AGT

Algebraic & Geometric Topology

Volume 25 (2025)

Cubulating drilled bundles over graphs

Mahan Mj

Biswajit Nag

msp

msp

Algebraic & Geometric Topology **25**:7 (2025) 4095–4146
DOI: 10.2140/agt.2025.25.4095
Published: 29 October 2025

Cubulating drilled bundles over graphs

MAHAN MJ
BISWAJIT NAG

We start with a Gromov-hyperbolic surface bundle E over a graph, and drill out essential simple closed curves from fibers to obtain a drilled bundle F. We prove that for such drilled bundles F, the fundamental group $\pi_1(F)$ is relatively hyperbolic with $\mathbb{Z} \oplus \mathbb{Z}$ peripheral groups. Combining the relative hyperbolicity of $\pi_1(F)$ thus obtained with a theorem of Wise, we establish virtually special cubulability of $\pi_1(F)$ provided that the maximal undrilled subbundles of F are cubulable.

20F65, 20F67; 22E40, 57M50

1. Introduction 4095
2. Bundles, drilled bundles, and graphs of groups 4100
3. Relative hyperbolicity 4101
4. Guessing geodesics 4118
5. Cubulation 4137
6. Virtual algebraic fibering 4141
References 4144

1 Introduction

In [24], Manning, Mj, and Sageev investigated the following question:

Question 1.1 *Let $1 \to H \to G \to F_n \to 1$ be an exact sequence of hyperbolic groups with F_n the free group on n generators. Is G virtually special cubulable?*

They gave sufficient conditions guaranteeing an affirmative answer. Surface-by-free groups as in Question 1.1 naturally arise as fundamental groups of surface bundles E over graphs $\mathcal{G}$. The main contribution of [24] was an explicit construction of a codimension-one quasiconvex subgroup of G along which G splits. Wise's quasiconvex hierarchy theorem [34] then furnishes cubulability of such groups. Groups G arising as in Question 1.1 may alternatively be described as a free product with amalgamation of hyperbolic groups along *nonquasiconvex hyperbolic subgroups*. For instance, if $n = 2$ and H is a hyperbolic surface group, then $G = G_1 *_H G_2$, where G_1 and G_2 are fundamental groups of hyperbolic 3-manifolds fibering over the circle. Note that H is not quasiconvex in G_1 and G_2. Thus the aim of Question 1.1 is to understand how far the quasiconvexity hypothesis in Wise's quasiconvex hierarchy theorem [34] can be relaxed. In other words, the class of hyperbolic groups arising as fundamental groups of surface bundles, as over graphs, ie G in Question 1.1, arise as a natural test case to relax the quasiconvexity hypothesis in Wise's theorem.

When $\mathcal{G}$ is a circle, Question 1.1 has an affirmative answer thanks to Kahn and Markovic's work on the surface subgroup problem [17], Bergeron and Wise's cubulability result [4], and Agol's theorem [1]. In this case, E is a hyperbolic 3-manifold M fibering over the circle. The existence of an embedded quasiconvex surface in such an M is thus not guaranteed when the first Betti number of M is one. On the other hand, for hyperbolic 3-manifolds with toroidal boundary components, the construction of embedded geometrically finite surfaces is much easier. In particular, if one drills out a simple closed curve σ from a fiber S of a fibered manifold M as above, then $M \setminus \sigma$ admits a complete hyperbolic structure by Thurston's theorem [18, Chapter 15], and each component of $S \setminus \sigma$ is geometrically finite. The fact that each component of $S \setminus \sigma$ is geometrically finite allows us to cut M along these components, and hence construct a geometrically finite hierarchy. This allows us to use Wise's relatively quasiconvex hierarchy theorem [34] even without assuming the existence of a plentiful supply of codimension-one geometrically finite subgroups.

We shall adopt the point of view that a surface bundle E over a graph $\mathcal{G}$ generalizes hyperbolic 3-manifolds fibering over the circle. We shall then create a setup where one can apply the above strategy for drilled hyperbolic 3-manifolds. Thus the main objects of study will be drilled surface bundles over graphs, where drilling corresponds to removing open neighborhoods of simple closed curves in fibers. More precisely (see Figure 1):

Definition 1.2 Let $\mathcal{G}$ be a connected graph, thought of as a 1-complex, and consider a bundle $\Pi\colon E \to \mathcal{G}$ with fiber S a surface. We refer to $\Pi\colon E \to \mathcal{G}$ as a *surface bundle over the graph* $\mathcal{G}$. The fiber $\Pi^{-1}(x)$ over $x \in \mathcal{G}$ will be denoted by S_x. If x is a vertex of $\mathcal{G}$, S_x will be called a singular fiber; otherwise it will be called a *regular fiber*.

Let $\sigma_i \subset S_{x_i}$ be a finite collection of essential simple closed curves in regular fibers S_{x_i}, so that for each regular fiber S_x, the collection of simple closed curves contained in S_x are disjoint. Let $\{N_\epsilon(\sigma_i)\}$ be a collection of small open tubular neighborhoods, missing singular fibers. We assume that:

(1) The closures $\{\overline{N_\epsilon(\sigma_i)}\}$ are disjoint.

(2) For $i \neq j$, σ_i and σ_j are not freely homotopic in E, nor is there a nontrivial free homotopy between σ_i and itself. This can equivalently be described as follows. Let $F = \big(E \setminus \bigcup_i N_\epsilon(\sigma_i)\big)$. Then any π_1-injective smooth map $(S^1 \times [0,1], S^1 \times \{0,1\}) \to \big(F, \bigcup_i \partial N_\epsilon(\sigma_i)\big)$ is homotopic, rel. boundary, into $\partial N_\epsilon(\sigma_i)$ for some i.

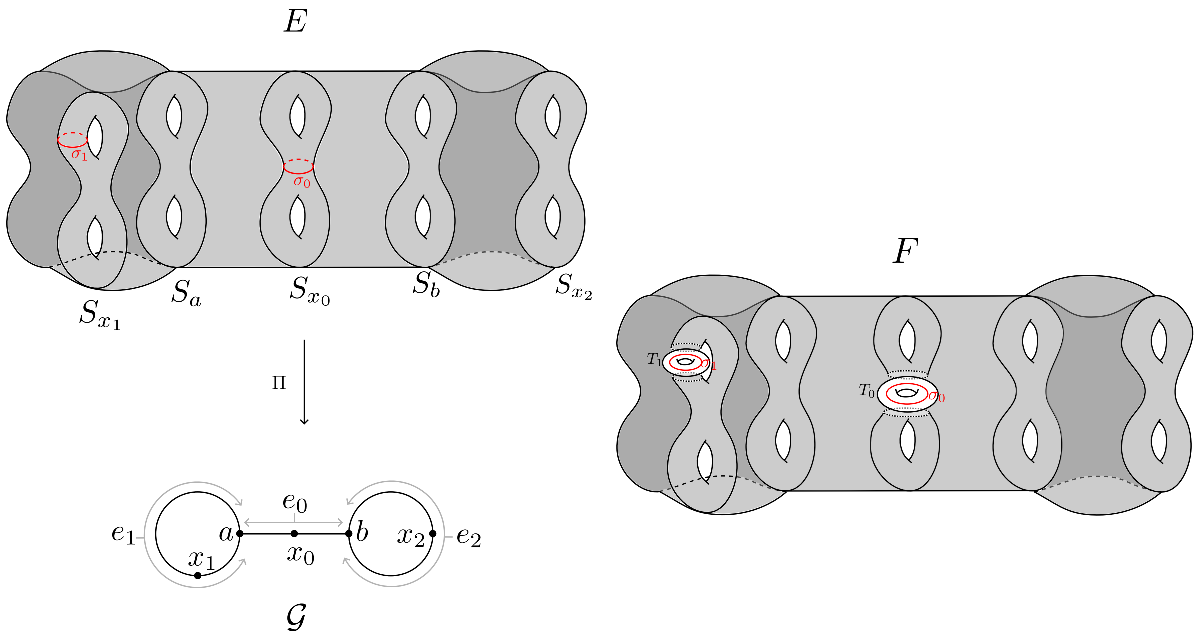

Figure 1: A drilled surface bundle over a graph: e_0 and e_1 are drilled edges, x_0 and x_1 are drilled points, S_{x_0} and S_{x_1} are drilled fibers, σ_0 and σ_1 are drilled curves, e_2 is an undrilled edge, and the closure $\bar{e}_2$ is an undrilled component. Also, S_a and S_b are singular fibers.

The complement $F = \bigl(E \setminus \bigcup_i N_\epsilon(\sigma_i)\bigr)$ will be referred to as a *drilled surface bundle over a graph*.

Drilled surface bundles over a graph are our principal objects of study. Our objective is to find sufficient conditions to cubulate them. Each torus $\{\partial N_\epsilon(\sigma_i)\}$ will be referred to as a *boundary torus* of F, and denoted by T_i. The union $\bigcup_i T_i$ will be called the *boundary* of F. The surfaces S_{x_i} (containing some σ_i) will be called *drilled fibers* and the curves σ_i will be called *drilled curves*. The points $x_i \in \mathcal{G}$ will be called *drilled points*. An edge e containing a drilled point will be called a *drilled edge*. Otherwise, we call it an *undrilled edge*.

A π_1-injective smooth map $(S^1 \times [0,1], S^1 \times \{0,1\}) \to \bigl(F, \bigcup_i \partial N_\epsilon(\sigma_i)\bigr)$ will be referred to as an *essential annulus* (see Figure 4). The second condition in Definition 1.2 thus says that F has no essential annuli.

We shall be particularly interested in the case that $\mathcal{G}$ is finite, and $\Gamma = \pi_1(E)$ is Gromov hyperbolic. In this special case, we note the further restriction on drilled curves that follows from the second condition in Definition 1.2. Let E_S be the cover of E corresponding to $\pi_1(S)$. Then we have:

Condition 1.3 *Let θ_1 and θ_2 be any two distinct elevations of the drilled curves $\sigma_i \subset E$ to E_S. Then θ_1 and θ_2 are not freely homotopic in E_S in the complement of other elevated curves. Equivalently, E_S has* ***no essential annuli*** (see Definition 3.4).

Condition 1.3 follows from Definition 1.2(2), viz that E has no essential annuli (see Definition 3.4). However, we make this explicit as much of this paper will involve E_S.

The bulk of this paper goes into establishing strong relative hyperbolicity of fundamental groups of drilled bundles (see Theorem 3.1):

Theorem 1.4 *Let E be a surface bundle over a graph $\mathcal{G}$ (see Definition 1.2) such that $\Gamma = \pi_1(E)$ is hyperbolic. Let F be a drilled surface bundle over $\mathcal{G}$ obtained by drilling E. Then $G(= \pi_1(F))$ is strongly hyperbolic relative to the collection of peripheral subgroups $\{\mathcal{P}_i = \pi_1(\partial N_\epsilon(\sigma_i))\}$.*

The drilling operation in this paper, and specifically in Theorem 1.4, is motivated, in part, by the drilling of simple closed geodesics from closed hyperbolic 3-manifolds [32, Chapter 4]. We also refer to [3; 2; 13] for related, but different, drilling constructions: in the first two, totally geodesic codimension-2 submanifolds are drilled out of CAT(-1) manifolds, and in the last, closed geodesics are drilled out of hyperbolic PD(3) groups.

With Theorem 1.4 in place, we can cut F along the components Σ of $S \setminus \bigcup_i \sigma_i$. The fundamental group $\pi_1(\Sigma)$ of each such component Σ is relatively quasiconvex in G (see Proposition 4.21). We can now apply Wise's relatively hyperbolic version of the quasiconvex hierarchy theorem [34, Theorem 15.1]. To state our main theorem, a bit more terminology needs to be set up. Let $\mathcal{K}_1, \dots, \mathcal{K}_n$ denote maximal subgraphs of $\mathcal{G}$ that contain no points x such that the fiber S_x is drilled. We refer to the $\mathcal{K}_i$ as undrilled components of $\mathcal{G}$ (see Figure 1). The restrictions of E to the undrilled components $\mathcal{K}_1, \dots, \mathcal{K}_n$ will be denoted by $E(\mathcal{K}_1), \dots, E(\mathcal{K}_n)$ of F (note that each such $E(\mathcal{K}_n)$ is naturally contained in F). Our main theorem is the following (see Theorem 5.5).

Theorem 1.5 *Let $\{E(\mathcal{K}_i)\}$ denote the restrictions of E to the undrilled components $\mathcal{K}_1, \dots, \mathcal{K}_n$ as above. If for each $\mathcal{K}_i$, $\pi_1(E(\mathcal{K}))$ is cubulable virtually special, then so is $G = \pi_1(F)$.*

Thus, if Question 1.1 has a positive answer (the undrilled case), then for the drilled groups G, cubulability follows. However, even in the absence of a definitive answer to Question 1.1, Theorem 1.5 furnishes a number of examples, as given below:

(1) when each edge of $\mathcal{G}$ contains an x such that S_x is drilled (Example 5.6),
(2) when undrilled components of $\mathcal{G}$ are either contractible or homotopy equivalent to a circle (Example 5.7),
(3) when the restrictions of E to the undrilled components satisfy the sufficient conditions in [24] (Example 5.8).[1]

Finally, we use Theorem 1.5 in conjunction with Kielak's theorem [20] to deduce that the cubulable virtually special groups virtually algebraically fiber (see Theorem 6.4).

Theorem 1.6 *Let F be drilled surface bundle over a finite graph $\mathcal{G}$ (see Definition 1.2) satisfying the hypotheses of Theorem 1.5. Then $G = \pi_1(F)$ virtually algebraically fibers.*

[1] We have removed the term "undrilled constituent" as it only occurred locally around this point only thrice.

To prove Theorem 1.6, we first establish that for any drilled surface bundle F over a graph, $\pi_1(F)$ has vanishing first l^2 Betti number (Proposition 6.3). This is done using work of Lott and Lück [21] and Fernós and Valette [12].

Scheme As indicated above, the main difficulty in establishing the cubulation Theorem 1.5 is the relative hyperbolicity result Theorem 1.4. Sections 3 and 4 are devoted to proving this result. Section 3 proves hyperbolicity of the universal cover of a *partially electrified space* (F, d_{pel}). Here, only the meridians of boundary tori in a drilled surface bundle F over $\mathcal{G}$ (see Definition 1.2) are electrified instead of the whole boundary tori to give a pseudometric d_{pel}. Theorem 3.14, the main theorem of Section 3, then shows that $(\widetilde{F}, d_{\mathrm{pel}})$ is hyperbolic (here, the lifted pseudometric is also denoted by d_{pel}).

Using Theorem 3.14, we then use a guessing geodesics argument in Section 4 and a theorem of Sisto to establish the relative hyperbolicity of G, ie Theorem 1.4.

Recurring notation Before we proceed, we collect below for the ease of reference the recurring notation we will use throughout the paper. Other notation, specific to a section, will be listed at the beginning of the relevant section.

- $\mathcal{G}$ is a finite graph. S is a closed hyperbolic surface. $\Pi\colon E \to \mathcal{G}$ is an S-bundle. H is the fundamental group of S. S_x is the fiber of Π over a point x. See Definition 1.2.
- $N_\epsilon(\sigma)$ is a small tubular neighborhood in E of a curve σ in a fiber S. $\partial N_\epsilon(\sigma)$ is the torus boundary of $N_\epsilon(\sigma)$. This torus is also denoted by T.
- F is the drilled bundle obtained by drilling E along the curves $\{\sigma_i\}$. See Definition 1.2.
- Γ denotes the fundamental group of E. G denotes the fundamental group of F.
- P_i denotes the $\mathbb{Z}^2$ subgroup of G corresponding to the torus $T_i = \partial N_\epsilon(\sigma_i)$ and $\mathcal{P}$ denotes the collection $\{P_i\}$.
- E_S is the cover of E corresponding to H. F_S is E_S but drilled along all the lifts of $\{\sigma_i\}$.
- M_r denotes a drilled atom. See Definition 3.2.
- d_{pel} is the partially electrocuted pseudometric on an appropriate space, defined in Section 3.4.
- $\mathcal{A}$ and $\mathcal{H}$ are used to denote annuli and hallways, respectively. See Definition 3.7.
- Given two lifts of fibers $\widetilde{S}_i and\ \widetilde{S}_j$ inside $\widetilde{F}$, π_{ij} denotes the projection map from the former to the latter.
- Σ denotes a maximal connected subsurface of S lying in the complement of the drilled curves.

Note to the reader The nature of the problem necessitates a fair bit of case-by-case analysis, especially in Section 4, where we deal with a host of topological and geometric objects. Some of these need to be named for quick referencing later. To make this somewhat easier to handle, we have hyperlinked many of the technical terms in subsequent appearances. Clicking on these will take the reader to the line where the term is introduced.

2 Bundles, drilled bundles, and graphs of groups

Consider the exact sequence

$$1 \to H \to \Gamma \to Q \to 1, \tag{2-1}$$

where $H = \pi_1(S)$ and $Q = \pi_1(\mathcal{G})$. Note that Γ has a graph of groups structure, where each edge and vertex group equals H and all edge-to-vertex maps are isomorphisms. Also, E has a graph of spaces structure, where each edge and vertex space equals S, and all edge-to-vertex maps are homeomorphisms. A description of $\pi_1 E$ in general may then be given as follows (see [24, Section 2] for instance). Choose a maximal tree $T \subset \mathcal{G}$. Assume, without loss of generality, that for any edge $e \subset T$, the gluing maps f_e^- are identity maps on S. Let $e_1, \dots, e_n$ be the edges in $\mathcal{G} \setminus T$. For each $i \in \{1, \dots, n\}$, write $f_i = f_{e_i}^+$. Also, set $\phi_i = (f_i)_* \colon \pi_1 S \to \pi_1 S$. Then $\pi_1 E$ is given by

$$\pi_1 E \cong \langle \pi_1 S, t_1, \dots, t_n \mid t_i^{-1} s t_i = \phi_i(s),\ \forall s \in \pi_1 S,\ i \in \{1, \dots, n\} \rangle.$$

Hyperbolic Γ It was shown by Hamenstädt [15] (see also [19; 28]) that, in the exact sequence (2-1), Γ is hyperbolic if and only if Q is a convex cocompact subgroup of the mapping class group in the sense of Farb and Mosher [11]. A useful fact that we will need is the following Scott–Swarup type theorem [29] due to Dowdall, Kent, and Leininger [9, Theorem 1.3]; see [26] for a different proof of the same result.

Theorem 2.1 *Let*

$$1 \to H \to \Gamma \to Q \to 1$$

be an exact sequence as above, so that Γ is hyperbolic, and $H = \pi_1(S)$. Then any finitely generated infinite-index subgroup of H is quasiconvex in Γ.

A graph of groups structure on $G = \pi_1(F)$

Recall that F is the drilled bundle obtained from the bundle $E \to \mathcal{G}$. We now describe a standard graph of groups description for $\pi_1(F)$. We shall define $G := \pi_1(F)$. Denote the barycentric subdivision of $\mathcal{G}$ by $\mathcal{G}^*$.

The vertex set $\mathcal{V}(\mathcal{G}^*)$ thus consists of

(1) edges e of $\mathcal{G}$, and

(2) vertices v of $\mathcal{G}$.

The edge set $\mathcal{E}(\mathcal{G}^*)$ of $\mathcal{G}^*$ is given by incidences between edges e and vertices v of $\mathcal{G}$, ie if the terminal vertex e^+ of $e \in \mathcal{G}$ is v, then we introduce an edge in $\mathcal{G}^*$ between e and v. Similarly for e^-.

