TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 370, Number 8, August 2018, Pages 5635–5704
http://dx.doi.org/10.1090/tran/7152
Article electronically published on March 22, 2018

DETECTING GEOMETRIC SPLITTINGS IN FINITELY PRESENTED GROUPS

NICHOLAS W. M. TOUIKAN

ABSTRACT. We present an algorithm which given a presentation of a group G without 2-torsion, a solution to the word problem with respect to this presentation, and an acylindricity constant κ outputs a collection of tracks in an appropriate presentation complex. We give two applications: the first is an algorithm which decides if G admits an essential free decomposition; the second is an algorithm which, if G is relatively hyperbolic, decides if it admits an essential elementary splitting.

1. INTRODUCTION

An important group invariant is whether or not it splits as a certain type of graph of groups. In this paper we prove an algorithmic analogue of Sela's κ-acylindrical super accessibility [Sel97, Theorem 4.3] for the class of one edged κ-acylindrical geometric splittings. In particular the main result, Theorem B, gives an algorithm that produces a list that contains a representative of every one edged κ-acylindrical geometric splitting of $\pi_1(C)$, up to equivalence in $\mathrm{Aut}(\pi_1(C))$. We give some corollaries of this theorem.

Theorem A. *There is an algorithm that takes as input a finite presentation $\langle X \mid R\rangle$ of a group G without 2-torsion and a solution to the word problem with respect to this presentation and decides whether or not the group G admits an* essential *free decomposition, i.e., a free decomposition*

$$G = H_1 * H_2$$

with $H_1 \neq \{1\} \neq H_2$.

This theorem is proved in Section 1.6. As a consequence we have the following corollary, whose proof we leave as an exercise. (Hint: if we can solve the word problem, then we can decide if a finitely generated group is abelian and we can decide, given a finite presentation, if an abelian group is cyclic.)

Corollary 1.1. *Let $G = \langle X \mid R\rangle$ be as in the statement of Theorem* A*. Then we can find a Grushko decomposition for G.*

In a sense, aside from the no 2-torsion assumption, this is the strongest result of this type possible: the restrictions on the input are as minimal as can be reasonably expected. This result also extends all previously known results (at least in the case without 2-torsion), which we now briefly survey.

Received by the editors March 8, 2011, and, in revised form, June 18, 2015, August 17, 2016, and December 10, 2016.

2010 *Mathematics Subject Classification.* Primary 20E06, 20F10; Secondary 57M05, 20E08.

Diao and Feighn in [DF05] showed how to find a Grushko decomposition of a fundamental group of a graph of free groups. Their techniques rely on Whitehead methods refined by Gersten and group actions on square complexes. Kharlampovich and Miasnikov in [KM05] showed how to find a Grushko decomposition of a fully residually free group by running their Elimination Process: the free decomposition becomes apparent by "separating the variables" in the defining equations.

Even in the presence of 2-torsion, Dahmani and Groves in [DG08a] are able to detect free splittings of certain relatively hyperbolic groups by generalizing an unpublished algorithm for hyperbolic groups due to Gerasimov. Their approach is to decide some connectivity criterion of the boundary of toral relatively hyperbolic groups. Our work implies this result in the 2-torsion-free case. Another result [GW09] due to Groves and Wilton, which works in the presence of 2-torsion, is that given a finite presentation of a group G and a solution to the word problem with respect to that presentation we can decide if G is free. If G is without 2-torsion, then this is an easy consequence of Corollary 1.1. At the end of Section 1.6 we will explain how to obtain this result in complete generality from the work in this paper.

It is also worth noting that Casals-Ruiz and Kazachkov used methods related to ours to describe solutions to equations over free products [CRK11].

The algorithm given in this paper is also well-suited to relatively hyperbolic groups. We have the following immediate corollary to Theorem C, which is proved in Section 1.7.

Corollary 1.2. *We can decide if a torsion-free relatively hyperbolic group with polycyclic parabolics has a trivial elementary JSJ decomposition, in the sense of* [Bow98, Bow01] *or* [GL11, Theorem 4].

This generalizes a result of Dahmani and Groves in [DG08b] for toral relatively hyperbolic groups in two ways. Firstly, the present approach works for a larger, wider class of groups. Secondly, it can detect splittings that are not in the class $\mathcal{Z}_{\max}$. Both of these earlier limitations arise from the fact that all previous algorithms to detect splittings in relatively hyperbolic groups depend on "equational" methods. In particular they will not work with nilpotent parabolics since we can't solve equations over nilpotent groups [Rom79] and they can't detect non-$\mathcal{Z}_{\max}$ splittings since Dehn twists around such groups give trivial automorphisms.

In an earlier preprint, *Effective Grushko decomposition* (`http://arxiv.org/abs/0906.3902v1`), the author claimed Theorem A *without* the no 2-torsion hypothesis. There is a gap in that proof: the argument is incomplete because the author did not take Möbius strips into consideration. In the second version of this paper we modified the argument so that it can handle κ-acylindricity and the existence of Möbius strips, at the cost of having to exclude 2-torsion. Otherwise, there was no gap in the second version of this paper, but it was horribly written. This third version attempts to rectify the issue and has more figures.

1.1. **Outline of the paper.** First we will give the basic definitions and results needed to make sense of the statement of Theorem B. After stating it and discussing some of its limitations, we apply it to detect free decompositions of finitely presented groups and elementary splittings of relatively hyperbolic groups. These applications also serve the role of providing a "tutorial" on how to use Theorem B.

Next we will define band complexes which are similar, but not identical, to the band complexes in [BF95]. Instead of using them to study minimal foliations, we will treat them as combinatorial objects to study Dunwoody patterns (see Definition 1.4). Next we will define transformations done to band complexes and tracks they carry. This constitutes the Rips machine of [BF95].

The Rips machine is designed to study a single lamination in a cell complex. The elimination process, inspired from works of Makanin and Razborov [Mak82, Raz87] as read from [KM98], is a branching search algorithm that constructs a finite rooted directed tree that decides the existence of certain types of laminations, in our case, tracks. We will show how to construct this tree one level at a time and give an analogy with splitting sequences for surface train tracks. We will then define various inadmissibility criteria which will forbid the elimination tree from growing at certain nodes.

As usual in this business, it will be relatively easy to handle the thin/Levitt/7-10 case as well as the surface/quadratic/12 case. The real difficulty is in handling the superquadratic/axial/15 case, and this is where most of the new ideas in this paper reside.

Eventually we will have given sufficiently many inadmissibility criteria, including a periodicity bound, to force the elimination tree to be finite. The leaves of this tree will give us the output of the main algorithm.

1.2. **Patterns, tracks, and geometric splittings.** We take it for granted that the reader is comfortable with Bass-Serre theory. The best reference, especially for this paper, would be [SW79]. Another standard reference is [Ser03]. We also assume the reader is well acquainted with polygonal 2-complexes, their fundamental groups, and the actions of fundamental groups on universal covers by deck transformations.

The graph of groups $\mathbb{X}$ has underlying graph X. We will write G *splits as a graph of groups* $\mathbb{X}$ or even $\mathbb{X}$ *is a splitting of* G instead of writing "G is the fundamental group of the graph of groups $\mathbb{X}$." We will also use the action of a group G on a simplicial tree T and the corresponding splitting $\mathbb{X}$, where $X = G\backslash T$, interchangeably. Finally all trees are assumed to be minimal.

Convention 1.3. In order to be sure to avoid any pathologies, we will restrict ourselves to the piecewise linear category of topological spaces.

Throughout this paper C will be a polygonal 2-complex. If $f : X \to Y$ is a continuous map we denote its functorial image $f_\sharp : \pi_1(X) \to \pi_1(Y)$, which is well-defined up to conjugacy. Let $Y \subset X$ be connected cell complexes. Consider the natural map

$$Y \xhookrightarrow{\quad i \quad} X$$

$$\pi_1(Y) \xrightarrow{\quad i_\sharp \quad} \pi_1(X)$$

where i denotes the inclusion map. We denote

$$\mathrm{Gp}(Y) = i_\sharp(\pi_1(Y)),$$

which gives a well-defined conjugacy class in $\pi_1(X)$.

Definition 1.4. Let C be a polygonal 2-complex. A *pattern* $P \subset C$ is an embedded 1-complex such that:

(i) for every 2-cell $D \subset C$, $P \cap D$ is a (possibly empty) *finite* collection of closed arcs joining *distinct* sides of D;
(ii) P does not meet $C^{(0)}$;
(iii) P has a regular neighbourhood $N(P) \subset C$ homeomorphic to $P \times [-1, 1]$.

Definition 1.5. A connected component of a pattern is called a *track*.

This definition of a pattern is slightly non-standard in that the last condition implies that our pattern is 2-*sided*. 2-sidedness, however, should be standard because it implies that the pattern is locally separating into two components. The Seifert-van Kampen Theorem immediately implies that the decomposition (which is essentially a graph of spaces)

$$C = N(P) \cup C \setminus P \tag{1}$$

splits $\pi_1(C)$ as a graph of groups $\mathbb{X}^P$ where the vertex groups are given by $\mathrm{Gp}(C_i)$ and the edge groups are given by $\mathrm{Gp}(t_j)$, where the C_i denote the connected components of $C \setminus P$ and the t_j denote the tracks in P respectively.

Proposition 1.6. *Let $P \subset C$ be a pattern and let $\widetilde{P}$ be the lift of P in the universal cover $\widetilde{C}$ of C. Each connected component of $\widetilde{P}$ separates $\widetilde{C}$ into two components. This gives rise to the $\pi_1(C)$-tree $T(P, C)$, whose vertices are connected components of $\widetilde{C} \setminus \widetilde{P}$, whose edges are connected components of $\widetilde{P}$, and such that the edge $\widetilde{t}$ is adjacent to the vertex $\widetilde{C_i}$ if $\widetilde{t}$ is contained in the closure of $\widetilde{C_i}$. Thus $T(P, C)$ can be obtained by a $\pi_1(C)$-equivariant identification map*

$$\pi : \widetilde{C} \twoheadrightarrow T(P, C). \tag{2}$$

Proof. The lift $\widetilde{P} \subset \widetilde{C}$ of P is again a pattern in $\widetilde{C}$. This gives a decomposition of $\widetilde{C}$ as in (1) which expresses $\pi_1(\widetilde{C})$ as a graph of groups $\mathbb{Y}$. If some component of $\widetilde{P}$ is not separating the underlying graph Y of $\mathbb{Y}$ contains a cycle contradicting the fact that $\pi_1(\widetilde{C}) = 1$. It therefore follows that the graph Y is a tree $T(P, C)$, which is easily seen to be a $\pi_1(C)$-tree.

π is obtained by collapsing each track neighbourhood $N(\widetilde{t}) = \widetilde{t} \times [-1, 1] \twoheadrightarrow [-1, 1]$ and by collapsing every connected component of **closure** $\left(\widetilde{C} \setminus N(\widetilde{P})\right)$ to a point. □

This next proposition follows immediately by thinking about the action of $\pi_1(C)$ on $T(P, C)$ induced by deck transformations and the meaning of the Seifert-van Kampen Theorem or simply by thinking of (1) as a *graph of spaces decomposition* à la [SW79].

Proposition 1.7. *Let $P \subset C$ be a pattern. Then the action of $\pi_1(C)$ on $T(P, C)$ gives the splitting of $\pi_1(C)$ as the graph of groups $\mathbb{X}^P$ induced by the decomposition* (1).

All that being sorted out, we can now make sense of the second and third words of the title of the paper.

Definition 1.8. For a pattern $P \subset C$ the tree $T(P,C)$ obtained in Proposition 1.6 is called the *Bass-Serre tree dual to* P or simply the *dual Bass-Serre tree.*

Definition 1.9. A splitting of $\pi_1(C)$ is *geometric* if it is represented by a pattern P, i.e., the Bass-Serre tree of the splitting is given by the action of $\pi_1(C)$ on $T(P,C)$. The pattern P is said to be *essential* if $T(P,C)$ is infinite.

The following fact is important since it implies that the class of geometric splittings is significant.

Theorem 1.10 (Restatement of [DS99, Lemma 2.2])**.** *Let $\pi_1(C)$ act minimally on a tree T. Then there exists a pattern $P \subset C$ such that there is a $\pi_1(\mathscr{C})$-equivariant surjective* simplicial *map called a* resolution

$$\rho : T(P,C) \to T.$$

In particular, the edge stabilizers of $T(P,C)$ are conjugate to subgroups of the edge stabilizers of T. Moreover if the action of $\pi_1(C)$ on T is nontrivial, then some track in P will be essential.

Immediately we get:

Corollary 1.11. *If $\pi_1(C)$ is freely decomposable, then some essential free decomposition is geometric.*

Corollary 1.12. *If C is a finite complex, then any Guirardel-Levitt JSJ deformation space* [GL09, Definition 4] *of $\pi_1(C)$ contains a geometric splitting.*

Proof. We refer the reader to the introduction of [GL09] for the terminology in this proof. Let T be some JSJ tree for $\pi_1(C)$ over some class of groups $\mathcal{A}$, i.e., a domination-maximal universally elliptic $\mathcal{A}$-tree. Theorem 1.10 implies the existence of a geometric tree $T(P,C) \to T$ that dominates T. Since the edge groups of $T(P,C)$ are contained in edge groups of T and $\mathcal{A}$ is assumed to be closed under taking subgroups, $T(P,C)$ is also an $\mathcal{A}$-tree. Since the edge groups of T are $\mathcal{A}$-universally elliptic, so must the edge groups of $T(P,C)$. It follows that $T(P,C)$ is also a domination-maximal universally elliptic $\mathcal{A}$-tree so the result follows. □

1.3. **Relative splittings.** Suppose we are given a finite collection of finite generating

$$S = \big\{\{h_i\}_{i\in I_n} \mid n = 1,\dots,m\big\}$$

of subgroups of $\pi_1(C)$ and that we want to study the geometric splittings of $\pi_1(C)$ in which the subgroups $\langle h_i\rangle$ are elliptic. Then we can make a new 2-complex $C_S \supset C$ with $\pi_1(C) \approx \pi_1(C_S)$ as follows (see Figure 1.) For each $S_n = \{h_i\}_{i\in I_n}$,

(1) Make a bouquet of circles B_n such that for each $h \in S_n$ there is a directed edge e_h in B_n.
(2) Attach the vertex v_n of B_n to the vertex v of C by an arc α_n.
(3) Attach a 2-cell so that the loop $\alpha_n * e_h * \alpha_n^{-1}$ is now homotopic to $h \in \pi_1(C,v)$.

We call the resulting 2-complex C_S. We note that $\pi_1(C) \approx \pi_1(C_s)$ because of the obvious deformation retraction $C_S \twoheadrightarrow C$. We now employ the following trick to restrict to relative geometric splittings.

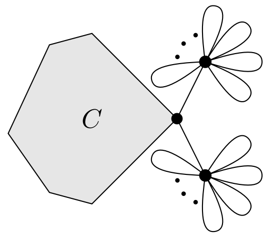

FIGURE 1. Attaching bouquets of circles to a presentation complex. We must then attach 2-cells to preserve the fundamental group.

Proposition 1.13. *Let $\mathcal{H}$ be the set of subgroups generated by the elements of S. Then every track t dual to a geometric splitting of $\pi_1(C)$ relative to $\mathcal{H}$ can be extended to a track $t' \subset C_S$ so that t' is disjoint from the edges e_h in item (1) of the construction of C_S.*

Sketch of proof. Let $\pi_1(C) \times T \to T$ be a geometric action with the subgroups $\mathcal{H}$ acting elliptically. Then we can extend the pattern $P \subset C \subset C_S$ to a pattern $P' \subset C_S$ such that P' has empty intersection with the edges e_h and such that we have a $\pi_1(C)$-equivariant isomorphism

$$T(P', C_S) \to T(P, C). \tag{3}$$

We do this by taking a resolution

$$\rho' : \widetilde{C_S} \to T(P, C)$$

which extends $\rho : \widetilde{C} \to T(P, C)$ such that the lifts of the vertices of v_i are mapped to vertices stabilized by appropriate conjugates of $\langle S_i \rangle$. We refer the reader to Section 2 of [DS99] for details on the resolution construction. It therefore follows that geometric splittings of $\pi_1(C)$ relative to $\mathcal{H}$ are given exactly by patterns in $\pi_1(C_S)$ that do not intersect the new edges e_h. □

1.4. **Equivalence under automorphisms.** Let

$$\begin{array}{rcl} \varphi : G \times T & \to & T \\ (g, x) & \mapsto & g \cdot x \end{array}$$

be an action of a group G on a tree T. Then for any $\alpha \in \mathrm{Aut}(G)$ we may *twist* φ by α to get a new action

$$\begin{array}{rcl} \varphi^\alpha : G \times T & \to & T \\ (g, x) & \mapsto & \alpha(g) \cdot x. \end{array}$$

Definition 1.14. Let $\varphi : G \times T \mapsto T$ and $\psi : G \times S \mapsto S$ be two actions of the group G on simplicial trees. Let $\alpha \in \mathrm{Aut}(G)$. We say the actions φ and ψ are $\mathrm{Aut}(G)$-*equivalent*, written $\varphi \sim_{\mathrm{Aut}(G)} \psi$, if there is a simplicial isomorphism $f : T \to S$ and an $\alpha \in \mathrm{Aut}(G)$ that makes the following diagram commutative:

$$\begin{array}{ccc} G \times T & \xrightarrow{\ \varphi^\alpha\ } & T \\ {\scriptstyle 1 \times f}\downarrow & & \downarrow{\scriptstyle f} \\ G \times S & \xrightarrow{\ \psi\ } & S \end{array}$$

Definition 1.15. Let P, P' be patterns in C. We say that *the patterns P and P' are* $\mathrm{Aut}(\pi_1(C))$*-equivalent*, written $P \sim_{\mathrm{Aut}(\pi_1(C))} P'$, if the natural actions $\pi_1(C) \times T(P, C) \to T(P, C)$ and $\pi_1(C) \times T(P', C) \to T(P', C)$ are $\mathrm{Aut}(\pi_1(C))$-equivalent.

In the case of closed surfaces patterns arise as multicurves, and automorphic equivalence of multicurves coincides with equivalence under homeomorphisms. For general 2-complexes (which may have trivial mapping class groups) these equivalences may not coincide.

1.5. **The main result.** A solution to the word problem in $\pi_1(C)$ is a procedure that decides if a loop (given as a sequence of directed edges in $C^{(1)}$) is nullhomotopic in C. In the case where C is a presentation complex such a loop uniquely defines a word in the prescribed generating set.

Theorem B. *There is an algorithm which takes as input a finite* 2*-complex* C *such that* $\pi_1(C)$ *has no* 2*-torsion, a solution to the word problem in* $\pi_1(C)$*, some positive integer* κ*, a finite collection*

$$S = \left\{ \{h_i\}_{i \in I_n} \mid n = 1, \ldots, m \right\}$$

of finite generating sets of subgroups $\mathcal{H} = \left\{ \langle h_i \rangle_{i \in I_n} \right\}_{n=1}^{m}$ *of* $\pi_1(C)$ *and outputs a finite collection of tracks* $t_1, \ldots, t_{n(C,\kappa,S)}$ *which lie in a complex* $C_S \supseteq C$ *(with equality if* $S = \emptyset$*) such that the isomorphism* $\pi_1(C_S) \approx \pi_1(C)$ *is given explicitly. These tracks give splittings of* $\pi_1(C)$ *relative to* $\mathcal{H}$ *with the following property: if* $\pi_1(C)$ *admits a geometric* κ*-acylindrical splitting relative to* $\mathcal{H}$ *represented by a track* t*, then there is some* $i \in \{1, \ldots, n(C, \kappa, S)\}$ *such that* $t \sim_{\mathrm{Aut}(\pi_1(C))} t_i$*.*

This theorem is proved in Section 7.6, where the main algorithm is given. The 2-complex C_s was defined in Section 1.3.

The result, as stated, is about splittings that are geometric, κ-acylindrical and with one edge group. This result is not the strongest possible, but it gives us all the applications we need.

It could be strengthened as follows: by Theorem 1.10 every κ-acylindrical tree T is resolved by a geometric tree $\rho : T(P, C) \twoheadrightarrow T$. The resolving tree $T(P, C)$, however, may not itself be κ-acylindrical. The analysis of the relationship between the trees $T(P, C)$ and T in [Del99] combined with the arguments of this paper actually gives a finite collection of tracks that *resolve* every κ-acylindrical tree. The geometric resolving splittings themselves may not be κ-acylindrical, but they are "locally" κ-acylindrical in a way that is good enough for our arguments. Although this would give us a full algorithmic version of Sela's super accessibility [Sel97, Theorem 4.3] for the class of κ-acylindrical one edged splittings, we have opted for a simpler formulation, thus removing a layer of notation. We hope the reader will agree that this is for the best.

For the sake of simplicity we have also restricted ourselves to one edged splitting or tracks instead of general patterns. This does not weaken the result because we allow relative splittings which enables us to produce refinements. Results such as [Wei02, Del99] then give explicit bounds on the number of components of the pattern.

It should also be noted that Theorem B does not necessarily enable us to detect whether $\pi_1(\mathscr{C})$ actually has a geometric κ-acylindrical splitting. To reach such a conclusion we must be able to further analyze the collection of tracks produced by the main algorithm. This means we must be able to solve more delicate algorithmic problems in the ambient group. The next two applications, especially the proof of Theorem C, will illustrate the necessary extra requirements.

Finally there is the issue of torsion. The current algorithm cannot handle actions on trees with arbitrarily long arcs with non-trivial pointwise stabilizers, even if these stabilizers are finite. Forbidding 2-torsion, for example, controls a problem that occurs with Möbius bands by bounding their width, which gives terminating conditions. If the algorithm were to run in the presence of 2-torsion, then it would still produce a (possibly empty) list of tracks and terminate. However, because the algorithm will have stopped prematurely this list may be missing some tracks.

There is no reason these torsion issues cannot be overcome. For example, [DG10] deals with laminations in band complexes with torsion. The author suspects that to solve this problem one would have to generalize band complexes to some version for 2-orbihedra with finite cell stabilizers.

1.6. **Computing Grushko decompositions.** This is an application of Theorem B with $\mathcal{H} = \emptyset$ and $\kappa = 0$. In this case we only need to be able to solve the word problem.

Proof of Theorem A. Let C be a presentation 2-complex form $\langle X \mid R \rangle$ and consider a maximal splitting of $\langle X \mid R \rangle$ over finite groups. If G admits an essential free decomposition, then by Corollary 1.11 there is a track $t \subset C$ that represents this splitting.

Free decompositions correspond exactly to 0-acylindrical actions on trees. We now apply the algorithm of Theorem B to get a finite collection of tracks,

$$t_1, \dots, t_n.$$

If there is a track t that represents an essential free splitting of G, then t is $\mathrm{Aut}(G)$-equivalent to some t_i in our finite collection. So G admits an essential free decomposition if and only if some t_i represents an essential free decomposition.

With our solution to the word problem we are able to check for each t_i:

- if $\mathrm{Gp}(t_i) = \{1\}$,
- if $C \setminus t_i$ is not connected, then both components must have non-trivial image in $\pi_1(C)$ via the inclusion map, or
- if $C \setminus t_i$ is connected, there is nothing to show,

and thus decide if t_i represents an essential free decomposition. □

We now give another method to decide if a finitely presented group G with decidable word problem is free (see [GW09].) First note that if we can solve the word problem, we can decide if a finitely presented group is abelian (check if the generators commute) and then, by linear algebra, we can compute its isomorphism type. In particular we can decide if G is isomorphic to $\mathbb{Z}$. We repeatedly apply the algorithm for Theorem A to attempt to compute the Grushko decomposition of G. If G is 2-torsion-free, then this will be the correct Grushko decomposition. Otherwise we still will obtain some (possibly trivial) free decomposition of G. We can then decide if each factor of this decomposition is isomorphic to $\mathbb{Z}$ or $\{1\}$. This will be the case if and only if G is free.

1.7. **Detecting splittings of relatively hyperbolic groups.** For this section we assume that the reader is familiar with relatively hyperbolic groups. The reader can consult [Far98] or [Hru10] for definitions. Let G be a finitely presented *torsion-free* group that is hyperbolic relative to the finitely generated subgroups $\mathcal{H} = \{H_1, \dots, H_m\}$. We assume that the groups $H_i \in \mathcal{H}$ are pairwise distinct and non-conjugate. Before continuing we need to give some definitions.

Definition 1.16. An element of $g \in G$ (respectively a subgroup $K \leq G$) is *parabolic* if there exists some $h \in G$ such that $h^{-1}gh \in H_i$ (respectively $h^{-1}Kh \leq H_i$) for some $i \in \mathcal{H}$.

Convention 1.17. We will assume in this section that all algorithms in a group are with respect to a presentation and that the (tuples of) elements of the input are given as (tuples of) words in the symmetrized generating set.

Definition 1.18. A splitting of G is *elementary* if all parabolic subgroups are elliptic and the edge groups are either trivial, infinite cyclic, or parabolic.

Definition 1.19. A triple $\big(\langle S \mid R\rangle, \texttt{CP}, \texttt{Gen}\big)$ where

(i) $\langle S \mid R\rangle$ is a finite group presentation,
(ii) $\texttt{CP}$ is an algorithm which solves the conjugacy problem with respect to $\langle S \mid R\rangle$, and
(iii) $\texttt{Gen}$ is an algorithm which decides whether or not a finite tuple generates $\langle S \mid R\rangle$

is called an *algorithmically tractable triple.*

Definition 1.20. A class $\mathcal{C}$ of finitely presented groups is called an *algorithmically tractable class of parabolics* if there is an algorithm which enumerates algorithmically tractable triples corresponding to the groups in $\mathcal{C}$.

It is worth pointing out that by [BCRS91] the class of *polycyclic-by-finite* groups is algorithmically tractable. We now collect some well-known facts about torsion-free relatively hyperbolic groups.

Proposition 1.21 ([Far98, Example 1, p. 819])**.** *For all $h \in G$, $\big(h^{-1}H_ih\big) \cap H_j \neq \{1\}$ if and only if $i = j$ and $h \in H_i$.*

This next result about elements of G follows from the work in [Bum04] but is stated explicitly in [Osi06]. The generalization to explicitly given subgroups of G follows applying Proposition 1.21.

Theorem 1.22 (cf. [Osi06, Theorem 5.6])**.** *Given $g \in G$ (respectively $K = \langle k_1, \dots, k_n\rangle \leq G$) if we are given a solution to the conjugacy problem for each $H_i, i = 1, \dots, n$, then we can decide whether there is some $h \in G$ such that $h^{-1}gh \in H_i$ (respectively $h^{-1}Kh \leq H_i$) for some $i \in \{1, \dots, n\}$, and find h if it exists.*

The following two facts are well known; however I couldn't find any precise references. They are stated here and proved.

Proposition 1.23. *Let $g \in G$ be a non-parabolic element. Then its centralizer $C(g)$ is infinite cyclic and malnormal.*

Proof. By Lemma 4.16 and Convention 4.19, the centralizer $C(g)$ of g is a word hyperbolic group. By [Aea91, Corollary 3.6], since G is torsion free $C(g)$ is infinite cyclic. Assume now for simplicity that $C(g) = \langle g \rangle$.

Suppose there was some $h \in G$ such that $h^{-1}\langle g \rangle h \cap \langle g \rangle \neq \{1\}$. Then by [Osi06, Corollary 4.26] there is some $l \in \mathbb{Z}_{\neq 0}$ such that $h^{-1}g^l h = g^{\pm l}$. This means that $h^2 \in C(g^l)$ and that $\langle g^l \rangle$ is normal in $\langle h, g^l \rangle$. Now as explained before $\langle k \rangle = C(\langle g^l \rangle) \geq C(g) = \langle g \rangle$, which implies that $k \in C(g)$, so $k \in \langle g \rangle$; thus $h^2 \in \langle g \rangle$. From this we get that $[\langle h, g^l \rangle : \langle g^l \rangle] \leq 2l$.

Now [Hem76, Lemma 11.4] states that if a group Q contains an infinite cyclic subgroup of finite index, then Q contains a finite subgroup K such that Q/K is either isomorphic to $\mathbb{Z}$ or $\mathbb{Z}_2 * \mathbb{Z}_2$. If $Q \leq G$, then Q is torsion-free, so K must be trivial and Q must be infinite cyclic. It follows that $\langle h, g^l \rangle$ is infinite cyclic, say $\langle h, g^l \rangle = \langle z \rangle$. Then $z \in C(g)$, so in particular $h \in C(g)$. □

Corollary 1.24. *A one edged elementary splitting of a torsion free relatively hyperbolic group G is 2-acylindrical.*

Proof. We first prove the following. *Claim: let e be an edge in a Bass-Serre T tree connecting the vertices u, v. Then at least one of the images of $G_e \leq G_u$ or $G_e \leq G_v$ is malnormal.*

Indeed, if the splitting in question is free, then the result holds. Suppose now that the edge group is $\langle g \rangle$ for some non-parabolic $g \in G$. Then $\langle g \rangle$ must be maximal cyclic in at least one of its images in the vertex groups; otherwise its centralizer is not cyclic because in the amalgam $\langle x \rangle *_{x^r=y^s} \langle y \rangle$ if $r, s \neq 1$, then $\langle xy, y^s \rangle \approx \mathbb{Z} \oplus \mathbb{Z} \leq C(y^s)$. Also the images of $\langle g \rangle$ in the associated subgroups cannot intersect since by [Osi06, Corollary 4.27] any Baumslag-Solitar group must be parabolic.

Suppose now that the edge group G_e is parabolic, but not maximal parabolic, and hence malnormal, in either G_u or G_v. Then we have parabolic proper overgroups $G_e < P_u \leq G_u$ and $G_e < P_v \leq G_v$. On one hand since $|P_u \cap P_v| = |G_e| = \infty$ they must lie in a common maximal parabolic subgroup P. On the other hand P does not act elliptically on T (it has a non-trivial induced splitting), contradicting the fact that the splitting is elementary. *This proves our claim.*

We now prove 2-acylindricity. Let u be a vertex in the Bass-Serre tree and let e, f be edges such that $e \cap f = \{u\}$. Since e, f are in the same G-orbit we have that G_e, G_f are conjugate. On the other hand if $e \neq f$, then if $G_v \geq G_e \cap G_f \neq \{1\}$, then G_e is not malnormal in G_v. Suppose towards a contradiction that there is some $g \in G_e$ such for some edge $h \subset T$ such that $e \cap h = \emptyset$ we have $g \cdot h = h$, and suppose moreover that there is some edge $f \subset T$ such that $e \cap f = \{u\}$ and $f \cap h = \{v\}$. Then we must have that $g \in G_e \cap G_f \cap G_h$. $g \in G_e \cap G_f$ implies that G_f is not malnormal in G_u, which means by our earlier claim that G_f must be malnormal in G_v so $G_f \cap G_h = \{1\}$, a contradiction. Therefore no element of $G \setminus \{1\}$ fixes an arc of T of length more than 2. □

We finally need the following.

Theorem 1.25 (Theorem 3 of [DG13])**.** *There exists an algorithm as follows. It takes an input of a finite presentation of a group G, a solution to its word problem, and a recursive class of finitely presented groups $\mathcal{C}$ (given by a Turing machine enumerating presentations of these groups).*

It terminates if and only if G is properly hyperbolic relative to subgroups that are in the class $\mathcal{C}$.

In this case, the algorithm outputs an (relative linear) isoperimetry constant K (in the sense of [Osi06, Definition 2.30]*), a generating set, and a finite presentation for each of the parabolic subgroups.*

Now we have our second application.

Theorem C. *Suppose we are given a finite presentation $\langle X \mid R\rangle$ of a torsion free group G that is relatively hyperbolic with finitely many parabolics that lie in an algorithmically tractable class of parabolics $\mathcal{C}$. Suppose also that we are given a solution to the word problem with respect to $\langle X \mid R\rangle$ and a finite collection S of finite generating sets for a set of subgroups $\mathcal{H}'$. Then we can decide if G admits an elementary splitting relative to $\mathcal{H}'$.*

Proof. We first note that by Theorem 1.10 if C is the presentation 2-complex associated to $\langle X \mid R\rangle$, then G admits an essential elementary splitting if and only if $\pi_1(C)$ admits an essential elementary *geometric* splitting. Any elementary splitting is, by Corollary 1.24, 2-acylindrical.

