mathcal{PMF}(\tau_0))$, and with τ_0 canonically carrying \mathcal{F}^u . Choosing (τ'_0, v'_0) as in the statement of the proposition, there is a mapping class Ψ such that $\Psi(\tau_i, v_i) = (\tau'_i, v'_i)$ for all i , and such that Ψ conjugates Φ to Φ' . It follows that Φ' is a pseudo-Anosov mapping class with unstable foliation $\mathcal{F}'^u = \Psi(\mathcal{F}^u)$, and moreover $\mathcal{PMF}(\mathcal{F}'^u) = \Psi(\mathcal{PMF}(\mathcal{F}^u)) = \bigcap_{i=0}^{\infty} \Psi(\Phi^i(\mathcal{PMF}(\tau_0))) = \bigcap_{i=0}^{\infty} \Phi'^i(\mathcal{PMF}(\tau'_0))$ and τ'_0 canonically carries \mathcal{F}'^u . That is, τ'_0 is a canonical invariant train track for Φ' . \diamond

As a consequence of this proposition, we define a splitting circuit $[\tau_0, v_0] \rightarrow \dots \rightarrow [\tau_n, v_n]$ to be *canonical* if it satisfies the conditions of the proposition with respect to some, and hence any choice of a lift (τ_0, v_0) of $[\tau_0, v_0]$.

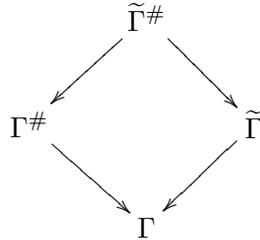
Here is our main result:

Theorem 9.2.2. *The language of canonical splitting circuits is a regular sublanguage of the language $\Lambda(\Gamma)$ of all loops in Γ .*

Proof. We prove this by constructing a killing automaton $\Gamma^\# \rightarrow \Gamma$ over Γ so that the language of canonical splitting circuits is the same as the language $\Lambda(\Gamma; \Gamma^\#)$ of loops in Γ no power of which lifts to $\Gamma^\#$. With our earlier discussion of the directed graphs $\hat{\Gamma}^\#$ and $\hat{\Gamma}$, it is clear how to proceed, and we can describe the proof quickly in terms of pushout diagrams.

There is a forgetful map $\tilde{\Gamma} \rightarrow \hat{\Gamma}$ which takes each vertex (τ, v) to the vertex τ , takes each splitting edge $(\tau, v) \succ (\tau', v')$ to the edge $\tau \succ \tau'$, and takes each change of cusp edge $(\tau, v) \rightarrow (\tau, v')$ to the vertex τ .

Construct a pushout diagram



To be precise, given a vertex τ of Γ , and vertices (τ, x) of $\Gamma^\#$ and (τ, v) of $\tilde{\Gamma}$ over τ , there is a vertex (τ, x, v) of $\tilde{\Gamma}^\#$ mapping to (τ, x) in $\Gamma^\#$ and to (τ, v) in $\tilde{\Gamma}$. For any vertex (τ, x) of $\Gamma^\#$ mapping to a vertex τ of Γ , and for any cusp change edge $(\tau, v) \rightarrow (\tau, v')$ in $\tilde{\Gamma}$ mapping to τ , there is an edge $(\tau, x, v) \rightarrow (\tau, x, v')$ in $\tilde{\Gamma}^\#$ mapping to the vertex (τ, x) in $\Gamma^\#$ and mapping to the edge $(\tau, v) \rightarrow (\tau, v')$ in $\tilde{\Gamma}$. For any edge $\tau \succ \tau'$ in Γ , and edges $(\tau, x) \succ (\tau', x')$ of $\Gamma^\#$ and $(\tau, v) \succ (\tau', v')$ of $\tilde{\Gamma}$ over $\tau \succ \tau'$, define an edge $(\tau, x, v) \succ (\tau', x', v')$ of $\tilde{\Gamma}^\#$ over the edges $(\tau, x) \succ (\tau', x')$ of $\Gamma^\#$ and $(\tau, v) \succ (\tau', v')$ of $\tilde{\Gamma}$. Thus, the map $\tilde{\Gamma}^\# \rightarrow \Gamma^\#$ forgets the cusp, and the map $\tilde{\Gamma}^\# \rightarrow \tilde{\Gamma}$ forgets the subtrack or C-splitting arc x .

The free action of \mathcal{MCG} on $\tilde{\Gamma}$ lifts to a free action on $\tilde{\Gamma}^\#$. The quotient graph $\Gamma^\#$ is well-defined, and the quotient map $\tilde{\Gamma}^\# \rightarrow \Gamma^\#$ is a covering map with deck transformation group \mathcal{MCG} . The directed map $\tilde{\Gamma}^\# \rightarrow \tilde{\Gamma}$ is equivariant with respect to the action of \mathcal{MCG} , and descends to a directed map $\Gamma^\# \rightarrow \Gamma$.

It remains to check that a splitting circuit is canonical if and only if it is in the language $\Lambda(\Gamma; \Gamma^\#) \dots$

◇

The killing graph.

9.3 Penner’s recipe

TO DO:

- Make it clearer the difference between full and fractional recipes.
- Give example, say, 5 punctured sphere, of full and fractional recipes.
- Plug this example into the “regular language” point of view.

- Explain the subtlety of Penner’s conjecture: the twist systems have arbitrarily high intersection number, so the train tracks have arbitrarily large number of bigons, so they do *not* obviously fall under our ordinary notion of train track expansions.

In his original description of the theory of pseudo-Anosov mapping classes [Thu88] Thurston gave a construction of pseudo-Anosov homeomorphisms which among other examples included the following: if c, d are two curves on S which intersect transversely and efficiently and fill the surface, then any word in the positive Dehn twist τ_c around c and the negative Dehn twist τ_d^{-1} around d is a pseudo-Anosov mapping class, as long as each of τ_c and τ_d^{-1} occurs at least once. Penner used train tracks to generalize this construction, replacing the single curves c, d by curve systems \mathcal{C}, \mathcal{D} . We describe Penner’s recipe here, together with a new and more general recipe that allows fractional Dehn twists. Behind this recipe is a certain subclass of train track expansions factored into “fractional Dehn twists”, and a special version of the canonical killing criterion that holds for such expansions.

Define a *twist system* of S to be a pair \mathcal{C}, \mathcal{D} , each a pairwise disjoint, pairwise nonisotopic family of essential simple closed curves, such that \mathcal{C} and \mathcal{D} intersect transversely and efficiently, and each component of $\mathcal{C}(S - (\mathcal{C} \cup \mathcal{D}))$ is a polygon with at most one puncture. It follows that each component has an even number of sides, either ≥ 4 sides and no puncture or ≥ 2 sides and one puncture. A *curve of \mathcal{C}, \mathcal{D}* is a curve of \mathcal{C} or a curve of \mathcal{D} . An *arc of \mathcal{C}, \mathcal{D}* is a component of $\mathcal{C} - \mathcal{D}$ or a component of $\mathcal{D} - \mathcal{C}$. The *singularity type* of \mathcal{C}, \mathcal{D} is the pair of sequences $(i_3, i_4, \dots; p_1, p_2, \dots)$ such that among the components of $\mathcal{C}(S - (\mathcal{C} \cup \mathcal{D}))$ there are i_k components which are nonpunctured $2k$ -gons for each $k \geq 3$, and there are p_k components which are once punctured $2k$ -gons for each $k \geq 1$; there may also be some number of nonpunctured 4-gons, which do not affect the singularity type.

Theorem 9.3.1 ([Pen88]; see also [Fat92]). *Let \mathcal{C}, \mathcal{D} be a twist system on S . Let $\tau_1, \tau_2, \dots, \tau_n$ be a sequence of homeomorphisms of S such that each τ_i is either a positive Dehn twist around a curve of \mathcal{C} or a negative Dehn twist around a curve of \mathcal{D} . Suppose furthermore that each curve of \mathcal{C}, \mathcal{D} occurs at least once as a twist curve in the sequence τ_1, \dots, τ_n . Then the homeomorphism $\phi = \tau_1 \circ \dots \circ \tau_n$ is isotopic to a pseudo-Anosov homeomorphism.*

For an additional conclusion to this theorem, namely that the singularity type of ϕ matches the singularity type of the twist system \mathcal{C}, \mathcal{D} , see Theorem 9.3.2. For a partial converse, see Proposition 9.3.3.

With some additional effort to formulate, but with no more effort to prove, we can generalize Penner’s recipe to “fractional Dehn twists”, as follows.

Consider a twist system \mathcal{C}, \mathcal{D} on S . Let $\mathcal{A}(\mathcal{C}, \mathcal{D})$ denote the set of arcs of \mathcal{C}, \mathcal{D} . Given a curve c of \mathcal{C}, \mathcal{D} , let $n(c)$ be the number of transverse intersection points of

c with other curves of \mathcal{C}, \mathcal{D} . For each rational number r with denominator $n(c)$ we shall define a *fractional Dehn twist* of exponent r around the curve c . This is not a mapping class, but just a relation between two twist systems. When r is an integer, the two twist systems are \mathcal{C}, \mathcal{D} and $\tau_c^r(\mathcal{C}), \tau_c^r(\mathcal{D})$ where τ_c is the positive Dehn twist around c .

Here is the intuition for the fractional Dehn twist of exponent $1/n(c)$ around c . Let's say that $c \in \mathcal{C}$. Consider a point $x \in c \cap \mathcal{D}$, and a short segment $I \subset \mathcal{D}$ containing x , with endpoints a, b . Under a full Left Dehn twist, the segment I is replaced by a segment which starts from a , moves along I until it approaches c , then turns Left and moves past $n(c)$ successive points of $c \cap \mathcal{D}$, returning to I once again, and then it continues along I until the opposite endpoint b . The same is done for each of the $n(c)$ points of $c \cap \mathcal{D}$. To define the fractional Dehn twist of exponent $1/n(c)$ along c , let the points of $c \cap \mathcal{D}$ be listed in cyclic order as $x_1, \dots, x_{n(c)}$, and choose a short segment $I_i \subset \mathcal{D}$ containing x_i , with endpoints a_i, b_i , so that a_1, \dots, a_n are on one side of c and b_1, \dots, b_n are on the other side. The segment I_i is replaced by a segment which starts at a_i , moves along I_i until it approaches c , then turns Left and moves along c to the next segment I_{i+1} (where the index is an element of $\mathbf{Z}/n(c)$), and then continues along I_{i+1} to b_{i+1} .

Here is a more formal definition of fractional Dehn twists along c . Let $R_n \subset S^1$ be the set of n^{th} roots of unity. Let A be an annulus neighborhood of c with a parameterization $A \approx S^1 \times [-1, 1]$ such that $c \approx S^1 \times 0$ and $(\mathcal{C} \cup \mathcal{D}) \cap A \approx (S^1 \times 0) \cup (R_{n(c)} \times [-1, 1])$. Let $h: [0, 1] \rightarrow [0, 1]$ be a smooth, monotonically decreasing map such that $h(x) = 1$ for x in a neighborhood of 0 and $h(x) = 0$ for x in a neighborhood of 1. Given $k \in \mathbf{Z}$ let $\phi_k: S \rightarrow S$ be the map which is the identity on $S - A$ and whose restriction to A has the form

$$\phi_k(z, t) = \begin{cases} (z, t) & \text{if } t \in [-1, 0) \\ \left(z \cdot \exp\left(\frac{2\pi i k h(t)}{n(c)}\right), t \right) & \text{if } t \in [0, 1] \end{cases}$$

Notice that $\phi_k: S \rightarrow S$ is discontinuous along c if k is not a multiple of n . Nevertheless, the pair $\mathcal{C}' = \phi_k(\mathcal{C}), \mathcal{D}' = \phi_k(\mathcal{D})$ is also a twist system on S , which we say is obtained from \mathcal{C}, \mathcal{D} by a *fractional Dehn twist, along the twist curve c , of exponent $\frac{k}{n(c)}$* . Note that the singularity type of $\mathcal{C}', \mathcal{D}'$ is identical to that of \mathcal{C}, \mathcal{D} . The elements of $\mathcal{A}(\mathcal{C}, \mathcal{D})$ contained in the twist curve c are said to be *killed* by the fractional Dehn twist, and all other elements of $\mathcal{A}(\mathcal{C}, \mathcal{D})$ *survive*. Note that the map ϕ_k induces a bijection between the elements of $\mathcal{A}(\mathcal{C}, \mathcal{D})$ that survive and a subset of $\mathcal{A}(\mathcal{C}', \mathcal{D}')$; if $b \in \mathcal{A}(\mathcal{C}, \mathcal{D})$ survives then $\phi_k(b) \in \mathcal{A}(\mathcal{C}', \mathcal{D}')$ is called the *descendant* of b in $\mathcal{A}(\mathcal{C}', \mathcal{D}')$.

Define a *positive word* of fractional Dehn twists to be a sequence of twist systems $(\mathcal{C}_1, \mathcal{D}_1), \dots, (\mathcal{C}_n, \mathcal{D}_n)$ such that for each $i = 1, \dots, n$ the system $(\mathcal{C}_{i+1}, \mathcal{D}_{i+1})$ is

obtained from $(\mathcal{C}_i, \mathcal{D}_i)$ by a fractional Dehn twist of positive exponent along some curve of \mathcal{C}_i or of negative exponent along some curve of \mathcal{D}_i .

Theorem 9.3.2. *Let $(\mathcal{C}_1, \mathcal{D}_1), \dots, (\mathcal{C}_n, \mathcal{D}_n)$ be a positive word of fractional Dehn twists, and let $\phi: S \rightarrow S$ be a homeomorphism such that $\phi(\mathcal{C}_1, \mathcal{D}_1) = (\mathcal{C}_n, \mathcal{D}_n)$. Consider the disjoint union*

$$\mathcal{A} = \mathcal{A}(\mathcal{C}_1, \mathcal{D}_1) \sqcup \dots \sqcup \mathcal{A}(\mathcal{C}_n, \mathcal{D}_n)$$

and let $d: \mathcal{A} \rightarrow \mathcal{A}$ be the partial bijection defined as follows: if $i = 1, \dots, n-1$ and if $b \in \mathcal{A}(\mathcal{C}_i, \mathcal{D}_i)$ survives the next fractional Dehn twist then $d(b) \in \mathcal{A}(\mathcal{C}_{i+1}, \mathcal{D}_{i+1})$ is the descendant of b ; and if $b \in \mathcal{A}(\mathcal{C}_n, \mathcal{D}_n)$ then $d(b) = \phi^{-1}(b) \in \mathcal{A}(\mathcal{C}_1, \mathcal{D}_1)$. Then the following is a sufficient condition for ϕ to be isotopic to a pseudo-Anosov homeomorphism whose singularity type matches that of $\mathcal{C}_1, \mathcal{D}_1$:

Arc Killing Criterion For each $b \in \mathcal{A}(\mathcal{C}_1, \mathcal{D}_1)$ there exists $k \geq 0$ and $i = 1, \dots, n-1$ such that $d^k(b) \in \mathcal{A}(\mathcal{C}_i, \mathcal{D}_i)$ is defined but $d^k(b)$ is killed by the next fractional Dehn twist and so $d^{k+1}(b)$ is not defined.

For a partial converse, see Proposition 9.3.3.

Remarks. The additional information about singularity type is implicit in [Pen88] and [Fat92].

In the context of Penner's recipe, namely full Dehn twists on the curves of a twist system \mathcal{C}, \mathcal{D} , the Arc Killing Criterion is clearly equivalent to the statement that every curve of \mathcal{C}, \mathcal{D} occurs at least once as a twist curve: if one arc is not killed, then none of the arcs in the same twist curve are killed, and that twist curve is not used.

In the context of Theorem 9.3.2, if all fractional Dehn twists are true Dehn twists, and if ϕ is the product of these Dehn twists, then the Arc Killing Criterion reduces to Penner's recipe. A slightly more general case occurs when the fractional Dehn twists are all true Dehn twists, but ϕ is not their product; instead ϕ is the composition of their product with a finite order diffeomorphism preserving $(\mathcal{C}_n, \mathcal{D}_n)$. This case also reduces to Penner's recipe: there is some power ϕ^n which does occur in Penner's recipe, and moreover the Arc Killing Criterion for ϕ implies the hypothesis of Penner's recipe for ϕ^n , which implies that ϕ^n is pseudo-Anosov, which implies in turn that ϕ is pseudo-Anosov.

Problem. It would be interesting to show that by allowing fractional Dehn twists which are not true Dehn twists one obtains pseudo-Anosov mapping classes ϕ such that neither ϕ nor any power ϕ^n occurs in Penner's original recipe (note that if

ϕ is produced by Penner's recipe then ϕ fixes each singularity and each separatix, but the same property can always be achieved by passing to a power). A proof of this might follow by considering the function of λ which, for each recipe, counts the number of conjugacy classes in that recipe whose expansion factor is at most λ . It would also be interesting to show by similar means that there exist pseudo-Anosov mapping classes which do not occur in the fractional Dehn twist version of Penner's recipe.

Proof. The bulk of the proof is to show sufficiency of the Arc Killing Criterion. We do this by showing that the appropriate splitting sequence satisfies the iterated canonical killing criterion, which then allows us to apply Theorems 9.1.1 and 9.1.2.

First we reduce to the case that each 4-gon component of $S - (\mathcal{C}_i \cup \mathcal{D}_i)$ has a puncture. If not, let $Q \subset S$ consist of one point in each nonpunctured 4-gon component of $S - (\mathcal{C}_1 \cup \mathcal{D}_1)$. Choosing the annulus defining each fractional Dehn twist to be disjoint from Q , it follows by induction that for each $i = 1, \dots, n$ the set Q consists of one point in each nonpunctured 4-gon component of $S - (\mathcal{C}_i \cup \mathcal{D}_i)$. The map ϕ takes nonpunctured 4-gon components of $S - (\mathcal{C}_1 \cup \mathcal{D}_1)$ to those of $S - (\mathcal{C}_n \cup \mathcal{D}_n)$, and so by isotoping ϕ relative to $\mathcal{C}_1 \cup \mathcal{D}_1$ we may assume $\phi(Q) = Q$. Let $S' = S - Q$. The hypotheses of the theorem stated for ϕ and for $\phi' = \phi|_{S'}$ are equivalent. If we prove the theorem for ϕ' , then ϕ' is pseudo-Anosov on S' , with singularity type identical to that of the twist system $\mathcal{C}_1 \cup \mathcal{D}_1$ on S' . This implies that ϕ is pseudo-Anosov on S , with singularity type identical to that of ϕ' except that there are $|Q|$ fewer two-pronged puncture singularities, which gives a singularity type identical to that of the twist system $\mathcal{C}_1 \cup \mathcal{D}_1$ on S . This completes the reduction.

Following [Pen88], to each twist system $(\mathcal{C}, \mathcal{D})$ we associate a train track $\tau(\mathcal{C}, \mathcal{D})$ as described in Figure 29. The fact that $\tau(\mathcal{C}, \mathcal{D})$ is indeed a train track follows from our assumption that there are no nonpunctured 4-gons (otherwise $\tau(\mathcal{C}, \mathcal{D})$ would only be a bigon track). The points of $\mathcal{C} \cap \mathcal{D}$ correspond bijectively with the sink branches of $\tau(\mathcal{C}, \mathcal{D})$; we visualize the sink branches as being very short, tantamount to nongeneric switches, though for compatibility with earlier results it is useful to leave them in their uncombed form, as sink branches. The arcs of the twist system \mathcal{C}, \mathcal{D} correspond bijectively with source branches of $\tau(\mathcal{C}, \mathcal{D})$; henceforth we will make use of this bijection without mention.

Any train path in $\tau(\mathcal{C}, \mathcal{D})$ with endpoints on the switches can be expressed as an alternating concatenation of source branches of and sink branches; visualizing the sink branches as being very short, we can therefore express a train path as a concatenation of source branches. Grouping these into \mathcal{C} -segments and \mathcal{D} -segments, it follows that any train path can be expressed uniquely as an alternating concatenation of \mathcal{C} -segments and \mathcal{D} -segments. Consider a sink branch b of $\tau(\mathcal{C}, \mathcal{D})$ (see

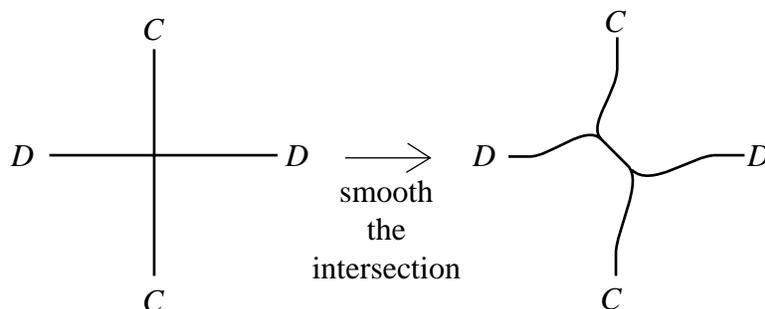


Figure 29: The train track $\tau(\mathcal{C}, \mathcal{D})$ of a twist system $(\mathcal{C}, \mathcal{D})$, obtained by a smoothing near each $x \in \mathcal{C} \cup \mathcal{D}$ to create a sink branch corresponding to x . Working in an oriented chart centered at x in which \mathcal{C} corresponds to the y -axis and \mathcal{D} corresponds to the x -axis, smooth the intersection to produce a nongeneric switch with two branch ends on each side, so that the tangent line along \mathcal{C} is never horizontal, the tangent line along \mathcal{D} is never vertical, and the tangent line at the nongeneric switch has negative slope. Then uncomb each nongeneric switch to form a sink branch.

Figure 29). Wherever the interior of a \mathcal{C} -segment of a train path passes over b , the train path has a Right crossing of b ; wherever the interior of a \mathcal{D} -segment passes over b , the train path has a Left crossing of b ; and wherever a train path switches from a \mathcal{C} -segment to a \mathcal{D} -segment as it passes over b , there is neither a Left nor a Right crossing of b . Two embedded train paths are said to *intersect transversely* at a sink branch b if one has a Left crossing and the other has a Right crossing of b .

Let γ be a curve of \mathcal{C}, \mathcal{D} , and let $\mathcal{C}', \mathcal{D}'$ be obtained by doing a fractional Dehn twist along γ , of positive exponent $\frac{1}{n(\gamma)}$ if $\gamma \in \mathcal{C}$ and negative exponent $-\frac{1}{n(\gamma)}$ if $\gamma \in \mathcal{D}$. There is a homotopic carrying $\tau(\mathcal{C}, \mathcal{D}) \rightsquigarrow \tau(\mathcal{C}', \mathcal{D}')$ which factors into a sequence of splittings

$$\tau(\mathcal{C}, \mathcal{D}) = \tau_0 \succ \cdots \succ \tau_n = \tau(\mathcal{C}', \mathcal{D}')$$

where $n = n(\gamma)$, one splitting along each sink branch in \mathcal{C} . If $\gamma \in \mathcal{C}$ then each of these is a R splitting, whereas if $\gamma \in \mathcal{D}$ then each is a L splitting. See Figure 30.

Now we investigate how a \mathcal{C} -splitting arc or proper subtrack of $\tau(\mathcal{C}, \mathcal{D})$ survives or is killed under the carrying $\tau(\mathcal{C}, \mathcal{D}) \rightsquigarrow \tau(\mathcal{C}', \mathcal{D}')$ induced by a fractional Dehn twist from \mathcal{C}, \mathcal{D} to $\mathcal{C}', \mathcal{D}'$. Let γ be the twist curve, realized as a curve carried on $\tau(\mathcal{C}, \mathcal{D})$. Suppose that the exponent of the twist has absolute value $k/n(\gamma)$; the numerator k plays a significant role in the analysis below. Note that γ contains Left crossings of all the sink branches it passes over if $\gamma \in \mathcal{C}$, whereas γ contains all Right crossings if $\gamma \in \mathcal{D}$.

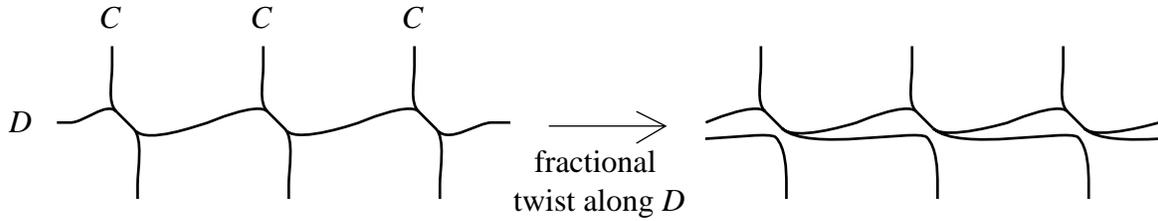


Figure 30: The carrying $\tau(\mathcal{C}, \mathcal{D}) \rightsquigarrow \tau(\mathcal{C}', \mathcal{D}')$ associated to a fractional Dehn twist from \mathcal{C}, \mathcal{D} to $\mathcal{C}', \mathcal{D}'$.

Case 1: A c-splitting arc α of $\tau(\mathcal{C}, \mathcal{D})$. Let $sb(\alpha)$ be the number of source branches of $\tau(\mathcal{C}, \mathcal{D})$ in α , so $sb(\alpha) = 0$ if $\alpha \cap \tau(\mathcal{C}, \mathcal{D})$ is sink branch and $sb(\alpha) > 0$ otherwise.

Case 1a: $\gamma \cap \alpha = \emptyset$. In this case α does not cross any sink branches in γ , and so by Proposition 6.4.1 α survives with descendant $\alpha' \subset \tau(\mathcal{C}', \mathcal{D}')$, $sb(\alpha') = sb(\alpha)$. Also, all arcs of \mathcal{C}, \mathcal{D} contained in α survive, and their descendants are the arcs of $\mathcal{C}', \mathcal{D}'$ contained in α' .

Case 1b: $\gamma \cap \alpha \neq \emptyset$, and $sb(\alpha) = 0$. In this case $\alpha \cap \tau(\mathcal{C}, \mathcal{D})$ is a sink branch contained in γ , and so it contains partial Left and Right crossings of this sink branch, implying that α is killed in $\tau(\mathcal{C}', \mathcal{D}')$, by Proposition 6.4.1.

Case 1c: $\gamma \cap \alpha \neq \emptyset$, and $sb(\alpha) > 0$. If γ has transverse intersection with some \mathcal{C} -segment or \mathcal{D} -segment of α , or if γ contains some terminal sink branch of α without containing the incident source branch of α , then α has a full or partial crossing of a sink branch of γ with the wrong parity, and so α is killed in $\tau(\mathcal{C}', \mathcal{D}')$, by Proposition 6.4.1. Suppose now that these possibilities do not occur.

Each component of $\alpha \cap \gamma$ is a \mathcal{C} or \mathcal{D} -segment. Suppose one of these segment has length ℓ . Each fractional Dehn twist along γ of exponent $\pm 1/n(\gamma)$ reduces the length of the segment by 1, and so after a fractional Dehn twist of exponent $\ell/n(\gamma)$ the length is reduced to zero corresponding to a transverse intersection; by the previous paragraph it now follows that α is killed if $\ell < k$. If all the segments of $\alpha \cap \gamma$ have length $\geq k$, then α survives with descendant $\alpha' \subset \tau(\mathcal{C}', \mathcal{D}')$, each component of $\alpha \cap \gamma$ corresponds to a component of $\alpha' \cap \gamma$ with length reduced by k , the length of each \mathcal{C} or \mathcal{D} -segment of α not contained in γ remains unaltered, and so we have $sb(\alpha') < sb(\alpha)$.

Now we turn to:

Case 2: A proper subtrack σ of $\tau(\mathcal{C}, \mathcal{D})$. Let $n(\sigma)$ be the total number of \mathcal{C} and \mathcal{D} -arcs contained in σ , equivalently the number of source branches of τ contained in σ .

Case 2a: $\gamma \cap \sigma = \emptyset$, or $\gamma \subset \sigma$. In either case σ survives with descendant $\sigma' \subset \tau(\mathcal{C}', \mathcal{D}')$, $n(\sigma') = n(\sigma)$, and each \mathcal{C}, \mathcal{D} arc contained in $\sigma - \gamma$ survives, with descendant a $\mathcal{C}', \mathcal{D}'$ arc of $\sigma' - \gamma$.

Case 2b: $\gamma \cap \sigma$ is a nonempty, proper subset of γ . If some component of $\gamma \cap \sigma$ is a sink branch then σ crosses that sink branch with the wrong parity, and has no crossing with the right parity, and so σ is killed by Proposition 6.1.2. So we may assume that each component δ of $\gamma \cap \sigma$ is a \mathcal{C} or \mathcal{D} segment of length ≥ 1 . Any sink branch b of τ contained in δ and incident to an endpoint of δ is said to be *terminal* in δ , and in this case we define the *valence* $v(b)$ to be the number of ends of source branches of τ in σ incident to b . Note that $v(b) = 2$ or 3 , and $v(b) = 2$ if and only if b is contained in the interior of a branch of σ ; $v(b) = 4$ is ruled out by the fact that b is terminal.

Case 2bi: There exists a component of $\gamma \cap \sigma$ with a terminal sink branch of valence 3. In this case σ contains a crossing of the wrong parity and no crossing of the correct parity, so σ is killed, by Proposition 6.1.2.

Case 2bii: For any component of $\gamma \cap \sigma$, its terminal sink branches each have valence 2. This is like case 1c above: if any component of $\gamma \cap \sigma$ has length $< k$ then σ dies; otherwise, σ survives with descendant $\sigma' \subset \tau(\mathcal{C}', \mathcal{D}')$, each component of $\gamma \cap \sigma$ has length reduced by k , and $n(\sigma') < n(\sigma)$.

Now we use the information on survival just described to prove Theorem 9.3.2. Under the hypotheses of the theorem it follows that

$$\tau(\mathcal{C}_1, \mathcal{D}_1) \succ \cdots \succ \tau(\mathcal{C}_n, \mathcal{D}_n)$$

and that $\phi(\tau(\mathcal{C}_1, \mathcal{D}_1)) = \tau(\mathcal{C}_n, \mathcal{D}_n)$. Now extend this to a ϕ -periodic sequence

$$\tau(\mathcal{C}_1, \mathcal{D}_1) \succ \cdots \succ \tau(\mathcal{C}_n, \mathcal{D}_n) \succ \tau(\mathcal{C}_{n+1}, \mathcal{D}_{n+1}) \succ \cdots \quad (9.1)$$

We may regard this as a splitting sequence by factoring $\tau(\mathcal{C}_i, \mathcal{D}_i) \succ \tau(\mathcal{C}_{i+1}, \mathcal{D}_{i+1})$ into splittings; but we will abuse terminology by suppressing the intermediate train tracks in this factorization.

We refer to the fractional Dehn twist from $\mathcal{C}_i, \mathcal{D}_i$ to $\mathcal{C}_{i+1}, \mathcal{D}_{i+1}$ as the i^{th} fractional Dehn twist. The descent maps of this sequence of fractional Dehn twists forms a sequence of partial bijections

$$\mathcal{A}(\mathcal{C}_1, \mathcal{D}_1) \xrightarrow{d} \dots \xrightarrow{d} \mathcal{A}(\mathcal{C}_n, \mathcal{D}_n) \xrightarrow{d} \mathcal{A}(\mathcal{C}_{n+1}, \mathcal{D}_{n+1}) \xrightarrow{d} \dots$$

Under the hypothesis of the theorem, it follows that for any $n \geq 0$ and any $\alpha \in \mathcal{A}(\mathcal{C}_n, \mathcal{D}_n)$ there exists $k \geq 0$ such that α survives with descendant $d^k(\alpha) \in \mathcal{A}(\mathcal{C}_{n+k}, \mathcal{D}_{n+k})$, but $d^k(\alpha)$ is killed in $\mathcal{A}(\mathcal{C}_{n+k+1}, \mathcal{D}_{n+k+1})$, and so α is killed in $\mathcal{A}(\mathcal{C}_{n+k+1}, \mathcal{D}_{n+k+1})$. Using this fact, we show explicitly that all splitting cycles and proper subtracks of $\tau(\mathcal{C}_1, \mathcal{D}_1)$ are eventually killed, which establishes item (4) in Theorem 9.1.1 and item (5) in Theorem 9.1.2, and by applying those theorems it follows that the sequence (9.1) is a canonical expansion.

Consider first a \mathcal{C} -splitting arc α_1 of $\tau(\mathcal{C}_1, \mathcal{D}_1)$, and as long as it survives let α_i be its descendant in $\tau(\mathcal{C}_i, \mathcal{D}_i)$. From case 1 above it follows that $n(\alpha_1) \geq n(\alpha_2) \geq \dots$, and so $n(\alpha_i)$ is constant when $i \geq I$, for some $I \geq 0$. By Case 1a it follows that for $i \geq I$, the \mathcal{C} and \mathcal{D} arcs of α_i which survive have as their descendants \mathcal{C} and \mathcal{D} arcs of α_{i+1} , and so eventually some arc of α_i is involved in the next fractional Dehn twist. This implies, by cases 1b and 1c, that either α_i is killed or $n(\alpha_{i+1}) < n(\alpha_i)$, but the latter is ruled out for $i \geq I$. Thus, α_1 is eventually killed.

Consider next a proper subtrack σ_1 of $\tau(\mathcal{C}_1, \mathcal{D}_1)$, and as long as it survives let σ_i be its descendant in $\tau(\mathcal{C}_i, \mathcal{D}_i)$. From case 2 above it follows that $n(\sigma_1) \geq n(\sigma_2) \geq \dots$, and so $n(\sigma_i)$ is constant when $i \geq I$, for some $I \geq 0$. Now we use the fact that σ_1 is proper to conclude that σ_I is proper, and so there exists a twist curve γ of $\tau(\mathcal{C}_I, \mathcal{D}_I)$, a component δ of $\gamma \cap \sigma_I$, and a terminal sink branch b of δ , such that $v(b) = 2$ or 3 . Let b'_I be a source branch of $\tau(\mathcal{C}_I, \mathcal{D}_I)$ incident to b but not contained in σ_I , and as long as $b'_I \in \mathcal{A}(\mathcal{C}_I, \mathcal{D}_I)$ survives for $i \geq I$ let $b'_i \in \mathcal{A}(\mathcal{C}_i, \mathcal{D}_i)$ be its descendant. Note by induction that $b'_i \not\subset \sigma_i$ but $b'_i \cap \sigma_i \neq \emptyset$. For some $i \geq I$, the arc b'_i is involved in the i^{th} fractional Dehn twist. By cases 2bi, 2bii, either σ_i dies or $n(\sigma_{i+1}) < n(\sigma_i)$, but the latter is ruled out for $i \geq I$. \diamond

Now we prove a partial converse to Theorem 9.3.2:

Proposition 9.3.3. *In the context of Theorem 9.3.2, suppose that no component of $S - (\mathcal{C}_1 \cup \mathcal{D}_1)$ is a nonpunctured 4-gon. If ϕ is a pseudo-Anosov mapping class whose singularity type equals that of $\mathcal{C}_1, \mathcal{D}_1$, then the Arc Killing Criterion holds.*

Proof. Since no component of $S - (\mathcal{C}_1 \cup \mathcal{D}_1)$ is a nonpunctured 4-gon, the splitting sequence 9.1 in the proof of Theorem 9.3.2 lives on the surface S itself, without the need to remove additional punctures.

Assuming that the Arc Killing Criterion fails, we must prove that the splitting sequence 9.1 fails the Canonical Killing Criterion, and we do this by exhibiting a proper subtrack which survives forever.

Consider a fractional Dehn twist $(\mathcal{C}, \mathcal{D}) \rightarrow (\mathcal{C}', \mathcal{D}')$ along a curve c of $(\mathcal{C}, \mathcal{D})$. Let α be an arc of $(\mathcal{C}, \mathcal{D})$ which is not killed, that is, $\alpha \not\subset c$. Let α' be its descendant in $(\mathcal{C}', \mathcal{D}')$. We may regard α, α' as source branches of $\tau(\mathcal{C}, \mathcal{D}), \tau(\mathcal{C}', \mathcal{D}')$, respectively. Under the fold map $\tau(\mathcal{C}', \mathcal{D}') \rightarrow \tau(\mathcal{C}, \mathcal{D})$, note that the inverse image of the interior of α is an open subarc of α' . This implies that $\sigma = \tau(\mathcal{C}, \mathcal{D}) - \text{int}(\alpha)$ carries $\sigma' = \tau(\mathcal{C}', \mathcal{D}') - \text{int}(\alpha')$, and evidently the carrying map is a homotopy equivalence.

It follows that, in the context of Theorem 9.3.2, if some arc of $\mathcal{C}_1, \mathcal{D}_1$ survives forever, then some proper subtrack of $\tau(\mathcal{C}_1, \mathcal{D}_1)$ survives forever. \diamond

9.4 One sink train track expansions

TO DO:

- Give credit to Kerckhoff for mentioning them in *Simplicial Systems*.

In this section we review [Mos86] which gives a recipe for constructing *all* pseudo-Anosov mapping classes up to a power, using what we call here “one sink train tracks”. To be precise, this recipe produces every pseudo-Anosov mapping class on S which has a fixed separatrix, but there always exists a power with a fixed separatrix, with exponent bounded in terms of the topology of S . The methods of [Mos86] were couched more in the language of “ideal arc systems” than in train track language, but here we will emphasize train tracks.⁷ Theorem 9.4.1 gives a recipe for constructing pseudo-Anosov mapping classes using one sink expansions. Behind this recipe is a concept of a “one sink train track expansions”, and a description of how the canonical killing criterion simplifies in this context. One sink train track expansions also occur, in the context of orientable measured foliations, in [Ker85] (see p. 270).

In Section 10 we show that every pseudo-Anosov mapping class is constructed in terms of one sink train track expansions, up to a power. In particular, every pseudo-Anosov mapping class which fixes some separatrix occurs. We will also review another result of [Mos86] which describes how one sink expansions can be used to virtually classify pseudo-Anosov mapping classes up to conjugacy.

A *one sink train track* is a completely combed, semigeneric, filling train track with one sink branch, together with a choice of transverse orientation on the sink branch. Since a completely combed train track has no transition branches, it follows that every branch except the one sink branch is a source branch.

⁷The usage of ideal arc systems in [Mos86] was the main reason for their continued usage in the unfinished monograph [Mos93].

Recurrence of a one sink train track τ is easy to check. Note that τ is orientable if and only if each source branch has its endpoints at opposite ends of the sink branch, in which case τ is clearly recurrent. If τ is nonorientable then there exists an endpoint s of the sink branch with the property that some source branch of τ has both of its ends located at s , and τ is recurrent if and only if *both* endpoints of the sink branch have this property.

The construction given in Proposition 3.6.1 shows that every arational measured foliation is carried by a one sink train track.

Let τ be a one-sink train track with sink branch b_s . We now describe Left and Right *one sink splittings* $\tau \succ \tau_L$, $\tau \succ \tau_R$ that depend naturally on the choice of transverse orientation on b_s , and which produce new one sink train tracks τ_L, τ_R . Figure 31 shows an example on a nonpunctured surface of genus 2 with two 4-cusped complementary regions; the reader can follow along with this figure in order to understand the construction of the splittings $\tau \succ \tau_L$, $\tau \succ \tau_R$. The naturality of the construction of one sink splittings allows us to associate a natural one sink splitting sequence to each choice of a one sink train track τ and a sequence of parities $d_1, d_2, \dots \in \{L, R\}$.