Then $G (= \pi_1(F))$ has a graph of groups structure (see Figure 2), where:

(1) The underlying graph is $\mathcal{G}^*$.

(2) The edge groups are all equal to $H = \pi_1(S)$.

(3) The vertex groups G_e (indexed by $e \in \mathcal{G}$) are given as follows. Since the bundle over the interior of $e \subset \mathcal{G}$ is $S \times (0, 1)$, G_e equals $\pi_1(M_e)$, where M_e is a possibly drilled copy of $S \times [0, 1]$.

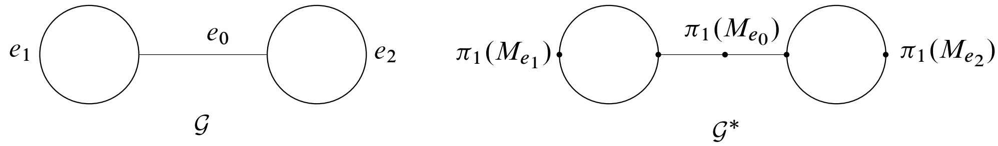

Figure 2: The graph of groups $\mathcal{G}^*$ for a given $\mathcal{G}$.

(4) The vertex groups G_v (indexed by the vertices v of $\mathcal{G}$) are all H.

(5) In particular, the edge-group to vertex-group maps are injective.

A convenient way to think of the graph of spaces associated to the above graph of groups description is as follows. For any edge $e \in \mathcal{G}$, the corresponding vertex space over $e \in \mathcal{G}^*$ is M_e, the possibly drilled copy of $S \times [0,1]$. For any vertex $v \in \mathcal{G}$, the corresponding vertex space over $v \in \mathcal{G}^*$ is $S \times \mathrm{star}(v)$, where $\mathrm{star}(v)$ denotes a small closed neighborhood of v in $\mathcal{G}$. The edge spaces are all given by S, and edge-to-vertex space inclusions are given by the inclusions of S into either a drilled copy of $S \times [0,1]$ or a copy of $S \times \mathrm{star}(v)$.

Let $P_i = \pi_1(T_i) = \mathbb{Z} + \mathbb{Z}$. Then there exists e such that $P_i \subset G_e = \pi_1(M_e)$. Whenever we are dealing with G equipped with the above graph of groups structure, we shall *implicitly* reindex the collections $\{T_i\}$ and $\{\mathbb{P}_i\}$ so that the collection of tori contained in M_e are indexed as $\{T_e^j\}$. (This notation will be used in Section 4 at the beginning of which we shall recall it.)

Lemma 2.2 *Suppose that Γ is hyperbolic. Let σ_i denote a finite collection of curves that are drilled from E to obtain F. Let $\mathcal{Q}_i \subset \Gamma$ denote the* (*conjugacy class of the*) *cyclic subgroup corresponding to σ_i. Then Γ is strongly hyperbolic relative to the collection $\{\mathcal{Q}_i\}$.*

Proof Since the σ_i are simple closed curves, and since Γ is hyperbolic, they denote primitive elements of Γ. Hence the collection $\{\mathcal{Q}_i\}$ denotes a malnormal quasiconvex family of subgroups. By a theorem of Bowditch [7, Theorem 7.11], Γ is strongly hyperbolic relative to the collection $\{\mathcal{Q}_i\}$. □

3 Relative hyperbolicity

Notation to be used in this section:

- Π_S is the projection $F_S \to \widetilde{\mathcal{G}}$.
- $\mathcal{A}_S$ denotes an essential annulus (see Definition 3.7) in F_S. Whenever applicable, B_S denotes the unique atomic wrapping annulus (see Definition 3.5) in $\mathcal{A}_S$. $\mathcal{A}_S^+$ and $\mathcal{A}_S^-$ are the two subannuli of $\mathcal{A}_S$ in the complement of B_S. The intersections of $\mathcal{A}_S^\pm$ with B_S are denoted by $\theta^\pm$.
- In the context of Theorem 3.17, Y is a graph of spaces with universal cover X with the tree of spaces structure $\Pi\colon X \to \mathcal{T}$. Then ρ, λ, H, and $2m$ stand respectively for constants for thinness, hyperbolicity, girth, and length of a hallway or an annulus.
- $\widetilde{\Pi}\colon \widetilde{F}_S \to \mathcal{T}$ is the tree of spaces structure for $\widetilde{F}_S$. X_v is used to denote vertex spaces of both X and $\widetilde{F}_S$.

We refer the reader to [10; 7] for generalities on relative hyperbolicity.

This section and the next are devoted to proving the following:

Theorem 3.1 *If Γ is hyperbolic, then G is (strongly) hyperbolic relative to the collection $\{P_i\}$.*

Theorem 3.1 says roughly that the result of drilling a hyperbolic 3-complex fibering over a graph gives a relatively hyperbolic 3-complex. This result is along the lines of earlier work of Belegradek and Hruska [3] (see also [2]), who prove similar results in a manifold context. The proof of Theorem 3.1 occupies the rest of this section and the next.

3.1 Scheme of the proof of Theorem 3.1

The proof of Theorem 3.1 consists of two steps as indicated at the end of the introduction.

Step 1 (Section 3) (hyperbolicity of $(\widetilde{F}, d_{\mathrm{pel}})$) The aim of this section is to prove Theorem 3.14: hyperbolicity of the partially electrified space $(\widetilde{F}, d_{\mathrm{pel}})$. Since the logic behind the proof has a number of ingredients involved, we lay out a sketch below.

Partial electrification (Section 3.4) We refer the reader to Section 3.4 for the precise notion of partial electrification. The construction that is relevant to this paper is the following. For each boundary torus T of F, there is a distinguished meridional direction, so that T is foliated by these meridians. Partial electrification electrifies each such meridian. This might be easier to see in the universal cover $\widetilde{T}$, where the elevations of each meridian are electrified. Since $\widetilde{T}$ may be identified with $\mathbb{R}^2$ foliated by parallel copies of $\mathbb{R}$ corresponding to elevations of the meridian, it follows that after electrifying each such copy of $\mathbb{R}$ in $\widetilde{T}$, the latter becomes quasi-isometric to $\mathbb{R}$, a hyperbolic space. This is in contrast with the usual electrification operation, where all of $\widetilde{T}$ is electrified to a space quasi-isometric to a point.

Decomposition into atoms We note next that E and F are built out of 3-manifolds in a natural way. The following definition captures this decomposition.

Definition 3.2 (drilled and undrilled atoms) Recall that $\Pi\colon E \to \mathcal{G}$ denotes the surface bundle E over the graph $\mathcal{G}$. Let $M_e = \Pi^{-1}(\bar{e}) \cap F$. If e is a drilled edge, we refer to M_e as a *drilled atom* of F; otherwise we refer to it as an *undrilled atom* of F (see Figure 3). Elevations of drilled (resp. undrilled) atoms of F to $\widetilde{F}$ will be referred to as drilled (resp. undrilled) atoms of $\widetilde{F}$.

If $\mathcal{L}$ is a maximal subgraph of $\mathcal{G}$, such that each edge of $\mathcal{L}$ is undrilled, $E(\mathcal{L}) = \Pi^{-1}(\mathcal{L}) \cap F$ will be called a *maximal undrilled subbundle* of F (as well as of E). Elevations of such $E(\mathcal{L})$ to $\widetilde{E}_r$ or $\widetilde{F}$ will be called *elevated maximal undrilled subbundles*.

Atoms and a combination theorem In Section 3.2 we shall see that work of Thurston ensures that drilling out geodesics in hyperbolic 3-manifolds gives new noncompact 3-manifolds with a complete hyperbolic structure of finite volume. These drilled 3-manifolds are the atoms used in building the drilled

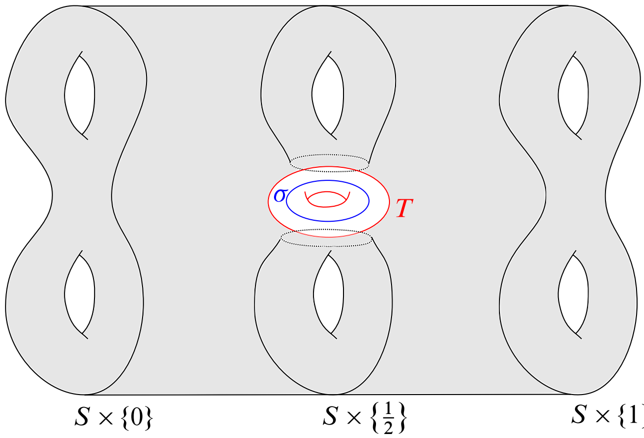

Figure 3: $S \times [0, 1]$ drilled along σ, or in other words, a drilled atom.

bundle F. It is not hard to see (Lemma 3.12) that after partial electrification, each atom has hyperbolic universal cover. It remains then to combine the partially electrified atoms using the underlying graph $\mathcal{G}^*$ to prove Theorem 3.14.

For this, we fall back upon the classical Bestvina–Feighn combination theorem (see Section 3.5.1) in Section 3.5 to piece the partially electrified atoms together. We shall use the annuli flare condition of Bestvina and Feighn to implement this.

It should be borne in mind here that the known combination theorems in the specific context of relative hyperbolicity such as [27] are inadequate for the purposes of this paper. This is the basic reason for using the more elaborate method of first proving a hyperbolic combination theorem in this section and then using it in the next using Sisto's generalization of Bowditch's guessing geodesics lemma to prove relative hyperbolicity.

Annuli In order to apply the Bestvina–Feighn combination theorem however, one needs to identify the essential annuli in F (see Definition 3.7). This is what is done in Sections 3.2 and 3.3. Of these, Section 3.2 identifies essential annuli in atoms (see Definition 3.4), and Section 3.3 concatenates them together.

After having identified essential annuli in F (see Definition 3.7), the combination theorem is applied to the partially electrified metric in Section 3.5.2 to prove Theorem 3.14.

Step 2 (Section 4) (guessing geodesics) The subsections of Section 4 follow a more straightforward linear logical order than Section 3. Our starting point is a necessary and sufficient condition for relative hyperbolicity due to Sisto building on earlier work of Bowditch and Hamenstädt. We recall this in Section 4.1.

Next, using the hyperbolic geodesics obtained as an output of Theorem 3.14 in Section 3, we guess a family of paths in $\widetilde{F}$ in Section 4.2. The key idea used here is that of electroambient quasigeodesics from [25]: a geodesic in the partially electrified $(\widetilde{F}, d_{\mathrm{pel}})$ can be lifted in a more or less canonical way to a path in $\widetilde{F}$

(with the ordinary, nonelectrified metric). In each atom, the lifted path is a genuine quasigeodesic. This follows from the fact that each drilled atom (see Definition 3.2) has a universal cover that is hyperbolic *relative* to the elevations of its boundary tori. However, relative hyperbolicity of $\widetilde{F}$ turns out to be considerably more difficult to establish. We need to piece together the electroambient quasigeodesics in atoms carefully to guess a path family in $\widetilde{F}$. A fair bit of case-by-case analysis is necessary to get the path family. We therefore provide the reader with a short guide through the cases in Section 4.2.1.

Section 4.3 then establishes that the path family constructed in Section 4.2 satisfies the usual property of quasigeodesics in a relatively hyperbolic space. With the stability property in place, Sections 4.4 and 4.5 check the conditions of Section 4.1 needed to establish relative hyperbolicity. Of these subsections, the conditions checked in Section 4.4 are fairly routine, whereas Section 4.5 is devoted to proving the analog of the thin triangles property. This completes the proof of Theorem 3.1.

Finally, Section 4.6 establishes the relative quasiconvexity of essential subsurfaces of drilled fibers.

3.2 Drilling 3-manifolds

As mentioned in Section 3.1, the aim of this subsection is to define, identify, and classify essential annuli (see Definition 3.4) in atoms (cf Definition 3.2). We refer the reader to Figure 4 for a quick idea of the various annuli that may arise.

Let M be either $S \times I$ or a hyperbolic 3-manifold fibering over the circle with fiber S. Thus, M fibers over a compact 1-manifold (possibly with boundary) with fiber S. We proceed to drill simple closed curves σ_i in some finite collection of fibers. A fiber S' will be referred to as *undrilled* if S' does not intersect any of the neighborhoods $N_\epsilon(\sigma_i)$ of the drilled curves. Let M_r denote the drilled manifold. Henceforth, in this paper, M_r will be equipped with

(1) a complete hyperbolic structure of finite volume when M fibers over the circle,

(2) a geometrically finite hyperbolic structure with convex boundary otherwise.

In the latter case, we shall think of M_r as the quotient of the convex hull of a geometrically finite representation $\rho\colon \pi_1(M_r) \to \mathrm{PSL}_2(\mathbb{C})$ such that the parabolics of $\rho(\pi_1(M_r))$ correspond precisely to the boundary tori of M_r. Since no two of the drilled curves are freely homotopic in M by Definition 1.2, it follows that M_r is atoroidal. Since M_r is an atoroidal Haken manifold, the existence of such hyperbolic structures in both cases follows from Thurston's theorem [18, Chapter 15].

Convention We shall choose hyperbolic structures as above on M_r and fix them for the rest of the discussion.

For such a hyperbolic structure on M_r, M_r is finite volume without boundary if M fibers over the circle. Otherwise the boundary consists of two copies of S, corresponding to the (singular) fibers (see Definition 1.2) over $0, 1 \in I$. Further, M_r has finitely many rank-two cusps (corresponding to the toroidal boundaries of regular neighborhoods of drilled curves). We start with the following observation:

Lemma 3.3 *Let M and M_r be as above where the set of drilled curves correspond to a finite nonempty collection $\sigma_i \in S_i$ of simple closed curves on fibers S_i. Any undrilled fiber S' is geometrically finite in M_r. Any connected components of $S_x \setminus \sigma$ of a drilled fiber is also geometrically finite in M_r.*

Proof This follows from the covering theorem [32, Theorem 9.2.2; 8]. □

Essential annuli in atoms We now describe the essential annuli in atoms.

Definition 3.4 An *immersed essential annulus* or simply an *essential annulus* $\mathcal{A}$ in a 3-manifold M with nonempty boundary is an immersion $i : (\mathcal{A}, \partial\mathcal{A}) \to (M, \partial M)$ such that

(1) $i_* : \pi_1(\mathcal{A}) \to \pi_1(M)$ is injective,

(2) i is not homotopic rel. boundary into ∂M.

When M is a 3-manifold fibering over the circle with a distinguished singular fiber S_x (see Definition 1.2), an *immersed essential annulus* or simply an *essential annulus* $\mathcal{A}$ is an immersion $i : (\mathcal{A}, \partial\mathcal{A}) \to (M, S_x)$ such that

(1) $i_* : \pi_1(\mathcal{A}) \to \pi_1(M)$ is injective,

(2) i is not homotopic rel. boundary into S_x.

Henceforth, we shall refer to immersed essential annuli simply as essential annuli. In the 3-manifold literature, essential annuli often refer to embedded essential annuli in $(M, \partial M)$. However, since we shall not have any special use for embedded annuli, we shall use the terminology from Definition 3.4 in this paper. The usage is consistent with the more group-theoretic notion in [5]; see Definition 3.7 below. Note that essential annuli $\mathcal{A}$ in fibered M are allowed to intersect S_x in the interior of $\mathcal{A}$ as well, ie $\mathcal{A}$ is allowed to "wrap around transverse to the fibers" of M multiple times so long as $\partial\mathcal{A}$ maps to S_x. Essential annuli in undrilled atoms (see Definition 3.2) are given by the following, *up to homotopy*:

(1) If $M = S \times I$, then essential annuli are of the form $\sigma \times I$, where σ is an essential, possibly immersed curve in S. After homotopy, we may assume that σ is a geodesic in some auxiliary hyperbolic structure on S.

(2) If M fibers over the circle, let $M_{\mathbb{Z}}$ denote the cover of M corresponding to $\pi_1(S)$, so that $M_{\mathbb{Z}}$ is homeomorphic to $S \times \mathbb{R}$. Then essential annuli in $M_{\mathbb{Z}}$ are concatenations of annuli $\mathcal{A}_i \subset S \times [i, i+1]$ as in item (1). Further, $\sigma_{i+1} = \mathcal{A}_i \cap S \times \{i+1\} = \mathcal{A}_{i+1} \cap S \times \{i+1\}$, so that $\mathcal{A}_i$ and $\mathcal{A}_{i+1}$ may be concatenated along the essential, possibly immersed curve σ_{i+1}.

Essential annuli in drilled atoms (see Definition 3.2) are a bit more involved. Let $M = S \times I$, and M_r be a drilled atom obtained from M (see Definition 3.2). Let $\sigma_1, \dots, \sigma_m$ denote the drilled curves on S. Realize the σ_i by geodesics in an auxiliary hyperbolic structure on S. Let Σ_0 denote the subsurface of S filled by $\bigcup_i \sigma_i$, ie adjoin to $\bigcup_i \sigma_i$ all simply connected complementary regions. Then essential annuli in M_r (see Definition 3.4) are of three kinds after homotopy:

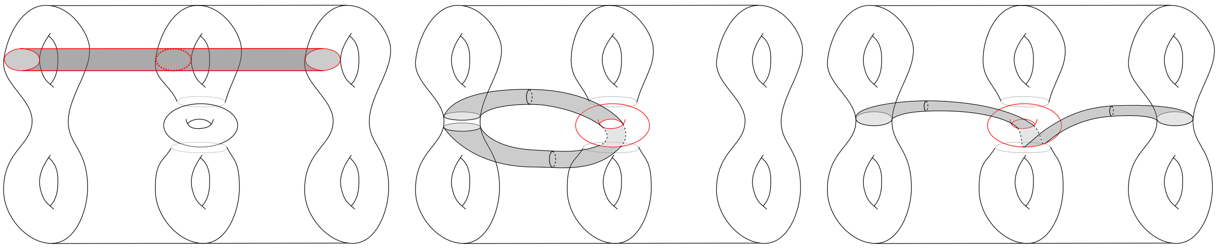

Figure 4: The three types of annuli in a drilled $S \times I$.

(1) essential annuli of the form $\gamma \times I$, respecting the product structure of $S \times I$, where γ is an essential, possibly immersed curve in S; in this case, γ necessarily lies in (a component of) $S \setminus \Sigma_0$,

(2) essential annuli with both boundary curves on either $S \times \{0\}$ or $S \times \{1\}$ with core curve homotopic to a multiple of one of the σ_i; in this case, $\mathcal{A}$ wraps around $N_\epsilon(\sigma_i)$ finitely many times,

(3) essential annuli with one boundary curve on $S \times \{0\}$, and one on $S \times \{1\}$ with core curve homotopic to a multiple of one of the σ_i, such that $\mathcal{A}$ wraps around $N_\epsilon(\sigma_i)$ finitely many times.

Definition 3.5 An annulus from case (2) above is referred to as an *atomic wrapping backtracking annulus*. An annulus from case (3) is referred to as an *atomic wrapping nonbacktracking annulus*.

Next, suppose that M fibers over the circle with monodromy Φ. Let $S_x \subset M$ denote the unique singular fiber (see Definition 1.2). Let $S_1, \dots, S_m$ denote drilled fibers and $\sigma_i \subset S_i$ for $i = 1, \dots, m$ denote drilled simple closed curves. Let M_r denote M after drilling. Let $M_{\mathbb{Z}}$ denote the cover of M corresponding to $\pi_1(S)$ and M_S denote the elevation of M_r to $M_{\mathbb{Z}}$. Let $\dots, S_x^{-1}, S_x^0, S_x^1, \dots$ denote the elevations of S_x to M_S. Without loss of generality, assume that the elevated annulus $\mathcal{A}_1$ starts on S_x^0. Also, let Σ_n denote the subsurface of S filled by $\bigcup_{i=1,\dots,m} \bigcup_{j=0,\dots,n-1} \Phi^j(\sigma_i)$, ie adjoin to $\bigcup_{i=1,\dots,m} \bigcup_{j=0,\dots,n-1} \Phi^j(\sigma_i)$ all simply connected complementary regions of S.