We first run the algorithm of Theorem 1.25 to find the finite collection $\{H_1, \dots, H_n\}$ (given by generating sets in $\langle X \mid R\rangle$) of parabolic subgroups. We then apply the algorithm of Theorem B with $\kappa = 2$ and

$$\mathcal{H} = \mathcal{H}' \cup \{H_1, \dots, H_n\}$$

with the collection of generating sets S. This gives us a finite collection of tracks $t_1, \dots, t_{n(C,\kappa,S)}$ that lie in C_S. It is now enough to check for each of these tracks if they represent an essential elementary splitting. Let t be one of these tracks. By Theorem 1.22 we can decide if $\mathrm{Gp}(t)$ is parabolic.

If $\mathrm{Gp}(t)$ is trivial, then as in the proof of Theorem A we can decide if it gives an essential splitting.

Suppose now that $\mathrm{Gp}(t)$ isn't parabolic. Since we can solve the word problem we can check whether $\mathrm{Gp}(t)$ is abelian. If it isn't, then it certainly cannot represent an elementary splitting. Otherwise $\mathrm{Gp}(t)$ is abelian. By Proposition 1.23, $\mathrm{Gp}(t)$ is contained in the centralizer of some non-parabolic element and is therefore contained in a non-parabolic maximal cyclic group. If t is a non-separating track, then it gives an essential elementary splitting. Otherwise t separates C_S, and G splits as a free product with amalgamation over $\mathrm{Gp}(t)$. To check if the splitting is essential it suffices to check, using the word problem, whether the generators of the vertex groups commute with $\mathrm{Gp}(t)$. Indeed, since we are assuming that $\mathrm{Gp}(t)$ is non-parabolic, we can assume that the vertex groups are non-parabolic; so by Proposition 1.23 if one of the vertex groups commutes with $\mathrm{Gp}(t)$, then it is at most a finite index cyclic overgroup of $\mathrm{Gp}(t)$. Deciding if the vertex group coincides with $\mathrm{Gp}(t)$ can now be solved using item (3) of Theorem 1.16 of [Osi06].

Suppose finally that $\mathrm{Gp}(t)$ is parabolic. Again, if t is non-separating the splitting is essential. Otherwise the splitting is essential if and only if $\mathrm{Gp}(t)$ doesn't equal one of the vertex groups. If neither of the vertex groups is parabolic, then the splitting is essential. Otherwise at most one of the vertex groups is parabolic, and we can decide if it is generated by $\mathrm{Gp}(t)$ using Theorem 1.22 and our solution to the generation problem given by the algorithmic tractability assumption (Definition 1.19(iii)). □

2. Band complexes

The algorithm of Theorem B is a procedure that will produce a rooted directed tree (i.e., a branching sequence) of band complexes. Band complexes first appeared in [BF95] to classify stable actions of finitely presented groups on $\mathbb{R}$-trees. Our version of band complexes differs in that they are combinatorial objects: they do not come with laminations; instead we will allow a band complex to *carry* multiple laminations or, in our case, tracks.

As combinatorial objects, our band complexes will contain the same amount of information as Makanin's generalized equations (cf. [KM98]).

2.1. Definitions and terminology.

Definition 2.1. A *band* $\mathbb{B}$ is a Cartesian product $J_{\mathbb{B}} \times [-1,1]$ where $J_{\mathbb{B}}$ is homeomorphic to a closed interval. The subsets $J_{\mathbb{B}} \times \{\pm 1\}$ are called *bases*. If $\mu = J_{\mathbb{B}} \times \{\pm 1\}$ is the base of a band, then we call the base $\overline{\mu} = J_{\mathbb{B}} \times \{\mp 1\}$ the *dual* of μ.

A band is therefore a rectangle with well-defined bases and a vertical direction.

Convention 2.2. The letters λ, μ, η, ν shall be used to denote bases, and $\overline{\lambda}$ will always denote the dual of λ. We shall denote by $\mathbb{B}(\lambda)$ the band that contains λ.

Definition 2.3. A *band complex* $\mathscr{C}$ is a 2-complex that is constructed in the following way:

1. Start with a simplicial graph Γ.
2. Attach the bases of the bands $\mathbb{B}_1, \dots \mathbb{B}_m$ to the interiors of edges of Γ via embeddings
$$g_{\mu_i} : \mu_i \hookrightarrow \Gamma \setminus \Gamma^{(0)},$$
where the $\{\mu_i\}$ is the set of bases of the bands.
3. Let
$$U = \big(\Gamma \cup \mathbb{B}_1 \cup \ldots \cup \mathbb{B}_m\big)/\sim$$
be the resulting identification space. We finally obtain $\mathscr{C}$ by attaching discs $D_1, \dots, D_l$ via immersions $f_i : \partial D_i \looparrowright U$ with the following requirement:
   (a) For all i, j, $f_i(\partial D_i) \cap \mathbb{B}_j$ can be expressed as a finite union of embedded arcs α_i that travel from one base of $\mathbb{B}_j$ to the other. Such arcs are called *connections*.
   (b) Connections are pairwise disjoint.
   (c) If a connection has non-trivial intersection with a side of a band $\mathbb{B}_j$, then it coincides with that side.

In the case of measured band complexes [BF95] the complicated requirement (3) above on the 2-cell attaching map is ensured if the 2-cell attaching maps intersect measured bands in *vertical* subsets.

Connections (as described in items (3a)-(3c) of the definition above) will occur exactly where *boundary connections* occur when working with generalized equations (see for example [KM98].) Controlling their cardinality is a key step in the repetition argument which deals with the thinning and superquadratic cases of the elimination process (Section 4.8).

Convention 2.4. Although a band complex is a 2-complex, whenever we mention a 2-cell we *really mean a 2-cell D_i that gets attached in step* (3) *of Definition* 2.3.

Thus, it is possible that a 2-cell D_i in a band complex never intersects any bands, in which case the image of the attaching map $f_i(\partial D_i)$ lies entirely in the underlying graph Γ in item (1) of Definition 2.3.

Convention 2.5. Formally speaking, a base μ isn't a subset of the band complex $\mathscr{C}$. That being said we will still write $x \in \mu$ for some point $x \in \mathscr{C}$ such that $x \in g_\mu(\mu)$. We will also write $\lambda \subset \mu$ if $g_\lambda(\lambda) \subset g_\mu(\mu)$. In the case where $g_\lambda(\lambda) = g_\mu(\mu)$ we will use the evocative symbol $\lambda \doteqdot \mu$ to avoid confusion. We will also treat the bands $\mathbb{B}_i$ as subsets of $\mathscr{C}$ when it is convenient.

Definition 2.6. We say that bases $\mu, \overline{\mu}$ are *matched bases* if $\mu \doteqdot \overline{\mu}$ and $\mathbb{B}(\mu)$ forms an annulus in $\mathscr{C}$. Otherwise a base is called *unmatched.*

Definition 2.7. A union of unmatched bases $U = \bigcup_{\mu \in S} \mu$ is called *strongly connected* if the union of the interior of the bases $U' = \bigcup_{\mu \in S} \textbf{interior}\,(\mu)$ is also connected (and therefore an interval). A maximal (with respect to inclusion) strongly connected union of unmatched bases is called a *maximal section.*

Maximal sections are almost the blocks in [BF95] and the closed sections in [KM98].

Definition 2.8 (Carrying a track)**.** Let $\mathscr{C}$ be a band complex and let $t \subset \mathscr{C}$ be a track. $\mathscr{C}$ *carries* t if t is contained in the union of the bands in $\mathscr{C}$ and furthermore:

(i) For each band $\mathbb{B}_j$, $t \cap \mathbb{B}_j$ consists of a union of pairwise disjoint embedded arcs travelling from one base of $\mathbb{B}_j$ to the other, and
(ii) $\mathbb{B}_j \cap t$ is disjoint from the connections in $\mathbb{B}_j$ as well as from its sides.

Definition 2.9 (Efficiently carrying)**.** For a base μ, let $S_\mu \subset \mu$ be the finite set containing the points of the form $x = \mu \cap c$, where c is a connection, and the points x that are the endpoints of bases. $\mathscr{C}$ *carries* t *efficiently* if for every base μ and every distinct $x, y \in S_\mu$ there is some point in $t \cap \mu$ that separates them.

Thus, if $\mathscr{C}$ carries a track t, then t is confined to the interior of the bands. The notion of carrying naturally generalizes to arbitrarily measured laminations, but since we will only be focusing on one leaf laminations we only need to deal with the hitting measure.

Definition 2.10 (Measure from a track)**.** Let $\mathscr{C}$ be a band complex, let t be a track carried by $\mathscr{C}$, and let $S \subset \mathscr{C}$ be a union of bases. We define the *hitting measure on* S *with respect to* t, denoted $|S|_t$, to be the cardinality of the intersection

$$|S|_t = |S \cap t|.$$

If μ is a base of $\mathscr{C}$, then we will sometimes call $|\mu|_t$ the *length of* μ *with respect to* t.

2.2. **Constructing (measured) band complexes from tracks.** Let C be a standard CW 2-complex and let $t \subset C$ be a track. We obtain a band complex $\mathscr{C}$ from C as follows.

For each 2-cell $D \subset C$, $t \cap D$ is a disjoint union of arcs travelling from one edge of ∂D to another edge of ∂D. Metrize the 1-skeleton $C^{(1)}$, giving each edge e_i a length of $|e_i \cap t| + 2$. Subdivide each 2-cell D into a union of bands and 2-cells such that a base μ has length $|\mu \cap t| = |\mu|_t$. Explicitly parameterize each band

$$\mathbb{B}(\mu) = [0, |\mu|_t] \times [-1, 1]$$

so that $t \cap \mathbb{B}(\mu) \cap t$ is a union of vertical sets as follows:

$$t \cap \mathbb{B}(\mu) = \bigcup_{i=1}^{|\mu|_t} \{i - \frac{1}{2}\} \times [-1, 1];$$

this is illustrated in Figure 2.

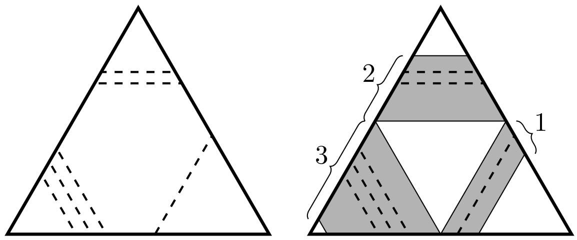

FIGURE 2. On the left, the intersection of a track t with a 2-cell U in a 2-complex. On the right how to construct the corresponding measured band complex efficiently carrying t by dividing U into three bands and four 2-cells.

This construction gives a band complex as in Definition 2.3. The underlying graph is the 1-skeleton $C^{(1)}$; the bases of the bands are embedded in this graph and avoid the vertices $C^{(0)}$. We further see that the remaining 2-cells have embedded (thus, immersed) boundaries that only intersect bands in their vertical sides. Furthermore, since no 2-cell attaching maps go through the interior of any of the bands, t is efficiently carried by the band complex $\mathscr{C}$.

Definition 2.11. A band complex $\mathscr{C}$ is *measured* if every base μ is metrized as a real closed interval $[0, |\mu|_t]$ (recall Definition 2.10). A subset v of a band $\mathbb{B}(\mu) = [0, |\mu|_t] \times [-1, 1]$ is called *vertical* if it is of the form

$$v = \{x_v\} \times [-1, 1].$$

We have proved:

Proposition 2.12. *For any track t contained in a 2-complex C, we can subdivide C into a measured band complex $\mathscr{C}$ in which t consists of a union of vertical sets. Moreover $\mathscr{C}$ carries t efficiently.*

Definition 2.13 (Combinatorial equivalence of band complexes)**.** Two band complexes $\mathscr{C}$ and $\mathscr{C}'$ are said to be *equivalent* if there is a homeomorphism $\mathscr{C} \to \mathscr{C}'$ that sends the underlying simplicial graph (Definition 2.3(1)) to the underlying simplicial graph, sends bands to bands, sends 2-cells to 2-cells, and for each of these objects restricts to a homeomorphism.

If we forget the measures on band complexes, we are left with only finitely many possibilities. Thus,

Proposition 2.14. *Let C be a finite 2-complex and let S be a finite collection of finite subsets of $\pi_1(C)$. Then there are only finitely many possible band complexes* (*up to the combinatorial equivalence of Definition* 2.13) *that arise from the possibly infinite collection of tracks $t \subset C_S$* (*Proposition* 1.13)*. Furthermore this list can be effectively constructed.*

3. Moves on band complexes carrying tracks

We will present moves that transform a band complex carrying a track into a new band complex carrying a new track. These moves are essentially the moves given in [BF95, §6.1]. Our treatment is slightly different since we want to explicitly realize each move as a continuous map $m : \mathscr{C} \to \mathscr{C}'$ sending the track t to some track $t' \subset \mathscr{C}'$, which we denote $(\mathscr{C}, t) \to (\mathscr{C}', t')$. This is accomplished with the zipping moves. They will be of use in a later section.

We will employ the convention of [KM98] and reuse the names of bases, as is customary in computer science.

3.1. **The basic moves on band complexes that carry a track.** Let $\mathscr{C}$ be a band complex efficiently carrying a track t. Suppose furthermore that $\mathscr{C}$ is measured (Definition 2.11) so that a base μ is metrized with length $|\mu|_t$ and $t \subset \mathscr{C}$ is a union of vertical sets.

We first define elementary moves $(\mathscr{C}, t) \to (\mathscr{C}', t')$ which transform the underlying band complex and track while preserving the fundamental group and dual Bass-Serre tree $T(t, \mathscr{C})$. These moves are actually π_1-isomorphic continuous maps $\mathscr{C} \to \mathscr{C}'$ that map t to t'.

Definition 3.1 (Type I zip)**.** Suppose we have a containment of bases $\lambda \subset \mu$ with $\overline{\lambda} \cap \overline{\mu} = \emptyset$ and $\overline{\lambda} \cap \mu = \emptyset$. The union of U vertical sets of $\mathbb{B}(\mu)$ that intersect λ is a rectangle homeomorphic to $\mathbb{B}(\lambda)$. A *type I zip of* $\mathbb{B}(\lambda)$ *into* $\mathbb{B}(\mu)$ consists of the operation of identifying $\mathbb{B}(\lambda)$ to U so that vertical sets are sent homeomorphically to vertical sets and $\overline{\lambda}$ is identified to the corresponding subset of $\overline{\mu}$.

Definition 3.2 (Type II zip, or squish)**.** Suppose we have the containments of bases $\lambda \subset \mu$ and $\overline{\lambda} \subset \overline{\mu}$. Suppose furthermore that there are vertical paths $\alpha \subset \mathbb{B}(\lambda)$ and $\beta \subset \mathbb{B}(\mu)$ such that the concatenation $\alpha * \beta$ is a nullhomotopic loop. A *type II zip of* $\mathbb{B}(\lambda)$ *into* $\mathbb{B}(\mu)$ is the operation of continuously identifying α to β and continuously extending this to an identification of $\mathbb{B}(\lambda)$ to a union of vertical subsets of $\mathbb{B}(\mu)$. This identification map must be injective when restricted to $\mathbb{B}(\mu)$ and $\mathbb{B}(\lambda)$ and must send vertical sets to vertical sets.

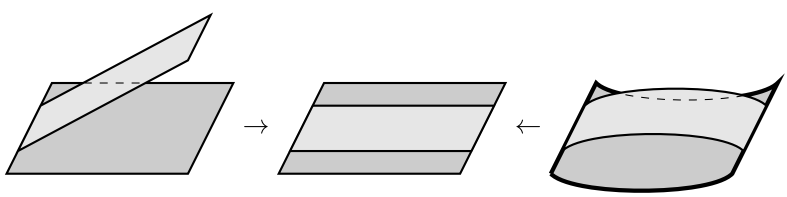

Figure 3. Type I and type II zips. The identifications must send tracks to tracks.

These zipping moves are not used in [BF95], but the type I zip is a step in the transfer (see Definition 3.10 later) or *M4 Slide* of [BF95]. The zipping moves will be necessary later when we will be "wrapping one band around another" (see Figure 19). Also the fact that they are given by explicit continuous maps is convenient.

Definition 3.3 (Collapse a band)**.** Let $\mathbb{B}(\mu)$ be a band such that $\mu \cap \overline{\mu} = \emptyset$. Then the *collapse of* $\mathbb{B}(\mu)$ *onto* μ is the operation of identifying each vertical subset v that intersects **interior** (μ) to the point $x_v = v \cap \mu$.

Definition 3.4 (Annulus)**.** An *annulus* $\mathcal{A}$ is a band $\mathbb{B}(\mu)$ such that $\mu \doteqdot \overline{\mu}$ whose image in the band complex $\mathscr{C}$ is homeomorphic to an annulus.

This next move is similar to the type II zip.

Definition 3.5 (Crush an annulus)**.** Let $\mathcal{A} \subset \mathscr{C}$ be an annulus such that $\mathrm{Gp}(\mathcal{A}) = \{1\}$. Then *crushing* $\mathcal{A}$ is the operation of identifying each vertical subset of $\mathcal{A}$ to its intersection with μ.

Definition 3.6 (Vertically subdivide a band)**.** Let $\mathbb{B} = [a,b] \times [-1,1]$ be a band in $\mathscr{C}$ and let $\{p\} \times [-1,1]$ be a vertical subset. *The operation of subdividing* $\mathbb{B}$ *along* $\{p\} \times [-1,1]$ consists of the following operations:

(1) Cut $\mathbb{B}$ along $\{p\} \times [-1,1]$ so that we get two bands $\mathbb{B}_1 = [a,p^-] \times [-1,1]$ and $\mathbb{B}_2 = [p^+,b] \times [-1,1]$.
(2) Attach a 2-cell along the loop $(p^- \times [-1,1]) * (p^+ \times [-1,1])$, where $*$ denotes concatenation. This 2-cell is called a *subdivision digon*.

Convention 3.7. We only allow band subdivision of $(t,\mathscr{C})$ if the resulting $t' \subset \mathscr{C}'$ is efficiently carried.

These basic operations may leave some messiness behind:

(i) After performing a zipping move a 2-cell may no longer have an immersed boundary.
(ii) After a collapse a 2-cell may have a free face, and perhaps the resulting band complex can be given as $\mathscr{C} = \mathscr{C}' *_p \alpha$, i.e., the connected sum at a point p of a band complex $\mathscr{C}'$ and a closed arc α.
(iii) After crushing an annulus or performing a type II zip, the boundary of a 2-cell may map onto a point or an interval resulting in a sphere.

We therefore introduce, as basic moves, the following cleaning operations.

Definition 3.8 (Delete superfluous cells)**.** If a 2-cell in $\mathscr{C}$ is a sphere, as may occur in (iii) above, or if it has a free face, remove it. Do the same for hanging arcs that occur in (ii) above.

Definition 3.9 (Tighten 2-cells)**.** If a 2-cell D no longer has an immersed boundary, then the attaching map $f_D : \partial D \to \mathscr{C}$ factors as

$$\partial D \twoheadrightarrow \left(S^1 *_{p_i} \tau_i\right) \overset{f'}{\looparrowright} \mathscr{C}, \tag{4}$$

where the middle term is a circle with some hanging trees τ_i, which arise from the "pinching" of the attaching map. This middle term is immersed into $\mathscr{C}$. We replace D by a 2-cell D' whose boundary is identified with S^1 in (4) and mapped to $\mathscr{C} \setminus D$ (abusing notation) via the immersion f' in (4). D' is called the *tightening* of D.

We leave it to the reader to verify that the result of a 2-cell removal and the tightening move $\mathscr{C} \to \mathscr{C}'$ can be realized by a continuous map. The transfer, given below and illustrated in Figure 4, is defined in terms of band subdivisions and zipping, but we will also treat it as an elementary move.

Definition 3.10 (The transfer)**.** Let $\mu \subset \lambda$ be bases such that $\mu \neq \overline{\lambda}$. The operation of *transferring* μ *from* λ *to* $\overline{\lambda}$ *across* $\mathbb{B}(\lambda)$ is the following sequence of operations.

(1) Subdivide the band $\mathbb{B}(\mu)$ horizontally:

$$\mathbb{B}(\mu) = \mathbb{B}(\mu_-) \cup \mathbb{B}(\mu_+)$$

with $\mu \doteqdot \mu_-$, $\overline{\mu_-} \doteqdot \mu_+$, and $\overline{\mu_+} \doteqdot \overline{\mu}$.

(2) Zip the band $\mathbb{B}(\mu_-)$ into $\mathbb{B}(\lambda)$. (By hypothesis, this is a type I zip).

(3) We rename the base μ_+ as μ.

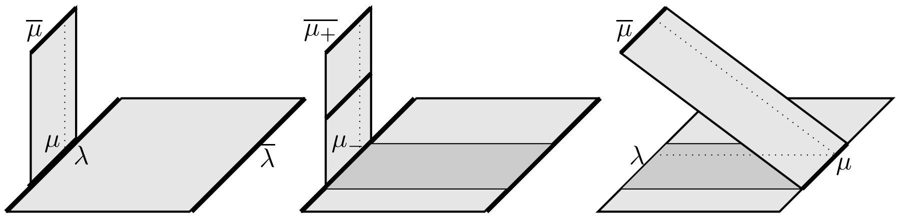

FIGURE 4. Transferring μ from λ to $\overline{\lambda}$ across $\mathbb{B}(\lambda)$. The dotted line shows a connection (Definition 2.3(3a)). After a transfer this connection may give rise to two connections.

3.2. **The preservation property.** We will now give a preservation result for our moves. This result is stated as a fact at the beginning of [BF95, §6]. In order to lay out the terminology that is necessary for our purposes, we will carefully state and prove the preservation property.

Proposition 3.11 (The preservation property)**.** *Let $m : (\mathscr{C}, t) \to (\mathscr{C}', t')$ be one of the basic moves given in Section* 3.1*. Then we have an isomorphism of fundamental groups*

$$m_\sharp : \pi_1(\mathscr{C}) \xrightarrow{\sim} \pi_1(\mathscr{C}')$$

and a simplicial isomorphism of dual Bass-Serre trees $m_T : T(t, \mathscr{C}) \to T(t', \mathscr{C}')$ induced by m. Furthermore this map is $m_\sharp$-equivariant in the following sense: letting $\pi_1(\mathscr{C})$ act naturally on $T(t, \mathscr{C})$ by deck transformations via the quotient map π (Proposition 1.6 (2)*), we have*

$$m_\sharp(g) \cdot m_T(x) = m_T(g \cdot x),$$

for all $g \in \pi_1(\mathscr{C})$ and all $x \in T(t, \mathscr{C})$.

Proof. We first prove the proposition for zipping, collapsing, and crushing moves.

We first show that the fundamental groups are isomorphic. Consider first either a type I zip (Definition 3.1) of $\mathbb{B}(\lambda)$ into $\mathbb{B}(\mu)$ (i.e., with $\lambda \subset \mu$) or the collapse of $\mathbb{B}(\mu)$ onto μ (Definition 3.3). These moves lift to $\pi_1(\mathscr{C})$-equivariant moves on $\widetilde{\mathscr{C}}$. Pick a basepoint $x \in \mu \subset \mathscr{C}$; by the disjointness criteria we see that no distinct lifts of x in $\widetilde{\mathscr{C}}$ are identified and that the resulting complex $\widetilde{\mathscr{C}'}$ remains simply connected. Since the lifts of x in $\widetilde{\mathscr{C}}$ are in bijective correspondence with $\pi_1(\mathscr{C})$; the isomorphism $\pi_1(\mathscr{C}) \approx \pi_1(\mathscr{C}')$ follows.

In the case of a type II zip (Definition 3.2) or an annulus crush (Definition 3.5) the π_1-triviality criteria ensure that we can find lifts $\widetilde{\mathbb{B}(\mu)}, \widetilde{\mathbb{B}(\lambda)}$ of $\mathbb{B}(\lambda), \mathbb{B}(\mu)$ (respectively) such that $\widetilde{\mathbb{B}(\mu)} \cup \widetilde{\mathbb{B}(\lambda)}$ is as in the right side of Figure 3 or that the annulus $\mathcal{A}(\mu)$ lifts to $\widetilde{\mathscr{C}}$. Arguing as before (taking a basepoint in μ) we obtain the isomorphism $\pi_1(\mathscr{C}) \approx \pi_1(\mathscr{C}')$.

In all cases the moves map bands to (interiors) of bands and vertical sets to vertical sets. Also points in the complement of the union of bands of $\mathscr{C}$ are sent to the complement of the union of band of $\mathscr{C}'$, and the restriction to the complement is injective. Lifting to $\widetilde{\mathscr{C}}$ we therefore see a bijective correspondence between the connected components of $\widetilde{t}$ and $\widetilde{t'}$. Furthermore if connected components $t_1, t_2 \subset \widetilde{t}$ are dual to edges that share a vertex in $T(t, \mathscr{C})$, then their images t_1', t_2' will be dual to edges in $T(t', \mathscr{C}')$ that share a vertex. The isomorphism of Bass-Serre trees follows, and $m_\sharp$-equivariance of the isomorphism follows from construction.

The proof for band subdivisions, superfluous 2-cell deletions, and tightenings is obvious. □

3.3. **Derived moves.** Having defined basic moves we shall now define the composite, or derived, moves that constitute the Rips machine. We first introduce the τ-complexity (originally Makanin's ξ-complexity [Mak82]), which is one of the main tools of our analysis. As we define the derived moves we will show why they do not increase this τ-complexity.

Definition 3.12. Let $\sigma \subset \mathscr{C}$ be a maximal section (Definition 2.7), and let $b(\sigma)$ be the number of unmatched bases contained in σ. We define the τ-complexity of a section to be

$$\tau(\sigma) = \max\big(b(\sigma) - 2, 0\big).$$

Definition 3.13 ([BF95, Definition 4.3], [KM98, §5])**.** Let $J \subset \mathscr{C}$ be a union of maximal sections. Then we define the *J-relative τ-complexity* to be

$$\tau(\mathscr{C}, J) = \sum_{\sigma \not\subset J} \tau(\sigma). \tag{5}$$

If $J = \emptyset$ write $\tau(\mathscr{C})$ instead of $\tau(\mathscr{C}, \emptyset)$.

Definition 3.14. For a point $x \in \mu$ we denote by $\gamma(x)$ the number of unmatched bases λ such that $\lambda \ni x$.

Definition 3.15. The *vertical length* of the attaching map $\partial D \looparrowright \mathscr{C}$ of a 2-cell D is the number of connected components of the preimages of the connections ((3a) of Definition 2.1). Equivalently, this is the number of times the attaching map travels through a band.

3.3.1. *The Möbius move.*

Definition 3.16. A dual pair $(\mu, \overline{\mu})$ such that $\mu \doteqdot \overline{\mu}$ and $\mathbb{B}(\mu)$ forms a Möbius band is called a *Möbius pair.*

This next move is described in [BF95, Lemma 6.4]. Note that since we require tracks to be two-sided, a track $t \subset \mathscr{C}$ can never intersect the core of a Möbius band. It follows that we can always subdivide a Möbius band along its core and the resulting band complex will still efficiently carry t.

Definition 3.17 (The Möbius move)**.** Given a Möbius pair $(\mu, \overline{\mu})$, we subdivide the band $\mathbb{B}(\mu) = \mu \times [-1, 1]$ along $\{m\} \times [-1, 1]$, the core of the Möbius band. Call the resulting bands $\mathbb{B}(\mu_0)$ and $\mathbb{B}(\mu_1)$. We then transfer the base μ_0 across the band $\mathbb{B}(\mu_1)$. The dual pair $(\mu_0, \overline{\mu_0})$ now forms an annulus $\mathcal{A}(\mu_0)$, and the pair $\mu_1, \overline{\mu_1}$ intersects at a point. If the annulus $\mathcal{A}(\mu_0)$ is π_1-trivial, we crush it. We rename $\mu_1, \overline{\mu_1}$ as $\mu, \overline{\mu}$ respectively.

Straightforward verification gives the following result:

Lemma 3.18. *Let $(\mathscr{C}, t)$ be a band complex efficiently carrying a track. Let $(\mu, \overline{\mu})$ form a Möbius pair. After the corresponding Möbius move $m : (\mathscr{C}, t) \to (\mathscr{C}, t')$ we have the following:*

(i) *The hitting measure of μ drops by a half, i.e.,*

$$|\mu|_{t'} = \frac{1}{2}|\mu|_t.$$

(ii) *The hitting measure of $|\mu|_{t'}$ is the same as the measures of the matched base $|\mu_0|_{t'}$.*

(iii) *The τ-complexity did not increase.*

3.3.2. *The thinning move.* The thinning move is applied whenever there is some point x in a base μ such that $\gamma(x) = 1$ (Definition 3.14). In [BF95] this is the move for Process I, but modified to keep the number of unmatched bases non-increasing (we prevent the "long bands" of [BF95] from occurring). In [KM98] this is the move applied in cases 7-10.

Definition 3.19. Let $x \in \mu$ be a point such that $\gamma(x) = 1$. Let $x \subset \mu_0 \subset \mu$ be the maximal segment that contains only points x' with $\gamma(x') = 1$. We call μ_0 a *maximal naked segment*.

Suppose that a base μ contains a naked segment. We now describe the *thinning move starting at μ*:

(1) *Subdivide μ.* Subdivide μ into segments

$$\mu = \mu_{-1} \cup \mu_0 \cup \mu_1$$

with μ_{-1} or μ_1 possibly empty and μ_0 a maximal naked segment of μ. This results in a subdivision of $\mathbb{B}(\mu)$. Denote by D_{-1} and D_1 the added subdivision digons.

(2) *Subdivide annuli and clean μ_0.* For every matched base pair $(\lambda, \overline{\lambda})$ such that λ intersects **interior** (μ_0) vertically subdivide

$$\mathbb{B}(\lambda) = \mathbb{B}(\lambda_{-1}) \cup \mathbb{B}(\lambda_0) \cup \mathbb{B}(\lambda_1)$$

so that $\lambda_0 \subset \mu_0$ and $\lambda_{\pm 1}$ doesn't intersect **interior** (μ_0).

Next, take all the resulting base pairs $(\lambda_0, \overline{\lambda_0})$ with $\lambda_0 \subset \mu_0$ and transfer λ_0 and $\overline{\lambda_0}$ from μ_0 to $\overline{\mu_0}$ through $\mathbb{B}(\mu_0)$. Now μ_0 doesn't intersect any other bases.

(3) *Collapse the naked segment and the added subdivision digons.* Collapse the band $\mathbb{B}(\mu_0)$ onto $\overline{\mu_0}$. Delete the subdivision digons D_{-1}, D_1 that were added in step 1 since they now have free faces.

(4) *Remove long bands.* A *long band* is a union of two bands

$$\mathbb{B}(\mu) \cup \mathbb{B}(\lambda)$$

with $\overline{\mu} \doteqdot \lambda$. If a long band is created, first transfer $\overline{\mu}$, and all other bases contained in λ, from λ to $\overline{\lambda}$ through $\mathbb{B}(\lambda)$ and then collapse the band $\mathbb{B}(\lambda)$ onto $\overline{\lambda}$.

(5) *Clean up.* Crush any remaining π_1-trivial annuli.

A proof of this next fact for generalized equations can be found in [KM98]; instead of adapting it we simply give another proof.

Lemma 3.20. *After applying a thinning move starting at a base μ, the τ-complexity did not increase. If no annuli were subdivided, the number of 2-cells did not increase, nor did the vertical lengths of 2-cell attaching maps.*

Proof. From (3), the number of 2-cells did not increase if no annuli were subdivided. Furthermore, since the only transfer move is immediately followed by a collapse of the transfer band, the vertical lengths of 2-cell attaching maps could not increase.

It remains to show that the τ-complexity doesn't increase. Suppose that we performed (2), so that μ_0 doesn't intersect any other bases (we could also do (1), (2) and then recombine the bands created in (1)). Then the number of maximal sections may have increased, but the number of unmatched bases did not; thus the τ-complexity did not increase.

Let σ be the maximal section containing μ and let σ' be the maximal segment containing $\overline{\mu}$, after the subdivision of matched bases.

Suppose first that μ is completely naked (i.e., $\mu_0 = \mu$). Then (1) doesn't occur and after the collapse of $\mathbb{B}(\mu)$ in (3), the total number of unmatched bases goes down by 2. It therefore clearly follows from Definition 3.13 (5) that the τ-complexity decreased.

Suppose now that μ gets subdivided into two segments μ_0, μ_1. Take σ to be co-initial with μ_0. After (1) and (2) σ gets subdivided into σ_0, σ_1 with $b(\sigma_0) = 1$ and $b(\sigma_1) = b(\sigma)$. In σ' the base $\overline{\mu}$ gets replaced by $\overline{\mu_0} \cup \overline{\mu_1}$. This increases $b(\sigma')$ by 1 (Definition 3.13 (5)), but in (3) we collapse $\mathbb{B}(\mu_0)$ onto $\overline{\mu_0}$. This deletes $\overline{\mu_0}$ from σ' so $b(\sigma')$ goes back down.