There is a natural choice of tangential orientation on the transversely oriented sink branch b_s , so that the direct sum of the tangential and transverse orientations equals the given orientation on the surface S . Choosing an oriented local coordinate system in which b_s is horizontal, with the transverse orientation pointing upward and the tangential orientation pointing rightward (as in Figure 31), there is a natural labelling of the ends of b_s , one as the Left endpoint and the other as the Right endpoint. Let $\mathcal{E} = \mathcal{E}(\tau)$ denote the set of ends of source branches of τ , decomposed as $\mathcal{E} = \mathcal{E}_L \cup \mathcal{E}_R$, where $\mathcal{E}_L = \mathcal{E}_L(\tau)$, $\mathcal{E}_R = \mathcal{E}_R(\tau)$ denote the ends of source branches incident to the Left and Right endpoints of b_s , respectively. The transverse orientation on b_s determines naturally a linear ordering on the sets \mathcal{E}_L , \mathcal{E}_R . Let $e_L = e_L(\tau) \in \mathcal{E}_L$ denote the last (topmost) branch end in the set \mathcal{E}_L , and $e_R = e_R(\tau)$ the last branch end in the set \mathcal{E}_R . Let b_L, b_R be the source branches containing the ends e_L, e_R respectively; note that $b_L \neq b_R$, for otherwise some side of some component of $\mathcal{C}(S - \tau)$ has no cusps, contradicting that τ is filling. There exists a splitting arc α determined naturally up to isotopy rel τ by the property that the two cusps into which the ends of α emerge are the unique cusps incident to the branch ends e_L, e_R , respectively. For each $d \in \{L, R\}$ there is a wide splitting of parity d along α , denoted $\tau \succ \tau'_d$.

Notice that τ'_L, τ'_R are each semigeneric with a unique sink branch, but neither is completely combed. In the example indicated in Figure 31, for each $d \in \{L, R\}$ the train track τ'_d has two transition branches, one mapping to a subpath of b_s under the fold map to τ , and the other mapping to a train path that contains $b_{\bar{d}}$. In general it is possible for the latter transition branch not to exist: it exists if

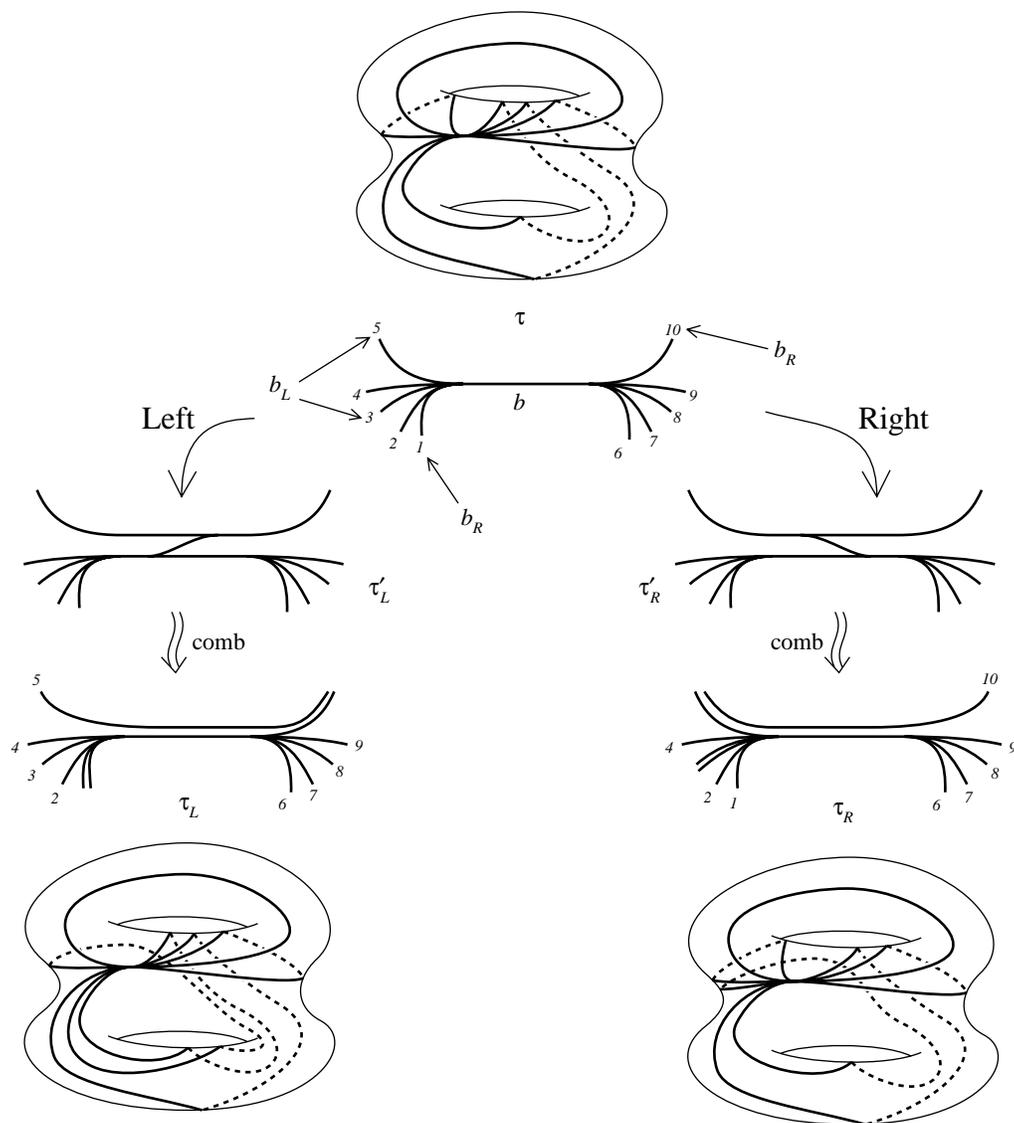


Figure 31: Splitting a one sink train track τ . The sink branch b has the upward transverse orientation, with branch ends on the Left indexed 1–5 and on the right indexed 6–10. b_L is the highest branch on the Left with two ends indexed 3, 5, and b_R is the highest branch on the Right with two ends indexed 1, 10. The splitting arc α separates ends 1–4 from end 5 on the Left and it separates ends 6–9 from end 10 on the right. Under a Left splitting, the source branch b_R has no descendant; descendants of other source branches with ends indexed 2–9 are indicated in τ_L . Under a Right splitting, b_L has no descendant; descendants of other source branches with ends indexed 1, 2, 4, 6–10 are indicated in τ_R .

and only if there are at least three branch ends on the \bar{d} end of b_s , in which case τ'_d has two transition branches as in the example of Figure 31; otherwise, τ'_d has only one transition branch, mapping to a subpath of b_s . Combing the transition branches of τ'_d produces a completely combed train track with one sink branch denoted τ_d . Any carrying map $\tau'_d \rightarrow \tau$ takes the sink branch of τ'_d diffeomorphically to a subinterval of the sink branch of τ , and any two carrying maps are homotopic through carrying maps (Proposition 3.5.2), from which it follows that the transverse orientation on the sink branch of τ pulls back naturally to a transverse orientation on the sink branch of τ'_d . Similarly, the transverse orientation on the sink branch of τ'_d pulls back naturally to a transverse orientation on the sink branch of τ_d . This completes the definition of the one sink train track τ_d and the splitting $\tau \succ \tau_d$. The source branch of τ_d which corresponds under the comb equivalence $\tau'_d \cong \tau_d$ with the postsplitting branch of τ'_d will be called the postsplitting branch of τ_d —in fact, this branch is $b_{\bar{d}}(\tau_d)$.

There is a particularly nice carrying map $f: \tau_d \rightarrow \tau$, the *one sink carrying map*, determined up to homotopy relative to the switches of τ_d by the following properties: f takes the sink branch of τ_d diffeomorphically to the sink branch b_s of τ ; there exists a source branch of τ_d whose image under f is the train path $b_L * b_s * b_R$; and every other source branch of τ_d is taken diffeomorphically to a source branch of τ . Note in particular that the postsplitting branch of τ_d , namely $b_{\bar{d}}(\tau_d)$, is taken diffeomorphically to $b_{\bar{d}}(\tau)$.

Next we define survival and descent of source branches under a one sink splitting. Let $\text{Src}(\tau)$ denote the set of source branches of a one sink train track τ . Given $d \in \{L, R\}$, we say that the source branch $b_{\bar{d}}$ is *killed* by the splitting $\tau \succ \tau_d$, and every other source branch of τ is said to *survive* the splitting. If b survives, then there is a unique source branch $b' \subset \tau_d$ such that $b \subset f(b')$, where f is the one sink carrying map, and we say that b' is the *descendant* of b and b is the *ancestor* of b' (note that for the one source branch $b_{\bar{d}}$ of τ which is killed in τ_d , there are two distinct branches of τ_d whose image under f contains $b_{\bar{d}}$, namely the postsplitting branch and the descendant of $b_{\bar{d}}$). Note that if $b \neq b_{\bar{d}}$ then the descendant of b is mapped diffeomorphically to b by f , whereas the descendant of $b_{\bar{d}}$ is mapped to the train path $b_L * b_s * b_R$. Letting $\delta(b) \in \text{Src}(\tau_d)$ be the descendant of $b \in \text{Src}(\tau)$, we obtain a partial bijection $\delta: \text{Src}(\tau) \rightarrow \text{Src}(\tau_d)$, whose domain is everything in $\text{Src}(\tau)$ except for $b_{\bar{d}}$, and whose image is everything in $\text{Src}(\tau_d)$ except the postsplitting branch.

Theorem 9.4.1 ([Mos86]). *Consider a homeomorphism $\phi: S \rightarrow S$ and a one sink splitting sequence $\tau_0 \succ \tau_1 \succ \cdots \succ \tau_n$ such that $\phi(\tau_0)$ is isotopic to τ_n . Let Src be the disjoint union $\text{Src}(\tau_0) \sqcup \cdots \sqcup \text{Src}(\tau_n)$, and let $\delta: \text{Src} \rightarrow \text{Src}$ be the partial bijection on Src which restricts to the descendant map $\delta: \text{Src}(\tau_{i-1}) \rightarrow \text{Src}(\tau_i)$ for each $i = 1, \dots, n$, and which restricts to the induced map of ϕ^{-1} from $\text{Src}(\tau_n)$*

to $\text{Src}(\tau_0)$. Then the following condition is necessary and sufficient for ϕ to be a pseudo-Anosov homeomorphism whose singularity type equals that of τ_0 :

Periodic source branch killing criterion For each $b \in \text{Src}$ there exists $k \geq 0$ and $i \in \{0, \dots, n - 1\}$ such that $\delta^k(b) \in \text{Src}(\tau_i)$ is defined but $\delta^{k+1}(b)$ is killed by the splitting $\tau_i \succ \tau_{i+1}$ and so $\delta^{k+1}(b)$ is not defined.

Proof. Extend the finite splitting sequence $\tau_0 \succ \tau_1 \succ \dots \succ \tau_n = \phi(\tau_0)$ to a ϕ -periodic, infinite, one sink splitting sequence $\tau_0 \succ \tau_1 \succ \dots$. Let $d_j \in \{L, R\}$ be the parity of $\tau_j \succ \tau_{j+1}$, and let $f_{j+1}: \tau_{j+1} \rightarrow \tau_j$ be a one sink carrying map. Given $I \leq J$ and $b \in \text{Src}(\tau_I)$, we say that b survives in τ_J if there exists a sequence $b = b_I, b_{I+1}, \dots, b_J$ with $b_i \in \text{Src}(\tau_i)$ for $i = I, \dots, J$ such that if $I + 1 \leq i \leq J$ then b_{i-1} survives the one sink splitting $\tau_{i-1} \succ \tau_i$ with descendant τ_i . If $b \in \text{Src}(\tau_I)$ does not survive in τ_J then we say that b is *killed* in τ_J . Clearly the ‘‘Periodic source branch killing criterion’’ stated above is equivalent to the following:

Source branch killing criterion For each i and $b \in \text{Src}(\tau_i)$ there exists $j > i$ such that b is killed in τ_j .

Applying item (4) in Theorem 9.1.1 and item (5) in Theorem 9.1.2, we must prove that the source branch killing criterion is equivalent to the statement that every proper subtrack and every C-splitting arc of τ_0 is eventually killed.

Indeed we shall prove more: discarding the assumption that $\tau_0 \succ \tau_1 \succ \dots$, we shall prove that the source branch killing criterion is equivalent to the canonical killing criterion, which says that for each i every proper subtrack and every C-splitting arc of τ_i is eventually killed.

Suppose that the source branch killing criterion holds: for each m and each $b \in \text{Src}(\tau_m)$ there exists $l \geq m$ such that b is killed in τ_l . Using this fact, we study how ends of source branches evolve under a one sink splitting sequence.

Consider a one sink splitting $\tau \succ \tau'$ of parity $d \in \{L, R\}$, with one sink carrying map $f: \tau' \rightarrow \tau$. The induced map $f: \mathcal{E}(\tau') \rightarrow \mathcal{E}(\tau)$ respects the decomposition $\mathcal{E} = \mathcal{E}_L \cup \mathcal{E}_R$ and respects the linear ordering on each side, so for each $d' \in \{L, R\}$ we have an induced map $f: \mathcal{E}_{d'}(\tau') \rightarrow \mathcal{E}_{d'}(\tau)$ which preserves the linear order. Moreover, order is *strictly* preserved in the sense that if $e_1 < e_2$ then $f(e_1) < f(e_2)$, with a single exception: if $\hat{e}_{\bar{d}}$ denotes the end of $b_{\bar{d}}$ opposite $e_{\bar{d}}$, and if $\hat{e}_{\bar{d}} \in \mathcal{E}_{d'}(\tau)$, then there are exactly two ends $e_1 \neq e_2 \in \mathcal{E}_{d'}(\tau')$ such that $f(e_1) = f(e_2) = \hat{e}_{\bar{d}}$; one of e_1, e_2 is an end of the postsplitting branch $b_{\bar{d}}(\tau')$, and the other is an end of the descendant of $b_{\bar{d}}(\tau)$. Also, the image of f contains every end in $\mathcal{E}(\tau)$ except for the topmost end e_d in $\mathcal{E}_d(\tau)$. To summarize, f is injective with the exception of a single pair of ends that are identified under f , and f is surjective with the exception of a single end that is missed by the image of f .

In Figure 31, for example, the end $\hat{e}_R \in \mathcal{E}_L(\tau)$ is indexed 1 and one can see its two unindexed pre-images in $\mathcal{E}_L(\tau_L)$; the end of $\mathcal{E}(\tau)$ indexed 5 has no preimage in $\mathcal{E}(\tau_L)$; and all other ends are in bijective correspondence under $\tau \succ \tau_L$ where the bijection is indicated with the indices 2–4, 6–10. Also, the end $\hat{e}_L \in \mathcal{E}_L(\tau)$ is indexed 3, and one can see its two unindexed pre-images in $\mathcal{E}_L(\tau_R)$; the end of $\mathcal{E}(\tau)$ indexed 10 has no preimage in $\mathcal{E}(\tau_R)$; and all other ends are in bijective correspondence under $\tau \succ \tau_R$ where the bijection is indicated with indices 1, 2, 4–9.

End evolution is related to survival and descent of source branches as follows. If $b \in \text{Src}(\tau)$ survives the d -splitting $\tau \succ \tau'$ with descendant $b' \in \text{Src}(\tau')$, then we have $b \neq b_{\bar{d}}$, and one of two possibilities occurs. If $b \neq b_d$ then f induces a bijection between the ends of b' and the ends of b . On the other hand, if $b = b_d$ then f maps one end of b' to one end of b_d ; the other end of b' maps to $\hat{e}_{\bar{d}}$, and the end e_d of b_d is not in the image of f .

We claim that in the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, every end eventually evolves to the top of the food chain and suddenly goes extinct, unless it stays on the very bottom for all time. To make sense of this claim, define an *evolution* of the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$ to be a sequence of the form e_0, e_1, e_2, \dots such that $e_j \in \mathcal{E}_d(\tau_j)$ for some fixed $d \in \{L, R\}$, $f_{j+1}(e_{j+1}) = e_j$, and the sequence is maximal in the sense that it is not a finite initial subsequence of any longer such sequence. For example, if we define $e_d^\infty(\tau_j)$ to be the bottom-most element of $\mathcal{E}_d(\tau_j)$ then $e_d^\infty(\tau_0), e_d^\infty(\tau_1), e_d^\infty(\tau_2), \dots$ is an infinite evolution.

The formal statement of the claim says:

End Evolution Claim: There is a unique infinite evolution, and for every finite evolution e_0, e_1, \dots, e_J , the final entry e_J is equal to $e_d(\tau_J)$ (the topmost element of $\mathcal{E}_d(\tau_J)$), and moreover the next splitting $\tau_J \succ \tau_{J+1}$ has parity d so that $e_d(\tau_J)$ has no pre-image under f_{J+1} .

To prove the claim, note that the set of all $e \in \mathcal{E}_d(\tau_j)$ which are in an infinite evolution forms an initial segment of the linear ordering on $\mathcal{E}_d(\tau_j)$, denoted $\mathcal{E}_d^\infty(\tau_j)$, containing at least the first (bottom-most) element $e_d^\infty(\tau_j) \in \mathcal{E}_d(\tau_j)$. Observing that the cardinality of $\mathcal{E}_d^\infty(\tau_j)$ is nondecreasing, it follows that the cardinality is constant for sufficiently large j , say $j \geq J$. Assuming that the stable cardinality is ≥ 2 we shall derive some contradiction. The only element of $\mathcal{E}(\tau_j)$ having more than one pre-image in $\mathcal{E}(\tau_{j+1})$ is $\hat{e}_{\bar{d}_j}$, and for $j \geq J$ it follows that $\hat{e}_{\bar{d}_j}$ is either not in $\mathcal{E}_d^\infty(\tau_j)$ or it is the topmost element of $\mathcal{E}_d^\infty(\tau_j)$ and only the bottom-most of its two preimages is in $\mathcal{E}_d^\infty(\tau_{j+1})$. By the assumption that $\mathcal{E}_d^\infty(\tau_j)$ has at least two elements it follows that its bottom-most element $e_d^\infty(\tau_j)$ is not equal to $\hat{e}_{\bar{d}_j}$. Also, clearly $e_d^\infty(\tau_j)$ is not equal to the topmost element $e_{\bar{d}_j}$ of $\mathcal{E}_{\bar{d}_j}(\tau_j)$. This implies that $e_d^\infty(\tau_j)$ is not an end of the source branch $b_{\bar{d}_j}$, which is the only source branch of τ_j that is killed in τ_{j+1} . Letting b_j^∞ be the source branch of τ_j containing $e_d^\infty(\tau_j)$, we

have shown that b_j^∞ survives the splitting $\tau_j \succ \tau_{j+1}$. Moreover, from the discussion above relating end evolution to descent of source branches, the descendant of b_j^∞ in τ_{j+1} is clearly b_{j+1}^∞ . It follows that b_j^∞ is not eventually killed, contradicting the hypothesis of the theorem.

Using the above claim about evolution of ends, we now verify that each proper subtrack and each C-splitting arc of τ_0 is eventually killed.

First we consider C-splitting arcs. Given a one sink train track τ with sink branch b_s , each C-splitting arc α satisfies $\alpha \cap \tau = b_s$, and α emerges into one cusp at the Left endpoint of b_s and one cusp at the Right endpoint. Given a one sink splitting $\tau \succ \tau'$ of parity $d \in \{L, R\}$ and a splitting arc α of τ , note that α is killed if and only if the the cusp into which α emerges at the d endpoint of b_s is incident to the end $e_d(\tau)$.

Consider now a C-splitting arc α_0 of τ_0 , and suppose that α_0 survives forever, with descendant α_i in τ_i . Let $a_i < b_i \in \mathcal{E}_L(\tau_i)$ be the ends incident to the cusp into which α_0 emerges on the Left endpoint of $b_s(\tau_i)$, and let $c_i < d_i \in \mathcal{E}_R(\tau_i)$ be the ends incident to the cusp into which α_i emerges on the Right endpoint. Since α_i survives the splitting $\tau_i \succ \tau_{i+1}$, it follows from the previous paragraph that α_i is not incident to the end $e_{d_i}(\tau_i)$. From this, together with the definition of descent for C-splitting arcs, it follows that $f_{i+1}(a_{i+1}) = a_i$ and similarly for b_i, c_i, d_i . This implies that b_0, b_1, b_2, \dots and d_0, d_1, d_2, \dots are infinite evolutions of the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$, on the Left and Right respectively, and neither is equal to the bottom-most evolution on the Left or Right. This contradicts the End Evolution Claim, and so α_0 is eventually killed.

Next we consider proper subtracks. Given a one sink splitting $\tau \succ \tau'$ of parity $d \in \{L, R\}$ and a proper subtrack $\sigma \subset \tau$, note that σ is killed if and only $b_d \subset \sigma$ and $b_{\bar{d}} \not\subset \sigma$.

Consider now a proper subtrack σ_0 of τ_0 , and suppose that σ_0 survives forever, with descendant σ_i in τ_i . We must derive a contradiction.

First we reduce to the case that there exists i such that either $b_L(\tau_i)$ or $b_R(\tau_i)$ is in σ_i . If this is not the case, then for all i the one sink carrying map f_{i+1} maps σ_{i+1} diffeomorphically onto σ_i . Pick any $e_0 \in \text{Src}(\tau_0)$ which is contained in σ_0 , and continuing by induction, having picked $e_i \in \text{Src}(\tau_i)$ which is contained in σ_i , then there exists $e_{i+1} \in \text{Src}(\tau_{i+1})$ which is contained in σ_{i+1} such that $f_{i+1}(e_{i+1}) = e_i$. It follows that e_0, e_1, \dots is an infinite evolution in the splitting sequence $\tau_0 \succ \tau_1 \succ \dots$. But we can pick e_0 lying in σ_0 so that e_0 is not the bottom-most end in either $\mathcal{E}_L(\tau_0)$ or $\mathcal{E}_R(\tau_0)$, because otherwise the only ends of τ_0 that lie in σ_0 are the two bottom-most ends, implying that those two ends lie in the same source branch, implying that some component of $\mathcal{C}(S - \tau_0)$ has a smooth closed curve in its boundary, contradicting that τ_0 fills (this is the same argument used before to conclude that that the two topmost branches b_L, b_R are distinct).

Suppose now that either $b_L(\tau_I)$ or $b_R(\tau_I)$ is in σ_I . We show by induction that for all $i \geq I$ either $b_L(\tau_i)$ or $b_R(\tau_i)$ is in σ_i . If there exists $d \in \{L, R\}$ such that $b_d(\tau_i) \subset \sigma_i$ but $b_{\bar{d}}(\tau_i) \not\subset \sigma_i$, since σ_i survives it follows $\tau_i \succ \tau_{i+1}$ has parity \bar{d} , which implies that $b_d(\tau_{i+1}) \subset \sigma_{i+1}$. On the other hand, if both $b_L(\tau_i)$ and $b_R(\tau_i)$ are in σ_i , and if $\tau_i \succ \tau_{i+1}$ has parity d , it follows that $b_{\bar{d}}(\tau_{i+1}) \subset \sigma_i$.

Since σ_0 is proper, there exists $b_0 \subset \tau_0$ such that $b_0 \not\subset \sigma_0$. We shall derive our ultimate contradiction by showing that b_0 survives infinitely. Assume by induction that b_0 survives with descendant $b_i \subset \tau_i$ such that $b_i \not\subset \sigma_i$. If $b_i = b_d(\tau_i)$ for $d \in \{L, R\}$ it follows by the previous paragraph that $\tau_i \succ \tau_{i+1}$ has parity d , and so b_i survives. On the other hand, if $b_i \neq b_d(\tau_i)$ for either $d \in \{L, R\}$ then b_i also survives. In either case b_i survives with descendant $b_{i+1} \in \text{Src}(\tau_{i+1})$. To complete the induction it remains to observe that $b_{i+1} \not\subset \sigma_{i+1}$, for if $b_{i+1} \subset \sigma_{i+1}$ then $b_i \subset f_{i+1}(b_{i+1}) \subset f_{i+1}(\sigma_{i+1}) = \sigma_i$.

This completes the proof of sufficiency of the source branch killing criterion.

The proof of necessity is similar to the proof of Proposition 9.3.3. Assume that some source branch b_0 of τ_0 survives forever, with descendant $b_i \subset \tau_i$. This means that under the carrying map $\tau_{i+1} \rightarrow \tau_i$, the preimage of the interior of b_i is a subarc of b_{i+1} . It follows that the sequence of subtracks $\sigma_i = \tau_i - \text{int}(b_i)$ is a sequence of surviving proper subtracks and their descendants. ◇

9.5 Remarks on the complexity of the Canonical Killing Criterion

Consider a splitting sequence $\tau_0 \succ \tau_1 \succ \dots$. In formulating the Canonical Killing Criterion, one associates to each train track τ a finite set which we shall denote $A(\tau)$, and to each splitting $\tau \succ \tau'$ one associates a “descendant” relation between $A(\tau)$ and $A(\tau')$, which can be iterated to define a descendant relation between $A(\tau)$ and $A(\tau')$ for any τ' homotopically carried by τ . The Canonical Killing Criterion says that for each i and each $a \in A(\tau_i)$ there exists $j \geq i$ such that a does not have a descendant in $A(\tau_j)$.

The complexity of the Canonical Killing Criterion can be measured by the size of the finite set $A(\tau)$. In the general form, $A(\tau)$ consists of all proper subtracks and C-splitting arcs of τ . A little thought shows that the cardinality $|A(\tau)|$ grows exponentially with the topological complexity of the surface. For example, if τ is the train track associated to a twist system \mathcal{C}, \mathcal{D} , then distinct proper subsets of the set of twist curves $\mathcal{C} \cup \mathcal{D}$ yield distinct proper subtracks of τ , and the cardinality of such subtracks is exponential in the genus and number of punctures.

In Sections 9.3 and 9.4 we showed that in certain cases the Canonical Killing Criterion can be reformulated in a particularly efficient manner. In Section 9.3 we showed that if the sequence can be factored in the form $\tau_{n_0} \succcurlyeq \tau_{n_1} \succcurlyeq \dots$ so that each τ_{n_i} is the train track associated to a twist system and each $\tau_{n_i} \succcurlyeq \tau_{n_{i+1}}$ is a Dehn twist

then one can formulate the canonical killing criterion using only a linear amount of data: the set $A(\tau_{n_i})$ can be taken as the set of curves in the twist system associated to τ_{n_i} , and clearly $|A(\tau_{n_i})|$ grows linearly in the genus and number of punctures. For sequences which can be factored as fractional Dehn twists, $A(\tau_{n_i})$ is the set of arcs of the twist system, which still grows linearly. In Section 9.4, where each τ_n is a one sink train track and each $\tau_n \succ \tau_{n+1}$ is a one sink splitting, the set $A(\tau_n)$ can be taken to be the set of source branches of τ_n , which grows linearly in the genus and number of punctures.

This raises the question: for a general splitting sequence, can the Canonical Killing Criterion be reformulated so that $|A(\tau)|$ grows at most linearly in the genus and number of punctures?

As remarked in Section 9.4 without proof, every arational measured foliation has a canonical one sink expansion, and in this case as we have just remarked the Canonical Killing Criterion has linear complexity. We have also seen in Section 7.1 that every arational measured foliation has a canonical one cusp expansion, and we shall show later in Section 11.1 that for one cusp splitting sequences the Canonical Killing Criterion also has linear complexity. Thus, if one is content with either one sink or one cusp expansions, then one gets linear complexity of the Canonical Killing Criterion in a sufficiently universal context to obtain expansions of every arational measured foliation.

TO DO:

- Describe how a killing criterion gives a killing graph $\tilde{\Gamma}^\#(S)$, mapped to $\tilde{\Gamma}(S)$.
- Describe the variations, with particular emphasis on the canonical killing criterion.
- Describe how to use these graphs to interpret when a splitting sequence satisfies a killing criterion.
- Describe the quotient as a finite automaton, and its associated killing automaton.
- Describe the language of canonical expansions as an infinitary regular language (I'm not quite sure what this means; perhaps it is the complement of an infinitary regular language?)

10 Classifying pseudo-Anosov conjugacy classes

TO DO:

- Recast the classification, so that it is completely described in terms of splitting circuits.
- First talk about the virtual centralizer.
- Then a theorem: two group elements are conjugate if and only if their virtual centralizers are conjugate, and when conjugated into the same virtual centralizer, the two elements become conjugate *within the virtual centralizer*. Same could be said for any conjugacy invariant subgroup containing an element, such as the centralizer, or the normalizer, etc.
- Then discussion of the structure of the virtual centralizer of a pseudo-Anosov, its translation subgroup, etc.
- Then splitting circuits.
- Then a conjugacy invariant for the virtual centralizer: for each direction, it has a particular element, namely the lowest power that fixes all separatrices and that moves in that direction; prove that this element is unique in the translation subgroup.
- Classify this element by its splitting circuits.
- Now go on to the circular expansion complex, and explain how it incorporates the previous information. In particular, explain how, given the circular expansion complex for one pseudo-Anosov, one can determine the virtual centralizer and its translation subgroup, and how one can determine the circular expansion complex for the smallest guy in the same direction which fixes all separatrices.

In this section we will give a complete classification of pseudo-Anosov conjugacy classes in \mathcal{MCG} . In Section 10.1 we explain the method for the torus, and in Sections 10.2 and 10.3 we give two classification methods on a general surface, using one cusp expansions in Section 10.2 and using expansion complexes in Section 10.3.

Before explaining the methods we will follow to classify pseudo-Anosov conjugacy classes, first we explain a folk theorem showing how to obtain this classification in a simple minded way using invariant train tracks. Suppose that τ is a canonical invariant train track for a mapping class Φ , choose a representative homeomorphism ϕ for Φ , and choose a carrying map $f: \phi(\tau) \rightarrow \tau$ with the proviso that f

takes switches to switches; this is easily obtained by homotoping an arbitrary carrying map $\phi(\tau) \rightarrow \tau$ through carrying maps. The composition $f \circ \phi: \tau \rightarrow \tau$ associates, to each branch b of τ , a train path $f \circ \phi \upharpoonright b$ which is a concatenation of branches denoted $p(b)$, called a *combinatorial train path* in τ . The pair (τ, p) is an example of a *combinatorial self carrying*, which is a train track and a map p from $\text{Br}(\tau)$ to the set of combinatorial train paths, such that for any two branches $b_1 * b_2$ forming a combinatorial train path, the path $p(b_1) * p(b_2)$ is also a combinatorial train path. Thus we have described a process which associates, to each pseudo-Anosov mapping class Φ , a set of combinatorial self carryings. For a given Φ and τ there are finitely many self-carryings, because there are finitely many choices of the map $f \circ \tau$ up to homotopy rel switches.

A *conjugacy* between two combinatorial self carryings (τ, p) , (τ', p') is a homeomorphism $\psi \in \text{Homeo}_+$ such that $\psi(\tau) = \tau'$ and, for each branch b of τ , we have $\psi(p(b)) = p'(b')$. There is clearly a simple algorithm which will decide whether two combinatorial self carryings are conjugate.

If Φ is a pseudo-Anosov mapping class, then the set of conjugacy classes of combinatorial self carryings associated to Φ is a finite set; let this set be denoted $T(\Phi)$. The finiteness of $T(\Phi)$ is a consequence of the fact that each invariant train track determines finitely many combinatorial self-carryings, and the fact that there are only finitely many Φ orbits of invariant train tracks of Φ ; the latter follows, for example, from the Stable Equivalence Theorem 7.2.3.

If Φ, Ψ are two pseudo-Anosov mapping classes, then Φ, Ψ are conjugate in \mathcal{MCG} if and only there is a bijection $T(\Phi) \leftrightarrow T(\Psi)$ such that any element of $T(\Phi)$ is conjugate to the corresponding element of $T(\Psi)$. Assuming that one can compute $T(\Phi)$ and $T(\Psi)$, the existence of such a bijection can clearly be checked algorithmically, thus giving the desired conjugacy classification of pseudo-Anosov mapping classes.

The difficulty with the approach just outlined is the data size, which is so large as to make this an impractical method. For example, if τ is an invariant train track for Φ , and if the carrying map $\phi(\tau) \rightarrow \tau$ factors into K alternating blocks of L and R parity, then the lengths of the train paths $p(b)$ are exponential in K .

The way around this difficulty is to package the information more efficiently.

As a first attempt, instead of describing the carrying map of a canonical invariant train track τ by using combinatorial self carryings, describe it by factoring the carrying map into splittings $\tau = \tau_0 \succ \tau_1 \succ \cdots \succ \tau_n = \phi(\tau)$. The combinatorial types of the train tracks in this splitting sequence form a splitting circuit, as explained in Section 9. If the mapping class Φ is pseudo-Anosov, then the set of such circuits is finite, and forms a complete conjugacy invariant of Φ . Unfortunately this method still has its inefficiencies: in constructing a splitting sequence from τ to $\phi(\tau)$, there are enough choices in which cusps to split that there may still be

exponentially many different splitting sequences from τ to $\phi(\tau)$.

In Sections 10.2 and 10.3 we offer two solutions to this difficulty, by describing much more efficient data structures for conjugacy classification of pseudo-Anosov mapping classes. In Section 10.2 the solution is to restrict to a very small subclass of splitting circuits, namely the one cusp canonical splitting circuits that arise from one cusp canonical expansions of the unstable foliation. In Section 10.3 the alternate solution is to package all splitting circuits for Φ into a single structure, the “circular expansion complex”, just as we used the expansion complex to package up all canonical expansions of an arational measured foliation (up to stable equivalence). Circular expansion complexes are somewhat unwieldy compared to one cusp splitting circuits, however circular expansion complexes have some theoretical advantages as well as practical advantages. One such advantage is the ability to quickly calculate the virtual normalizer and virtual centralizer of a pseudo-Anosov mapping class. Indeed, these subgroups play an crucial role in the classification using one cusp splitting circuits.

We begin by describing our methods in the case of the torus.

10.1 The torus

In Section 1.3 we described the torus $T = \mathbf{R}^2/\mathbf{Z}^2$, its mapping class group $\mathcal{MCG} \approx \mathrm{SL}(2, \mathbf{Z})$, Teichmüller space $\mathcal{T} \approx \mathbf{H}^2$, projective measured foliation space $\mathcal{PMF} \approx \mathbf{RP}^1$, measured foliations \mathcal{F}_r of constant slope $r \in \mathbf{R}^*$, essential simple closed curves c_r of constant slope $r \in \mathbf{Q}^*$, and filling bigon tracks $\tau_{[p,q]}$, $p = \frac{a}{b}, q = \frac{c}{d} \in \mathbf{Q}^*$ with $ad - bc = \pm 1$.

As mentioned briefly at the end of Section 1.3, if we puncture the torus T at a single base point, say the image of the origin under the universal covering map $\mathbf{R}^2 \rightarrow T$, we obtain the once punctured torus T_1 , and the inclusion map $T_1 \hookrightarrow T$ induces an isomorphism of mapping class groups, as well as an equivariant homeomorphism of compactified Teichmüller spaces. We can therefore transfer all of the above notation from T over to T_1 without any ambiguity, with the slight caveat that we always choose curves c_r and train tracks $\tau_{[p,q]}$ to be disjoint from the base point so that they are included in T_1 . The reason for using T_1 is that it has negative Euler characteristic, and so $\tau_{[p,q]}$ is a true train track, not a bigon track.

We shall describe the conjugacy classification in $\mathcal{MCG}(T_1) = \mathrm{SL}(2, \mathbf{Z})$. The center of $\mathrm{SL}(2, \mathbf{Z})$ is the order 2 subgroup $\{I, -I\}$, each element of which forms a conjugacy class all on its own. Outside of these two special cases, one uses the trace, a conjugacy invariant, to break the conjugacy classification into three broad cases: *elliptic elements* with $|\mathrm{Tr}(M)| < 2$, *parabolic elements* with $|\mathrm{Tr}(M)| = 2$, and *hyperbolic* or *Anosov elements* with $|\mathrm{Tr}(M)| > 2$.

Note that multiplication by $-I$ changes the sign of the trace, and $M, N \in$

$\mathrm{SL}(2, \mathbf{Z})$ are conjugate if and only if $-M, -N$ are conjugate. It therefore suffices to classify the elements of non-negative trace. We will describe explicitly all conjugacy classes with traces between -2 and $+2$, and after that all conjugacy classes with trace > 2 .

Elliptic and parabolic elements. Elements $M \in \mathrm{SL}(2, \mathbf{Z})$ with $|\mathrm{Tr}(M)| \leq 1$ all have finite order, and they are completely classified by the rotation number of their action on the circle of rays in \mathbf{R}^2 . Elements of trace 0 have order 4, and there are two conjugacy classes, one of rotation number $+1/4$ represented by $\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$, the other of rotation number $-1/4$ represented by $\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix}$. Elements of trace 1 have order 6, and there are two conjugacy classes, one of rotation number $+1/6$ represented by $\begin{pmatrix} 0 & -1 \\ 1 & 1 \end{pmatrix}$, and the other of rotation number $-1/6$ represented by $\begin{pmatrix} 1 & 1 \\ -1 & 0 \end{pmatrix}$. Elements of trace -1 have order 3, and there are two conjugacy classes, one of rotation number $+1/3$ represented by $\begin{pmatrix} -1 & -1 \\ 1 & 0 \end{pmatrix}$, and the other of rotation number $-1/3$ represented by $\begin{pmatrix} 1 & 1 \\ -1 & 0 \end{pmatrix}$.

Each $M \in \mathrm{SL}(2, \mathbf{Z})$ with $\mathrm{Tr}(M) = 2$ is conjugate to a unique matrix of the form $\begin{pmatrix} 1 & n \\ 0 & 1 \end{pmatrix}$, and the integer n is a complete conjugacy invariant. Each $M \in \mathrm{SL}(2, \mathbf{Z})$ with $\mathrm{Tr}(M) = -2$ is conjugate to a unique matrix of the form $\begin{pmatrix} -1 & n \\ 0 & -1 \end{pmatrix}$ and again n is a complete conjugacy invariant.

We leave the proofs of the elliptic and parabolic classifications to the reader.

Anosov elements. It remains to describe the conjugacy classification for elements $A \in \mathrm{SL}(2, \mathbf{Z})$ with $\mathrm{Tr}(A) > 2$. There are two positive eigenvalues $\lambda > 1$ and $\lambda^{-1} < 1$. Let r be the slope of the eigenvector with eigenvalue λ . The number r is called the *expanding eigendirection* of A , and it is an attracting fixed point with respect to the fractional linear action of A on \mathbf{R}^* . Using the equation $A(r) = r$ it follows easily that the number r is a quadratic irrationality, and hence by a standard theorem its continued fraction expansion $r = [n_0, n_1, n_2, \dots]$ is eventually periodic. Let $2j$ be the smallest positive even integer that is a period for $[n_0, n_1, n_2, \dots]$, and for a sufficiently large even integer i define the *primitive even period loop* for r to be the sequence of integers $(n_i, n_{i+1}, \dots, n_{i+2j-1})$, well-defined up to even cyclic permutation. Even cyclic permutation means that an even number of terms are removed from the beginning of the sequence and attached to the end.

Theorem 10.1.1. *Given $A \in \mathrm{SL}(2, \mathbf{Z})$ satisfying $\mathrm{Tr}(A) > 2$, and with expanding eigendirection r , the trace of A together with the even cyclic permutation class of the primitive even period loop of the continued fraction expansion of r form a complete conjugacy invariant for A .*

Rather than proving this theorem in “elementary” terms, expressed purely in terms of continued fractions, we prefer to give a train track proof of this theorem, in

order to motivate the generalization to arbitrary finite type surfaces. In particular, the proof we give will generalize to prove Theorems 10.2.7 and 10.3.2, our main theorems on classification of pseudo-Anosov conjugacy classes.

Moreover, the proof we give for Theorem 10.1.1 could be translated directly into “elementary” terms by simply replacing the train track $\tau_{[\frac{a}{b}, \frac{c}{d}]}$, $ad - bc = \pm 1$ by the *Farey interval* $[\frac{a}{b}, \frac{c}{d}]$, and then doing everything in terms of these Farey intervals.

Train track expansions on the torus. Every recurrent, filling, completely combed train track on T_1 is isotopic to some $\tau_{[p,q]}$.

A key fact needed in the proof of Theorem 10.1.1 will be that if τ, τ' are any two recurrent, filling, completely combed traintracks on T_1 then there exists $C \in \text{SL}(2, \mathbf{Z})$ such that $M_C(\tau) = \tau'$ up to isotopy. To see why this is true, let $\tau = \tau_{[\frac{a}{b}, \frac{c}{d}]}$ and $\tau' = \tau_{[\frac{a'}{b'}, \frac{c'}{d}]}$. Then the matrix

$$C = \begin{pmatrix} a' & b' \\ c' & d' \end{pmatrix} \begin{pmatrix} a & b \\ c & d \end{pmatrix}^{-1}$$

will do the trick.