Lemma 3.6 *There exists $N \in \mathbb{N}$ such that $\Sigma_n = S$ for all $n \geq N$.*

Proof Since Φ is pseudo-Anosov, there exists $N \in \mathbb{N}$ such that

$$d_{\mathcal{C}(S)}(\sigma_i, \Phi^n(\sigma_i)) \geq 3$$

for all $n \geq N$, where $d_{\mathcal{C}(S)}$ denotes distance in the curve complex of S. Hence, σ_i and $\Phi^n(\sigma_i)$ fill S for all for all $n \geq N$. □

Essential annuli (see Definition 3.4) in M_r for M a fibered 3-manifold will be described in greater detail and in a more general setting in Section 3.3 below. For now, it suffices to say that an essential annulus in M_r lifts to an essential annulus in M_S starting and ending on S_x^i and S_x^j for some $i, j \in \mathbb{Z}$.

3.3 Essential annuli in F

As mentioned in Section 3.1, the aim of this subsection is to concatenate the atomic essential annuli (see Definition 3.4) of Section 3.2 above to obtain essential annuli in F (see Definition 3.7).

We use the classification of essential annuli (see Definition 3.7) in M_r described above to identify essential annuli in the drilled bundle F. Since $F \subset E$, an essential annulus $\mathcal{A} \subset F$ may also be regarded as a subset of E. The formalism we use is due to Bestvina and Feighn [5].

Definition 3.7 [5, page 87] Let $\mathcal{G}$ be a graph, and Y a graph of spaces with base graph $\mathcal{G}$, so that the maps of edge-spaces to vertex spaces are injective at the level of the fundamental group. Let $X = \widetilde{Y}$ be the universal cover, and $\mathcal{T}$ the resulting tree of spaces, whose vertex and edge spaces are universal covers of vertex and edge spaces of Y. Let m be a positive integer and I denote the closed unit interval. A *hallway* in X is a map $\mathcal{H}\colon [-m, m] \times I \to X$ with the following properties:

1. $\mathcal{H}^{-1}(\cup X_e) = \{-m, -m+1, \dots, m\} \times I$,
2. $\mathcal{H}$ is transverse to $\bigcup X_e$ relative to the previous condition,
3. for each $i \in \{-m, -m+1, \dots, m\}$, the image of $\mathcal{H}$ restricted to $\{i\} \times I$ is a geodesic in the corresponding edge space.

We say that such a hallway has *length* $2m$. The *girth* of the hallway is the length of the curve $\mathcal{H}(\{0\} \times I)$. A hallway is *essential* if the projection of $\mathcal{H}$ onto the base tree $\mathcal{T}$ is a path which does not backtrack.

A map $\mathcal{A}\colon [-m, m] \times S^1 \to Y$ is an *annulus* if the lift $\tilde{\mathcal{A}}\colon [-m, m] \times I \to X$ is a hallway. An annulus $\mathcal{A}$ is *essential* if the hallway $\tilde{\mathcal{A}}$ is essential.

An essential annulus $\mathcal{A}\colon [0, m] \times S^1 \to Y$ is said to *start at* Y_v if $\mathcal{A}(\{0\} \times S^1) \subset Y_v$.

We now relate the notion of essential annuli (see Definition 3.4) in 3-manifolds with those in Definition 3.7 using Definition 3.2. Recall from Section 2 that, F admits a graph of spaces structure, with underlying graph $\mathcal{G}^*$, where the vertex spaces of F are drilled or undrilled atoms (see Definition 3.2). Let $\widetilde{\mathcal{G}}$ denote the universal cover of $\mathcal{G}^*$.

Let F_S be the cover of F corresponding to the kernel of the map $\Pi_* : \pi_1(F) \to \pi_1(\mathcal{G})$. Let $\Pi_S : F_S \to \widetilde{\mathcal{G}}$ denote the projection map from F_S to $\widetilde{\mathcal{G}}$. Then F_S admits a graph of spaces structure, where:

1. Each vertex space is a copy of $S \times I$ with, possibly, some curves drilled. The interior of each vertex space is given by one of the following:
   - $M \setminus S_x$ if M is an undrilled atom of F fibering over the circle (see Definition 3.2),
   - $M_r \setminus S_x$ if M_r is a drilled atom of F (see Definition 3.2) fibering over the circle,
   - $M \setminus S \times \{0, 1\}$ if M is an undrilled atom of F fibering over I (see Definition 3.2),
   - $M_r \setminus S \times \{0, 1\}$ if M_r is a drilled atom of F (see Definition 3.2) fibering over I.
2. Each edge space is S.

Thus, F_S is a drilled graph bundle over the graph $\widetilde{\mathcal{G}}$ (see Definition 1.2). Note that E_S also has a similar graph of spaces decomposition over $\widetilde{\mathcal{G}}$ where all the vertex spaces are copies of $S \times I$, and each edge space is S. With this structure in mind, F_S embeds into E_S as a graph of spaces over $\widetilde{\mathcal{G}}$. We now describe essential annuli in F_S in terms of the description of in atoms given in Section 3.2.

Let $\mathcal{A}$ denote an essential annulus in F (see Definition 3.7). Let $\mathcal{A}_S$ denote an elevation of $\mathcal{A}$ to F_S. Then the image $\Pi_S(\mathcal{A}_S)$ is an unparametrized geodesic $[a,b]$ in $\widetilde{\mathcal{G}}$. Let $a = v_0, v_1, \dots, v_n = b$ denote the vertices of $[a,b]$. Let F_S^i denote the vertex space of F_S corresponding to v_i. Then $\mathcal{A}_S$ is given by a concatenation of essential annuli in the atoms F_S^i. We shall show now that $\mathcal{A}_S$ comes broadly in two flavors: backtracking and nonbacktracking.

Definition 3.8 Suppose $\mathcal{A}_S$ is an essential annulus (see Definition 3.7) of the form $\gamma \times [p,q]$, respecting the topological product structure of F_S, such that

(1) γ is an essential, possibly immersed curve in S,

(2) $\gamma \times \{0\}$ lies on S_x^0, and $\gamma \times \{n\}$ lies on S_x^n.

Then we shall refer to such an essential annulus (see Definition 3.7) as a *nonbacktracking annulus*.

In this case, σ necessarily lies in (a component of) $S \setminus \Sigma_n$ (for instance, by the annulus theorem). In this case, there are no atomic wrapping annuli in the sense of Section 3.2 contained in $\mathcal{A}_S$.

Next, suppose $\mathcal{A}_S$ is an essential annulus (see Definition 3.7) containing an atomic wrapping annulus in the sense of Definition 3.5. Recall that E_S denotes the surface bundle over $\widetilde{\mathcal{G}}$ (see Definition 1.2) corresponding to the cover of E corresponding to $\pi_1(S)$, so that F_S is obtained from E_S by drilling a family of nonhomotopic simple closed curves. Then, there exists a drilled curve η in E_S such that $\mathcal{A}_S$ wraps around the boundary of $N_\epsilon(\eta)$ in F_S. Hence, the core curve of $\mathcal{A}_S$ is homotopic to a nontrivial power of η. We observe the following:

Lemma 3.9 *Suppose $\mathcal{A}_S$ is an essential annulus (see Definition 3.7) containing an atomic wrapping annulus (see Definition 3.5). Then $\mathcal{A}_S$ contains exactly one atomic wrapping annulus.*

Proof Since $\mathcal{A}_S$ is an essential annulus containing an atomic wrapping annulus, it has core curve freely homotopic to a (nontrivial power of a) drilled curve η. Note next that, by Definition 1.2 and Condition 1.3, no two distinct elevations of drilled curves to E_S are freely homotopic. Hence, $\mathcal{A}_S$ contains exactly one atomic wrapping annulus. □

Essential annuli (see Definition 3.7) $\mathcal{A}_S$ in F_S containing an atomic wrapping annulus (see Definition 3.5) are of two kinds, depending on the nature of the atomic wrapping annulus. Let $B_S \subset \mathcal{A}_S$ denote the unique atomic wrapping annulus contained in $\mathcal{A}_S$ with core curve homotopic to η^m for $m \neq 0$.

(1) Both boundary curves of B_S lie on a single singular fiber (see Definition 1.2) S_y. In this case, B_S, and hence $\mathcal{A}_S$, wraps around $N_\epsilon(\eta)$ finitely many times. (Note that B_S is an atomic wrapping backtracking annulus in the sense of Definition 3.5.) We refer to $\mathcal{A}_S$ as an *annulus with backtracking*.

(2) There exists a drilled atom M_e (see Definition 3.2) with boundary surfaces S_y and S_z, such that the distinct boundary curves of B_S lie on the distinct boundary surfaces S_y and S_z. Again, B_S, and hence $\mathcal{A}_S$, wraps around $N_\epsilon(\eta)$ finitely many times, but otherwise is required to respect the product structure on M_e. (Note that B_S is an atomic wrapping, nonbacktracking annulus tracking annulus.) We shall refer to $\mathcal{A}_S$ as an annulus *with wrapping but no backtracking*.

Definition 3.10 An essential annulus (see Definition 3.7) $\mathcal{A}_S$ in F_S is said to *intersect* a drilled atom (see Definition 3.2) $M_r \subset F_S$ if $\mathcal{A}_S$ contains an elementary annulus $B_S \subset M_r$. (Note that B_S may be a nonbacktracking annulus (see Definition 3.8), an atomic wrapping backtracking annulus, or an atomic wrapping nonbacktracking annulus (see Definition 3.5).)

Lemma 3.11 *There exists $D \in \mathbb{N}$ such that the following holds: Let $\mathcal{A}_S \subset F_S \subset E_S$ be an essential annulus intersecting (see Definition 3.10) drilled atoms M_a and M_b (see Definition 3.2), where $a, b \in \widetilde{\mathcal{G}}$. Then*

$$d_{\widetilde{\mathcal{G}}}(a,b) \le D.$$

Proof This is similar to Lemma 3.6. It was shown by Kent and Leininger [19] and Hamenstädt [15] that the following are equivalent:

(1) $\pi_1(E)$ is hyperbolic,

(2) $\pi_1(\mathcal{G})$ acts on the curve complex $\mathcal{C}(S)$ with qi-embedded orbits,

(3) $\pi_1(\mathcal{G})$ is a convex-cocompact subgroup of the mapping class group $\mathrm{MCG}(S)$ in the sense of [11].

Since $\pi_1(E)$ is hyperbolic by assumption, it follows that $\pi_1(\mathcal{G})$ acts on the curve complex $\mathcal{C}(S)$ with qi-embedded orbits. Let $i: \widetilde{\mathcal{G}} \to \mathcal{C}(S)$ be the induced qi-embedding. Hence, there exists $D \in \mathbb{N}$ such that if $d_{\widetilde{\mathcal{G}}}(a,b) > D$, then $d_{\mathcal{C}(S)}(i(a), i(b)) > 3$. Let σ_a and σ_b denote drilled curves in M_a and M_b, respectively, so that σ_a and σ_b fill S.

If the core curve of $\mathcal{A}_S$ intersects both M_a and M_b in the sense of Definition 3.10, it must be disjoint from the subsurface of S filled by σ_a and σ_b, an impossibility. □

3.4 Partial electrification

Recall that the bundle $\Pi: E \to \mathcal{G}$ restricts to a map $\Pi_r: F \to \mathcal{G}$.

Partial electrification of $\partial N_\epsilon(\sigma_i)$ Let M_r denote a drilled atom in F (see Definition 3.2), and let $\partial N_\epsilon(\sigma_i)$ denote the boundary of a drilled curve in M_r. Choose a homeomorphism of $\partial N_\epsilon(\sigma_i)$ with $S^1 \times S^1$, and assume, without loss of generality, that each $S^1 \times \{t\}$ is a meridian. Equip $S^1 \times S^1$ with a product metric, where the first factor is given the zero metric, and the second factor the standard round metric of radius one. We refer to the resulting path-pseudometric on $\partial N_\epsilon(\sigma_i)$ as *the partially electrified path-pseudometric* and denote it by d_{pel}.

The reader at this stage might wish to refer back to Section 3.1 (paragraph titled "Partial electrification") for the heuristic idea behind the notion of partial electrification introduced above. Further, the paragraph titled "Atoms and a combination theorem" in Section 3.1 says briefly why we need this technique.

Partial electrification of drilled atoms Recall that any drilled M_r of F has been equipped with the restriction of a complete hyperbolic metric. If M_r is obtained by drilling $S \times I$, then the surface boundary components $S \times \{0, 1\}$ are convex. Removing a small neighborhood of the cusps we obtain boundary tori $\partial N_\epsilon(\sigma_i)$ equipped with flat Euclidean metrics. Abusing notation slightly, we continue to refer to the resulting compact 3-manifold with boundary also as M_r. Rescaling the hyperbolic metric if necessary, we may assume that $\partial N_\epsilon(\sigma_i)$ has a product metric as in the previous paragraph. Replace each such metric by the partially electrified path-pseudometric d_{pel} described in the previous paragraph. We now consider a family of paths, each of which is given by a concatenation of pieces that either

(1) have interior disjoint from the boundary tori $\{\partial N_\epsilon(\sigma_i)\}$ (the length of such a piece is given by its hyperbolic length), or

(2) lie entirely in some boundary torus $\{\partial N_\epsilon(\sigma_i)\}$ (the length of such a path is given by its length with respect to d_{pel} on $\{\partial N_\epsilon(\sigma_i)\}$).

Then the length of a path is given by the sum of the above pieces. The resulting path-pseudometric on M_r is referred to as *the partially electrified path-pseudometric on* M_r and is also denoted by d_{pel}.

If the atom M_r of F is obtained from an atom M of E that fibers over the circle, then any elevation M_S of M_r to F_S is a cyclic cover of M_r corresponding to the natural epimorphism to $\mathbb{Z}$ inherited from M. Then M_S is a concatenation of a $\mathbb{Z}$'s worth of atoms of F_S. Each atom M_e of F_S is equipped with the inherited path metric from M_S, and the resulting path-pseudometric on M_e is referred to as *the partially electrified path-pseudometric on* M_e and is also denoted by d_{pel}.

The universal cover of M_r will be denoted by $\widetilde{M}_r$. The lift of the partially electrified path-pseudometric on M_r to $\widetilde{M}_r$ is referred to as *the partially electrified path-pseudometric on* $\widetilde{M}_r$ and is also denoted by d_{pel}. Similarly for M_e.

Lemma 3.12 *There exist $\delta \geq 0$, $K \geq 1$, and $\epsilon \geq 0$ such that the following hold: Let M_e denote an atom of F_S. Let S_x denote a surface boundary component of M_e. Let $\mathcal{P}$ denote the collection of elevations of the tori $\{\partial N_\epsilon(\sigma_i)\}$ to $\widetilde{M}_e$. Then:*

(1) *$\widetilde{M}_e$ is strongly hyperbolic relative to $\mathcal{P}$.*

(2) *$(\widetilde{M}_e, d_{\mathrm{pel}})$ is δ-hyperbolic.*

(3) *The inclusion of $\widetilde{S}_x$ into $(\widetilde{M}_e, d_{\mathrm{pel}})$ is a (K, ϵ)-qi-embedding for any elevation $\widetilde{S}_x$ of S_x.*

Proof Since a hyperbolic structure on M_e may be chosen so that it has convex boundary, $\widetilde{M}_e$ is strongly hyperbolic relative to $\mathcal{P}$. (This is a consequence of the main theorem of [10].) The second conclusion is then a special case of [25, Lemma 1.20].

Let d denote the metric on $\widetilde{M}_e$ lifted from the intrinsic path-metric on M_e. Let d_e denote the electric metric on $\widetilde{M}_e$ after electrifying the elements of $\mathcal{P}$ as in [10]. For $u, v \in \widetilde{S}_x$, let γ_{uv}, γ_{uv}^e, and γ_{uv}^p denote geodesics with respect to d, d_e, and d_{pel}, respectively. The second conclusion will follow from two facts:

(1) By [25, Lemma 1.21; 10, Lemma 4.5 and Proposition 4.6], γ_{uv}, γ_{uv}^e, and γ_{uv}^p track each other (uniformly, independent of u and v) away from $\mathcal{P}$. (See Lemma 4.4 below for a slightly more general statement.)

(2) The nearest-point projections of elements of $\mathcal{P}$ equipped with d_{pel} onto $\widetilde{S}_x$ are either uniformly bounded in diameter, or uniform quasi-isometric embeddings.

In fact, any $P \in \mathcal{P}$ is stabilized by a conjugate of $\pi_1(N_\epsilon(\sigma_i)) = \mathbb{Z}^2$ for some i. Let $\mathrm{stab}(P)$ denote the stabilizer of P. Then, $\mathrm{stab}(P) \cap \widetilde{S}_x$ is either trivial or infinite cyclic.

In the former case, γ_{uv}^e (after a small isotopy if necessary) may be assumed to be disjoint from P. In the latter case, if γ_{uv}^e does enter and exit P at y and z, respectively, then there exist $y_1, z_1 \in \widetilde{S}_x$ such that the geodesic joining y_1 and z_1 lies uniformly close to an elevation of σ_i. It follows that γ_{uv} and γ_{uv}^p track each other throughout their lengths. The third conclusion follows. □

Partial electrification of F and $\widetilde{F}$ The metric on each of the atoms after drilling (and before partial electrification) is denoted by d. Equip F with a C^0 piecewise Riemannian metric that is bi-Lipschitz to the metric d on each of the atoms. We refer to this metric on F also by d. Now, consider rectifiable paths $\sigma\colon [0,1] \to F$ consisting of finitely many pieces, each of which is contained in an atom. The length of σ is declared to be the sum of the lengths of these subpaths. Replacing d on each atom by the partially electrified metric d_{pel} on that atom, we obtain a partially electrified path pseudometric, also denoted by d_{pel}, on F.

Lifting d and d_{pel} to the universal cover $\widetilde{F}$, we obtain a metric d and a partially electrified path pseudometric d_{pel}, respectively. The distance between a pair of point u, v is then obtained by passing to an infimum over all paths σ as above joining u and v.

Remark 3.13 The partial electrification construction above is adapted from [27] (see [25, Section 1.3] for a summary).

3.5 Partial electrification and relative hyperbolicity

As mentioned in Section 3.1, the aim of this subsection is to use the classification of essential annuli from Sections 3.2 and 3.3 and the Bestvina–Feighn combination theorem (Theorem 3.17 below) to prove Theorem 3.14 below. Note that the statement of Theorem 3.14 involves the partially electrified metric defined in Section 3.4.

Let $\widetilde{F}$ denote the universal cover of F. We lift the pseudometric d_{pel} to $\widetilde{F}$ to obtain a partially electrified metric on $\widetilde{F}$ denoted again by d_{pel}. The following is the main theorem of this section.

Theorem 3.14 *$(\widetilde{F}, d_{\mathrm{pel}})$ is hyperbolic.*

To prove Theorem 3.14, we shall use the Bestvina–Feighn combination theorem (Theorem 3.17).

3.5.1 The Bestvina–Feighn combination theorem

Definition 3.15 [5] Let $\mathcal{G}$ be a graph, and Y a graph of spaces with base graph $\mathcal{G}$, so that the maps of edge-spaces to vertex spaces are injective at the level of the fundamental group. Let $X = \widetilde{Y}$ be the universal cover, and $\mathcal{T}$ the resulting tree of spaces, whose vertex and edge spaces are universal covers of vertex and edge spaces of Y. Suppose that the following hold:

(1) Vertex spaces $\{X_v\}$ and edge spaces $\{X_e\}$ are all δ-hyperbolic metric spaces for some $\delta > 0$.