Suppose finally that μ gets subdivided into μ_{-1}, μ_0, μ_1. After (1) the maximal section σ gets split into $\sigma_{-1}, \sigma_0, \sigma_1$, with $b(\sigma_0) = 1$ and $b(\sigma_{\pm 1}) \geq 2$. If we look at the contribution of what is left of σ we have a decrease in the contribution of τ-complexity of at least

$$\tau(\sigma) - \big(\tau(\sigma_{-1}) + \tau(\sigma_0) + \tau(\sigma_1)\big) \geq 1.$$

On the other hand, if we look at σ' we see that $\overline{\mu}$ gets subdivided into three bases and $\overline{\mu_0}$ gets deleted after the collapse in (3). We therefore have an increase in the contribution to the τ complexity of resulting sections that constitute σ', which may have been subdivided, of at most $\tau(\sigma') + 1$; thus the total τ-complexity did not increase.

Note further that in all the cases above, if we were working with a J-relative complexity, with $\overline{\mu} \subset J$, then the J-relative τ complexity also did not increase. □

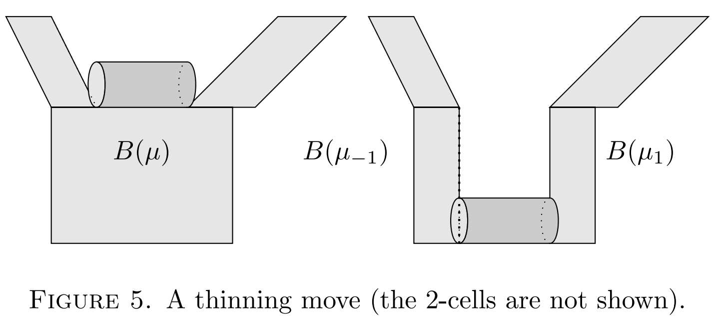

FIGURE 5. A thinning move (the 2-cells are not shown).

3.3.3. *The entire transformation.* This is the entire transformation given before [KM98, Case 12]; it also constitutes the move used in Process II in [BF95]. One of the disadvantages of the topological setting is that dealing with closures of maximal sections is awkward due to the fact that endpoints of bases can lie in the closures of distinct maximal sections.

Definition 3.21. An *ordering* $<$ *on a band complex* $\mathscr{C}$ is an ordering $<$ on the union U of maximal sections (Definition 2.7) that is compatible with some embedding $U \hookrightarrow \mathbb{R}$. An unmatched base μ whose interior is $<$-coinitial is called a *leading base.* An endpoint of μ is called initial (terminal) if it is the limit of a $<$-decreasing ($<$-increasing) sequence of points in **interior** (μ).

Convention 3.22. Band complexes will always be assumed to be equipped with an ordering.

We now describe the *entire transformation with carrier* μ. Let μ be a maximal leading base.

(1) *Subdivide matched bases.* If the base μ intersects any annuli $\mathcal{A}(\lambda_i)$, we vertically subdivide them so that the resulting annuli are either contained in μ or do not intersect the interior of μ.
(2) *Move bases to the right.* Transfer every other leading base $\lambda \subset \mu$ (except μ) onto $\overline{\mu}$ through $\mathbb{B}(\mu)$.
(3) *Collapse the naked initial segment.* Let μ_0 be closure of the $<$-coinitial maximal naked subsegment of μ. Subdivide $\mathbb{B}(\mu)$ into $\mathbb{B}(\mu_0) \cup \mathbb{B}(\mu_1)$. Denote by D_1 the added subdivision digon. Collapse $\mathbb{B}(\mu_0)$ onto $\overline{\mu_0}$ and delete D_1 because it has a free face.
(4) *Rename and clean up.* We rename μ_1 as μ and crush any π_1-trivial annuli.

Definition 3.23. The leading base μ given in the definition of the entire transformation is called the *carrier base.*

An illustration of the result of an entire transformation is given in Figure 7. This next result follows from a counting argument.

Lemma 3.24 ([BF95, Proposition 7.5])**.** *After applying an entire transformation the* τ*-complexity did not increase.*

Unlike in the thinning case, the vertical lengths of 2-cell attaching maps may increase.

3.4. **The Rips machine.** The *Rips machine*, which was first described in [BF95], is a geometric adaptation of Makanin's algorithm which takes a measured band complex $\mathscr{C}$ and produces a sequence called the *Rips sequence*:

$$\mathscr{C} = \mathscr{C}_0 \to \mathscr{C}_1 \to \dots$$

of measured band complexes, constructed inductively. We assume that $\mathscr{C}$ is equipped with an order $<$, as given in Definition 3.21.

Lemma 3.25. *Let* $\mathscr{C}$ *be equipped with an order* $<$ *and let* $m : (\mathscr{C}, t) \to (\mathscr{C}', t')$ *be a Möbius move, a thinning move, or an entire transformation. Then there is a naturally induced order* $<'$ *on* $\mathscr{C}'$*.*

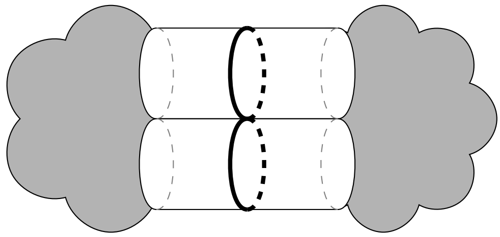

Figure 6. A band complex in terminal form; i.e., all bases are matched. As a graph of spaces, the edge space is clearly visible. There is a unique track, drawn in black, that is efficiently carried by this band complex. (Recall that tracks are connected by definition.)

Proof. A Möbius move may split a maximal section σ (i.e., delete a point) into two maximal sections σ_0, σ_1 $<$ therefore restricts to $<'$ on the new union of maximal sections.

Denote by U, U' the union of maximal sections in $\mathscr{C}, \mathscr{C}'$, respectively. If a thinning move is applied, then some maximal base is collapsed onto its dual so that $\mathscr{C} \to \mathscr{C}'$ is actually a retraction. It follows that after splitting some maximal sections the restriction $U \to U'$ is also a retraction, so there is a natural restriction $\mathscr{C}'$. For entire transformations, the initial subdivisions and transfers will at most split U into more open intervals. This is immediately followed by a collapse, so there is a well-defined $<'$ as before. □

We now describe the Rips sequence for $\mathscr{C}$ induced by a track t efficiently carried by $\mathscr{C}$.

(1) Crush any π_1-trivial annuli and delete any superfluous 2-cells.
(2) If there is point $\mathscr{C}_i$ with $\gamma(x) = 1$, we apply a thinning move collapsing the $<$-minimal maximal naked segment (Definition 3.19) to obtain $\mathscr{C}_{i+1}$.
(3) Otherwise, if there are any unmatched bases,
    (a) if possible, apply a Möbius move on a $<$-minimal Möbius band or
    (b) apply an entire transformation, then
    tighten all 2-cells attaching maps.
(4) Once all the bases are matched stop.

We note that our choice of ordering $<$ on $\mathscr{C}$ is by no means canonical. However once it is made, the Rips sequence becomes deterministic.

Proposition 3.26. *Let $t \subset C$ be a track in a 2-complex and let $\mathscr{C}$ be the corresponding measured band complex given in Proposition* 2.12. *Then after finitely many steps the Rips machine terminates on a band complex $(\mathscr{C}_T, t_T)$, with all bases matched and coinciding. The interior of the union of the bases is a regular neighbourhood of t_T homeomorphic to $t \times (-1, 1)$, t is a wedge sum of circles, and there are no connections in the interior of the bands.*

Proof. Since all bases start off with finite integer valued length, the Rips machine eventually stops since every step decreases the length of some base by a positive integer.

All bases are matched, and they must all coincide, since t_T is connected and it's efficiently carried by $\mathscr{C}_T$. Now, because all the bases are matched, if μ is any base

$|\mu|_{t_T}$, the number of connected components of t_T must equal 1. Efficient carrying also excludes the possibility of connection in the interior of a band. $\square$

Proposition 3.11 implies that the final dual Bass-Serre tree $T(t_T, \mathscr{C}_T)$ is $\pi_1(\mathscr{C})$-equivariantly isomorphic to the original $T(t, C)$. The following definition is important for the next section.

Definition 3.27. A band complex $\mathscr{C}$ is in *terminal form* if it is as described in conclusion of Proposition 3.26; see Figure 6.

4. The elimination process

We will now turn our attention to the set of all tracks that are efficiently carried by a band complex.

Let $\mathscr{C}$ be a band complex equipped with an ordering $<$ on the union of maximal sections (Definition 3.21). Then, given a track $t \subset \mathscr{C}$, the Rips machine (Section 3.4) will perform a specific derived transformation $(\mathscr{C}, t) \to (\mathscr{C}', t')$. The type of transformation, either an annulus crush, a thinning move, a Möbius move, or an entire transformation, is determined by $\mathscr{C}$ and the ordering $<$, but not the track t it carries.

Although the type of transformation doesn't depend on the track t carried by $\mathscr{C}$, the combinatorial equivalence class of the resulting $\mathscr{C}'$ does depend on the track t. For example consider Figure 7, which shows two different combinatorial outcomes coming from two different tracks carried by the same band complex. Since we want to study the set of all tracks carried by $\mathscr{C}$ we must consider all these combinatorial outcomes simultaneously.

4.0.1. *The elimination tree* $\mathfrak{T}(C)$. Given a band complex $\mathscr{C}$ with a track $t \subset \mathscr{C}$, the Rips machine gives a sequence $(\mathscr{C}, t) \to \cdots \to (\mathscr{C}_T, t_T)$ with $(\mathscr{C}_T, t_T)$ in terminal form. If we want to consider all tracks carried by $\mathscr{C}$, then we must have a branching sequence or, in other words, a rooted directed tree.

Definition 4.1. Let $\mathscr{C}$ be a band complex. Denote by $\textbf{tracks}(\mathscr{C})$ the set of tracks efficiently carried by $\mathscr{C}$.

Definition 4.2 (Combinatorially equivalent derived transformations)**.** Let t_1 and t_2 be two tracks efficiently carried by a band complex $\mathscr{C}$. The derived transformations $(\mathscr{C}, t_1) \to (\mathscr{C}_1', t_1')$ and $(\mathscr{C}, t_2) \to (\mathscr{C}_2', t_2')$ are *combinatorially equivalent* if there is a commuting homeomorphism e,

$$\begin{array}{ccc} & \mathscr{C} & \\ \swarrow & & \searrow \\ \mathscr{C}_1' & \xrightarrow{\quad e \quad} & \mathscr{C}_2' \end{array}$$

that is a combinatorial equivalence of band complexes in the sense of Definition 2.13.

Convention 4.3. In section 3.4 it was convenient to consider band complexes as being measured in order to precisely describe continuous quotient maps. For the rest of the paper, unless stated otherwise, band complexes $\mathscr{C}$ without tracks will be considered equal if they are combinatorially equivalent in the sense of Definition 2.13.

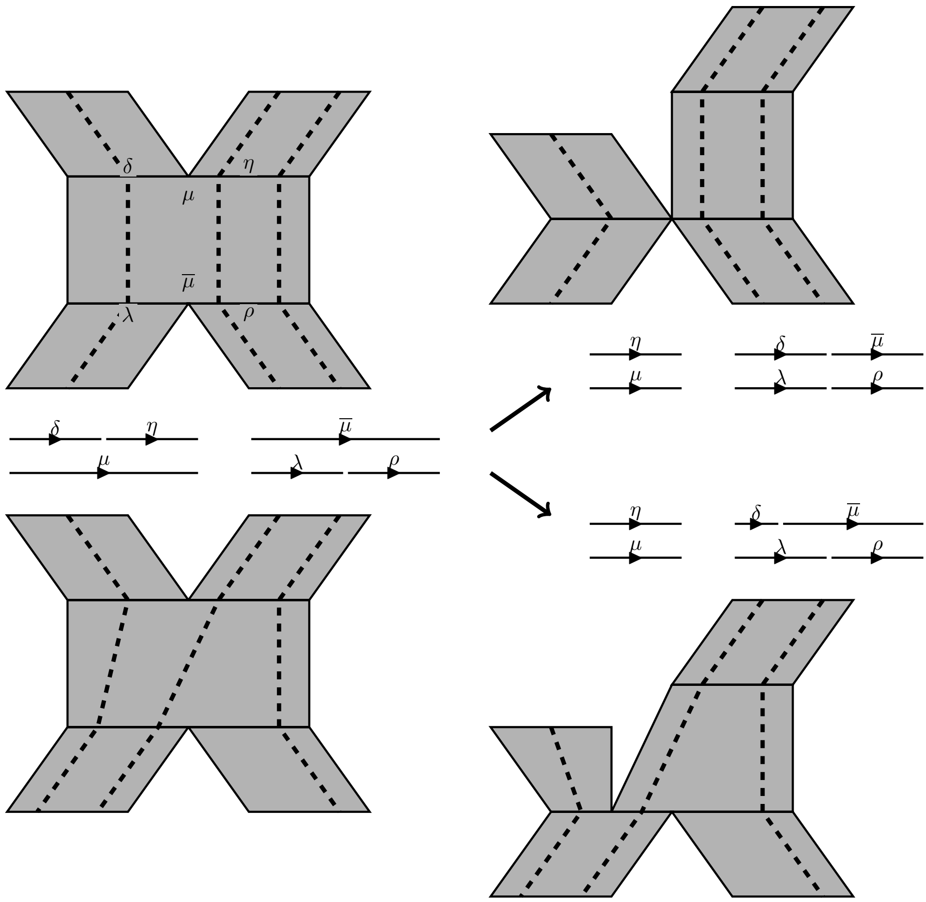

FIGURE 7. Two different tracks carried by the same band complex $\mathscr{C}$ give distinct combinatorial outcomes after applying an entire transformation (see Section 3.3.3). Here the carrier μ moves base δ onto its dual. The base diagrams illustrate the ordering $<$ (see Definition 3.21) on the union of maximal sections.

The entire transformations shown in Figure 7 are not combinatorially equivalent. We note that, in the notation of the above definition, the tracks t_1', t_2' need no longer give combinatorially equivalent derived transformations of $\mathscr{C}_1' = \mathscr{C}_2'$.

Definition 4.4 (Derived transformations of a band complex)**.** Let $\mathscr{C}$ be a band complex equipped with an ordering $<$ of the union of its maximal sections (Definition 3.21). For every $t \in \mathbf{tracks}(\mathscr{C})$ there is a corresponding derived transformation $(\mathscr{C}, t) \to (\mathscr{C}', t')$ with resulting band complex $\mathscr{C}'$. The type of this transformation, either an annulus crush, a Möbius move, a thinning move, or an entire transformation, depends on the underlying band complex $\mathscr{C}$ (and the ordering). Let $\mathscr{C}_1, \ldots, \mathscr{C}_n$ denote the finite set of combinatorial equivalence classes of the resulting band complexes. A *derived transformation of a band complex* $\mathscr{C}$ is the operation that produces the finite collection of continuous maps

$$\begin{array}{ccccc} & & \mathscr{C} & & \\ & \swarrow & & \searrow & \\ \mathscr{C}_1 & & \cdots & & \mathscr{C}_n \end{array} \tag{6}$$

We remark that we can algorithmically construct the set $\mathscr{C}_1, \ldots, \mathscr{C}_n$.

Lemma 4.5. *If the $<$ is an ordering on $\mathscr{C}$, then all its children obtained by a derived transformation have well-defined induced orderings.*

Proof. The induced ordering given in Lemma 3.25 does not depend on the track, only on the continuous map between the underlying band complexes. □

Convention 4.6. For the rest of the paper, unless stated otherwise, we will assume that a band complex $\mathscr{C}$ comes equipped with such an ordering $<$ of the union of its maximal sections.

The *elimination process* for a 2-complex C is the construction of the elimination tree $\mathfrak{T}(C)$, a directed rooted tree defined inductively as follows:

(1) The root of $\mathfrak{T}(C)$ is the polygonal complex C.
(2) The set of children of C is the finite collection of band complexes $\mathscr{C}_1, \ldots, \mathscr{C}_{n_C}$ provided by Proposition 2.14 that can efficiently carry all tracks of C. For each $\mathscr{C}_i, 1 \leq i \leq n_C$, we equip the union of maximal sections (Definition 2.7) with an ordering $<$ as in Definition 3.21.
(3) If a band complex $\mathscr{C}_v$ in $\mathfrak{T}(C)$ is in terminal form (Definition 3.27), then it is called a *terminal leaf.*
(4) If a band complex $\mathscr{C}_v$ in $\mathfrak{T}(C)$ cannot be brought to terminal form via derived moves because the union of the bands is not connected or all bases are matched, but there are 2-cell attaching maps that intersect the interior of the bands, then it is called an *inadmissible leaf.*
(5) Otherwise we continue to grow $\mathfrak{T}(C)$ at a band complex $\mathscr{C}_v$ by adding its descendants with a derived transformation (Definition 4.4). Equip the union of bases of each descendant $\mathscr{C}_{v'}$ of $\mathscr{C}_v$ with the induced order $<$ given by Lemma 4.5.

As will be explained in the next section $\mathfrak{T}(C)$ gives a way to encode the set of tracks that can lie in the polygonal complex C. It follows that in general it is infinite.

4.1. **The sets of tracks in a band complex organized by open neighbourhoods.**

Lemma 4.7. *Let $\mathscr{C} \to \mathscr{C}'$ be one of the continuous maps of the derived transformation on $\mathscr{C}$ (Definition (4.4) (6)). Suppose that $\mathscr{C}'$ is equipped with a measure (Definition 2.11). Then there is a well-defined pullback measure on $\mathscr{C}$. Furthermore, as long as the union of the interiors of the bands in $\mathscr{C}'$ is connected, the maximal measure of each base μ of $\mathscr{C}$ is no more than the sum of the measures of the bases in $\mathscr{C}'$ that are in the image of μ via the map $\mathscr{C} \to \mathscr{C}'$.*

Proof. It is now enough to consider the basic moves in Section 3.1, i.e., the zips, collapses, annulus crushes, vertical and horizontal subdivisions. In all cases given such a transformation $\hat{\mathscr{C}} \to \mathscr{C}'$ as a continuous map, there is a unique pullback measure we can put on $\hat{\mathscr{C}}$. The upper bound on the measure of the bases of $\mathscr{C}$ is obvious from the definitions of the derived moves. ☐

Corollary 4.8 (Going backwards)**.** *Let $\mathscr{C} \to \mathscr{C}'$ be one of the continuous maps of the derived transformation on $\mathscr{C}$. Suppose that $\mathscr{C}'$ efficiently carries a track t (Definition 2.9). Then there is a unique track $\hat{t}$ that is efficiently carried by $\mathscr{C}$ such that $(\mathscr{C}, \hat{t}) \to (\mathscr{C}', t)$ is a derived transformation in the Rips machine.*

Proof. Since $\mathscr{C}'$ efficiently carries t, we can put a measure on $\mathscr{C}'$ (Definition 2.11) that corresponds to t. Lemma 4.7 gives a pullback measure on $\mathscr{C}$, which induces a track $\hat{t} \subset \mathscr{C}$ which induces the derived transformation. Furthermore it is routine to check for the basic moves in Section 3.1 that $\hat{t}$ is indeed efficiently carried by $\mathscr{C}$, provided t is efficiently carried by $\mathscr{C}'$. ☐

Any band complex in terminal form (Definition 3.27) efficiently carries a unique track. We will now show how repeatedly going backwards enables us to use $\mathfrak{T}(C)$ to organize the collection of tracks that are efficiently carried by C.

Proposition 4.9. *There is a bijective correspondence between the set of tracks in C and $\{\mathscr{C}_l\}$, the set of terminal leaves (Section 4.0.1 (3)) of $\mathfrak{T}(C)$.*

Proof. Let $\mathscr{C}$ be a band complex at the top level of $\mathfrak{T}(C)$. As a topological space it is homeomorphic to C. Any track efficiently carried by $\mathscr{C}$ is obviously a track in C.

If $t \subset \mathscr{C}$ is a track, by Proposition 3.26 the Rips sequence for $(\mathscr{C}, t)$ will give a path in $\mathfrak{T}(C)$ from $\mathscr{C}$ to some band complex in terminal form $\mathscr{C}_l$. This map from tracks to leaves is injective since, having fixed $<$, the outcome of a particular derived move $(\mathscr{C}, t) \to (\mathscr{C}', t')$ depends only on the track t.

On the other hand, let $\mathscr{C}_l$ be a terminal leaf of $\mathfrak{T}(C)$. We may metrize all its bases to have length 1, so that each band is explicitly parameterized as $[0,1] \times [-1,1]$. $\mathscr{C}_l$ efficiently carries the track t_l which intersects each band as $\{\frac{1}{2}\} \times [-1,1]$. Now starting at $(\mathscr{C}_l, t_l)$ and repeatedly going backwards (Corollary 4.8) in $\mathfrak{T}(C)$, we obtain $(\mathscr{C}, t)$, where $t \subset \mathscr{C}$ is an efficiently carried track. ☐

If $\mathscr{C}$ is some band complex not in terminal form (as in Section 4 (3)), then any efficiently carried track $t \subset \mathscr{C}$ gives rise to one of the children in the elimination tree

$$\begin{array}{ccc} & \mathscr{C} & \\ \swarrow & & \searrow \\ \mathscr{C}_1 & \cdots & \mathscr{C}_n \end{array}$$

It follows that there are injective maps

$$\iota_j : \mathbf{tracks}(\mathscr{C}_j) \hookrightarrow \mathbf{tracks}(\mathscr{C}) \tag{7}$$

whose images give a cover

$$\mathbf{tracks}(\mathscr{C}) = \bigcup_j \iota_j \left(\mathbf{tracks}(\mathscr{C}_j)\right).$$

Definition 4.10. Let $\mathscr{C}_v$ be a band complex in an elimination tree $\mathfrak{T}(C)$. We denote the *track neighbourhood*

$$\mathrm{N}_{\mathfrak{T}(C)}(\mathscr{C}_v) \subset \mathbf{tracks}(\mathscr{C})$$

to be the set of tracks carried by $\mathscr{C}$ obtained by composing the maps (7) going from $\mathscr{C}_v$ all the way back to $\mathscr{C}$.

Equivalently if $\mathscr{C}_v$ is a band complex in $\mathfrak{T}(C)$, then there is a natural inclusion $\mathfrak{T}(\mathscr{C}_v) \subset \mathfrak{T}(C)$. Proposition 4.9 immediately gives the inclusion $\mathbf{tracks}(\mathscr{C}_v) \hookrightarrow \mathbf{tracks}(\mathscr{C})$ obtained by iterating (7).

4.2. **Analogies with surface train tracks.** If the reader has some familiarity with surface train tracks, the following analogies may be helpful.

If Σ is a surface, then a train track $\tau \subset \Sigma$ is analogous to a band complex structure $\mathscr{C}$ on a 2-complex C. If we assign positive integer weights to the branches of a train track τ satisfying the switch equations, then we get a multicurve in Σ, which is analogous to a pattern in C. For us a track in a 2-complex is analogous to a simple closed curve.

The assignment of weights to branches of a train track gives rise to a splitting sequence

$$\tau = \tau_0 \to \cdots$$

which will eventually split τ into a multicurve if the weights are positive integers. This is analogous to a Rips process.

On the other hand, if we put a measure on a train track, then we can consider all possible train tracks

$$\begin{array}{ccc} & \tau & \\ \swarrow & & \searrow \\ \tau_1 & \cdots & \tau_n \end{array}$$

that can be obtained from τ via a splitting move. Iterating, this gives an analogue to the elimination tree. In fact we will get an actual tree if we impose some kind of order $<$ which specifies which switch to split at each step. If eventually the train track has split itself into a simple closed curve τ_T, then we have a train track in terminal form. Assigning weight 1 to the branch and working backwards (i.e., using folding sequences) gives us a "complicated" simple closed curve in Σ.

If we were to consider the set of projectivized measured laminations, then irrational laminations would give infinite splitting sequences. Furthermore the neighbourhoods of Definition 4.10 are somewhat analogous to open neighbourhoods in the Hausdorff topology on laminations. Indeed, two laminations are "close" if the corresponding splitting sequences coincide for a long time. This all carries through to measured laminations on cell complexes, but this technology is not needed, and the ordering $<$ will cause us to stay stuck in a single minimal component.

4.3. **Inadmissibility from κ-acylindricity.** Up to now the κ-acylindricity of the dual Bass-Serre tree $T(t, \mathscr{C})$ has not been used at all.

Definition 4.11. A track t efficiently carried by a band complex $\mathscr{C}$ is called a *κ-track* if the dual tree $T(t, \mathscr{C})$ is κ-acylindrical. We denote by $\mathbf{tracks}_\kappa(\mathscr{C})$ the set of κ-tracks efficiently carried by $\mathscr{C}$.

We give two extra criteria to exclude vertices $\mathscr{C}_v$ of $\mathfrak{T}(C)$ because their track neighbourhood $\mathrm{N}_{\mathscr{C}}(\mathscr{C}_v)$ cannot contain any κ-tracks.

Lemma 4.12. *If $\pi_1(\mathscr{C})$ has no elements of order 2 and $t \in \mathbf{tracks}_\kappa(\mathscr{C})$, and if $\mu \doteqdot \overline{\mu}$, then either $(\mu, \overline{\mu})$ forms an annulus that can be crushed (which decreases the number of bands) or $|\mu|_t \leq \kappa$, where $|\mu|_t$ is the hitting measure.*

Proof. By the κ-acylindricity assumption any element that fixes an arc of length $\kappa+1$ in the dual tree $T(t,\mathscr{C})$ must be trivial. By assumption $\mathbb{B}(\mu)$ either forms an annulus $\mathcal{A}(\mu)$ or a Möbius strip $\mathcal{M}(\mu)$. In both cases $\mathrm{Gp}(\mathcal{A}(\mu))$ or $\mathrm{Gp}(\mathcal{M}(\mu))$ is generated by an element $g \neq 1$, since clause (1) in the Rips machine (Section 3.4) crushes π_1-trivial annuli and $\mathrm{Gp}(\mathcal{M}(\mu))$ must act non-trivially on $T(t,\mathscr{C})$.

We pass to the universal cover and consider the equivariant map π of Proposition 1.6 (2), and we see that in the annulus case g fixes an arc of $T(t,\mathscr{C})$ of length $|\mu|_t$. In the Möbius strip case $g^2 \neq 1$ fixes an arc of length $|\mu|_t$ in $T(t,\mathscr{C})$. In the Möbius band case this forces $|\mu|_t$ to be at most κ. In the annulus case, if $|\mu|_t > \kappa$, then we can crush it since $\mathrm{Gp}(\mathcal{A}(\mu)) = \{1\}$. □

As an immediate corollary we have:

Proposition 4.13. *Let $\mathscr{C}_v$ be a band complex in $\mathfrak{T}(C)$. If along some path*

$$p : \mathscr{C}_u \to \cdots \to \mathscr{C}_v$$

in $\mathfrak{T}(C)$ either some annulus gets subdivided more than $\kappa+1$ times or some base μ that formed a Möbius pair with its dual gets shortened or vertically subdivided more than $\kappa+1$ times, then $\mathrm{N}_{\mathfrak{T}(C)}(\mathscr{C}_v)$ doesn't contain any κ-acylindrical tracks.

Definition 4.14. We call a path p in $\mathfrak{T}(C)$ such as the one given in Proposition 4.13 *κ-inadmissible.*

4.4. **Automorphic minimality and repetitions.**

Definition 4.15 (Size and minimality)**.** The *size* of a track $t \subset \mathscr{C}$ in a band complex is the finite sum

$$\mathrm{Size}(t) = \sum_{\mu} |\mu \cap t|$$

where μ ranges over the bases of $\mathscr{C}$. A track t is called *automorphically minimal* if among all other tracks t' such that $t \sim_{\mathrm{Aut}(\pi_1(\mathscr{C}))} t'$, $\mathrm{Size}(t) \leq \mathrm{Size}(t')$.

This next lemma is easy to prove from the definitions of the basic moves.

Lemma 4.16. *Let $m : \mathscr{C} \to \mathscr{C}'$ be a basic transformation. If t'_1 is a track efficiently carried by $\mathscr{C}'$, then there is a corresponding track t_1 efficiently carried by $\mathscr{C}$ such that $m : (\mathscr{C}, t_1) \to (\mathscr{C}', t'_1)$. Furthermore if t'_2 is efficiently carried by $\mathscr{C}'$ and $\mathrm{Size}(t'_1) < \mathrm{Size}(t'_2)$, then $\mathrm{Size}(t_1) < \mathrm{Size}(t_2)$, where t_2 is the track efficiently carried by $\mathscr{C}$ corresponding to t_2.*

This notion of automorphic minimality may seem convoluted, but the proof of the following proposition may clear things up for the reader.

Proposition 4.17 (Repetitions and minimality)**.** *Let $\mathscr{C}_v$ be a band complex in $\mathfrak{T}(C)$. If along some path*

$$p : \mathscr{C}_u \to \cdots \to \mathscr{C}_v$$

in $\mathfrak{T}(C)$ there are two combinatorially equivalent band complexes (*Definition* 2.13) *$\mathscr{C}_u \overset{\approx}{\to} \mathscr{C}_v$, i.e., a* repetition, *then $\mathrm{N}_{\mathfrak{T}(C)}(\mathscr{C}_v)$ cannot contain any minimal tracks.*

Proof. Suppose towards a contradiction that there was a track $t \subset \mathscr{C}$ in $\mathrm{N}_{\mathfrak{T}(C)}(\mathscr{C}_v)$ that was minimal. Let $q_v : \mathscr{C}_v \to \cdots \to \mathscr{C}_l$ be the path to the terminal leaf in $\mathfrak{T}(C)$ corresponding to t (recall Proposition 4.9).

Let p_u be the concatenation of paths p and q_v, i.e., $p_u : \mathscr{C}_u \to \cdots \to \mathscr{C}_v \to \cdots \to \mathscr{C}_l$. Since $\mathscr{C}_u \approx \mathscr{C}_v$, we can attach the path q_v to $\mathscr{C}_u$ to get a corresponding path

$q_u : \mathscr{C}_u \to \cdots \to \mathscr{C}_{l'}$, where $\mathscr{C}_{l'} \approx \mathscr{C}_l$. Let $t' \subset \mathscr{C}$ be the track corresponding to $\mathscr{C}_{l'}$, and denote by t'_i (respectively t_i) the image of t' (respectively t) in $\mathscr{C}_i$, should there be such an image. Derived moves always decrease the lengths of bases; thus, in $\mathscr{C}_u$, $\mathrm{Size}(t'_u) < \mathrm{Size}(t_u)$. Working backwards in $\mathfrak{T}(C)$, i.e., repeatedly applying Corollary 4.8, all the way back to a direct descendant $\mathscr{C}$ of C yields $\mathrm{Size}(t') < \mathrm{Size}(t)$.

On one hand, by the definitions of q_v and q_u, we have $(\mathscr{C}_v, t_v) \overset{\approx}{\to} (\mathscr{C}_u, t'_u)$. On the other hand we have a sequence of derived moves

$$(\mathscr{C}_u, t_u) \to \cdots \to (\mathscr{C}_v, t_v).$$

This gives a composition of continuous maps

$$\varphi : \mathscr{C}_u \to \cdots \to \mathscr{C}_v \overset{\approx}{\to} \mathscr{C}_u,$$

which by Proposition 3.11 induces an isomorphism on π_1; hence $\varphi_\sharp \in \mathrm{Aut}(\pi_1(\mathscr{C}_u))$. Furthermore, by construction $\varphi(t_u) = t'_u$, so again Proposition 3.11 gives us that $T(t_u, \mathscr{C}_u)$ is $\varphi_\sharp$-equivariantly isomorphic to $T(t'_u, \mathscr{C}_u)$, where the action on the second tree is given by $(g, x) \mapsto \varphi_\sharp(g) \cdot x$. It follows that t' and t are automorphically equivalent, contradicting the minimality of t. □

Definition 4.18. A path p in $\mathfrak{T}(C)$ that satisfies the hypotheses of Proposition 4.17 is called *repetition-inadmissible.*

4.5. **Restricted elimination processes.** In order to construct $\mathfrak{T}(C)$ we will sometimes have to construct auxiliary elimination trees that are rooted at band complexes $\mathscr{C}$ in $\mathfrak{T}(C)$.