As described in Section 1.3, there are unique splittings of parity R, L, and C on $\tau_{[p,q]}$; the resulting train tracks are easily described easily in terms of Farey summation. It follows that for every measured foliation \mathcal{F}_r carried on a train track τ_0 there exists a unique train track expansion of \mathcal{F}_r based at τ_0 consisting entirely of completely combed train tracks, and this expansion is unique up to isotopy subject to the requirement that if r is rational then the expansion ends in the simple closed curve c_r . From the Stable Equivalence Theorem 7.2.3 applied to the surface T_1 (or from well known facts about stable equivalence of continued fraction expansions), it follows that when r is irrational, any two train track expansions of \mathcal{F}_r are stably equivalent. In the case that $r \in [0, \infty]$ and $\tau_0 = \tau_{[0, \infty]}$, the continued fraction expansion of r can be read off from the RL sequence of the train track expansion of r , as noted in Section 1.3.

Consider now an Anosov matrix $A \in \text{SL}(2, \mathbf{Z})$ with $\text{Tr}(A) > 2$, and corresponding mapping class M_A on T_1 . Let r be the expanding eigendirection of A , and λ the expanding eigenvalue, and it follows that $M_A(\mathcal{F}_r) = \lambda\mathcal{F}_r$, and that the projective class of \mathcal{F}_r is an attracting fixed point in \mathcal{PMF} . Choosing any train track expansion $\tau_0 \succ \tau_1 \succ \dots$ for \mathcal{F}_r , the splitting sequence $M_A(\tau_0) \succ M_A(\tau_1) \succ \dots$ is a train track expansion for $\lambda\mathcal{F}_r$ and hence also for \mathcal{F}_r . It follows that $\tau_0 \succ \tau_1 \succ \dots$ is stably equivalent to $M_A(\tau_0) \succ M_A(\tau_1) \succ \dots$, that is, for some $m, n \geq 0$ and each $i \geq 0$ the train tracks $\tau_{m+i}, M_A(\tau_{n+i})$ are isotopic, and the parities of the splittings $\tau_{m+i} \succ \tau_{m+i+1}$ and $M_A(\tau_{n+i}) \succ M_A(\tau_{n+i+1})$ are identical. As we saw in Section 9.1, since \mathcal{F}_r is the unstable foliation it follows that $n < m$, and by truncating the sequence at τ_n we may assume that $n = 0$ and that $M_A(\tau_i) = \tau_{m+i}$ for all

$i \geq 0$. Let D_i be the parity of the splitting $\tau_{i-1} \succ \tau_i$, so the sequence D_1, D_2, \dots is periodic with period m (which might not be the primitive period). Consider the period loop (D_1, \dots, D_m) indexed by the \mathbf{Z}/m ; we call this the *parity loop*, and it depends ostensibly on the choice of τ_0 .

We shall reduce Theorem 10.1.1 to the following:

Theorem 10.1.2. *The parity loop (D_1, \dots, D_m) , up to cyclic permutation, is a complete conjugacy invariant for A .*

Proof. First we check that the parity loop (D_1, \dots, D_m) , up to cyclic permutation, is well-defined independent of the choice of base train track τ_0 . For any other base train track τ'_0 , the expansion of \mathcal{F}_r based at τ'_0 is stably equivalent to that based at τ_0 , and hence the parity sequences are stably equivalent. Moreover, for some m, n, m', n' we have $M_A(\tau_{n+i}) = \tau_{m+i}$ and $M_A(\tau'_{n'+i}) = \tau'_{m'+i}$ for all $i \geq 0$, and by stable equivalence of the sequences τ, τ' it follows that $m - n = m' - n'$. This implies that the parity loops associated to τ_0 and to τ'_0 are identical up to cyclic permutation.

Next we check that the parity loop is a conjugacy invariant. If $A, A' \in \text{SL}(2, \mathbf{Z})$ are conjugate, and both have trace > 2 , choose $C \in \text{SL}(2, \mathbf{Z})$ so that $C^{-1}A'C = A$. Let r, r' be the expanding eigendirections of A, A' , respectively, and so we have $c(r) = r'$, and so $M_C(\mathcal{F}_r) = \mathcal{F}_{r'}$. Choose a base train track τ_0 for an expansion $\tau_0 \succ \tau_1 \succ \dots$ of \mathcal{F}_r , and taking $\tau'_i = M_C(\tau_i)$ it follows that $\tau'_0 \succ \tau'_1 \succ \dots$ is an expansion of $\mathcal{F}_{r'}$. It also follows that the parity of the splitting $\tau_{i-1} \succ \tau_i$ equals the parity of $\tau'_{i-1} \succ \tau'_i$, and that the translation distance of M_A acting on $\tau_0 \succ \tau_1 \succ \dots$ equals the translation distance of $M_{A'}$ acting on $\tau'_0 \succ \tau'_1 \succ \dots$. Thus, the period loops of M_A and of $M_{A'}$ are identical.

Finally we check that the parity loop is a complete conjugacy invariant.

Define

$$A_L = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \quad A_R = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}$$

Note that $A_L(\tau_{[0,\infty]}) = \tau_{[1,\infty]}$ and $A_R(\tau_{[0,\infty]}) = \tau_{[0,1]}$. It follows that for $r \in [0, \infty]$, if \mathcal{F}_r has train track expansion $\tau_{[0,\infty]} = \tau_0 \succ \tau_1 \succ \dots$ with a LR sequence D_1, D_2, \dots , then $\mathcal{F}_{A_L(r)}$ has a train track expansion

$$\tau_{[0,\infty]} \succ^L \tau_{[1,\infty]} = A_L(\tau_0) \succ A_L(\tau_1) \succ \dots$$

and similarly for $\mathcal{F}_{A_R(r)}$. To put it another way, the effect of acting by A_L on train track expansions starting at $\tau_{[0,\infty]}$ is to prepend an L to the LR sequence, and the effect of acting by A_R is to prepend an R.

Given any parity loop (D_1, \dots, D_m) consisting of at least one L and one R, we shall show that any element of $\text{SL}(2, \mathbf{Z})$ with trace > 2 and with parity loop

(D_1, \dots, D_m) is conjugate to the matrix

$$A(D_1, D_2, \dots, D_m) = A_{D_1} A_{D_2} \cdots A_{D_m}$$

Consider the infinite train track sequence

$$\tau_{[\frac{0}{0}, \frac{1}{0}]} = \tau_0 \succ^{D_1} \tau_1 \succ^{D_2} \dots \succ^{D_m} \tau_m \succ \dots$$

where we extend the period loop (D_1, D_2, \dots, D_m) to a sequence of period m indexed by $i \geq 1$. Note that

$$M_{A(D_1, \dots, D_m)} \tau_i = \tau_{m+i}$$

and hence (D_1, \dots, D_m) is the parity loop of the matrix $A(D_1, \dots, D_m)$. Now consider any $A \in \text{SL}(2, \mathbf{Z})$ with trace > 2 , with expanding eigendirection r' , and with the same period loop (D_1, \dots, D_m) . As we saw above, by truncating any train track expansion for $\mathcal{F}_{r'}$ we obtain an expansion of the form

$$\tau'_0 \succ^{D_1} \tau'_1 \succ^{D_2} \dots \succ^{D_m} \tau'_m \succ \dots$$

whose RL sequence is the periodic extension of (D_1, \dots, D_m) . As noted earlier, there exists $C \in \text{SL}(2, \mathbf{Z})$ so that $C(\tau_0) = \tau'_0$; since the splitting sequences $\tau_0 \succ \tau_1 \succ \dots$ and $\tau'_0 \succ \tau'_1 \succ \dots$ have the same RL sequence, it follows that $C(\tau_i) = \tau'_i$ for all $i \geq 0$. We thus have $CA(\tau_0) = A'C(\tau_0)$ and so $C^{-1}A'^{-1}CA$ fixes the train track τ_0 . The only two mapping classes that can fix a filling track track on T_1 are $\pm I$. Since A, A' have positive trace, it is easy to check that $C^{-1}A'^{-1}CA = -I$ is impossible, and hence $C^{-1}A'^{-1}CA = I$, proving that A, A' are conjugate. \diamond

Proof of Theorem 10.1.1. To carry out the proof we must simply translate between the invariants described in Theorem 10.1.1 and those in Theorem 10.1.2.

Consider a parity loop (D_1, \dots, D_m) with at least one R and one L. By cyclic permutation we may assume that $D_1 = \text{L}$ and $D_m = \text{R}$, and so m is even. Let n_0 be the number of initial Ls, n_1 the number of following Rs, \dots , n_{m-1} the number of final Rs. First note that if $A \in \text{SL}(2, \mathbf{Z})$ is conjugate to $A(D_1, \dots, D_m)$ with expanding eigendirection r then the continued fraction expansion of r is eventually periodic with period loop (n_0, \dots, n_{m-1}) . Since m is even, this immediately implies that the primitive even period loop is determined up to even cyclic permutation by the even cyclic permutation class of (n_0, \dots, n_{m-1}) . This proves that the even cyclic permutation class of the primitive even period loop is a conjugacy invariant of A .

Conversely, suppose that A, A' have the same trace > 2 , and expanding eigendirections r, r' with the same primitive even period loop (n_0, \dots, n_{m-1}) , m even. By conjugating A, A' we may assume that the continued fraction expansions of r, r' are

both periodic with period loop (n_0, \dots, n_{m-1}) , in which case we have $r = r'$. It suffices to show that $A = A'$.

A finite sequence is *primitive* if it cannot be expressed as an iterate of a shorter finite sequence. Letting D_1, \dots, D_k be the LR sequence starting with n_0 Ls, followed by n_1 Rs, etc., up to n_{m-1} Rs, it follows that (D_1, \dots, D_k) is a primitive sequence, because (n_0, \dots, n_{m-1}) cannot be expressed as an iterate of a shorter sequence of even length. Consider the matrix $A(D_1, \dots, D_k)$. Since $A(r) = r$ and $A'(r) = r$, it follows that the parity loops of both A and A' are iterates of (D_1, \dots, D_k) , which implies that both A and A' are powers of $A(D_1, \dots, D_k)$. However, distinct powers of $A(D_1, \dots, D_k)$ have distinct traces, and so $A = A'$. \diamond

Remark on data size. If $A \in \text{SL}(2, \mathbf{Z})$ with $\text{Tr}(A) > 2$ then A is conjugate to a positive matrix, and there is a finite number of positive matrices in the conjugacy class of A , to be precise: all matrices of the form $A(D_1, \dots, D_m)$ where (D_1, \dots, D_m) is any RL sequence in the same cyclic permutation class as the parity loop of A . The set of these matrices is therefore a complete invariant of M up to conjugacy. However, the size of the data needed to store this set of matrices is considerably larger than the size of the data needed to describe the parity loop of M . Consider an example for which the parity loop is primitive, of length N , and has all constant parity blocks of length 1 or 2. In this case the number of positive matrices in the conjugacy class is N , and each entry of each matrix has a number of decimal digits which is linear in N , thereby requiring storage of size on the order of N^2 decimal digits. On the other hand, the parity loop itself has size N . Thus, one obtains a much more tractable invariant—one which takes up much less storage space and which is still a complete invariant—by using the parity loop. Not the least advantage of the parity loop is the ease of comparison needed to efficiently solve the conjugacy problem: once the parity loop of two matrices of trace > 2 has been computed, comparing the loops for cyclic equivalence is a snap.

As noted in the introductory remarks to Section 10, data size is also an important factor in obtaining efficient conjugacy invariant of pseudo-Anosov in general.

10.2 Pseudo-Anosov conjugacy using one cusp splitting circuits

In this section we show how to use one cusp splitting circuits to obtain a complete classification of pseudo-Anosov conjugacy classes. Our goal is to make the classification so explicit that, in Section 11.1, we will (almost) be able to enumerate pseudo-Anosov conjugacy classes, as was done in Section 10.1 for the torus (the meaning of “almost” enumeration will be explained in the beginning of Section 11.1).

Here is an outline of the method.

Consider a pseudo-Anosov mapping class ϕ with unstable foliation \mathcal{F} , choose a separatrix ℓ of ϕ , and let $\tau_0 \succ \tau_1 \succ \dots$ be a canonical one cusp expansion of ϕ associated to the separatrix ℓ . If ϕ fixes ℓ then by One Cusp Stable Equivalence ?? we may truncate the sequence and choose M so that $\phi(\tau_i) = \tau_{i+M}$ for all $i \geq 0$. We therefore obtain a circuit of combinatorial types of one cusp train tracks $[\tau_0] \succ \dots \succ [\tau_M] = [\tau_0]$, and we will prove in Theorem ?? that the set of such circuits, one for each separatrix fixed by ϕ , is a complete invariant of the conjugacy class of ϕ .

This leaves us with the problem of classifying pseudo-Anosov conjugacy classes which do not have a fixed separatrix. To solve this problem, we use the virtual centralizer $\text{VC}(\phi)$ of a pseudo-Anosov mapping class $\phi \in \mathcal{MCG}$, which is the subgroup of all $\psi \in \mathcal{MCG}$ such that ψ commutes with some power of ϕ . We shall prove:

Proposition 10.2.1. *Two pseudo-Anosov mapping classes $\phi, \psi \in \mathcal{MCG}$ are conjugate if and only if the subgroups $\text{VC}(\phi), \text{VC}(\psi)$ are conjugate, say $\text{VC}(\phi) = \alpha^{-1} \text{VC}(\psi) \alpha$, and the elements $\phi, \alpha^{-1} \psi \alpha$ are conjugate in the group $\text{VC}(\phi)$.*

Thus, we reduce the conjugacy classification of pseudo-Anosov mapping classes to two problems: the conjugacy classification of maximal two-ended virtually pseudo-Anosov subgroups; and the conjugacy classification within such subgroups. The latter problem is straightforward: in general, conjugacy classification in a two-ended group G is easy, assuming one has an explicit description of G .

For the former problem, given a pseudo-Anosov mapping class ϕ , we shall pick out of $\text{VC}(\phi)$ a particular pseudo-Anosov element $\bar{\phi}$ which is canonical in the sense that $\text{VC}(\phi)$ and $\text{VC}(\phi')$ are conjugate subgroups of \mathcal{MCG} if and only if $\bar{\phi}, \bar{\phi}'$ are conjugate elements of \mathcal{MCG} . Moreover, $\bar{\phi}$ will have the property that it fixes every separatrix of its unstable foliation, and so the conjugacy class of $\bar{\phi}$ is completely described by its set of one cusp canonical splitting circuits. Thus we obtain a classification of $\bar{\phi}$, which gives a classification of $\text{VC}(\phi)$, which gives the desired classification of ϕ .

The virtual centralizer of a pseudo-Anosov mapping class. In any group G , define the *virtual centralizer* $\text{VC}(g)$ of $g \in G$ to be the subgroup of all $h \in G$ such that for some $m \neq 0$ we have $h^{-1} g^m h = g^m$. Let $\langle g \rangle$ be the infinite cyclic group generated by g . Define the *virtual normalizer* $\text{VN}(g)$ to be the subgroup of all $h \in \mathcal{MCG}$ such that $\langle g \rangle \cap h^{-1} \langle g \rangle h$ has finite index in both $\langle g \rangle$ and $h^{-1} \langle g \rangle h$.

Proposition 10.2.1 is a special case of the following elementary fact, which has the effect of factoring conjugacy classification into two steps:

Proposition 10.2.2. *In any group G , two elements g, g' are conjugate if and only if the subgroups $\text{VC}(g), \text{VC}(g')$ are conjugate and, for any $h \in G$ such that $h^{-1} \text{VC}(g') h = \text{VC}(g)$, the elements $h^{-1} g' h$ and g are conjugate in the group $\text{VC}(g)$. \diamond*

Indeed, this proposition holds with $\text{VC}(g)$ replaced by $\text{VN}(g)$ or by any subgroup $H(g)$ naturally associated to g and containing g , such as the ordinary centralizer or normalizer; here “natural” means that for any group isomorphism $f: G \rightarrow G'$ and any $g \in G$ we have $f(H(g)) = H(f(g))$.

Recall that a finitely generated group V is *two-ended* if its Cayley graph has two ends; by Stallings’ theory of ends [Sta68] this is true if and only if V contains a finite index, infinite cyclic subgroup. If V is two ended then the action of V on its Cayley graph extends to an action on its set of ends, whose kernel is a characteristic subgroup of index at most 2 in V called the *translation subgroup* $T(V)$. There is a commutative diagram with exact rows

$$\begin{array}{ccccccc}
 1 & \longrightarrow & K(V) & \longrightarrow & T(V) & \longrightarrow & \mathbf{Z} \longrightarrow 1 \\
 & & \parallel & & \downarrow & & \downarrow \\
 1 & \longrightarrow & K(V) & \longrightarrow & V & \longrightarrow & C \longrightarrow 1
 \end{array}$$

with the following properties: either $C \approx \mathbf{Z}$ or $C \approx D_\infty$, the infinite dihedral group; the injection $\mathbf{Z} \rightarrow C$ has index 1 or 2; and $K(V)$ is a characteristic finite subgroup of V called the *kernel* of V . Note that $V = T(V)$ if and only if C is infinite cyclic. Also, if C is infinite dihedral, then each element of $V - T(V)$ has square lying in $K(V)$ and hence has finite order. Thus, all infinite order elements of V lie in $T(V)$. Given a group G and $g \in G$ for which $\text{VN}(g)$ is two-ended, the subgroup $T(\text{VN}(g))$ is equal to $\text{VC}(g)$.

Consider a pseudo-Anosov mapping class $\Phi \in \mathcal{MCG}$, with repelling fixed point ξ^r and attracting fixed point ξ^a in \mathcal{PMF} . We denote $K(\Phi) = K(\text{VN}(\Phi))$. The following proposition is essentially from [BLM83], with a few details added:

Proposition 10.2.3. *If $\Phi \in \mathcal{MCG}$ is pseudo-Anosov, then $\text{VN}(\Phi)$ is two-ended and $\text{VC}(\Phi)$ is the translational subgroup of $\text{VN}(\Phi)$. The group $\text{VN}(\Phi)$ is equal to the subgroup of all $\Psi \in \mathcal{MCG}$ such that $\Psi\{\xi^r, \xi^a\} = \{\xi^r, \xi^a\}$, and the action of $\text{VN}(\Phi)$ on the set $\{\xi^r, \xi^a\}$ is conjugate to the action of $\text{VN}(\Phi)$ on its set of ends. It follows that $\text{VC}(\Phi)$ is equal to the subgroup of all $\Psi \in \mathcal{MCG}$ such that $\Psi(\xi^r) = \xi^r$ and $\Psi(\xi^a) = \xi^a$. \diamond*

The group $\text{VN}(\Phi)$ can be made to act on the surface S itself in a particularly nice manner. Let $\mathcal{F}^s, \mathcal{F}^u$ be the stable and unstable foliations of a pseudo-Anosov representative ϕ of Φ , and so $P[\mathcal{F}^s] = \xi^r$ and $P[\mathcal{F}^u] = \xi^a$. The measured foliations $\mathcal{F}^s, \mathcal{F}^u$ determine a singular Euclidean structure on S : away from the singular points this structure is described by the metric $ds^2 = (d\mathcal{F}^s)^2 + (d\mathcal{F}^u)^2$, where $d\mathcal{F}^s$ and $d\mathcal{F}^u$ denote the transverse measures on \mathcal{F}^s and \mathcal{F}^u . Define $\widehat{\text{Aff}}(\phi)$ to be the set of homeomorphisms of S which are affine with respect to the singular Euclidean

structure ds and which preserve the collection of horizontal and vertical leaves, possibly interchanging horizontal with vertical. Define $\text{Aff}(\phi)$ to be the subgroup of index at most two in $\widehat{\text{Aff}}(\phi)$ which preserves horizontal and preserves vertical. Define $\text{Isom}(\phi)$ to be the finite subgroup of $\text{Aff}(\phi)$ which acts isometrically on S . We have:

Proposition 10.2.4. *The quotient map $\text{Homeo}_+ \rightarrow \mathcal{MCG}$ restricts to an injection of $\text{Aff}(\phi)$ with image $\text{VN}(\Phi)$, taking $\text{Aff}(\phi)$ to $\text{VC}(\Phi)$, and taking $\text{Isom}(\phi)$ to $K(\Phi)$. \diamond*

It follows that $\text{Aff}(\phi)$ acts on the set $S^u(\phi)$ of separatrices of \mathcal{F}^u . The kernel of this action is an infinite cyclic subgroup of $\text{Aff}(\phi)$ that we denote $Z(\phi)$, and there is a unique generator of this group denoted $\zeta(\phi)$ which is pseudo-Anosov with stable foliation \mathcal{F}^s and unstable foliation \mathcal{F}^u (the other generator has stable foliation \mathcal{F}^u and unstable foliation \mathcal{F}^s). If we choose ϕ' to be any other pseudo-Anosov representative of Φ , with stable and unstable foliations $\mathcal{F}'^s, \mathcal{F}'^u$, then there exists $\psi \in \text{Homeo}_0$ and constants $\alpha, \beta > 0$ such that $\psi(\mathcal{F}'^s) = \alpha\mathcal{F}'^s$ and $\psi(\mathcal{F}'^u) = \beta\mathcal{F}'^u$, and hence ψ conjugates the action of $\text{Aff}(\phi)$ to the action of $\text{Aff}(\phi')$. This implies that the quotient map $\text{Homeo}_+ \rightarrow \mathcal{MCG}$ takes $Z(\phi)$ to a well defined subgroup of $\text{VC}(\Phi)$ denoted $Z(\Phi)$, and it takes $\zeta(\phi)$ to a well-defined element of $\text{VN}(\Phi)$ denoted $\zeta(\Phi)$. Moreover we have:

Proposition 10.2.5. *Given pseudo-Anosov mapping classes $\Phi, \Phi' \in \mathcal{MCG}$, the subgroups $\text{VC}(\Phi), \text{VC}(\Phi')$ are conjugate if and only if the mapping classes $\zeta(\Phi), \zeta(\Phi')$ are conjugate.*

Proof. If Φ, Φ' are conjugate, one can choose pseudo-Anosov representatives ϕ, ϕ' and $\psi \in \text{Homeo}_+$ conjugating ϕ to ϕ' , and it follows that ψ conjugates $\zeta(\phi)$ to $\zeta(\phi')$, and so the mapping class Ψ of ψ conjugates $\zeta(\Phi)$ to $\zeta(\Phi')$. Conversely, suppose that $\zeta(\Phi), \zeta(\Phi')$ are conjugate, and it follows that $\text{VC}(\zeta(\Phi)), \text{VC}(\zeta(\Phi'))$ are conjugate, but clearly $\text{VC}(\Phi) = \text{VC}(\zeta(\Phi))$ and $\text{VC}(\Phi') = \text{VC}(\zeta(\Phi'))$. \diamond

Combining Propositions 10.2.1 and 10.2.5, we have reduced pseudo-Anosov conjugacy classification to two problems: classifying the elements $\zeta(\Phi)$ up to conjugacy, and classifying elements of a two-ended, translational group up to conjugacy.

The conjugacy problem in a two ended, translational group G has an easy solution as follows. The group G is a semidirect product $K(G) \rtimes_f Z$ where Z is infinite cyclic, and $f: Z \rightarrow \text{Aut}(K(G))$ is an action of Z on $K(G)$ which determines the semidirect product by the usual formula $(k, z) \cdot (k', z') = (k \cdot f(z)(k'), z \cdot z')$. Let z be a generator of Z . The kernel of the action f is an infinite cyclic subgroup of Z generated by z^k for some $k \geq 1$, and clearly z^k is in the center of G . It follows that if (k, z^p) is conjugate to (l, z^q) , then there is a conjugator of the form (m, z^r) where

$0 \leq r < k$. In order to test conjugacy in G , therefore, we are restricted to a finite set of possible conjugators that need to be tested.

Remark. A little speculation and/or calculation should convince one that “most” pseudo-Anosov mapping classes Φ have the property that $\text{VC}(\Phi)$ is infinite cyclic; it would be interesting to prove a concrete formulation of this statement. The property that $\text{VC}(\Phi)$ is infinite cyclic is clearly a conjugacy invariant of Φ .

One cusp splitting circuits. Given a pseudo-Anosov mapping class Φ , we have associated to Φ a pseudo-Anosov mapping class $z(\Phi)$ with the same stable and unstable foliations, characterized as the smallest mapping class acting as the identity on the separatrices of the stable and unstable foliations. We have reduced the conjugacy classification of Φ to the conjugacy classification of $z(\Phi)$.

What we shall do in this section is to use one cusp splitting circuits to give a complete conjugacy classification of all pseudo-Anosov mapping classes Φ which have *at least one* fixed separatrix.

Consider a pseudo-Anosov mapping class Φ represented by a homeomorphism ϕ with stable and unstable foliations $\mathcal{F}^s, \mathcal{F}^u$ and expansion factor $\lambda > 0$, and let $\ell \in s(\mathcal{F}^u)$ be a separatrix fixed by ϕ . Let τ_0 be any train track that canonically carries \mathcal{F}^u . Let $v \in \text{cusps}(\tau_0)$ be the cusp corresponding to ℓ . Let $\tau_0 \succ \tau_1 \succ \cdots$ be the one cusp expansion of \mathcal{F}^u based at τ_0 with starting cusp v_0 .

Lemma 10.2.6. *With the notation as above, the splitting sequence $\tau_0 \succ \tau_1 \succ \cdots$ is eventually ϕ -periodic, that is, there exists integers $I \geq 0, N \geq 1$ such that for all $i \geq I$ we have $\phi(\tau_i) = \tau_{i+N}$ up to isotopy.*

Proof. Applying ϕ to $\tau_0 \succ \tau_1 \succ \cdots$ we obtain a one cusp expansion $\phi(\tau_0) \succ \phi(\tau_1) \succ \cdots$ of $\phi(\mathcal{F}^u) = \lambda\mathcal{F}^u$, based at the cusp $\phi(v_0) \in \text{cusps}(\phi(\tau_0))$ that corresponds to $\phi(\ell)$. Since $\mathcal{PMF}(\mathcal{F}^u) = \mathcal{PMF}(\lambda\mathcal{F}^u)$ it follows that $\phi(\tau_0) \succ \phi(\tau_1) \succ \cdots$ is the one cusp splitting of \mathcal{F}^u based at τ_0 , with starting cusp $\phi(v_0)$ corresponding to $\phi(\ell)$. Since $\phi(\ell) = \ell$, the cusps v_0 and $\phi(v_0)$ correspond via the separatrices of \mathcal{F}^u . Applying Theorem 7.1.5, it follows that the splitting sequences $\tau_0 \succ \tau_1 \succ \cdots$ and $\phi(\tau_0) \succ \phi(\tau_1) \succ \cdots$ are stably equivalent. Choose integers $a, b \geq 0$ so that for all $i \geq 0$ the train tracks τ_{a+i} and $\phi(\tau_{b+i})$ are isotopic.

We must show that $a > b$. It cannot happen that $a = b$ because in that case $\phi(\tau_a) = \tau_a$ implying by Corollary 3.15.3 that the mapping class Φ has finite order, a contradiction. If $a < b$ then from the equation $\phi^{-1}(\tau_n) = \tau_{n+(b-a)}$ for $n \geq a$ it follows that τ_a is an invariant train track for ϕ^{-1} , and that $\tau_a \succ \tau_{a+1} \succ \cdots$ is a ϕ^{-1} periodic splitting sequence, and hence by applying Theorem 9.1.1 we conclude that $\cap_i \mathcal{PMF}(\tau_{a+i}) = \mathcal{PMF}(\mathcal{F}^s)$, because \mathcal{F}^s is the unstable foliation of ϕ^{-1} . But

by applying the Expansion Convergence Theorem 5.1.1 we obtain $\cap_i \mathcal{PMF}(\tau_i) = \mathcal{PMF}(\mathcal{F}^u)$, contradicting the obvious fact that $\mathcal{PMF}(\mathcal{F}^u) \cap \mathcal{PMF}(\mathcal{F}^s) = \emptyset$. \diamond

From the lemma, it follows that the train track τ_I is an invariant train track for ϕ and the splitting sequence $\tau_I \succ \tau_{I+1} \succ \dots$ is ϕ -periodic with period N . The construction of this sequence depended on the choice of a separatrix $\ell \in s(\mathcal{F}^u)$, and we call this a *one cusp ϕ -periodic expansion of \mathcal{F}^u , associated to the separatrix ℓ* . Although the sequence $\tau_I \succ \tau_{I+1} \succ \dots$ depended on the choice of τ_0 , clearly it is independent of τ_0 up to stable equivalence.

Remark. Taking ϕ -periodicity into account, we can say a little more: if $\tau'_{I'} \succ \tau'_{I'+1} \succ \dots$ is another one cusp ϕ -periodic expansion of \mathcal{F}^u associated to ℓ , one of the two splitting sequences $\tau_I \succ \tau_{I+1} \succ \dots$ and $\tau'_{I'} \succ \tau'_{I'+1} \succ \dots$ is obtained from the other (up to reindexing) by truncating an initial segment. It follows that a one cusp, ϕ -periodic expansion of \mathcal{F}^u associated to ℓ extends to a well-defined *bi-infinite ϕ -periodic expansion of \mathcal{F}^u associated to ℓ* , which we think of as a “train track axis” of ϕ .

Now we show how to associate a “splitting circuit” to each ϕ -periodic one cusp train track axis. This circuit is a cyclically ordered sequence of combinatorial types of one cusp train tracks, together with a parity assigned to each successive pair of train tracks in the circuit.

Let $\tau_0 \succ \tau_1 \succ \dots$ be a ϕ -periodic, one cusp expansion of \mathcal{F}^u , associated to a ϕ -invariant separatrix ℓ of \mathcal{F}^u . Choose cusps $v_i \in \text{cusps}(\tau_i)$ corresponding to ℓ , so $(\tau_{i-1}, v_{i-1}) \succ (\tau_i, v_i)$ are one cusp splittings. Let N be the periodicity of the sequence. It follows that $\phi^k(\tau_i, v_i) = (\tau_{i+kN}, v_{i+kN})$ for each $i, k \in \mathbf{Z}$, and hence for any congruence class $I \in \mathbf{Z}/N$, for all integers $i \in I$ the one cusp train tracks (τ_i, v_i) are all combinatorially equivalent; let T_I denote this combinatorial equivalence class of one cusp train tracks. It also follows that for i, i' in the same congruence class mod N the parities of the splittings $\tau_{i-1} \succ \tau_i$ and $\tau_{i'-1} \succ \tau_{i'}$ are equal; for a congruence class $I \in \mathbf{Z}/N$ let $D_I \in \{L, R\}$ denote the associated parity. The pair cyclic sequences T_I, D_I ($I \in \mathbf{Z}/N$) is called the *one cusp splitting circuit* associated to ϕ and ℓ , and we denote this data as

$$T_0 \succ^{d_1} T_1 \succ^{d_2} \dots \succ^{d_N} T_N = T_0$$

Two one cusp splitting circuits are considered to be equivalent if they have the same length N and their terms are identical up to a cyclic permutation of the index set \mathbf{Z}/N .

Since any two cusp expansion of \mathcal{F}^u associated to ℓ are stably equivalent, it follows that their one cusp splitting circuits are equivalent.

The following theorem combines with Lemma ?? to give a complete almost conjugacy classification of pseudo-Anosov mapping classes, but it does a little more: it gives a complete conjugacy classification of pseudo-Anosov mapping classes Φ that have at least one fixed separatrix. The gist of the theorem is that the one cusp splitting circuits associated to the fixed separatrices form a complete conjugacy invariant of Φ . The proof is basically an adaptation of the proof of unmarked stable equivalence for one cusp expansions 7.1.7.

Theorem 10.2.7 (Conjugacy of pseudo-Anosovs with fixed separatrices).

Let Φ be a pseudo-Anosov mapping class represented by a pseudo-Anosov homeomorphism ϕ that fixes at least one separatrix of its unstable foliation \mathcal{F}^u . Given $\Phi' \in \mathcal{MCG}$, the following are equivalent:

- (1) Φ and Φ' are conjugate
- (2) Φ' is represented by a pseudo-Anosov homeomorphism ϕ' that fixes at least one separatrix of its unstable foliation \mathcal{F}'^u , and the one cusp splitting circuit of ϕ associated to some fixed separatrix $\ell \in s(\mathcal{F}^u)$ is equivalent to the one cusp splitting circuit of ϕ' associated to some fixed separatrix $\ell' \in s(\mathcal{F}'^u)$.
- (3) Φ' is represented by a pseudo-Anosov homeomorphism that fixes at least one separatrix of its unstable foliation, and there is a bijective correspondence between the set of one cusp splitting circuits of Φ and the set of one cusp splitting circuits of Φ' , such that corresponding circuits are equivalent.

Proof. The implication (3) \implies (2) is immediate.

The proof of (1) \implies (3) is a consequence of naturality of one cusp expansions, as follows. Assuming $\Phi = \Psi\Phi'\Psi^{-1}$, we have $\Psi(\mathcal{F}^u) = \mathcal{F}'^u$, inducing a bijection $s(\mathcal{F}^u) \rightarrow s(\mathcal{F}'^u)$ that takes Φ fixed separatrices to Φ' fixed separatrices. If $\ell \in s(\mathcal{F}^u)$ is fixed by Φ , and if $\tau_0 \succ \tau_1 \succ \dots$ is the one cusp expansion associated to ℓ , it follows that $\ell' = \Psi(\ell)$ is fixed by Φ' and that $\Phi(\tau_0) \succ \Phi(\tau_1) \succ \dots$ is the one cusp expansion associated to ℓ' . But clearly the one cusp splitting circuits associated to these two expansions are equivalent: the mapping class Φ provides the needed combinatorial equivalencies between one cusped train tracks.

To prove (2) \implies (1), consider one cusp expansions of $\mathcal{F}^u, \mathcal{F}'^u$ associated to fixed separatrices ℓ, ℓ' of Φ, Φ' , respectively:

$$\begin{aligned} \tau_0 \succ \tau_1 \succ \dots \\ \tau'_0 \succ \tau'_1 \succ \dots \end{aligned}$$

and suppose that the associated splitting circuits of these two sequences are both equivalent to

$$T_0 \succ^{d_1} T_1 \succ^{d_2} \dots \succ^{d_N} T_N = T_0$$

We show that Φ, Φ' are conjugate. Choose cusp sequences $v_i \in \text{cusps}(\tau_i)$, $v'_i \in \text{cusps}(\tau'_i)$, $i \geq 0$, so that $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ and $(\tau'_i, v'_i) \succ (\tau'_{i+1}, v'_{i+1})$ are one cusp splittings for each $i \geq 0$. Up to truncation and reindexing of the splitting sequences it follows that for each congruence class I modulo N and each $i \in I$ the one cusp train tracks (τ_i, v_i) and (τ'_i, v'_i) are in the combinatorial equivalence class T_I , and the one cusp splittings $\tau_{i-1} \succ \tau_i$ and $\tau'_{i-1} \succ \tau'_i$ have the same parity D_i . Choose a combinatorial equivalence θ_i from (τ_i, v_i) to (τ'_i, v'_i) . The same argument that was used in the proof of Theorem 7.1.7 (2) \implies (1) shows that the mapping classes of the θ_i are all identical; let Θ denote this mapping class. We have

$$\Theta\Phi(\tau_0, v_0) = \Theta(\tau_N, v_N) = (\tau'_N, v'_N)$$

and

$$\Phi'\Theta(\tau_0, v_0) = \Phi'(\tau'_0, v'_0) = (\tau'_N, v'_N)$$

and so by Lemma 3.15.2 it follows that $\Theta\Phi = \Phi'\Theta$. ◇

10.3 Pseudo-Anosov conjugacy using circular expansion complexes

In the previous section, we used stable equivalence for one cusp expansions to prove, essentially, that the one cusp expansions associated to unstable foliations of pseudo-Anosov homeomorphisms are eventually periodic in the sense of combinatorial type. The periodicity data was used to obtain a not-quite-complete conjugacy classification of pseudo-Anosov mapping classes: we obtained a complete classification up to almost conjugacy, which by Lemma ?? is a uniformly bounded finite-to-one conjugacy classification; we also obtained a complete classification of a *subclass* of pseudo-Anosov mapping classes, namely those with a fixed separatrix, and this subclass contains a uniformly bounded power of every pseudo-Anosov mapping class.

In order to obtain a complete conjugacy classification, we use the expansion complex to tie up into one package all expansions (up to stable equivalence) of the unstable foliation, not just the one cusp expansions. Stable equivalence of expansion complexes will imply that expansion complexes of pseudo-Anosov unstable foliations are eventually periodic in the appropriate sense, and we use the periodicity data to obtain a complete conjugacy classification of pseudo-Anosov mapping classes.

Eventual periodicity of expansion complexes. Consider a pseudo-Anosov mapping class Φ represented by a pseudo-Anosov homeomorphism ϕ with unstable foliation \mathcal{F}^u . Fix a train track τ canonically carrying \mathcal{F}^u . Let Ξ be the expansion complex of \mathcal{F}^u based at τ , each of whose closed cells $c(\sigma)$ is labelled by the isotopy class of a decorated train track σ . We shall prove Lemma 10.3.1 which expresses eventual Φ -periodicity of Ξ .

Recall from Section 7.7 that for each face relation $c(\sigma) \subset c(\sigma')$ in Ξ there is a forest collapse map $q_{\sigma'\sigma}: \sigma' \rightarrow \sigma$, well defined up to homotopy through forest collapse maps.

Consider now $\Phi(\Xi)$, which is the expansion complex of $\Phi(\mathcal{F}^u)$ based at $\Phi(\tau)$. Since $\mathcal{PMF}(\mathcal{F}^u) = \mathcal{PMF}(\Phi(\mathcal{F}^u))$ it follows from the Stable Equivalence Theorem 7.2.3 or from its corollary 7.2.6 that Ξ and $\Phi(\Xi)$ are stably equivalent. We may therefore fix subcomplexes $\hat{\Xi} \subset \Xi$, $\hat{\Xi}' \subset \Xi$, each a neighborhood of infinity, such that Φ takes $\hat{\Xi}$ isomorphically to $\hat{\Xi}'$, by a map taking a cell $c(\sigma)$ to the cell $c(\Phi(\sigma))$. Moreover, from the proof of Unmarked Stable Equivalence 7.7.1, for any face relation $c(\sigma) \subset c(\sigma')$ in $\hat{\Xi}$ the following diagram commutes up to homotopy through unmarked face collapse maps:

$$\begin{array}{ccc}
 \sigma' & \xrightarrow{\phi} & \Phi(\sigma') \\
 q_{\sigma'\sigma} \downarrow & & \downarrow q_{\Phi(\sigma')\Phi(\sigma)} \\
 \sigma & \xrightarrow{\phi} & \Phi(\sigma)
 \end{array}
 \tag{*}$$

Lemma 10.3.1 (Eventual periodicity of expansion complexes). *The subcomplex $\hat{\Xi}$ of Ξ can be chosen so that it is a connected neighborhood of infinity, and so that $\Phi(\hat{\Xi}) \subset \hat{\Xi}$. For all $\xi \in \hat{\Xi}$, the sequence $\Phi^n(\xi)$ diverges to infinity in Ξ as $n \rightarrow \infty$.*

We shall apply this lemma repeatedly in what follows.

Proof. Start with a subcomplex $\hat{\Xi} \subset \Xi$, a neighborhood of infinity, so that $\Phi(\hat{\Xi}) \subset \Xi$, as constructed above using stable equivalence of Ξ and $\Phi(\Xi)$. We describe how to alter $\hat{\Xi}$ so as to guarantee that $\Phi(\hat{\Xi}) \subset \hat{\Xi}$.