(2) There exist $K \geq 1$ and $\epsilon \geq 0$ such that the maps of edge-spaces to vertex spaces for X are all (K, ϵ)-quasi-isometric embeddings.

Then Y is said to be a graph of hyperbolic metric spaces satisfying the qi-embedded condition.

Recall the notion of hallways and annuli from Definition 3.7.

Definition 3.16 [5, page 87] For $\lambda > 1$, a hallway $\mathcal{H}$ is said to be λ-*hyperbolic* if

$$l(\mathcal{H}(\{0\} \times I)) \leq \frac{1}{\lambda} \max\bigl\{l(\mathcal{H}(\{-m\} \times I)), l(\mathcal{H}(\{m\} \times I))\bigr\},$$

where l denotes the length of the path in the corresponding edge space.

Let $\rho > 0$. Given $i \in \{-m, -m+1, \dots, m\}$, denote the vertex space that $\mathcal{H}([i, i+1] \times I)$ lies in by X_{v_i}. The hallway is ρ-*thin* if for all such i and for any $t \in I$, $d_{X_{v_i}}(\mathcal{H}(i, t), \mathcal{H}(i+1, t)) < \rho$.

The *girth* (resp. *length*) of the annulus $\mathcal{A}$ is the girth (resp. *length*) of the induced hallway $\tilde{\mathcal{A}}$.

Similarly, the rest of the terminology, ie hyperbolicity, thinness, essentiality, for the annulus $\mathcal{A}$, is defined via $\tilde{\mathcal{A}}$.

The annuli flare condition The graph of spaces Y (with base graph $\mathcal{G}$) satisfies the *annuli flare condition* if there exist $\lambda > 1$ and $m \geq 1$ such that the following holds: given any $\rho > 0$, there exists a constant $H(\rho)$ so that whenever $\mathcal{A}$ is a ρ-thin essential annulus (see Definition 3.7) of length $2m$ and girth at least $H(\rho)$, $\mathcal{A}$ is λ-hyperbolic. The graph of spaces Y satisfies the *weak annuli flare condition* if there are numbers $\lambda > 1, m > 1$, and H such that any 4δ-thin essential annulus (see Definition 3.7) of length $2m$ and girth at least H is λ-hyperbolic. The following statement gives the Bestvina–Feighn combination theorem in a consolidated form.

Theorem 3.17 [5, Theorem 3.2 of the correction] *Let $\Pi\colon Y \to \mathcal{G}$ be a graph of spaces whose vertex and edge spaces have δ-hyperbolic universal covers for some $\delta > 0$. If $\Pi\colon Y \to \mathcal{G}$ satisfies the following conditions, then the universal cover X of Y is hyperbolic*:

(1) *the quasi-isometrically embedded condition* (*see Definition 3.15*),

(2) *the annuli flare condition or the weak annuli flare condition.*

The following definition adapts [23, Definition 4.26] to our context.

Definition 3.18 We say that a vertex space Y_v flares in all directions weakly if there are numbers $\lambda > 1$, $m > 1$, and H such that if $\mathcal{A}\colon S^1 \times [0,m] \to Y$ is any 4δ-thin essential annulus (see Definition 3.7) of length m and girth at least H *starting at* Y_v (in the sense of Definition 3.7), then the associated lifted hallway $\mathcal{H}$ satisfies

$$l(\mathcal{H}(\{0\} \times I)) \le \frac{1}{\lambda} l(\mathcal{H}(\{m\} \times I)).$$

We shall need a modification of Theorem 3.17 to guarantee global quasiconvexity of a vertex space. The following now adapts [23, Proposition 4.27] to our context.

Corollary 3.19 *Let $\Pi\colon Y \to \mathcal{G}$ be a graph of spaces satisfying the conditions of Theorem 3.17. Further, let $Y_v \subset Y$ be a vertex space that flares in all directions (see Definition 3.18) weakly. Then $\widetilde{Y}_v$ is quasiconvex in $\widetilde{Y}$.*

Proof Since the context of [23, Proposition 4.27] is slightly different, we sketch the mild modifications necessary. We recall a construction from [22, Section 3] here. It follows from [22, Theorem 3.8] that there exists $C > 0$ such that the following holds. Given any geodesic $\lambda \subset (\widetilde{Y}_v, d_v)$, there exists a "ladder" $\mathcal{L}_\lambda \subset \widetilde{Y}$ containing λ, such that $\mathcal{L}_\lambda$ is C-quasiconvex. Hyperbolicity of $(\widetilde{Y}, d)$ now guarantees that for all such geodesics $\lambda \subset (\widetilde{Y}_v, d_v)$, $\mathcal{L}_\lambda$ is hyperbolic. The construction of $\mathcal{L}_\lambda$ in [22, Section 3] now shows that $\widetilde{\Pi}\colon \widetilde{Y} \to \widetilde{\mathcal{G}}$ restricts to $\widetilde{\Pi}_\lambda\colon \mathcal{L}_\lambda \to \mathcal{T}$, where $\mathcal{T} \subset \widetilde{\mathcal{G}}$ is a tree. Further, for any vertex $w \in \mathcal{T}$, $\widetilde{\Pi}_\lambda^{-1}(w)$ is a geodesic segment in the vertex space $\widetilde{\Pi}^{-1}(w) \subset \widetilde{Y}$. Thus, $\mathcal{L}_\lambda$ is a tree of spaces, where each vertex space is isometric to an interval. The hypothesis that $Y_v \subset Y$ be a vertex space that flares in all directions weakly guarantees that $\mathcal{L}_\lambda$ flares in all directions also. Hence λ is a quasigeodesic (with uniform constants) in $\mathcal{L}_\lambda$. Since $\mathcal{L}_\lambda$ is C-quasiconvex, λ is a quasigeodesic (with uniform constants) in $\widetilde{Y}$. □

Remark 3.20 More generally, if $Z \subset Y_v$ is a subspace such that

(1) the inclusion induces an injection at the level of fundamental groups,

(2) $\widetilde{Z}$ is qi-embedded in $\widetilde{Y}_v$,

then an auxiliary vertex w and an edge $e = [w, v]$ may be added to the base graph $\mathcal{G}$, so that $Y_w = Y_e = Z$. Definition 3.18 and Corollary 3.19 may thus be applied to such subspaces Z of Y_v as well. If Z flares in all directions (see Definition 3.18) weakly, then $\widetilde{Z}$ is quasiconvex in $\widetilde{Y}$ by Corollary 3.19.

We also record the following, where we assume implicitly that there is a cocompact group action so that the annuli flare condition makes sense. (We also refer the reader to [5, page 90 and Section 4 of the correction] and for an analogous hallways flare condition.)

Lemma 3.21 *Let $\Pi_{\mathcal{T}}\colon X \to \mathcal{T}$ be a tree of spaces obtained as a universal cover of a graph of compact spaces. Suppose that each vertex space X_v and edge space X_e of X is δ-hyperbolic. Further, suppose that the following conditions are satisfied:*

(1) *the quasi-isometrically embedded condition* (*see Definition 3.15*),

(2) *the annuli flare condition or the weak annuli flare condition.*

Let $\mathcal{T}_0$ be a subtree of $\mathcal{T}$ and $X_0 = \Pi_{\mathcal{T}}^{-1}(\mathcal{T}_0)$. Then X_0 is hyperbolic.

Proof We note that any hallway in X_0 is also a hallway in X. In particular, $\Pi_{\mathcal{T}}\colon X_0 \to \mathcal{T}_0$ is a tree of spaces satisfying

(1) the quasi-isometrically embedded condition (see Definition 3.15),

(2) the annuli flare condition or the weak annuli flare condition.

The corollary is now a direct consequence of Theorem 3.17. □

3.5.2 Hyperbolicity of partially electrified bundle: proof of Theorem 3.14 To prove Theorem 3.14, it suffices to check the two conditions of Theorem 3.17.

Hyperbolicity of vertex spaces and the quasi-isometrically embedded condition This follows from Lemma 3.12.

Identifying ρ-thin annuli It remains to prove the annuli flare condition. We recall the description of essential annuli (see Definition 3.7) in F_S from Section 3.3. Let D be as in Lemma 3.11. We choose m in the annuli flaring condition so that $2m > D$. Hence, any essential annulus (see Definition 3.7) $\mathcal{A}_S$ in F_S can intersect (see Definition 3.10) *at most one drilled atom* (see Definition 3.2). Thus, any essential annulus (see Definition 3.7) $\mathcal{A}_S$ in F_S of length $2m$ can be of exactly one of the following three mutually exclusive types:

Case 1 $\mathcal{A}_S$ is a nonbacktracking annulus (see Definition 3.8) with core curve having free homotopy type distinct from any of the drilled curves,

Case 2 $\mathcal{A}_S$ contains an atomic wrapping nonbacktracking annulus B_S. Here, B_S wraps around $\partial N_\epsilon(\sigma)$ for some σ. The core curve of $\mathcal{A}_S$ is then freely homotopic to a (nontrivial power of) σ.

Case 3 $\mathcal{A}_S$ contains an atomic wrapping backtracking annulus B_S wrapping around $\partial N_\epsilon(\sigma)$ for some σ. The core curve of $\mathcal{A}_S$ is then freely homotopic to a (nontrivial power of) σ.

In cases 2 and 3 above, $\mathcal{A}_S$ is the concatenation of three pieces:

(1) The first is the atomic wrapping annulus B_S (with or without backtracking); see Definition 3.5.

(2) The other two are nonbacktracking annuli (see Definition 3.8) $\mathcal{A}_S^{\pm}$, such that $\mathcal{A}_S^{\pm} \cap B_S$ consist of curves $\theta^{\pm}$ that are freely homotopic in F_S. If B_S is an atomic wrapping nonbacktracking annulus, then there exist distinct singular fibers (see Definition 1.2) $S_x^{\pm}$ of F_S (bounding the atom of F_S containing B_S) such that $\theta^{\pm} \subset S_x^{\pm}$. If B_S is an atomic wrapping backtracking annulus (see Definition 3.5), then there exists a single singular fiber (see Definition 1.2) S_x of F_S (a boundary component of the atom of F_S containing B_S) such that $\theta^{\pm} \subset S_x$, and $\theta^{\pm}$ are freely homotopic within S_x.

Since B_S wraps around $\partial N_\epsilon(\eta)$ for some drilled curve η, and since $\partial N_\epsilon(\eta)$ is an elevation of one of finitely many tori in F, the core curve of $\mathcal{A}_S$ is the same as the core curve of B_S, and hence is of the form γ^n for some $n \in \mathbb{N}$, where γ is one of the drilled curves in E.

Checking the annuli flare condition **Case 1** ($\mathcal{A}_S$ is a nonbacktracking annulus (see Definition 3.8) with core curve having free homotopy type distinct from any of the drilled curves) We start with the following converse to the Bestvina–Feighn combination theorem.

Proposition 3.22 *E_S satisfies the weak annuli-flare condition.*

Proof Since E is hyperbolic, this is a special case of [28, Proposition 5.8], where it is shown that E satisfies a flaring condition. This is equivalent to hyperbolicity of hallways in E_S and implies the weak annuli-flare condition. □

Corollary 3.23 *Nonbacktracking annuli* (see Definition 3.8) *with core curve having free homotopy type distinct from any of the core curves satisfy the annuli-flare condition*: *more precisely*, *there exist* $\lambda > 1$ *and* $m > 1$ *such that if* $\mathcal{A}_S$ *is a nonbacktracking annulus* (see Definition 3.8) *with girth at least* 1, *it satisfies the weak annuli flare condition.*

Proof Note that $\mathcal{A}_S \subset F_S \subset E_S$. Since $\mathcal{A}_S$ is a nonbacktracking annulus (see Definition 3.8) in F_S, it is an essential annulus (see Definition 3.7) in E_S. Proposition 3.22 now gives the desired conclusion. □

Cases 2 and 3 ($\mathcal{A}_S$ is a wrapping annulus with core curve freely homotopic to a power of one of the drilled curves) We shall give a unified proof of these two cases. Let B_S denote the atomic wrapping annulus (see Definition 3.5) contained in $\mathcal{A}_S$. Let σ denote the drilled curve such that B_S wraps around $\partial N_\epsilon(\sigma)$. Choose an orientation of σ and $n \in \mathbb{N}$ such that the core curve of B_S (and hence that of $\mathcal{A}_S$) is freely homotopic to σ^n. Let M_e denote the drilled atom (see Definition 3.2) of F_S containing B_S. Let $S_x^\pm$ denote the boundary components of M_e, and $\mathcal{A}_S^\pm$ denote the two components of $\mathcal{A}_S \setminus \operatorname{Int}(B_S)$. Then each of the annuli $\mathcal{A}_S^\pm$ is an essential annulus (see Definition 3.7) in E_S with one boundary curve in $S_x^+ \cup S_x^-$. (If $\mathcal{A}_S$ is backtracking, then both boundary curves lie in the same surface boundary component. If $\mathcal{A}_S$ is without backtracking, then the boundary curves lie in different surface boundary components.) Let $\theta^\pm$ denote the boundary curve of $\mathcal{A}_S^\pm$ on $S_x^+ \cup S_x^-$.

It will help to explicate the special case where $n = 1$, ie the core curve of $\mathcal{A}_S$ is freely homotopic to σ. Then each of $\mathcal{A}_S^\pm$ is a flaring annulus, with $l(\theta^\pm)$ uniformly close to the girth of $\mathcal{A}_S$. This follows from the fact that the length $l(\theta^\pm)$ is uniformly bounded. Hence, in the formulation of Definition 3.16,

$$l(\mathcal{H}(\{0\} \times I)) \le \frac{1}{\lambda} l(\mathcal{H}(\{m\} \times I)),$$

for all m large enough (ie we can ignore the negative direction of the hallway from Definition 3.16). Further, note that the d_{pel}-length of the annulus B_S is uniformly bounded. This is because the meridian of $\partial N_\epsilon(\sigma)$ that B_S wraps around has length zero. Hence the concatenation $\mathcal{A}_S^+ \cup B_S \cup \mathcal{A}_S^-$ satisfies the weak annuli flare condition.

For general $n \in \mathbb{N}$, any elevation of $\theta^{\pm}$ (freely homotopic to σ^n) to the universal cover $\widetilde{E}_S$ gives a uniform quasigeodesic. This follows from the fact that σ is an elevation of one of the (finitely many) drilled curves. Hence, again, $l(\theta^{\pm})$ is uniformly close to the girth of $\mathcal{A}_S$. Again, $\mathcal{A}_S^{\pm}$ satisfies the one-sided flare condition

$$l(\mathcal{H}(\{0\} \times I)) \leq \frac{1}{\lambda} l(\mathcal{H}(\{m\} \times I))$$

in the formulation of Definition 3.16. The same argument from the previous paragraph now shows that the concatenation $\mathcal{A}_S^{+} \cup B_S \cup \mathcal{A}_S^{-}$ satisfies the weak annuli flare condition.

Thus, the sufficient conditions of Theorem 3.17 are satisfied by essential annuli $\mathcal{A}_S$ (see Definition 3.7) in F_S, equipped with partially electrified pseudometric d_{pel} (see Section 3.4). Theorem 3.14 follows. □

Remark 3.24 One place where the partial electrification metric d_{pel} is essential in the above proof is to conclude that the d_{pel}-length of B_S is uniformly bounded.

3.5.3 Consequences of Theorem 3.14 The proof of Theorem 3.14 above gives some additional information that we shall need below. Let $\widetilde{\Pi}\colon \widetilde{F}_S \to \mathcal{T}$ be the tree of spaces for the universal cover $\widetilde{F}_S$ $(=\widetilde{F})$ of F_S. Let $\mathcal{T}_0 \subset \mathcal{T}$ denote a subtree. Let $\widetilde{\Pi}^{-1}(\mathcal{T}_0) = X_0$, so that $\widetilde{\Pi}\colon X_0 \to \mathcal{T}_0$ is a tree of spaces. The following definition collects together some notions that will be used in Corollary 3.26 below.

Definition 3.25 Let $v \in \mathcal{T}_0$ be a vertex such that

(1) $\widetilde{M}_v$ is an atom,

(2) $\widetilde{S}_x$ is a boundary component of $\widetilde{M}_v$ that is not contained in the boundary of any other vertex space of $\widetilde{\Pi}\colon X_0 \to \mathcal{T}_0$.

Then we say that $\widetilde{M}_v$ is a *boundary atom of* X_0 and $\widetilde{S}_x$ is a *boundary component of* X_0. If $\widetilde{M}_v$ is drilled (resp. undrilled), we say that $\widetilde{M}_v$ is a *drilled* (*resp. undrilled*) *boundary atom of* X_0 and $\widetilde{S}_x$ is a *drilled* (*resp. undrilled*) *boundary component of* X_0.

Corollary 3.26 *Let X_0 be as above. Then X_0 is hyperbolic, and any drilled boundary component $\widetilde{S}_x$ of X_0 is d_{pel}-quasiconvex. Here d_{pel} denotes the restriction of the partially electrified pseudometric d_{pel}* (see Section 3.4) *from X.*

Proof Hyperbolicity of X_0 follows from Lemma 3.21 after noting that we have checked the weak annuli flare condition for F_S. Next, we attach an auxiliary vertex space $\widetilde{S}_x \times I$ to X_0 along $\widetilde{S}_x$ and an auxiliary edge e to $\mathcal{T}_0$ to v to get a tree of spaces $X_0 \cup_{\widetilde{S}_x} \widetilde{S}_x \times I \to \mathcal{T}_0 \cup_v e$. Here, $\widetilde{S}_x \times \{1\} \subset \widetilde{S}_x \times I$ is attached to X_0 along $\widetilde{S}_x \subset X_0$, so that $\widetilde{S}_x \subset X_0$ becomes the edge-space X_e corresponding to the new edge e. Further, let v_0 denote the extra vertex introduced in $\mathcal{T}_0 \cup_v e$ (corresponding to $\{0\} \in I$). Thus, the new vertex space $\widetilde{S}_x \times I$ is X_{v_0}. To prove d_{pel}-quasiconvexity of $\widetilde{S}_x$, it suffices, by Corollary 3.19, to show that essential annuli starting on X_{v_0} (in the sense of Definition 3.7) flare in all directions in the sense of

Definition 3.18. But such annuli are simply essential annuli (see Definition 3.7) in the bundle $E \to \mathcal{G}$ (before drilling). Further, their core curves lie in $S_x \setminus \sigma_i$ for one of the finitely many drilled curves σ_i. The conclusion now follows from Theorem 2.1, which guarantees that each of these finitely many proper essential subsurfaces of S_x have uniformly quasiconvex elevations, and hence that any essential annulus (see Definition 3.7) starting on such a subsurface flares in all directions. □

Another consequence of Corollary 3.19 that we record is the following.