Let $J \subset \mathscr{C}$ be a union of maximal sections. We redefine the order $<$ so that the maximal sections in J are terminal. The restricted elimination tree is used to study how the bases of $\mathscr{C}$ can be moved into J. If $\mathscr{C} \to \mathscr{C}'$ is a derived transformation which is a continuous map, then J has a well-defined image in $\mathscr{C}'$ which we also denote by J. $\mathfrak{T}(\mathscr{C}, J)$ is constructed as follows:

1. The root of $\mathfrak{T}(\mathscr{C}, J)$ is $\mathscr{C}$.
2. If every base of a band complex $\mathscr{C}_l$ in $\mathfrak{T}(\mathscr{C}, J)$ is contained in J, then $\mathscr{C}_l$ is called a *J-terminal leaf.*
3. *Inadmissible leaves* are defined the same way as for $\mathfrak{T}(C)$.
4. Otherwise we apply a corresponding derived transformation, either a Möbius move on $(\mu, \overline{\mu})$, a thinning move starting at μ, or an entire transformation with leading base μ to create the children of $\mathscr{C}_v$. Equip the union of bases of every child $\mathscr{C}'_v$ of $\mathscr{C}_v$ with the induced order $<$.

Restricted elimination processes will be required for some subprocesses of our main algorithm. It is obvious that the various inadmissibility criteria for a standard elimination tree $\mathfrak{T}(C)$ also hold for restricted elimination trees.

Convention 4.19. Many statements about restricted elimination trees will also follow for the standard elimination tree by replacing $\mathfrak{T}(C)$ by the elimination trees $\mathfrak{T}(\mathscr{C}_1, \emptyset), \ldots, \mathfrak{T}(\mathscr{C}_{n_C}, \emptyset)$ where $\mathscr{C}_1, \ldots, \mathscr{C}_{n_C}$ are the children of C (see step (2) in Section 4.0.1). These elimination trees are contained in $\mathfrak{T}(C)$. We will therefore assume that results about restricted elimination processes will apply to the standard elimination process, the latter being a special case.

4.6. **The infinite branches of $\mathfrak{T}(C)$.** One of Makanin's key observations is that every infinite branch of $\mathfrak{T}(C)$ stabilizes into one of three cases. The main ingredient is the following lemma.

Lemma 4.20. *If after applying an entire transformation which didn't involve a Möbius move we are in the situation where we must make a thinning move, then (J-relative) τ-complexity decreased.*

In the case of band complexes, this fact is explained between Proposition 7.5 and Proposition 7.6 of [BF95]. The proof consists of a straightforward complexity counting argument. This next result is ubiquitous whenever the elimination process/Rips machine is involved.

Theorem 4.21 (Fundamental classification)**.** *Every infinite branch $\mathfrak{b}$ in $\mathfrak{T}(C)$ ($\mathfrak{T}(\mathscr{C}, J)$) has a tail $\mathfrak{b}_v = \mathscr{C}_v \to \cdots$ of one of the following forms:*

- (i) Thinning: *Every derived transformation along $\mathfrak{b}_v$ is a thinning move.*
- (ii) Quadratic: *Every derived transformation along $\mathfrak{b}_v$ is an entire transformation. For all but finitely many points in the (J-complement of the) union of bases of the band complexes $\mathscr{C}_w$ along $\mathfrak{b}_v$ we have $\gamma(x) = 2$ (Definition 3.14).*
- (iii) Superquadratic: *Every derived transformation along $\mathfrak{b}_v$ is an entire transformation. There is a whole open interval of points in the (J-complement of the) union of bases of the band complexes $\mathscr{C}_w$ along $\mathfrak{b}_v$ such that $\gamma(x) \geq 3$.*

Furthermore if we require $\mathrm{N}_{\mathfrak{T}(C)}(\mathscr{C}_v)$ to contain a κ-track, then we may assume that no Möbius moves or annulus subdivisions occur along $\mathfrak{b}_v$.

In (ii) above, we would really like to say that every point in the union of bases is contained in exactly two unmatched bases. Points on the boundary of bases, however, may be contained in up to four distinct bases, but there are only finitely many of them.

Proof. If the (J-restricted) τ-complexity is 0, because we are not allowing long bands (see step 4 of the thinning move, Section 3.3.2), all bases are matched (all bases moved onto J) so we are at a leaf. It therefore follows by Lemma 4.20 that the infinite branch $\mathfrak{b}$ eventually always consists of thinning moves or eventually always consists of entire transformations.

Suppose now that $\mathfrak{b}_v$ is not of thinning type. If $\mathscr{C}_v$ is of quadratic type, then after applying an entire transformation $\mathscr{C}_v'$ is still quadratic. The trichotomy now follows. The fact that Möbius moves and annulus subdivisions stop occurring follows from Lemma 4.12. □

König's Lemma states that every infinite rooted tree with vertices of finite valency must have an infinite branch. This classification of infinite branches is the foundation of the construction of a finite subtree of $\mathfrak{T}(C)$ containing all the leaves corresponding to minimal κ-acylindrical tracks.

4.7. **The admissible subtree $\mathfrak{A}(\mathscr{C}, J)$.**

Definition 4.22. The *admissible elimination tree* $\mathfrak{A}(\mathscr{C}, J) \subset \mathfrak{T}(\mathscr{C}, J)$ is the subtree obtained by forbidding κ-inadmissible and repetition-inadmissible subpaths (Definitions 4.14 and 4.18). We similarly define the admissible elimination tree $\mathfrak{A}(C) \subset \mathfrak{T}(C)$.

This next proposition enables us to restrict the search for tracks in the algorithm for Theorem B to admissible elimination trees.

Proposition 4.23. *Let t be an automorphically minimal κ-track efficiently carried by $\mathscr{C}$. Any path $\mathscr{C} \to \cdots$ in $\mathfrak{T}(\mathscr{C}, J)$ induced by $t \subset \mathscr{C}$ must be contained in $\mathfrak{A}(\mathscr{C}, J)$.*

Proof. Otherwise Proposition 4.14 or 4.18 leads to a contradiction of the hypotheses. □

Proposition 4.24. *For every n, the subtree of radius n of $\mathfrak{A}(\mathscr{C}, J)$ can be effectively constructed.*

Proof. For any band complex the collection of children (see Definition 4.4 (6)) can be constructed effectively, and the various inadmissibility conditions can be verified effectively. □

4.8. **Reduction to the superquadratic case.** We show that infinite thinning or quadratic branch in $\mathfrak{T}(\mathscr{C}, J)$ contains a repetition. From this it will follow that the admissible elimination tree $\mathfrak{A}(\mathscr{C}, J)$ does not have any infinite thinning or quadratic branches.

Lemma 4.25 (cf. [KM98, Lemma 15])**.** *Any sufficiently long thinning path $\mathscr{C}_v \to \cdots$ contains a repetition, i.e., a subpath $\mathscr{C}_u \to \cdots \to \mathscr{C}_w$ with $\mathscr{C}_u \approx \mathscr{C}_w$.*

Proof. By Theorem 4.21 we may assume that no more annulus subdivisions occur. By Lemma 3.20, the number of 2-cells and the vertical lengths of the 2-cell attaching maps are non-increasing. Furthermore since the τ-complexity is bounded and there are no maximal sections σ with $\tau(\sigma) = 0$ there is a bound on the number of bases.

A band complex is obtained by gluing bands to a graph Γ and then attaching other 2-cells. Since derived transformations do not change anything in the exterior of union of the bands, the number of bands and 2-cells remains bounded, and the combinatorial lengths of the attaching immersions of the 2-cells is bounded; any sufficiently long thinning path will have a repetition. □

This next lemma will also be used later in Section 6.2.

Lemma 4.26 (cf. [KM98, Case 14])**.** *Any sufficiently long quadratic path $\mathscr{C}_v \to \cdots$ contains a repetition.*

Proof. By Lemma 3.24, the τ complexity doesn't increase, and we may assume that no Möbius moves or annulus subdivisions occur. It remains to show that the vertical lengths of 2-cell attaching maps remain bounded; the result will then follow as in the previous proof.

Consider Figure 8 with leading bases μ, λ, where μ is the carrier. The only way the attaching map of a 2-cell will decrease in vertical length is if it has a subpath as in the shaded path on the left of Figure 8 that travels through $\mathbb{B}(\lambda) \cup \mathbb{B}(\mu)$. The only way for a segment in the boundary of a 2-cell to get "stretched" is if it travels through the vertical sides of $\mathbb{B}(\lambda)$ and $\mathbb{B}(\eta)$. In this case we have a segment σ of length 2 that gets stretched to a segment σ' of length 3. After the entire transformation, however, η and μ are now leading bases. This means that after the next entire transformation there is a subsegment of length 2 σ' that gets shortened again back to length 1.

It therefore follows that the vertical lengths of the boundaries of 2-cells remain bounded throughout the quadratic path, and the result follows. □

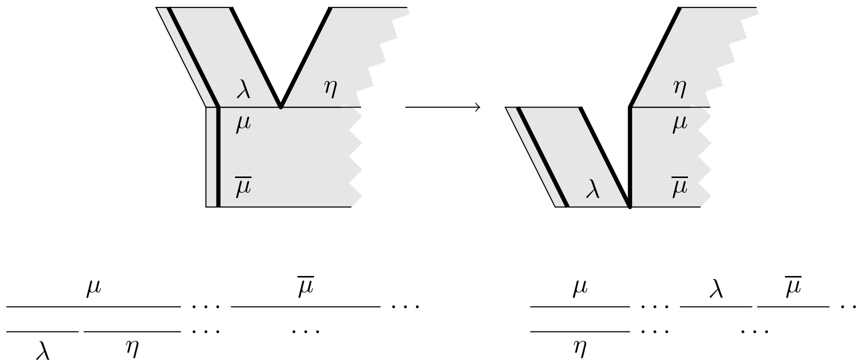

FIGURE 8. Above, the attaching maps of 2-cells under an entire transformation in the quadratic case. Below, the resulting ordered base configurations.

These two lemmas immediately imply the following corollary, which pretty much sets the tone for the rest of the paper.

Corollary 4.27. *The admissible elimination subtree* $\mathfrak{A}(\mathscr{C}, J) \subset \mathfrak{T}(\mathscr{C}, J)$ *does not have any infinite paths of thinning or quadratic type. Equivalently, all infinite paths in* $\mathfrak{A}(\mathscr{C}, J)$ *have superquadratic tails.*

5. Overlapping pairs and periodic mergers

We start our attack of the superquadratic case by examining overlapping pairs and by introducing a new move: the periodic merger (precisely defined in Section 5.3). Throughout this section π will denote the map $\pi : \widetilde{\mathscr{C}} \to T(t, \mathscr{C})$ given in Proposition 1.6 (2).

Definition 5.1. A dual pair $(\mu, \overline{\mu})$ is an *overlapping pair* if **interior** $(\mu \cap \overline{\mu}) \neq \emptyset$ and the pair is *orientation preserving*, i.e., if the image of $\mathbb{B}(\mu) \subset \mathscr{C}$ does not contain an embedded Möbius band.

Convention 5.2. When $(\mu, \overline{\mu})$ is an overlapping pair we will assume that $\mu < \overline{\mu}$, where $<$ is the ordering on $\mathscr{C}$ (Definition 3.21).

Definition 5.3. Let $\mathbb{B}(\mu)$ be a band such that $(\mu, \overline{\mu})$ is an overlapping pair. Let $p \in \mu$ be the $<$-initial point (Definition 3.21). The *tubular loop* τ_μ is the loop $\alpha * \beta$ where α is the path in the side of $\mathbb{B}(\mu)$ starting at p and going from μ to $\overline{\mu}$ and β is the path in μ connecting the endpoint of α to p. See Figure 9.

We note that our definition of a tubular loop is an oriented based loop. This gives rise to an element $g \in \pi_1(\mathscr{C})$ that we will call a *μ-tubular element*. For the rest of the paper we will avoid mentioning the basepoint.

Suppose that $\mathscr{C}$ carries a track t. Figure 10 illustrates the action of the μ-tubular element τ_μ on $T(t, \mathscr{C})$ given by the quotient map π of Proposition 1.6. $p \in \mathscr{C}$ is as in Definition 5.3, and $\widetilde{p} \in \widetilde{\mu}$ is a lift in the universal cover $\widetilde{\mathscr{C}}$ of $p \in \mu \subset \mathscr{C}$. It

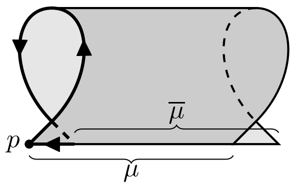

FIGURE 9. An overlapping pair. The dark loop indicated on the left is the tubular loop τ_μ.

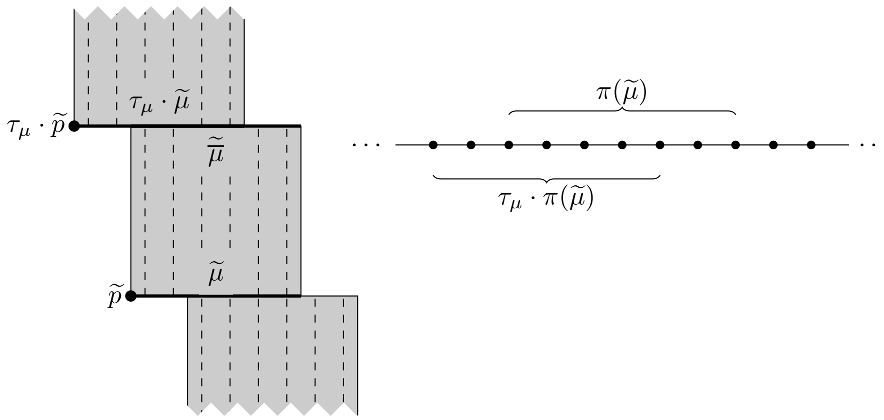

FIGURE 10. On the left, the action of a tubular element on the universal cover by deck transformations; on the right the corresponding action on the tree $T(P, C)$. Tracks are shown as dashed lines.

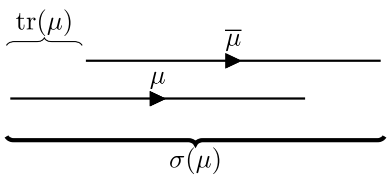

FIGURE 11. The translation length $\mathrm{tr}(\mu)$ and $\sigma(\mu)$, the section corresponding to μ given in Definition 5.5.

is evident, for any track t efficiently carried by $\mathscr{C}$, that τ_μ must fix some axis of $T(t, \mathscr{C})$. Closer examination immediately yields:

Lemma 5.4. *If $\mathscr{C}$ efficiently carries a track t and $(\mu, \overline{\mu})$ is an overlapping pair, then the μ-tubular element τ_μ acts on $T(t, \mathscr{C})$ hyperbolically with translation length denoted*

$$\mathrm{tr}(\mu) = |\mu \setminus \overline{\mu}|_t.$$

One of the principal features of a group acting acylindrically on a tree is that infinite line stabilizers are cyclic. It follows that if there are two overlapping pairs that themselves overlap sufficiently, the corresponding tubular elements must fix a common axis and therefore must lie in a common cyclic subgroup. In Section 5.3 we

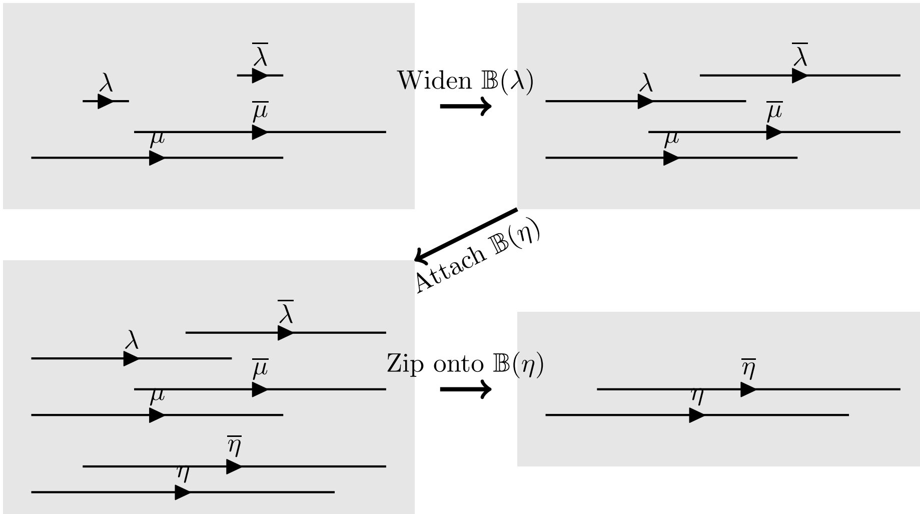

FIGURE 12. A periodic merger (defined in Section 5.3). In this example $\mathrm{tr}(\tau_\mu) = 2, \mathrm{tr}(\tau_\lambda) = 3$, and $\mathrm{tr}(\tau_\eta) = \gcd(2,3) = 1$.

will describe the periodic merger, a move from [AHT06], which will replace these two overlapping pairs by a single overlapping pair. This is illustrated in Figure 12. There are two subtleties to this operation which do not occur when deciding if an equation has a solution or if we simply want to count orbits as in [AHT06].

- We must merge these bands into one while preserving the fundamental group of the band complex as well as the dual Bass-Serre tree.
- We must be able to *algorithmically* produce a finite list containing all combinatorial outcomes of a periodic merger for $(\mathscr{C}, t)$, where t ranges over $\mathbf{tracks}_\kappa(\mathscr{C})$ (Definition 4.11).

There is one outstanding difficulty: given two commuting elements $g, h \in \pi_1(\mathscr{C})$, deciding if they lie in a common cyclic subgroup. Only being able to solve the word problem in $\pi_1(\mathscr{C})$ is insufficient to solve this problem in general. To overcome this impasse we will use the author's generalized Bulitko Lemma [Tou14] in a way that is completely different from its usual purpose.

5.1. **Interactions with tubular elements: Entanglement.**

Definition 5.5. If $(\mu, \overline{\mu})$ is an overlapping pair, then we denote the *section corresponding to* $(\mu, \overline{\mu})$ as

$$\sigma(\mu) = \mu \cup \overline{\mu}.$$

Suppose now that there is another band $\mathbb{B}(\lambda)$ whose unmatched bases both lie in $\sigma(\mu)$, as in Figure 13.

Convention 5.6. We write $(\lambda, \overline{\lambda}) \subset \sigma(\mu)$ to signify $\lambda \cup \overline{\lambda} \subset \sigma(\mu)$. We will always assume that base pairs $(\lambda, \overline{\lambda})$, $(\mu, \overline{\mu})$ are unmatched.

To study how these bands interact we have the following:

Definition 5.7. Let $(\mu, \overline{\mu})$ be overlapping and let $(\lambda, \overline{\lambda}) \subset \sigma(\mu)$ with $\lambda < \overline{\lambda}$. Let p be as in Definition 5.3. The *μ-relative loop* $\rho^\lambda_{\sigma(\mu)}$ is the concatenation $\alpha * \beta * \gamma$

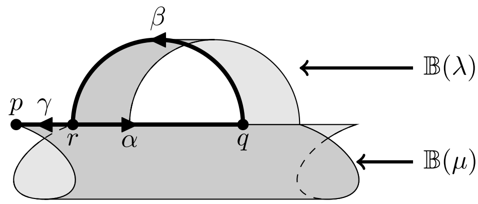

FIGURE 13. The band $\mathbb{B}(\lambda)$ has both bases lying in the section $\sigma(\mu)$. The μ-relative loop $\rho^{\lambda}_{\sigma(\mu)}$ is the loop $\alpha * \beta * \lambda$.

shown in Figure 13, where α is the path from p to q in $\sigma(\mu)$, β is the path from q to r travelling along a side of $\mathbb{B}(\lambda)$, and γ is the path in $\sigma(\mu)$ from r to p.

Again fixing a lift $\widetilde{p}$ of p and $\widetilde{\mu}$ of μ in the universal cover $\widetilde{\mathscr{C}}$ of $\mathscr{C}$, we can describe the deck transformation given by $\rho^{\lambda}_{\sigma(\mu)}$ (see Figure 14). From this we immediately get:

Lemma 5.8. *If $\mathscr{C}$ efficiently carries a track t, then in $T(t, \mathscr{C})$ the length of the arc*

$$\operatorname{axis}(\tau_\mu) \cap \left(\rho^{\lambda}_{\sigma(\mu)} \cdot \operatorname{axis}(\tau_\mu)\right)$$

is at least $|\lambda|_t$.

Definition 5.9. We define $\operatorname{tr}_{\sigma(\mu)}(\lambda)$ to be the measure of the arc between the leftmost point of λ and the leftmost point of $\overline{\lambda}$, i.e., the length $|\alpha|_t - |\gamma|_t$ as shown in Figure 13. We say the dual pair $(\lambda, \overline{\lambda})$ is *orientation preserving* if the holonomy $\lambda \to \overline{\lambda}$ extends to an orientation preserving homeomorphism of $\sigma(\mu)$ (relative to its endpoints).

For the following three lemmas assume that $t \in \mathbf{tracks}_\kappa(\mathscr{C})$.

Definition 5.10. If $(\mu, \overline{\mu})$ is an overlapping pair, $(\lambda, \overline{\lambda}) \subset \sigma(\mu)$ with $(\lambda, \overline{\lambda})$ orientation preserving, and the commutator $[\tau_\mu, \rho^{\lambda}_{\sigma(\mu)}] = 1$, then we say that dual pairs $(\mu, \overline{\mu})$ and $(\lambda, \overline{\lambda})$ are *entangled.*

The following is obvious but necessary for computational considerations.

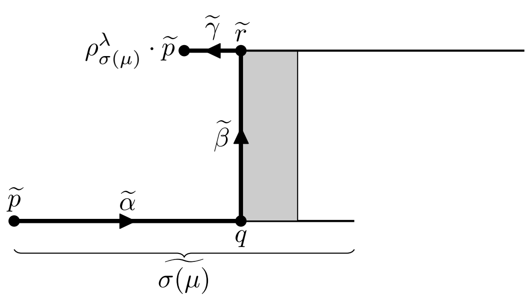

FIGURE 14. The deck transformation corresponding to $\rho^{\lambda}_{\sigma(\mu)}$. The lifts of α, β, and γ of Figure 13 are shown.

Lemma 5.11. *Let $\mathscr{C}$ be a band complex with $(\mu,\overline{\mu})$ an overlapping pair, and let $(\lambda,\overline{\lambda}) \subset \sigma(\mu)$ such that $(\lambda,\overline{\lambda})$ is orientation preserving. If we can solve the word problem in $\pi_1(\mathscr{C})$, then we can decide if $(\mu,\overline{\mu})$ and $(\lambda,\overline{\lambda})$ are entangled.*

Lemma 5.12. *Let $(\mu,\overline{\mu})$ be an overlapping pair, suppose $(\lambda,\overline{\lambda}) \subset \sigma(\mu)$, and that $(\lambda,\overline{\lambda})$ is orientation preserving and unmatched. If $[\tau_\mu,\rho^\lambda_{\sigma(\mu)}] = 1$, i.e., $(\mu,\overline{\mu})$ and $(\lambda,\overline{\lambda})$ are entangled, then $\rho^\lambda_{\sigma(\mu)}$ acts hyperbolically on $T(t,\mathscr{C})$,*

$$\mathrm{axis}(\rho^\lambda_{\sigma(\mu)}) = \mathrm{axis}(\tau_\mu),$$

and the translation length of $\rho^\lambda_{\sigma(\mu)}$ is $\mathrm{tr}_{\sigma(\mu)}(\lambda)$.

Proof. Because $[\tau_\mu,\rho^\lambda_{\sigma(\mu)}] = 1$,

$$\rho^\lambda_{\sigma(\mu)} \cdot \mathrm{axis}(\tau_\mu) = \mathrm{axis}\left(\rho^\lambda_{\sigma(\mu)}\tau_\mu(\rho^\lambda_{\sigma(\mu)})^{-1}\right) = \mathrm{axis}(\tau_\mu).$$

It therefore follows that $\langle\rho^\lambda_{\sigma(\mu)}\rangle$ fixes a bi-infinite arc in $T(t,\mathscr{C})$. From Figure 14, $\rho^\lambda_{\sigma(\mu)}$ translates this arc by $\mathrm{tr}(\rho^\lambda_{\sigma(\mu)})$ (as defined in Definition 5.9); it therefore follows that $\mathrm{axis}(\tau_\mu)$ is the minimal invariant subtree for $\langle\rho^\lambda_{\sigma(\mu)}\rangle$. ☐

Lemma 5.13. *Let $(\mu,\overline{\mu})$ be an overlapping pair, let $(\lambda,\overline{\lambda}) \subset \sigma(\mu)$, and let $(\lambda,\overline{\lambda})$ be orientation preserving. If $|\lambda|_t \geq \mathrm{tr}(\mu)+\kappa+1$, then $[\tau_\mu,\rho^\lambda_{\sigma(\mu)}] = 1$; i.e., $(\mu,\overline{\mu})$ and $(\lambda,\overline{\lambda})$ are entangled.*

Proof. Let

$$I = \mathrm{axis}(\tau_\mu) \cap \mathrm{axis}\big(\rho^\lambda_{\sigma(\mu)}\tau_\mu(\rho^\lambda_{\sigma(\mu)})^{-1}\big).$$

By Lemma 5.8 we have $|I| \geq \mathrm{tr}(\mu)+\kappa+1$. Let J be a co-initial (or co-final) subsegment of I with$|J| \geq \kappa+1$. Without loss of generality (up to choosing if J is co-initial or co-final) we may assume that

$$\tau_\mu \cdot J \subset I \subset \mathrm{axis}\big(\rho^\lambda_{\sigma(\mu)}\tau_\mu(\rho^\lambda_{\sigma(\mu)})^{-1}\big),$$

and since λ is orientation preserving and since translation length is invariant under conjugation we have

$$\big(\rho^\lambda_{\sigma(\mu)}\tau_\mu^{-1}(\rho^\lambda_{\sigma(\mu)})^{-1}\big)\cdot(\tau_\mu \cdot J) = J,$$

which by κ-acylindricity implies that $[\tau_\mu,\rho^\lambda_{\sigma(\mu)}] = 1$. ☐

5.2. **The Bulitko trick.** If we are given a band complex $\mathscr{C}$ with an overlapping pair $(\mu,\overline{\mu})$ that is entangled with $(\lambda,\overline{\lambda})$, then, by Lemma 5.12, if $\mathbf{tracks}_\kappa(\mathscr{C}) \neq \emptyset$, $\langle\tau_\mu,\rho^\lambda_{\sigma(\mu)}\rangle$ must be cyclic. In particular there is $\langle g\rangle = \langle\tau_\mu,\rho^\lambda_{\sigma(\mu)}\rangle$ so that $g^n = \tau_\mu$ and $g^m = \rho^\lambda_{\sigma(\mu)}$ with n,m minimal in absolute value. The Bulitko trick will either enable us to find n,m or will certify that $\mathbf{tracks}_\kappa(\mathscr{C}) = \emptyset$.

Let $\phi : \pi_1(\mathscr{C}) \to \mathbb{H}$ be a homomorphism to the fundamental group of a one edged κ-acylindrical graph of groups. Then as long as some element of $\langle g\rangle$ is sent to a hyperbolic element n,m will still be the minimal integers such that $\phi(g)^n = \phi(\tau_\mu)$ and $\phi(g)^m = \phi(\rho^\lambda_{\sigma(\mu)})$.

Since $[\rho^\lambda_{\sigma(\mu)},\tau_\mu] = 1$ then any $\phi : \pi_1(\mathscr{C}) \to \mathbb{H}$ will send τ_μ and $\rho^\lambda_{\sigma(\mu)}$ to elements fixing a common axis, provided their images are hyperbolic. In this case, the acylindrical Bulitko Lemma provides a computable function depending only on τ_μ

and $\rho^{\lambda}_{\sigma(\mu)}$ that bounds n, m given above. Before giving the lemma we must first present the necessary terminology.

Let $\mathbb{H}$ act κ-acylindrically on a based tree (T, t_0), let p be some hyperbolic element of $\mathbb{H}$ (with respect to the given splitting of $\mathbb{H}$), and let $L \subset T$ denote its axis. Consider the set of segments

$$\{L' \subset L \mid L = \bigcup_{n\in\mathbb{Z}} p^n L'\}.$$

A minimal element of this set with respect to inclusion is called a *fundamental domain* of L.

For an element $h \in \mathbb{H}$, let $[v_0, h \cdot v_0]$ denote the geodesic between v_0 and $h \cdot v_0$. Let $g \in \mathbb{H}$. If a segment

$$\sigma = [v_0, h \cdot v_0] \cap g \cdot L$$

is non-empty, then we call it an L-periodic subsegment of $[v_0, h \cdot v_0]$. The L-*periodicity* of σ is the integer

$$\lfloor \frac{|\sigma|}{|L_0|} \rfloor,$$

where L_0 is a fundamental domain of L. We can now state

Theorem 5.14 (Acylindrical Bulitko Lemma, Theorem 1.3 of [Tou14])**.** *There exists a computable function $\mathfrak{n} : \mathbb{N} \times \mathbb{N} \times \mathbb{N} \to \mathbb{N}$ such that for any non-trivial homomorphism $\phi : G \to \mathbb{H}$, where the group G has a finite presentation $\langle Y \mid S\rangle$ and the group $\mathbb{H}$ has a κ-acylindrical splitting with based Bass-Serre tree (T, t_0), and for any hyperbolic element in $p \in \mathbb{H}$ (denote its axis $L \subset T$), there exists a homomorphism $\phi^* : G \to \mathbb{H}$ such that for all $y \in Y$:*

- *if $[t_0, \phi(y) \cdot t_o]$ has no L-periodic subsegments, then $\phi(y) = \phi^*(y)$; and*
- *if $[t_0, \phi(y) \cdot t_o]$ has L-periodic subsegments, then there is a bijective correspondence between the L-periodic subsegments of $[t_0, \phi(y)\cdot t_o]$ and $[t_0, \phi(y)^* \cdot t_o]$, but the L-periodicity of all the periodic subsegments of $[t_0, \phi(y)^* \cdot t_o]$ is at most $\mathfrak{n}(|Y|, |S|, \kappa)$.*

Proposition 5.15. *Let $(\mu, \overline{\mu})$ be an overlapping pair in a band complex $\mathscr{C}$ and suppose it is entangled with $(\lambda, \overline{\lambda})$. There is an algorithm which terminates with one of the two following outputs:*

(i) *It gives an element g such that $\langle g\rangle = \langle \tau_\mu, \rho^{\lambda}_{\sigma(\mu)}\rangle$.*
(ii) *It (correctly) certifies that $\mathbf{tracks}_\kappa(\mathscr{C}) = \emptyset$.*

Proof. Let $G = \pi_1(\mathscr{C})$. Since $\mathscr{C}$ is an explicitly given cell complex and since we can solve the word problem in $\pi_1(\mathscr{C})$, it is possible to give a finite presentation $\langle Y \mid S\rangle$ of G where τ_μ and $\rho^{\lambda}_{\sigma(\mu)}$ are included in the generating set Y.

Let $\mathfrak{n}$ be the computable function given by Theorem 5.14 and let $M = \mathfrak{n}(|Y|, |S|, \kappa)$. For every pair n_i, m_j of absolute value less than M let (u_i, v_j) be a pair such that $u_i n_i + v_i m_j = \gcd(n_i, m_i)$. Let $g_{ij} = (\tau_\mu)^{u_i} (\rho^{\lambda}_{\sigma(\mu)})^{v_j}$ and check whether $(g_{ij})^{n_i} = \tau_\mu$ and $(g_{ij})^{m_j} = \rho^{\lambda}_{\sigma(\mu)}$. If we find some g_{ij} satisfying item (i), then we stop. Otherwise if all these verifications are negative we know that $\mathbf{tracks}_\kappa(\mathscr{C}) = \emptyset$.

Indeed suppose towards a contradiction that none of the g_{ij} are roots of τ_μ and $\rho^{\lambda}_{\sigma(\mu)}$, but that there is some $t \in \mathbf{tracks}_\kappa(\mathscr{C})$. By Lemma 5.12 there is some $g \in G$ such that $\langle g\rangle$ is the maximal cyclic group stabilizing $L = \mathrm{axis}(\tau_\mu) = \mathrm{axis}\Big(\rho^{\lambda}_{\sigma(\mu)}\Big)$.