We claim that there exists a cell $c(\sigma)$ such that $c(\phi^n(\sigma))$ is a cell of $\hat{\Xi}$ for all $n \geq 0$. To prove the claim, choose a separatrix ℓ of \mathcal{F}^u and choose K so that $\phi^K(\ell) = \ell$. The coordinate axis of $\hat{\Xi}$ corresponding to ℓ is labelled by an elementary factorization of the one cusp expansion of \mathcal{F} based at τ associated to ℓ , denoted $\tau = \tau_0 \succ \tau_1 \succ \dots$. By Lemma 10.2.6, this sequence is eventually ϕ^K periodic, that is, there exists $N, I \geq 0$ such that $\phi^K(\tau_i) = \tau_{i+N}$ for all $i \geq I$. For each $i = I, \dots, N - 1$ the sequence of cells $c(\tau_i), c(\tau_{i+N}), c(\tau_{i+2N}), \dots$ lies in a single coordinate axis of Ξ , diverging to infinity in that axis, and hence the sequence eventually lies in $\hat{\Xi}$. It follows that there exists $I \geq 0$ such that $c(\tau_i) \in \hat{\Xi}$ for all $i \geq I$. Taking $\sigma = \tau_I$ proves the claim.

For each i consider the subcomplex $\hat{\Xi}(\phi^i(\sigma))$ consisting of all cells $c(\sigma') \subset \hat{\Xi}$ such that $\phi^i(\sigma) \succ \sigma'$. By Corollary 7.2.4, $\hat{\Xi}(\phi^i(\sigma))$ is a neighborhood of infinity. Now we

redefine $\hat{\Xi}$ to be

$$\hat{\Xi} = \bigcup_{i=0}^{\infty} \hat{\Xi}(\phi^i(\sigma)) = \bigcup_{i=0}^{K-1} \hat{\Xi}(\phi^i(\sigma))$$

This is clearly a connected neighborhood of infinity, and $\Phi(\hat{\Xi}) \subset \hat{\Xi}$. For each cell $c(\sigma)$, if $\phi^n(c(\sigma))$ did not diverge to infinity then there would be an infinite subsequence n_i such that $\phi^{n_i}(c(\sigma))$ is constant, contradicting Corollary 3.15.3 which says that the stabilizer of a filling train track is a finite subgroup of \mathcal{MCG} . It follows that $\phi^n(\xi)$ diverges to infinity for each $\xi \in \Xi$. \diamond

Here is a somewhat more informative proof that $\phi^n(\xi)$ diverges to infinity for each $\xi \in \Xi$, which will be useful in what follows.

Consider the length function $L: \Xi \rightarrow [0, \infty)$ defined by $L(\xi) = \sum_i \text{Length}(\xi_i)$. Divergence to infinity of a point $\xi \in \Xi$ is equivalent to divergence to infinity of the length function $L(\xi)$. For $t > 0$ the level set $L^{-1}(t)$ inherits a stratification from Ξ , and with respect to this stratification $L^{-1}(t)$ is a simplex. Choose $T > 0$ so that $L^{-1}[T, \infty) \subset \hat{\Xi}$. Note that for $n \geq 0$ the set $\Phi^n(L^{-1}[T, \infty)) = L^{-1}[\lambda^n T, \infty)$ is a subset of $\hat{\Xi}$, where $\lambda > 1$ is the expansion factor of Φ . It follows that if $\xi \in L^{-1}[T, \infty)$ then $\Phi^n(\xi)$ diverges to infinity as $n \rightarrow \infty$. Each of the finitely many cells of $\hat{\Xi}$ which are not contained in $L^{-1}[T, \infty)$ will be mapped into $L^{-1}[T, \infty)$ after some finite power of Φ .

The simplices $L^{-1}(t)$ form the fibers of a fibration of $L^{-1}[T, \infty)$ over $[T, \infty)$, giving this domain the structure of a (simplex) \times (half-line).

Circular expansion complexes. The circular expansion complex of Φ , denoted Ω , will be a stratified topological space, equipped with a labelled cell structure and a labelled face relation.

As a cell-complex, $\Omega = \hat{\Xi}/\Phi$ is the quotient of $\hat{\Xi}$ by the action of the semigroup generated by Φ ; let $Q: \hat{\Xi} \rightarrow \Omega$ denote the quotient map. The cells of Ω fit together to define a CW-decomposition of Ω . We shall slightly alter our notation for cells, reserving a lower case c for cells of Ω and an upper case C for cells of Ξ . The cell of Ξ labelled by a decorated train track σ is thus denoted $C(\sigma)$.

Note that Q need not be injective on each closed cell of $\hat{\Xi}$, and so the closure of an open cell of Ω need not be homeomorphic to a closed cell, that is, Ω need not be a cell complex. However, this situation seems extremely rare, and it might be hard to construct a concrete example where Ω is not a cell complex. One would need to construct an example with a closed cell $C \subset \hat{\Xi}$ and two faces $C(\sigma'), C(\sigma'') \subset C$ such that $\Phi^n(\sigma') = \sigma''$ for some $n \geq 0$.

Ignoring for a moment the cell structure of Ω , the topology and the stratification of Ω can be described alternately as follows. Note that Ω is identified with the

quotient of $L^\infty[T, \infty)$ by Φ , with the non-negative powers of Φ acting as deck transformation semigroup. A fundamental domain for this semi-group action is $L^{-1}[T, \lambda T]$, with Φ taking the simplex fiber $L^{-1}(T)$ isomorphically to the simplex fiber $L^{-1}(\lambda T)$, mapping vertices to vertices according to the action of Φ on unstable separatrices. The fibration of $L^{-1}[T, \infty)$ over $[T, \infty)$ descends to a fibration of Ω over the circle $S^1_\lambda = \mathbf{R}/\log(\lambda) \cdot \mathbf{Z}$, as described in the following commutative diagram:

$$\begin{array}{ccc} L^{-1}[T, \infty) & \xrightarrow{Q} & \Omega \\ L \downarrow & & \downarrow \\ [T, \infty) & \xrightarrow{\log} & S^1_\lambda \end{array}$$

The strata of Ξ intersected with $\hat{\Xi}$ descend to strata of Ω . Each stratum is a subcomplex. Each fiber of the fibration $\Omega \rightarrow S^1_\lambda$, intersected with the strata, gives that fiber the structure of a simplex, and so as a stratified space Ω has the structure of a simplex bundle over the circle. Note that the permutation on the vertices of this simplex, under the monodromy map of the fibration, is conjugate to the action of Φ on unstable separatrices.

Now we discuss the labelling on cells and on face relations of Ω .

Each cell c of Ω is labelled with a decorated train track as follows. Choose a lift $C = C(\sigma)$ of c , meaning a cell C of $\hat{\Xi}$ whose interior projects via Q homeomorphically to the interior of c . Now label the cell c with the same decorated train track σ that labels C ; the cell c labelled by σ is denoted $c(\sigma)$. Note that σ is *not* well-defined up to isotopy independent of the choice of the chosen lift C , although it is well-defined up to combinatorial equivalence: cells in the same Φ orbit of $\hat{\Xi}$ are labelled by decorated train tracks which are combinatorially equivalent via powers of Φ .

Each face relation $c(\sigma) \subset c(\sigma')$ of Ω is labelled with an unmarked forest collapse $q_{\sigma'\sigma}: \sigma \rightarrow \sigma'$ as follows. Consider the chosen lifts $C(\sigma), C(\sigma') \subset \hat{\Xi}$. By construction there exists an integer n , not necessarily positive, such that $C(\sigma)$ is a face of $\Phi^n(C(\sigma')) = C(\Phi^n(\sigma'))$. It follows that there is a (marked) forest collapse $q: \Phi^n(\sigma') \rightarrow \sigma$, and we define the unmarked forest collapse $q_{\sigma'\sigma}$ as the composition

$$\sigma' \xrightarrow{\Phi^n} \Phi^n(\sigma') \xrightarrow{q} \sigma$$

$\underbrace{\hspace{10em}}_{q_{\sigma'\sigma}}$

This completes the description of the circular expansion complex Ω of a pseudo-Anosov mapping class Φ . Notice that Ω depends ostensibly on choices, namely, the choice of the the expansion complex Ξ , and the choice of the lifted cells $C = \tilde{c}$ used to define the labelling on Ω . However, we shall show that Ω is a well-defined

independent of these choices, with the appropriate definition of “isomorphism” of circular expansion complexes. Not only that, but Ω is a well-defined invariant of the conjugacy class of Φ , and in fact it is a complete invariant, as we will prove in Theorem 10.3.2 below.

Remarks. Notice that unlike the situation with one cusp splitting circuits, we do *not* just label a cell with the whole combinatorial type of a decorated train track: we must choose a particular representative of this combinatorial type, in order to have a domain and range for the unmarked collapse map $q_{\sigma'\sigma}$. This would all be unnecessary if we had assurances that the decorated train tracks σ, σ' had trivial stabilizers in \mathcal{MCG} , for then *any* choice of representatives of the combinatorial types of σ, σ' would yield a unique unmarked forest collapse $q_{\sigma'\sigma}$ up to homotopy through forest collapses, and hence there would be no loss in labelling the cells c, c' by the combinatorial types of σ, σ' rather than by the isotopy types of σ, σ' . However, the possibility of finite stabilizers in \mathcal{MCG} forces the labelling structure to be described in terms such as we have done. Readers familiar with the theory of orbihedra [Hae91] may discern that another way to approach this issue is by working in the category of orbihedra.

Consider now two circular expansion complexes Ω, Ω' . An *isomorphism* from Ω to Ω' consists of a stratification preserving CW-isomorphism $G: \Omega \rightarrow \Omega'$, and for each cell $c(\sigma) \subset \Omega$ with image $c(\sigma') = G(c(\sigma)) \subset \Omega'$, a decoration preserving combinatorial equivalence $g_{\sigma\sigma'}: \sigma \rightarrow \sigma'$, so that for each face relation $c(\sigma_1) \subset c(\sigma_2)$ in Ω , setting $c(\sigma'_1) = G(c(\sigma_1))$ and $c(\sigma'_2) = G(c(\sigma_2))$, the following diagram commutes up to homotopy through unmarked forest collapses:

$$\begin{array}{ccc}
 \sigma_2 & \xrightarrow{g_{\sigma_2\sigma'_2}} & \sigma'_2 \\
 q_{\sigma_2\sigma_1} \downarrow & & \downarrow q_{\sigma'_2\sigma'_1} \\
 \sigma_1 & \xrightarrow{g_{\sigma_1\sigma'_1}} & \sigma'_1
 \end{array}$$

Using this notion of isomorphism we next show that the circular expansion complex Ω of a pseudo-Anosov mapping class Φ is well-defined up to isomorphism, independent of the choices in the construction of Ω . There are two choices to consider: choice of the lifts of cells $c \subset \Omega$; and choice of the expansion complex Ξ of the unstable foliation of Φ .

Consider a cell $c(\sigma) \subset \Omega$ and its chosen lift $C(\sigma) \subset \hat{\Xi}$. Let Ω' be the circular expansion complex obtained from Ω by changing the chosen lift C to another lift $C' = C(\sigma') \subset \hat{\Xi}$. Thus, the labelling on c is changed from σ to σ' , and all unmarked

forest collapse maps labelling face relations $c_1 \subset c$ and $c \subset c_2$ are changed appropriately; all labels not involving c are unchanged. We shall prove that the identity map $G: \Omega \rightarrow \Omega'$ is an isomorphism. There exists an integer n such that $\Phi^n(\sigma) = \sigma'$ and $\Phi^n(C) = C'$; and so $g_{\sigma\sigma'} = \Phi^n$ is a decoration preserving isomorphism from σ to σ' , as required to define an isomorphism. Consider a face $c_1 = c(\sigma_1)$ of c in Ω , and let $C_1 = C(\sigma_1) \subset \hat{\Xi}$ be the chosen lift of c_1 . There exists an integer m such that C_1 is a face of $\Phi^m(C)$, and we have an unmarked forest collapse $q: \Phi^m(\sigma) \rightarrow \sigma_1$ that labels the face relation $c_1 \subset c$ in Ω . Changing the label on c from σ to $\Phi^n(\sigma) = \sigma'$, note that C_1 is a face of $\Phi^{m-n}(C') = \Phi^m(C)$, and so the unmarked forest collapse $q': \Phi^{m-n}(\sigma') \rightarrow \sigma_1$ is the label of the face relation $c_1 \subset c$ in Ω' . Clearly q' is homotopic to $\Phi^n \circ q$ through unmarked forest collapses—in fact they are the same map—as required for an isomorphism. A similar argument when c is a face of a cell $c_2 \subset \Omega$ establishes that the identity map $G: \Omega \rightarrow \Omega'$ is an isomorphism.

We have shown that Ω is well-defined up to isomorphism when we change the chosen lift of a single cell, and by induction on the number of such changes it follows that Ω is well-defined independent of the choice of lifts.

Next we show that Ω is well-defined when we change the expansion complex Ξ of \mathcal{F}^u . Suppose we choose another expansion complex Ξ' , based at a train track τ' . Choose subcomplexes $\hat{\Xi} \subset \Xi$, $\hat{\Xi}' \subset \Xi'$, neighborhoods of infinity, so that $\Phi(\hat{\Xi}) \subset \hat{\Xi}$ and $\Phi(\hat{\Xi}') \subset \hat{\Xi}'$. Let $\Omega = \hat{\Xi}/\Phi$, $\Omega' = \hat{\Xi}'/\Phi$ be the circular expansion complexes. Applying the Stable Equivalence Theorem 7.2.3 we obtain subcomplexes $\hat{\hat{\Xi}} \subset \hat{\Xi}$ and $\hat{\hat{\Xi}}' \subset \hat{\Xi}'$, each a neighborhood of infinity, and a (marked) isomorphism $F: \hat{\hat{\Xi}} \rightarrow \hat{\hat{\Xi}}'$. Note that this isomorphism commutes with Φ , and so F induces a stratification preserving, cellular isomorphism from Ω to Ω' . To check that cell labels and face relation labels are preserved, we make use of the freedom of choosing lifts of cells: define labels on Ω by choosing lifts of cells to lie in $\hat{\hat{\Xi}} \cap \hat{\Xi} \cap F^{-1}(\hat{\hat{\Xi}}' \cap \hat{\Xi}')$, and define Ω' by choosing lifts of cells to be the images under F of the lifts of cells in Ω , necessarily lying in $F(\hat{\hat{\Xi}} \cap \hat{\Xi}) \cap \hat{\hat{\Xi}}' \cap \hat{\Xi}'$. It follows that Ω, Ω' are isomorphic. This completes the proof that circular expansion complexes are well-defined.

Here is our main classification theorem:

Theorem 10.3.2. *Two pseudo-Anosov mapping classes $\Phi, \Phi' \in \text{MCG}$ are conjugate in MCG if and only if the circular expansion complexes of Φ and Φ' are isomorphic.*

The proof of this theorem shares several features with the proof of unmarked stable equivalence of expansion complexes 7.7.1.

Proof. Let $\mathcal{F}^u, \mathcal{F}'^u$ be unstable foliations for Φ, Φ' . Choose a train tracks τ, τ' canonically carrying $\mathcal{F}^u, \mathcal{F}'^u$, let Ξ, Ξ' be the expansion complexes of \mathcal{F}^u based at

τ, τ' , and choose subcomplexes $\hat{\Xi} \subset \Xi$ and $\hat{\Xi}' \subset \Xi'$, neighborhoods of infinity, so that $\Phi(\hat{\Xi}) \subset \hat{\Xi}$ and $\Phi'(\hat{\Xi}') \subset \hat{\Xi}'$. We have circular expansion complexes $\Omega = \hat{\Xi}/\Phi$, $\Omega' = \hat{\Xi}'/\Phi'$.

Suppose that Φ, Φ' are conjugate, $\Phi' = \Psi\Phi\Psi^{-1}$ for $\Psi \in \mathcal{MCG}$. It follows that $\mathcal{F}'^u = \Psi(\mathcal{F}^u)$ is an unstable foliation for Φ' . We may choose $\Xi' = \Psi(\Xi)$ based at $\tau' = \Psi(\tau)$. Moreover, Ψ is an unmarked stable equivalence between Ξ and Ξ' . Since $\Phi'\Psi = \Psi\Phi$, we may choose $\hat{\Xi}' = \Psi(\hat{\Xi}) \subset \Xi'$. It follows that Ψ induces an isomorphism between circular expansion complexes $\Omega = \hat{\Xi}/\Phi$ and $\Omega' = \hat{\Xi}'/\Phi'$.

We now prove the converse: if Ω, Ω' are isomorphic then Φ, Φ' are conjugate. The idea of the proof is that the isomorphism $G: \Omega \rightarrow \Omega'$ lifts to an unmarked stable equivalence between Ξ and Ξ' , which according to Theorem 7.7.1 is induced by some mapping class Ψ , and we will have $\Phi' = \Psi\Phi\Psi^{-1}$. The details are a bit trickier and we do not, in fact, quote Theorem 7.7.1, although we do use one important step from the proof.

First, we want to lift $G: \Omega \rightarrow \Omega'$ to a stratification preserving CW-isomorphism $F: \hat{\Xi} \rightarrow \hat{\Xi}'$ with the property that $\Phi' \circ F = F \circ \Phi$. But we can't always get what we want, because the quotient maps $Q: \hat{\Xi} \rightarrow \Omega$ and $Q': \hat{\Xi}' \rightarrow \Omega'$ are not true covering maps, and so it seems that we cannot use lifting theory ([Spa81], chapter 2). However, we get what we need by dropping the requirement that F be a surjection, requiring F only to be a stratification preserving, CW-isomorphism onto a subcomplex of $\hat{\Xi}'$. To do this, fix a base cell $c \subset \hat{\Xi}$. If the cell $Q(c) \subset \Omega$ is lifted via Q' to a cell c' of $\hat{\Xi}'$ so that c' is sufficiently close to infinity, then the proof of lifting theory provides the needed map F so that $F(c) = c'$. Details are left to the reader. Note that $F(\hat{\Xi})$ is a neighborhood of infinity in $\hat{\Xi}'$.

Next we want to lift the combinatorial equivalencies $g_{\sigma\sigma'}$, relating cell labels of Ω and labels of Ω' , to combinatorial equivalencies $f_{\sigma\sigma'}$ that relate cell labels of $\hat{\Xi}$ to cell labels of $F(\hat{\Xi})$. The chosen lifts of cells of Ω lie in $\hat{\Xi}$, and we may assume that the chosen lifts of cells of Ω' lie in $\hat{\Xi}'$. In fact, we may assume that for each cell $c(\sigma)$ of Ω , if $c' = c(\sigma') = G(c) \subset \Omega$, and if $C = C(\sigma) \subset \hat{\Xi}$ and $C' = C(\sigma') \subset \hat{\Xi}'$ are the chosen lifts, then $F(C) = C'$. The isomorphism between Ω and Ω' provides us with a combinatorial equivalence $g_{\sigma\sigma'}: \sigma \rightarrow \sigma'$. This gives us a combinatorial equivalence from some representative of each Φ orbit of cells in $\hat{\Xi}$ to some representative of each Φ' orbit of cells in $\hat{\Xi}'$. Now we use equivariance to extend this as follows. Consider an arbitrary cell C of $\hat{\Xi}$, let $c = Q(C) \subset \Omega$, and let $c' = G(c) \subset \Omega'$. Choose an integer n so that $\Phi^n(C) \subset \hat{\Xi}$ is the chosen lift of c . It follows that $C' = \Phi'^n(F(C)) = F(\Phi^n(C))$ is the chosen lift of c' . Letting σ, σ' be the decorated train tracks labelling the cells C, C' , it follows that $\Phi^n(\sigma)$ and $\Phi'^n(\sigma')$ are the labels of c, c' , and hence we have a combinatorial equivalence

$$g = g_{\Phi^n(\sigma), \Phi'^n(\sigma')} : \Phi^n(\sigma) \rightarrow \Phi'^n(\sigma')$$

By composition we define a combinatorial equivalence $f_{\sigma\sigma'} : \sigma \rightarrow \sigma'$:

$$\begin{array}{ccccccc}
 \sigma & \xrightarrow{\Phi^n} & \Phi^n(\sigma) & \xrightarrow{g} & \Phi'^n(\sigma') & \xrightarrow{\Phi'^{-n}} & \sigma' \\
 & & & & & \searrow & \\
 & & & & & & f_{\sigma\sigma'}
 \end{array}$$

We now wish to show that F and the system of combinatorial equivalences $f_{\sigma\sigma'}$ form an unmarked stable equivalence from Ξ to Ξ' , and to do this we must verify that the maps $f_{\sigma\sigma'}$ respect face relations. The verification is aided by the commutative diagram below. In this diagram and the accompanying discussion, as a shorthand each cell of Ξ or Ξ' is represented by its labelling train track, e.g. we write σ as a shorthand for $C(\sigma)$.

Consider a face relation $C(\sigma_1) \subset C(\sigma_2)$ in $\hat{\Xi}$ and the image face relation $C(\sigma'_1) = F(C(\sigma_1)) \subset C(\sigma'_2) = F(C(\sigma_2))$ in $F(\hat{\Xi})$. We must prove commutativity of the square labelled ? in the diagram below.

The orbit of $C(\sigma_1)$ contains a cell $C(\sigma_1^n) = C(\Phi^n(\sigma_1))$ which is the chosen lift of the cell $Q(C(\sigma_1)) \subset \Omega$, and the orbit of $C(\sigma'_1)$ contains a cell $C(\sigma_1'^n) = C(\Phi'^n(\sigma'_1))$ which is the chosen lift of $Q(C(\sigma'_1))$. The isomorphism $G: \Omega \rightarrow \Omega'$ provides us with a combinatorial equivalence $g_1: \sigma_1^n \rightarrow \sigma_1'^n$. Square 1 commutes by definition of $f_{\sigma_1\sigma_1'}$.

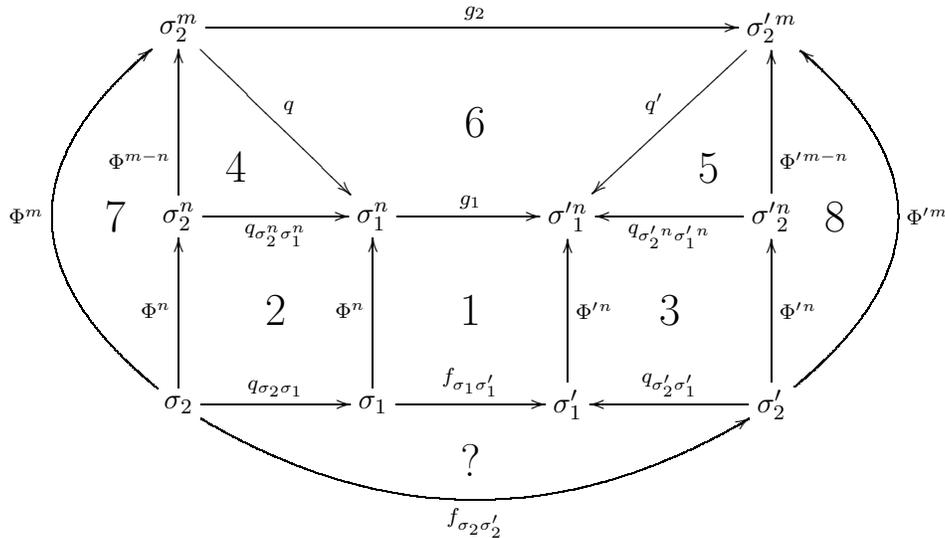
Letting $\sigma_2^n = \Phi^n(\sigma_2)$ and $\sigma_2'^m = \Phi'^m(\sigma_2')$, squares 2 and 3 commute by naturality of (marked) forest collapse maps.

The orbit of $C(\sigma_2)$ contains a cell $C(\sigma_2^m) = C(\Phi^m(\sigma_2))$ which is the chosen lift of the cell $Q(C(\sigma_2)) \subset \Omega$, and the orbit of $C(\sigma_2')$ contains a cell $C(\sigma_2'^m) = C(\Phi'^m(\sigma_2'))$ which is the chosen lift of $Q(C(\sigma_2'))$. The isomorphism $G: \Omega \rightarrow \Omega'$ provides us with a combinatorial equivalence $g_2: \sigma_2^m \rightarrow \sigma_2'^m$. The exterior square in the diagram, which is unlabelled, commutes by definition of $f_{\sigma_2\sigma_2'}$.

Since $\Phi^{m-n}(\sigma_2^n) = \sigma_2^m$, the labelling of face relations of Ω provides us with an unmarked forest collapse $q: \sigma_2^m \rightarrow \sigma_1^n$ which, by definition, is the composition $q_{\sigma_2^n\sigma_1^n} \circ \Phi^{n-m}$, and hence triangle 4 commutes. Similarly, since $\Phi'^{m-n}(\sigma_2'^n) = \sigma_2'^m$, Ω' provides us with an unmarked forest collapse $q': \sigma_2'^m \rightarrow \sigma_1'^n$ which equals $q_{\sigma_2'^n\sigma_1'^n} \circ \Phi'^{n-m}$, and hence triangle 5 commutes.

Commutativity of square 6 follows because $G: \Omega \rightarrow \Omega'$ is an isomorphism, all the cells involved in this diagram are the chosen lifts of cells in Ω or Ω' , and the requirement that the combinatorial equivalences of an isomorphism respect face relations.

Commutativity of triangles 7 and 8 is obvious, and the desired commutativity of square ? follows.



This completes the proof that F is an unmarked stable equivalence from Ξ to Ξ' . By the same argument used in the proof of Theorem 7.7.1, all of the combinatorial equivalences $f_{\sigma\sigma'}: \sigma \rightarrow \sigma' = F(\sigma)$ are in the same mapping class, denoted Ψ . Now choose a cell $C(\sigma) \subset \hat{\Xi}$ that lies on an axis of Ξ , and hence the decoration of σ contains a unique nondegenerate cusp v . It follows that $\sigma' = F(C(\sigma)) = C(\Psi(\sigma))$ lies on an axis of Ξ' and hence has a unique nondegenerate cusp v' . Since $f_{\sigma\sigma'}$ preserves decorations, the bijection $\Psi: \text{cusps}(\sigma) \rightarrow \text{cusps}(\sigma')$ takes v to v' . The cell $\Phi(C(\sigma)) = C(\Phi(\sigma))$ is also on an axis, the undecorated cusp of $\Phi(\sigma)$ is $\Phi(v)$, and the cell $\Phi'(C(\sigma')) = C(\Phi'(\sigma'))$ is on an axis, and the undecorated cusp of $\Phi'(\sigma')$ is $\Phi'(v')$. Finally, since $f_{\Phi(\sigma)\Phi'(\sigma')}$ preserved decorations it follows that $\Psi: \text{cusps}(\Phi(\sigma)) \rightarrow \text{cusps}(\Phi'(\sigma'))$ takes $\Phi(v)$ to $\Phi'(v')$. We have shown that $\Psi\Phi(\sigma, v) = \Phi'\Psi(\sigma, v)$. Applying Lemma 3.15.2 it follows that $\Psi\Phi = \Phi'\Psi$, and so Φ, Φ' are conjugate in \mathcal{MCG} . \diamond

Two cusp splitting circuits.

TO DO:

- Describe the strata $\Omega^{(n)}$.
- Discuss how $\Omega^{(1)}$ is just the one cusp splitting circuits for the first return maps of all separatrices. So, $\Omega^{(1)}$ does not suffice to classify pseudo-Anosov conjugacy classes. Discuss what it would take to come up with an explicit counterexample, but I don't have one up my sleeve.

- Discuss $\Omega^{(2)}$. All splitting circuits contained in here are called “two cusp splitting circuits”.
- Theorem: $\Omega^{(2)}$ is a complete conjugacy invariant.
- Proof: Just need connectivity.
- Why this is somewhat better than all of Ω . Can one do better? Is there some scheme which augments $\Omega^{(1)}$ with a minimal amount of extra data to make the classification complete? Discussion of one such scheme . . .

11 Enumeration of pseudo-Anosov conjugacy classes

In this section we enumerate the canonical one cusp splitting circuits that are used in Section 10 to classify pseudo-Anosov conjugacy classes. Because these circuits only enumerate a “finite index” collection of pseudo-Anosov conjugacy classes, and because the enumeration is “finite-to-one”, we do not obtain a bijective enumeration of pseudo-Anosov conjugacy classes by this method, however we will be content here with an “almost enumeration”.

The first step in the enumeration is to study the canonical killing criterion in the context of one cusp splitting sequences, in order to prove that the canonical killing automaton over the one cusp train track automaton has degree bounded by a linear function of the genus and number of punctures; this is carried out in Section 11.1.

Then over the next several sections we describe how to enumerate combinatorial types of one cusp train tracks, with a goal towards explicit construction of the one cusp train track automata and their canonical killing automata.

11.1 The canonical killing criterion for one cusp splitting sequences

One cusp splitting circuits are used in Theorem 10.2.7 to classify pseudo-Anosov mapping classes up to almost conjugacy, in fact to classify up to conjugacy those pseudo-Anosov mapping classes that fix at least one separatrix. Our goal in the next few sections is to show how to enumerate such one cusp splitting circuits in a particularly efficient fashion, using finite deterministic automata. This enumeration does not quite suffice to obtain an exact enumeration of conjugacy classes of pseudo-Anosov mapping classes: not every pseudo-Anosov mapping class Φ fixes a separatrix, although a uniformly bounded power of Φ does; and a pseudo-Anosov homeomorphism Φ may fix more than one separatrix, yielding more than one splitting circuit, but we will not have the technology to compare distinct splitting circuits for conjugacy until a later section.

get rid of
“almost
conjugacy”.
Rewrite
FDA com-
ments, in
light of
putting
FDAs ear-
lier: “we
will ex-
plicitly
construct
the FDAs.”

In Section 7.1 we studied stable equivalence of one cusp train track expansions. We defined a one cusp train track to be an ordered pair (τ, v) where τ is a generic train track and $v \in \text{cusps}(\tau)$. In that section we also defined three relations between one cusp train tracks: “one cusp elementary splittings”; “one cusp slides”; and “one cusp splittings” which are a kind of wide splitting that factors uniquely into one cusp slides followed by a one cusp elementary splitting. The objects of study in Section 7.1 were one cusp splitting sequences, that is, sequences of the form $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ where each $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ is a one cusp splitting.

In this section we prove Theorem 11.1.10 which gives an efficient reformulation of the Canonical Killing Criterion for one cusp splitting sequences. In particular, we shall formulate the Canonical Killing Criterion in this context with linear complexity, as remarked in Section 9.5. We will need extensive technical preliminaries before we can state the theorem.

Stable train tracks. It turns out that in order to understand a one cusp expansion up to stable equivalence, one need only consider a subclass of one cusp train tracks, called “stable train tracks”, as we now explain.

Recall that the cusps of a train track τ are identified with the cusps of $\mathcal{C}(S - \tau)$. Consider a generic, filling, one cusp train track (τ, v) . Let s^v denote the switch of τ at which the cusp v is located, let b^v be the branch on the one-ended side of s^v , and let c^v denote the component of $\mathcal{C}(S - \tau)$ having v as a cusp. An *open side* of $\mathcal{C}(S - \tau)$ is a component of $\partial\mathcal{C}(S - \tau) - \text{cusps}(\mathcal{C}(S - \tau))$, and a *closed side* is the closure of an open side. Let $\bar{\partial}^v$ be the union of the cusps and open sides of $\partial\mathcal{C}(S - \tau)$ whose closures do not contain v , and let $\partial^v = \partial\mathcal{C}(S - \tau) - \bar{\partial}^v$. To be explicit: $\bar{\partial}^v$ contains the complete boundary of each component of $\mathcal{C}(S - \tau)$ distinct from c^v ; if c^v has $n \geq 3$ cusps then $\bar{\partial}^v$ contains the $n - 2$ closed sides of c^v that are disjoint from v , and ∂^v consists of v and the two open sides of c^v incident to v ; if c^v has exactly 2 cusps then $\bar{\partial}^v$ contains, as an isolated point, the cusp of c^v distinct from v , and ∂^v consists of the rest of ∂c^v ; and if c^v has exactly 1 cusp then $\partial c^v = \partial^v$. The case where c^v has exactly 2 cusps is somewhat exceptional and must be handled with care, which is why ∂^v and $\bar{\partial}^v$ are defined the way they are. We shall let $\partial(\bar{\partial}^v)$ denote either the two cusps forming the manifold boundary of $\bar{\partial}^v$ when c^v has ≥ 3 cusps, or the isolated cusp of $\bar{\partial}^v$ when c^v has 2 cusps, or the empty set when c^v has 1 cusp.

A *stable train track* is a generic, filling, one cusp train track (τ, v) such that the overlay map $\mathcal{C}(S - \tau) \rightarrow S$ is injective on the set $\bar{\partial}^v$. Note that stability is a combinatorial invariant of one cusp train tracks. To emphasize the usage, the phrase “stable train track” necessarily refers to a one cusp train track (τ, v) .

The following results, particularly the corollary, explain the terminology “sta-

ble”:

Lemma 11.1.1. *If (τ, v) is a stable train track and if $(\tau, v) \succ (\tau', v')$ is a one cusp slide or a one cusp elementary split then (τ', v') is stable.*

Applying this lemma inductively to the elementary factorization of a one cusp splitting sequence we obtain:

Corollary 11.1.2. *If $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ is a one cusp splitting sequence, and if (τ_I, v_I) is stable for some I , then (τ_i, v_i) is stable for all $i \geq I$.* \diamond

Proof of Lemma 11.2.1. Let $(\tau, v) \succ (\tau', v')$ be a one cusp slide or elementary split, and suppose that (τ, v) is stable. Consider the overlay maps $f: \mathcal{C}(S - \tau) \rightarrow S$, $f': \mathcal{C}(S - \tau') \rightarrow S$. From the local model for a slide (Figure 11) or elementary split (Figure 13) one sees easily that for each neighborhood U of v in $\mathcal{C}(S - \tau)$ there is a diffeomorphism $g: \mathcal{C}(S - \tau') \rightarrow \mathcal{C}(S - \tau)$ with the following property: if $x \in \mathcal{C}(S - \tau') - g^{-1}(U)$ then $f'(x) = f(g(x))$. We can choose U so that $U \cap \bar{\partial}^v = \emptyset$, and so $g^{-1}(U) \cap \bar{\partial}^{v'} = \emptyset$. Since f is injective on $\bar{\partial}^v$ it follows that f' is injective on $\bar{\partial}^{v'}$. \diamond

Next we show that stable train tracks do, in fact, occur in one cusp expansions:

Proposition 11.1.3. *If $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ is a one cusp, canonical expansion of an arational measured foliation \mathcal{F} , then for all sufficiently large i the one cusp train track (τ_i, v_i) is stable.*

Proof. By Corollary 11.2.2 we just have to prove that some (τ_i, v_i) is stable, although the proof will show directly that (τ_i, v_i) is stable for all sufficiently large i .

Let ν_0 be a tie bundle over τ_0 , with tie foliation \mathcal{F}_v equipped with a positive Borel measure, and with \mathcal{F} chosen in its equivalence class to be a canonical model with a surjective carrying inclusion $\mathcal{F} \hookrightarrow \nu_0$, so \mathcal{F} is the horizontal foliation. Let ℓ be the separatrix of \mathcal{F} corresponding to the cusp v_0 . Let ξ_t be the initial segment of ℓ of length t . By Lemma 7.1.4 it follows that the elementary move sequence associated to the growing separatrix family ξ_t is the one cusp elementary factorization of $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$. Let ν_t be the tie bundle obtained from ν_0 by slicing along ξ_t , let τ_t be the quotient train track, and let v_t the the cusp of τ_t associated to the boundary point of ξ_t . We must show that (τ_t, v_t) is stable for sufficiently large t . Let $\bar{\partial}^v \nu_t$ be the union of cusps and open sides of ν_t that are disjoint from v_t . It suffices to show that for sufficiently large t , the map $\nu_t \rightarrow \tau_t$ is injective on the set $\bar{\partial}^v \nu_t$, or equivalently, each tie of ν_t intersects $\bar{\partial}^v \nu_t$ in at most one point. To guarantee this, simply choose $t = \text{Length}(\xi_t)$ sufficiently large so that each tie of ν_0 which intersects $\bar{\partial}^v \nu_0$ also intersects the interior of ξ_t . \diamond

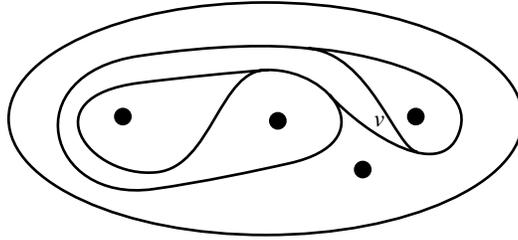


Figure 32: A stable train track, obtained from splitting Figure 6(b). This example is of type 1. A one cusp slide on this example will produce a type 2 stable train track.

The import of Proposition 11.2.3 for our present needs is that in order to study periodic one cusp expansions it suffices to focus on stable train tracks; there is no need to study one cusp train tracks that are not stable.

It should be clear that not every generic, filling train track τ has a cusp v such that (τ, v) is stable: consider for example Figure 6; consider also any one sink train track; and consider also any train track associated to a twist system. In Section 11.2 we shall describe a method for constructing all stable train tracks.

The proof of Proposition 11.2.3 suggests a method for constructing stable train tracks: take an arbitrary one cusp train track (τ, v) , and start doing one cusp splits and slides, trying to maneuver so as to eventually split through every branch of τ . Once this is accomplished, the resulting one cusp train track (τ', v') is guaranteed to be stable. Carrying this process out on the train track in Figure 6(b) results in Figure 32.

Uniqueness of the stable cusp. Given a generic, filling train track τ , we have seen that there may or may not exist $v \in \text{cusps}(\tau)$ for which (τ, v) is stable. One question which arises is: if v exists, how unique is it? The following result gives an (almost) complete answer:

Proposition 11.1.4. *If τ is a generic filling train track then there are exactly zero, one, or two cusps v such that (τ, v) is a stable train track. If there are two cusps, then one of the following happens:*

- (1) *There exists a punctured bigon c of $\mathcal{C}(S - \tau)$ such that for each of the two cusps v of c , (τ, v) is a stable train track.*
- (2) *There exists a component c of $\mathcal{C}(S - \tau)$ with ≥ 3 sides, and there exist two adjacent cusps v_1, v_2 of c , such that (τ, v_1) and (τ, v_2) are stable train tracks.*

As a converse to (1), if (τ, v) is stable, if v is a cusp of a punctured bigon c of $\mathcal{C}(S - \tau)$, and if v' is the other cusp of c , then (τ, v') is also stable.

Stable train tracks which satisfy item (2) in this proposition are somewhat exceptional. In fact they can be safely ignored: if $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ is a one cusp expansion of an arational measured foliation then only finitely many of the (τ_i, v_i) satisfy (2). One can prove this using the methods of Section 11.2.

Proof. Let $p: \mathcal{C}(S - \tau) \rightarrow S$ be the overlay map.

First observe that if (τ, v) is a stable train track then $p(\bar{\partial}^v) \subset p(\partial^v)$, and so $p(\partial^v) = \tau$.

Suppose that $v_1 \neq v_2 \in \text{cusps}(\tau)$ are such that (τ, v_1) and (τ, v_2) are both stable.

If $\partial^{v_1} \cap \partial^{v_2} = \emptyset$ then $\partial^{v_1} \subset \bar{\partial}^{v_2}$ and $\partial^{v_2} \subset \bar{\partial}^{v_1}$, and so from the above observation it follows that $p(\partial^{v_1}) = p(\partial^{v_2}) = \tau$. However, by definition of stability both $\bar{\partial}^{v_1}$ and $\bar{\partial}^{v_2}$ inject into S , and so $p(\partial^{v_1}) = p(\partial^{v_2}) = \tau$ is a nonsmooth 1-manifold, which is absurd.

The intersection $\partial^{v_1} \cap \partial^{v_2}$ must therefore contain at least an open side of $\partial\mathcal{C}(S - \tau)$, which implies that v_1, v_2 are adjacent cusps on the same component c of $\partial\mathcal{C}(S - \tau)$. It follows that v_1, v_2 are the only cusps v of τ for which (τ, v) is stable, and items (1) and (2) are immediate.