Corollary 3.27 *Let $\mathcal{T}_1 \subset \mathcal{T}$ denote a subtree of finite diameter (but not necessarily locally finite). Let $\widetilde{\Pi}^{-1}(\mathcal{T}_1) = X_1$, so that $\widetilde{\Pi}\colon X_1 \to \mathcal{T}_1$ is a tree of spaces. Then X_1 is hyperbolic, and for any singular fiber S_x* (see Definition 1.2) *with $\widetilde{S}_x \subset X_1$, $\widetilde{S}_x$ is d_{pel}-quasiconvex in X_1.*

Proof Since any essential hallway has length bounded by the diameter of $\mathcal{T}_1$, hyperbolicity of X_1 follows from Theorem 3.17. The same reason guarantees d_{pel}-quasiconvexity of $\widetilde{S}_x$ by Corollary 3.19. □

Lemma 3.28 *Let X_1, S_x, and $\widetilde{S}_x$ be as above. Then $\widetilde{S}_x$ with its intrinsic metric is qi-embedded in X_1. Also, $\widetilde{S}_x$ is properly embedded in $(\widetilde{F}, d_{\mathrm{pel}})$.*

Proof Suppose that S_x is a boundary component of an atom M_e (drilled or undrilled) of F_S (in the sense of Definition 3.25). By Lemma 3.12, $\widetilde{S}_x$ is qi-embedded in $\widetilde{M}_e$ equipped with d_{pel} (the case where M_e is undrilled is obvious). Let $\widetilde{M}_e$ denote the vertex space X_v for v a vertex of $\mathcal{T}$. Next, let $\mathcal{T}_1 \subset \mathcal{T}$ denote a subtree of finite diameter (but not necessarily locally finite) containing v. Then Corollary 3.27 shows that $\widetilde{S}_x$ is d_{pel}-quasiconvex in $\widetilde{\Pi}^{-1}(\mathcal{T}_1) = X_1$.

A reprise of the proof of Lemma 3.12 now shows that $\widetilde{S}_x$, equipped with its intrinsic metric is qi-embedded in X_1. Since the finite diameter of $\mathcal{T}_1$ was arbitrary, the second conclusion follows. □

Corollary 3.29 *There exist $K \geq 1$ and $\epsilon > 0$ such that the following hold. Let γ be a geodesic in $(\widetilde{F}, d_{\mathrm{pel}})$. Let $\widetilde{S}_x$ be the elevation of a singular fiber bounding a drilled atom (see Definition 3.2) $\widetilde{M}_r$ of $\widetilde{F}$. Suppose further that u and v are entry and exit points on $\widetilde{S}_x$ of γ into and out of $\widetilde{M}_r$, respectively. Let $\gamma \mid [u, v]$ denote the subpath of γ between u and v, and $\beta(u, v)$ the geodesic in $\widetilde{S}_x$ (with its intrinsic metric) between u and v. Let*

$$\gamma'(u, v) = (\gamma \setminus \gamma \mid [u, v]) \cup \beta(u, v)$$

be obtained from γ by replacing $\gamma \mid [u, v]$ by $\beta(u, v)$. Then $\gamma'(u, v)$ is a (K, ϵ)-quasigeodesic in $(\widetilde{F}, d_{\mathrm{pel}})$.

Proof Let X_1 denote the closure of the component of $X \setminus \widetilde{S}_x$ containing $\widetilde{M}_r$. By Corollary 3.26, $\widetilde{S}_x$ is quasiconvex in (X_1, d_{pel}) with uniform constants (independent of $\widetilde{S}_x$ and X_1). Hence, there exists D such that $\gamma \mid [u, v]$ lies in a D-neighborhood of $\widetilde{S}_x$ in (X_1, d_{pel}). Hence, by Lemma 3.28, $\beta(u, v)$ is a (K_1, ϵ_1)-quasigeodesic in (X_1, d_{pel}). The corollary follows. □

Corollary 3.29 allows us to replace d_{pel}-geodesics by uniform d_{pel}-quasigeodesics that do not backtrack from drilled atoms (see Definition 3.2) in the following sense.

Definition 3.30 A d_{pel}-quasigeodesic γ' in $(\widetilde{F}, d_{\text{pel}})$ is said to be a d_{pel}-quasigeodesic *without backtracking from drilled atoms* (see Definition 3.2) if it satisfies the following: if γ' enters a drilled atom $\widetilde{M}_r$ (see Definition 3.2) through a boundary component $\widetilde{S}_1$, then it can only leave $\widetilde{M}_r$ through a boundary component $\widetilde{S}_2 \neq \widetilde{S}_1$.

Corollary 3.29 now allows us to observe that for any d_{pel}-geodesic γ, there exists a (K, ϵ)-quasigeodesic γ' without backtracking from drilled atoms (see Definition 3.2) joining the endpoints of γ. In fact, γ' is obtained from γ by

(1) carrying out replacements of all backtracking segments in drilled atoms (see Definition 3.2),

(2) isotoping the resulting path slightly to make it disjoint from $\widetilde{S}_x$ as in Corollary 3.29.

The paragraph titled "Checking the annuli flare condition" in the proof of Theorem 3.14 gives the following further conclusion.

Corollary 3.31 *There exists $L \geq 1$ such that the following holds: Let $[a, b] \subset \widetilde{\mathcal{G}}$ denote a geodesic of length at least L such that M_a and M_b are drilled atoms* (see Definition 3.2) *of F_S. Let $M_{[a,b]}$ denote the 3-manifold given by $\Pi_S^{-1}([a, b])$. Also, let S_a and S_b denote the boundary components of $M_{[a,b]}$* (see Definition 3.25). *Let $\pi_1(S_a), \pi_1(S_b)$ denote the subgroups of $\pi_1(M_{[a,b]})$ carried by S_a and S_b. Then $\pi_1(S_a) \cap \pi_1(S_b) = \{1\}$.*

Proof Since M_a and M_b are drilled atoms (see Definition 3.2), $\widetilde{S}_a$ and $\widetilde{S}_b$ are d_{pel}-quasiconvex in $\widetilde{M}_{[a,b]}$, equipped with the d_{pel}-metric by Corollary 3.26. Further, the d_{pel}-quasiconvexity constant is uniform, independent of a and b.

If $\pi_1(S_a) \cap \pi_1(S_b) \neq \{1\}$, then there exists a loop $\alpha_a \subset S_a$ freely homotopic to an $\alpha_b \subset S_b$. By uniform d_{pel}-quasiconvexity, any geodesic in the free homotopy class of α_a (resp. α_b) must lie close to S_a (resp. S_b). This forces the existence of an L such that $d_{\widetilde{\mathcal{G}}}(a, b) < L$. □

4 Guessing geodesics

We refer the reader back to the paragraph titled "Step 2" in Section 3.1 for a summary of the contents of the various subsections of this section and the logical order we follow.

Notation to be used in this section:

- Given a hyperbolic drilled atom M_r (see Definition 3.2), λ, $\hat{\lambda}$, λ_{ea}, and λ_p denote in $\widetilde{M}_r$ respectively the geodesic, the electric geodesic, the electroambient quasigeodesic (see Definition 4.3), and the d_{pel}-geodesic, all with the same endpoints.

- We use P to denote a lift of a boundary torus of F to $\widetilde{F}$, and $\mathcal{P}$ is used to denote the collection of all such lifts. $\mathcal{P}_r$ denotes the collection of such lifts contained in $\widetilde{M}_r$.

- For z lying on a drilled curve σ in M_r, m_z is a specially chosen meridian of $\partial N_\epsilon(\sigma)$.

- For two lifts of boundary surfaces $\widetilde{S}_1$ and $\widetilde{S}_2$ in $\widetilde{M}_r$, Z_1 and Z_2 denote respectively the image of the projection of $\widetilde{S}_2$ onto $\widetilde{S}_1$ and vice versa.
- For $[a,b]$ a geodesic in the tree $\mathcal{T}$ for the tree of spaces $\widetilde{F}_S \to \mathcal{T}$, $X_{[a,b]}$ is the subtree of spaces over $[a,b]$.
- $\mathcal{P}_{[a,b]}$ is the collection of lifts of $\{\partial N_\epsilon(\sigma_i)\}$ in $X_{[a,b]}$.
- $\mathcal{B}$ denotes an elevation of a maximal undrilled subbundle (see Definition 3.2) of F_S.
- $\mathcal{F}$ is the path family defined in Section 4.2.7 to which the guessing geodesics lemma (Theorem 4.2) is applied. Paths in this family are denoted by η.
- A is used to denote undrilled atoms (see Definition 3.2), and B is used for *undrilled molecules*.

4.1 The guessing geodesics lemma

We shall need a necessary and sufficient condition for relative hyperbolicity due to Sisto [30] building on earlier work of Bowditch [6] and Hamenstädt [16].

Definition 4.1 Let (X, d_X) be a geodesic metric space, and $\mathcal{P}$ a collection of subsets. The collection $\mathcal{P}$ is said to be *mutually bounded* if for each $K \geq 0$, there exists B such that $\operatorname{diam}(N_K(P) \cap N_K(Q)) \leq B$, for all $P \neq Q \in \mathcal{P}$.

In [30], Sisto refers to the mutual boundedness criterion above as condition (α_1). We shall need the following:

Theorem 4.2 [30, Theorem 4.2] *Let (X, d_X) and $\mathcal{P}$ be as above. Suppose that for all $x, y \in X$ we are given*

(1) *a path $\eta(x,y)$ connecting them,*
(2) *a closed subset $\theta(x,y) \subset \eta(x,y)$.*

Suppose that there exists $D \geq 0$ such that the following are satisfied:

(1) *If $d_X(x,y) \leq 2$ then $\operatorname{diam}(\theta(x,y)) \leq D$.*
(2) *Let d_H denote Hausdorff distance. Then for all $x', y' \in \eta(x,y)$, we have*

$$d_H(\theta(x',y'), \theta(x,y)|[x',y'] \cup \{x',y'\}) \leq D,$$

where $\theta(x,y)|[x',y'] = \theta(x,y) \cap \eta(x,y)|[x',y']$.

(3) *For all $x, y, z \in X$, $\theta(x,y) \subset N_D(\theta(x,z) \cup \theta(z,y))$.*
(4) *If $x', y' \in \eta(x,y)$ do not both lie on the same $P \in \mathcal{P}$, then there exists $z \in \theta(x,y)$ between x' and y'.*
(5) *The elements of $\mathcal{P}$ are mutually bounded.*

(6) *For all* $k \geq 0$, *there exists* $K \geq 0$ *such that the following holds. If for some* $P \in \mathcal{P}$
- $d_X(x, P) \leq k$,
- $d_X(y, P) \leq k$, *and*
- $d_X(x, y) \geq K$,

then
- $\theta(x, y) \subset B_K(x) \cup B_K(y)$, *and*
- *there exists* $z \in \theta(x, y) \cap N_D(P)$.

Then X *is strongly hyperbolic relative to* $\mathcal{P}$.

Theorem 4.2 says roughly that if we can guess a family of paths in X that satisfy the conditions required of geodesics in a relatively hyperbolic space, then X itself is relatively hyperbolic.

4.2 Path families

Recall that $\mathcal{P}$ denotes the collection of elevations of $\partial N_\epsilon(\sigma_i)$ in $\widetilde{F}$, where σ_i ranges over the finitely many drilled curves in E. Recall that $\widetilde{F}$ admits three natural pseudometrics in our setup:

(1) the path-metric d lifted from F (recall that the metric d on F is the natural path-metric induced from E),

(2) the electric metric d_e obtained from d by electrifying the collection $\mathcal{P}$ [10],

(3) the partially electrified pseudometric d_{pel} constructed in Section 3.4.

We now recall a construction from [25, Definition 1.13].

Definition 4.3 Let $\hat{\lambda}$ denote an electric geodesic in $(\widetilde{F}, d_e)$ joining $a, b \in \widetilde{F}$. Modify $\hat{\lambda}$ to a path λ_{ea} as follows. First, λ_{ea} coincides with $\hat{\lambda}$ away from the elevations of $\partial N_\epsilon(\sigma_i)$ to $\widetilde{F}$. Next, for any $\widetilde{\partial N_\epsilon(\sigma_i)}$ that $\hat{\lambda}$ intersects, let x_i and y_i denote the entry and exit points. Join x_i and y_i by a geodesic in $\widetilde{\partial N_\epsilon(\sigma_i)}$ equipped with its flat Euclidean metric. The resulting path λ_{ea} will be called an *electroambient quasigeodesic* in $\widetilde{F}$.

We recall the following consequence of [25, Lemma 1.21; 10, Lemma 4.5 and Proposition 4.6] for easy reference.

Lemma 4.4 *Let* $\widetilde{M}_r$ *denote a drilled atom of* $\widetilde{F}$ (see Definition 3.2). *Let* $a, b \in \widetilde{M}_r$. *Let* $\mathcal{P}_r$ *denote the collection of elevations of* $\widetilde{\partial N_\epsilon(\sigma_i)}$ *to* $\widetilde{M}_r$. *Then* $\widetilde{M}_r$ *is strongly hyperbolic relative to the collection* $\mathcal{P}_r$.

Let d, d_e, *and* d_{pel} *denote the metric*, *electric* (*pseudo*)*metric*, *and the partially electrified* (*pseudo*)*metric* (see Section 3.4) *respectively on* $\widetilde{M}_r$. *Let* λ, $\hat{\lambda}$, λ_{ea}, *and* λ_p *denote respectively the geodesic*, *the electric geodesic*, *the electroambient quasigeodesic* (see Definition 4.3), *and the geodesic with respect to the* d_{pel} *on* $\widetilde{M}_r$ *joining* a *and* b. *Then* λ, $\hat{\lambda}$, λ_{ea}, *and* λ_p *track each other away from* $\mathcal{P}_r$.

Proof It follows from [10] that $\widetilde{M}_r$ is strongly hyperbolic relative to the collection $\mathcal{P}_r$, since M_r admits the structure of a complete hyperbolic manifold with convex boundary.

Lemma 4.5 and Proposition 4.6 of [10] guarantee that λ and $\hat{\lambda}$ track each other away from $\mathcal{P}_r$. The construction of λ_{ea} guarantees that λ_{ea} and $\hat{\lambda}$ agree exactly, away from $\mathcal{P}_r$. Finally, [25, Lemma 1.21] guarantees that λ_p and $\hat{\lambda}$ track each other away from $\mathcal{P}_r$. □

4.2.1 Connectors: a scheme The rest of this subsection is devoted towards constructing a family of paths $\mathcal{F}$. In Section 4.2.2 we construct paths that approximate the shortest paths connecting two different elevations of boundary components of a drilled atom (see Definition 3.2). The paths thus constructed are called *connectors*.

In Section 4.2.3, we extend this family to one for *every* pair of points in elevations of boundary components of a drilled atom (see Definition 3.2). The paths thus constructed are called *extended connectors*.

In Section 4.2.4, we are interested in a concatenation of atoms along a geodesic segment in the Bass–Serre tree. We describe how to concatenate extended connectors in atoms to such concatenations of atoms.

Apart from drilled atoms (see Definition 3.2), one needs to consider elevations $\mathcal{B}$ of maximal undrilled subbundle (see Definition 3.2). Let $\widetilde{S}$ be a boundary component of the elevation $\widetilde{M}_r$ of a drilled atom (see Definition 3.2). Further assume that $\widetilde{S}$ is contained in an elevation $\mathcal{B}$ of a maximal undrilled subbundle (see Definition 3.25). Then we need to consider geodesics from $\widetilde{S}$ to itself lying within $\mathcal{B}$. These are geodesics in the intrinsic hyperbolic metric on $\mathcal{B}$ starting and ending on $\widetilde{S}$. Section 4.2.5 is devoted to such paths.

Section 4.2.6 generalizes the discussion of Section 4.2.5 to the case where there are distinct drilled atoms (see Definition 3.2) abutting $\mathcal{B}$ in distinct boundary components (see Definition 3.25) $\widetilde{S}$ and $\widetilde{S}'$.

With the above constituent connectors in place, we finally define the path family in Section 4.2.7.

4.2.2 Connectors in drilled atoms Recall that for any drilled atom M_r (see Definition 3.2) of F_S, $\widetilde{M}_r$ is strongly hyperbolic relative to $\mathcal{P}_r$ by Lemma 3.12 (see Lemma 3.12 for notation). We assume that M_r has been equipped with a complete hyperbolic metric with convex boundary (as in the proof of the first conclusion of Lemma 3.12). Let $\widetilde{S}_1$ and $\widetilde{S}_2$ denote two boundary components of $\widetilde{M}_r$. We are interested in the nearest-point projections of $\widetilde{S}_1$ and $\widetilde{S}_2$ on each other. The nearest-point projection of $\widetilde{S}_1$ (resp. $\widetilde{S}_2$) onto $\widetilde{S}_2$ (resp. $\widetilde{S}_1$) will be denoted by π_{12} (resp. π_{21}). Three possible cases arise:

Case 1 (product region connectors) There exists a maximal essential proper subsurface Σ of one of the surface boundary components of M_r such that $\Sigma \times [0,1]$, equipped with the standard product structure, embeds in M_r, with $\Sigma \times \{0,1\} \subset \partial M_r$. Further, there exists an elevation $\widetilde{\Sigma} \times [0,1] \subset \widetilde{M}_r$, such that $\widetilde{\Sigma} \times \{0\} \subset \widetilde{S}_1$, and $\widetilde{\Sigma} \times \{1\} \subset \widetilde{S}_2$. In this case, $\pi_{12}(\widetilde{S}_1)$ lies in a uniformly bounded neighborhood of $\widetilde{\Sigma} \times \{1\} \subset \widetilde{S}_2$, and $\pi_{21}(\widetilde{S}_2)$ lies in a uniformly bounded neighborhood of $\widetilde{\Sigma} \times \{0\} \subset \widetilde{S}_1$. For any $z \in \Sigma$, we refer to an elevation of $\{z\} \times [0,1]$ to $\widetilde{M}_r$ as a *product region connector* between $\widetilde{S}_1$ and $\widetilde{S}_2$. Note that in this case, $\widetilde{S}_1$ and $\widetilde{S}_2$ are necessarily elevations of distinct boundary components of M_r.

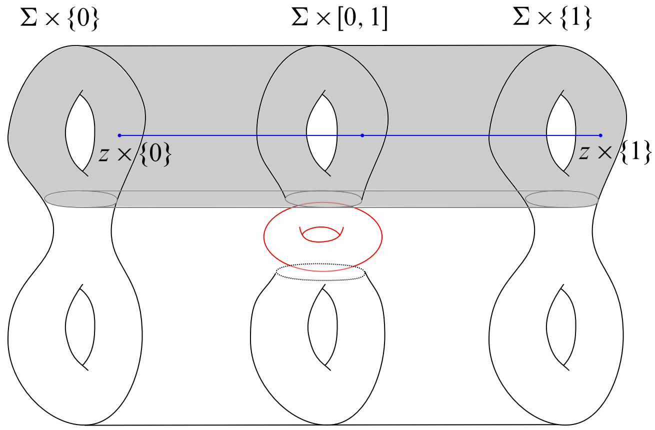

Figure 5: Product region connector in M_r.

Definition 4.5 If $\widetilde{S}_1$ and $\widetilde{S}_2$ are connected by a product region connector, we say that $\widetilde{S}_1$ and $\widetilde{S}_2$ are a *product pair* of elevations.

We emphasize that for $\widetilde{S}_1$ and $\widetilde{S}_2$ a product pair, the nearest-point projection of one onto the other lies in a uniformly bounded neighborhood of $\widetilde{\Sigma} \times \{0\}$ for a maximal essential proper subsurface Σ.