Let n, m be the integers such that $g^n = \tau_\mu$ and $g^m = \rho^\lambda_{\sigma(\mu)}$. By hypothesis $|n|$ or $|m|$ is greater than M. Let $\mathbb{H} = \pi_1(\mathscr{C})$, let T be $T(t, \mathscr{C})$, let ϕ be the identity, let $p = g$, and let L denote the axis of g. We can choose the basepoint t_0 of T so that $t_0 \in L$, which implies that the geodesics $[t_0, \phi(\rho^\lambda_{\sigma(\mu)}) \cdot t_0]$ and $[t_0, \phi(\tau_\mu) \cdot t_0]$ consist of a single L-periodic segment. By Theorem 5.14 there exists an endomorphism ϕ^* such that $[t_0, \phi^*(\rho^\lambda_{\sigma(\mu)}) \cdot t_0]$ and $[t_0, \phi^*(\tau_\mu) \cdot t_0]$ contain a single non-trivial L-periodic segment. This implies the hyperbolicity of $\phi^*(\tau_\mu)$ and $\phi^*(\rho^\lambda_{\sigma(\mu)})$; thus the restriction $\phi^*|_{\langle g \rangle}$ is injective. On the other hand, the bound on L periodicity implies that $\phi^*(\tau_\mu) = g^{n_0}$ and $\phi^*(\rho^\lambda_{\sigma(\mu)}) = g^{m_0}$ with $|n_0|, |m_0| \leq M$. Now we must have $\phi^*(g) = g^r$, which implies that $|n_0| = |r||n|$ and $|m_0| = |r||m|$, which contradicts the assumption that $|n|, |m| > M$. □

5.3. **Merging entangled pairs.** Let $\mathscr{C}$ be a measured band complex and let $t \in \mathbf{tracks}_\kappa(\mathscr{C})$. Suppose we have an overlapping pair $(\mu, \overline{\mu})$ entangled with $(\lambda, \overline{\lambda}) \subset \sigma(\mu)$. Then it will sometimes be possible to merge the bands $\mathbb{B}(\lambda)$ and $\mathbb{B}(\mu)$ into a new band $\mathbb{B}(\eta)$.

Simply attaching a new band with overlapping bases to $\mathscr{C}$ inside the segment $\sigma(\mu)$ will add a cyclic free factor to $\pi_1(\mathscr{C})$. We must therefore also attach a 2-cell to encode that τ_η is a root of τ_μ and $\rho^\lambda_{\sigma(\mu)}$. In order to do so we may first have to widen $\mathbb{B}(\lambda)$.

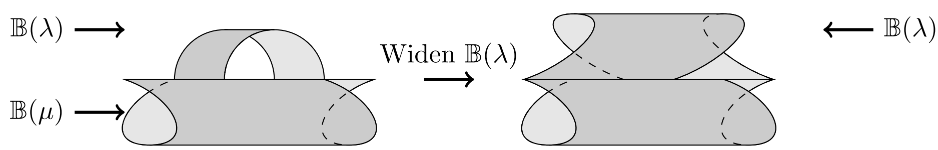

FIGURE 15. We widen the band $\mathbb{B}(\lambda)$ so that it becomes co-initial and co-terminal with $\sigma(\mu)$.

Consider the operation of widening a band $\mathbb{B}(\lambda)$ illustrated in Figure 15. We do this so that the resulting base λ is co-initial with $\sigma(\mu)$ and $\overline{\lambda}$ is co-final. The inverse of a widening is a deformation retraction, so it preserves $\pi_1(\mathscr{C})$. Furthermore the element $\rho^\lambda_{\sigma(\mu)}$ of the fundamental group is unchanged. If $[\rho^\lambda_{\sigma(\mu)}, \tau_\mu] = 1$, then for any track t efficiently carried by $\mathscr{C}$ both $\rho^\lambda_{\sigma(\mu)}$ and τ_μ have the same axis by Lemma 5.12. Figure 16 depicts what happens when we pass to the universal cover. Because $\rho^\lambda_{\sigma(\mu)}$ has the same axis as τ_μ there is a natural way to extend the pattern in $\widetilde{\mathscr{C}}$ to a pattern of $\widetilde{\mathscr{C}'}$ so that the resulting dual trees $T(t, \mathscr{C}), T(t', \mathscr{C}')$ are equivariantly isomorphic. No 2-cells were added so t' is efficiently carried. We summarize in the following lemma.

Lemma 5.16. *Let μ and λ be as above and let $t \in \mathbf{tracks}_\kappa(\mathscr{C})$. We can widen $\mathbb{B}(\lambda)$ so that, in the new band complex $\mathscr{C}'$, λ is co-initial with $\sigma(\mu)$ and $\overline{\lambda}$ is co-final with $\sigma(\mu)$. Furthermore the dual trees $T(t, \mathscr{C})$ and $T(t', \mathscr{C}')$ are equivariantly isomorphic. In particular* $\mathrm{tr}(\lambda), \mathrm{tr}(\mu)$ *remain invariant, and $t' \in \mathbf{tracks}_\kappa(\mathscr{C}')$.*

Lemma 5.17. *Let $(\mu, \overline{\mu})$ be an overlapping pair entangled with $(\lambda, \overline{\lambda})$ and suppose furthermore that $(\lambda, \overline{\lambda})$ is itself overlapping. Then after widening $\mathbb{B}(\lambda)$ as in Lemma 5.16 we have $|\lambda|_t \geq \mathrm{tr}(\mu)$.*

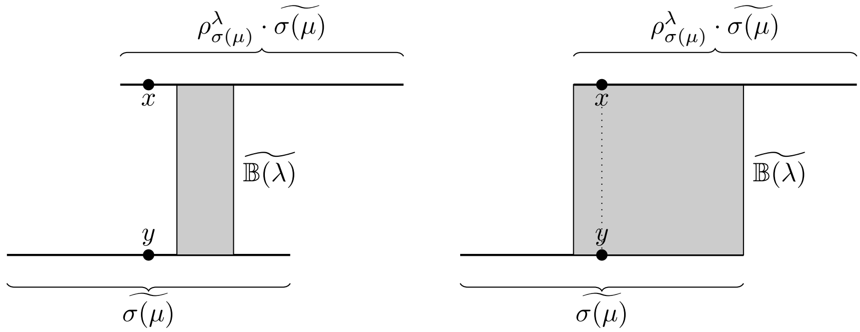

FIGURE 16. The widening move corresponds to equivariantly widening every lift $\widetilde{\mathbb{B}(\lambda)}$ of $\mathbb{B}(\lambda)$ in the universal cover. If x, y lie in the same track $\widetilde{t}$ in $\widetilde{\mathscr{C}}$, then in $\widetilde{\mathscr{C}'}$ we extend the track to pass through the widened $\widetilde{\mathbb{B}(\lambda)}$.

Proof. If λ is co-initial and co-final with $\sigma(\mu)$ but $|\lambda|_t < \mathrm{tr}(\mu)$, then λ can't overlap with its dual. $\square$

Before continuing we need the following fact, which follows by meditating on the Euclidean algorithm.

Lemma 5.18. *Let n, m be positive integers and let $d = \gcd(n, m)$. Then without loss of generality there are integers $u, v \in \mathbb{Z}_{\geq 0}$ such that $d = un - vm$. Moreover we have non-decreasing sequences of integers $0 = v_0 \leq v_1 \leq \cdots \leq v_{u+v} = v$ and $0 = u_0 \leq u_1 \leq \cdots \leq u_{u+v} = u$ with*

$$u_i + v_i + 1 = u_{i+1} + v_{i+1}$$

such that the following inequalities hold:

$$0 \leq u_i n + v_i m \leq m + n. \tag{8}$$

This fact motivates the following observation.

Lemma 5.19. *Let $(\mu, \overline{\mu})$ be an overlapping pair entangled with $(\lambda, \overline{\lambda})$. If $|\lambda|_t \geq \mathrm{tr}(\mu)$, then*

$$|\sigma(\mu)| \geq \mathrm{tr}(\mu) + \mathrm{tr}_{\sigma(\mu)}(\lambda).$$

Proof. By hypothesis $|\lambda|_t \geq \mathrm{tr}(\mu)$ and

$$\lambda \cup \overline{\lambda} \subset \sigma(\mu) \Rightarrow |\lambda|_t + \mathrm{tr}_{\sigma(\mu)}(\lambda)| \leq |\sigma(\mu)|,$$

which give the required inequality. $\square$

Suppose we are in the situation of Lemma 5.19 and that we have widened $\mathbb{B}(\lambda)$ as in Lemma 5.16. We will illustrate the attachment of $\mathbb{B}(\eta)$ with a concrete example. Suppose that $\mathrm{tr}(\mu) = 3$, $\mathrm{tr}_{\sigma(\mu)}(\lambda) = 7$, $|\mu|_t = 8$, and $|\lambda|_t = 4$; see Figure 17. $\gcd(3, 7) = 1$, we pick the linear combination $5 * 3 - 2 * 7 = 1$, which we rewrite as the series $3 + 3 + 3 - 7 + 3 + 3 - 7 = 1$, and we do this because each initial subsum is positive and at most $10 = 3 + 7$ as in Lemma 5.18 (8). This lemma implies that such a series can be found for any pair of bands satisfying our hypotheses. Now $\alpha_{\mu\lambda}$ in Figure 17 is a simple path that is a concatenation of vertical sets prescribed

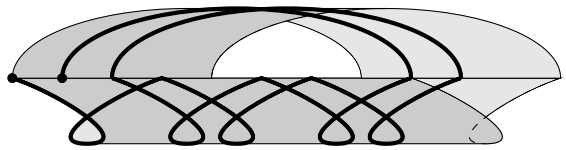

FIGURE 17. The vertical path $\alpha_{\mu\lambda}$ contained in bands $\mathbb{B}(\lambda), \mathbb{B}(\mu)$ corresponding to the series $3+3+3-7+3+3-7=1=\gcd(3,7)$.

by the associated series. Its endpoints in $\sigma(\mu)$ have distance $\gcd(\mathrm{tr}(\mu), \mathrm{tr}_{\sigma(\mu)}(\lambda))$, which in our example is 1. By Lemmas 5.18 and 5.19, we can always find such an $\alpha_{\mu\lambda}$ for any pair of bands that satisfies the hypotheses. This is why we needed to widen $\mathbb{B}(\lambda)$ in the first place; if it were too narrow we wouldn't be able to construct $\alpha_{\mu\lambda}$.

In $\pi_1(\mathscr{C})$ the elements τ_μ and $\rho^\lambda_{\sigma(\mu)}$ commute, so the product implied by the path $\alpha_{\mu\lambda}$ is

$$(\tau_\mu)^3(\rho^\lambda_{\sigma(\mu)})^{-1}(\tau_\mu)^2(\rho^\lambda_{\sigma(\mu)})^{-1} = (\tau_\mu)^5(\rho^\lambda_{\sigma(\mu)})^{-2};$$

$(\tau_\mu)^5(\rho^\lambda_{\sigma(\mu)})^{-2}$ is an element that translates $\mathrm{axis}(\tau_\mu) = \mathrm{axis}(\rho^\lambda_{\sigma(\mu)})$ by a distance of $\gcd(\mathrm{tr}(\mu), \mathrm{tr}_{\sigma(\mu)}(\lambda))$. Noting that (signed) translation length gives an embedding from the stabilizer of a bi-infinite line in a κ-acylindrical tree to $\mathbb{Z}$, we conclude that $g = (\tau_\mu)^5(\rho^\lambda_{\sigma(\mu)})^{-2}$ is the element such that $\langle g\rangle = \langle \tau_\mu, \rho^\lambda_{\sigma(\mu)}\rangle$ that will (up to sign ± 1) be produced by the algorithm given in Proposition 5.15.

We now attach a new band $\mathbb{B}(\eta)$ so that $\sigma(\eta) = \sigma(\mu)$ and

$$\mathrm{tr}(\eta) = \gcd(\mathrm{tr}(\mu), \mathrm{tr}_{\sigma(\mu)}(\lambda)).$$

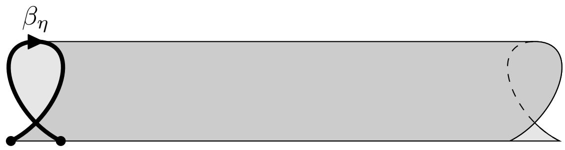

FIGURE 18. Adding $\mathbb{B}(\eta)$ and the arc β_η.

In our example this forces $|\eta|_t = 10$. Let β_η be the path shown in Figure 18. Simply attaching $\mathbb{B}(\eta)$ to $\mathscr{C}$ gives the fundamental group $\pi_1(\mathscr{C}) * \langle \tau_\eta\rangle$. We also attach a 2-cell B along the simple closed path $\alpha_{\mu\lambda} * \beta_\eta$ to get the resulting fundamental group

$$(\pi_1(\mathscr{C}) * \langle \tau_\eta\rangle) / \langle\langle \tau_\eta = g\rangle\rangle \approx \pi_1(\mathscr{C})$$

by defining $\tau_\eta = g$. First note that the new band complex $\mathscr{C}'$ can be seen as containing $\mathscr{C}$ and that the track $t \subset \mathscr{C} \subset \mathscr{C}'$ naturally extends to a track $t' \subset \mathscr{C}'$ and that the trees $T(t, \mathscr{C}) \approx T(t', \mathscr{C}')$ are equivariantly isomorphic. Further note that by the way the attaching map $\alpha_{\mu\lambda} * \beta_\eta$ is defined, $\mathscr{C}'$ carries t' efficiently.

The next step is to zip the bands $\mathbb{B}(\mu)$ and $\mathbb{B}(\lambda)$ onto $\mathbb{B}(\eta)$. In our example, since $\tau_\mu = (\tau_\eta)^3$, $\mathbb{B}(\mu)$ should "wrap" three times around $\mathbb{B}(\eta)$. To accomplish this we *horizontally* subdivide $\mathbb{B}(\mu)$ into 3 bands and successively perform three zipping operations (Definitions 3.1 and 3.2). This is best visualized in the universal cover; see Figure 19. Thus, we have produced a new band complex $\mathscr{C}'$ efficiently carrying

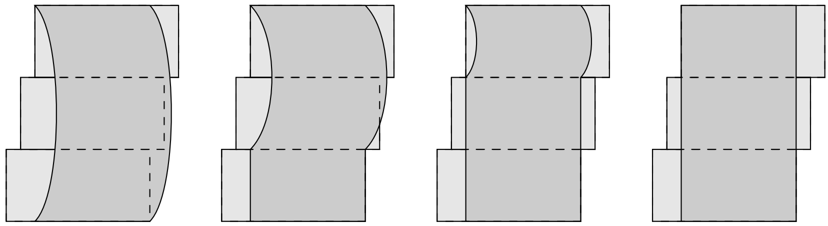

FIGURE 19. Zipping $\mathbb{B}(\mu)$ onto $\mathbb{B}(\eta)$, as seen from the universal cover. Here $\tau_\mu = (\tau_\eta)^3$.

a track t', and by Proposition 3.11 we have preserved the fundamental group and the dual Bass-Serre tree. Although we considered a specific example, this discussion is sufficiently general to make the following claim.

Proposition 5.20 (The periodic merger)**.** *Let $\mathscr{C}$ be a band complex with an overlapping pair $(\mu, \overline{\mu})$ and an overlapping pair $(\lambda, \overline{\lambda}) \subset \sigma(\mu)$ such that $(\mu, \overline{\mu})$ and $(\lambda, \overline{\lambda})$ are entangled. Suppose that $t \in$ **tracks**$_\kappa(\mathscr{C})$. If, after widening $\mathbb{B}(\lambda)$, as in Lemma 5.16, we have $|\lambda|_t \geq \mathrm{tr}(\mu)$, then:*

- *There is a continuous map $m : \mathscr{C} \to \mathscr{C}'$ where $\mathscr{C}'$ with $m(t) = t' \subset \mathscr{C}'$ where t' is a track efficiently carried by $\mathscr{C}'$.*
- *The induced map $m_\sharp : \pi_1(\mathscr{C}) \to \pi_1(\mathscr{C}')$ is an isomorphism, and there is an $m_\sharp$-equivariant isomorphism of dual Bass-Serre trees $T(t, \mathscr{C}) \to T(t', \mathscr{C}')$.*
- *The resulting band complex $\mathscr{C}'$ has two fewer bases.*

Proof. $\mathscr{C}'$ is obtained by first perhaps widening $\mathbb{B}(\lambda)$ so that it is co-initial and co-terminal with $\sigma(\mu)$, then attaching a band $\mathbb{B}(\eta)$ so that $\sigma(\eta) = \sigma(\mu)$, attaching a 2-cell, and finally horizontally subdividing and zipping $\mathbb{B}(\mu)$ and $\mathbb{B}(\lambda)$ onto $\mathbb{B}(\eta)$. The resulting composition of operations preserves fundamental groups and dual trees. □

5.4. **A modification to $\mathfrak{T}(C)$: Adding periodic mergers to the elimination process.** We now turn our attention to band complexes, viewed as combinatorial objects.

Definition 5.21. Let $\mathscr{C}_v$ be a band complex in $\mathfrak{T}(C)$. Then we say that $\mathscr{C}_v$ is *merging inadmissible* if it contains an overlapping pair $(\mu, \overline{\mu})$ that is entangled with $(\lambda, \overline{\lambda})$, but the algorithm of Proposition 5.15 certifies that $\mathscr{C}_v$ cannot efficiently carry a κ-track.

Merging inadmissibility can be verified algorithmically. Indeed, given $(\lambda, \overline{\lambda}) \subset \sigma(\mu)$, Lemma 5.11 states that entanglement can be computed and merging inadmissibility is certified from the output of the algorithm of Proposition 5.15. By definition, this only depends on the underlying band complex and not on the track it carries.

Whether entangled base pairs can actually be merged and the outcome of this operation depend on the track t carried by $\mathscr{C}$. Merging inadmissibility, however, guarantees that no matter the track being carried by $\mathscr{C}$, such a merging is impossible.

Now we have proved that if $(\mu, \overline{\mu})$, $(\lambda, \overline{\lambda})$ are a pair of entangled overlapping pairs with $(\lambda, \overline{\lambda}) \subset \sigma(\mu)$, then for any $t \in$ **tracks**$_\kappa(\mathscr{C}_v)$, after widening $\mathbb{B}(\lambda)$ as in Lemma

5.16, we will be able to apply a periodic merger by Lemma 5.19. It therefore follows that we can discard merging inadmissible band complexes since they cannot carry κ-tracks.

This next lemma simply follows from the fact that we can enumerate the combinatorial outcomes of all such periodic mergers since we widened some band by a controlled amount, added a band and a 2-cell with an attaching map of length M, and applied the N zipping operation, where M, N are bounded by the output of the algorithm of Proposition 5.15.

Lemma 5.22. *Let $\mathscr{C}$ be a band complex containing overlapping pairs $(\mu,\overline{\mu}),(\lambda,\overline{\lambda})$ that are entangled and such that $(\lambda,\overline{\lambda}) \subset \sigma(\mu)$. Then we can effectively construct a finite set of band complexes*

$$\begin{array}{ccc} & \mathscr{C} & \\ \swarrow & & \searrow \\ \mathscr{C}_1 & \cdots & \mathscr{C}_{m_{\mathscr{C}}} \end{array}$$

containing all possible outcomes $m : \mathscr{C} \to \mathscr{C}'$ of merging $\mathbb{B}(\mu)$ and $\mathbb{B}(\lambda)$ as described by Proposition 5.20*, with the track t ranging over* $\mathbf{tracks}_\kappa(\mathscr{C})$*.*

We note that although some of the band complexes produced by Lemma 5.22 may not correspond to *any* periodic mergers, the resulting band complexes will have the same fundamental group and a lower τ-complexity. In particular if any of them admit a κ-track, then so must the original $\mathscr{C}$. There is thus no danger of introducing "false positives" by giving $\mathscr{C}$ illegitimate children. We now include periodic mergers in our elimination tree.

Definition 5.23. The elimination tree $\mathfrak{T}_+(\mathscr{C}, J)$ is constructed inductively similarly as $\mathfrak{T}(\mathscr{C}, J)$ in Section 4.5 except with a new clause that takes precedence over item (5) given in Section 4.0.1.

(5m) If $\mathscr{C}_v$ is a non-terminal band complex in $\mathfrak{T}_+(\mathscr{C} J)$ that contains overlapping pairs $(\mu,\overline{\mu})$ and $(\lambda,\overline{\lambda})$ satisfying the premises of Lemma 5.22, then define as its children the collection of band complexes given by Lemma 5.22.

Obviously all the κ-inadmissibility and repetition inadmissibility criteria on paths $\mathfrak{T}(\mathscr{C}, J)$ also apply to $\mathfrak{T}_+(\mathscr{C}, J)$, as does the classification in Theorem 4.21. For the remainder of this paper we will use $\mathfrak{T}_+(\mathscr{C}, J)$ as our elimination tree.

6. Overlapping pairs must occur and stabilize

For this section let $\mathscr{C}$ be a band complex, and let $t \subset \mathscr{C}$ be an automorphically minimal κ-track efficiently carried by $\mathscr{C}$. The corresponding Rips sequence gives a path in $\mathfrak{T}_+(\mathscr{C}, J)$. Throughout this section we will fix a superquadratic subpath

$$p(t_1) : (\mathscr{C}_1, t_1) \to \cdots \to (\mathscr{C}_P, t_P) \tag{9}$$

where, in particular, the track t_1 is an automorphically minimal κ-track efficiently carried by $\mathscr{C}_1$. We will further assume that the J-relative τ-complexity remains constant throughout p. The purpose of this section is to show that if p is sufficiently long, then it must have a tail in which some base μ is repeatedly the carrier. This is called μ-periodicity (Definition 6.14).

We prove this by first defining a quantity called the excess that remains constant throughout p. We then consider the quadratic part of $\mathscr{C}_1$ and use this to show that a union of bases called the participating segments has a length bounded above by some computable multiple of the excess. This multiple is computed using a restricted elimination process. It will be clear from the definition that the excess is bounded above by some constant multiple of the length of the longest base in each $\mathscr{C}_i$ occurring in p. Finally we will define something called a C-T cycle with the property that whenever it occurs, a considerable portion of the participating segments gets cut out. Our bound on the total length of the participating segments will prevent these cycles from occurring too often. The critical detail is that this bound does not depend on the actual track t. It only depends on the sequence of underlying band complexes that occur along the path p and the assumption that t is an automorphically minimal κ-track efficiently carried by $\mathscr{C}$.

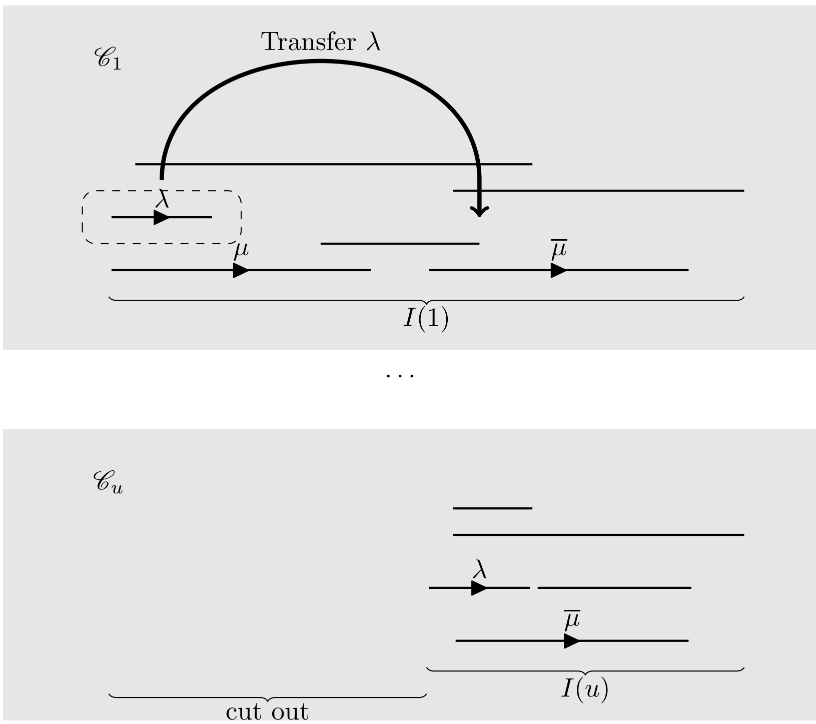

FIGURE 20. In $\mathscr{C}_1$, μ is a carrier and λ gets transferred. In $\mathscr{C}_u$, λ is again leading. At least $|\mu|_{t_1}$ got cut out of $I(1)$.

We will give an example of what is meant by "cut out". Suppose that the carrier μ in $\mathscr{C}_1$ does not overlap with its dual, that some base λ gets transferred, and that eventually in $\mathscr{C}_u$ λ is again a leading base; see Figure 20. Let the interval $I(u)$ be the natural image of $I(1)$ in $\mathscr{C}_u$. On one hand we have $|I(u)|_{t_u} < |I(1)|_{t_1}$, on the other hand we have $|I(1)|_{t_1} - |I(u)|_{t_u} \geq |\mu|_{t_1}$; i.e., we cut out at least $|\mu|_{t_1}$ from $I(1)$.

6.1. **The excess invariant.** The excess invariant given in Definition 6.2 is originally due to Makanin [Mak82]. It also occurs in [BF95, KM98].

Definition 6.1. Let p be the path given in (9). By $C(p)$ we denote *the set of bases that are carriers at some point along* p, and by $T(p)$ we denote *the set of bases that* *are transferred at some point along* p. We denote the *participating segments*

$$\sigma(p) = \bigcup_{\mu \in C(p) \cup T(p)} (\mu \cup \overline{\mu}) .$$

Recall that we are reusing symbols, so that if λ denotes a base in $\mathscr{C}_1$, then it also naturally denotes a base in every $\mathscr{C}_u$ that occurs along p. With this in mind it makes sense to reuse notation to denote the corresponding subset $\sigma(p)$ of each $\mathscr{C}_u$ in p the same way. If there is a danger of confusion we will explicitly write $\sigma(p) \subset \mathscr{C}_u$. Noting that entire transformations are compositions of zipping moves and collapses (cf. Section 3.1), the continuous map $\mathscr{C}_u \to \mathscr{C}_{u'}$ actually induces a *retraction* of $\sigma(p) \subset \mathscr{C}_u$ onto $\sigma(p) \subset \mathscr{C}_{u'}$. It follows that $\sigma(p) \subset \mathscr{C}_{u'}$ naturally embeds into $\sigma(p) \subset \mathscr{C}_u$. Along the path $p(t_1)$ given in (9) we will write $|\sigma(p)|_{t_u}$ to denote the hitting measure of $\sigma(p) \subset \mathscr{C}_u$ with respect to t_u. By the embedding above we have a chain of proper inequalities

$$|\sigma(p)|_{t_1} > \cdots > |\sigma(p)|_{t_P} > 0. \tag{10}$$

Definition 6.2. Suppose now a subset σ is a union of bases and a track t is carried by $\mathscr{C}$. We call the following quantity *excess*:

$$\psi_t(\sigma) = \left(\sum_{\mu \subset \sigma} |\mu|_t \right) - 2|\sigma|_t.$$

Excess measures how far the σ is from being quadratic as quantified by the hitting measure. From the definition of the entire transformation, a straightforward counting argument gives the following.

Lemma 6.3 (Excess is invariant; cf. [BF95, (7.6.1)])**.** *Let* $p(t_1)$ *be the Rips process given in* (9)*; i.e., all moves are entire transformations, and* J*-relative* τ*-complexity is constant. Then we have equalities*

$$\psi_{t_1}(\sigma(p)) = \cdots = \psi_{t_w}(\sigma(p)) .$$

6.2. **Bounding the quadratic part.** We can decompose $\sigma(p) \subset \mathscr{C}_1$ into

$$\sigma(p) = Q(p) \cup SQ(p) \subset \mathscr{C}_1,$$

where $Q(p)$, the *quadratic part of* $\sigma(p)$, is the closure of the set of points that are contained in exactly two bases. We define $SQ(p)$, the *superquadratic part of* $\sigma(p)$, to be the closure of $\sigma(p) \setminus Q(p)$ (all points in $SQ(p)$ lie in at most three bases). It is worth noting that this decomposition is almost disjoint (intersection consists of finitely many points) and that some bases may lie partially in $Q(p)$ and partially in $SQ(p)$.

If we forget the tracks t_i carried by $\mathscr{C}_i$ in $p(t_1)$, then we have an induced superquadratic path

$$p : \mathscr{C}_1 \to \cdots \to \mathscr{C}_P \tag{11}$$

with constant J-relative τ complexity that lies in $\mathfrak{T}(\mathscr{C}, J)$. Note that we are not requiring $\mathscr{C}_p$ to be terminal.

Lemma 6.4. *We can construct a computable function* f_{rep}*, depending only on the band complex* $\mathscr{C}$ *and a path* p *as in* (11)*, with positive integer values such that for any automorphically minimal track* t *efficiently carried by* $\mathscr{C}_1$ *we have the inequality*

$$|Q(p)|_t \leq f_{\mathrm{rep}}(\mathscr{C}_v, p) \cdot |SQ(p)|_t. \tag{12}$$

Proof. Subdivide each band (Definition 3.6) so that each base either lies entirely in $Q(p)$ or intersects $Q(p)$ with empty interior. $Q(p)$ is now a union of maximal sections (Definition 2.7). Let J_Q be the complement of $Q(p)$. Consider the J_Q-restricted admissible elimination tree $\mathfrak{A}(\mathscr{C}_1, J_Q)$. Since $\sigma(p)$ is a union of bases that is closed under taking duals, the leaves of $\mathfrak{A}(\mathscr{C}_1, J_Q)$ will be band complexes in which every base in $Q(p)$ is eventually moved onto $SQ(p)$. On the other hand, since $\mathfrak{A}(\mathscr{C}_1, J_Q)$ is J_Q-relatively quadratic, by Corollary 4.27 the subtree $\mathfrak{A}(\mathscr{C}_1, J_Q)$ is finite and therefore algorithmically constructible.

Going backwards in $\mathfrak{A}(\mathscr{C}_1, J_Q)$ from every admissible leaf to the root $\mathscr{C}_1$ (recall Section 4.1) and repeatedly applying the upper bound of Lemma 4.7 give us a finite (algorithmically constructible) set of upper bounds for lengths of the bases in $Q(P)$ in $\mathscr{C}_1$ in terms of $|SQ(p)|_t$. The maximum over this set can be used to compute an upper bound of $|Q(p)|_t$ in terms of $|SQ(p)|_t$ for any automorphically minimal track t efficiently carried by $\mathscr{C}_1$. □

Whenever a base λ_u is the carrier in $(\mathscr{C}_u, t_u) \to (\mathscr{C}_{u'}, t_{u'})$ in (9), after transferring other bases, it gets shortened. The bound given by Lemma 6.4 and the invariance of excess tell us that, although bases get shorter, their lengths remain bounded below throughout p. This will force them to overlap.

Lemma 6.5. *Let $p(t_1)$ be as in (9) and let $\lambda_u \in C(p) \cup T(p)$ be the base in $\mathscr{C}_u$ such that $|\lambda_u|_{t_u}$ is maximal. The initial length $|\sigma(p)|_{t_1}$ is always bounded by*

$$|\sigma(p)|_{t_1} \leq (f_{\text{rep}}(\mathscr{C}_1, p) + 1)\, N^2 |\lambda_u|_{t_u},$$

where N is the number of bases in $\mathscr{C}_u$.

Proof. In $(\mathscr{C}_u, t_u)$, every $\mu \in C(p) \cup T(p)$ has length at most $|\lambda_u|_{t_u}$. It therefore follows that $|\sigma(p)|_{t_u} \leq N \cdot |\lambda_u|_{t_u}$. Now since there are no more than N bases, each point in $\sigma(p)$ is contained by at most N bases. The strict upper bound for the invariant excess

$$\psi_{t_1}(\sigma(p)) = \psi_{t_u}(\sigma(p)) < N^2 \cdot |\lambda_u|_{t_u},$$

as well as the bound $|SQ(p)|_{t_1} \leq \psi_{t_1}(\sigma(p))$, combine with Lemma 6.4 to give the desired inequality since $|\sigma(p)|_{t_1} = |Q(p)|_{t_1} + |SQ(p)|_{t_1}$. □

6.3. **Orientation reversing overlaps: Two lemmas.** The proof in the next section requires two additional lemmas.

Lemma 6.6. *Suppose $\pi_1(\mathscr{C})$ has no elements of order 2 and $t \in$ $\mathbf{tracks}_\kappa(\mathscr{C})$. Then for any dual pair $(\mu, \overline{\mu})$ such that $\mu \cap \overline{\mu} \neq \emptyset$, that is, orientation reversing, we have*

$$|\mu \cap \overline{\mu}|_t < \kappa + 1.$$

Proof. Since $\mu \cap \overline{\mu} \neq \emptyset$ and $(\mu, \overline{\mu})$ is orientation reversing, the image of $\mathbb{B}(\mu)$ in $\mathscr{C}$ contains a Möbius band. Let the simple closed curve γ be homotopic to the core of this band. We may view γ as an element of $\pi_1(\mathscr{C})$. If $\mathscr{C}$ efficiently carries a track t, then γ must invert some segment in $T(t, \mathscr{C})$; thus $\gamma \neq 1$.