The converse to (1) follows immediately from the observation that if c is a punctured bigon with cusps v, v' then $\partial^v = \partial^{v'}$ and $\bar{\partial}^v = \bar{\partial}^{v'}$. \diamond

Types of stable train tracks. Let (τ, v) be a one cusp train track. We fix more notation that will be used hereafter in any discussion of stable train tracks.

Drawing the branch b^v vertically with s^v at the top and the switch orientation pointing downward, the branches on the two-ended side of s^v located to the Left and Right of s^v are denoted b^L and b^R . To be precise, going around the point v in counterclockwise order the branches are encountered in the cyclic order b^R, b^L, b^v . The switch s^v has three overlay pullbacks, one of which is a cusp of $\mathcal{C}(S - \tau)$ identified with v , and the other two are regular points of $\partial\mathcal{C}(S - \tau)$. We distinguish the two regular pullbacks of s^v as the Left pullback \tilde{s}^L and the Right pullback \tilde{s}^R ; to be precise, there is a continuous pullback of b^L with endpoint \tilde{s}^L and a continuous pullback of b^R with endpoint \tilde{s}^R .

Assuming that the one cusp train track (τ, v) is stable, at least one of the two points \tilde{s}^L, \tilde{s}^R must be in ∂^v . If two of \tilde{s}^L, \tilde{s}^R are in ∂^v then we say that (τ, v) is of *type 2*. If exactly one of \tilde{s}^L, \tilde{s}^R is in ∂^v then (τ, v) is of *type 1*, and we further distinguish type 1L when $\tilde{s}^L \in \bar{\partial}^v$ and type 1R when $\tilde{s}^R \in \bar{\partial}^v$. Note that the type is a combinatorial invariant.

For each branch b of τ , the overlay pre-image of $\text{int}(b)$ consists of two smooth arcs in $\partial\mathcal{C}(S - \tau)$, at least one of which is in ∂^v , since (τ, v) is stable. If exactly one of these arcs is in ∂^v then we say that b is a *type 1 branch*, and if both are in ∂^v then b is a *type 2 branch*. For example, the branch b^v has type 2 if and only if (τ, v) has type 2.

For each cusp $v' \neq v$ of τ , the switch $s^{v'}$ has three overlay pre-images in $\partial\mathcal{C}(S - \tau)$, two regular points and one cusp identified with v' , and since $v' \in \bar{\partial}^v$ it follows that the two regular points are in $\bar{\partial}^v$.

The type 1 branches of a stable train track can be described as follows:

Proposition 11.1.5. *Let (τ, v) be a stable train track. If (τ, v) is of type 2 then the overlay image of each side of $\bar{\partial}^v$ is a single type 1 source branch of τ . If (τ, v) is of type 1 then the same is true for all but one side of $\bar{\partial}^v$. The exception, when (τ, v) has type 1D, $D \in \{L, R\}$, is the unique side α containing \tilde{s}^D , and in this case the overlay image of α is a union of two type 1 branches, namely $b^D \cup b^v$. Every type 1 branch of τ is one of the branches described above.*

Proof. Let $p: \mathcal{C}(S - \tau) \rightarrow S$ be the overlay map. Suppose that $p(\text{int}(\alpha))$ contains a switch s' , and let v' be the cusp located at s' ; we regard v' as a cusp of $\mathcal{C}(S - \tau)$. If $s' \neq s^v$ then $v' \in \bar{\partial}^v$, but v' is identified with a point of $\text{int}(\alpha)$, contradicting injectivity of $\bar{\partial}^v$. It follows that $p(\text{int}(\alpha))$ contains at most one switch and that switch must be s^v ; moreover, in this case clearly (τ, v) is of type 1. The remaining contentions of the proposition are straightforward. \diamond

Splits and slides of stable train tracks. If (τ, v) is a one cusp train track, then the branch b^v has an inflow at s^v and so b^v is either a transition branch or a sink branch. If b^v is a transition branch then there is a one cusp slide $(\tau, v) \succcurlyeq (\tau', v')$, and Proposition 11.1.5 implies that this always happens when (τ, v) is of type 1. If b^v is a sink branch then there is a one cusp split $(\tau, v) \succ (\tau', v')$ of each parity $D \in \{L, R\}$. We now analyze these possibilities more carefully when (τ, v) is stable. The case analysis is more complicated when $\partial(\bar{\partial}^v) \neq \emptyset$, so on first reading it may help to focus on the case where $\partial(\bar{\partial}^v) = \emptyset$, which is to say that c^v is a once punctured monogon.

First we handle the case of a type 1 stable train track. To state the result, and for later purposes, we need to define the parity of a slide move. Consider a generic train track τ and a transition branch b . Orient b from its inflow s_0 to its outflow s_1 . Choose an oriented coordinate system defined on a neighborhood of b , taking b to a vertical segment in the plane pointing downward from s_0 to s_1 . At the outflow s_1 there is another branch b' on the two-ended side of s_1 , and the branch b' is located on either the Left or the Right side of b ; choosing $D \in \{L, R\}$ so that b'

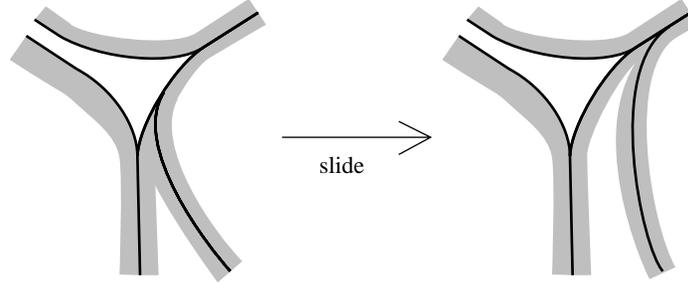


Figure 33: A type 1 stable train track (τ, v) . The only one cusp elementary move $(\tau, v) \succ (\tau', v')$ is a slide move, resulting in a type 2 stable train track (τ', v') . The shading shows sides of ∂^v and $\partial^{v'}$. This example shows a slide of parity R.

is on the D-side of b , we define D to be the parity of the transition branch b . Also, if $\tau \succ \tau'$ is the slide move along b , then we define D to be the parity of this slide move. Figures 33 and 35(b) show slide moves of parity R, and Figure 11 shows one of parity L.

Proposition 11.1.6. *Suppose that (τ, v) is a type 1 stable train track, and so by Proposition 11.1.5 we know that b^v is a transition branch and there is a one cusp slide move $(\tau, v) \succ (\tau', v')$. Then (τ', v') has type 2. Moreover, assuming that (τ, v) has type $1D$, for $D \in \{L, R\}$, the following are true:*

- b^v is a transition branch of parity D and the slide move $(\tau, v) \succ (\tau', v')$ has parity D.
- The branch b^D is of type 1 and the branch $b^{\bar{D}}$ is of type 2.

Proof. Choose $D \in \{L, R\}$ so that (τ, v) has type $1D$. Let α be the side of $\bar{\partial}^v$ containing \bar{s}^D , so the overlay image of α is $b^D \cup b^v$. It immediately follows that b^D has type 1 and $b^{\bar{D}}$ has type 2. Letting s_1 be the switch at the end of b^v opposite s^v , it follows that b^v has an outflow at s_1 , and so b^v is a transition branch and $(\tau, v) \succ (\tau', v')$ is therefore a slide move. Moreover, the cusp v_1 located at s_1 is clearly on the D-side of b^v , and so the slide move has parity D.

Regarding v_1 as a cusp of $\mathcal{C}(S - \tau)$, we have $v_1 \in \alpha \subset \bar{\partial}^v$ and so by stability the two regular pullbacks of s_1 are in ∂^v . This implies that the branch on the one-ended side of s_1 is of type 2. Choosing a point x_1 in the interior of this branch, we can do the slide so that the two regular preimages of $s^{v'}$ are identified with the two regular preimages of x_1 , and so (τ', v') has type 2. See Figure 33. \diamond

It follows from Proposition 11.1.6 that in any sequence of one cusp splits and

slides on stable train tracks, there can never be two consecutive type 1 train tracks. This motivates the following definition.

Given a pair of type 2, stable train tracks (τ, v) , (τ', v') , we say that $(\tau, v) \succ (\tau', v')$ is a *type 2 move* if:

- $(\tau, v) \succ (\tau', v')$ is a one cusp elementary split or slide.
- $(\tau, v) \succ (\tau', v')$ factors as $(\tau, v) \succ (\tau'', v'') \succ (\tau', v')$, each a one cusp elementary split or slide, where (τ'', v'') is of type 1.

In the second case it follows from Proposition 11.1.6 that $(\tau'', v'') \succ (\tau', v')$ is a one cusp slide, but it is still ostensibly possible for $(\tau, v) \succ (\tau'', v'')$ to be either a one cusp slide or a one cusp elementary split. In fact only the latter can occur, so a type 2 move *cannot* be a concatenation of two one cusp slides, as the following proposition shows.

Proposition 11.1.7. *Given a type 2 stable train track (τ, v) , each type 2 move $(\tau, v) \succ (\tau', v')$ satisfies one of the following:*

Type 2 slide: $(\tau, v) \succ (\tau', v')$ is a one cusp slide.

Elementary Type 2 splitting: $(\tau, v) \succ (\tau', v')$ is a one cusp elementary splitting.

Nonelementary type 2 splitting: *There exists $D \in \{L, R\}$ and there exists a type 1 stable train track (τ'', v'') such that $(\tau, v) \succ (\tau', v')$ factors as $(\tau, v) \succ (\tau'', v'') \succ (\tau', v')$, a one cusp elementary splitting of parity D followed by a slide of parity D .*

The three possibilities in Proposition 11.1.7 are depicted in Figure 34.

Proposition 11.1.7 is an immediate consequence of a more precise statement, Proposition 11.1.8. In understanding this statement and its proof, it may be best on first reading to consider the special case that $\partial(\bar{\partial}^v) = \emptyset$, which occurs precisely when c^v is a once punctured monogon. In this case Proposition 11.1.8 simplifies to the statement that every type 2 move is a nonelementary type 2 splitting.

Let \hat{s}^v denote the endpoint of b^v opposite s^v and let \hat{v} denote the switch of τ located at \hat{s}^v . In the case that b^v is a sink branch, recalling the definition of L and R crossings of a sink branch, and also recalling the definition of b^L and b^R above, it follows that the interior of b^L intersects the L crossing of b^v and the interior of b^R intersects the R crossing; we define \hat{b}^L , resp. \hat{b}^R , to be the branches on the two-ended side of \hat{s} whose interiors intersect the L, resp. R crossings of b^v .

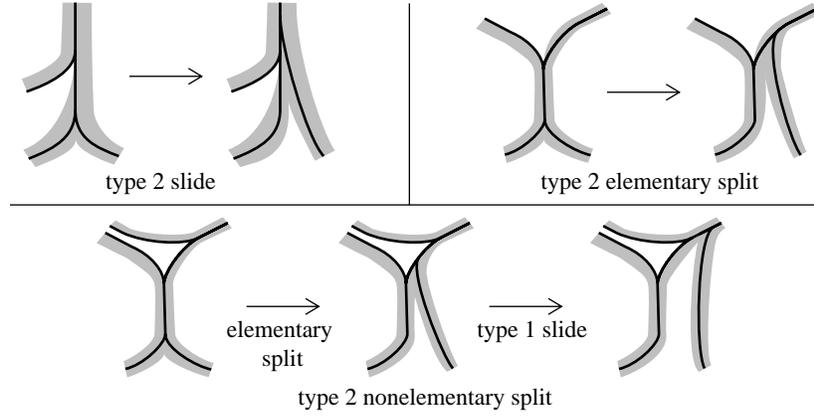


Figure 34: Type 2 moves, all of parity $D = R$. Shading shows sides of ∂^v (\hat{b}^D is shown of type 1, but it could also be of type 2). If b^v is a transition branch then there is a type 2 slide $(\tau, v) \cong (\tau', v')$. If b^v is a sink branch and \hat{b}^D is a type 2 branch then there is a type 2 elementary split $(\tau, v) \succ (\tau', v')$ of parity D . If b^v is a sink branch and \hat{b}^D is a type 1 branch then there is a type 2 nonelementary split $(\tau, v) \succ (\tau', v')$ of parity D .

Proposition 11.1.8. *Consider a type 2 stable train track (τ, v) , so by Proposition 11.1.5 each type 1 branch of (τ, v) is the overlay image of a side of $\bar{\partial}^v$. The type 2 moves $(\tau, v) \succ (\tau', v')$ are completely described as follows:*

- If b^v is a transition branch then $\hat{v} \in \partial(\bar{\partial}^v)$, the branch on the one-ended side of \hat{s} is of type 2, and there is a type 2 slide $(\tau, v) \cong (\tau', v')$.
- If b^v is a sink branch, then for each $D \in \{L, R\}$, if \hat{b}^D is of type 2 then there is a type 2 elementary split $(\tau, v) \succ (\tau', v')$ of parity D , whereas if \hat{b}^D is of type 1 then there is a type 2 nonelementary split $(\tau, v) \succ (\tau', v')$ factoring as an elementary split of parity D followed by a slide of parity D . Moreover:
 - If $\hat{v} \in \partial(\bar{\partial}^v)$ and c^v has 2 cusps then both of \hat{b}^L, \hat{b}^R are of type 2.
 - If $\hat{v} \in \partial(\bar{\partial}^v)$ and c^v has ≥ 3 cusps then one of \hat{b}^L, \hat{b}^R are of type 2 and the other is of type 1.
 - If $\hat{v} \notin \partial(\bar{\partial}^v)$ then both of \hat{b}^L, \hat{b}^R are of type 1.

Proof. To follow the proof, refer to Figure 34. Recall that since (τ, v) is type 2, the branch b^v is type 2.

Suppose that b^v is a transition branch. Since $\hat{v} \neq v$ it follows that $\hat{v} \in \bar{\partial}^v$, but since b^v has type 2 it follows that one of the sides of $\mathcal{C}(S - \tau)$ incident to \hat{v} is in ∂^v , and so $\hat{v} \in \partial(\bar{\partial}^v)$. By stability it follows that the two regular points of $\partial\mathcal{C}(S - \tau)$ mapping to \hat{s} are in ∂^v , and so the branch on the one-ended side of \hat{s} is of type 2. It follows that $(\tau, v) \succ (\tau', v')$ is a type 2 slide.

Suppose next that b^v is a sink branch. Note that the two regular points in the overlay pre-image of \hat{s} are both contained in ∂^v , because they are contained in the same sides, respectively, as the two regular points \bar{s}^L, \bar{s}^R in the overlay pre-image of s^v . If $\hat{v} \in \partial(\bar{\partial}^v)$ and if c^v has exactly 2 cusps then it follows that both sides of $\mathcal{C}(S - \tau)$ incident to \hat{v} are in ∂^v and so both \hat{b}^L, \hat{b}^R are of type 2. If $\hat{v} \in \partial(\bar{\partial}^v)$ and if c^v has ≥ 3 cusps then one side of $\mathcal{C}(S - \tau)$ incident to \hat{v} is in ∂^v and the other is in $\bar{\partial}^v$, and so one of \hat{b}^L, \hat{b}^R is of type 2 and the other is of type 1. If $\hat{v} \notin \partial(\bar{\partial}^v)$ then both sides of $\mathcal{C}(S - \tau)$ incident to \hat{v} are in $\bar{\partial}^v$, and so both \hat{b}^L, \hat{b}^R are of type 1. To complete the proof, given $D \in \{L, R\}$, it suffices to observe that if \hat{b}^D is of type 2 then the one cusp splitting of parity D on (τ, v) results in a type 2 stable train track, whereas if \hat{b}^D is of type 1 then the one cusp splitting of parity D on (τ, v) results in a type 1 stable train track. \diamond

One cusp fold maps. Given an arbitrary one cusp slide or elementary split $(\tau, v) \succ (\tau', v')$ we define a special kind of carrying map $F: \tau' \rightarrow \tau$ called a *one cusp fold* (see Figure 35), defined by the property that for each switch s' of τ' distinct from $s^{v'}$, the image $s = F(s')$ is also a switch, and the map F takes some neighborhood of s' in τ' diffeomorphically to a neighborhood of s in τ . In other words, the only cusp that F folds is the one cusp v' . Note that $F(s^{v'})$ is an interior point of a uniquely determined branch of τ . Existence of a one cusp fold $F: \tau' \rightarrow \tau$ is clear from the definition of slides and elementary splittings as described in Sections 3.12 and 3.13, and moreover F is unique up to homotopy through one cusp folds, by a homotopy which is stationary on all switches of τ' except $s^{v'}$.

Given stable train tracks $(\tau, v), (\tau', v')$ and a one cusp fold $F: \tau' \rightarrow \tau$, we claim that each side b' of $\bar{\partial}^{v'}$ is mapped diffeomorphically onto a side b of $\bar{\partial}^v$, inducing a bijection of sides. To see why this is true, for each neighborhood U of v in $\mathcal{C}(S - \tau)$ there exists a neighborhood U' of v' in $\mathcal{C}(S - \tau')$ such that F lifts to a partial diffeomorphism $\tilde{F}: \mathcal{C}(S - \tau') \rightarrow \mathcal{C}(S - \tau)$, defined on $\mathcal{C}(S - \tau') - U'$ and with image $\mathcal{C}(S - \tau) - U$. Noting that $\bar{\partial}^{v'} \cap U' = \emptyset$ and $\bar{\partial}^v \cap U = \emptyset$, it follows that \tilde{F} induces a diffeomorphism from $\bar{\partial}^{v'}$ to $\bar{\partial}^v$, proving the claim.

Given a one cusp train track (τ, v) define $\text{Br}(\tau, v) = \text{Br}(\tau) - \{b^v\}$. In the stable case, for $i = 1, 2$ define $\text{Br}_i(\tau, v) \subset \text{Br}(\tau, v)$ to be the set of all type i branches of τ except b^v .

Proposition 11.1.9. *Given a one cusp slide or elementary split $(\tau, v) \succ (\tau', v')$,*

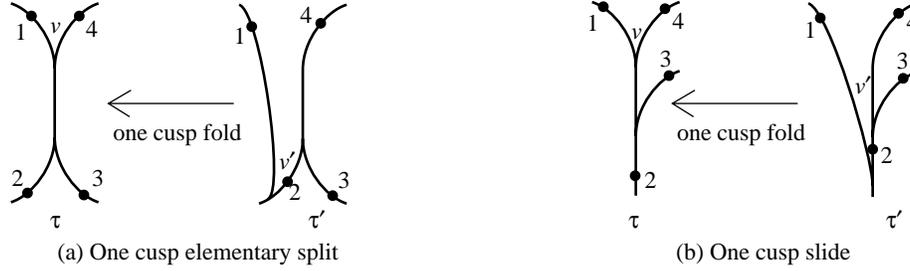


Figure 35: The one cusp fold $F: \tau' \rightarrow \tau$ associated to a one cusp split or slide $(\tau, v) \rightsquigarrow (\tau', v')$ maps the switches of τ' except for $s^{v'}$ bijectively to the switches of τ except for s^v , taking a neighborhood of each switch diffeomorphically to a neighborhood of its image. In the stable case, F induces a bijection between $\partial^{v'}$ and ∂^v , mapping each side of $\partial^{v'}$ diffeomorphically to the corresponding side of ∂^v . Also, there is a unique bijection $\text{Br}(\tau) - \{b^v\} \leftrightarrow \text{Br}(\tau') - \{b^{v'}\}$ with the property that if b corresponds to b' then the image of $\text{int}(b')$ under the fold has nontrivial intersection with $\text{int}(b)$; corresponding branches under this bijection are labelled with the same integer.

there is a unique bijection $\text{Br}(\tau, v) \leftrightarrow \text{Br}(\tau', v')$ with the property that for any one cusp fold $F: \tau' \rightarrow \tau$:

- $b \leftrightarrow b' \implies F(\text{int}(b')) \cap \text{int}(b) \neq \emptyset$.

Assuming moreover that (τ, v) and (τ', v') are stable, it follows that if $b \leftrightarrow b'$ then b, b' have the same type, and so the bijection $\text{Br}(\tau, v) \leftrightarrow \text{Br}(\tau', v')$ restricts to bijections $\text{Br}_i(\tau, v) \leftrightarrow \text{Br}_i(\tau', v')$ for each $i = 1, 2$.

Proof. Define a relation \sim between the sets $\text{Br}(\tau, v)$ and $\text{Br}(\tau', v')$ where $b \sim b'$ if and only if b' maps partly over b , meaning that $F(\text{int}(b')) \cap \text{int}(b) \neq \emptyset$. This relation is well-defined independent of F , because F is unique up to homotopy through one cusp folds. Notice that $b \sim b'$ is *not* a bijection: using the labelling from Figure 35, there are branches $b_1 \neq b_2 \in \text{Br}(\tau, v)$ and $b'_1 \neq b'_2 \in \text{Br}(\tau', v')$ such that $b_1 \sim b'_1$, $b_2 \sim b'_1$, $b_2 \sim b'_2$. However, this is the only deviation from bijection: every branch of $\text{Br}(\tau, v)$ except b_1, b_2 is mapped partly over by a unique branch of $\text{Br}(\tau', v')$, and every branch of $\text{Br}(\tau', v')$ except b'_1, b'_2 maps partly over a unique branch of $\text{Br}(\tau, v)$. It follows that by removing one relator from \sim , namely the relator $b_2 \sim b'_1$, we obtain a uniquely determined bijection \leftrightarrow .

Suppose now that both (τ, v) and (τ', v') are stable. Fix a neighborhood U of v in $\mathcal{C}(S - \tau)$ disjoint from $\bar{\partial}^v$, and choose a neighborhood U' of v' in $\mathcal{C}(S - \tau')$ disjoint from $\bar{\partial}^{v'}$ and a diffeomorphism $\tilde{F}: \mathcal{C}(S - \tau') - U' \rightarrow \mathcal{C}(S - \tau) - U$ that

lifts the fold map F , and so that \tilde{F} induces a diffeomorphism from $\tilde{\partial}^{v'}$ to $\bar{\partial}^v$. Consider corresponding branches $b \leftrightarrow b'$. Notice that with the single exception of the related branches $b_2 \leftrightarrow b'_2$ labelled 2 in Figure 35, one can choose $x \in \text{int}(b)$ and $x' \in \text{int}(b')$ so that $F(x') = x$, and so that both overlay preimages of x are in $\mathcal{C}(S - \tau) - U$ and both overlay preimages of x' are in $\mathcal{C}(S - \tau') - U'$. Since $\tilde{F}: \mathcal{C}(S - \tau') - U' \rightarrow \mathcal{C}(S - \tau) - U$ takes $\tilde{\partial}^{v'}$ diffeomorphically to $\bar{\partial}^v$ it easily follows that b has type 2 if and only if b' has type 2, with the possible exception of the related pair $b_2 \leftrightarrow b'_2$. Choose $x \in \text{int}(b_2)$ and $x' \in \text{int}(b'_2)$ so that $F(x') = x$. There exists one pair of lifts $\tilde{x} \in \partial\mathcal{C}(S - \tau) - U$, $\tilde{x}' \in \partial\mathcal{C}(S - \tau') - U'$ such that $\tilde{F}(\tilde{x}) = \tilde{x}'$, and so $\tilde{x} \in \bar{\partial}^v$ if and only if $\tilde{x}' \in \tilde{\partial}^{v'}$. There is another pair of lifts $\hat{x} \in \partial\mathcal{C}(S - \tau) - U$, $\hat{x}' \in \partial\mathcal{C}(S - \tau) \cap U'$, and so $\tilde{F}(\hat{x}')$ is not defined. However, notice that $\hat{x}' \in \partial^{v'}$. In order to complete the proof we must show that $\hat{x} \in \partial^v$. Arguing by contradiction, suppose that $\hat{x} \in \bar{\partial}^v$. By examining Figure 35 it is clear that there exists $\mathbb{D} \in \{\text{L}, \text{R}\}$ such that the point $\tilde{s}^{\mathbb{D}}$, one of the two the regular overlay pullbacks of s^v , is on the same side of $\mathcal{C}(S - \tau)$ as \hat{x} , and it follows that $\tilde{s}^{\mathbb{D}} \in \bar{\partial}^v$. This shows that (τ, v) is of type 1, and by applying Proposition 11.1.6 it follows that b^v is a transition branch and $(\tau, v) \succ (\tau', v')$ is a slide move. Moreover, letting s_1 be the outflow switch of b^v and b'' the other branch on the two-ended side of s_1 , it is clear from Figure 35 that b'' is on the $\bar{\mathbb{D}}$ side of b^v , and so (τ, v) is of type $1\bar{\mathbb{D}}$ and the slide move $(\tau, v) \succ (\tau', v')$ is of parity $\bar{\mathbb{D}}$. However, applying Proposition 11.1.6 again it follows that $\tilde{s}^{\bar{\mathbb{D}}} \in \bar{\partial}^v$. Since s^v has two lifts $\tilde{s}^{\mathbb{D}}$, $\tilde{s}^{\bar{\mathbb{D}}}$ that are in $\bar{\partial}^v$, we have contradicted the assumption that (τ, v) is stable. \diamond

TO DO:

- Must also consider killing of type 1 branches: they are killed simply when they are split.
- Remark on the difference between killing of type 1 and type 2 branches: type 1 are killed when they are split, type 2 are killed when they do the splitting. Why is this difference? I do not know, go ask your dad.

One cusp survival and descent. For any sequence of one cusp splittings $(\tau_0, v_0) \succ \dots \succ (\tau_n, v_n)$, by composing the bijections $\text{Br}(\tau_0, v_0) \leftrightarrow \dots \leftrightarrow \text{Br}(\tau_n, v_n)$ we obtain a bijection denoted with the same symbol $\text{Br}(\tau_0, v_0) \leftrightarrow \text{Br}(\tau_n, v_n)$.

We shall define killing, survival, and descent for type 2 moves between stable train tracks, rather than for elementary one cusp moves.

Consider a type 2 stable train track (τ, v) and a type 2 move $(\tau, v) \succ (\tau', v')$ of parity $\mathbb{D} \in \{\text{L}, \text{R}\}$. There are two branches of $\text{Br}(\tau, v)$ which are killed by this move: in Figures 35(a,b), they are the two branches of τ labelled 1 and 2. The

where is parity of a slide move discussed? That might be the good place to describe the two exceptional transition branches

branch labelled 1 is just $b^{\bar{D}}$. The branch labelled 2 is \hat{b}^D , using the notation from before the statement of Proposition 11.1.8, and it is characterized uniquely by the following properties: \hat{b}^D is incident to the endpoint \hat{s}^v of b^v opposite v , and if $b^{\bar{D}} \rightsquigarrow b' \in \text{Br}(\tau', v')$ then $\hat{b}^D \sim b'$, that is, b' maps partly over \hat{b}^D via the one cusp fold $\tau' \rightarrow \tau$. The two branches $b^{\bar{D}}, \hat{b}^D \in \text{Br}(\tau, v)$ each depend naturally on the parity D . Every branch $b \in \text{Br}(\tau, v) - \{b^{\bar{D}}, \hat{b}^D\}$ survives, and its descendant is the unique branch $b' \in \text{Br}(\tau', v')$ such that $b \rightsquigarrow b'$. Note that $b^{\bar{D}}$ is a type 2 branch of τ . When $(\tau, v) \succ (\tau', v')$ is a type 2 nonelementary split, then \hat{b}^D is of type 1; whereas when $(\tau, v) \succ (\tau', v')$ is a type 2 slide or elementary split, the branch \hat{b}^D is one of the two transition branches of type 2, contained in ∂C^v and incident to $\partial(\bar{\partial}^v)$.

Given a sequence of type 2 moves on stable train tracks $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$, we define a new killing criterion:

explain these two exceptional transition branches more clearly earlier

The branch killing criterion: For every I , each branch $b \in \text{Br}(\tau_I, v_I)$ is eventually killed in the following sense: there exists $J \geq I$ and a sequence $b = b_I, \dots, b_J$, with $b_i \in \text{Br}(\tau_i, v_i)$ for $I \leq i \leq J$, such that for $I < i \leq J$ the branch b_{i-1} survives the splitting $(\tau_{i-1}, v_{i-1}) \succ (\tau_i, v_i)$ with descendant b_i , but b_J is killed by the splitting $(\tau_J, v_J) \succ (\tau_{J+1}, v_{J+1})$.

Theorem 11.1.10. *If $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ is a sequence of type 2 moves on stable train tracks, then the sequence satisfies the canonical killing criterion if and only if it satisfies the branch killing criterion.*

Proof. Let $D_i \in \{L, R\}$ be the parity of the move $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$. The branches killed by this move are $b_i^{D_i}$ and $\hat{b}_i^{D_i}$. Let $F_i: \tau_{i+1} \rightarrow \tau_i$ denote the fold map which is either the one cusp fold map when $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ is a type 2 slide or elementary splitting, or is the concatenation of two one cusp fold maps when $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ is a type 2 nonelementary splitting.

Let us fix an indexing $B(\tau_0, v_0) = \{b_0^1, \dots, b_0^J\}$, and by induction we extend this to an indexing $B(\tau_i, v_i) = \{b_i^1, \dots, b_i^J\}$ so that for each $j = 1, \dots, J$ and each $i \geq 0$ we have $b_i^j \rightsquigarrow b_{i+1}^j$. In each train track τ_i , the only unindexed branch is the branch b^{v_i} on the one-ended side of the switch s_{v_i} . The relation \rightsquigarrow respects the classification of branches into type 1 and 2 by Proposition 11.1.9, and so we can choose the indexing and choose $K \geq J$ so that for all $i \geq 0$, $B_1(\tau_i, v_i) = \{b_i^1, \dots, b_i^K\}$ and $B_2(\tau_i, v_i) = \{b_i^{K+1}, \dots, b_i^J\}$. **MAYBE SAY MORE HERE: THE RELATION \rightsquigarrow RESPECTS TYPE 1, TYPE 2, AND THE EXCEPTIONAL TRANSITION BRANCHES. BUT THIS HAS TO BE STATED WITH THE AUGMENTED BRANCHES, WHERE b^v IS AUGMENTED ONTO ANY BRANCH WITH ENDPOINT s^v . WAIT, THIS AUGMENTED BIT, IT'S NOT TRUE FOR THE KILLED BRANCH!!!! THAT'S PART OF THE HEART OF THE ARGUMENT, ISN'T IT?**

Branch killing criterion \implies proper subtrack killing criterion. Suppose that the branch killing criterion is satisfied, but there exists $I \geq 0$ and a proper, nonempty subtrack $\sigma_I \subset \tau_I$ that survives forever with descendant $\sigma_i \subset \tau_i$ for $i \geq I$; from this we shall derive a contradiction. The fold map $F_i: \tau_{i+1} \rightarrow \tau_i$ restricts to a homotopic carrying map $\sigma_{i+1} \rightarrow \sigma_i$.

For $i \geq I$ let $N_i = N_i(\sigma_i)$ be the number of branches of $B(\tau_i, v_i)$ that are contained in σ_i . We claim that for each $i \geq I$ and $j = 1, \dots, K$, if $b_{i+1}^j \subset \sigma_{i+1}$ then $b_i^j \subset \sigma_i$; this follows because $F_i(b_{i+1}^j) \subset \sigma_i$, but $F_i(b_{i+1}^j)$ contains a point in the interior of b_i^j and so $b_i^j \subset \sigma_i$. As a consequence of the claim, the sequence N_i is nonincreasing, and so by increasing I if necessary the sequence N_i becomes constant for $i \geq I$, and moreover if $i \geq I$ and $j = 1, \dots, K$, then $b_i^j \subset \sigma_i$ if and only if $b_{i+1}^j \subset \sigma_{i+1}$.

Now we break into two cases, depending on whether or not σ_I contains every type 2 branch of τ_I .

Case 1: σ_I contains every type 2 branch of τ_I . It follows that for all $i \geq I$, σ_i contains every type 2 branch of τ_i . Since σ_I is proper, there exists a type 1 branch b_I^j of τ_I such that $b_I^j \not\subset \sigma_I$. It follows that for all $i \geq I$, b_i^j is a type 1 branch of τ_i such that $b_i^j \not\subset \sigma_i$. By the branch killing criterion, there exists $i \geq I$ such that $b_i^j = \hat{b}_i^{\text{D}_i}$ is killed by the move $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$. Since $\hat{b}_i^{\text{D}_i}$ is of type 1, this move is a type 2 nonelementary splitting. Let $b_i^{j'} = b_i^{\bar{\text{D}}_i}$ be the other branch of τ_i which is killed by the splitting, a type 2 branch. It follows that $b_{i+1}^{j'}$ is a type 2 branch of τ_{i+1} and so $b_{i+1}^{j'} \subset \sigma_{i+1}$. However, $F_i(b_{i+1}^{j'}) \supset \hat{b}_i^{\text{D}_i}$, by the characterization of the branch $\hat{b}_i^{\text{D}_i}$ killed by the parity D_i , type 2 move $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$. This shows that $b_i^j = \hat{b}_i^{\text{D}_i} \subset F_i(\sigma_{i+1}) = \sigma_i$, a contradiction.

Case 2: σ_I does not contain every type 2 branch of τ_I . Nevertheless σ_I , like any subtrack of τ_I , does contain at least one type 2 branch of τ_I .

The branch killing criterion implies that for each $j = K+1, \dots, J$ there exist infinitely many i such that b_i^j is killed by $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$: if this were not true then for some $I' \geq I$ the branch $b_{I'}^j$ would survive forever with descendant $b_i^j \subset \tau_i$, $i \geq I'$. Also, we have just shown that for each j , the statement $b_i^j \subset \sigma_i$ is either true for all $i \geq I$ or false for all $i \geq I$. It follows that:

- There exists $i \geq I$ and $j \neq j' \in \{K+1, \dots, J\}$, such that the branch $b_i^{\bar{\text{D}}_i} = b_i^j$, which is killed by the move $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$, satisfies $b_i^j \subset \sigma_i$, but the branch $b_{i+1}^{\bar{\text{D}}_{i+1}} = b_{i+1}^{j'}$, which is killed by the move $(\tau_{i+1}, v_{i+1}) \succ (\tau_{i+2}, v_{i+2})$, satisfies $b_{i+1}^{j'} \not\subset \sigma_{i+1}$.

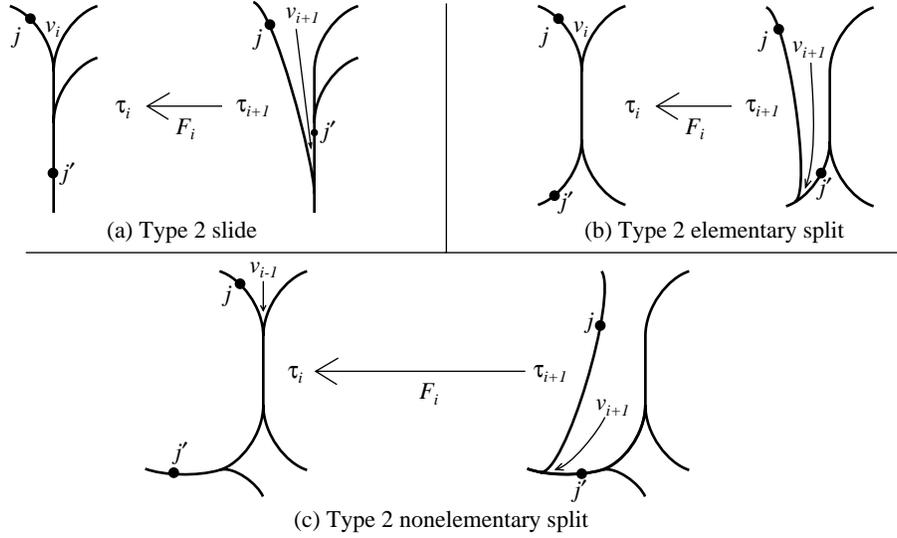


Figure 36: A subtrack σ_i killed by a move $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ of parity D; the figure shows $D = R$. As argued in the text, if σ_i survives with descendant σ_{i+1} we may reduce to the case that the branch $b_i^L = b_i^j \subset \sigma_i$, whereas $b_{i+1}^R = b_{i+1}^{j'} \not\subset \sigma_{i+1}$. By stabilizing the cardinality of $B(\tau_i, v_i) \cap \sigma_i$, it follows that $b_{i+1}^j \subset \sigma_{i+1}$ and $b_i^{j'} \not\subset \sigma_i$. This contradicts that b_{i+1}^j contains a point that maps to an interior point of $b_i^{j'}$.

It follows that $b_{i+1}^j \subset \sigma_{i+1}$ and $b_i^{j'} \not\subset \sigma_i$. Referring to Figure 36, the fold map $F_i: \tau_{i+1} \rightarrow \tau_i$ maps an interior point of b_{i+1}^j to an interior point of $b_i^{j'}$, implying that $b_i^{j'} \subset F_i(\sigma_{i+1}) = \sigma_i$, a contradiction.

This completes the proof that the branch killing criterion implies the subtrack killing criterion.

Branch killing criterion \implies c-splitting arc killing criterion. Suppose that the branch killing criterion is satisfied but there exists $I \geq 0$ and a c-splitting arc $\alpha_I \subset \tau_I$ which survives forever with descendant $\alpha_i \subset \tau_i$, $i \geq I$. In this case although it is evidently true that $F_i(\alpha_{i+1}) \subset \alpha_i$, the reverse inclusion does not necessarily hold, in fact it fails precisely when s^{v_i} is an endpoint of α_i . Nevertheless, the inclusion $F_i(\alpha_{i+1}) \subset \alpha_i$ is sufficient to prove, by the same method as above, that for each $j = 1, \dots, J$, if $b_{i+1}^j \subset \alpha_{i+1}$ then $b_i^j \subset \alpha_i$. The number $N_i = N_i(\alpha_i)$ of branches in $B(\tau_i, v_i)$ that are contained in α_i is therefore a nonincreasing sequence, and so we may increase I if necessary so that N_i is constant when $i \geq I$. It follows that for $i \geq I$ and $j = 1, \dots, J$, we have $b_i^j \subset \alpha_i$ if and only if $b_{i+1}^j \subset \alpha_{i+1}$.

Note that α_I cannot contain every type 2 branch of τ_I ; for example, both of the branches b_I^L, b_I^R are of type 2, but they cannot both be contained in τ_I . Arguing as in case 2 above, we may choose $j, j' \in \{K+1, \dots, J\}$ so that the branch $b^{\bar{D}_i} = b_i^j$, which is killed by the splitting $\tau_i \succ \tau_{i+1}$, is contained in α_i , but the branch $b^{\bar{D}_{i+1}} = b_{i+1}^{j'}$, which is killed by the splitting $\tau_{i+1} \succ \tau_{i+2}$, is *not* contained in α_{i+1} . It follows that $b_{i+1}^j \subset \alpha_{i+1}$ and $b_i^{j'} \not\subset \alpha_i$. As before, F_i takes some point of b_{i+1}^j to an interior point of $b_i^{j'}$, leading to the contradiction that $b_i^{j'} \subset \alpha_i$.

This completes the proof of one direction of Theorem 11.1.10, that the branch killing criterion implies the canonical killing criterion.

Canonical killing criterion \implies branch killing criterion. Assume that the branch killing criterion fails, so there exists $I \geq 0$ and $j = 1, \dots, K$ such that for each $i \geq I$ the branch $b_i^j \in B(\tau_i, v_i)$ survives $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ with descendant $b_{i+1}^j \in B(\tau_{i+1}, v_{i+1})$. Letting D_i be the parity of $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$, we have $b_i^j \neq b_i^{D_i}$. We shall produce either a proper, nonempty subtrack or a C-splitting arc of τ_I which survives forever.

Define a train path α_i as follows:

$$\alpha_i = \begin{cases} b_i^j & \text{if } b_i^j \neq b_i^{D_i} \\ b_i^j * b^{v_i} & \text{if } b_i^j = b_i^{D_i} \end{cases}$$

In the second case the concatenation takes place at the switch s^{v_i} . In particular, the train path α_i traverses either one or two branches, and in the latter case the interior switch of α_i is s^{v_i} . The train path α_i is an embedding except possibly that its endpoints are identified. Neither end of α_i is located at the switch s^{v_i} , for if $b_i^j \neq b_i^{D_i}$ then $\alpha_i = b_i^j$ has no end located at s^{v_i} , and if $b_i^j = b_i^{D_i}$ then $\alpha_i = b_i^j * b^{v_i}$ has s^{v_i} as an interior point.