Case 2 (annular connectors) Recall that M_r is obtained from $S \times I$ after drilling finitely many curves $\{\sigma_i\}$. For any $z \in \sigma_i \times \{0,1\}$, for some i, suppose that $z_I = z \times I \setminus \mathrm{Int}(N_\epsilon(\sigma_i)) \subset S \times I$ is contained in M_r, ie z_I does not intersect any $N_\epsilon(\sigma_j)$ for $j \neq i$. Let $z_I^{\pm}$ denote the two components of z_I, and let m_z denote the meridian of $N_\epsilon(\sigma_i)$ passing through $z_I^{\pm} \cap \partial N_\epsilon(\sigma_i)$. Then an *annular connector* η in M_r starts and ends at $\{z\} \times \{0,1\}$, traverses a connected component of z_I, wraps around m_z finitely many times, and finally traverses a (possibly same) connected component of z_I. If η starts and ends at the same point $\{z\} \times \{0\}$ (or $\{z\} \times \{1\}$), then it wraps around m_z (the meridian) k times for some $k \in \mathbb{Z}$; otherwise if it starts and ends at distinct points in $\{z\} \times \{0,1\}$, it wraps around m_z (the meridian) $k + \frac{1}{2}$ times for some $k \in \mathbb{Z}$.

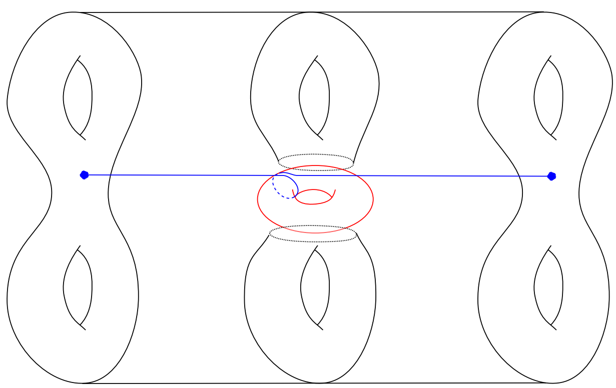

Figure 6: An annular connector in M_r.

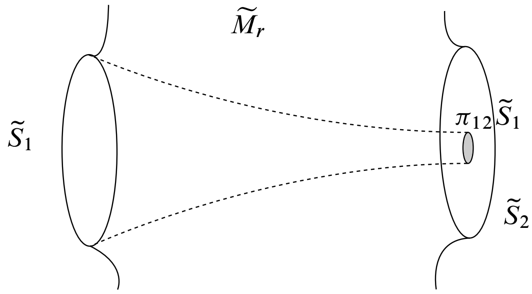

Figure 7: A mutually bounded pair.

Any elevation of an annular connector in M_r is called an *annular connector in* $\widetilde{M}_r$. If the number of times an annular connector wraps around m_z does not equal $\pm\frac{1}{2}$, we call it and its elevations *strict annular connectors*. Let $\widetilde{S}_1$ and $\widetilde{S}_2$ denote the elevations of the boundary components of M_r passing through the endpoints of an annular connector $\tilde{\eta}$. We observe the following:

Lemma 4.6 *Suppose that $\tilde{\eta}$ is a strict annular connector between $\widetilde{S}_1$ and $\widetilde{S}_2$ joining $\tilde{z}_1 \in \widetilde{S}_1$ to $\tilde{z}_2 \in \widetilde{S}_2$. Let $\widetilde{\sigma}$ and $\widetilde{\sigma}'$ denote the elevations of σ_i through $\tilde{z}_1$ and $\tilde{z}_2$, respectively. Then $\pi_{12}(\widetilde{S}_1)$ (resp. $\pi_{21}(\widetilde{S}_2)$) lies in a uniformly bounded neighborhood of $\widetilde{\sigma}'$ (resp. $\widetilde{\sigma}$).*

Proof Let z_1 (the image of $\tilde{z}_1$ under the covering projection) be the basepoint for $\pi_1(M_r)$, and let z_σ be an annular connector that wraps around m_z only half a time. Identify $\pi_1(S_2)$ with loops based at z_1 preceded and succeeded by z_σ with appropriate orientations.

Then the result follows from the fact that $\pi_1(S_l) \cap \pi_1(S_k)^g = \mathbb{Z}$ for $g \in \pi_1(M_r)$ representing any strict annular connector starting and ending at z_1, and $1 \le l, k \le 2$. □

Definition 4.7 If $\widetilde{S}_1$ and $\widetilde{S}_2$ are connected by a strict annular connector, we say that $\widetilde{S}_1$ and $\widetilde{S}_2$ are an *annular pair* of elevations.

Case 3 (the mutually bounded case)

Lemma 4.8 *If $\widetilde{S}_1$ and $\widetilde{S}_2$ is a pair of distinct elevations that are neither a product pair* (see Definition 4.5), *nor an annular pair* (see Definition 4.7), *then $\pi_{12}(\widetilde{S}_1)$ is uniformly bounded in diameter*, *ie $\widetilde{S}_1$ and $\widetilde{S}_2$ are mutually bounded in the sense of Definition 4.1.*

Proof Let $H_i < \pi_1(M_r)$ denote the stabilizers of $\widetilde{S}_i$ for $i = 1, 2$. The Lemma then follows from the fact that $H_1 \cap H_2 = \{1\}$ (see, for instance, the proof of [31, Corollary 3]) (see). □

Figure 7 illustrates this case.

4.2.3 Extended connectors in drilled atoms We now describe a preferred family of quasigeodesics in $\widetilde{M}_r$ connecting $x \in \widetilde{S}_1$ to $y \in \widetilde{S}_2$. Denote $\pi_{12}\widetilde{S}_1$ by Z_2 and $\pi_{21}\widetilde{S}_2$ by Z_1. Then there are the following preferred product regions in $\widetilde{M}_r$:

(1) If $\widetilde{S}_1$ and $\widetilde{S}_2$ are a boundary components (see Definition 3.25), then there exists a proper essential subsurface Σ of S and elevations $\widetilde{\Sigma}_i$ for $i = 1, 2$ of boundary components of the product region $\Sigma \times [0, 1]$ such that Z_1 (resp. Z_2) is coarsely $\widetilde{\Sigma}_1$ (resp. $\widetilde{\Sigma}_2$). Also Z_i are (coarsely) the boundary components of $\widetilde{\Sigma} \times [0, 1]$, where the $[0, 1]$ direction has uniformly bounded length. We normalize the length of $z \times [0, 1]$ to one for all $z \in \widetilde{\Sigma}$.

(2) If $\widetilde{S}_1, \widetilde{S}_2$ is an annular pair, then Z_1 and Z_2 are coarsely elevations of (curves parallel to) σ_i. Further, Z_1 and Z_2 are boundary components of a flat strip $\widetilde{\sigma}_i \times [0, kl(m_z) + a]$, where k equals the number of times the annular connector wraps around the meridian m_z, $l(m_z)$ denotes the length of the meridian m_z, and a equals the sum of the lengths of $z_I^{\pm}$, assuming without loss of generality that they are equal.

(3) Otherwise, if $\widetilde{S}_1, \widetilde{S}_2$ is a mutually bounded pair (as in Lemma 4.8), then Z_1 and Z_2 are coarsely points, ie Z_1 and Z_2 have uniformly bounded diameter.

In all three cases, $\widetilde{S}_i$ is strongly hyperbolic relative to Z_i for $i = 1, 2$. Further, there is a natural product $Z \times [a_1, a_2]$ embedded in $\widetilde{M}_r$ with $Z \times \{a_i\} = Z_i$ for $i = 1, 2$. The interval $[a_1, a_2]$ has the following properties:

1. If $\widetilde{S}_1$ and $\widetilde{S}_2$ are boundary components (see Definition 3.25), $[a_1, a_2]$ has length one.
2. If $\widetilde{S}_1, \widetilde{S}_2$ is an annular pair, $[a_1, a_2]$ has length equal to $kl(m_z) + a$.
3. If $\widetilde{S}_1, \widetilde{S}_2$ is a mutually bounded pair, $[a_1, a_2]$ has length equal to the length of an electroambient quasigeodesic (see Definition 4.3) in $\widetilde{M}_r$ joining $\tilde{z}_1$ and $\tilde{z}_2$.

Next, for any $x_i \in \widetilde{S}_i$ with $i = 1, 2$, let $y_i \in Z_i$ denote a nearest-point projection (in the intrinsic metric on $\widetilde{S}_i$) of x_i onto Z_i. Identifying Z_i for $i = 1, 2$ with $Z \times \{0\} \subset Z \times [0, 1]$ and $Z \times \{1\} \subset Z \times [0, 1]$, respectively, we have the following preferred family of paths joining x_1 and x_2 in $\widetilde{M}_r$:

Projecting both y_1 and y_2 to the Z-factor, we get points that we call y_1 and y_2 again to avoid cluttering notation. Let $[y_1, y_2]$ denote the geodesic in Z joining y_1 and y_2. Let p be any point on $[y_1, y_2]$. Then the preferred collection of paths joining $x_1, x_2 \in \widetilde{M}_r$ are given by the concatenation of the following segments:

1. the geodesic $[x_1, y_1] \subset \widetilde{S}_1$ joining x_1 and y_1,
2. the geodesic $[y_1, p \times \{0\}] \subset (Z \times \{0\})(= Z_1) \subset \widetilde{S}_1$ joining y_1 and $p \times \{0\}$ in $(Z \times \{0\})$,
3. the vertical interval $p \times [0, 1]$ traveling from $p \times \{0\}$ to $p \times \{1\}$,
4. the geodesic $[p \times \{1\}, y_2] \subset (Z \times \{1\})(= Z_2) \subset \widetilde{S}_2$ joining $p \times \{1\}$ and y_2 in $(Z \times \{1\})$,
5. the geodesic $[y_2, x_2] \subset \widetilde{S}_2$ joining y_2 and x_2.

Note that there is only one vertical interval $p \times [0, 1]$ traveling from $p \times \{0\}$ to $p \times \{1\}$ in each member of the family given above. Let $\mathcal{F}(M_r, x_1, x_2)$ denote the above family.

The construction above shows the following:

Lemma 4.9 *Each $\alpha \in \mathcal{F}(M_r, x_1, x_2)$ tracks the d_{pel}-geodesic and the electroambient quasigeodesic* (see Definition 4.3) *between x_1 and x_2 in the intrinsic metric on $\widetilde{M}_r$.*

Proof The fact that α tracks the electroambient quasigeodesic (see Definition 4.3) *along elements of $\mathcal{P}_r$* follows from the fact that the nearest-point projection of any $P \in \mathcal{P}_r$ onto $\widetilde{S}_1$ is either uniformly bounded or (coarsely) an elevation of σ_i to $\widetilde{S}_1$. More precisely, in the second case, there exists an elevation $\widetilde{\sigma}_i \subset \widetilde{S}_1$ such that the nearest-point projection of $P \in \mathcal{P}_r$ onto $\widetilde{S}_1$ lies in a uniformly bounded neighborhood of $\widetilde{\sigma}_i$, and hence the concatenations $[x_1, y_1] \cup [y_1, p \times \{0\}]$ and $[p \times \{1\}, y_2] \cup [y_2, x_2]$ used to define α have a maximal subpath each parallel to P.

Away from elements of $\mathcal{P}_r$, this is a consequence of Lemmas 3.12 and 4.4. □

4.2.4 Extended connectors in concatenated drilled atoms Lemma 4.9 can be extended to a 3-manifold obtained by concatenating finitely many atoms. As before, let $\Pi \colon \widetilde{F} \to \mathcal{T}$ denote the Bass–Serre tree of $\widetilde{F} = \widetilde{F}_S$, with vertex spaces X_v given by elevated atoms.

Corollary 3.31 can be strengthened slightly as follows:

Corollary 4.10 *There exists $L \geq 1$ such that the following holds: Let $[a, b] \subset \mathcal{T}$ denote a geodesic of length at least L such that X_a and X_b are elevated drilled atoms* (see Definition 3.2) *of F_S. Let $X_{[a,b]}$ denote the 3-manifold given by $\Pi^{-1}([a, b])$. Also, let $\widetilde{S}_a \subset X_a$ and $\widetilde{S}_b \subset X_b$ denote boundary components* (see Definition 3.25) *of $X_{[a,b]}$ in the sense of Definition 3.25. Then:*

1. *$X_{[a,b]}$ is strongly hyperbolic relative to the collection $\mathcal{P}_{[a,b]}$ (the elements of $\mathcal{P}$ contained in it).*
2. *There exists D depending only on $d_{\mathcal{T}}(a, b)$ such that $\pi_{ab}(\widetilde{S}_a)$ has diameter bounded by D (here, π_{ab} denotes the nearest-point projection of $\widetilde{S}_a$ on $\widetilde{S}_b$).*

Proof The first conclusion has the same proof as the first conclusion of Lemma 3.12. The second now follows from Corollary 3.31. □

The construction of extended connectors in $X_{[a,b]}$ can now be carried out exactly as in Section 4.2.3. We denote the family thus constructed by $\mathcal{F}(X_{[a,b]}, x_1, x_2)$. Corollary 4.10 gives us the following immediate consequence.

Lemma 4.11 *We continue with the setup of Corollary 4.10. There exists $L \geq 1$ such that if $d_{\mathcal{T}}(a, b) \geq L$, then preferred extended connectors $\alpha \in \mathcal{F}(X_{[a,b]}, x_1, x_2)$ are concatenations of three pieces, where the middle piece α_m is necessarily a cobounded connector, and the first and last ones are geodesics in $\widetilde{S}_a$ and $\widetilde{S}_b$.*

Thus, the endpoints of α_m are coarsely well-defined, ie there exists D depending only on $d_{\mathcal{T}}(a, b)$, and z_1 and z_2 in $\widetilde{S}_a$ and $\widetilde{S}_b$, respectively, such that the endpoints of α_m lie on $\widetilde{S}_a$ and $\widetilde{S}_b$ at a distance of at most D from z_1 and z_2.

Since $X_{[a,b]}$ is strongly hyperbolic relative to the collection $\mathcal{P}_{[a,b]}$, as are all drilled atoms $\widetilde{M}_c$ (see Definition 3.2) for $c \in [a,b]$, the restriction of $\gamma \in \mathcal{F}(X_{[a,b]}, x_1, x_2)$ to $\widetilde{M}_c$ may be perturbed by a uniformly bounded amount, so that $(\gamma \cap \widetilde{M}_c) \in \mathcal{F}(M_c, x_1, x_2)$ for some $x_1, x_2 \in \partial \widetilde{M}_c$.

Remark 4.12 Lemma 4.11 implies that α_m is a coarsely well-defined electroambient quasigeodesic (see Definition 4.3) in $X_{[a,b]}$, and *any preferred connector* between $\widetilde{S}_a$ and $\widetilde{S}_b$ coarsely contains it. Note however that the parameter D determining coarseness depends on $d_{\mathcal{T}}(a,b)$.

Let $[a,b] \subset \mathcal{T}$ denote a geodesic of length at least L such that X_a and X_b are elevated drilled atoms of F_S (see Definition 3.2). Let $X_{[a,b]}$ denote the 3-manifold given by $\Pi^{-1}([a,b])$. Assume further that the only drilled atoms in $X_{[a,b]}$ are X_a and X_b. Let $\widetilde{S}_a$ and $\widetilde{S}_b$ denote the boundary components (see Definition 3.25) of $X_{[a,b]} \setminus \mathrm{Int}(X_a \cup X_b)$, so that $\widetilde{S}_a$ and $\widetilde{S}_b$ are also boundary components (see Definition 3.25) of X_a and X_b, respectively. We refer to $\widetilde{S}_a$ and $\widetilde{S}_b$ as *internal boundary components* of X_a and X_b, respectively. All other boundary components will be referred to as *external boundary components*.

Lemma 4.13 *Let $X_{[a,b]}$, X_a, and X_b be as above. Then there exists $L \geq 2$ such that if $d_{\mathcal{T}}(a,b) \geq L$, the following holds. For any external boundary components $\widetilde{S}_1$ and $\widetilde{S}_2$ of $X_{[a,b]}$, there exists a coarsely well-defined geodesic λ joining internal boundary components $\widetilde{S}_a$ and $\widetilde{S}_b$ in $X_{[a,b]} \setminus \mathrm{Int}(X_a \cup X_b)$ such that any $\alpha \in \mathcal{F}$ connecting $\widetilde{S}_1$ and $\widetilde{S}_2$ coarsely contains λ. Further, the coarseness is uniform, independent of $X_{[a,b]}$.*

Proof Let π_{1a} and π_{2b} denote nearest-point projections of $\widetilde{S}_1$ onto $\widetilde{S}_a$, and $\widetilde{S}_2$ onto $\widetilde{S}_b$, respectively. Then the images of π_{1a} and π_{2b}, given by W_1 and W_2, respectively, are given by an elevation each of a proper essential subsurfaces of $\widetilde{S}_a$ and $\widetilde{S}_b$, respectively. Hence, by Theorem 2.1, there exists $L \geq 2$ such that if $d_{\mathcal{T}}(a,b) \geq L$, there exists a coarsely well-defined shortest geodesic σ in $X_{[a,b]} \setminus \mathrm{Int}(X_a \cup X_b)$ joining W_1 and W_2. Any $\alpha \in \mathcal{F}$ connecting $\widetilde{S}_1$ and $\widetilde{S}_2$ must coarsely join points in W_1 and W_2, and hence must coarsely contain σ. □

4.2.5 Controlling backtracking in elevated subbundles Suppose that $\widetilde{M}_r$ is a drilled atom of $\widetilde{F}$ (see Definition 3.2). Let $\widetilde{S}_1$, $\widetilde{S}_2$, and $\widetilde{S}_3$ denote three distinct boundary components of $\widetilde{M}_r$, and let $\mathcal{B}$ be the elevation of a maximal undrilled subbundle of F_S (Definition 3.2) such that $\mathcal{B} \cap \widetilde{M}_r = \widetilde{S}_3$. Let π_{13} and π_{23} be nearest-point projections of $\widetilde{S}_1$ and $\widetilde{S}_2$ onto $\widetilde{S}_3$. Then the images of π_{13} and π_{23} are either

(1) given by an elevation of a proper essential subsurface of S_3 (as in Section 4.2.2), or

(2) uniformly bounded in diameter.

Further, by Theorem 2.1, $\pi_{13}(\widetilde{S}_1)$ and $\pi_{23}(\widetilde{S}_2)$ are uniformly quasiconvex in $\mathcal{B}$, ie there exists $C \geq 1$ such that for any $\mathcal{B}$, $\widetilde{S}_1$, $\widetilde{S}_2$, and $\widetilde{S}_3$ as above, $\pi_{13}(\widetilde{S}_1)$ and $\pi_{23}(\widetilde{S}_2)$ are C-quasiconvex in $\mathcal{B}$. Hence, there exists D_0 such that if $d(\pi_{13}(\widetilde{S}_1), \pi_{23}(\widetilde{S}_2)) \geq D_0$, there is a coarsely unique shortest path α in $\mathcal{B}$ joining $\pi_{13}(\widetilde{S}_1)$ and $\pi_{23}(\widetilde{S}_2)$. Thus, due to quasiconvexity of $\pi_{13}(\widetilde{S}_1)$ and $\pi_{23}(\widetilde{S}_2)$, there exists $C_1 \geq 0$ such that exactly one of the following holds:

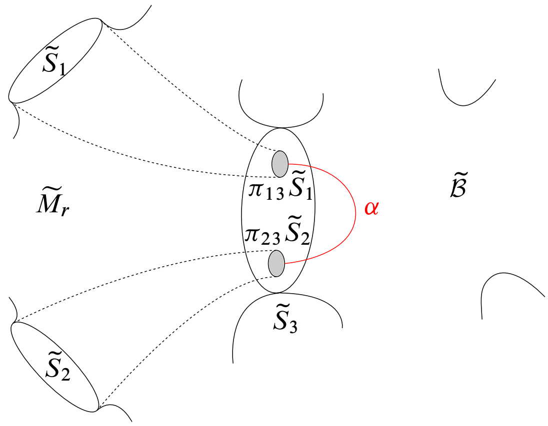

Figure 8: Allowable backtrack in $\mathcal{B}$.