By assumption $\gamma^2 \neq 1$, but an analysis of the action of γ^2 on $\widetilde{\mathscr{C}}$ and $T(t, \mathscr{C})$ similar to the one shown in Figure 10 shows that γ^2 fixes an arc of length $|\mu \cap \overline{\mu}|_t$ in $T(t, \mathscr{C})$. κ-acylindricity therefore ensures the required bound. □

Lemma 6.7. *Let $(\mu, \overline{\mu})$ be an overlapping pair and suppose that $(\lambda, \overline{\lambda})$ is orientation reversing and $\lambda \cup \overline{\lambda} \subset \sigma(\mu)$. If $\pi_1(\mathscr{C})$ has no 2-torsion, then for any $t \in \mathbf{tracks}_\kappa(\mathscr{C})$*

$$|\lambda|_t < \mathrm{tr}(\mu) + \kappa + 1.$$

Proof. Suppose towards a contradiction that $|\lambda|_t \geq \mathrm{tr}(\mu) + \kappa + 1$. Assume that λ is to the left of $\overline{\lambda}$ and let q be the leftmost point of λ. By successively transferring $\overline{\lambda}$ through $\mathbb{B}(\mu)$ we can arrange it so that the leftmost point of $\overline{\lambda}$ is moved to a distance of less than $\mathrm{tr}(\mu)$ to the right of q. It therefore follows that $|\overline{\lambda} \cap \lambda|_t \geq \kappa + 1$, which contradicts Lemma 6.6. □

6.4. **Cutting off too much.** We now introduce C-T cycles which are guaranteed to "shorten" $\sigma(p) \subset \mathscr{C}_v$ by some fixed amount each time they occur. The idea of a C-T cycle is inspired from Case 3 in the proof of [AHT06, Theorem 12].

Definition 6.8. Let p be as in (9). A *C-T cycle* is a subpath $\mathscr{C}_j \to \cdots \to \mathscr{C}_k$ such that for each $\lambda \in C(p)$ there is some $j \leq l \leq k$ such that λ is the carrier in some $\mathscr{C}_l$ and one of the following occurs:

- If λ either doesn't overlap with its dual or is orientation reversing, then for some $l < l' \leq k$, some base δ that was transferred by λ in $\mathscr{C}_l \to \mathscr{C}_{l+1}$ is a leading base again in $\mathscr{C}_{l'}$.
- If $(\lambda, \overline{\lambda})$ form an overlapping pair, then there are $l < l' < l'' \leq k$ such that λ gets carried in $\mathscr{C}_{l'}$ (i.e., it ceases to be a carrier) and is a leading base again in $\mathscr{C}_{l''}$. Also there is some base δ that was transferred in $\mathscr{C}_l \to \mathscr{C}_{l+1}$ that is again a leading base in $\mathscr{C}_{l'''}$ for some $l < l''' \leq k$.

Lemma 6.9. *Let $p(t_1)$ be a path as in (9) and let $\mathscr{C}_j \to \cdots \to \mathscr{C}_k$ be a C-T cycle in $p(t_1)$. Then either every base in $C(p) \cup T(p)$ has length at most* $\max\{2\kappa, 1\}$ *or*

$$|\sigma(p)|_{t_j} - |\sigma(p)|_{t_k} \geq \frac{|\lambda|_{t_j}}{2(2+\kappa)}, \tag{13}$$

where λ is the longest base in $\mathscr{C}_j$.

Proof. First note that until λ is the carrier in some $\mathscr{C}_l$ where $j \leq l \leq k$ we have

$$|\lambda|_{t_j} = |\lambda|_{t_j+1} = \cdots = |\lambda|_{t_l}.$$

By definition of a C-T cycle, there is such a $j \leq l \leq k$ where λ is the carrier in $\mathscr{C}_l$. Assume that $|\lambda|_{t_j} > \max\{2\kappa, 1\}$. We will now show (13). The proof divides into cases; decimals denote subcases.

Case 1: *λ does not overlap with its dual.* In this case let δ be some base that is carried by λ. Since it is moved by at least $|\lambda|_{t_j}$ to the right, by definition of a C-T cycle there are some $l < l' \leq k$ where δ is a leading base in $\mathscr{C}_{l'}$ again, which means that $|\sigma(p)|_{t_l} - |\sigma(p)|_{t'_l} \geq |\lambda|_{t_j}$ so the result holds. (See Figure 20.)

Case 2: *λ has non-trivial intersection with $\overline{\lambda}$.*

Case 2.1: *$(\lambda, \overline{\lambda})$ is orientation reversing.* Let δ be some base that is carried by λ. Then δ is moved $2|\lambda|_{t_l} - |\lambda \cap \overline{\lambda}|_{t_l} - |\delta|_{t_l}$ to the right. On one hand $|\delta|_{t_l} \leq |\lambda|_{t_l}$; on the other hand by Lemma 6.6 $|\lambda \cap \overline{\lambda}|_{t_l} \leq \kappa$. We therefore conclude that δ is moved at least $|\lambda|_{t_l} - \kappa$ to the right. When δ is again a leading base in $\mathscr{C}_{l'}, l < l' \leq k$, then

$$|\sigma(p)|_{t_l} - |\sigma(p)|_{t_{l'}} \geq |\lambda|_{t_l} - \kappa - 1 \geq \frac{|\lambda|_{t_l}}{2},$$

so (13) holds.

Case 2.2: $(\lambda, \overline{\lambda})$ *form an overlapping pair.* There is some maximal l' such that in $\mathscr{C}_l \to \cdots \to \mathscr{C}_{l'-1}$, λ is always the carrier base. Note moreover that $\mathrm{tr}(\lambda)$ remains constant.

Case 2.2.1: $|\lambda|_{t_l} \leq (2\kappa+2)\mathrm{tr}(\lambda)$. Let δ be some base that is carried by λ. Then it is moved to the right by $\mathrm{tr}(\lambda)$, and by hypothesis there is some $l < l''' \leq k$ such that δ is leading again; hence

$$|\sigma(p)|_{t_l} - |\sigma(p)|_{t_{l'}} \geq \frac{|\lambda|_{t_l}}{2\kappa+2},$$

and (13) holds.

Case 2.2.2: $|\lambda|_{t_l} > (2\kappa+2)\mathrm{tr}(\lambda)$.

Case 2.2.2.1: $|\sigma(p)|_{t_l} - |\sigma(p)|_{t_{l'}} \geq \dfrac{|\lambda|_{t_l}}{2(2\kappa+2)}$. (13) immediately holds.

Case 2.2.2.2: $|\sigma(p)|_{t_l} - |\sigma(p)|_{t_{l'}} < \dfrac{|\lambda|_{t_l}}{2(2\kappa+2)}$. In particular, in $\mathscr{C}_{l'}$,

$$|\lambda|_{t_{l'}} > |\lambda|_{t_l}\left(1 - \frac{1}{2(2\kappa+2)}\right) = |\lambda|_{t_l}\left(\frac{4\kappa+3}{2(2\kappa+2)}\right) > (2\kappa+3/2)\mathrm{tr}(\lambda) \tag{14}$$

by the Case 2.2.2 assumption on $|\lambda|_{t_l}$.

Case 2.2.2.2.1: *The carrier η in $\mathscr{C}_{l'}$ does not overlap with its dual.* In this case λ is moved to the right by at least $|\lambda|_{t_{l'}}$, and we note that for all $\kappa \geq 0$ we have

$$1 - \frac{1}{2(2\kappa+2)} > 1/2 > \frac{1}{2(2\kappa+2)};$$

thus when λ is a leading base again in $\mathscr{C}_{l''}$, by (14) we have

$$|\sigma(p)|_{t_{l'}} - |\sigma(p)|_{t_{l''}} > \frac{|\lambda|_{t_l}}{2},$$

so (13) holds.

Case 2.2.2.2.2: *The carrier η in $\mathscr{C}_{l'}$ has non-trivial intersection with $\overline{\eta}$ and is orientation reversing.* Note that

$$|\eta|_{t_{l'}} > |\lambda|_{t_{l'}} \geq |\lambda|_{t_l}\left(1 - \frac{1}{2(2\kappa+2)}\right).$$

As in Case 2.1 we deduce that η carries λ more than $|\lambda|_{t_{l'}} - \kappa$ to the right. So in $\mathscr{C}_{l''}$ when λ is leading again at least $|\lambda|_{t_{l'}} - \kappa$ was cut from $\sigma(p)$. Suppose towards a contradiction that

$$|\lambda|_{t_{l'}} - \kappa < \frac{|\lambda|_{t_l}}{2(2\kappa+2)}.$$

Then since we are in Case 2.2.2.2 we have

$$|\lambda|_{t_l}\left(1-\frac{1}{2(2\kappa+2)}\right)-\kappa<|\lambda|_{t_{l'}}-\kappa.$$

Combining these gives

$$\begin{aligned} & |\lambda|_{t_l}\left(1-\frac{1}{2(2\kappa+2)}\right)-\kappa && <\frac{|\lambda|_{t_l}}{2(2\kappa+2)} \\ \Rightarrow\quad & |\lambda|_{t_l}\left(1-\frac{2}{2(2\kappa+2)}\right) && <\kappa \\ \Rightarrow\quad & |\lambda|_{t_l}\left(\frac{4\kappa+2}{4\kappa+4}\right) && <\kappa \\ \Rightarrow\quad & |\lambda|_{t_l}<\left(1+\frac{2}{(4\kappa+4)}\right)\kappa && <\frac{3}{2}\kappa, \end{aligned}$$

which contradicts our assumption that $|\lambda|_{t_l}>2\kappa$.

Case 2.2.2.2.3: $(\eta,\overline{\eta})$ *form an overlapping pair where* η *in* $\mathscr{C}_{l'}$ *is the carrier.* We may assume that after repeatedly getting shortened in $\mathscr{C}_l\to\cdots\to\mathscr{C}_{l'-1}$, λ is still long enough in $\mathscr{C}_{l'}$ for $(\lambda,\overline{\lambda})$ to be an overlapping pair. Indeed suppose this was not the case. By not overlapping we have $|\lambda|_{t_{l'}}<\mathrm{tr}(\lambda)$, and substituting into the Case 2.2.2 assumption yields

$$|\lambda|_{t_{l'}}<\frac{|\lambda|_{t_l}}{(2\kappa+2)}.$$

This contradicts (14) since

$$1-\frac{1}{2(2\kappa+2)}>\frac{1}{2\kappa+2}.$$

Case 2.2.2.2.3.1: (*Recall Definition* 5.5) *In* $\mathscr{C}_{l'}$, $\sigma(\eta)\subset\sigma(\lambda)$. This means, since we are in Case 2.2.2.2, that

$$|\eta|_{t_{l'}}>(2\kappa+3/2)\mathrm{tr}(\lambda)>\mathrm{tr}(\lambda)+\kappa+1,$$

so by Lemma 5.13 τ_λ and $\rho^{\eta}_{\sigma(\lambda)}$ commute, so the overlapping pairs $(\mu,\overline{\mu})$ and $(\lambda,\overline{\lambda})$ are entangled. By Lemma 5.17 we can apply the periodic merger of Proposition 5.20 to merge the bands $\mathbb{B}(\eta)$ and $\mathbb{B}(\mu)$. Since we are working in the elimination tree $\mathfrak{T}_+(C)$, by (5m) of Definition 5.23 *we must merge the bands* $\mathbb{B}(\eta)$ *and* $\mathbb{B}(\lambda)$*, which decreases the complexity.* This contradicts the assumption that the τ-complexity remains constant throughout p.

Case 2.2.2.2.3.2: *In* $\mathscr{C}_{l'}$, $\sigma(\eta)\supset\sigma(\lambda)$. We finally distinguish two subcases:

Case 2.2.2.2.3.2.1: $(\kappa+2)\mathrm{tr}(\eta)\le|\lambda|_{t_{l'}}$. Again as in Case 2.2.2.2.3.1 we can perform a periodic merger.

Case 2.2.2.2.3.2.2: $(\kappa+2)\mathrm{tr}(\eta)>|\lambda|_{t_{l'}}$. In this case λ gets moved by $\mathrm{tr}(\eta)$ to the right. λ is again a leading base in $\mathscr{C}_{l''}$; then we will have cut at least $\mathrm{tr}(\eta)$

from $\sigma(p)$ in passing from $\mathscr{C}'_l$ to $\mathscr{C}_{l''}$. On one hand, since this is a subcase of Case 2.2.2.2, we have

$$\mathrm{tr}(\eta) > \frac{|\lambda|_{t_{l'}}}{\kappa+2} > \frac{|\lambda|_{t_l}}{\kappa+2}\left(1-\frac{1}{2(2\kappa+2)}\right),$$

where the last inequality is from (14), and we can estimate

$$\begin{aligned} & \frac{2(2\kappa+2)}{\kappa+2}\frac{1}{2} && \geq 1 \\ \Rightarrow\quad & \frac{2(2\kappa+2)}{\kappa+2}\left(1-\frac{1}{2(2\kappa+2)}\right) && > 1 \\ \Leftrightarrow\quad & \frac{1}{\kappa+2}\left(1-\frac{1}{2(2\kappa+2)}\right) && > \frac{1}{2(2\kappa+2)} \\ \Leftrightarrow\quad & \frac{|\lambda|_{t_l}}{\kappa+2}\left(1-\frac{1}{2(2\kappa+2)}\right) && > \frac{|\lambda|_{t_l}}{2(2\kappa+2)} \\ \Rightarrow\quad & \mathrm{tr}(\eta) && > \frac{|\lambda|_{t_l}}{2(2\kappa+2)}. \end{aligned}$$

(13) therefore holds and, furthermore, all possibilities have been exhausted. □

We can now combine Lemmas 6.5 and 6.9.

Corollary 6.10. *Let p be a path as in* (9)*, i.e., a path induced by an automorphically minimal κ-track efficiently carried by $\mathscr{C}_1$. Then at most*

$$N^2(f_{\mathrm{rep}}(\mathscr{C}_1,p)+1)(2\kappa+2)$$

disjoint C-T cycles can occur.

Proof. By Lemmas 6.9 and 6.5, whenever a C-T cycle $\mathscr{C}_j \to \cdots \to \mathscr{C}_k$ occurs we can bound from below the difference

$$|\sigma(p)|_{t_j} - |\sigma(p)|_{t_k} \geq \frac{|\sigma(p)|_{t_1}}{N^2(f_{\mathrm{rep}}(\mathscr{C}_1,p)+1)(2\kappa+2)}.$$

Since $|\sigma(p)|_{t_1} > |\sigma(p)|_{t_2} > \cdots > 0$ the desired bound on the number of C-T cycles follows. □

6.5. **C-T-inadmissibility and a reduction to μ-periodicity.** Given a path p in a (restricted) elimination tree $\mathfrak{T}_+(\mathscr{C},J)$ we can define the sets $C(p)$ and $T(p)$ (Definition 6.1) and therefore corresponding C-T cycles (Definition 6.8).

Definition 6.11. If a path $p : \mathscr{C}_v \to \cdots \to \mathscr{C}_u$ in $\mathfrak{T}_+(\mathscr{C},J)$ contains more disjoint C-T cycles than the computable bound given by Corollary 6.10, then it is called *C-T-inadmissible.*

Definition 6.12 (Admissible)**.** A subtree of $\mathfrak{T}_+(\mathscr{C},J)$ is said to be *admissible* if it doesn't contain any leaves that are inadmissible (recall (4) in Section 4), κ-inadmissible paths (Definition 4.14), repetition inadmissible paths (Definition 4.18), or C-T inadmissible paths. We denote by $\mathfrak{A}_+(\mathscr{C},J) \subset \mathfrak{T}_+(\mathscr{C},J)$ the maximal admissible subtree and call it the *admissible elimination tree.*

Proposition 6.13. *Any path $p : \mathscr{C} \to \cdots$ in $\mathfrak{T}_+(\mathscr{C},J)$ induced by an automorphically minimal κ-track efficiently carried by $\mathscr{C}$ must also lie in $\mathfrak{A}_+(\mathscr{C},J)$.*

Proof. This follows immediately from Proposition 4.23 and Corollary 6.10. □

Definition 6.14. Let μ be the carrier base in $\mathscr{C}$. A path

$$p : \mathscr{C}_u \to \cdots$$

in $\mathfrak{T}_+(\mathscr{C}, J)$ is called *μ-periodic* if, throughout p, μ is the carrier base and $(\mu, \overline{\mu})$ forms an overlapping pair.

Proposition 6.15. *Any infinite path of $\mathfrak{A}_+(\mathscr{C}, J)$ has a tail*

$$\mathscr{C}_u \to \cdots$$

that is λ-periodic for some base λ.

Proof. Let p be some infinite path of $\mathfrak{A}_+(\mathscr{C}, J)$. We can form the sets $C(p)$ and $T(p)$, since every tail p' of p gives $C(p') \subset C(p)$ and $T(p') \subset T(p)$. Passing to a tail of p we may assume that each base in $C(p)$ is the carrier infinitely often and every base in $T(p)$ is carried infinitely often.

We may further assume that the (relative) τ-complexity remains constant, that no annulus subdivisions or Möbius moves occur, and that p is superquadratic. If $C(p)$ consists of more than one element, then infinitely many C-T cycles (Definition 6.8) occur so p is not contained in an admissible subtree.

It therefore follows that some base λ must repeatedly be the carrier throughout p. Now if λ doesn't overlap with $\overline{\lambda}$ after N (the total number of bases) entire transformations λ can no longer be a maximal leading base. It follows that $(\lambda, \overline{\lambda})$ form an overlapping pair. ☐

Therefore, if we can find a computable bound on the number of times in a row the same base λ can be a carrier base in some λ-periodic path $\mathscr{C}_v \to \cdots$ induced by a minimal κ-track efficiently carried by $\mathscr{C}_v$, then we will be able to effectively construct a finite subtree of $\mathfrak{A}_+(C)$, whose leaves give a set of tracks containing all automorphically minimal κ-tracks efficiently carried by $\mathscr{C}$.

7. Bounding the periodicity of overlapping pairs

Definition 7.1. We say that a maximal leading base λ is a *principal overlapping carrier* if $(\lambda, \overline{\lambda})$ form an overlapping pair and there are no other overlapping pairs $(\mu, \overline{\mu})$ that can be merged with $(\lambda, \overline{\lambda})$.

By Proposition 6.15, in any sufficiently long admissible branch in $\mathfrak{T}_+(C)$, the situation depicted in Figure 21 will occur. At each such entire transformation although $|\lambda|$ decreases, $\mathrm{tr}(\lambda)$ is invariant. Taking inspiration from words, where a large initial segment of a word overlaps with a terminal segment, e.g., *abcabcabcab*, we have the following:

Definition 7.2. Let $\mathscr{C}$ be a band complex efficiently carrying a track t and suppose that $(\lambda, \overline{\lambda})$ form an overlapping pair. We define the *periodicity of* λ to be the positive integer

$$\mathrm{period}_t(\lambda) = \left\lfloor \frac{|\lambda|_t}{\mathrm{tr}(\lambda)} \right\rfloor.$$

Lemma 7.3. *Let λ be a principal overlapping carrier in a band complex $\mathscr{C}$ efficiently carrying a track t with N bases. Then in the Rips process induced by t starting at $(\mathscr{C}, t)$, λ can be the maximal leading base at most $N\mathrm{period}_t(\lambda)$ times in a row.*

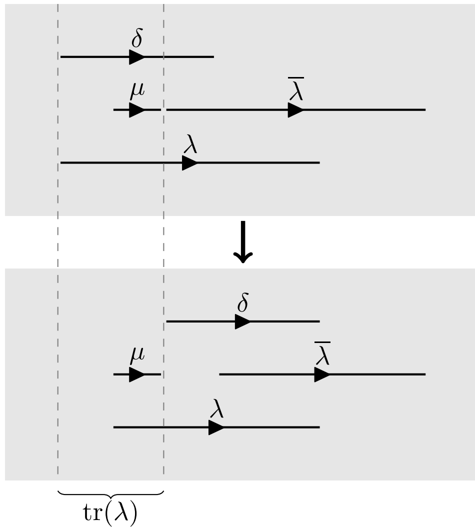

FIGURE 21. An entire transformation where the leading base λ forms a principal overlapping carrier. The transfer base δ is moved to the right by $\mathrm{tr}(\lambda)$. If δ is eventually a leading base again, then $|\lambda|$ will have decreased by $\mathrm{tr}(\lambda)$.

Proof. Each time a base δ is carried by λ it is moved to the right by $\mathrm{tr}(\lambda)$. The next time that base is carried the base λ will be shortened by $\mathrm{tr}(\lambda)$ so the advertised bound holds. □

It therefore follows that the periodicity of λ bounds the number of consecutive times λ can be a carrier in a Rips process. Equivalently, this is the maximal length of a λ-periodic path. Given a band complex $\mathscr{C}$ that has a principal overlapping carrier λ, we will compute an upper bound for $\mathrm{period}_t(\lambda)$ that holds for every automorphically minimal κ-track t. If $\mathrm{period}_t(\lambda)$ exceeds this bound, we prove the existence of another automorphically equivalent track t' efficiently carried by $\mathscr{C}$ such that $\mathrm{Size}(t') < \mathrm{Size}(t)$. It is worth noting that, since we are requiring *automorphic* equivalence, simply applying Bulitko's Lemma is not good enough, as the latter only bounds minimal periodicity in possibly non-injective homomorphic images.

To compute this bound we will construct *auxiliary elimination trees.* This will require all the machinery developed up until now as well as a few new ideas. The auxiliary tree will in fact be a rooted tree of trees. Off the leaves of this tree we will be able to read an upper bound for the periodicity.

Proposition 6.15 combined with this periodicity bound for every occurring principal overlapping carrier in $\mathfrak{A}_+(C)$ will finally enable us to construct a finite subtree guaranteed to give us all the tracks we need for Theorem B.

7.1. **Periodic block form, and the tree $\mathfrak{T}_{\mathrm{PBF}}(\mathscr{C}, J)$.**

Definition 7.4. Suppose that for some overlapping pair $(\lambda, \overline{\lambda})$, $\sigma(\lambda)$ is a maximal section (recall Definition 2.7). Then we call $\sigma(\lambda)$ a *periodic block.* A band complex such that every base lies in some periodic block is said to be in *periodic block form.*

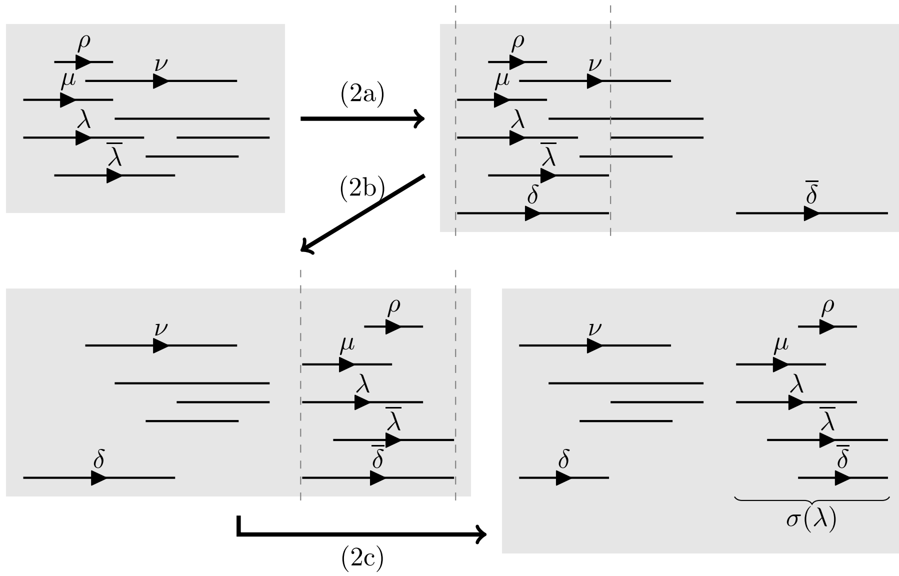

FIGURE 22. Creating a new periodic block for the principal overlapping carrier λ. Depicted is the sequence (2a) - (2c) in the construction of $\mathfrak{T}_{\mathrm{PBF}}(\mathscr{C}, J)$.

We now describe another J-restricted elimination process which constructs the tree $\mathfrak{T}_{\mathrm{PBF}}(\mathscr{C}, J)$. This process brings a band complex $\mathscr{C}$ into periodic block form and is a variation of the construction of $\mathfrak{A}_+(\mathscr{C}, J)$, only this time, whenever we encounter a principal overlapping carrier, we add a new band $\mathbb{B}(\delta)$, transfer $(\lambda, \overline{\lambda})$ all the way to the right using entire transformations, and enlarge J to $J \cup \sigma(\lambda)$; see Figure 22. By Proposition 6.15, we will have constructed a finite admissible tree in which every leaf is a band complex whose bases either lie in J or in some periodic block. We build $\mathfrak{T}_{\mathrm{PBF}}(\mathscr{C}, J)$ as follows:

(1) We perform a J-restricted elimination process, adding only admissible band complexes (recall Definition 6.12).
(2) If at some point as we grow our tree there is a $\mathscr{C}_i$ with a principal overlapping carrier λ, then instead of doing an entire transformation we do the following:
  (a) We attach a band $\mathbb{B}(\delta)$ to $\mathscr{C}_i$ by identifying $\delta \doteqdot \sigma(\lambda)$, so that $\overline{\delta}$ doesn't meet any other bases. We extend the ordering $<_i$ so that the maximal section corresponding to $\overline{\delta}$ is terminal (Definition 3.21).
  (b) We transfer all the bases contained in $\sigma(\lambda)$ through $\mathbb{B}(\delta)$ onto $\overline{\delta}$.
  (c) We collapse the naked initial segment of δ. This gives the $\mathscr{C}_i'$.
  (d) We set $J' = J \cup \sigma(\lambda)$ and we continue growing our tree at $\mathscr{C}_i'$ by returning to step (1) but with J' in place of J.
(3) If after step (2d) we have moved all the bases onto J', then we stop.

By direct inspection we verify:

Lemma 7.5. *If* $(\mathscr{C}_i, J)$ *is quadratic or superquadratic, then after performing steps* (2a) - (2d)*, the resulting band complex is still quadratic or superquadratic* (*respectively*)*.*

Lemma 7.6. *If we add a band* $\mathbb{B}(\delta)$ *to* $\mathscr{C}_1$ *as in* (2a)*, then transfer all the bases contained in* $\sigma(\lambda)$ *onto* $\overline{\delta}$ *as in* (2b) *to produce* $\mathscr{C}_i'$*. Then* $\tau(\mathscr{C}_i) \geq \tau(\mathscr{C}_i')$*. Moreover if* J *and* J' *are as in* (2d) *above, then we have a strict inequality*

$$\tau(\mathscr{C}_i, J) > \tau(\mathscr{C}_i', J').$$

Proof. We keep track of the τ-complexity. Adding the band $\mathbb{B}(\delta)$ adds a base to a maximal section σ with $b(\sigma) \geq 2$ (recall Definitions 2.7 and 3.13), which increases the τ-complexity by 1 and creates another maximal section $\sigma(\overline{\delta})$ with only one base, which doesn't contribute to the τ-complexity. We then move the base λ onto $\sigma(\overline{\delta})$, which decreases $b(\sigma)$ by 1. Now $b(\sigma(\overline{\delta})) = 2$, so it still contributes 0 to the τ-complexity. So far we have added 1 and removed 1 from the τ-complexity. Since λ was assumed to be a leading base and $\delta \supset \lambda$, doing the rest of (2b) and (2c) amounts to a sequence of transformations which do not increase the τ-complexity.

After all this $\sigma(\lambda)$ is a maximal section with $\tau(\sigma(\lambda)) > 0$ and $\sigma(\lambda) \cap J = \emptyset$, which implies that $\tau(\mathscr{C}_i, J) > \tau(\mathscr{C}_i', J')$. □

Corollary 7.7. $\mathfrak{T}_{\mathrm{PBF}}(\mathscr{C}, J)$ *is finite. Furthermore, if* $t \subset \mathscr{C}$ *is an efficiently carried automorphically minimal* κ*-track, then the path* $\mathscr{C} \to \cdots \to \mathscr{C}_p$ *to a band complex in periodic block form induced by* t *is contained in* $\mathfrak{T}_{\mathrm{PBF}}(\mathscr{C}, J)$*.*

Proof. We first prove the first statement. Suppose towards a contradiction that this is not the case; then $\mathfrak{T}_{\mathrm{PBF}}(\mathscr{C}, J)$ has an infinite branch $\mathfrak{b}$. Since we are constructing an admissible elimination tree in the sense of Definition 6.12, by Proposition 6.15 this infinite branch can be assumed to start with some $\mathscr{C}_u$ with a principal overlapping carrier λ. By the definition of $\mathfrak{T}_{\mathrm{PBF}}(\mathscr{C}, J)$, this means we must construct a new periodic block, item (2), which by Lemma 7.6 strictly decreases the relative τ-complexity, so this event can only happen finitely many times, contradicting the fact that $\mathfrak{b}$ is infinite.

The second claim follows immediately from Proposition 6.13. □

7.2. **Normalized periodic block form.** Once a band complex $\mathscr{C}$ is in periodic block form, it will be possible to perform periodic block mergers, which decreases the number of periodic blocks and the τ-complexity. Furthermore it will enable us to put a partial order on the periodic blocks, called a periodic hierarchy. First we give another version of entanglement.

Definition 7.8. If $(\mu, \overline{\mu})$ is an overlapping pair such that $\sigma(\mu)$ is a maximal section (Definition 2.7), then $(\mu, \overline{\mu})$ is called a *block overlapping pair.*

In particular, if $\mathscr{C}$ is in periodic block form, then every maximal section is in fact a block overlapping pair.

Definition 7.9. Let $(\lambda, \overline{\lambda})$ and $(\mu, \overline{\mu})$ be disjoint block overlapping pairs. Suppose there is a band $\mathbb{B}(\delta)$ connecting $\sigma(\lambda)$ and $\sigma(\mu)$, i.e., $\delta \subset \sigma(\lambda)$ and $\overline{\delta} \subset \sigma(\mu)$. For any $p \in \mathbb{B}(\delta)$ we can define δ-relative tubular elements $\tau_\lambda^\delta, \tau_\mu^\delta$ as in Figure 23. $(\lambda, \overline{\lambda})$ *and* $(\mu, \overline{\mu})$ *are entangled by* $\mathbb{B}(\delta)$ if

$$[\tau_\lambda^\delta, \tau_\mu^\delta] = 1.$$

Lemma 7.10. *Let* $(\mu, \overline{\mu})$ *and* $(\lambda, \overline{\lambda})$ *be entangled by* $\mathbb{B}(\delta)$*. If* $[\tau_\mu^\delta, \tau_\lambda^\delta] = 1$*, then* τ_μ^δ *and* τ_λ^δ *have the same axis in the dual Bass-Serre tree* $T(t, \mathscr{C})$ *and lie in a common cyclic subgroup.*

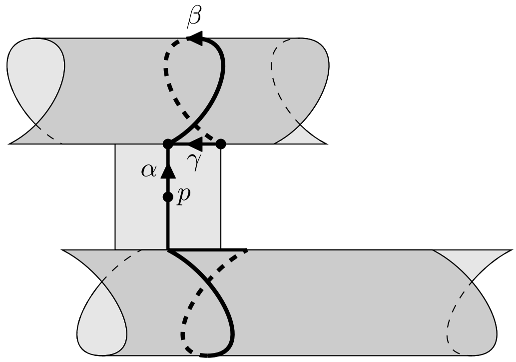

FIGURE 23. δ-relative tubular elements. The relative tubular element τ_λ^δ is the loop based at p given by $\alpha * \beta * \gamma * \alpha^{-1}$.

Sketch of proof. This is proved in the same way as Lemma 5.12. □

Lemma 7.11. *In the band complex $\mathscr{C}$, let $(\mu, \overline{\mu})$ and $(\lambda, \overline{\lambda})$ be block overlapping pairs, let $\mathbb{B}(\delta)$ satisfy $\delta \subset \sigma(\mu)$, and let $\overline{\delta} \subset \sigma(\lambda)$. If for some $t \in \mathbf{tracks}_\kappa(\mathscr{C})$,*

$$|\delta|_t > \max\{\mathrm{tr}(\mu) + \kappa, \mathrm{tr}(\lambda) + \kappa\},$$

then $(\mu, \overline{\mu})$ and $(\lambda, \overline{\lambda})$ are entangled by $\mathbb{B}(\delta)$. Furthermore the tubular elements τ_λ^δ and τ_μ^δ have the same axis in $T(t, \mathscr{C})$.