The proof is an extremely delicate sifting of cases, some of which lead to contradictions of the canonical killing criterion, others of which lead to various other contradictions. The top level of the case analysis depends on properties of α_I . Each endpoint of the train path α_I is either an inflow or an outflow, and so by mimicking the classification of branches we can classify α_I as either a sink path (two inflows), a transition path (an inflow and an outflow), or a source path (two outflows). If α_I is a transition path or source path then its endpoints might be identified, but if α_I is a sink path then it is embedded.

Case 1: α_I is a sink path. In other words, α_I is a C-splitting arc of τ_i . We prove that α_I survives forever, showing by induction that for $i \geq I$ the path α_i is a C-splitting arc that survives the splitting $\tau_i \succ \tau_{i+1}$ with descendant α_{i+1} . Suppose

by induction that α_i is a c-splitting arc. There are two subcases, depending on whether $\alpha_i = b_i^j \neq b_i^{D_i}$ or $\alpha_i = b_i^{D_i} * b^{v_i}$. In each of these subcases, α_i survives: if $\alpha_i = b_i^j \neq b_i^{D_i}$ then $s^{v_i} \notin \partial\alpha_i$ which implies that α_i does not contain a partial \bar{D}_i crossing of b^{v_i} ; and if $\alpha_i = b_i^{D_i} * b^{v_i}$ then it is immediately obvious that α_i does not contain a partial \bar{D}_i crossing of b^{v_i} . In either case, $s^{v_i} \notin \partial\alpha_i$. Letting $\alpha' \subset \tau_{i+1}$ be the descendant of α_i , we must show that $\alpha' = \alpha_{i+1}$. Note that $s^{v_{i+1}} \notin \partial\alpha'$, because the cusps at the endpoints of α' correspond to the cusps at the endpoints of α_i , and $s^{v_i} \notin \partial\alpha_i$. Since α_i contains exactly one branch in $\text{Br}(\tau_i, v_i)$, namely b_i^j , we have $N_i(\alpha_i) = 1$. By the earlier discussion, in the proof of “Branch killing criterion \implies c-splitting arc killing criterion”, since α_i survives with descendant α' we have $N_i(\alpha_i) \geq N_{i+1}(\alpha')$, and $F_i(\alpha') = \alpha_i$. Since $s^{v_{i+1}} \notin \partial\alpha'$ it follows that $\alpha' \neq b^{v_{i+1}}$ and so $N_{i+1}(\alpha') \geq 1$. This proves that $N_{i+1}(\alpha') = 1$ and again by the earlier discussion we have $b_{i+1}^j \subset \alpha'$. The set $\text{Cl}(\alpha' - b_{i+1}^j)$ is either empty or it consists of the arc $b^{v_{i+1}}$. If $\text{Cl}(\alpha' - b_{i+1}^j) = \emptyset$ then $\alpha' = b_{i+1}^j$ and, knowing that $s^{v_{i+1}} \notin \partial\alpha'$, we conclude that $\alpha_{i+1} = b_{i+1}^j = \alpha'$. If $\text{Cl}(\alpha' - b_{i+1}^j) = b^{v_{i+1}}$ then the endpoint of $b^{v_{i+1}}$ in the interior of α' must be $s^{v_{i+1}}$, because we know that $s^{v_{i+1}} \notin \partial\alpha'$, and it follows that $\alpha_{i+1} = b_{i+1}^j * b^{v_{i+1}} = \alpha'$. This proves that α_I survives forever and so the canonical killing criterion fails.

Case 2: α_I is a source path. We prove by induction that for each $i \geq I$, $\alpha_i = b_i^j$ is a type 1 branch identified with the overlay image of a side of $\bar{\partial}^{v_i} = \bar{\partial}^{v_i} \mathcal{C}(S - \tau_i)$, which is then used to contradict the canonical killing criterion. In the induction, most of the work goes into the basis step, proving that $\alpha_I = b_I^j$ is the overlay image of a side of $\bar{\partial}^{v_I}$, and to do this there are several subcases to rule out by various contradictions. The basis step of the induction is established, and the inductive step proved, in case 2Bii, which then breaks into further subcases each producing either a violation to the canonical killing criterion or some other contradiction.

Recall the definition of the types of source branches: LL, RR, or LR. The type of the source path α_I is similarly defined: for each end η of α_I there is a parity $D_\eta \in \{L, R\}$ such that η is on the D_η side of the cusp incident to η ; the type of α_I is $D_\eta, D_{\eta'}$ where η, η' are the two ends of α_I .

We need some notation; see Figure 37. Let v, v' be the cusps at the two ends of α_I , located at switches $s^v, s^{v'}$. Let $b^v, b^{v'}$ be the branches on the one-ended sides of $s^v, s^{v'}$, respectively. Let $a^v, a^{v'}$ be the branches with interiors disjoint from α_I on the two-ended sides of $s^v, s^{v'}$, respectively. Since $s^v, s^{v'} \neq s^{v_I}$ it follows that v, v' , regarded as cusps of $\mathcal{C}(S - \tau_I)$, are in $\bar{\partial}^{v_I}$. This implies that $b^v, b^{v'}$ are type 2 branches.

Case 2A: α_I is a source path of type LL or RR. See Figure 37(A). From the hypothesis of Case 2A, the cusps v, v' are on opposite sides of α_I , as are the branches $\alpha^v, \alpha^{v'}$, and moreover $v \neq v'$. If $\alpha_I = b_I^j$, that is if α_I contains no switch, it also follows that α_I is a type 2 branch; on the other hand, if $\alpha_I = b^{D_i} * b^{v_I}$ then, since (τ_I, v_I) is a type 2 stable train track, it follows that each of b^{D_i}, b^{v_I} are type 2 branches. In either case it follows that v, v' are each in $\partial(\bar{\partial}^{v_I})$, and since $v \neq v'$ it follows that $\bar{\partial}^{v_I}$ is a topological arc with endpoints v, v' and that $\alpha^v, \alpha^{v'}$ are type 1 branches. In this case we derive a contradiction. We first remark that in this case α_I is embedded, with its endpoints at distinct switches. Assuming that α_I is of type LL, it follows that, in the boundary orientation that $\partial\mathcal{C}(S - \tau_I)$ inherits from the oriented surface $\mathcal{C}(S - \tau_I)$, each of v, v' are the initial endpoints of the oriented topological arc $\bar{\partial}^{v_I}$. In other words, $v = v'$, a contradiction. Under the assumption that α_I is of type RR, it follows that v, v' are the terminal endpoints of $\bar{\partial}^{v_I}$, reaching the same contradiction that $v = v'$.

Case 2B: α_I is a source path of type LR. There are two subcases, depending on whether $\alpha_I = b_I^{D_I} \cup b^{v_I}$ or $\alpha_I = b_I^j \neq b_I^{D_I}$.

Case 2Bi: $\alpha_I = b_I^{D_I} \cup b^{v_I}$. In this case the cusps v, v' and the branches $\alpha^v, \alpha^{v'}$ are all on the same side of α_I . The cusp v_I can be on either side of α_I , and so we break into two further subcases.

Case 2Bia: v_I is on the side of α_I opposite the cusps v, v' . See Figure 37(Bia). The fact that (τ_I, v_I) is of type 2 implies that both of the branches $b_I^{D_I}, b^{v_I}$ are of type 2. However, since v_I is on the side of α_I opposite v, v' it follows that α_I is the image of a side of $\partial\mathcal{C}(S - \tau_I)$ denoted $\tilde{\alpha}$ whose endpoints are the two cusps v, v' , and since $v, v' \in \bar{\partial}^{v_I}$ it follows that $\tilde{\alpha} \subset \bar{\partial}^{v_I}$. This implies that the two branches $b_I^{D_I}, b^{v_I}$ comprising α_I are of type 1, a contradiction.

Case 2Bib: v_I is on the same side of α_I as the cusps v, v' . See Figure 37(Bib). Notice that b^{v_I} is a transition branch and the type 2 move $(\tau_I, v_I) \succ (\tau_{I+1}, v_{I+1})$ is a type 2 slide. Choosing the notation so that v' is the cusp at the end of b^{v_I} opposite v_I , since $b_I^{D_I}$ is on the D_I side of v_I it follows that $\alpha^{v'}$ is on the \bar{D}_I side of v' , from which it follows that the transition branch b^{v_I} has parity \bar{D}_I and so the slide $(\tau_I, v_I) \succ (\tau_{I+1}, v_{I+1})$ has parity \bar{D}_I and the branch $b_I^{D_I}$ is killed, but we have assumed that the parity is D_I and $b_I^{D_I}$ survives, a contradiction. Figure 37(Bib) shows the case $D_I = R$, but the slide evidently has parity L.

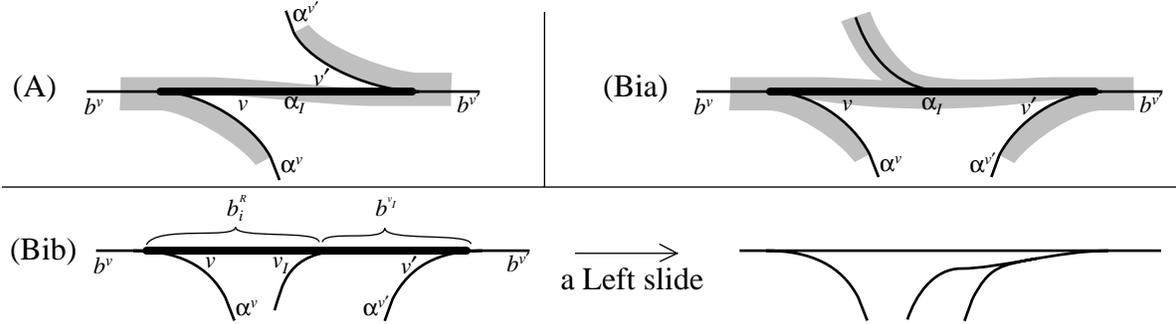


Figure 37: α_I is a source path of type LL. In (A) and (Bia), shading indicates ∂^{v_I} , no shading indicates $\bar{\partial}^{v_I}$. In (A), s^{v_I} may be an interior point of α_I , with incident branch b_i^R . In (A) the contradiction is that $v = v' \in \partial(\bar{\partial}^{v_I})$, but evidently $v \neq v'$. In (Bia) the contradiction is that $b_I^{D_I}$ has type 1, contradicting that (τ_I, v_I) has type 2. In (Bib) the contradiction is that the slide $(\tau_I, v_I) \succ (\tau_{I+1}, v_{I+1})$ has parity \bar{D}_I , but we have assumed that the parity of $(\tau_I, v_I) \succ (\tau_{I+1}, v_{I+1})$ is D_I ; the diagram shows $D_I = R$.

Case 2Bii: $\alpha_I = b_I^j$. It follows that b_I^j is a type 1 branch, and is the image of a side of $\bar{\partial}^{v_I}$, completing the basis step of the induction.

We are at last ready for the induction step: assume that $\alpha_i = b_i^j$ is a type 1 branch identified with a side of $\bar{\partial}^{v_i}$. Consider the parity D_I type 2 move $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$. We know that b_i^j is not killed by this move, and so $b_i^j \neq \hat{b}^{D_i}$. This implies that $F_i: \tau_{i+1} \rightarrow \tau_i$ maps b_{i+1}^j bijectively to b_i^j , and so b_{i+1}^j is also a type 1 branch identified with a side of $\bar{\partial}^{v_{i+1}}$, implying that $\alpha_{i+1} = b_{i+1}^j$. This completes the induction.

Notice in the course of the induction we have shown that $F_i^{-1}(b_i^j) = b_{i+1}^j$. It follows that if we define $\sigma_i = \tau_i - \text{int}(b_i^j)$ then $F_i(\sigma_{i+1}) = \sigma_i$.

Let c_i be the component of $\mathcal{C}(S - \tau_i)$ which contains the side of $\bar{\partial}^{v_i}$ mapping to b_i^j . The one cusp fold $F_i: \tau_{i+1} \rightarrow \tau_i$ lifts to a cusp preserving diffeomorphism $c_{i+1} \cap \bar{\partial}^{v_{i+1}} \rightarrow c_i \cap \bar{\partial}^{v_i}$. Let $n \geq 1$ be the number of cusps of c_i , for all $i \geq I$.

Case 2Biiia: $n \geq 2$. The endpoints of b_i^j are therefore at distinct switches for all $i \geq I$, and so σ_i is a subtrack of τ_i . Moreover, clearly σ_i survives the move $(\tau_i, v_i) \succ (\tau_{i+1}, v_{i+1})$ with descendant σ_{i+1} . Thus, σ_I is a proper, nonempty subtrack of τ_I that survives forever, contradicting the canonical killing criterion.

Case 2Biiib: $n = 1$. The branch b_i^j is the overlay image of the boundary of a one-cusped monogon component of $\mathcal{C}(S - \tau_i)$, and in this case σ_i is unfortunately

not a subtrack of τ_i , so we must find another way to show that the canonical killing criterion fails.

Let s_i be the switch at which both ends of b_i^j are located. Let b'_i be the branch on the one-ended side of s_i , a type 2 branch. Notice that $b'_i \neq b^{v_I}$, because if $b'_i = b^{v_I}$ then $b_i^j = \hat{b}_i^L = \hat{b}_i^R$ and so b_i^j is killed by the splitting $\tau_i \succ \tau_{i+1}$ regardless of the parity of this splitting. Notice also that $b'_i \rightsquigarrow b'_{i+1}$ for all $i \geq I$, and so we may choose $j' \in \{1, \dots, J\}$ so that $b'_i = b_i^{j'}$ for all $i \geq I$.

We now break into further subcases depending on the nature of b'_I .

Case 2Biibx: b'_I is a sink branch. In other words, b'_I is a C-splitting arc, and in this case we show that each $b_i^{j'}$ is a C-splitting arc and that $b_i^{j'}$ survives forever with descendant $b_i^{j'}$ in τ_i , contradicting the canonical killing criterion.

Assuming by induction that $b_i^{j'}$ is a sink branch, since $b_i^{j'} \neq b^{v_I}$ we can now argue as in Case 1 that $b_i^{j'}$, regarded as a C-splitting arc, survives the splitting $\tau_i \succ \tau_{i+1}$. Let b'' be its descendant in τ_{i+1} , a C-splitting arc. Clearly $b'' \supset b_{i+1}^{j'}$, and we must prove that $b'' = b_{i+1}^{j'}$. The only way this could fail is if the switch contained in the interior of b'' is $s^{v_{i+1}}$, but this implies that $b_{i+1}^{j'} = b^{v_{i+1}}$ which we have already seen is impossible.

Case 2Biiby: b'_I is a transition branch, with switch $s'_I \neq s^{v_I}$ opposite s_I . A similar argument as in Case 2Biibx proves that for all $i \geq I$, $b_i^{j'}$ is a transition branch with switch $s'_i \neq s^{v_i}$ opposite s_i , and moreover $F_i^{-1}(b_i^{j'} \cup b_i^j) = b_{i+1}^{j'} \cup b_{i+1}^j$. We may now define a proper, nonempty subtrack $\sigma_i = \text{Cl}(\tau_i - (b_i^{j'} \cup b_i^j))$, and clearly σ_I survives forever with descendant $\sigma_i \subset \tau_i$, contradicting the canonical killing criterion.

Case 2Biibz: b'_I is a transition branch, with switch s^{v_I} opposite s_I . It is possible in this case that b'_I is killed, and if this happens then b'_{I+1} is also a transition branch with switch $s^{v_{I+1}}$ opposite s_{I+1} . If this continues infinitely, then $\sigma_i = \text{Cl}(\tau_i - (b_i^{j'} \cup b_i^j))$ is a subtrack of τ_i , and σ_I survives forever with descendant $\sigma_i \subset \tau_i$, a contradiction.

Otherwise, there exists $I' \geq I$ such that $b_{I'}^{j'}$ is a transition branch with switch $s^{v_{I'}}$ opposite $s_{I'}$, but $b_{I'}^{j'}$ is not killed. If $b^{v_{I'}}$ is a sink branch it follows that $b_{I'+1}^{j'}$ is a sink branch that falls into case 2Biibx, and if $b^{v_{I'}}$ is a transition branch it follows that $b_{I'+1}^{j'}$ is a transition branch that falls into case 2Biiby, both of which lead to contradictions.

Change notation: s_i^* for the distinguished switch, v_i^* for the distinguished cusp. Use * in general for these objects, starting from the beginning of the one cusp theory, and use it in diagrams too.

Case 3: α_I is a transition path. The first step is showing that each α_i is a transition path; this is accomplished in Lemma 11.1.11 below. Afterwards we use this fact to show that the Canonical Killing Criterion fails.

Assuming that α_i is indeed a transition path, we classify α_i into one of three types as follows. Let s_i^{in} be the inflowing switch and s_i^{out} the outflowing switch of α_i . The path α_i has a *flow orientation* uniquely characterized by agreeing with the switch orientations at s_i^{in} and s_i^{out} . Let v_i^{in} , etc., denote the cusps at the corresponding switches. There are three types of α_i to keep in mind: type T for transition; type TT for transition–transition; and type SS for source–sink.

Type T: $s^{v_i} \notin \alpha_i$. In this case $\alpha_i = b_i^j$ is a transition branch.

Type TT: $s^{v_i} \in \alpha_i$ with switch orientation agreeing with the flow orientation. In this case $\alpha_i = b_i^{\text{D}i} * b^{v_i}$, $b_i^{\text{D}i}$ is a transition branch with ends at s_i^{in} and s^{v_i} , and b^{v_i} is a transition branch with ends at s^{v_i} and s_i^{out} .

Type SS: $s^{v_i} \in \alpha_i$ with switch orientation opposite the flow orientation. In this case $\alpha_i = b_i^{\text{D}i} * b^{v_i}$, $b_i^{\text{D}i}$ is a source branch with ends at s_i^{out} and s^{v_i} , and b^{v_i} is a sink branch with ends at s^{v_i} and s_i^{in} .

Before proceeding, we first claim that, after increasing I if necessary, each τ_i is recurrent for $i \geq I$. To see why, supposing that some τ_j is nonrecurrent, then τ_i is nonrecurrent for all $i \geq j$, and letting σ_i be the maximal recurrent subtrack, it follows that σ_i is a proper, nonempty subtrack of τ_i for $i \geq j$. Moreover, $\sigma_i \succ \sigma_{i+1}$ for all $i \geq j$, and hence this carrying is eventually homotopic, and so some σ_i survives forever, contradicting the canonical carrying criterion. By increasing I we may henceforth assume that each τ_i is recurrent for $i \geq I$.

Now we analyze the paths α_i for all $i \geq I$.

Lemma 11.1.11. *After increasing I by at most 1, the following holds. For each $i \geq I$, α_i is an embedded transition path, and if its type is TT or SS then the cusps v_i and v_i^{out} are on opposite sides of α_i . Moreover, α_i is a union of type 2 branches.*

Proof. The proof is illustrated in Figure 38, broken into cases (a–h).

First we show that if α_i is a transition path then it is embedded. For type T (case a) and type TT (case b) this follows from recurrence of τ_i . For type SS (case c), suppose to the contrary that α_i is not an embedded train path, and so it must be a smooth closed curve embedded in τ_i , regarded as a subtrack of τ_i . Since b_i^j is not killed, it follows that the subtrack α_i is not killed, and its descendant in τ_{i+1} is also an embedded smooth closed curve, but that curve contains only one switch and hence is a sink loop, contradicting recurrence of τ_{i+1} .

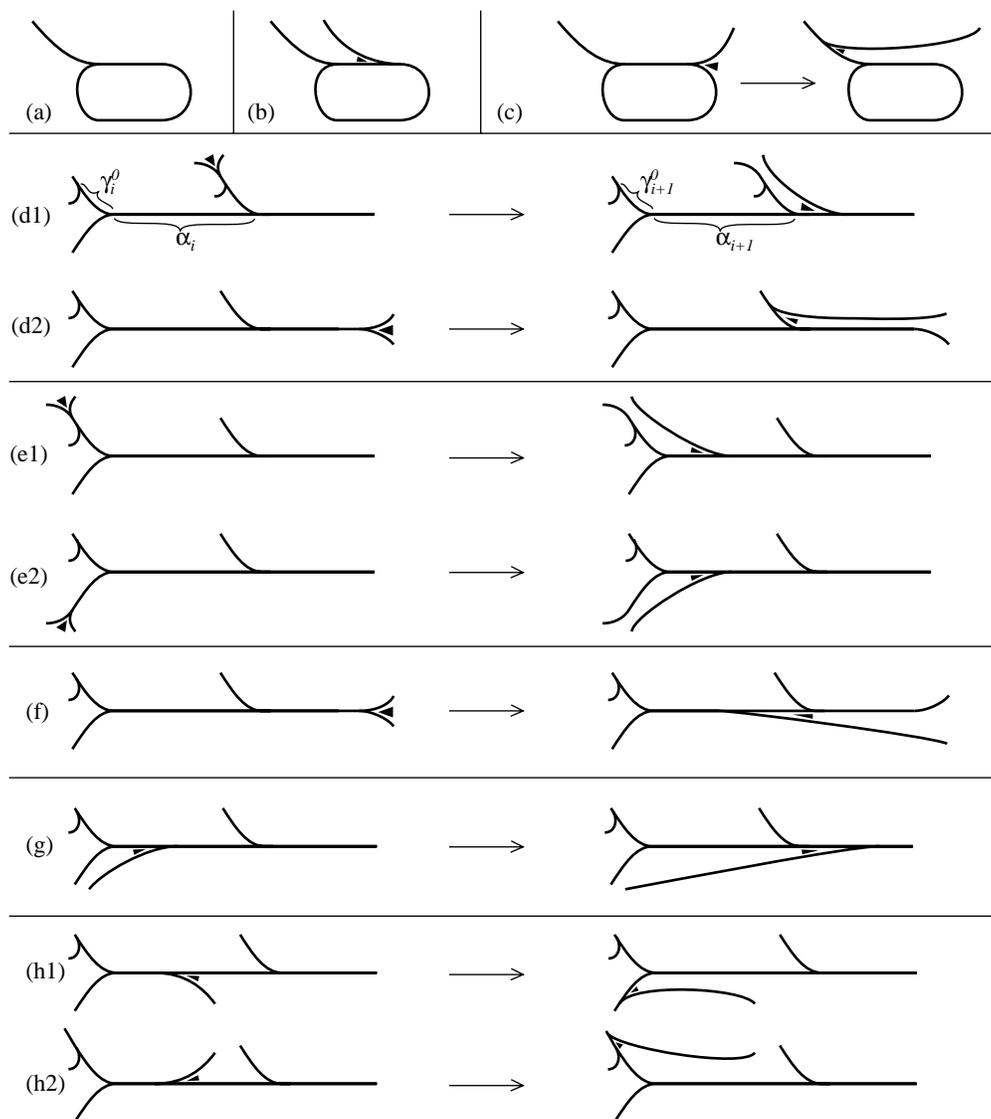


Figure 38: If α_i is a transition path then so is α_{i+1} . α_i must be embedded: if not, then in types T and TT τ_i is not recurrent (a,b), and in type SS τ_{i+1} is not recurrent (c). If α_i is of type T then α_{i+1} either (d) of type T, or (e) of type TT, or (f) of type SS. If α_i is of type TT, then α_{i+1} of type T (g). If α_i is of type SS, then α_{i+1} is of type T (h). The source branch γ_i^0 survives in all cases except (e1) and (h2). The paths α_i, α_{i+1} and the adjacent source branches $\gamma_i^0, \gamma_{i+1}^0$, are labelled explicitly in (d1), and are left implicit in (d2)–(h2).

Next we show by induction that each α_i is an embedded transition path for $i \geq I$. Assuming this is so for α_i , we analyze the three possible types T, TT, SS for α_i in separate cases.

Cases (d–f): α_i is of type T: Let U be a regular neighborhood of α_i in τ_i . Note that U contains four branch ends incident to α_i : both ends on the two-ended side of v^{in} , the end on the one-ended side of v^{out} , and the end disjoint from α_i on the two-ended side of v^{out} .

One case occurs when F_i is injective over U (meaning that F_i is injective on $F_i^{-1}(U)$), in which case $F_i^{-1}(\alpha_i) = \alpha_{i+1}$ is a transition path of type T (not shown in Figure 38).

By examining the diagrams for type 2 moves in Figure 34, one sees that if F_i is not injective over U , then F_i fails to be injective on some subset E of the ends incident to α_i , and we can make a case analysis based on this subset. If E is the pair of ends incident to v^{out} , then α_{i+1} is a transition path of type T (case (d)); there are two subcases, (d1) and (d2), depending on which direction the the cusp v_i points. If E is a single end on the two-ended side of v^{in} then α_{i+1} is a transition path of type TT (case (e)); there are two subcases, (e1) and (e2), one for each of the two ends on the two-ended side of v^{in} . Note that in case (e1), the train track τ_{i+1} has the property that v_{i+1} and v_{i+1}^{out} are on the same side of the transition path α_{i+1} ; see case (g) below for the proof that this cannot happen. If E is the single end on the one-ended side of v^{out} , then α_{i+1} is a transition path of type SS (case (f)). These are all of the possible cases: if F_i fails to be injective over two or more ends incident to α_i , then it must be the two ends incident to v^{out} and disjoint from α_i .

Case (g): α_i is of type TT. In this case we first establish that v_i and v_i^{out} are on opposite sides of α_i , but this is an immediate consequence of the fact that $b_i^{\text{D}_i} \subset \alpha_i$ and that the type 2 slide $\tau_I \cong \tau_{i+1}$ has parity D_i .

It now follows immediately that, under the parity D_i type 2 slide $\tau_I \cong \tau_{i+1}$, the fold map F_i takes b_{i+1}^j bijectively to α_i , and that $\alpha_{i+1} = b_{i+1}^j$ is a transition path of type T.

Case (h): α_i is of type SS. In this case, since the parity is D_i and $b_i^{\text{D}_i} = \alpha_i$ survives the splitting $\tau_i \succ \tau_{i+1}$, it follows that α_{i+1} is a transition path of type T. There are again two subcases, (h1) where v_i and v_i^{out} are on the same side of α_i , and case (h2) where they are on the opposite side.

Although it is possible that α_I itself has type SS with the cusps v_I and v_I^{out} on the same side of α_I , as shown in Figure 38(h2), for any $i \geq I$ the only way that

α_{i+1} is of type SS is in case (f), and in that case the cusp $v_{i+1} \in \text{int}(\alpha_{i+1})$ is on the opposite side of α_{i+1} from the cusp located at the outflowing switch of α_{i+1} . Thus, for $i \geq I + 1$, if α_i is of type SS then the cusps v_i and v_i^{out} are on opposite sides of α_i .

Finally, if $v_i \in \alpha_i$ then all branches incident to v_i are of type 2 since (τ_i, v_i) is of type 2, and if $v_i \notin \alpha_i$ then α_i is a transition branch which is always of type 2. \diamond

Now we use Lemma 11.1.11 to reach a contradiction in Case 3. From the fact that in types SS or TT the cusps v_i, v_i^{out} are on opposite sides of α_i , it follows that α_i lifts to a smooth arc $\tilde{\alpha}_i \subset \partial \mathcal{C}(S - \tau_i)$ with one endpoint at the cusp v_i^{out} . Since α_i is a union of type 2 branches it follows that $\tilde{\alpha}_i \subset \partial^{v_i}$, and so $v_i^{\text{out}} \in \partial(\bar{\partial}^{v_i})$. Let β_i be the side of ∂^{v_i} containing $\tilde{\alpha}_i$, with one endpoint on v_i^{out} and the other endpoint on v_i . Subdivide β_i into subarcs concatenated at the points which project to switches $\neq s^{v_i}$. This concatenation can be written in the form

$$\beta_i = \tilde{\alpha}_i * \gamma_i^0 * \rho_i^1 * \dots * \rho_i^K * \gamma_i^K * \tilde{b}^{D'}$$

with the following properties. The subarc adjacent to the endpoint v_i^{out} is $\tilde{\alpha}_i$. The subarc adjacent to the endpoint v_i is a lift $\tilde{b}_i^{D'}$ of $b_i^{D'}$ for some $D' \in \{L, R\}$. The remaining subarcs alternate between γ 's and ρ 's. The overlay image of each $\gamma_i^i, i = 0, \dots, k$, is a type 1 source branch of τ_i . The overlay image of each $\rho_i^k, k = 1, \dots, K$, is a C-splitting arc, consisting of either a single type 2 sink branch, or a concatenation of two type 2 branches, one of which is a transition branch identified with b_i^R or b_i^L and the other a sink branch identified with b^{v_i} . The key fact we need from all this is that the subarc of β_i adjacent to $\tilde{\alpha}_i$ is γ_i^0 , whose overlay image is a type 1 source branch.

We claim that the source branch γ_i^0 survives forever, throwing us into Case 2, whose subcases each lead to some contradiction. The proof that γ_i^0 survives forever follows the same case analysis as the proof of Lemma 11.1.11. The train track τ_i must be in one of the cases (d)–(h) excepting cases (e1) and (h2) (case (e1) is ruled out because v_{i+1} and v_{i+1}^{out} must be on opposite sides of α_{i+1} , and case (h2) is ruled out because v_i and v_i^{out} must be on opposite sides of α_i), or τ_i is in the unnamed case in which both α_i and α_{i+1} are of type T. Among all these cases, however, cases (e1) and (h2) are the only ones in which γ_i^0 is killed, as Figure 38 shows. Thus, γ_i^0 survives forever.

This completes the proof of Theorem 11.1.10. \diamond

11.2 Combinatorial types of stable train tracks

We shall show that there is a special class of one cusp train tracks, called *stable* train tracks, which have the property that if $\tau_0 \succ \tau_1 \succ \dots$ is a one cusp canonical

expansion of an arational measured foliation then τ_n is stable for all sufficiently large n . The “stable” property is a combinatorial invariant, and hence each one cusp splitting circuit that classifies a pseudo-Anosov conjugacy class consists entirely of combinatorial types of stable train tracks.

This section describes stable one cusp train tracks, and classifies their combinatorial types. The next section uses this classification to describe an enumeration of one cusp splitting circuits of pseudo-Anosov mapping classes.

We describe correspondence between stable one cusp train tracks and a certain class of decorated cell divisions on S , called “stable cell divisions with distinguished prongs” or “stable CDPs”. This correspondence holds on the level of isotopy types as well as combinatorial types; see Proposition 11.2.5. Then we describe the combinatorial types of stable CDPs in terms of “chord diagrams”, following [Mos83], [Mos93], and [Mos96]. In particular, [Mos96] gives an in depth description of chord diagrams for triangulations with a single vertex. Our discussion of chord diagrams here will be for cell divisions with a single vertex, a sufficiently similar context that we shall appeal to [Mos96] for many of the details.

Stable train tracks. Consider a filling, one cusp train track (τ, v) . We identify v with a cusp in the cusped surface $\mathcal{C}(S - \tau)$. Let s^v denote the switch of τ at which the cusp v is located, and let c^v denote the component of $\mathcal{C}(S - \tau)$ having v as a cusp. Let $\partial^v = \partial^v \mathcal{C}(S - \tau)$ be the union of sides of $\mathcal{C}(S - \tau)$ incident to v , consisting of either one or two sides of c^v depending on whether c^v has one side or ≥ 2 sides. Let $\bar{\partial}^v = \bar{\partial}^v \mathcal{C}(S - \tau)$ be the union of all the sides of $\mathcal{C}(S - \tau)$ that are *not* incident to v , so $\bar{\partial}^v$ contains the complete boundary of each component of $\mathcal{C}(S - \tau)$ distinct from c^v , and if c^v has $n \geq 3$ sides then $\bar{\partial}^v$ contains the $n - 2$ sides of c^v that are disjoint from v .

A *stable train track* is a one cusp train track (τ, v) such that the overlay map $\mathcal{C}(S - \tau) \rightarrow S$ is injective on the set $\bar{\partial}^v$. Note that stability is a combinatorial invariant: if (τ, v) is a stable train track and ϕ is a homeomorphism then $(\phi(\tau), \phi(v))$ is a stable train track. To make the usage clear, a “stable train track” is always a one cusp train track (τ, v) , that is, a train track τ decorated with a particular cusp v .

The following results, particularly the corollary, explain the terminology “stable”:

Lemma 11.2.1. *If (τ, v) is a stable train track and if $(\tau, v) \gg (\tau', v')$ is a one cusp slide or a one cusp elementary split then (τ', v') is stable.*

Applying this lemma inductively to the elementary factorization of a one cusp splitting sequence we obtain:

Corollary 11.2.2. *If $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ is a one cusp splitting sequence, and if (τ_I, v_I) is stable for some I , then (τ_i, v_i) is stable for all $i \geq I$. \diamond*

Proof of Lemma 11.2.1. Let $(\tau, v) \succ (\tau', v')$ be a one cusp slide or elementary split, and suppose that (τ, v) is stable. Consider the overlay maps $f: \mathcal{C}(S - \tau) \rightarrow S$, $f': \mathcal{C}(S - \tau') \rightarrow S$. From the local model for a slide (Figure 11) or elementary split (Figure 13) one sees easily that there is a diffeomorphism $g: \mathcal{C}(S - \tau') \rightarrow \mathcal{C}(S - \tau)$ and a neighborhood $U \subset \mathcal{C}(S - \tau')$ of v' disjoint from $\bar{\partial}^{v'}$ such that if $x \in \mathcal{C}(S - \tau') - U$ then $f'(x) = f(g(x))$. In other words, except near the cusp v' , the overlay map f' makes no identifications that are not already made by f . Since f makes no identifications on $\bar{\partial}^v$, and since $U \cap \bar{\partial}^{v'}(S - \tau') = \emptyset$, it follows that f' makes no identifications on $\bar{\partial}^{v'}(S - \tau')$, and hence (τ', v') is stable. \diamond

Next we show that stable train tracks do, in fact, occur in one cusp expansions:

Proposition 11.2.3. *If $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ is a one cusp expansion of an arational measured foliation \mathcal{F} , then for all sufficiently large i the one cusp train track (τ_i, v_i) is stable.*

Proof. By Corollary 11.2.2 we just have to prove that some (τ_i, v_i) is stable, although the proof will show directly that (τ_i, v_i) is stable for all sufficiently large i .

Let ν_0 be a tie bundle over τ_0 , with tie foliation \mathcal{F}_v equipped with a positive Borel measure, and with \mathcal{F} chosen in its equivalence class to have a surjective carrying inclusion $\mathcal{F} \hookrightarrow \nu_0$, so \mathcal{F} is the horizontal foliation. Let ℓ be the separatrix of \mathcal{F} corresponding to the cusp v_0 . Let ξ_t be the initial segment of ℓ of length t . By Lemma 7.1.4 it follows that the elementary move sequence associated to the growing separatrix family ξ_t is the one cusp elementary factorization of $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$. Let ν_t be the tie bundle obtained from ν_0 by slicing along ξ_t , let τ_t be the quotient train track, and let v_t be the cusp of τ_t associated to the boundary point of ξ_t . We must show that (τ_t, v_t) is stable for sufficiently large t . Letting $\bar{\partial}^v \nu_t$ be the union of sides of ν_t that are disjoint from v_t , we must show that the map $\bar{\partial}^v \nu_t \rightarrow \tau_t$ is injective for sufficiently large t , or equivalently, that each tie of ν_t intersects $\bar{\partial}^v \nu_t$ in at most one point. To guarantee this, simply choose $t = \text{Length}(\xi_t)$ sufficiently large so that each tie of ν_0 which intersects $\bar{\partial}^v \nu_0$ also intersects the interior of ξ_t . \diamond

The import of Proposition 11.2.3 for our present needs is that in order to study one cusp splitting circuits it suffices to understand the combinatorial types of one cusp train tracks which are stable; there is no need to study nonstable ones.

The structure of stable train tracks. We shall give a general construction of stable train tracks, in terms of cell divisions of S with a particularly simple structure.

The construction is summarized in Proposition 11.2.5. As we will see below, this construction will make it easy to classify stable train tracks up to combinatorial equivalence.

For any one cusp train track (τ, v) we let s^v be the switch of τ at which the cusp v , and let b^v be the branch on the one-ended side of s^v .

First we classify stable one cusp train tracks (τ, v) into two types. Let $s_1, s_2 \in \partial\mathcal{C}(S - \tau)$ be the two regular boundary points which are pullbacks of s^v via the overlay map. Note that at least one one of points s_1, s_2 is contained in ∂^v : they cannot both be contained in $\bar{\partial}^v$, because they both map to s^v and the overlay map is injective on $\bar{\partial}^v$. We say that (τ, v) is of *type 1* if one of s_1, s_2 is contained in ∂^v , and (τ, v) is of *type 2* if both are. Note that the type is a combinatorial invariant.

Proposition 11.2.4. *If (τ, v) is of type 1 then b^v is a transition branch and the one cusp slide move $(\tau, v) \asymp (\tau', v')$ along b^v produces a type 2 one cusp train track (τ', v') .*

Proof. Let α be the side of $\bar{\partial}^v$ that contains one of the regular points $s_1 \in \partial\mathcal{C}(S - \tau)$ that maps to s^v . Following along α in the direction of v , at the end of α we reach a cusp v_1 of $\mathcal{C}(S - \tau)$. The segment $[s_1, v_1]$ embeds in τ , and moreover its interior contains no switches of τ , because such a switch would pull back via the overlay map to some cusp v'' of $\mathcal{C}(S - \tau)$, but $v'' \neq v$ and so $v'' \in \bar{\partial}^v$, contradicting that $\bar{\partial}^v$ injects into τ . Thus, the image of $[s_1, v_1]$ is equal to b^v , and it is a transition branch. Moreover, the two regular points of $\partial\mathcal{C}(S - \tau)$ that are identified with v_1 must be in ∂^v , by injectivity of $\bar{\partial}^v \rightarrow \tau$, and so if s is a point of τ just beyond v_1 then s pulls back via the overlay map to two points of ∂^v . We can slide along b^v past the cusp v_1 to the point x , to produce a new one cusp train track (τ', v') , such that the two regular pullbacks of v' are identified with the two pullbacks of x , which are in ∂^v , and hence (τ', v') is of type 2. \diamond

From Proposition 11.2.4 it follows that if we can construct all type 2 stable train tracks then we can also construct all type 1 stable train tracks. We now concentrate on the type 2 construction. The method is to show that each type 2 train track can be obtained by a simple construction from a simple class of cell divisions of S .

For present purposes a *cell division* C of S is a CW-decomposition of S , with k -skeleton denoted $C^{(k)}$, that satisfies the following provisions:

- No open 1-cell of C contains a puncture.
- Each open 2-cell of C contains at most one puncture.
- Each nonpunctured 2-cell of C has at least three sides.

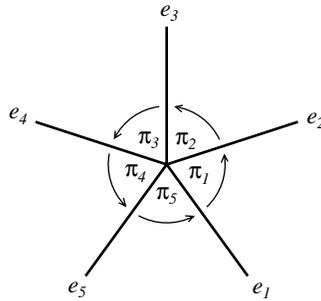


Figure 39: Ends e_i and prongs $\pi_i = (e_i, e_{i+1})$, $i \in \mathbf{Z}/5$, around a 0-cell v of valence 5. Arrows represent the successor map that defines the circular ordering on $\text{Ends}(v)$.