(1) There is a coarsely well-defined geodesic α joining $\pi_{13}(\widetilde{S}_1)$ and $\pi_{23}(\widetilde{S}_2)$ in $\mathcal{B}$. In this case, any d_{pel}-geodesic between $\widetilde{S}_1$ and $\widetilde{S}_2$ in $\widetilde{M}_r \cup \mathcal{B}$ C_1-coarsely contains α, ie any d_{pel}-geodesic between $\widetilde{S}_1$ and $\widetilde{S}_2$ in $\widetilde{M}_r \cup \mathcal{B}$ contains a subpath tracking α within distance C_1 of it. In such a case the d_{pel}-geodesic is said to *have an allowable backtrack in* $\mathcal{B}$.

(2) $\pi_{13}(\widetilde{S}_1) \cup \pi_{23}(\widetilde{S}_2)$ is C_1-quasiconvex in $\mathcal{B}$. In this case, a d_{pel}-geodesic λ_p between $\widetilde{S}_1$ and $\widetilde{S}_2$ is said to *have nonallowable backtracks in* $\mathcal{B}$ if $\lambda_p \cap \mathcal{B} \neq \varnothing$. By perturbing any such d_{pel}-geodesic between $\widetilde{S}_1$ and $\widetilde{S}_2$ by a bounded amount, we obtain a d_{pel}-quasigeodesic that does not intersect $\mathcal{B}$ at all.

Remark 4.14 The constant C_1 above may be chosen to be uniform, ie independent of the choice of $\mathcal{B}$, $\widetilde{M}_r$, $\widetilde{S}_1$, $\widetilde{S}_2$, and $\widetilde{S}_3$ (as there are only finitely many such possibilities up to the action of $\pi_1(F)$).

4.2.6 Controlling small connectors in elevated subbundles We generalize the above discussion in Section 4.2.5 to the case where there are two elevations $\widetilde{S}_3$ and $\widetilde{S}_3'$ coming from *different* elevated drilled atoms (see Definition 3.2) abutting a common elevated maximal undrilled subbundle $\mathcal{B}$ (see Definition 3.2).

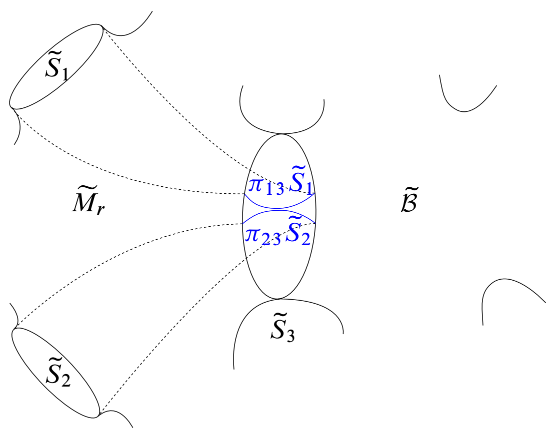

Figure 9: Nonallowable backtrack in $\mathcal{B}$.

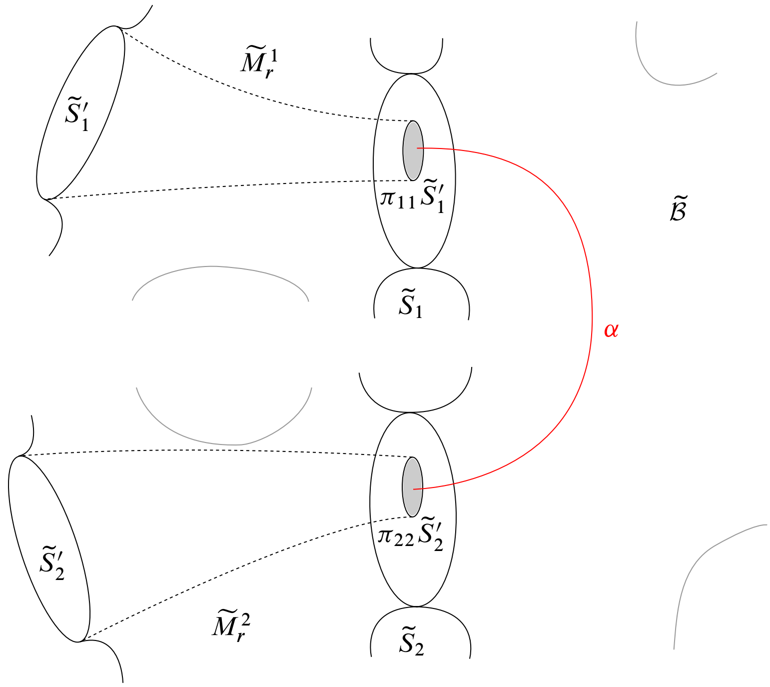

Figure 10: A long connector in an undrilled elevation.

Let $\widetilde{M}_r^1 \neq \widetilde{M}_r^2$ be drilled atoms of $\widetilde{F}$. Let $\mathcal{B}$ be the elevation of a maximal undrilled subbundle of F_S (Definition 3.2) such that $\mathcal{B} \cap \widetilde{M}_r^i = \widetilde{S}_i$ for $i = 1, 2$. Let $\widetilde{S}_i' \subset \widetilde{M}_r^i$ for $i = 1, 2$ denote boundary components of $\widetilde{M}_r$ (different from $\widetilde{S}_i$). Let π_{11} and π_{22} be nearest-point projections of $\widetilde{S}_1'$ and $\widetilde{S}_2'$ onto $\widetilde{S}_1$ and $\widetilde{S}_2$, respectively. Then the images of π_{ii} for $i = 1, 2$ are either

(1) given by an elevation of a proper essential subsurface of S_i (as in Section 4.2.2), or

(2) uniformly bounded in diameter.

By Theorem 2.1, $\pi_{11}(\widetilde{S}_1')$ and $\pi_{22}(\widetilde{S}_2')$ are uniformly quasiconvex in $\mathcal{B}$. As in Section 4.2.5 above, there exist $C_1 \geq 0$ and $D_1 \geq 0$ such that exactly one of the following holds:

(1) There is a coarsely well-defined geodesic α joining $\pi_{11}(\widetilde{S}_1')$ and $\pi_{22}(\widetilde{S}_2')$ in $\mathcal{B}$. In this case, any d_{pel}-geodesic between $\widetilde{S}_1'$ and $\widetilde{S}_2'$ in $\widetilde{M}_r \cup \mathcal{B}$ C_1-coarsely contains α.

(2) $\pi_{11}(\widetilde{S}_1') \cup \pi_{22}(\widetilde{S}_2')$ is C_1-quasiconvex in $\mathcal{B}$. Further, the distance between $\pi_{11}(\widetilde{S}_1')$ and $\pi_{22}(\widetilde{S}_2')$ in $\mathcal{B}$ is at most D_1.

We shall refer to paths α as in item (1) above as *long connectors in undrilled elevations*.

If, on the other hand $\pi_{11}(\widetilde{S}_1') \cup \pi_{22}(\widetilde{S}_2')$ is C_1-quasiconvex in $\mathcal{B}$ as in item (2) above, then the triple $\widetilde{M}_r^1, \mathcal{B}, \widetilde{M}_r^2$ will be called a *short connector triple*.

4.2.7 Defining the path family We now define the family $\mathcal{F}$ of paths that will feed into Theorem 4.2. Recall that $(\widetilde{F}, d_{\mathrm{pel}})$ is hyperbolic by Theorem 3.14. For $x, y \in \widetilde{F}$, define $\eta_p(x, y)$ to be a geodesic in $(\widetilde{F}, d_{\mathrm{pel}})$. Now, replace $\eta_p(x, y)$ by a d_{pel}-quasigeodesic that

(1) does not backtrack from drilled atoms (in the sense of Definition 3.30) as in Corollary 3.29,

(2) does not have nonallowable backtracks in elevations of maximal undrilled subbundles.

Abusing notation slightly, we continue to denote the nonbacktracking d_{pel}-quasigeodesic by η_p. For each drilled atom $\widetilde{M}_r$ (see Definition 3.2) and each connected component ζ_r of $\eta_p \cap \widetilde{M}_r$, replace ζ_r by an electroambient geodesic joining its endpoints. Denote the resulting path in $\widetilde{F}$ joining the endpoints of η_p by η.

We refer to η as the electroambient path obtained from η_p by *de-electrification*. Next, define $\mathcal{F}_p$ to be a collection $\{\eta_p(x, y)\}$ of uniform quasigeodesics in $(\widetilde{F}, d_{\mathrm{pel}})$ without backtracking in drilled atoms (cf Definition 3.30), one for every pair $x, y \in \widetilde{F}$. Define $\mathcal{F}$ to be the collection $\{\eta(x, y)\}$ obtained from the collection $\{\eta_p(x, y)\}$ by de-electrification. For each such $\eta(x, y)$ define $\theta(x, y)$ to be the closure of $\eta(x, y) \setminus \bigcup_{P \in \mathcal{P}} P$. Thus, $\theta(x, y)$ is obtained from $\eta(x, y)$ by removing the interiors of the intersections with elements of $\mathcal{P}$.

In short, $\mathcal{F}$ is obtained from geodesics in $(\widetilde{F}, d_{\mathrm{pel}})$ by

(1) removing backtracking in drilled atoms (see Definition 3.30), and removing nonallowable backtracks in elevations of maximal undrilled subbundle (see Definition 3.2),

(2) subsequent de-electrification.

Remark 4.15 The purpose of replacing a d_{pel}-geodesic by a d_{pel}-quasigeodesic that does not backtrack from drilled atoms is to minimize intersections with elevated singular fibers (see Definition 1.2) $\widetilde{S}_x$ in $\widetilde{F}$. This ensures that the only allowable backtracking is in elevations of maximal undrilled subbundles (see Definition 3.2).

4.3 Stability

The aim of this subsection is to prove the following stability condition, which is the main technical result of this section:

Proposition 4.16 *Given $D > 0$, there exists $C > 0$ such that the following holds: Let $\eta(x, y), \eta(u, v) \in \mathcal{F}$ such $d(x, u) \leq D$ and $d(y, v) \leq D$. Then $\eta(x, y)$ and $\eta(u, v)$ track each other in a C-neighborhood of each other.*

We first observe a version of Proposition 4.16 in atoms.

Lemma 4.17 *Given $D > 0$, there exists $C > 0$ such that the following holds: Let $\eta(x, y), \eta(u, v) \in \mathcal{F}$ be such that:*

1. *There exists an atom $\widetilde{\mathbb{M}}$ equal to $\widetilde{M}_r$ (drilled) or $\widetilde{M}$ (undrilled) such that $\eta(x, y), \eta(u, v)$ are contained in $\widetilde{\mathbb{M}}$.*
2. *$x, u \in \widetilde{S}_x$ and $y, v \in \widetilde{S}_y$, where $\widetilde{S}_x, \widetilde{S}_y$ are elevations of singular fibers (see Definition 1.2), or equivalently, boundary components of $\widetilde{\mathbb{M}}$.*
3. *$d(x, u) \le D$ and $d(y, v) \le D$.*

Then $\eta(x, y)$ and $\eta(u, v)$ track each other in a C-neighborhood of each other.

Proof This is a consequence of Lemma 4.4, which guarantees that $\eta(x, y)$ and $\eta(u, v)$ track each other away from the elements of $\mathcal{P}$ in $\widetilde{\mathbb{M}}$. Further, since each element of $\mathcal{P}$ is a flat $\mathbb{R}^2$, geodesics in each element of $\mathcal{P}$ track each other provided they start and end nearby. The lemma now follows from the construction of electroambient quasigeodesics (see Definition 4.3). □

Before starting with the proof of Proposition 4.16, we point out that the main idea below is to divide an element $\eta \in \mathcal{F}$ into pieces that satisfy the property that its endpoints are coarsely well-defined.

Proof of Proposition 4.16 Assume, without loss of generality, that $x, u \in \widetilde{S}_1 \subset X_a$ and $y, v \in \widetilde{S}_2 \subset X_b$, where $a, b \in \mathcal{T}$. There exists an indexing set $\mathcal{I}$ giving a sequence of vertices $a = a_0, \dots, a_n = b$, possibly with repetition, such that $\eta(x, y)$ traverses $\widetilde{M}_{a_i}$ in order. Since $\eta(x, y) \in \mathcal{F}$, it does not

1. backtrack in drilled atoms (see Definition 3.30),
2. have nonallowable backtracks in elevations of maximal undrilled subbundle (see Definition 3.2).

If there exists a maximal undrilled subbundle $\mathcal{B}$ (see Definition 3.2) in which $\eta(x, y)$ has an allowable backtrack (Section 4.2.5), we collect together all the (necessarily consecutive) vertices in the sequence $\{a_i\}$ that correspond to atoms contained in $\mathcal{B}$ and replace them by a single vertex B_j for some j. We refer to such B_j as an *undrilled molecule*. Thus, if some such B_j occurs, then there exists $a_i \in \mathcal{I}$ such that B_j occurs in a unique subsequence of the form $a_i B_j a_i$ in $\mathcal{I}$.

Again, if there exists a maximal undrilled subbundle $\mathcal{B}$ (see Definition 3.2) in which $\eta(x, y)$ has a long connector (Section 4.2.6), then also we collect together all the (necessarily consecutive) vertices in the sequence $\{a_i\}$ that correspond to (elevated) atoms contained in $\mathcal{B}$ and replace them by a single vertex B_k for some k. We also refer to such a B_k as an *undrilled molecule*. Thus, if some such B_k occurs, then there exists $a_i \neq a_s \in \mathcal{I}$ such that B_k occurs in a unique subsequence of the form $a_i B_k a_s$ in $\mathcal{I}$.

The construction of undrilled molecules now allows us obtain a new finite sequence

$$\mathcal{J} = a_{1,1}, \dots, a_{1,m_1}, B_1, a_{2,1}, \dots, a_{2,m_2}, B_2, \dots.$$

Note that the only possible repetition allowable in this sequence are of the following form. If η has an allowable backtrack in B_j for some j, then there is a triple of the form $a_{j,m_j}, B_j, a_{j+1,1}$ with $a_{j,m_j} = a_{j+1,1}$ corresponding to the same elevated drilled atom (see Definition 3.2).

By the properties of an allowable backtrack (Section 4.2.5) or long connectors in undrilled elevations (Section 4.2.6), both $\eta(x,y)$ and $\eta(u,v)$ coarsely contain a subpath in B_j for all undrilled molecules B_j (here we are conflating the index B_j with the elevated maximal undrilled subbundle (see Definition 3.2) it indexes). Let $\widetilde{S}_j^-$ and $\widetilde{S}_j^+$ denote the boundary components of B_j (see Definition 3.25) through which $\eta(x,y)$ and $\eta(u,v)$ enter and leave B_j. Then there exist

(1) $z_j^-(x,y) \in \eta(x,y) \cap \widetilde{S}_j^-$ and $z_j^+(x,y) \in \eta(x,y) \cap \widetilde{S}_j^+$,

(2) $z_j^-(u,v) \in \eta(u,v) \cap \widetilde{S}_j^-$ and $z_j^+(u,v) \in \eta(u,v) \cap \widetilde{S}_j^+$,

such that

(1) $z_j^-(x,y)$ and $z_j^-(u,v)$ lie at a distance of at most $2C_1$ from each other on $\widetilde{S}_j^-$,

(2) $z_j^+(x,y)$ and $z_j^+(u,v)$ lie at a distance of at most $2C_1$ from each other on $\widetilde{S}_j^+$.

Let $\eta_j(x,y)$ denote the subpath of $\eta(x,y)$ between $z_j^-(x,y)$ and $z_j^+(x,y)$. Let $\eta_j(u,v)$ denote the subpath of $\eta(u,v)$ between $z_j^-(u,v)$ and $z_j^+(u,v)$. By hyperbolicity of each B_j, $\eta_j(x,y)$ and $\eta_j(u,v)$ track each other in a C_1'-neighborhood of each other, where C_1' depends only on C_1 and the hyperbolicity constant of B_j. Hence C_1' is uniform.

To prove Proposition 4.16, it therefore suffices to assume that the finite sequence $\mathcal{J}$ constructed from $\mathcal{I}$ *does not contain any undrilled molecule* B_j.

A caveat is in order. We note that if there is a short connector triple $\widetilde{M}_r^1, \mathcal{B}, \widetilde{M}_r^2$ as in Section 4.2.6, then the atoms of the elevated maximal undrilled subbundle (see Definition 3.2) $\mathcal{B}$ are *not* combined into a single undrilled molecule. Further, there exists uniform $C_2 \geq 1$ (independent of $\mathcal{B}$, $\eta(x,y)$, and $\eta(u,v)$) such that after a uniformly bounded perturbation if necessary, both $\eta(x,y)$ and $\eta(u,v)$

(1) intersect the same set of undrilled atoms (see Definition 3.2) $A_1,\dots,A_k$ of $\mathcal{B}$ in order without backtracking in any of the atoms (ie after leaving any of the undrilled atoms A_i, $\eta(x,y)$, and $\eta(u,v)$ do not return to it),

(2) $k \leq C_2$ (this follows from the quasiconvexity property of the union of projections used to define short connector triples in Section 4.2.6).

We summarize this by saying that $\eta(x,y)$ and $\eta(u,v)$ have no backtracking in short connector triples.

We now return to the sequence of atoms $\mathcal{J} = a_0,\dots,a_m$ where

(1) $\mathcal{J}$ has no undrilled molecule, and hence

(2) $\mathcal{J}$ is the vertex sequence of a *geodesic* in the Bass–Serre tree $\mathcal{T}$ of $\widetilde{F}$.

The absence of backtracking in short connector triples along with no backtracking in drilled atoms (see Definition 3.30) guarantees that $\mathcal{J}$ is the vertex sequence of a geodesic.

Let L be the maximum of C_2 and the constants in Lemmas 4.11 and 4.13. If there is a sequence of more than L contiguous undrilled blocks in $\mathcal{J}$, then choose a maximal sequence $a_k, \dots, a_l$ with $l-k \geq L$ indexing such blocks. Then, by Lemma 4.13, there exists a coarsely well-defined α connecting boundary components $\widetilde{S}_k$ and $\widetilde{S}_l$ (the *internal boundary components* occurring in Lemma 4.13) coarsely contained in both $\eta(x,y)$ and $\eta(u,v)$. Hence there exist $z_k(x,y) \in \eta(x,y) \cap \widetilde{S}_k$ and $z_k(u,v) \in \eta(u,v) \cap \widetilde{S}_k$ such that $z_k(x,y)$ and $z_k(u,v)$ lie at a distance of at most $2C_1$ from each other on $\widetilde{S}_k$, where C_1 is the uniform constant of coarseness from Lemma 4.13. We may therefore assume henceforth that there does not exist a sequence of more than L contiguous undrilled blocks in $\mathcal{J}$.

Next, let $a_{k_1}, \dots, a_{k_L}$ be a subsequence of $\mathcal{J}$ such that

(1) each a_{k_i} indexes a drilled atom (see Definition 3.2),

(2) $k_{i+1} > k_i$,

(3) for any $i \in 1, \dots, L$, every j strictly between a_{k_i} and $a_{k_{i+1}}$ indexes an undrilled atom (see Definition 3.2).

Then, by the simplification in the above paragraph, $a_{k_{i+1}} - a_{k_i} \leq L$. Thus, any subsequence of L drilled atoms in $\mathcal{J}$ interpolated only by undrilled atoms (see Definition 3.2) has length at most L^2. We shall refer to such a subsequence of $\mathcal{J}$ as a *subsequence of successive L drilled atoms*. Note that in a subsequence of successive L drilled atoms, the drilled atoms (see Definition 3.2) need not be contiguous.