Sketch of proof. This is proved in the same way as Lemma 5.13. □

We can't simply widen some band $\mathbb{B}(\lambda)$ and still have a well-formed band complex and preserve the dual tree. We could do this in Lemma 5.16 by studying what happened in the dual tree and by noting that the result was still a well-formed band complex. A similar analysis for block overlapping pairs gives the following:

Lemma 7.12. *Let $(\mu, \overline{\mu})$ be a block overlapping pair in a band complex $\mathscr{C}$ and let $t \in \mathbf{tracks}_\kappa(\mathscr{C})$. Then we can obtain a new band complex $\mathscr{C}' \supset \mathscr{C}$, equipped with an efficiently carried κ-track t', by widening $\mathbb{B}(\mu)$. After widening, μ has an initial naked segment and we recover $\mathscr{C}' \to \mathscr{C}$ by collapsing (Definition 3.3) this initial segment. The dual Bass-Serre trees $T(t', \mathscr{C}')$ and $T(t, \mathscr{C})$ are therefore equivariantly isomorphic. Furthermore $\mathrm{tr}(\mu)$ remains invariant.*

Lemma 7.13. *Let $(\mu, \overline{\mu})$ and $(\lambda, \overline{\lambda})$ be block overlapping pairs entangled by $\mathbb{B}(\delta)$ in a band complex $\mathscr{C}$. Then for any κ-track t efficiently carried by $\mathscr{C}$ we can obtain a new band complex $\mathscr{C}' \supset \mathscr{C}$ by first widening $\mathbb{B}(\lambda), \mathbb{B}(\mu)$ by at most $|\lambda|_t + |\mu|_t$ and then widening $\mathbb{B}(\delta)$ so that*

(i) *either $\sigma(\lambda) \subset \delta$ or $\sigma(\mu) \subset \overline{\delta}$, and*
(ii) *the track t extends to a track $t \subset t'$ efficiently carried by $\mathscr{C}'$ such that there is a $\pi_1(\mathscr{C})$-equivariant isomorphism of dual Bass-Serre trees*

$$T(t, \mathscr{C}) \xrightarrow{\sim} T(t', \mathscr{C}').$$

We note that λ, μ only need to be increased to length at most $|\lambda|_t + |\mu|_t$ for this to work.

7.2.1. *The periodic block merger.* We now describe the *periodic block merger*; see Figure 24. Let $\mathscr{C}$ have block overlapping pairs $(\mu, \overline{\mu})$, $(\lambda, \overline{\lambda})$ entangled by $\mathbb{B}(\delta)$. We do the following:

(1) Widen bands $\mathbb{B}(\lambda), \mathbb{B}(\mu)$, and then $\mathbb{B}(\delta)$ to obtain a band complex $\mathscr{C}'$ as given in Lemma 7.13.
(2) Assume that $\overline{\delta} \supset \sigma(\lambda)$. We treat $\overline{\delta}$ as the carrier base and move all the bases contained in $\sigma(\lambda)$ onto δ via entire transformations.
(3) The dual pairs $(\lambda, \overline{\lambda}), (\mu, \overline{\mu})$ are now entangled in the sense of Definition 5.10 and both form overlapping pairs, so by Lemma 5.17 we can apply the periodic merger given in Proposition 5.20, merging $\mathbb{B}(\lambda), \mathbb{B}(\mu)$ into some new $\mathbb{B}(\eta)$.

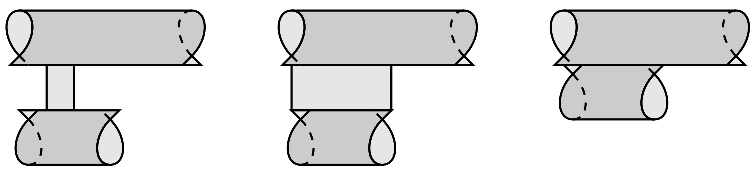

FIGURE 24. Steps (1) and (2) of type periodic block merger.

We now record the following observation.

Lemma 7.14. *Let $\mathscr{C}$ be a band complex containing block overlapping pairs $(\lambda, \overline{\lambda})$, $(\mu, \overline{\mu})$ entangled by $\mathbb{B}(\delta)$. Then we can effectively construct a finite set of band complexes*

$$\begin{array}{ccc} & \mathscr{C} & \\ \swarrow & & \searrow \\ \mathscr{C}_1 & \cdots & \mathscr{C}_n \end{array}$$

containing all possible combinatorial outcomes $m : (\mathscr{C}, t) \to (\mathscr{C}', t')$ of applying a periodic block merger as t ranges over $\mathbf{tracks}_\kappa(\mathscr{C})$.

7.2.2. *Normalization: Merging away entanglement.* It may be that $(\mu, \overline{\mu})$ is a block overlapping pair that is entangled with $(\lambda, \overline{\lambda})$ but $|\lambda|_t < \mathrm{tr}(\mu)$, for example if $\lambda, \overline{\lambda}$ are very short and near the extremities of $(\mu, \overline{\mu})$. In this situation it is not possible to directly apply a periodic merger. However since block overlapping pairs can always be widened we have the following:

Lemma 7.15. *Let $\mathscr{C}$ have a block overlapping pair $(\mu, \overline{\mu})$ entangled with $(\lambda, \overline{\lambda})$. For any $t \in \mathbf{tracks}_\kappa(t)$ we can widen $\mathbb{B}(\mu)$ by $\mathrm{tr}(\mu)$ so that after widening $\mathbb{B}(\lambda)$ we have $|\lambda|_t \geq \mathrm{tr}(\mu)$.*

Proof. By Lemma 7.12, we can widen $(\mu, \overline{\mu})$ (by at most $\mathrm{tr}(\mu)$) so that λ is at distance more than $\mathrm{tr}(\mu)$ from the endpoints of $\sigma(\mu)$. It follows that after widening as in Lemma 5.16, λ is sufficiently long. □

Corollary 7.16. *Let $\mathscr{C}$ have a block overlapping pair $(\mu, \overline{\mu})$ entangled with $(\lambda, \overline{\lambda})$. For any $t \in \mathbf{tracks}_\kappa(t)$, after perhaps widening $\mathbb{B}(\mu)$ by at most $\mathrm{tr}(\mu)$ and $\mathbb{B}(\lambda)$ as in Lemma 5.16, we can perform a periodic merger (Proposition 5.20) of $\mathbb{B}(\mu)$ and $\mathbb{B}(\lambda)$.*

Definition 7.17. Let $\mathscr{C}$ be a band complex in periodic block form and let $t \in \mathbf{tracks}_\kappa(\mathscr{C})$. Then we can repeatedly apply periodic block mergers and periodic mergers as described in Corollary 7.16 so no block overlapping pair is entangled with any other base pair. If such is the case, then we call corresponding $m : (\mathscr{C}, t) \to (\mathscr{C}', t')$ the *normalization of* $(\mathscr{C}, t)$ and we say that $\mathscr{C}'$ is in *normalized periodic block form.*

We now give the combinatorial equivalent. Recall that entanglement is algorithmically decidable and depends only on the band complex $\mathscr{C}$.

Definition 7.18. Let $\mathscr{C}$ be a band complex in periodic block form. Then its *normalized children* is the collection of band complexes

$$\begin{array}{ccc} & \mathscr{C} & \\ \swarrow & & \searrow \\ \mathscr{C}_1 & \cdots & \mathscr{C}_n \end{array}$$

obtained by enumerating all possible combinatorial outcomes of normalizations $m : (\mathscr{C}, t) \to (\mathscr{C}', t')$ where t ranges over $\mathbf{tracks}_\kappa(\mathscr{C})$.

7.3. **Periodic hierarchies and maximal periodic blocks.** Throughout this section $\mathscr{C}$ will be a band complex in normalized periodic block form.

Definition 7.19. Let $\mathscr{C}$ be in normalized periodic block form, let $t \in \mathbf{tracks}_\kappa(\mathscr{C})$, and let $\sigma(\lambda)$ be a periodic block. A base $\delta \subset \sigma(\lambda)$ is *t-long in $\sigma(\lambda)$* if $|\delta|_t \geq \mathrm{tr}(\lambda) + \kappa + 1$. Otherwise it is called *t-short.*

Lemma 7.20. *Let $\mathscr{C}, t, \lambda$, and δ be as in Definition* 7.19*. Then if $\overline{\delta}$ also lies in $\sigma(\lambda)$, δ must be t-short in $\sigma(\lambda)$.*

Proof. If $(\delta, \overline{\delta})$ is orientation reversing, then this follows immediately from Lemma 6.7. If $(\delta, \overline{\delta})$ is orientation preserving but t-long in $\sigma(\lambda)$, then it is entangled with $(\lambda, \overline{\lambda})$ but not merged. By Corollary 7.16 this contradicts the assumption that $\mathscr{C}$ is in normalized periodic block form. $\square$

It is possible for a base to be long in one periodic block, but its dual must lie in another periodic block and it must be short in that periodic block. If a base and its dual are both long in their respective periodic blocks, then we can perform a periodic block merger.

Definition 7.21 (Periodic hierarchies)**.** Let the band complex $\mathscr{C}$ be in normalized periodic block form. A *periodic hierarchy* $\mathcal{H}$ is a partial order $<_\mathcal{H}$ on the set of periodic blocks that is generated as follows:

(i) If $\mathbb{B}(\delta)$ has bases lying in periodic blocks $\sigma(\lambda_1)$ and $\sigma(\lambda_2)$, then we may either declare $\sigma(\lambda_1) <_\mathcal{H} \sigma(\lambda_2)$, $\sigma(\lambda_2) <_\mathcal{H} \sigma(\lambda_1)$ or that $\sigma(\lambda_1)$ and $\sigma(\lambda_2)$ are incomparable.

(ii) We extend (i) to a partial order, if possible.

A periodic block $\sigma(\lambda)$ is *$\mathcal{H}$-maximal* if it is maximal with respect to the partial order. If $t \in \mathbf{tracks}_\kappa(\mathscr{C})$, then we define the *induced periodic hierarchy $\mathcal{H}(t)$* to be generated by setting $\sigma(\lambda_1) <_{\mathcal{H}(t)} \sigma(\lambda_2)$ if and only if δ is t-long in $\sigma(\lambda_1)$ in (i).

It is obvious that periodic hierarchies, being finite combinatorial objects, can be effectively listed. What is less obvious is whether the definition of an induced periodic hierarchy actually gives a periodic hierarchy.

Lemma 7.22. *If $\mathscr{C}$, t and $\mathcal{H}(t)$ are as in Definition* 7.21*, then $<_{\mathcal{H}(T)}$ gives a partial order on the set of periodic blocks; thus $\mathcal{H}(t)$ is a periodic hierarchy.*

Proof. Suppose that for some $\mathbb{B}(\delta)$, both δ and $\overline{\delta}$ are long in $\sigma(\lambda_1)$ and $\sigma(\lambda_2)$, respectively. Then by Lemma 7.11 the block overlapping pairs are entangled and can be merged, contradicting the assumption that $\mathscr{C}$ is normalized, so (i) of Definition 7.21 is satisfied. Thus (i) gives a directed graph Γ without loops of length 2 with periodic blocks as vertices. If Γ has a directed cycle, then this would imply that for some $(\lambda, \overline{\lambda})$, $\mathrm{tr}(\lambda) < \mathrm{tr}(\lambda)$, which is absurd. We can therefore extend $\mathcal{H}$ to a partial order giving (ii) of Definition 7.21. □

Corollary 7.23. *If $(\lambda, \overline{\lambda})$ is an $\mathcal{H}(t)$-maximal block overlapping pair, then every base $\delta \subset \sigma(\lambda)$ is t-short.*

7.4. **Bounding the periodicity of maximal periodic blocks.**

Definition 7.24. Suppose $\sigma(\lambda)$ is a periodic block and suppose that we can vertically subdivide $\mathbb{B}(\lambda)$ into three bands

$$\mathbb{B}(\lambda_1), \qquad \mathbb{B}(\lambda_2), \qquad \mathbb{B}(\lambda_3)$$

such that:

- $(\lambda_2, \overline{\lambda_2})$ form an overlapping pair.
- No bases other than $\overline{\lambda_1}$ and λ_3 intersect $\sigma(\lambda_2)$.
- $\mathbb{B}(\lambda_2)$ contains no connections (recall Definition 2.3(3a)).

Then we call $\mathbb{B}(\lambda_2)$ a *clean tube.*

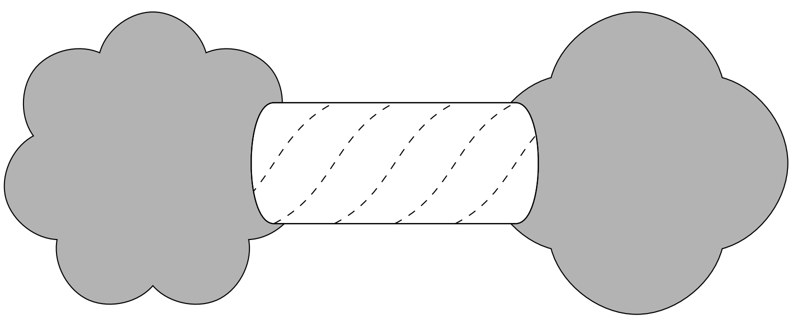

FIGURE 25. A clean tube in a band complex is literally an embedded $S^1 \times [-1, 1]$. The track t spirals around the clean tube. Unwinding it by a Dehn twist will decrease the size.

The significance of clean tubes is illustrated in Figure 25. Although the next result is obvious from this picture, it is important to state it carefully to get explicit bounds.

Lemma 7.25. *Let t be an automorphically minimal κ-track efficiently carried by $\mathscr{C}$, and let $\sigma(\lambda)$ be a periodic block. Then for any clean tube $\mathbb{B}(\lambda_2) \subset \mathbb{B}(\lambda)$,*

$$|\lambda_2|_t \leq 2\mathrm{tr}(\lambda).$$

Proof. Suppose towards a contradiction that for some clean tube $\mathbb{B}(\lambda_2)$ we have

$$|\lambda_2|_t > 2\mathrm{tr}(\lambda).$$

Let $I \subset \lambda_2 \cap \overline{\lambda_2}$ be an interval of length $\mathrm{tr}(\lambda)$ that is of distance $\mathrm{tr}(\lambda)$, the leftmost endpoint of $\sigma(\lambda_2)$. Let $p \in t \cap I$ be the rightmost point of $t \cap I$. If we follow the connected component of $t \cap \mathbb{B}(\lambda_2)$ that contains p and intersects I again in p', we see that the distance between p and p' in λ_2 is exactly $\mathrm{tr}(\lambda)$. Recall the notation of Definition 2.1 and consider the map

$$J_{\mathbb{B}(\lambda_2)} \times [-1,1] \to \mathscr{C}$$

with $J_{\mathbb{B}(\lambda_2)} \times \{1\} \to \lambda_2$ and $J_{\mathbb{B}(\lambda_2)} \times \{-1\} \to \overline{\lambda_2}$. The preimage of I has two connected components $I_{\pm 1} \subset J_{\mathbb{B}(\lambda_2)} \times \{\pm 1\}$. Let α be the straight line in $J_{\mathbb{B}(\lambda_2)} \times [-1,1]$ between the rightmost point of I_1 and the rightmost point of I_{-1} and let β be the line in $J_{\mathbb{B}(\lambda_2)} \times [-1,1]$ between the leftmost point of I_1 and the leftmost point of I_{-1}. α and β are chosen to be transverse to the preimage of t. Let Q be the quadrilateral in $J_{\mathbb{B}(\lambda_2)} \times [-1,1]$ enclosed by $I_{\pm 1}, \alpha, \beta$ (see Figure 26). Then via $J_{\mathbb{B}(\lambda_2)} \times [-1,1] \to \mathscr{C}$, Q is mapped to an annulus $A \subset \mathscr{C}$ such that $Q \subset \mathbb{B}(\lambda_2)$ with α, β mapping onto each component of ∂A.

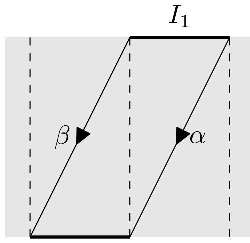

FIGURE 26. The quadrilateral Q inside $\mathbb{B}(\lambda_2)$. The annulus A is obtained by identifying I_1 and I_{-1}. The track is drawn as dashed lines.

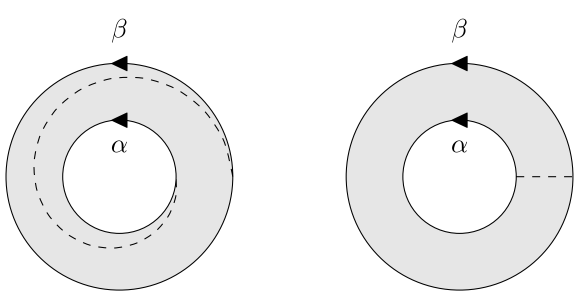

FIGURE 27. On the left the image A of Q after identifying I_1 and I_2 parameterized as (15) (drawn to scale). On the right the result of the Dehn twist τ_A. The dashed line represents $t_m \cap A$. Note that τ_A restricts to the identity on ∂A.

We parameterize this annulus A as

$$\{r\exp(i\theta) | 1 \leq r \leq 2, 0 \leq \theta \leq 2\pi\} \subset \mathbb{C}. \tag{15}$$

By construction, $t_A = t \cap A$ corresponds to the curve

$$\begin{array}{rcl} t_A : [0,1] & \to & A \\ s & \mapsto & (2-s)\exp(i(2\pi s)). \end{array}$$

If we make a Dehn τ_A twist around A, then post-composing gives

$$\begin{array}{rcl} \tau_A \circ t_A : [0,1] & \to & A \\ s & \mapsto & (2-s). \end{array}$$

Such a Dehn twist is illustrated in Figure 27. Since the homeomorphism τ_A restricts to the identity on ∂A, it extends to a homeomorphism of $\mathscr{C}$ which we will also denote τ_A. Consider the new track $\tau_A(t)$. On one hand, $\tau_A(t)$ is still efficiently carried by $\mathscr{C}$. On the other hand, after perturbing by an isotopy we have $\tau_A(t) \cap \sigma(\lambda) = t \cap \sigma(\lambda) - 1$, but otherwise for every subset $S \subset \mathscr{C} \setminus \mathbb{B}(\lambda_2)$ we have $S \cap t = S \cap \tau_A(t)$. It follows that

$$\mathrm{Size}(\tau_A(t)) < \mathrm{Size}(t). \tag{16}$$

The homeomorphism $\tau_A : \mathscr{C} \to \mathscr{C}$ lifts to a homeomorphism of $\widetilde{\tau_A} : \widetilde{\mathscr{C}} \to \widetilde{\mathscr{C}}$ sending the lift $\widetilde{t}$ to the lift $\widetilde{\tau_A(t)}$. Furthermore τ_A maps naturally to an element of $\mathrm{Aut}(\pi_1(\mathscr{C}))$. It follows that the trees $T(t, \mathscr{C})$ and $T(\tau_A(t), \mathscr{C})$ are $(\tau_A)_\sharp$-equivariantly isomorphic so that $t \sim_{\mathrm{Aut}(\pi_1(\mathscr{C}))} \tau_A(t)$. $\tau_A(t)$ is therefore obviously a κ-track, and with (16) we see that t is not automorphically minimal, which is a contradiction. □

Informally, if clean tubes are longer than twice the translation length, there would be enough room to make an annulus, as shown in Figures 26, 27. This would enable us to shorten the track t by a Dehn twist, contradicting automorphic minimality. Clean tubes therefore must be short. We now use this observation to give a combinatorial bound for periodicity.

Corollary 7.26 (Periodicity bound for maximal periodic blocks)**.** *Let $t \subset \mathscr{C}$ be an efficiently carried automorphically minimal κ-track and let $\sigma(\lambda)$ be an $\mathcal{H}(t)$-maximal periodic block. Let $|C|$ denote the number of connected components of connection preimages in $\mathbb{B}(\lambda) = J_{\mathbb{B}(\lambda)} \times [-1,1]$ and let B denote the number of bases, other than $\lambda, \overline{\lambda}$, contained in $\sigma(\lambda)$. Then*

$$\mathrm{period}_t(\lambda) \leq (6 + 2\kappa)B + 2|C| + 2.$$

Proof. Parameterize $\mathbb{B}(\lambda)$ as $J_\lambda \times [-1,1]$ so that the preimage of t and of every connection is contained in a union of vertical lines. Every base $\delta \subset \mathbb{B}(\lambda)$ has a preimage with connected components $\delta^+ \subset J_\lambda \times \{1\}$ and $\delta^- \subset J_\lambda \times \{-1\}$. Let C denote the preimage of the connections. Consider the complement

$$Y = J_{\mathbb{B}(\lambda)} \times [-1,1] \setminus \left(C \cup \left(\bigcup_\delta \left(\delta^\pm \times [-1,1]\right) \right) \right),$$

where δ runs over the bases contained in $\sigma(\lambda)$. On one hand Y has at most $|C| + 2B + 1$ connected components; on the other hand every maximal clean tube contained in $\mathbb{B}(\lambda)$ is the preimage of one of these components. Thus every connected component of Y has length at most $2\mathrm{tr}(\lambda)$ by Lemma 7.25. Furthermore, by Corollary 7.23

and Definition 7.19, every $\delta^{\pm}$ has width at most $\mathrm{tr}(\lambda)+\kappa$. We therefore have the bound

$$\begin{aligned} |\lambda|_t &\leq 2B\left(\mathrm{tr}(\lambda)+\kappa\right)+2\mathrm{tr}(\lambda)\left(|C|+2B+1\right) \\ &\leq 2B\left(\mathrm{tr}(\lambda)+\mathrm{tr}(\lambda)\kappa\right)+2\mathrm{tr}(\lambda)\left(|C|+2B+1\right) \\ &\leq \mathrm{tr}(\lambda)\left((6+2\kappa)B+2|C|+2\right), \end{aligned}$$

from which we immediately obtain the advertised bound. □

The significance of this bound is that it depends only on the combinatorial band complex $\mathscr{C}$ and the combinatorial periodic hierarchy $\mathcal{H}(t)$. We obtain the following computable function.

Definition 7.27. Let $\mathscr{C}$ be in normalized periodic block form, let $\mathcal{H}$ be a periodic hierarchy on $\mathscr{C}$, and let $\sigma(\lambda)$ be an $\mathcal{H}$-maximal periodic block. We define

$$\mathrm{period}_{\mathcal{H}}(\lambda)=(6+2\kappa)B+2|C|+2,$$

where $B, |C|$ are as in Corollary 7.26.

Unfortunately it may be that the periodic block whose periodicity we are interested in is not $\mathcal{H}$-maximal. We deal with this in the next section.

7.5. **Bounding the periodicity of principal overlapping carriers: Auxiliary trees.** Throughout this section we will use the following notation. If $\mathscr{C}$ is a band complex, then the term J in the pair $(\mathscr{C}, J)$ will always denote a union of block overlapping pairs. If $\mathscr{C}$ happens to be in normalized periodic block form and is equipped with a periodic hierarchy $\mathcal{H}$, then we will denote the corresponding triple $(\mathscr{C}, J; \mathcal{H})$.

We will try to bound the periodicity of the block overlapping pairs that constitute J. It may happen however that two block overlapping pairs get merged at some point. Because of this we will use the following naming convention.

Convention 7.28 (Renaming merged bases in auxiliary trees)**.** Suppose that two bands $\mathbb{B}(\lambda), \mathbb{B}(\mu)$ get merged onto some band $\mathbb{B}(\eta)$ in a periodic merger. Then, as far as naming bases is concerned, we will consider $\eta=\lambda=\mu$; i.e., we will allow a base to have multiple names.

This renaming convention is justified, since on one hand, we want to bound the periodicity of an overlapping pair that gets merged, so we must keep track of what it got merged with. On the other hand, by the following result, which is an immediate consequence of Proposition 5.20, we are guaranteed that any periodicities we compute will be overestimates.

Lemma 7.29. *Let $(\lambda, \overline{\lambda})$ be an overlapping pair such that $\mathbb{B}(\lambda)$ gets zipped onto $\mathbb{B}(\eta)$ via a periodic merger $\mathscr{C} \to \mathscr{C}'$ mapping an efficiently carried κ-track $t \subset \mathscr{C}$ into the efficiently carried κ-track $t' \subset \mathscr{C}'$. Then*

$$\mathrm{period}_t(\lambda) \leq \mathrm{period}_{t'}(\eta).$$

In particular periodicities of block overlapping pairs are non-decreasing when passing to normalizations.

Definition 7.30. Let $(\mathscr{C}, J)$ be a band complex with a principal overlapping pair $(\lambda, \overline{\lambda})$. Then we define the *auxiliary children* of $(\mathscr{C}, J)$ to the collection:

$$\begin{array}{ccc} & (\mathscr{C}, J) & \\ \swarrow & & \searrow \\ (\mathscr{C}_1, J \cup \sigma(\lambda); \mathcal{H}_1) & \cdots & (\mathscr{C}_N, J \cup \sigma(\lambda); \mathcal{H}_N) \end{array}$$

where

$$\{(\mathscr{C}_1, J \cup \sigma(\lambda); \mathcal{H}_1), \ldots, (\mathscr{C}_N, J \cup \sigma(\lambda); \mathcal{H}_N)\}$$

is obtained by first taking the leaves of $\mathfrak{T}_{\mathrm{PBF}}(\mathscr{C}, J)$, then taking their normalized children (Definition 7.18), and finally taking all combinatorial possibilities for periodic hierarchies (Section 7.3).

In Section 7.3 we bounded the periodicity of a maximal block overlapping pair, but it may be that $\sigma(\lambda)$, where λ is a principal overlapping pair in $(\mathscr{C}, J)$, is not a maximal periodic block in some of the auxiliary descendants $(\mathscr{C}_i, J \cup \sigma(\lambda), \mathcal{H}_i)$, due to the choice of periodic hierarchy $\mathcal{H}_i$. It follows that simply passing to auxiliary children isn't sufficient to bound $\mathrm{period}_t(\lambda)$ where t ranges over the automorphically minimal κ-tracks efficiently carried by $\mathscr{C}$.

In Section 7.5.1 we will construct from $(\mathscr{C}, J)$ a finite auxiliary tree $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, J \cup \sigma(\lambda))$. If $\mathscr{C}$ efficiently carries an automorphically minimal κ-track t, then we will construct an induced tree $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, t, J \cup \sigma(\lambda))$ in Section 7.5.2. This induced tree will be proved to contain a combinatorial witness for an upper bound of $\mathrm{period}_t(\mu)$ for some base μ such that $\sigma(\mu) \subset J$ is a periodic block. We will also have a containment

$$\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, t, J \cup \sigma(\lambda)) \subset \mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, J \cup \sigma(\lambda)).$$

From this it will follow that the construction of $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, \sigma(\lambda))$ will give a way to bound $\mathrm{period}_t(\lambda)$ in $\mathscr{C}$ where t ranges over the automorphically minimal κ-tracks.

7.5.1. *The auxiliary tree* $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, J \cup \sigma(\lambda))$. We construct $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, J \cup \sigma(\lambda))$ with the following recursive algorithm. The reader may skip ahead to Figure 28 to get an idea of what this tree is supposed to look like.

(1) If $(\lambda, \overline{\lambda})$ is a principal overlapping pair in $(\mathscr{C}, J)$, then we declare $(\mathscr{C}, J)$ to be the *root of* $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, J \cup \sigma(\lambda))$. Write $J' = J' \cup \sigma(\lambda)$.
(2) The $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, J')$-children of the root $(\mathscr{C}, J)$ are the auxiliary children of $\mathscr{C}$ (Definition 7.30). These are connected to the root by *auxiliary edges*. Further descendants are added as follows:
   (a) If some periodic block $\sigma(\mu) \subset J'$ in an auxiliary child $(\mathscr{C}', J'; \mathcal{H}')$ of $(\mathscr{C}, J)$ is $\mathcal{H}'$-maximal, then $(\mathscr{C}', J'; \mathcal{H}')$ is called a *witnessing terminal*. We stop growing $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, J')$ at $(\mathscr{C}', J')$.
   (b) Otherwise we modify the order $<'$ on $\mathscr{C}'$ (Definition 3.21) so that some $\mathcal{H}'$-maximal periodic block $\sigma(\delta)$ is initial with δ the carrier base.
   We start building $\mathfrak{T}_+(\mathscr{C}, J')$ rooted at $(\mathscr{C}', J')$. For every path originating from the root we forbid the base δ from being the carrier base more than $B \cdot \mathrm{period}_{\mathcal{H}'}(\delta)$ times in a row, where B denotes the number of bases in $\mathscr{C}'$. Once δ ceases to be the carrier base we forget about $\mathcal{H}'$ and go to (3) below.

(3) We continue growing $\mathfrak{T}^{\text{aux}}(\mathscr{C}, J')$ using the following rules:
   (a) If the leading base μ in some $(\mathscr{C}'', J')$ does not form a principal overlapping pair, then continue growing $\mathfrak{A}_+(\mathscr{C}', J')$ at $(\mathscr{C}'', J')$ by adding admissible descendants (Definition 6.12), if there are any.
   If $(\mathscr{C}'', J')$ has no admissible children and some of its bases are not in J, then $(\mathscr{C}'', J')$ is *halted by inadmissibility.*
   Otherwise, if the elimination process stops because all the bases of $\mathscr{C}''$ were moved into J', then we take the auxiliary children

$$\{(\mathscr{C}''_r, J'; \mathcal{H}_r)\}$$

   of $(\mathscr{C}'', J')$ (which is already in periodic block from, but may not be normalized), equipped with periodic hierarchies. All these children are *witnessing terminals* as in (2a).
   (b) If the carrier base μ in some $(\mathscr{C}'', J')$ is a principal overlapping carrier we first construct $\mathfrak{T}^{\text{aux}}(\mathscr{C}'', J'')$ rooted at $(\mathscr{C}'', J')$, where $J'' = J' \cup \sigma(\mu)$. This is the recursion.
   Next we take

$$W = W(\mathscr{C}'', J'', \mu) = \{(\mathscr{C}_i, J_i; \mathcal{H}_i)\}$$

   to be the set of all witnessing terminals of $\mathfrak{T}^{\text{aux}}(\mathscr{C}'', J'')$ in which $\sigma(\mu)$ is an $\mathcal{H}_i$-maximal block overlapping pair (recall Convention 7.28).
   If this set is empty, then $(\mathscr{C}'', J')$ is declared to be a *halted terminal* and no further descendants are added. Otherwise the following number is defined and computable:

$$\text{period}_{(\mathscr{C}'', J')}(\mu) = \max_{(\mathscr{C}_i, J_i; \mathcal{H}_i) \in W} \text{period}_{\mathcal{H}_i}(\mu). \tag{17}$$

   We continue growing $\mathfrak{A}_+(\mathscr{C}', J')$ at $(\mathscr{C}'', J')$, but we forbid μ from being the carrier base more than $B \cdot \text{period}_{(\mathscr{C}'', J')}(\mu)$ times in a row in every path originating at $(\mathscr{C}'', J')$. Once μ ceases to be a carrier base we go back to (3a) or (3b) as appropriate.

Definition 7.31. Let $\mathscr{C}'$ be a band complex occurring in $\mathfrak{T}^{\text{aux}}(\mathscr{C}, J)$. The *depth of $\mathscr{C}'$ in $\mathfrak{T}^{\text{aux}}(\mathscr{C}, J)$* is the number of auxiliary edges in $\mathfrak{T}^{\text{aux}}(\mathscr{C}, J)$ connecting $\mathscr{C}$ and $\mathscr{C}'$.

Lemma 7.32. *The maximal depth of a descendant of $\mathscr{C}$ in $\mathfrak{T}^{\text{aux}}(\mathscr{C}, J)$ is at most the relative τ-complexity $\tau(\mathscr{C}, J)$.*

Proof. By Lemma 7.6 all the leaves of $\mathfrak{T}_{\text{PBF}}(\mathscr{C}, J,)$ have strictly smaller τ complexity relative to $J \cup \sigma(\lambda)$. Furthermore, periodic mergers never increase τ-complexity. Finally if the relative τ complexity $\tau(\mathscr{C}', J') = 0$, then by (3a) all its children are witnessing terminals. □

We think of auxiliary edges as being vertical; thus

Definition 7.33. A subtree of $\mathfrak{T}^{\text{aux}}(\mathscr{C}, J)$ sitting inside some $\mathfrak{A}_+(\mathscr{C}', J')$ or, equivalently, without auxiliary edges is called *horizontal.*

Proposition 7.34. *Let λ be a principal overlapping pair in a band complex $\mathscr{C}$ occurring in $\mathfrak{A}_+(\mathscr{C})$. Then $\mathfrak{T}^{\text{aux}}(\mathscr{C}, \sigma(\lambda))$ is finite.*

Proof. We first show that for any $\mathscr{C}''$ occurring in some horizontal subtree $\mathfrak{A}_+(\mathscr{C}', J') \subset \mathfrak{T}^{\text{aux}}(\mathscr{C}, \sigma(\lambda))$ with a principal overlapping carrier μ, the auxiliary tree $\mathfrak{T}^{\text{aux}}(\mathscr{C}'', J' \cup \sigma(\mu))$ is finite.