Note that a 0-cell of C may or may not be a puncture. Given a cell division C , let $V(C)$ be the set of 0-cells or vertices, $E(C)$ the set of 1-cells or edges, and $F(C)$ the set of 2-cells or faces. For each $e \in E(C)$ let $\text{Ends}(e)$ be the set of ends of e , a two element set; formally, $\text{Ends}(e)$ is the set of ends of $\text{int}(e)$ in the sense of Freudenthal. Let $\text{Ends}(C) = \bigcup_{e \in E(C)} \text{Ends}(e)$. Each $\eta \in \text{Ends}(e) \subset \text{Ends}(C)$ is located at a particular 0-cell v , meaning that as a point on e diverges to the end η then, in S , the point converges to v . For each $v \in V(C)$ let $\text{Ends}(v)$ be the subset of $\text{Ends}(C)$ of all ends η located at v . The set $\text{Ends}(v)$ inherits a circular ordering from the orientation on S (see Figure 39). We regard this circular ordering as being defined by a permutation of $\text{Ends}(v)$ with one cycle, called the *successor map*, denoted $e \mapsto \text{Succ}(e) = e'$, whose inverse map is called the *predecessor map* $e' \mapsto \text{Pred}(e') = e$. Define a *prong* at v to be an ordered pair $(e, e') \in \text{Ends}(v)$ where $e' = \text{Succ}(e)$. Let $\text{Pr}(v)$ be the set of prongs of v , with the circular ordering inherited from $\text{Ends}(v)$: given $e, e', e'' \in \text{Ends}(v)$ with $e' = \text{Succ}(e)$ and $e'' = \text{Succ}(e')$, the successor of (e, e') in $\text{Pr}(v)$ is (e', e'') . For each 2-cell c of C with n sides, as we traverse the boundary of c in counterclockwise order we encounter n prongs, and the circularly ordered set of these prongs is denoted $\text{Pr}(c)$.

Next we define the *dual pretrack* of a cell division C , denoted $\tau^\perp(C)$ (see Figure 40). First, for each $e \in E(C)$ let $b^\perp(e)$, the *dual branch* of e , be an arc that intersects $C^{(1)}$ transversely in a single point located in the interior of e . Next, consider a 2-cell c of C , which we regard formally as a component of $\mathcal{C}(S - C^{(1)})$. List the n sides of c in circular order as e_1, \dots, e_n . Let π_i be the prong between the sides e_i, e_{i+1} , $i \in \mathbf{Z}/n$. The overlay pullback of $\bigcup_{e \in E(C)} b^\perp(e)$ to c consists of n “half-branches” b'_1, \dots, b'_n , where b'_i has one endpoint in the interior of e_i and one endpoint in the interior of c . Let $g(c)$ be a cusped n -gon in the interior of c , with one cusp at the interior endpoint of each b'_i forming a generic train track switch at

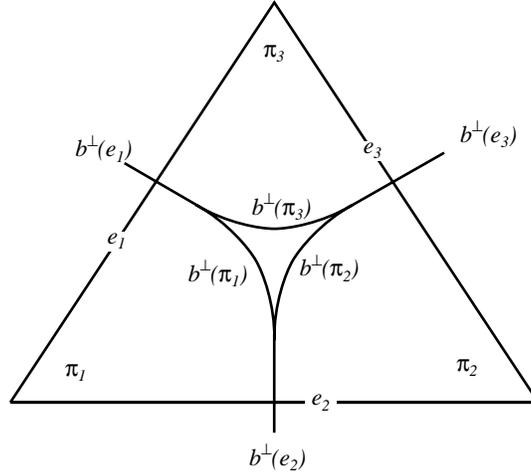


Figure 40: The dual track $\tau^\perp(C)$ of a cell division C has dual branches $b^\perp(e)$, $b^\perp(\pi)$ for each 1-cell $e \in E(C)$ and each prong $\pi \in \text{Pr}(C)$.

that point, and with the interior of $g(c)$ containing the puncture of c if there is one. For each prong $\pi_i \in \text{Pr}(c)$ there exists a side $b^\perp(\pi_i)$ of $g(c)$, called the *dual branch* of π_i , such that $b'_i \cup b^\perp(\pi_i) \cup b'_{i+1}$ is a smooth path. Identifying $g(c)$ with its overlay image in S , we now define

$$\tau^\perp(C) = \left(\bigcup_{e \in E(C)} b^\perp(e) \right) \cup \left(\bigcup_{c \in F(C)} g(c) \right) = \left(\bigcup_{e \in E(C)} b^\perp(e) \right) \cup \left(\bigcup_{\pi \in \text{Pr}(C)} b^\perp(\pi) \right)$$

It follows that for c as above, $\tau^\perp(C) \cap c = b'_1 \cup \dots \cup b'_n \cup g(c)$, which we call the *dual track* of c .

The components of $\mathcal{C}(S - \tau^\perp(C))$ and their indices may be listed as follows. In the interior of each punctured n -gon c of C there is a punctured, cusped n -gon $g(c)$, with $\iota(g(c)) = -\frac{n}{2} < 0$. In the interior of each nonpunctured n -gon c of C there is a nonpunctured, cusped n -gon $g(c)$, with $\iota(g(c)) = 1 - \frac{n}{2} < 0$, where the inequality holds by the requirement that $n \geq 3$. Containing each puncture 0-cell v of C there is a punctured nullgon $g(v)$ with $\iota(g(v)) = 0$. Containing each nonpuncture 0-cell v of C there is a nonpunctured nullgon $g(v)$ with $\iota(g(v)) = 1$. The existence of non-negative index components of $\mathcal{C}(S - \tau^\perp(C))$ implies that $\tau^\perp(C)$ is *not* a train track.

Consider now an ordered subset $\Pi \subset \text{Pr}(C)$. The pair (C, Π) is called a *cell division with distinguished prongs* or *CDP*. The ordering on Π will be significant only in the stable setting below; otherwise one may ignore the ordering and think

of Π as just a subset of $\text{Pr}(C)$. Define

$$\tau^\perp(C, \Pi) = \tau^\perp(C) - \bigcup_{\pi \in \Pi} \text{int}(b(\pi))$$

Notice that a necessary and sufficient condition for $\tau^\perp(C, \Pi)$ to be a pretrack is that for each 2-cell c of C and any two successive prongs π, π' in the circularly ordered set $\text{Pr}(c)$, at least one of π, π' is not in Π . In particular, if Π contains at most one prong in each 2-cell and Π does not contain the prong of any punctured monogon, then $\tau^\perp(C, \Pi)$ is a pretrack.

Assuming now that $\tau^\perp(C, \Pi)$ is a pretrack, there is a simple computation that determines the indices of the components of $\mathcal{C}(S - \tau^\perp(C, \Pi))$, from which one can decide whether $\tau^\perp(C, \Pi)$ is a train track by checking that all indices are negative. Namely, construct a bipartite graph Γ with one vertex for each 0-cell and each 2-cell of C , and with an edge connecting a 0-cell v to a 2-cell c for each prong $\pi \in \text{Pr}(v) \cap \text{Pr}(c) \cap \Pi$. The components of $\mathcal{C}(S - \tau^\perp(C, \Pi))$ correspond bijectively to the components of Γ . Given a component of Γ , by adding the indices of the components of $\mathcal{C}(S - \tau^\perp(C))$ corresponding to vertices of the given component of Γ , we obtain the index of the corresponding component of $\mathcal{C}(S - \tau^\perp(C, \Pi))$. Assuming that $\tau^\perp(C, \Pi)$ is a train track, there is another simple computation that allows one to determine whether $\tau^\perp(C, \Pi)$ is filling: each component of Γ should be a tree and should contain at most one vertex corresponding to a puncture 0-cell of C .

Consider a CDP (C, Π) where Π is an ordered subset of $\text{Pr}(C)$ consisting of one or two prongs, the first called the *distinguished prong* and denoted π_* and the second—if it exists—called the *secondary prong* and denoted $\pi_\#$. Let c_* be the 2-cell of C containing π_* . If $\pi_\#$ exists let $c_\#$ be the 2-cell of C containing $\pi_\#$. We say that (C, Π) is a *stable* CDP if the following properties hold:

- c_* is a nonpunctured triangle of C called the *distinguished triangle*.
- If the 0-cell of C is not a puncture then $\pi_\#$ *must* exist.
- If $\pi_\#$ exists then:
 - $c_\# \neq c_*$, and
 - $c_\#$ is not punctured, and hence $c_\#$ has at least three sides (by definition of a cell division).

Note that these conditions guarantee that $\tau = \tau^\perp(C, \Pi)$ is a filling train track (see Figure 41). In fact the singularity type of τ is computable from the combinatorics of (C, Π) , as follows: τ has a nonpunctured cusped n -gon for every nonpunctured n -gon of C outside of c_* and $c_\#$; τ has a punctured cusped n -gon for every punctured n -gon

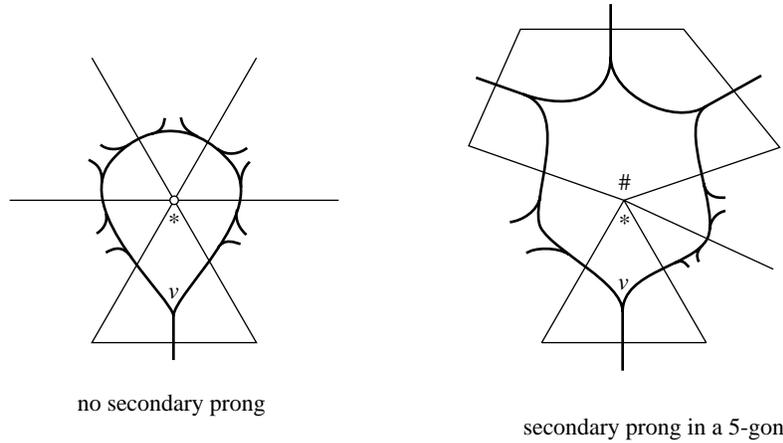


Figure 41: If (C, Π) is a stable CDP then its dual track $\tau = \tau^\perp(C, \Pi)$ has a naturally determined cusp v such that (τ, v) is a type 2 stable train track. If the secondary prong $\pi_\#$ does not exist then the component c^v of $\mathcal{C}(S - \tau)$ containing v is a cusped monogon. If the secondary prong $\pi_\#$ exists and is contained in an n -gon $c_\#$ then c^v is a cusped $(n - 1)$ -gon.

of C ; and τ has one other complementary component P containing the 0-cell of C , either a cusped monogon when the secondary prong does not exist, and otherwise a cusped $(n - 1)$ -gon if $c_\#$ has n sides. Letting v be the unique cusp of τ that is contained in c_* , we say that (τ, v) is the *dual one cusp train track* of (C, Π) .

Note the following naturality properties of the dual one cusp train track. Given a stable CDP (C, Π) :

- The isotopy type of the dual one cusp train track (τ, v) is well-defined, and depends only on the isotopy type of (C, Π) .
- For any $\phi \in \text{Homeo}_+$, $\phi(\tau, v)$ is (isotopic to) the dual one cusp train track of $\phi(C, \Pi)$.

Note also that (τ, v) is a type 2 stable train track: each component of $S - C$, except for c_* , and except for $c_\#$ when it has three sides, contains the embedded image of a component of $\bar{\partial}^v \mathcal{C}(S - \tau)$, and these images are disjoint because the components of $\mathcal{C}(S - C)$ are disjoint. Type 2 follows because the cusp v is in c_* and thus is isolated from the images of the components of $\bar{\partial}^v \mathcal{C}(S - \tau)$. This proves part of the following:

Proposition 11.2.5. *The map which assigns to each stable CDP its dual one cusp train track induces a bijection between the set of isotopy (resp. combinatorial) types*

of stable CDPs and isotopy (resp. combinatorial) types of type 2 stable train tracks.

Proof. We have described a map which associates, to each stable CDP, a type 2 stable train track. We must construct an inverse map: for every type 2 stable train track (τ, v) we construct a dual stable CDP (C, Π) . The naturality and duality provisions that we need will follow easily from the construction:

- (C, Π) is well-defined up to isotopy, depending only on the isotopy type of (τ, v) .
- If $\phi \in \text{Homeo}_+$ then $\phi(C, \Pi)$ is the dual stable CDP of $\phi(\tau, v)$.
- The dual of the dual of a stable train track (τ, v) is isotopic to (τ, v) , and the dual of the dual of a stable CDP (C, Π) is isotopic to (C, Π) .

To start with, ignoring the stable setting for the moment, we associate to every generic, filling train track τ a CDP (C', Π') whose dual is τ . The cell division C' has one 0-cell in each component of $\mathcal{C}(S - \tau)$, corresponding with the puncture in that component if there is one; there is one 1-cell $e^\perp(b)$ dual to each branch b of τ ; and there is one 2-cell containing each switch of τ , and by necessity this 2-cell is a nonpunctured triangle. Note that for each triangle t of C' there is a unique prong $\pi'(t) \in \text{Pr}(t)$ such that $\tau \cap t$ is obtained from the dual track of t by removing the branch dual to the prong $\pi'(t)$. Letting $\Pi' = \{\pi'(t) \mid t \in F(C')\}$, clearly (C', Π') is a CDP, called the *dual CDP* of τ . The dual CDP of τ is well-defined up to isotopy, depending only on the isotopy type of τ , and it is natural with respect to the action of Homeo_+ . It follows that the dual CDP induces a bijection between the set of isotopy (resp. combinatorial) types of generic, filling train tracks τ and the set of isotopy (resp. combinatorial) types of CDPs (C', Π') having the property that each 2-cell of C' is a nonpunctured triangle, and for each 0-cell $v \in V(C)$, the set $\Pi' \cap \text{Pr}(v)$ has at least one element if v is a puncture and at least three elements if v is not a puncture.

Suppose now that (τ, v) is a type 2 stable train track. Let (C', Π') be the dual CDP of the generic train track τ . We shall remove certain edges of C' to produce a cell division C .

Recall that each side α of $\bar{\partial}^v$ maps bijectively to a source branch of τ via the overlay map, and so there is a 1-cell e_α^\perp of C' dual to α .

Define C by removing from C' each 1-cell e_α^\perp dual to a side α of $\bar{\partial}^v$. We have several tasks: prove C is a cell division, which requires identifying the 2-cells of C ; identify the distinguished prong π_* ; identify the secondary prong $\pi_\#$, in particular saying when $\pi_\#$ exists; prove that (C, Π) is a stable CDP; and prove the naturality and duality provisions.

First we identify the distinguished triangle and distinguished prong of C . Let c_* be the triangle of C' containing the switch s^v . None of the three branches incident to s^v are images of sides of $\bar{\partial}^v$, by definition of a type 2 stable train track, and it follows that c_* is a complementary component of C . In particular, we can define the distinguished prong to be $\pi_* = \pi'(c_*)$.

Next we identify a further collection of 2-cells of C . Consider a component P of $\mathcal{C}(S - \tau)$, an n -cusped polygon for some $n \geq 1$, such that ∂P is completely contained in $\bar{\partial}^v$. The overlay map embeds P in S . Let $\alpha_1, \dots, \alpha_n$ be the sides of P , let $v_i = \alpha_i \cap \alpha_{i+1}$, and let β_i be the branch on the one-ended side of v_i . Note that $\alpha_i \subset \bar{\partial}^v$ but $\beta_i \not\subset \bar{\partial}^v$, for each $i = 1, \dots, n$. Each $e_{\alpha_i}^\perp$ is therefore removed from C , but none of the $e_{\beta_i}^\perp$ are removed. It follows that $e_{\beta_1}^\perp, \dots, e_{\beta_n}^\perp$ form the sides of an n -gon component c_P^\perp of $\mathcal{C}(S - C^{(1)})$ containing P , and containing no puncture except for the one in P if there is one. Moreover, $c_P^\perp \cap \tau$ is clearly a dual track of c_P^\perp .

We still have to identify the secondary prong $\pi_\#$ and the 2-cell $c_\#$ containing $\pi_\#$, if they exist. Consider the component c^v of $\mathcal{C}(S - \tau)$ having v as one of its cusps, and suppose that P is an n -cusped polygon for $n \geq 1$. If $n = 1$ then the secondary prong does not exist, and if $n \geq 2$ then it does exist. We consider the two cases $n \geq 3, n = 2$ separately.

If $n \geq 3$ then c^v contains the unique component of $\bar{\partial}^v$ which is not a closed curve, and instead is an arc consisting of a concatenation of $n - 2$ sides $\alpha_1, \dots, \alpha_{n-2}$. Let v_0, \dots, v_{n-2} be the cusps at the ends of these arcs, with $\partial\alpha_i = \{v_{i-1}, v_i\}$. Let β_i be the branch on the one-ended side of v_i . Note that each $e_{\alpha_i}^\perp$ is removed from C' , but none of the $e_{\beta_i}^\perp$ are removed, and so C' has an $(n + 1)$ -sided 2-cell $c_\#$ whose sides are identified with $e_{\beta_0}^\perp, \dots, e_{\beta_n}^\perp$. The secondary prong $\pi_\#$ is identified with the prong at which $e_{\beta_0}^\perp, e_{\beta_n}^\perp$ join. Note that the interior of $c_\#$ is unpunctured.

If $n = 2$ let v' be the cusp of c^v distinct from v , and let $c_\#$ be the unpunctured triangle of C' containing v' . Let $\beta_0, \beta_1, \beta_2$ be the three branches of τ incident to v' with β_1 on the one-ended side of v' , so the three sides of $c_\#$ are $e_{\beta_0}^\perp, e_{\beta_1}^\perp, e_{\beta_2}^\perp$. None of $\beta_0, \beta_1, \beta_2$ are identified with sides of $\bar{\partial}^v$: β_1 is not a source branch, and β_0, β_2 could only be identified with sides of ∂^v . It follows that none of $e_{\beta_0}^\perp, e_{\beta_1}^\perp, e_{\beta_2}^\perp$ are removed, and so $c_\#$ is a nonpunctured triangular 2-cell of C . The secondary prong $\pi_\#$ is identified with the prong at which $e_{\beta_0}^\perp, e_{\beta_2}^\perp$ join.

In both of the cases $n \geq 3$ and $n = 2$, it is clear from the description that $\tau \cap c_\#$ is obtained from a dual track of $c_\#$ by removing the dual branch of the cusp $\pi_\#$.

We claim that every component of $S - C^{(1)}$ occurs in the accounting above, and to prove this we must show that every triangle t of C' is contained in one of the components of $S - C^{(1)}$ described above. Let v_t be the cusp of τ contained in t . If $v_t = v$ then $c = c_*$, so suppose that $v \neq v_c$. If v_t is a cusp of a component P of

$\mathcal{C}(S - \tau)$ such that $\partial P \subset \bar{\partial}^v$ then clearly $t \subset c_P^\perp$. Finally, if v_t is a cusp of c^v distinct from v then clearly $t \subset c_\#$.

This completes the proof that (C, Π) is a stable CDP, and it is clear from the construction that the naturality and duality provisions hold. \diamond

I keep dropping the "type 2". Perhaps the terminology should be simplified.

Combinatorial types of stable train tracks. By Proposition 11.2.5, in order to classify combinatorial types of stable train tracks it suffices to classify combinatorial types of stable CDPs. We do this as follows.

Consider a stable CDP (C, Π) . The counterclockwise circular ordering on $\text{Ends}(C)$ is characterized by the *successor map* $\text{Succ}: \text{Ends}(C) \rightarrow \text{Ends}(C)$, a permutation with a single cycle. The inverse permutation of Succ is the *predecessor map* $\text{Pred}: \text{Ends}(C) \rightarrow \text{Ends}(C)$. Define the *opposite end map* $\text{Opp}: \text{Ends}(C) \rightarrow \text{Ends}(C)$ to be the transposition which interchanges the two ends of any 1-cell of C . Recall that a prong $\pi \in \text{Pr}(C)$ is formally an ordered pair $\pi = (e, e')$ such that $e' = \text{Succ}(e)$. Each prong $\pi \in \text{Pr}(C)$ is a corner of some 2-cell $c \in F(C)$, and the next corner of c in counterclockwise order around c from π , denoted $\text{Next}(\pi)$, is given by the formula

$$\text{Next}(e, \text{Succ}(e)) = (\text{Pred} \circ \text{Opp}(e), \text{Opp}(e))$$

In other words, the natural bijection $\text{Ends}(C) \leftrightarrow \text{Pr}(C)$ given by $e \leftrightarrow (e, \text{Succ}(e))$ conjugates the permutation $\text{Pred} \circ \text{Opp}$ with the "next corner" map $\text{Next}: \text{Pr}(C) \rightarrow \text{Pr}(C)$. It follows that the 2-cells of C are in natural one-to-one correspondence with the cycles of the permutation $\text{Pred} \circ \text{Opp}$. Define Punct to be the union of the cycles of $\text{Pred} \circ \text{Opp}$ corresponding to the corners of punctured 2-cells of C . The quintuple $(\text{Ends}(C), \text{Succ}, \text{Opp}, \text{Punct}, \Pi)$ is called the *combinatorial diagram* of the CDP (C, Π) .

Given two stable CDPs $(C, \Pi), (C', \Pi')$, an *isomorphism* between their combinatorial diagrams $(\text{Ends}(C), \text{Succ}, \text{Opp}, \text{Punct}, \Pi)$ and $(\text{Ends}(C'), \text{Succ}', \text{Opp}', \text{Punct}', \Pi')$ is defined to be a bijection $\text{Ends}(C) \xrightarrow{f} \text{Ends}(C')$ that preserves the structure of the combinatorial diagrams, that is: $f \circ \text{Succ} = \text{Succ}' \circ f, f \circ \text{Opp} = \text{Opp}' \circ f, f(\text{Punct}) = \text{Punct}'$, and $f(\Pi) = \Pi'$.

Proposition 11.2.6. *Given two type 2 stable train tracks $(\tau, v), (\tau', v')$ with dual stable CDPs $(C, \Pi), (C', \Pi')$, the following are equivalent:*

- (1) (τ, v) and (τ', v') are combinatorially equivalent.
- (2) (C, Π) and (C', Π') are combinatorially equivalent.
- (3) The combinatorial diagrams of (C, Π) and (C', Π') are isomorphic.

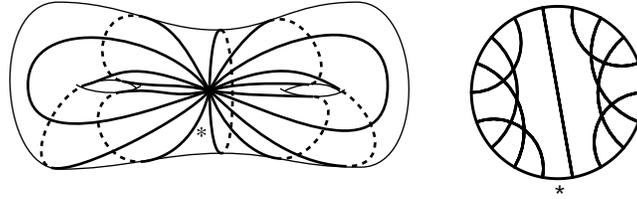


Figure 42: A chord diagram for a CDP (C, π_*) on a once-punctured surface of genus 2. There are 9 1-cells, 18 ends, and 18 prongs. There are 6 2-cells, each a nonpunctured triangle, one of them the distinguished triangle. The dual stable train track (τ, v) has singularity type consisting of a once-punctured monogon with cusp v , and 5 nonpunctured trigons.

Proof. The equivalence of (1) and (2) was established in Proposition 11.2.5. The implication (2) \implies (3) is obvious.

To prove (3) \implies (2) we must construct a combinatorial equivalence $F: (C, \Pi) \rightarrow (C', \Pi')$. Choose a handle decomposition of S dual to C , and another dual to C' . Since $f \circ \text{Succ} = \text{Succ}' \circ f$, we can construct F taking the unique 0-handle D_0 of C to the unique 0-handle D'_0 of C' by an orientation preserving homeomorphism so that $F(D_0 \cap C^{(1)}) = D'_0 \cap C'^{(1)}$, and so that F induces the map $\text{Succ}: \text{Ends}(C) \rightarrow \text{Ends}(C')$. Since $f \circ \text{Opp} = \text{Opp}' \circ f$, we may extend F to a homeomorphism from the 0-handles union the 1-handles of C to the 0-handles union the 1-handles of C' , taking $C^{(1)}$ to $C'^{(1)}$. We may now extend F over the 2-cells, and since $f(\text{Punct}) = \text{Punct}'$ this extension can be defined to be a bijection on punctures. Since $f(\Pi) = \Pi'$ it follows that $F(\Pi) = \Pi'$, and hence F is a combinatorial equivalence from (C, Π) to (C', Π') . \diamond

Chord diagrams. With Proposition 11.2.6 it becomes clear how to enumerate combinatorial types of stable train tracks: simply enumerate combinatorial diagrams of stable CDPs.

To visualize the combinatorial diagram of a stable CDP (C, Π) , the *chord diagram* is a useful device. With practice, the chord diagram can be used to decide by inspection when two combinatorial diagrams are isomorphic. Figure 42 shows an example on a once-punctured surface of genus 2, and Figure 43 shows an example on a thrice-punctured torus.

First we describe the chord diagram of C itself. Draw a round circle in the plane and mark one point on the circle for each element of $\text{Ends}(C)$; the correspondence should preserve circular ordering, where we use the usual counterclockwise ordering for the circle, and the permutation Succ for $\text{Ends}(C)$. Then connect two points by

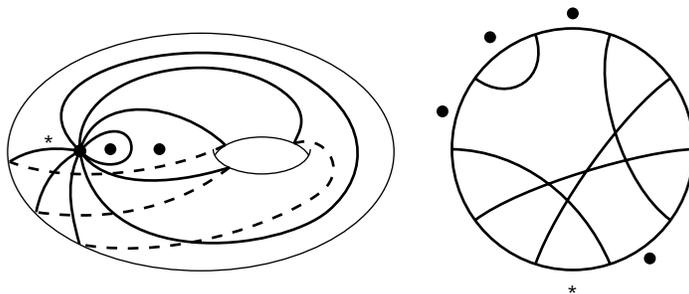


Figure 43: A stable CDP on a thrice-punctured torus, with chord diagram. There are two nonpunctured triangles (one the distinguished triangle), a punctured triangle, and a punctured monogon. The dual stable train track (τ, v) has singularity type consisting of two punctured monogons, one with cusp v , a punctured trigon, and two nonpunctured trigons.

a chord if they correspond to two ends which are interchanged by the transposition Opp . It is often easier to see what is going on by drawing hyperbolic chords — circular arcs perpendicular to the outer circle — instead of Euclidean chords. If S has no puncture in the interior of any 2-cell of C then $\text{Punct} = \emptyset$ and this completes the description of the chord diagram of C . More generally, elements of $\text{Pr}(C)$ are represented in the chord diagram by gaps, which are circular segments between adjacent chord endpoints, and we indicate the elements of Punct by drawing a solid dot adjacent to each appropriate gap, called a *gap dot*, so a punctured n -gon will yield n gap dots around the periphery of the circle. This completes the description of the chord diagram of C .

To describe the chord diagram of (C, Π) , represent the distinguished prong π_* by drawing a $*$ adjacent to the appropriate gap, and for the secondary prong $\pi_\#$, when it exists, draw a $\#$ adjacent to the appropriate gap.

The representation of punctured n -gons by gap dots is somewhat inefficient when n is large; for example, the two punctured n -gons in Figure 43 are represented by four gap dots. On a multi-punctured sphere, the punctured regions of a stable CDP (C, Π) can be represented in a particularly efficient way in the chord diagram. The chord diagram of C can be viewed as a literal copy of C : the chords divide the interior of the circle into regions, and the whole closed disc bounded by the circle can be mapped to the sphere by collapsing the boundary circle to the unique 0-cell of C , taking each chord to 1-cell, and taking each region of the disc to a 2-cell of C . Hence, we can indicate a punctured region simply by putting a dot in the appropriate complementary region of chords. Figure 44 shows two examples of an stable CDP on an 8-punctured sphere using this notation.

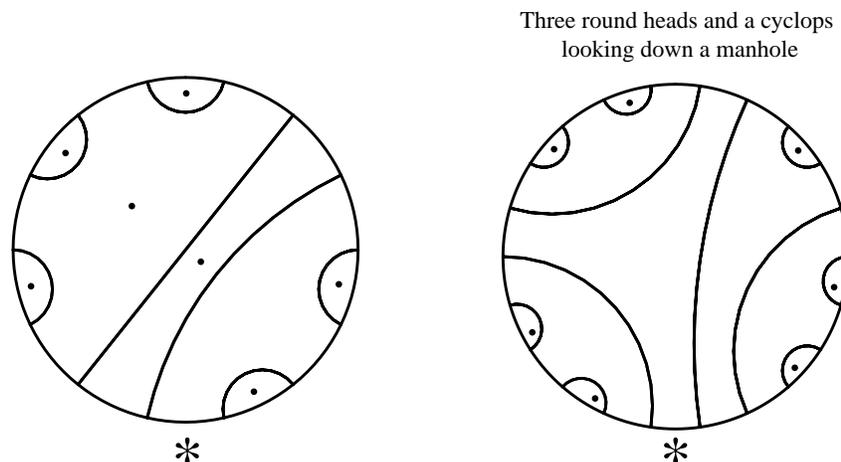


Figure 44: Some chord diagrams in genus 0 with 8 punctures. The first example has 5 punctured monogons, 1 punctured bigon, and 1 punctured 4-gon, in addition to the distinguished triangle. The second has 7 punctured monogons and 5 nonpunctured triangles, including the distinguished triangle.

Another nice feature of CDPs for multipunctured sphere is that one can superimpose the dual stable train track onto the chord diagram quite clearly, as shown in Figure 45.

Now we use these ideas to enumerate certain combinatorial types in low genus.

Genus 0 with 4 punctures. On a 4-punctured sphere, every recurrent, stable train track has the same singularity type, namely 4 once-punctured monogons, and there is a unique combinatorial type of stable train track with this singularity type, shown in Figure 45(b).

Genus 0 with 5 punctures. We give a complete enumeration of combinatorial types of stable, filling train tracks on the 5-punctured sphere. There are two possible singularity types of filling train tracks: 5 punctured monogons and 1 nonpunctured triangle; or 4 punctured monogons and 1 punctured bigon.

Consider a recurrent, stable train track (τ, v) with 5 punctured monogons and one nonpunctured triangle, with v located in a punctured monogon. The corresponding stable CDP (C, π_*) has 4 punctured monogons and 2 non-punctured triangles including the distinguished triangle. To obtain the possible chord diagrams for C , note that sides of the 4 punctured monogons bound a nonpunctured 4-gon,

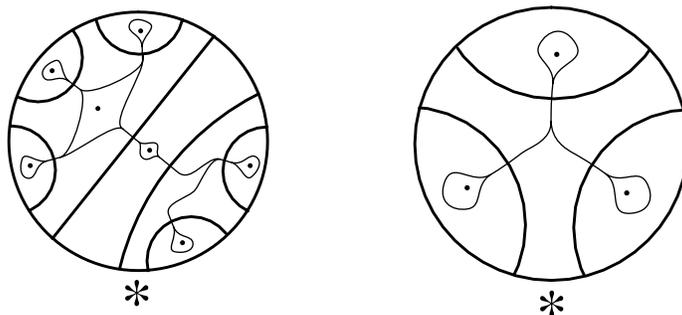


Figure 45: Some genus 0 chord diagrams and their dual train tracks. (a) shows the dual train track for the first example in Figure 44. (b) shows the unique combinatorial type of a recurrent, stable train track on a 4-punctured sphere.

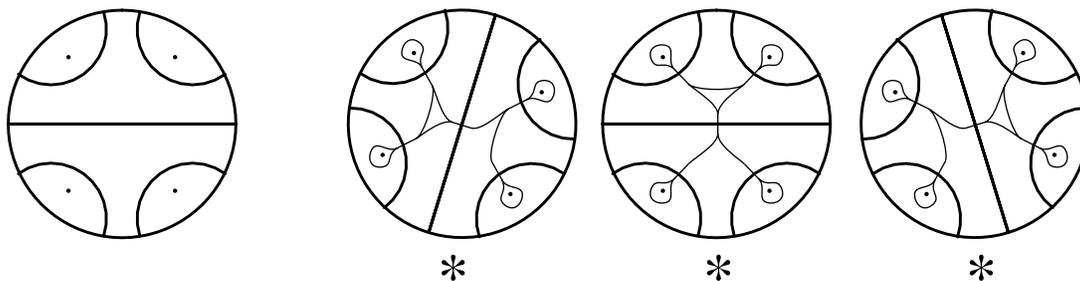


Figure 46: On a 5-punctured sphere, there are 3 combinatorial types of stable train tracks with 5 punctured monogons and 1 nonpunctured trigon. Each of these 3 examples has 3 possible positions for the secondary prong, giving 9 combinatorial types of stable train tracks with 4 punctured monogons and 1 punctured bigon.

which can be triangulated into 2 triangles in two different ways. However, these two triangulations are equivalent via a one quarter rotation of the 4-gon. Hence there is again a unique chord diagram for C , shown in Figure 46. This time there are six possible positions for the distinguished prong, but up to the 2-fold rotational symmetry of the chord diagram there are only 3 possible nonisomorphic chord diagrams, all shown in Figure 46, with their dual stable train tracks.

Consider next a stable train track (τ, v) with 4 punctured monogons and one punctured bigon containing the cusp v . The corresponding stable CDP $(C, (\pi_*, \pi_\#))$ again has 4 punctured monogons and 2 nonpunctured triangles, one triangle containing the distinguished prong π_* and the other containing the secondary prong $\pi_\#$. Once the choice of the distinguished prong is made among the 3 possible choices, the remaining nonpunctured triangle has three possible choices for the secondary prong, making $3 \cdot 3 = 9$ combinatorial types.

Consider next the case where (τ, v) has 4 punctured monogons and one punctured bigon, with the cusp v contained in a punctured monogon. The dual stable CDP (C, π_*) has three punctured monogons, a punctured bigon, and the distinguished triangle. The combinatorial type of C is uniquely determined and there are 3 positions for the distinguished prong, giving 3 combinatorial types altogether. The chord diagrams are left to the reader.

Finally consider the case where (τ, v) has 5 punctured monogons and a nonpunctured triangle, with the cusp v located in the nonpunctured triangle. The dual stable CDP $(C, (\pi_*, \pi_\#))$ has 5 punctured monogons, one nonpunctured triangle containing π_* , and one nonpunctured 4-gon containing the secondary prong $\pi_\#$. There are 3 positions for the distinguished prong, and 4 positions for the secondary prong, giving $3 \cdot 4 = 12$ combinatorial types altogether, whose chord diagrams are left to the reader.

This completes the enumeration, giving $3 + 9 + 3 + 12 = 27$ combinatorial types of filling, stable train tracks on a 5-punctured sphere.

Genus 0 with 6 and 7 punctures. We enumerate only the combinatorial types of stable, filling train tracks (τ, v) with the “principle” singularity type, namely a punctured monogon around each puncture and the rest nonpunctured trigons, and with v located in a punctured monogon.

When (τ, v) has 6 punctured monogons and 3 nonpunctured trigons, the dual stable CDP (C, π_*) has 5 punctured monogons and 3 nonpunctured triangles, with v located in one of the nonpunctured triangles. There is one combinatorial type for C : the boundaries of the 5 punctured monogons form a nonpunctured 5-gon, which can be subdivided into 3 nonpunctured triangles in 5 different ways, but these 5 ways differ by rotation of the 5-gon, leading to a unique combinatorial type. There

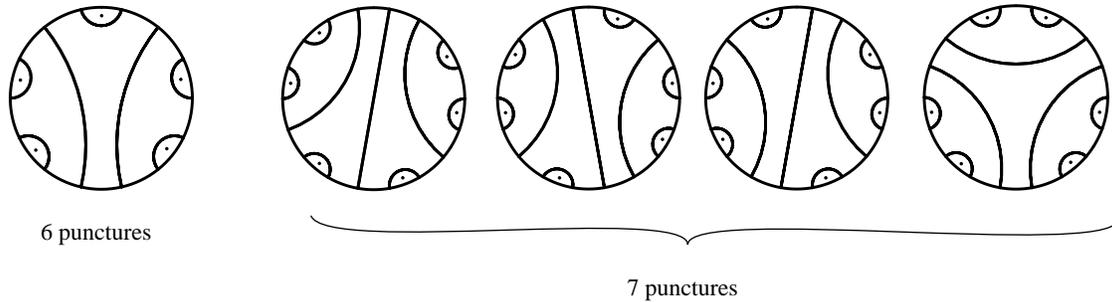


Figure 47: Chord diagrams for stable CDPs (C, π_*) with one puncture the vertex, a punctured monogon around every other puncture, and the rest nonpunctured triangles. With 6 punctures there is a unique combinatorial type, and with 7 punctures there are four combinatorial types, two of which differ by an orientation reversing homeomorphism.

are $3 \cdot 3 = 9$ positions for the distinguished prong π_* , yielding 9 combinatorial types.

When (τ, v) has 7 punctured monogons and 4 nonpunctured trigons, the dual stable CDP (C, π_*) has 6 punctured monogons and 4 nonpunctured triangles, with v located in one of the nonpunctured triangles. There are 4 combinatorial types for C : the boundaries of the 6 punctured monogons form a nonpunctured 6-gon which can be subdivided into 4 triangles, and up to rotation of the 6-gon there are 4 distinct ways to do this triangulation. The combinatorial type of (C, π_*) depends on where π_* is located. One of the 4 types for C has no rotational symmetry, and the choice of π_* gives $3 \cdot 4 = 12$ combinatorial types. Two of the 2 types for C have $\mathbf{Z}/2$ rotational symmetry, and the choice of π_* gives an additional $2 \cdot \frac{1}{2} \cdot 3 \cdot 4 = 12$ combinatorial types. The remaining type for C has $\mathbf{Z}/3$ rotational symmetry, and the choice of π_* gives $\frac{1}{3} \cdot 3 \cdot 4 = 4$ combinatorial types. Altogether there are $12 + 12 + 4 = 28$ combinatorial types.

Genus 1. On a once-punctured torus, there is a unique combinatorial equivalence class of recurrent stable train tracks (τ, v) , shown in Figure 48. It has a punctured bigon.

Consider next a twice punctured torus, and a stable train track (τ, v) with two punctured monogons and two punctured trigons, and with v located in a punctured monogon. The dual CDP (C, π_*) has one punctured monogon and three nonpunctured triangles, one of which is the distinguished triangle. There is a unique combinatorial type for C , which can be seen as follows. If the boundary edge of the punctured monogon of C is removed, the resulting cell division has a punctured bigon and two non-punctured triangles. If the puncture in the bigon is removed

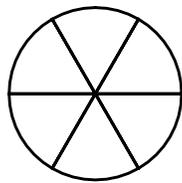


Figure 48: On a once-punctured torus, there is a unique combinatorial type of recurrent, stable train tracks.

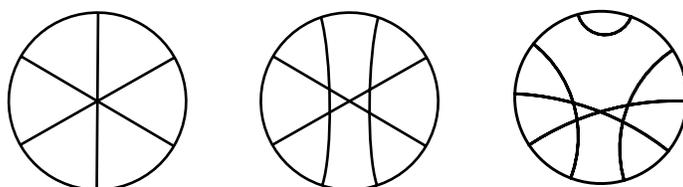


Figure 49: On a twice-punctured torus, if (C, π_*) is a stable CDP with one punctured monogon and three nonpunctured triangles, the combinatorial type of C is unique. The chord diagram is obtained from the unique chord diagram for a triangulation of a once punctured torus by splitting one chord and inserting a punctured monogon.

and the bigon is collapsed by identifying its two boundary arcs, the result is a cell division C' of a once punctured torus with two triangles, whose combinatorial type is unique. This suggests that C can be reconstructed from C' by choosing an end $\eta \in \text{Ends}(C')$, splitting the 1-cell containing η into two parallel 1-cells bounding a bigon, thereby splitting η into one of the prongs of the bigon, and inserting a punctured monogon at this prong. This operation can be performed on the chord diagram of C' , as illustrated in Figure 13.25. There are 6 ends η to choose from, but they are equivalent up to a 6-fold rotation of C' , hence there is a unique combinatorial type for C . The three nonpunctured triangles of C have 9 prongs altogether, from which to choose the distinguished prong π_* .