We prove this by induction on $\tau(\mathscr{C}'', J')$. If $\tau(\mathscr{C}'', J') = 0$, then if it is not itself yet a witnessing terminal, then its auxiliary children are halted terminals and the result follows. Otherwise if $\tau(\mathscr{C}'', J') = 1$, by Lemma 7.6, all auxiliary children must have $(J' \cup \sigma(\lambda))$-relative τ-complexity equal to 0, so they are witnessing terminals.

Now we suppose that all auxiliary trees are finite for all relative τ complexities less than n and that $\tau(\mathscr{C}'', J') = n$. Any auxiliary descendant $(\mathscr{C}''', J''; \mathcal{H}''')$ will have $\tau(\mathscr{C}''', J'') < n$ by Lemma 7.6. We construct $\mathfrak{A}_+(\mathscr{C}''', J'')$ for each auxiliary descendant of $(\mathscr{C}'', J')$ according to rules (3a) and (3b). Whenever a principal overlapping carrier η occurs in some $\mathscr{C}^{(4)}$, the corresponding auxiliary tree built in (3b) is finite by the induction hypothesis. Either $\mathscr{C}^{(4)}$ is a halted terminal or we can compute the finite $\text{period}_{(\mathscr{C}^{(4)}, J'')}(\mu)$ (as given in (17) of step (3b)). By Proposition 6.15, this prevents horizontal subtrees from having infinite branches; thus by König's Lemma they are finite. It follows that $\mathfrak{T}^{\text{aux}}(\mathscr{C}'', J' \cup \sigma(\lambda))$ is finite. The result now follows by induction. ☐

Having established that $\mathfrak{T}^{\text{aux}}(\mathscr{C}, J)$ is finite and therefore effectively constructible, we can now define the following computable function.

Definition 7.35. For a band complex $\mathscr{C}$ with a principal overlapping pair $(\lambda, \overline{\lambda})$ we define

$$\text{period}_{\mathscr{C}}(\lambda) = \max_{(\mathscr{C}_i, J_i, \mathcal{H}_i) \in S} \text{period}_{\mathcal{H}_i}(\lambda),$$

where S is the set of witnessing terminals $(\mathscr{C}_i, J_i, \mathcal{H}_i)$ of $\mathfrak{T}^{\text{aux}}(\mathscr{C}, \sigma(\lambda))$, where, following the renaming Convention 7.28, $\sigma(\lambda)$ is $\mathcal{H}_i$-maximal.

It remains to show that this $\text{period}_{\mathscr{C}}(\lambda)$ gives an upper bound for $\text{period}_t(\lambda)$, where t is an automorphically minimal κ-track efficiently carried by $\mathscr{C}$. This will be done by studying the induced tree.

7.5.2. *The induced tree* $\mathfrak{T}^{\text{aux}}(\mathscr{C}, t, J \cup \sigma(\lambda))$. A triple $(\mathscr{C}, t, J)$ will denote a band complex $\mathscr{C}$, an automorphically minimal κ-track t efficiently carried by $\mathscr{C}$, and a union J of periodic blocks. If $\mathscr{C}$ is in normalized periodic block form, in the 3+1 tuple $(\mathscr{C}, t, J; \mathcal{H}(t))$, $\mathcal{H}(t)$ will denote the periodic hierarchy induced by t (Definition 7.21).

In any (restricted) elimination tree $\mathfrak{T}$ rooted at $(\mathscr{C}, J)$ a track $t \subset \mathscr{C}$ efficiently carried by $\mathscr{C}$ induces a directed path in $\mathfrak{T}$. By Corollary 7.7 and by Definition 7.17 the following makes sense:

Definition 7.36. Let $(\mathscr{C}, t, J)$ be a band complex with a principal overlapping pair $(\lambda, \overline{\lambda})$. Then its *induced auxiliary child* is given by the labelled graph

$$\begin{array}{c}(\mathscr{C}, J, t)\\ \downarrow \\ (\mathscr{C}', t', J \cup \sigma(\lambda); \mathcal{H}(t')),\end{array}$$

where $(\mathscr{C}', t, J \cup \sigma(\lambda); \mathcal{H}(t'))$ is the auxiliary descendant of $(\mathscr{C}, J)$ (Definition 7.30) induced by $t \subset \mathscr{C}$.

The algorithm to construct $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, t, J \cup \sigma(\lambda))$ is analogous to the algorithm to construct $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, J \cup \sigma(\lambda))$ given in Section 7.5.1. The numbering of the clauses is intended to coincide.

(1) If $(\lambda, \overline{\lambda})$ is a principal overlapping pair in $(\mathscr{C}, t, J)$, then we declare $(\mathscr{C}, t, J)$ to be the *root of* $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, t, J \cup \sigma(\lambda))$. Write $J' = J' \cup \sigma(\lambda)$.

(2) The $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, t, J')$-child of the root $(\mathscr{C}, t, J)$ is the induced auxiliary child of $\mathscr{C}$. These are connected to the root by an *auxiliary edge*. Further descendants are added as follows:
   (a) If some periodic block $\sigma(\mu) \subset J'$ in an auxiliary descendant $(\mathscr{C}', t', J'; \mathcal{H}(t'))$ of $(\mathscr{C}, t, J)$ is $\mathcal{H}(t')$-maximal, then $((\mathscr{C}', J'; \mathcal{H}(t'))$ is called a *witnessing terminal*. We stop growing $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}, t, J')$.
   (b) Otherwise we change the order $<$ on $(\mathscr{C}', t', J'; \mathcal{H}(t'))$ so that some $\mathcal{H}(t')$-maximal periodic block $\sigma(\delta)$ is initial with δ the carrier base.
   We start our Rips process, building the path in $\mathfrak{A}_+(\mathscr{C}, J')$ rooted at $(\mathscr{C}', J')$ induced by the track $t' \subset \mathscr{C}'$. By Corollary 7.26, δ is not the carrier base more than $B \cdot \mathrm{period}_{\mathcal{H}'}(\delta)$ times in a row, where B denotes the number of bases in $\mathscr{C}'$. Once δ ceases to be the carrier base we go to (3) below.

(3) We continue growing the path induced by t in $\mathfrak{A}_+(\mathscr{C}, J')$ using the following rules:
   (a) If the carrier base μ in some $(\mathscr{C}'', t'', J')$ is not a principal overlapping carrier, then we add its descendant in $\mathfrak{A}_+(\mathscr{C}, J')$ as usual.
   If all the bases $\mathscr{C}''$ are moved into J', then we take the induced auxiliary child
   $$(\mathscr{C}''', t''', J'; \mathcal{H}(t'''))$$
   of $(\mathscr{C}'', J')$ (which is already in periodic block from but may not be normalized). Again we call $(\mathscr{C}''', t''', J'; \mathcal{H}(t'''))$ a *witnessing terminal* as in (2a).
   (b) If the carrier base μ in some $(\mathscr{C}'', t'', J')$ forms a principal overlapping pair with its dual we first construct the induced auxiliary tree $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}'', t'', J'')$ rooted at $(\mathscr{C}'', J')$, where $J'' = J' \cup \sigma(\mu)$.
   Next we take
   $$R = R(\mathscr{C}'', t'', J'', \mu) = \{(\mathscr{C}_i, t_i, J_i; \mathcal{H}(t_i))\}$$
   to be the set of all witnessing terminals of $\mathfrak{T}^{\mathrm{aux}}(\mathscr{C}'', J'')$ in which $\sigma(\mu)$ is an $\mathcal{H}(t_i)$-maximal block overlapping pair (recall Convention 7.28). If this set is empty, then $(\mathscr{C}'', J')$ is declared to be a *halted terminal* and no further descendants are added. Otherwise the following number is defined and computable (recall Definition 7.27):
   $$\mathrm{period}_{(\mathscr{C}'', t'', J')}(\mu) = \max_{(\mathscr{C}_i, t_i, J_i; \mathcal{H}(t_i)) \in R} \mathrm{period}_{\mathcal{H}(t_i)}(\mu). \tag{18}$$
   We now continue growing the path at $(\mathscr{C}'', t'', J'')$ in $\mathfrak{A}_+(\mathscr{C}', J')$ induced by $t' \subset \mathscr{C}'$. By Corollary 7.26 and Lemma 7.29, μ will not be the carrier base more than $B \cdot \mathrm{period}_{(\mathscr{C}'', t'', J')}(\mu)$ times in a row.
   Once μ ceases to be a carrier base we go back to (3a) or (3b) as appropriate.

The induced auxiliary tree can be thought of as being constructed one vertex at a time (see Figure 28), as opposed to a branching process. Another important distinction is that induced auxiliary trees do not have terminals that are halted by inadmissibility.

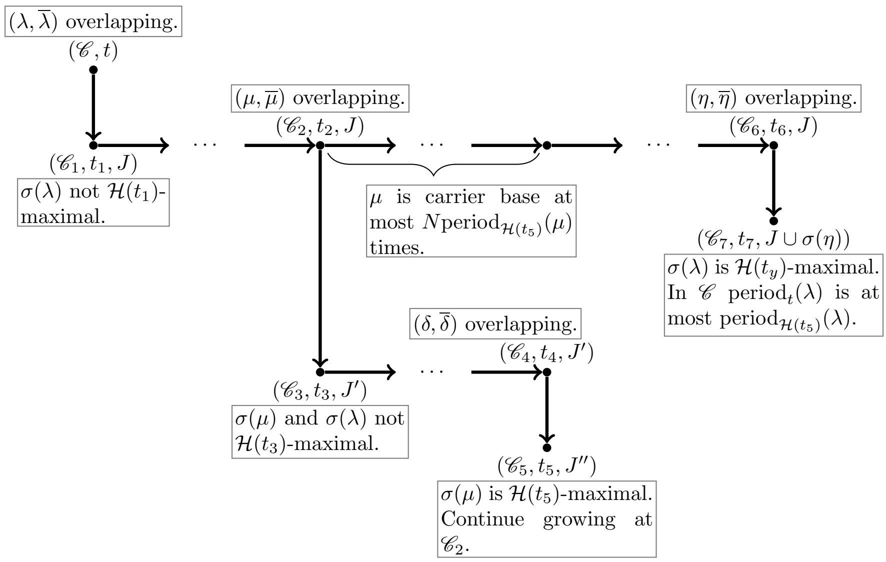

FIGURE 28. The tree $\mathfrak{T}^{\text{aux}}(\mathscr{C},t,\sigma(\lambda))$; here $J=\sigma(\lambda)$, $J'=\sigma(\lambda)\cup\sigma(\mu)$, and $J''=J'\cup\sigma(\delta)$. The band complexes are numbered in order of appearance. Whenever an overlapping pair occurs, an auxiliary edge is constructed. $\mathscr{C}_4$ and $\mathscr{C}_6$ are halted terminals. $\mathscr{C}_5$ and $\mathscr{C}_7$ are witnessing terminals.

Lemma 7.37. $\mathfrak{T}^{\text{aux}}(\mathscr{C},t,J)$ *is finite.*

Proof. By Lemma 7.32 there is a bound on the number of auxiliary edges in any path. The finiteness of horizontal paths follows from the definition of a Rips process; i.e., Size(t) (Definition 4.15) effectively bounds the length of such paths. □

Lemma 7.38. *If* $(\lambda,\overline{\lambda})$ *is a principal overlapping pair in* $\mathscr{C}$*, then there is a witnessing terminal* $(\mathscr{C}_T,t_T,J_T;\mathcal{H}(t_T))$ *in* $\mathfrak{T}^{\text{aux}}(\mathscr{C},t,\sigma(\lambda))$ *in which* $\sigma(\lambda)$ *(following renaming convention* 7.28*) is an* $\mathcal{H}(t_T)$*-maximal periodic block such that*

$$\text{period}_t(\lambda)\le\text{period}_{\mathcal{H}(t_T)}(\lambda).$$

Proof. We first must show that such a witnessing terminal $(\mathscr{C}_T,t_T,J_t;\mathcal{H}(t_T))$ exists. Suppose towards a contradiction that this is not the case. Then the top level elimination tree in $\mathfrak{T}^{\text{aux}}(\mathscr{C},t,\sigma(\lambda))$ must end in a halted terminal $(\mathscr{C}_1,t_1,\sigma(\lambda))$; otherwise all the bases are moved onto $\sigma(\lambda)$, and this periodic block will be maximal (since it's unique).

Let $(\lambda_1,\overline{\lambda_1})$ be the overlapping pair in $\mathscr{C}_1$ and let $J_2=\sigma(\lambda)\cup\sigma(\lambda_1)$. Let $(\mathscr{C}_2,t_2,J_2)$ be the auxiliary descendant, as in (2), of $\mathscr{C}_1$. The horizontal path in $\mathfrak{T}^{\text{aux}}(\mathscr{C},t,\sigma(\lambda))$ starting at $(\mathscr{C}_2,t_2,J_2)$ must end with a halted terminal; otherwise it ends in a witnessing terminal $(\mathscr{C}_T, t_T, J_2; \mathcal{H}(t_T))$ with either $\sigma(\lambda)$ or $\sigma(\lambda)_1$ $\mathcal{H}(t_T)$-maximal. This contradicts either the hypothesis that $\mathscr{C}_1$ is a halted terminal or the hypothesis that $\mathfrak{T}^{\text{aux}}(\mathscr{C}, t, \sigma(\lambda))$ has no witnessing terminals with $\sigma(\lambda)$ maximal.

Continuing in this fashion we obtain a sequence of halted terminals of increasing depth

$$(\mathscr{C}_1, t_1, J_1), (\mathscr{C}_2, t_2, J_2), \ldots, (\mathscr{C}_F, t_F, J_F),$$

where $(\lambda, \overline{\lambda})_i$ is the principal overlapping pair in $\mathscr{C}_i$ and $J_{i+1} = J_i \cup \lambda_i$. Since $\tau(\mathscr{C}_i, J_i) > \tau(\mathscr{C}_{i+1}, J_{i+1})$ this sequence is finite, which forces some principal overlapping pair λ_j in $\mathscr{C}_F$ to be $\mathcal{H}(t_F)$-maximal, contradicting the fact that $\mathscr{C}_j$ is a halted terminal.

It therefore follows that there is a witnessing terminal $(\mathscr{C}_T, t_T, J_t; \mathcal{H}(t_T))$ in $\mathfrak{T}^{\text{aux}}(\mathscr{C}, t, \sigma(\lambda))$ in which $\sigma(\lambda)$ is maximal. The desired upper bound on $\text{period}_t(\lambda)$ now follows immediately from Definition 7.27, Lemma 7.29, and Corollary 7.26. □

So far we have shown that the induced auxiliary tree contains a witnessing terminal whose combinatorial periodicity bounds the actual periodicity. We now bound the periodicity for all tracks.

Proposition 7.39. *Let λ be a principal overlapping pair in a band complex $\mathscr{C}$ in $\mathfrak{A}_+(\mathscr{C})$. Then for all automorphically minimal κ-tracks t efficiently carried by $\mathscr{C}$ the following holds:*

$$\text{period}_t(\lambda) \leq \text{period}_{\mathscr{C}}(\lambda),$$

where $\text{period}_{\mathscr{C}}(\lambda)$ *is the computable function given by Definition* 7.35.

Proof. We first show that for any automorphically minimal κ-track t efficiently carried by $\mathscr{C}$ we have a natural containment:

$$\mathfrak{T}^{\text{aux}}(\mathscr{C}, t, \sigma(\lambda)) \subset \mathfrak{T}^{\text{aux}}(\mathscr{C}, \sigma(\lambda)). \tag{19}$$

We will show this by analyzing how $\mathfrak{T}^{\text{aux}}(\mathscr{C}, t, \sigma(\lambda))$ is constructed by adding one band complex at a time.

Going through the construction algorithms point-by-point, by Proposition 6.13, the only problem that could arise is in step (3b) of the construction of the auxiliary trees. It could be that for some $(\mathscr{C}', t', J')$ with principal overlapping pair $(\mu, \overline{\mu})$, $(\mathscr{C}', t', J')$ is not a halted terminal and $\text{period}_{(\mathscr{C}', t', J')}(\mu)$ from (18) is greater than $\text{period}_{(\mathscr{C}', J')}(\mu)$ from (17).

Note however that in the recursive construction of $\mathfrak{T}^{\text{aux}}(\mathscr{C}, t, \sigma(\lambda))$, we must first construct $\mathfrak{T}^{\text{aux}}(\mathscr{C}', t', J' \cup \lambda)$ before adding a "horizontal" child of $(\mathscr{C}', t', J')$. It follows, by the definition and properties of auxiliary children, that the next vertex added to $\mathfrak{T}^{\text{aux}}(\mathscr{C}, t, \sigma(\lambda))$ is still contained in $\mathfrak{T}^{\text{aux}}(\mathscr{C}, \sigma(\lambda))$.

Continuing in this manner it is obvious (the reader is, of course, free to supply their own argument by induction on relative τ-complexity) that the set R appearing in (18) of case (3b) in the construction of the induced auxiliary tree is a subset of W appearing in (17) of case (3b) of the construction of the auxiliary tree. We conclude that

$$\text{period}_{(\mathscr{C}', t', J')}(\mu) \leq \text{period}_{(\mathscr{C}', J')}(\mu).$$

(19) now follows; thus by Lemma 7.38 a witnessing terminal of $\mathfrak{T}^{\text{aux}}(\mathscr{C}, \sigma(\lambda))$ bounds $\text{period}_t(\lambda)$ from above. □

7.6. **The proof of Theorem B: A description of the main algorithm.** Suppose we are given a finite 2-complex C such that $\pi_1(C)$ has no elements of order 2, a solution to the word problem for $\pi_1(C)$, an acylindricity constant κ, and a finite collection

$$S = \{\{h_i\}_{i\in I_n} \mid n = 1, \dots, m\}$$

of finite generating sets of subgroups $\mathcal{H} = \{\langle h_i\rangle_{i\in I_n}\}_{n=1}^{m}$.

We start by replacing C by C_S given in Section 1.3. This can be done algorithmically. Using the construction of Section 2.2, we see that for any track $t \subset C$ there is a corresponding band complex that carries it efficiently. These band complexes can be effectively enumerated; thus it is possible to construct the finite set $\{\mathscr{C}_1, \dots, \mathscr{C}_{n_C}\}$ of band complexes given by Proposition 2.14 (see also Section 4.0.1 (2)). This gives the first level of our elimination tree. We will now define the ultimate elimination tree $\mathfrak{T}^P{}_+(C)$ in this paper. Here is the final inadmissibility criterion.

Definition 7.40. Let $\mathscr{C}$ be a band complex. A path

$$p : \mathscr{C}_u \to \cdots \to \mathscr{C}_v$$

in $\mathfrak{A}_+(\mathscr{C})$ (recall section 6.5) is called *periodicity-inadmissible* if

(1) p is a μ-periodic path (Definition 6.14) for some base μ in $\mathscr{C}_u$, and
(2) the length of p is greater than $N \cdot \text{period}_{\mathscr{C}_u}(\mu)$ (Definition 7.35), where N is the number of bases in $\mathscr{C}_u$.

$\mathfrak{T}^P{}_+(C)$ is constructed identically to $\mathfrak{A}_+(C)$, but we also forbid periodicity inadmissible paths. By Proposition 7.34 we can compute the periodicity bound and thus effectively decide whether a path is periodicity-inadmissible, so the resulting tree is finite.

To help the reader, however, we will give here a more explicit construction of $\mathfrak{T}^P{}_+(C)$ that will summarize the important results of this paper. We start with our root, the polygonal 2-complex C. We add the descendants $\mathscr{C}_1, \dots, \mathscr{C}_{n_C}$. We then build $\mathfrak{T}^P{}_+(C)$ "one generation at a time" as follows:

(1) For every admissible vertex without descendants we add the descendants as described in Definition 5.23.
(2) If a freshly added descendant is merging inadmissible (Definition 5.21), we declare it inadmissible and stop adding its descendants.
(3) We now consider every directed path constructed so far in our elimination tree. If a path $\mathscr{C}_u \to \cdots \to \mathscr{C}_v$ is either
   - κ-inadmissible (Definition 4.14),
   - repetition inadmissible (Definition 4.18),
   - C-T-inadmissible (Definition 6.11) or
   - periodicity inadmissible,

   then we declare the last vertex $\mathscr{C}_v$ to be inadmissible and stop adding its descendants.

We will now argue that tree $\mathfrak{T}^P{}_+(C)$ is finite and can be algorithmically constructed. First note that at every step we can algorithmically construct the set of descendants of a band complex (Definition 4.4); in particular $\mathfrak{T}^P{}_+(C)$ has finite branching. By König's lemma it is therefore enough to show that $\mathfrak{T}^P{}_+(C)$ has no directed infinite paths.

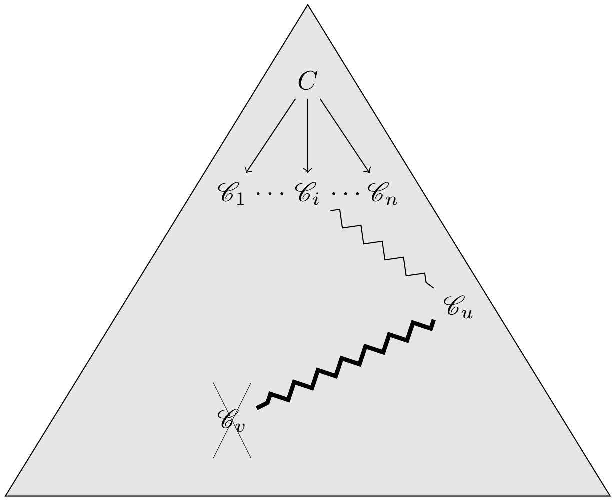

FIGURE 29. As we build $\mathfrak{T}^P{}_+(C)$ we have a path (drawn thick) in which $\mathscr{C}_u$ and $\mathscr{C}_v$ are equal. This path is therefore repetition inadmissible and $\mathscr{C}_v$ has no further descendants.

Suppose towards a contradiction that this was the case. By Theorem 4.21 any such infinite branch must be either thinning, quadratic, or superquadratic. In the thinning and quadratic cases, such a branch must either have a κ-inadmissible subpath or contain a repetition (see Section 4.8). In the superquadratic case any infinite branch must either have a tail with infinitely many C-T cycles or the tail must be μ-periodic for some base μ (see the proof of Proposition 6.15). In both of these cases such a tail will have either a C-T-inadmissible or a periodicity inadmissible initial segment. It therefore follows that $\mathfrak{T}^P{}_+(C)$ has no infinite branches. Furthermore the four inadmissibility criteria are algorithmically verifiable; thus $\mathfrak{T}^P{}_+(C) \subset \mathfrak{T}(C)$ is algorithmically constructible.

The leaves of $\mathfrak{T}^P{}_+(C)$ give a subset of all possible tracks of C (recall Section 4.1). We will now show that this subset contains a representative of every automorphically minimal κ-track.

Suppose that there was some automorphically minimal κ-track $t \subset C$ that induced a path $p : C \to \mathscr{C}_1 \to \cdots \to \mathscr{C}_l$ in $\mathfrak{T}(C)$ with $\mathscr{C}_l$ terminal, which isn't contained in $\mathfrak{T}^P{}_+(C)$. Then, by definition of $\mathfrak{T}^P{}_+(C)$, p must either contain a κ-inadmissible, a repetition inadmissible, a C-T-inadmissible, or a periodicity inadmissible subpath. Propositions 4.13, 4.17, Corollary 6.10, and Proposition 7.39 cover each of these cases and contradict the assumption that t is an automorphically minimal κ-track. It follows that the admissible leaves of $\mathfrak{T}^P{}_+(C)$ give a set of tracks

$$\left\{t_1, \ldots, t_{n(C,\kappa,S)}\right\}$$

in C that satisfy the requirements of Theorem B. $\square$

ACKNOWLEDGMENTS

I wish to thank Olga Kharlampovich for her answers to some very technical questions, enabling me to adapt the ideas in [KM98] to obtain the results of sections

6 and 7. I am grateful to Alexei Miasnikov for his numerous (ultimately successful) attempts to explain the elimination process in advanced courses and for drawing my attention to tracks. I am also grateful to Ilya Kazachkov, Montserrat Casals-Ruiz, Martin Dunwoody, Gilbert Levitt, François Dahmani, Vincent Guirardel, Daniel Groves, and Henry Wilton for discussions and encouragement. The second version of this paper was written while the author was an NSERC postdoctoral fellow at CIRGET in UQAM and at the Oxford Mathematical Institute. Although much of the third version of the paper was written on the agonizingly slow Amtrak Adirondack train that connects New York and Montreal, I feel no gratitude towards Amtrak. Finally I thank the anonymous referee for actually reading this paper and giving precise and relevant feedback. The paper is much better now.

References

[Aea91] J. M. Alonso, T. Brady, D. Cooper, V. Ferlini, M. Lustig, M. Mihalik, M. Shapiro, and H. Short, *Notes on word hyperbolic groups*, edited by Short, Group theory from a geometrical viewpoint (Trieste, 1990), World Sci. Publ., River Edge, NJ, 1991, pp. 3–63. MR1170363

[AHT06] Ian Agol, Joel Hass, and William Thurston, *The computational complexity of knot genus and spanning area*, Trans. Amer. Math. Soc. **358** (2006), no. 9, 3821–3850, DOI 10.1090/S0002-9947-05-03919-X. MR2219001

[BCRS91] Gilbert Baumslag, Frank B. Cannonito, Derek J. Robinson, and Dan Segal, *The algorithmic theory of polycyclic-by-finite groups*, J. Algebra **142** (1991), no. 1, 118–149, DOI 10.1016/0021-8693(91)90221-S. MR1125209

[BF95] Mladen Bestvina and Mark Feighn, *Stable actions of groups on real trees*, Invent. Math. **121** (1995), no. 2, 287–321, DOI 10.1007/BF01884300. MR1346208

[Bow98] Brian H. Bowditch, *Cut points and canonical splittings of hyperbolic groups*, Acta Math. **180** (1998), no. 2, 145–186, DOI 10.1007/BF02392898. MR1638764

[Bow01] B. H. Bowditch, *Peripheral splittings of groups*, Trans. Amer. Math. Soc. **353** (2001), no. 10, 4057–4082, DOI 10.1090/S0002-9947-01-02835-5. MR1837220

[Bum04] Inna Bumagin, *The conjugacy problem for relatively hyperbolic groups*, Algebr. Geom. Topol. **4** (2004), 1013–1040, DOI 10.2140/agt.2004.4.1013. MR2100689

[CRK11] Montserrat Casals-Ruiz and Ilya V. Kazachkov, *On systems of equations over free products of groups*, J. Algebra **333** (2011), 368–426, DOI 10.1016/j.jalgebra.2010.09.044. MR2785952

[Del99] Thomas Delzant, *Sur l'accessibilité acylindrique des groupes de présentation finie* (French, with English and French summaries), Ann. Inst. Fourier (Grenoble) **49** (1999), no. 4, 1215–1224. MR1703085

[DF05] Guo-An Diao and Mark Feighn, *The Grushko decomposition of a finite graph of finite rank free groups: an algorithm*, Geom. Topol. **9** (2005), 1835–1880, DOI 10.2140/gt.2005.9.1835. MR2175158

[DG08a] François Dahmani and Daniel Groves, *Detecting free splittings in relatively hyperbolic groups*, Trans. Amer. Math. Soc. **360** (2008), no. 12, 6303–6318, DOI 10.1090/S0002-9947-08-04486-3. MR2434288

[DG08b] François Dahmani and Daniel Groves, *The isomorphism problem for toral relatively hyperbolic groups*, Publ. Math. Inst. Hautes Études Sci. **107** (2008), 211–290, DOI 10.1007/s10240-008-0014-3. MR2434694

[DG10] François Dahmani and Vincent Guirardel, *Foliations for solving equations in groups: free, virtually free, and hyperbolic groups*, J. Topol. **3** (2010), no. 2, 343–404, DOI 10.1112/jtopol/jtq010. MR2651364

[DG13] François Dahmani and Vincent Guirardel, *Presenting parabolic subgroups*, Algebr. Geom. Topol. **13** (2013), no. 6, 3203–3222, DOI 10.2140/agt.2013.13.3203. MR3248731

[DS99] M. J. Dunwoody and M. E. Sageev, *JSJ-splittings for finitely presented groups over slender groups*, Invent. Math. **135** (1999), no. 1, 25–44, DOI 10.1007/s002220050278. MR1664694

[Far98] B. Farb, *Relatively hyperbolic groups*, Geom. Funct. Anal. **8** (1998), no. 5, 810–840, DOI 10.1007/s000390050075. MR1650094

[GL09] Vincent Guirardel and Gilbert Levitt, *JSJ decompositions: definitions, existence, uniqueness. I: The JSJ deformation space*, preprint, arXiv:0911.3173 (2009).

[GL11] Vincent Guirardel and Gilbert Levitt, *Trees of cylinders and canonical splittings*, Geom. Topol. **15** (2011), no. 2, 977–1012, DOI 10.2140/gt.2011.15.977. MR2821568

[GW09] Daniel Groves and Henry Wilton, *Enumerating limit groups*, Groups Geom. Dyn. **3** (2009), no. 3, 389–399, DOI 10.4171/GGD/63. MR2516172

[Hem76] John Hempel, 3-*Manifolds*, Ann. of Math. Studies, No. 86, Princeton University Press, Princeton, N. J.; University of Tokyo Press, Tokyo, 1976. MR0415619

[Hru10] G. Christopher Hruska, *Relative hyperbolicity and relative quasiconvexity for countable groups*, Algebr. Geom. Topol. **10** (2010), no. 3, 1807–1856, DOI 10.2140/agt.2010.10.1807. MR2684983

[KM98] O. Kharlampovich and A. Myasnikov, *Irreducible affine varieties over a free group. II. Systems in triangular quasi-quadratic form and description of residually free groups*, J. Algebra **200** (1998), no. 2, 517–570, DOI 10.1006/jabr.1997.7184. MR1610664

[KM05] Olga Kharlampovich and Alexei G. Myasnikov, *Effective JSJ decompositions*, Groups, languages, algorithms, Contemp. Math., vol. 378, Amer. Math. Soc., Providence, RI, 2005, pp. 87–212, DOI 10.1090/conm/378/07012. MR2159316

[Mak82] G. S. Makanin, *Equations in a free group* (Russian), Izv. Akad. Nauk SSSR Ser. Mat. **46** (1982), no. 6, 1199–1273, 1344. MR682490

[Osi06] Denis V. Osin, *Relatively hyperbolic groups: intrinsic geometry, algebraic properties, and algorithmic problems*, Mem. Amer. Math. Soc. **179** (2006), no. 843, vi+100, DOI 10.1090/memo/0843. MR2182268

[Raz87] A. Razborov, *On systems of equations in a free group*, Ph.D. thesis, Steklov Math. Institute, Moscow, 1987.

[Rom79] V. A. Roman′kov, *Universal theory of nilpotent groups* (Russian), Mat. Zametki **25** (1979), no. 4, 487–495, 635. MR534291

[Sel97] Z. Sela, *Acylindrical accessibility for groups*, Invent. Math. **129** (1997), no. 3, 527–565, DOI 10.1007/s002220050172. MR1465334

[Ser03] Jean-Pierre Serre, *Trees,* translated from the French original by John Stillwell, corrected 2nd printing of the 1980 English translation. Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2003. MR1954121

[SW79] Peter Scott and Terry Wall, *Topological methods in group theory*, Homological group theory (Proc. Sympos., Durham, 1977), London Math. Soc. Lecture Note Ser., vol. 36, Cambridge Univ. Press, Cambridge-New York, 1979, pp. 137–203. MR564422

[Tou14] Nicholas W. M. Touikan, *Bulitko's lemma for acylindrical splittings*, J. Algebra **406** (2014), 251–271, DOI 10.1016/j.jalgebra.2014.02.024. MR3188337

[Wei02] Richard Weidmann, *The Nielsen method for groups acting on trees*, Proc. London Math. Soc. (3) **85** (2002), no. 1, 93–118, DOI 10.1112/S0024611502013473. MR1901370

Department of Mathematical Sciences, Stevens Institute of Technology, Hoboken, New Jersey 07030

Email address: nicholas.touikan@gmail.com