The method used for the twice-punctured torus can be adapted to the thrice punctured torus, to enumerate combinatorial types of stable train tracks (τ, v) with three punctured monogons and three nonpunctured triangles, with v in a punctured monogon; there are 60 such types. The combinatorial types of the dual CDPs (C, π_*) , each with two punctured monogons and four nonpunctured triangles, are constructed as follows. Start from the unique chord diagram for a one vertex triangulation of a torus. Choose two ends of the chord diagram, not necessarily distinct,

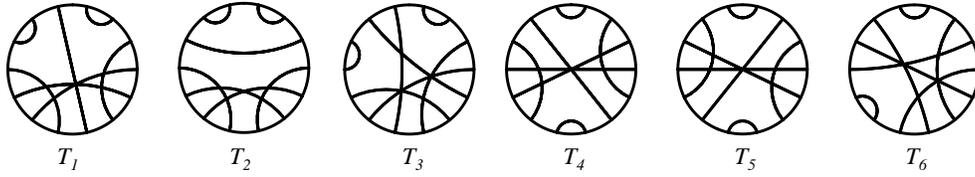


Figure 50: On a thrice-punctured torus, there are 60 combinatorial types of stable CDPs (C, π_*) with two punctured monogons and three nonpunctured triangles, obtained from 6 combinatorial types for C itself by inserting a distinguished prong.

as positions for two punctured monogons to be inserted. Then split the chords and insert the monogons. If both monogons are inserted at ends of the same chord, there will be a 4-gon to be triangulated. Figure 51 shows the 6 different chord diagrams obtained in this manner, labelled T_1 through T_6 . T_1 and T_2 are obtained by inserting monogons at the same chord end, and triangulating the resulting 4-gon in two different ways. T_3 is obtained by inserting the monogons at adjacent chord ends. T_4 is obtained by inserting the monogons at chord ends separated by a single chord end. T_5 and T_6 are obtained by inserting monogons at opposite ends of the same chord, and triangulating the resulting 4-gon. For each chord diagram, there are 12 possible positions for the distinguished prong. Diagrams $T_1 - T_4$ have no symmetries, and diagrams T_5, T_6 have 2-fold rotational symmetry. Hence there are $4 \cdot 12 + 2 \cdot (12/2) = 60$ distinct combinatorial types.

Genus 2 with one puncture. On a once-punctured surface of genus 2, there are 105 combinatorial types for a stable train track (τ, v) with a punctured monogon containing v and 5 nonpunctured triangles. In the dual stable CDP (C, π_*) , C is a one vertex triangulation of a genus 2 surface, of which there are 9 different combinatorial types, shown in Figure ???. We refer the reader to [Mos96] for this enumeration. Of these types, 3 have no rotational symmetry, 5 have 2-fold rotational symmetry, and 1 has 3-fold rotational symmetry. Each have 18 prongs, giving the count of $3 \cdot 18 + 5 \cdot (18/2) + 1 \cdot (18/3) = 105$.

These 105 combinatorial types are the exact same objects as the 105 states in the “principal layer” of the word acceptor for the automatic structure of the mapping class group of a once punctured surface of genus 2, as described in [Mos96].

11.3 Enumeration of one cusp splitting circuits

Each one cusp splitting on a stable train track produces a stable train track, and hence one can construct a directed graph $\tilde{\Gamma} = \tilde{\Gamma}(S)$ whose vertices are the isotopy

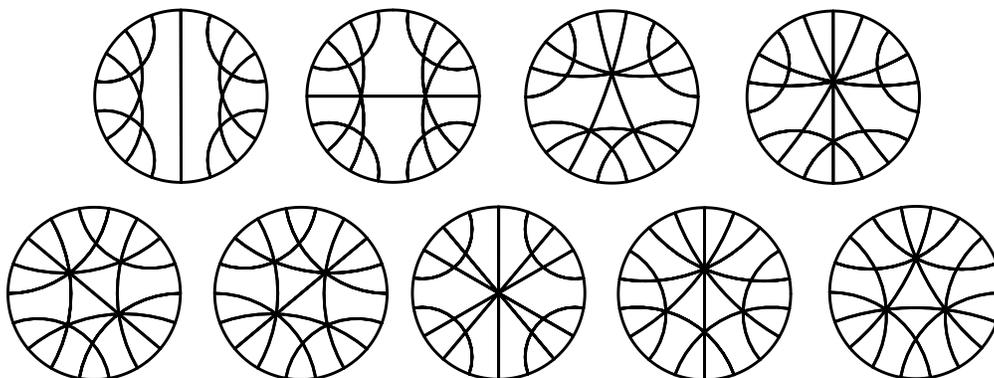


Figure 51: The 9 combinatorial types of one vertex triangulations of a genus 2 surface.

types of one cusp train tracks and whose directed edges are the one cusp splittings. The mapping class group acts on $\tilde{\Gamma}$, and by Lemma ?? the action is free. Also, there are only finitely many orbits of the action, and hence there is a finite quotient graph $\Gamma = \Gamma(S)$ whose vertices are the combinatorial types of stable train tracks, called the *canonical digraph* of S . We give an algorithmic construction of Γ , by using chord diagrams to describe the combinatorial types of one cusp train tracks.

The directed loops in Γ are precisely the one cusp splitting circuits, although not all such loops arise from pseudo-Anosov mapping classes with fixed separatrices. Given a directed loop γ in Γ , we say that γ is a *canonical pseudo-Anosov loop* if there exists a pseudo-Anosov mapping class Φ and a fixed separatrix ℓ such that γ is the one cusp splitting circuit associated to Φ and ℓ .

In order to achieve our enumeration we must answer the following question in an effective manner:

- Given a directed loop γ in Γ , how do we tell whether γ is a canonical pseudo-Anosov loop?

The first answer to this question is that the Arationality Theorem 6.3.2 can be applied, yielding a version of the canonical killing criterion applied to directed loops γ in Γ , with the conclusion that γ satisfies the canonical killing criterion if and only if γ is a canonical pseudo-Anosov loop. In the course of this we will give an equivalent formulation of the canonical killing criterion that applies to one cusp splitting sequences and circuits, and which is analogous to the criterion developed in Section 9.4 for one switch splitting sequences.

The second answer to this question is to consider the collection Λ of directed loops as a regular language over the alphabet consisting of the directed edges of Γ ;

we prove that the subcollection Λ_{ps} of canonical pseudo-Anosov loops is a regular sublanguage. This is accomplished by constructing a finite directed graph Γ' and a directed map $d: \Gamma' \rightarrow \Gamma$ and proving that a directed loop γ in Γ is in Λ_{ps} if and only if no iterate of γ lifts to Γ' ; only a uniformly bounded number of iterates need to be checked, which suffices to prove that Λ_{ps} is regular.

We will give several examples of components of Γ and the portion of Γ' lying over that component, which immediately leads to rich collections of pseudo-Anosov conjugacy classes.

Finally, we shall show how the canonical diagraph encodes not only one cusp splitting circuits, but also the one switch splitting circuits described in Section ??; in fact, we shall describe a duality between one cusp splitting circuits and one switch splitting circuits, which relates a pseudo-Anosov conjugacy class to its inverse.

Good elementary moves: no secondary prong. Consider a stable CDP (C, π) with distinguished prong π , distinguished triangle t , and no secondary prong. Recalling that, formally, a prong is an ordered pair of ends of C such that one succeeds the other in the circular order, we denote $\pi = (e^L, e^R)$; in diagrams we draw t with π pointing upward, and using the counterclockwise ordering on ends as we look at the page it follows that e^R is the end to the Right of π and e^L is the end to the Left. Given $D \in \{L, R\}$ let h^D be the 1-cell with end e^D . Let h^{Opp} be the third side of t , opposite the distinguished prong π .

We need the observation that no two sides of t are identified, because this would force the existence of at least two vertices in C . It follows that each cell of C opposite a side of t is distinct from t .

Given a cell division C and a 1-cell h of C such that the cells on either side of h are distinct and at most one is punctured, removal of h produces a cell division C' , and there is a natural *coalescence map* $p: \text{Pr}(C) \rightarrow \text{Pr}(C')$ defined as follows: for each prong $\pi = (e, e') \in \text{Pr}(C)$ such that $e, e' \notin \text{Ends}(h)$, it follows that $e, e' \in \text{Ends}(C')$ and that $(e, e') \in \text{Pr}(C')$, and we define $p(\pi) = \pi$. For each end e of h , let $(e', e), (e, e'') \in \text{Pr}(C)$ be the prongs incident to e' . It follows that $(e', e'') \in \text{Pr}(C')$, and we define $p(e', e) = p(e, e'') = (e', e'')$. Thus, the coalescence map is two-to-one on the four prongs incident to ends of h , and one-to-one otherwise.

Given a parity $D \in \{L, R\}$, we describe an operation $(C, \pi) \rightarrow (C', \pi')$ called a *good elementary move* or *GEM* of parity D as follows. Let c^D denote the cell of C which is attached to t along h^D . As observed above, $t \neq c^D$. Thus, when we remove h^D we obtain a cell division C'' , in which t and c^D join into a single 2-cell c'' of C'' . Define π'' to be the image of π under coalescence. Since $\pi'' \in \text{Pr}(c'')$ and c'' is not a triangle, it follows that (C'', π'') is not a CDP. However, there is up to isotopy a unique way to add a 1-cell h'' to C'' , so that π'' is a prong of a nonpunctured

triangle t'' whose opposite side is h'' ; let C' be this cell division. There is moreover a unique prong $\pi' \in \text{Pr}(C')$ whose image under the coalescence map $\text{Pr}(C'') \rightarrow \text{Pr}(C')$ is equal to π'' . Under these circumstances we say that $(C, \pi) \rightarrow (C', \pi')$ is a GEM of parity D .

We have proved the first part of the following:

Lemma 11.3.1. *For each CDP (C, π) with no secondary prong and for each parity $D \in \{L, R\}$:*

Outgoing GEM *There exists a CDP (C', π') such that $(C, \pi) \rightarrow (C', \pi')$ is a good elementary move of parity D .*

Incoming GEM *There exists a CDP (C'', π'') such that $(C'', \pi'') \rightarrow (C, \pi)$ is a good elementary move of parity D .*

For both the incoming and outgoing GEMS, the isotopy class of (C', π') depends only on the isotopy class of (C, π) and the parity D .

Proof. To prove the second part one simply works the construction in the reverse: remove the 1-cell of C opposite π and identify the appropriate 1-cell to insert to produce C'' . ◇

Good elementary moves: with secondary prong. Next we consider good elementary moves in contexts where a secondary prong exists. Things are more complicated, because the location of the secondary prong can effect whether incoming and outgoing GEMs exist. In order to understand this, given a good elementary move $(C, \pi) \rightarrow (C', \pi')$ without secondary prong we explain the relation between prongs of C and prongs of C' .

Given a stable CDP with secondary prong $(C, (\pi, \rho))$, if we remove ρ then (C, π) may or may not be a stable CDP, depending on whether the unique 0-cell v of C is a puncture; to proceed with the discussion we shall, if necessary, imagine that v is a puncture, so we can imagine that (C, π) is a CDP.

Let t be the distinguished triangle of (C, π) , and define $\text{Pr}^*(C, \pi) = \text{Pr}(C) - \text{Pr}(t)$. Any secondary prong must be an element of $\text{Pr}^*(C, \pi)$. Consider a GEM $(C, \pi) \rightarrow (C', \pi')$ of parity D , and let C'' be the intervening cell division, obtained from C by removing h^D and from C' by removing h'^{Opp} . The coalescence map $\text{Pr}(C, \pi) \rightarrow \text{Pr}(C'')$ is injective and its image is everything except for a single prong denoted π_C . The coalescence map $\text{Pr}(C', \pi') \rightarrow \text{Pr}(C'')$ is injective and its image is everything except for a single prong denoted $\pi_{C'}$. Note that $\pi_C \neq \pi_{C'}$. The unique inverse image of π_C in $\text{Pr}^*(C, \pi)$ is denoted π_{in}^D ; it is incident to the distinguished prong π along the end e^D . The unique inverse image of $\pi_{C'}$ in $\text{Pr}^*(C', \pi')$ is denoted π_{out}^D . By composing the coalescence map $\text{Pr}_{out}^{*d} = \text{Pr}^*(C, \pi) - \{\pi_{out}^D\} \rightarrow \text{Pr}(C'')$ with

the inverse of the coalescence map $\text{Pr}_{in}^{*d} = \text{Pr}^*(C, \pi) - \{\pi_{in}^D\} \rightarrow \text{Pr}(C'')$ we obtain a bijection $\text{Pr}_{in}^{*d}(C, \pi) \leftrightarrow \text{Pr}_{out}^{*d}(C', \pi')$.

Consider now a secondary prong $\rho \in \text{Pr}^*(C, \pi)$, so that $(C, (\pi, \rho))$ is a stable CDP. If $\rho \neq \pi_{in}^D$ then let ρ' be the corresponding prong in $\text{Pr}^*(C', \pi')$, and we define $(C, (\pi, \rho)) \rightarrow (C, (\pi, \rho'))$ to be a GEM of parity D . We have proved the first part of:

Lemma 11.3.2. *Given a stable CDP (C, Π) with distinguished prong π and secondary prong ρ , and given $D \in \{L, R\}$, the following hold:*

Outgoing GEM *If $\rho \neq \pi_{in}^D$ then there exists a CDP (C', π') such that $(C, \pi) \rightarrow (C', \pi')$ is a good elementary move of parity D .*

Proof. The second part follows easily from the definitions. ◇

Good elementary moves and one cusp splittings. Given a good elementary move $(C, \Pi) \rightarrow (C', \Pi')$, we shall show that the dual one cusp train tracks $\tau = \tau^\perp(C, \Pi)$ and $\tau' = \tau^\perp(C', \Pi')$ are related by one cusp elementary moves.

Here is the key observation which justifies the concept of “stability”:

Lemma 11.3.3. *Given a good elementary move $(C, \pi) \rightarrow (C', \pi')$ of parity D , letting (τ, v) and (τ', v') be the dual one cusp train tracks of (C, π) , (C', π') respectively, there is a sequence*

$$(\tau', v') \succ (\tau'', v'') \succ (\tau, v)$$

consisting of a one cusp elementary splitting of parity D followed by a one cusp slide. ◇

Proof. The proof is a picture, given in Figure ?? ◇

THERE ARE NOT QUITE RIGHT AS STATED, BECAUSE OF THE REFACTORIZATION THAT IS NEEDED: “SPLITS LAST” RATHER THAN “SPLITS FIRST”. AND EVEN THEN IT’S NOT QUITE RIGHT BECAUSE SOME GEMS ARE ONLY SLIDE MOVES, DEPENDING ON WHERE THE SECONDARY PRONG IS LOCATED.

Corollary 11.3.4. *If (τ', v') is a stable one cusp train track of type I , and if $(\tau', v') \succ (\tau, v)$ is the one cusp splitting of parity D , then (τ, v) is stable and the elementary factorization of $(\tau', v') \succ (\tau, v)$ consists of a one cusp slide followed by a one cusp elementary splitting of parity D .* ◇

Corollary 11.3.5. *If $(\tau_0, v_0) \succ (\tau_1, v_1) \succ \dots$ is a one cusp splitting sequence, and if (τ_I, v_I) is stable for some I , then (τ_i, v_i) is stable for all $i \geq I$.* ◇

TO DO:

- Start by describing elementary moves with no secondary prong, “split slide” moves.
- Describe the map on prong sets outside of the distinguished triangle.
- Describe the special prongs in the domain and range.
- Now introduce the secondary prong.
- One thing I’m worried about: did I leave out considerations of the secondary prong in the one switch killing criterion????
- Describe elementary moves, classified according to their effect on the dual train track:

Slide type: Secondary cusp incident to h^{Opp} ; a slide leading directly to another type I stable train track.

Split type: Secondary cusp incident to a cusp incident to h^{Opp} ; an elementary split leading directly to another type I stable train track.

Split–slide type: Neither of the above; the dual branch $b(h^{\text{Opp}})$; an elementary split leading to a type II stable train track; then a slide leading to a type I. stable train track.

- Explain how the type II stable train tracks are the ones that are truly stable.
- Prove that every separatrix has a stable one cusp train track, and hence by stable equivalence *all* of the train tracks that appear in a one cusp train track axis are stable.
- Conclusion: Proposition: a one cusp splitting sequence consists eventually of stable train tracks; but are they type II? But when the secondary prong intervenes, you can go directly from a type I to another type I without passing a type II, so that’s not right either ...

The canonical diagraph.**Canonical pseudo-Anosov loops.**

TO DO:

- Deciding the canonical killing criterion using finite automata: canonical digraphs and degeneracy graphs.
- We have to give a simple version of the canonical killing criterion for one cusp splitting sequences. Ultimately it will be dual to the version for one switch splitting sequences. I wonder if it can be proved equivalent? Except that they go in the opposite direction ...

Examples.

TO DO:

- Many examples, taken from [Mos93].

One cusp — one switch duality. A START ON THE CANONICAL DIGRAPH

Let $\tilde{\Gamma}_c$ be the graph of one cusp train tracks and one cusp splittings. To be precise, the vertices of $\tilde{\Gamma}_c$ are comb equivalence classes of pairs (τ, v) where τ is a generic, filling, recurrent train track and $v \in \text{cusps}(\tau)$, and where (τ, v) is comb equivalent to (τ', v') if τ is equivalent to τ' and v, v' correspond under the bijection $\text{cusps}(\tau) \leftrightarrow \text{cusps}(\tau')$ induced by the comb equivalence. Let $[\tau, v]$ denote the comb equivalence class of (τ, v) . Define a directed edge $[\tau, v] \succ [\tau', v']$ if there are representatives $(\tau, v), (\tau', v')$ of these comb equivalence classes so that $(\tau, v) \succ (\tau', v')$ is a one cusp splitting.

We claim that the natural action of \mathcal{MCG} on the directed graph $\tilde{\Gamma}_c$ is free. It suffices to prove that the action on vertices is free. Suppose that $\Phi[\tau, v] = [\tau, v]$, and so $\phi(\tau, v) = (\tau', v')$ where (τ, v) and (τ', v') are comb equivalent and ϕ is a representative of Φ . We must prove that ϕ is isotopic to the identity; this is a strengthening of Lemma 3.15.3. Choose a tie bundle ν over τ and let $\psi: S \rightarrow S$ be a carrying inclusion of τ' into ν . Replacing ϕ with $\psi \circ \phi$ we may assume that $\phi \dots$

check [Mos86] to see how the pictures are drawn, with the transverse orientation upward or downward. It seems to make a difference in how the chord diagrams are drawn.

TO DO:

- Put this into the “comb equivalence” section.
- Proof: show directly that two comb equivalent but nonisotopic train tracks are

WHERE WILL THIS GO???

- One cusp and two cusp expansions. Structure of circular expansion complexes.
 - Show that the one cusp property is stable under splitting along any splitting involving that cusp. Thus we get one cusp splitting sequences, as expansions of each train track.
- The 2-strata suffice to classify up to conjugacy; a simpler structure to keep track of, rather than the whole circular expansion complex.
- The 1-strata do not suffice to classify up to conjugacy, however, they do classify almost conjugacy, as long as you keep track of multiplicity.
- Relate one cusp expansions to one sink expansions; connection with [Mos93]. Proof that every pseudo-Anosov homeomorphism with a fixed separatrix has a one sink expansion and a one cusp expansion.
- Enumeration of pseudo-Anosov conjugacy classes, using the canonical killing criterion.
- OLD DISCUSSION OF CANONICAL DIGRAPHS AND DEGENERACY DIGRAPHS

Given a finite splitting sequence $\tau_0 \succ \cdots \succ \tau_n$ one wants to know whether the rational killing criterion holds, i.e. whether $\mathcal{MF}^\perp(\tau_0) \times \mathcal{MF}(\tau_n) \subset \mathcal{FP}$. From the definition it is clear that this question is decidable, but we shall show that it is decidable in a very elementary sense, using finite deterministic automata. To be precise, consider a directed graph $\tilde{\Gamma}(S)$ whose vertices are isotopy classes of train tracks and whose directed edges are isotopy classes of splittings. The mapping class group $\mathcal{MCG}(S)$ acts on $\tilde{\Gamma}(S)$ by directed isomorphisms, with quotient digraph $\Gamma(S) = \tilde{\Gamma}(S)/\mathcal{MCG}(S)$; this is the *canonical digraph* of S (actually, this definition doesn't quite work because the action of $\mathcal{MCG}(S)$ on $\tilde{\Gamma}(S)$ can have nontrivial stabilizer groups; we shall give a more formal definition which gets around this problem). The canonical digraph $\Gamma(S)$ is finite: there are only finitely many orbits of isotopy classes of train tracks and splittings, up to the action of $\mathcal{MCG}(S)$. A splitting sequence $\tau_0 \succ \cdots \succ \tau_n$ descends to a directed path $[\tau_0] \succ \cdots \succ [\tau_n]$ in $\Gamma(S)$. Now the decision problem for the rational killing criterion may be precisely stated and solved:

Theorem 11.3.6 (Degeneracy Theorem). *There exists a finite directed graph $\Gamma'(S)$ and a directed map $\Gamma'(S) \mapsto \Gamma(S)$ such that a splitting sequence $\tau_0 \succ \cdots \succ \tau_n$ satisfies the rational killing criterion if and only if the directed path $[\tau_0] \succ \cdots \succ [\tau_n]$ in $\Gamma(S)$ does not lift to $\Gamma'(S)$.*

The directed graph $\Gamma'(S)$ is called the *degeneracy digraph*. It will also be constructed as the quotient $\tilde{\Gamma}'(S)/\mathcal{MCG}(S)$ of some directed graph $\tilde{\Gamma}'(S)$ on which the mapping class group acts, and the directed map $\Gamma'(S) \mapsto \Gamma(S)$ will be the quotient of an $\mathcal{MCG}(S)$ -equivariant finite-to-one directed map $\tilde{\Gamma}'(S) \mapsto \tilde{\Gamma}(S)$. The directed graph $\tilde{\Gamma}'(S)$ will, in a sense, keep track of potential degeneracies. That is, given a train track τ , a *degeneracy* of τ is an essential curve which has zero intersection with some element of both $\mathcal{MF}^\perp(\tau)$ and some element of $\mathcal{MF}(\tau)$. We shall show that $\tau_0 \succ \cdots \succ \tau_n$ fails the rational killing criterion if and only if there is a certain degeneracy of each τ_i , $i = 1, \dots, n$, which somehow survives each splitting in the sequence. Moreover, we will show that for each train track τ one need only keep track of a special finite set of degeneracies: the sequence $\tau = \tau_0 \succ \cdots \succ \tau_n$ fails the rational killing criterion if and only if there a degeneracy of τ in the special finite set which survives the entire splitting sequence.

Before launching into the formal definition and proof we shall show what happens on the torus.

11.4 Example: the once-punctured torus (in progress)

11.5 Formal definition of the canonical digraph (in progress)

11.6 Proof of the Degeneracy Theorem (in progress)

11.7 Further examples (in progress)

- The five punctured sphere.

12 The conjugacy problem in \mathcal{MCG} (in progress)

In this section we augment the results of the previous section with work of Nielsen and a few additional ideas to obtain complete conjugacy invariants in \mathcal{MCG} . We then show how to compute these invariants for any mapping class, thus solving the conjugacy problem in \mathcal{MCG} .

12.1 Conjugacy invariants of finite order mapping classes, after Nielsen (in progress)

Finite order conjugacy classes in \mathcal{MCG} were completely classified by Nielsen in a series of papers. Nielsen proved that each finite order conjugacy class in \mathcal{MCG} is represented by a finite order homeomorphism S [Nie86a]. Furthermore, that representative is unique up to topological conjugacy (reference??). Thus, classifying finite order conjugacy classes in \mathcal{MCG} is equivalent to classifying finite order, orientation preserving homeomorphisms of S up to topological conjugacy. The latter classification was achieved by Nielsen in [Nie86b]. The invariants are simply the local rotation numbers attached to the nonregular periodic points ...

12.2 Conjugacy invariants of algebraically finite mapping classes, after Nielsen (in progress)

A mapping class is *algebraically finite* if it has a reducing system on whose complement the mapping class has finite order. Nielsen gave a complete conjugacy classification of algebraically finite mapping classes [Nie86c] ...

12.3 Conjugacy invariants of the Thurston reduction (in progress)

12.4 Complete conjugacy invariants (in progress)

12.5 Tailoring the Bestvina-Handel algorithm (in progress)

Bestvina and Handel described in [BH95] an algorithm which, given $\phi \in \mathcal{MCG}$, decides whether ϕ is finite order, reducible, or pseudo-Anosov, and in the latter case produces an invariant train track of ϕ . We can augment their algorithm to obtain an algorithm that computes complete conjugacy invariants for any mapping class.

12.6 Two cusp expansion complexes.

Given a pseudo-Anosov mapping class Φ , the Bestvina-Handel algorithm outputs an invariant train track τ for Φ , and in fact the algorithm produces a splitting sequence

from τ to $\Phi(\tau)$. From this we need to construct a complete conjugacy invariant for Φ , perhaps a circular expansion complex. We shall focus more on the more easily computed 2-strata of the circular expansion complex. This requires a study of the 2-strata of a general expansion complex, which requires an understanding of “two cusp expansions” of an arational measured foliation.

References

- [AF91] P. Arnoux and A. Fathi, *Un exemple de difféomorphisme pseudo-Anosov*, C. R. Acad. Sci. Paris Sér. I Math. **312** (1991), no. 2, 241–244.
- [Bau92] M. Bauer, *Examples of pseudo-Anosov homeomorphisms*, Trans. AMS **330** (1992), no. 1, 333–359.
- [BH95] M. Bestvina and M. Handel, *Train tracks for surface homeomorphisms*, Topology **34** (1995), 109–140.
- [BLM83] J. Birman, A. Lubotzky, and J. McCarthy, *Abelian and solvable subgroups of the mapping class groups*, Duke Math. J. **50** (1983), no. 4, 1107–1120.
- [Bon97a] Francis Bonahon, *Geodesic laminations with transverse Hölder distributions*, Ann. Sci. École Norm. Sup. (4) **30** (1997), no. 2, 205–240.
- [Bon97b] Francis Bonahon, *Transverse Hölder distributions for geodesic laminations*, Topology **36** (1997), no. 1, 103–122.
- [Bon99] F. Bonahon, *Closed curves on surfaces*, available at <http://math.usc.edu/~fbonahon>, 1999.
- [Bon01] Francis Bonahon, *Geodesic laminations on surfaces*, Laminations and foliations in dynamics, geometry and topology (Stony Brook, NY, 1998), Contemp. Math., vol. 269, Amer. Math. Soc., Providence, RI, 2001, pp. 1–37.
- [CB] A. Casson and S. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, Handwritten notes, University of Texas at Austin, Volume 1 Fall 1982, Volume 2 Spring 1983.
- [CB88] A. Casson and S. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, London Mathematical Society Student Texts, vol. 9, Cambridge University Press, 1988.
- [Fat87] Albert Fathi, *Dehn twists and pseudo-Anosov diffeomorphisms*, Invent. Math. **87** (1987), no. 1, 129–151.
- [Fat92] A. Fathi, *Démonstration d'un théorème de Penner sur la composition des twists de Dehn*, Bull. Soc. Math. France **120** (1992), no. 4, 467–484.

- [FLP⁺79] A. Fathi, F. Laudenbach, V. Poenaru, et al., *Travaux de Thurston sur les surfaces*, Astérisque, vol. 66–67, Société Mathématique de France, 1979.
- [FM99] B. Farb and L. Mosher, *Precompact Teichmüller geodesics and Schottky subgroups of mapping class groups*, in preparation, 1999.
- [Gil76] Jane Gilman, *On conjugacy classes in the Teichmüller modular group*, Michigan Math. J. **23** (1976), no. 1, 53–63.
- [GM91] F. Gardiner and H. Masur, *Extremal length geometry of Teichmüller space*, Complex Variables Theory Appl. **16** (1991), 209–237.
- [Gol84] W. Goldman, *The symplectic nature of fundamental groups of surfaces*, Adv. in Math. **54** (1984), 200–225.
- [Hae91] A. Haefliger, *Complexes of groups and orbihedra*, Group theory from a geometrical viewpoint (Trieste, 1990), World Sci. Publishing, River Edge, NJ, 1991, pp. 504–540.
- [Hir76] M. Hirsch, *Differential topology*, Springer, 1976.
- [HTC96] H. Hamidi-Tehrani and Z.-H. Chen, *Surface diffeomorphisms via train tracks*, Topology Appl. **73** (1996), 141–167.
- [HU79] J. E. Hopcroft and J. D. Ullman, *Introduction to automata theory, languages, and computation*, Addison Wesley, 1979.
- [Kea75] M. Keane, *Interval exchange transformations*, Math. Z. **141** (1975), 25–31.
- [Kea77] M. Keane, *Non-ergodic interval exchange transformations*, Israel J. Math. **26** (1977), no. 2, 188–196.
- [Ker85] S. Kerckhoff, *Simplicial systems for interval exchange maps and measured foliations*, Ergodic Th. & Dyn. Sys. **5** (1985), no. 2, 257–271.
- [KN76] H. Keynes and D. Newton, *A “minimal”, non-uniquely ergodic interval exchange transformation*, Math. Z. **148** (1976), no. 2, 101–105.
- [Knu73] D. Knuth, *The art of computer programming, vol. 1, Fundamental algorithms*, Addison Wesley, 1973.
- [Mas82] H. Masur, *Interval exchange transformations and measured foliations*, Ann. of Math. **115** (1982), 169–200.

- [Mas92] H. Masur, *Hausdorff dimension of the set of nonergodic foliations of a quadratic differential*, Duke Math. J. **66** (1992), no. 3, 387–442.
- [McM02a] C. T. McMullen, *Billiards and Teichmüller curves on Hilbert modular surfaces*, preprint, 2002.
- [McM02b] C. T. McMullen, *Teichmüller geodesics of infinite complexity*, preprint, 2002.
- [MM99] H. Masur and Y. Minsky, *Geometry of the complex of curves, I. Hyperbolicity*, Invent. Math. **138** (1999), no. 1, 103–149.
- [Mos83] L. Mosher, *Pseudo-Anosovs on punctured surfaces*, Ph.D. thesis, Princeton University, 1983.
- [Mos86] L. Mosher, *The classification of pseudo-Anosovs*, Low-dimensional topology and Kleinian groups (Coventry/Durham, 1984), LMS Lecture Note Ser., vol. 112, 1986, pp. 13–75.
- [Mos93] L. Mosher, *Topological invariants of measured foliations*, preprint, 1993.
- [Mos95] L. Mosher, *Mapping class groups are automatic*, Ann. of Math. **142** (1995), 303–384.
- [Mos96] L. Mosher, *A user's guide to the mapping class group: once punctured surfaces*, Geometric and computational perspectives on infinite groups, DIMACS Ser. Discrete Math. Theoret. Comput. Sci., Amer. Math. Soc., 1996, pp. 101–174.
- [Mos03] L. Mosher, *What is a train track?*, Notices of the AMS ??? (2003), ???, to appear.
- [MS91] H. Masur and J. Smillie, *Hausdorff dimension of sets of nonergodic measured foliations*, Ann. of Math. **134** (1991), no. 3, 455–543.
- [Nie86a] J. Nielsen, *Mapping classes of finite order*, Jacob Nielsen: Collected Mathematical Papers, Volume 2, Birkhäuser, 1986, Translation by John Stillwell of: Abbildungsklassen endlicher Ordnung. Acta. Math. 75 (1942), 23–115., pp. 65–103.
- [Nie86b] J. Nielsen, *The structure of periodic surface transformations*, Jacob Nielsen: Collected Mathematical Papers, Volume 2, Birkhäuser, 1986, Translation by John Stillwell of: Die Struktur periodischer Transformationen von Flächen. Math.-fys. Medd. Danske Vid. Selsk. 15, nr. 1 (1937), pp. 65–103.

- [Nie86c] J. Nielsen, *Surface transformation classes of algebraically finite type*, Jacob Nielsen: Collected Mathematical Papers, Volume 2, Birkhäuser, 1986, Originally published in Mat.-fys. Medd. Danske Vid. Selsk. 21, nr. 2 (1944), pp. 65–103.
- [Pen88] R. C. Penner, *A construction of pseudo-Anosov homeomorphisms*, Trans. AMS **310** (1988), 179–197.
- [Pen92] R. C. Penner, *Combinatorics of train tracks*, Annals of Mathematics Studies, vol. 125, Princeton University Press, 1992.
- [PP87] A. Papadopolous and R. Penner, *A characterization of pseudo-Anosov foliations*, Pac. J. Math. **130** (1987), 359–377.
- [Rau79] G. Rauzy, *Échanges d'intervalles et transformations induites*, Acta Arith. **34** (1979), no. 4, 315–328.
- [Spa81] E. Spanier, *Algebraic topology*, Springer, New York—Berlin, 1981.
- [Sta68] J. Stallings, *On torsion free groups with infinitely many ends*, Ann. of Math. **88** (1968), 312–334.
- [Tak94] I. Takarajima, *On a construction of pseudo-anosov diffeomorphisms by sequences of train tracks*, Pacific J. Math **166** (1994), 123–191.
- [Thu87] W. P. Thurston, *Geometry and topology of 3-manifolds*, Princeton University Lecture Notes, 1987.
- [Thu88] W. P. Thurston, *On the geometry and dynamics of diffeomorphisms of surfaces*, Bull. AMS **19** (1988), 417–431.
- [Vee78] W. A. Veech, *Interval exchange transformations*, J. Analyse Math. **33** (1978), 222–272.
- [Vee82] W. A. Veech, *Gauss measures for transformations on the space of interval exchange maps*, Ann. of Math. (2) **115** (1982), no. 1, 201–242.
- [Vee89] W. A. Veech, *Teichmüller curves in moduli space, Eisenstein series and an application to triangular billiards*, Invent. Math. **97** (1989), no. 3, 553–583, Erratum: volume 103 number 2, 1991, page 447.
- [Wal82] Peter Walters, *An introduction to ergodic theory*, Graduate Texts in Mathematics, vol. 79, Springer-Verlag, New York, 1982.

Index

- C-splitting arc, 104, 140

- C, 98
- \nearrow , 70
- $\mathcal{MF}(\tau)$, 80
- $\mathcal{PMF}(\tau)$, 80
- \approx , 94
- \mathcal{C} , 53
- $\cdot \setminus \cdot$, 219
- D, 98
- \bar{D} , 98
- \nearrow^F , 106
- \nearrow^H , 106
- L, 98
- \mathcal{MCG} , 33
- $\nu(\xi)$, 115
- \bar{d} , 137
- P, 32
- R, 98
- $\sigma(\xi)$, 115
- \sphericalangle , 92
- γ , 97
- S, 32
- \bar{S} , 32
- T, 32
- ancestor
 - of a subtrack, 137
 - of splitting cycle, 141
- arational, 9, 13, 37
 - partial, 45
- awl, 84

- basin, 59
- branch, 51
 - post-splitting, 98
 - sink, 58
 - source, 58
 - transition, 58
- canonical model, 43
- carrying
 - bijection, 73
 - canonical, 73
 - full
 - of a measured foliation, 73
 - of a train track, 106
 - homotopic, 106
 - injection, 70, 73
 - map
 - of a measured foliation, 73
 - of a train track, 70
 - of a class in \mathcal{MF} , 80
- carrying equivalence, 92, 95
- carrying image, 76
- central splitting, 103
 - elementary, 98
- Choquet simplex, 15
- closed curve
 - essential, 32
 - peripheral, 32
 - trivial, 32
- collapsing rectangle, 98
- comb equivalence, 94, 95
 - of splittings, 101
- combinatorial equivalence, 54
- combinatorial invariant, 54
- combinatorial length, 51
- combinatorial type, 109
- complete, 84
 - expanding separatrix family, 118
 - splitting sequence, 112
- complete combing, 95
- completion, 53
- conformal structure, 32

- corner
 - reflex, 52
 - regular, 52
- cross bundle, 53
- crossing, 99, 137
 - partial, 142
- cusp
 - reflex, 52
 - regular, 52
- cutting out, 219
- descendant
 - of a subtrack, 137
 - of splitting cycle, 141
 - of C-splitting arc, 148
- diagonal, 84
- diagonal extension, 85
- dumbbell, 89
- elementary factorization, 122
- end, 51
- equivalence of measured foliations, 13, 42, 350
 - topological, 22, 164
 - unmeasured, 15, 22, 42, 124
- essential
 - curve family, 112
- essential subsurface, 56
- fiber of carrying map, 71
- filling, 84
- fold map, 104
- forest, 194
 - collapse map, 194
 - unmarked, 204
 - collapsible, 194
- fractional Dehn twist, 251
- generic
 - switch, 51
 - train track, 14, 51
- hyperbolic structure, 32
- infinitely split, 130
- inflow, 58
- invariant weight, 81
- isotopy, 54
- join, 133
- killing
 - of splitting cycle, 141
 - of subtracks, 137
 - of C-splitting arc, 148
- killing criterion
 - canonical, 149
 - iterated canonical, 149
 - iterated rational, 134
 - rational, 134, 147
 - splitting cycle, 135, 147
 - subtrack, 135, 146
- leaf, 36
 - closed, 35
 - cycle, 36
 - nonsingular, 35
- leaf cycle, 13
- leaf segment, 36
 - nonsingular, 36
- legal, 85
- line of descent, 152
- mapping class group, 33
- measured foliation, 35, 36
 - arational, 37
 - canonical, 43
 - component, 44
 - partial, 39
 - total, 39
- moduli space, 33
- multi-ended side, 51
- one-ended side, 51

- outflow, 58
- overlay map, 53
- parity, 98
 - of a slide move, 296
 - of a transition branch, 296
- parity splitting
 - elementary, 98
- perturbable transverse orientation, 36
- pinch map, 104
- pinch points, 41
- pinched subsurface, 41
- pinching, 17
- polar cross bundle, 54
- postsplitting branch, 104
- pretrack, 51
- pseudo-Anosov, 49
- puncture, 32
- quadratic differential, 33
- recurrent, 88
- reducible, 49
- relative isotopy, 102
 - of splitting cycles, 141
- saddle connection, 13, 36
- semigeneric
 - switch, 51
- separatrix
 - finite, 36
 - infinite, 36
 - proper, 39
- shunt, 116
- singularity
 - canonical, 43
- singularity type of
 - arational measured foliation, 44
 - filling train track, 84
 - pseudo-Anosov homeomorphism, 49
 - twist system, 250
- sink
 - branch, 58
 - loop, 59
- slide, 92
 - one cusp, 164
- source branch, 58
 - LL, RR, LR, 59
- splitting
 - central, 97, 104
 - elementary, 97
 - forced, 100
 - general, 101
 - left, 97
 - notation, 98
 - one cusp, 165
 - one sink, 259
 - parity, 97
 - right, 97
 - wide, 102
- splitting arc, 102, 104
 - c, 140
- splitting circuit, 234, 244, 248
- splitting cycle, 140
- splitting sequence, 14, 112
 - one cusp, 165
 - periodic, 235
- stable under extension, 146
- stabilizer
 - of train track, 111
- subtrack, 56
- support
 - of a partial measured foliation, 39
 - of a train track, 56
- surface with cusps, 53
- surface-with-corners, 52
- survival
 - of splitting cycle, 141
 - of subtracks, 137
 - of C-splitting arc, 148
- survive, 261

- switch, 51
 - generic, 51
 - semigeneric, 51
- switch condition, 81

- tie, 68
- tie bundle, 68
- track with terminals, 51, 216
 - transient, 217
- train line, 61
- train path, 51
- train ray, 61
- train track, 54
 - canonical invariant, 235, 237
 - decorated, 171
 - expansion, 14, 112
 - invariant, 235
 - one cusp, 110, 164, 247
 - one sink, 258
 - stable, 291, 314
- transition
 - branch, 58
 - forest, 59
 - tree, 59
- twist system, 250

- weight function, 81
- Whitehead
 - collapse, 36
 - equivalence, 36

Glossary (in progress)

Equivalence. The equivalence relation on measured foliations on S is generated by isotopy and Whitehead moves. The equivalence relation on partial measured foliations is generated by isotopy, Whitehead moves, and partial fulfillment. Measured foliation space $\mathcal{MF} = \mathcal{MF}(S)$ may be described either as equivalence classes of measured foliations or equivalence classes of partial measured foliations. Unmeasured equivalence of partial measured foliations is generated by ordinary equivalence and by the relation of having the same underlying unmeasured partial foliation. Topological equivalence of partial measured foliations is generated by unmeasured equivalence and the relation of differing by an element of $\text{Homeo}_+(S)$.

LEE MOSHER:

DEPARTMENT OF MATHEMATICS, RUTGERS UNIVERSITY, NEWARK

NEWARK, NJ 07102

MOSHER@ANDROMEDA.RUTGERS.EDU