For such a subsequence of successive L drilled atoms, let $a = a_{k_1}$, $b = a_{k_L}$, and $X_{[a,b]}$, $\widetilde{S}_a$, and $\widetilde{S}_b$ be as in Lemma 4.11. Then, by Lemma 4.11, there exists a coarsely well-defined α connecting boundary components (see Definition 3.25) $\widetilde{S}_a$ and $\widetilde{S}_b$ coarsely contained in both $\eta(x,y)$ and $\eta(u,v)$.

Finally divide $\mathcal{J}$ into a subsequence of successive L drilled atoms as follows: let

$$\mathcal{J} = a_0 = a_{n_0}, \dots, a_{n_1}, \dots, a_{n_s}, \dots, a_m$$

be such that

(1) each a_{n_i} for $i = 1, \dots, s$ is a drilled atom (see Definition 3.2),

(2) the subsequence of $\mathcal{J}$ between a_{n_i} and $a_{n_{i+1}}$ (both included) has exactly L drilled atoms for $i = 0, \dots, s$.

We also call the initial sequence $a_{n_0}, \dots, a_{n_1}$ a subsequence of successive L drilled atoms. For each such subsequence of successive L drilled atoms, there exists a coarsely well-defined α connecting boundary components (see Definition 3.25) $\widetilde{S}_a$ and $\widetilde{S}_b$ as before, and coarsely contained in both $\eta(x,y)$ and $\eta(u,v)$.

Setting $c = a_0$ and $d = a_{n_s}$, it follows that there exist C', and $y', v' \in \widetilde{S}_d$, and subpaths $\eta(x,y')$ and $\eta(u,v')$ of $\eta(x,y)$ and $\eta(u,v)$, respectively,

(1) starting at x and u, respectively,

(2) ending at y' and v', respectively,

such that

(1) y' and v' lie at a distance of at most C' from each other on $\widetilde{S}_d$,

(2) $\eta(x, y')$ and $\eta(u, v')$ track each other in a C'-neighborhood of each other,

(3) C' is independent of x, y, u, v, and $\mathcal{J}$.

Let $\eta(y', y)$ and $\eta(v', v)$ denote the subpaths of $\eta(x, y)$ and $\eta(u, v)$, respectively,

(1) starting at y' and v', respectively,

(2) ending at y and v, respectively.

Let $p = a_{n_s}$ and $q = a_m$. Then $X_{[p,q]}$ is a 3-manifold given by a concatenation of at most L^2 atoms. The proof of Lemma 3.12 now furnishes relative hyperbolicity for each such $X_{[p,q]}$. Since there are only finitely many possibilities, the constants of relative hyperbolicity are uniform. The tracking properties of $\eta(y', y)$ and $\eta(v', v)$ now follow from relative hyperbolicity. Combining this with the tracking properties of $\eta(x, y')$ and $\eta(u, v')$ already established, Proposition 4.16 follows. □

4.4 Checking conditions 1, 2, and 4–6 of Theorem 4.2

We shall now show that the family $\mathcal{F}$ defined above satisfy the conditions of Theorem 4.2.

Theorem 4.2(1) After rescaling F if necessary, we might as well assume that the distance between any singular fiber (see Definition 1.2) of F and the boundary $\partial(N_\epsilon(\sigma_i))$ of the neighborhood of any drilled curve is at least 4. Hence, if $d_X(x, y) \leq 2$, then the condition follows from strong relative hyperbolicity of $\widetilde{M}_r$ where M_r is any drilled atom (see Definition 3.2) in F (Lemma 4.4).

Theorem 4.2(2) This is a consequence of the proof of stability of elements of $\mathcal{F}$, Proposition 4.16.

Theorem 4.2(4) This follows from the fact that any element of $\mathcal{F}$ starting on $P_1 \in \mathcal{P}$ and ending on an element $P_2 \neq P_1$ of $\mathcal{P}$ necessarily has points in the complement of $\bigcup_{P \in \mathcal{P}} P$.

Theorem 4.2(5) This follows from Corollary 4.10. Indeed, for any $P_1, P_2 \in \mathcal{P}$, there exists $[a_1, a_2] \in \mathcal{T}$ such that $P_i \in X_{a_i}$ for $i = 1, 2$. The strong relative hyperbolicity of $X_{[a_1,a_2]}$ relative to the collection of elements of $\mathcal{P}$ contained in it (with constants depending only on $d_{\mathcal{T}}(a_1, a_2)$) furnishes (5).

Theorem 4.2(6) This also follows from Corollary 4.10. Indeed, as in (5) above, we can choose $[a_1, a_2] \in \mathcal{T}$ such that $d_{\mathcal{T}}(a_1, a_2) \leq 2k$, where k is as in (6). Then strong relative hyperbolicity of $X_{[a_1,a_2]}$ relative to the collection of elements of $\mathcal{P}$ contained in it furnishes the constant K required by (6).

4.5 Thin triangles in $\mathcal{F}$

It remains to prove the thin triangle condition, ie Theorem 4.2(3). Let $a, b, c \in \widetilde{F}$.

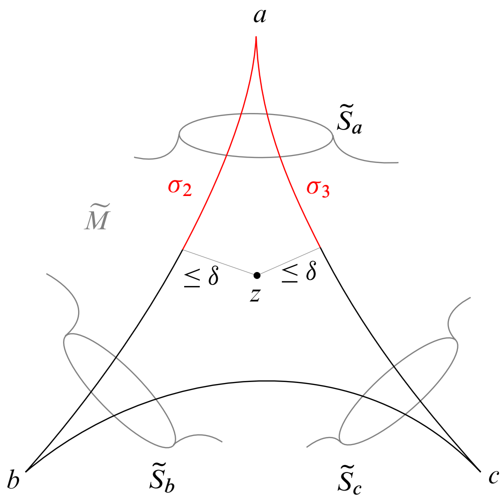

Figure 11: The case when the median lies in an undrilled atom.

Let γ_1^d, γ_2^d, and γ_3^d be sides of a quasigeodesic triangle in $(\widetilde{F}, d_{\mathrm{pel}})$ with vertices a, b, and c used for constructing elements of $\mathcal{F}$. Let $\gamma_1, \gamma_2, \gamma_3 \in \mathcal{F}$ denote the elements of the path family constructed from them. Let a (resp. b, c) be the vertex opposite γ_1 (resp. γ_2, γ_3). Let z denote a centroid.

Case 1 (z lies in an undrilled atom $\widetilde{M}$ in $\widetilde{F}$; see Figure 11) Thinness of triangles follows from stability, Proposition 4.16. Indeed z lies close to each of γ_1, γ_2, and γ_3 (in the usual unelectrified metric on $\widetilde{F}$), as each boundary component, and hence $\widetilde{M}$, is properly embedded in $(\widetilde{F}, d_{\mathrm{pel}})$; see Lemma 3.28. Hence for a (or b, or c) there exist a pair of paths σ_2, σ_3 (given by subpaths of γ_2 and γ_3, respectively) starting from a and ending close to z in $\widetilde{M}$. By Proposition 4.16, σ_2 and σ_3 track each other (with uniform constants). A similar argument works for b and c completing this case.

Case 2 (z lies in a drilled atom $\widetilde{M}_r$ in $\widetilde{F}$; see Figure 12) There are three boundary components $\widetilde{S}_a$, $\widetilde{S}_b$, and $\widetilde{S}_b$ of $\widetilde{M}_r$ (recall that the vertices of the triangle are a, b, and c) and subpaths β_1^d, β_2^d, and β_3^d of γ_1^d, γ_2^d, and γ_3^d such that the following holds. In $\widetilde{M}_r$ equipped with d_{pel}, the subpaths β_1^d, β_2^d, and β_3^d pass close to z. We assume that β_1^d and β_2^d have one endpoint each on $\widetilde{S}_c$, β_3^d and β_1^d have one endpoint each on $\widetilde{S}_b$, and β_2^d and β_3^d have one endpoint each on $\widetilde{S}_a$.

This gives a hexagon in $\widetilde{M}_r$, where the other three sides (other than β_1^d, β_2^d, and β_3^d) are geodesics in $\widetilde{S}_a$, $\widetilde{S}_b$, and $\widetilde{S}_c$, joining the two intersection points of the segments. Call these sides the *complementary geodesics*, and denote them by α_a, α_b, and α_c, respectively. Call the subpaths β_1^d, β_2^d, and β_3^d the *internal geodesics*.

Suppose that at least one of the complementary geodesics, say α_a, is long. Note that $\alpha_a \subset \widetilde{S}_a$ joins x_2 and x_3, where x_2 and x_3 are respectively the intersection points of β_2^d and β_3^d with $\widetilde{S}_a$. Let $p_2 \in \beta_2^d$ and $p_3 \in \beta_3^d$ be points that are δ-close to z in the d_{pel}-metric, where δ is a constant depending only on the hyperbolicity constant of $(\widetilde{F}, d_{\mathrm{pel}})$; see Theorem 3.14. Since $\widetilde{S}_a$ is quasi-isometrically embedded

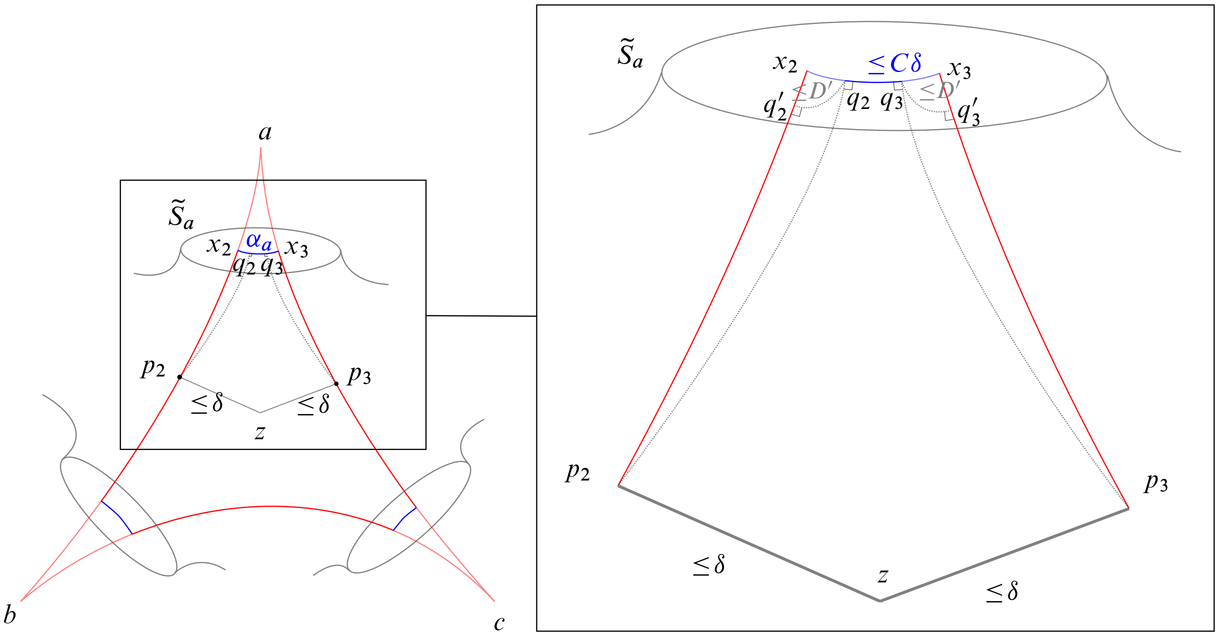

Figure 12: The case when the median lies in a drilled atom.

(Lemma 3.12) in $(\widetilde{M}_r, d_{\text{pel}})$, α_a is a quasigeodesic in $(\widetilde{M}_r, d_{\text{pel}})$. (The quasigeodesic constants are uniform as there are only finitely many possibilities for drilled atoms; see Definition 3.2.) Let π_a denote a nearest-point projection of $(\widetilde{M}_r, d_{\text{pel}})$ onto α_a. Let $\pi_a(p_i) = q_i$ for $i = 2, 3$. Then the d_{pel}-length of the subsegment $\alpha_a(q_2, q_3)$ of α_a joining q_2 and q_3 is at most 8δ (see [22, Lemma 3.2] for instance for a proof of this standard fact). Since α_a is a quasigeodesic in $(\widetilde{M}_r, d_{\text{pel}})$, the length of $\alpha_a(q_2, q_3)$ in the intrinsic metric on $\widetilde{S}_a$ is at most $C\delta$, where C depends only on the uniform quasigeodesic constant for α_a, and is therefore uniform.

Let $[x_2, q_2]$ (resp. $[x_3, q_3]$) denote the geodesic in $\widetilde{S}_a$ joining x_2 and q_2 (resp. x_3 and q_3). Also, let $[q_2, p_2]_{\text{pel}}$ (resp. $[q_3, p_3]_{\text{pel}}$) denote the d_{pel}-geodesic segments in $\widetilde{M}_r$ joining q_2 and p_2 (resp. q_3 and p_3). Then

(1) $\beta_2' = [x_2, q_2] \cup [q_2, p_2]_{\text{pel}}$ is a d_{pel}-quasigeodesic with uniform constants (cf [22, Lemma 3.2]),

(2) $\beta_3' = [x_3, q_3] \cup [q_3, p_3]_{\text{pel}}$ is a d_{pel}-quasigeodesic with uniform constants.

Let β_2 (resp. β_3) denote the subpaths of β_2^d (resp. β_3^d) joining x_2 and p_2 (resp. x_3 and p_3). By relative hyperbolicity of $\widetilde{M}_r$ (Lemma 3.12), and the tracking properties in Lemma 4.4, there exists $q_2' \in \beta_2$ (resp. $q_3' \in \beta_3$) such that the distance (in the unelectrified metric d on $\widetilde{M}_r$) between q_2 and q_2' (resp. q_3 and q_3') is uniformly bounded by a uniform constant D'. Hence $d(q_2', q_3') \leq C\delta + 2D'$.

Let γ_2' (resp. γ_3') be the subpath of γ_2 from a to q_2' (resp. a to q_3'). By Proposition 4.16, γ_2' and γ_3' track each other within a uniform distance C'' of each other.

Removing the initial subpath of β_2^d (resp. β_3^d) between x_2 and q_2' (resp. x_3 and q_3'), we can replace the complementary geodesic α_a by a geodesic α_a' of length at most $C\delta + 2D'$ joining q_2' and q_3'. Carrying out such replacements for all the long complementary geodesics, we obtain a d_{pel}-quasigeodesic hexagon whose complementary geodesics are uniformly bounded in length. Let β_f^1, β_f^2, and β_f^3 be the internal geodesics of the resulting hexagon. By strong relative hyperbolicity of $\widetilde{M}_r$ with respect to $\mathcal{P}_r$, the internal geodesics β_f^1, β_f^2, and β_f^3 satisfy Theorem 4.2(3).

Combining this with the tracking properties of pairs such as γ_2' and γ_3' proved above, using Proposition 4.16, it follows that $\gamma_1, \gamma_2, \gamma_3 \in \mathcal{F}$ satisfy Theorem 4.2(3).

Proof of Theorem 3.1 That the family $\mathcal{F}$ constructed in Section 4.2.7 satisfies Theorem 4.2(1), (2), and (4)–(6) has been established in Section 4.4. The thin triangles condition, ie Theorem 4.2(3), has been checked above in this subsection. Hence, by Theorem 4.2, $\widetilde{F}$ is strongly hyperbolic relative to the collection $\mathcal{P}$. □

4.6 Relative quasiconvexity

Let S_0 be a drilled fiber in E. Further, after isotoping the drilled curves if necessary, we can assume that the collection of drilled curves $\sigma_1, \dots, \sigma_k$ in S_0 is maximal, ie no other drilled curve may be isotoped into S_0 in the complement of $\bigcup_{i=1,\dots,k} N_\epsilon(\sigma_i)$.

Definition 4.18 The collection of drilled fibers is *reduced* if the collection of drilled curves in any singular fiber (see Definition 1.2) is maximal in the above sense.

Then $S_0 \setminus \left(\bigcup_i \sigma_i\right)$ consists of finitely many components $\Sigma_1, \dots, \Sigma_m$. By Theorem 2.1, we have the following:

Lemma 4.19 *Each $\pi_1(\Sigma_i)$ is quasiconvex in $\pi_1(E)$.*

Next, we consider F and $(\widetilde{F}, d_{\text{pel}})$. Then we have:

Lemma 4.20 *Let Σ be a component of $S_0 \setminus \left(\bigcup_i \sigma_i\right)$ as above for S_0 a drilled fiber. Then there exists $C \geq 1$ such that any elevation $\widetilde{\Sigma}$ is quasiconvex $(\widetilde{F}, d_{\text{pel}})$.*

Proof The same argument as in Section 3.5.2 in the paragraph "Identifying ρ-thin annuli" identifies the collection of essential annuli (see Definition 3.7) with core curve homotopic to a curve in Σ. Corollary 3.23 now establishes that Σ flares in all directions in the sense of Definition 3.18. Hence, by Corollary 3.19 and Remark 3.20, there exists $C \geq 1$ such that any elevation $\widetilde{\Sigma}$ is quasiconvex $(\widetilde{F}, d_{\text{pel}})$. □

We finally have the following:

Proposition 4.21 *Let Σ be a component of $S_0 \setminus \left(\bigcup_i \sigma_i\right)$ as above for S_0 a drilled fiber. Then there exists $C' \geq 1$ such that any elevation $\widetilde{\Sigma}$ is relatively C'-quasiconvex.*

Proof Let $\mathcal{T}$ denote the Bass–Serre tree of $\widetilde{F}$, $\Pi\colon \widetilde{F} \to \mathcal{T}$ denote the tree of spaces structure, and v be the vertex such that $\widetilde{S}_0 \subset \widetilde{M}_v$. Let $\mathcal{T}_0$ denote the C-neighborhood of v in $\mathcal{T}$, where C is as in Lemma 4.20. Also, let $X_0 = \Pi^{-1}(\mathcal{T}_0)$. By Lemma 4.20 the d_{pel}-geodesic joining a pair of points $x, y \in \widetilde{\Sigma}$ lies in X_0. Further, the proof of Lemma 3.28 establishes that $\widetilde{\Sigma}$ with its intrinsic metric is qi-embedded in (X_0, d_{pel}).

The proof of Theorem 3.1 applied to X_0 establishes strong-relative hyperbolicity of X_0 relative to the collection $\mathcal{P}_0$ consisting of the elements of $\mathcal{P}$ contained in X_0. Since $\widetilde{\Sigma}$ with its intrinsic metric is qi-embedded in (X_0, d_{pel}), the construction of the path family $\mathcal{F}$ in Section 4.2.7 shows that we can take $\eta(x, y)$ to lie on $\widetilde{\Sigma}$ for $x, y \in \widetilde{\Sigma}$. Hence, $\widetilde{\Sigma}$ is relatively C'-quasiconvex for some C'. □

5 Cubulation

Notation to be used in this section:

- η is used to denote subdivision intervals. F_η is the drilled bundle restricted to η. $\{\Sigma_{\eta,i}\}$ are the subsurface components in the complement of the drilled curves of F_η.
- For $\mathcal{K}$ a maximal subgraph of $\mathcal{G}$ not containing any drilled edges, $F(\mathcal{K})$ denotes the drilled subbundle of F restricted to $\mathcal{K}$.
- $\mathcal{G}_{\mathcal{K}}$ is the canonical reduced form (see Definition 5.2) of $\mathcal{G}$.

We refer the reader to [14; 33; 34] for details on virtually special CAT(0) cube-complexes. However, before attempting to cubulate $G = \pi_1(F)$, we first describe another graph of groups structure on G.

5.1 Another graph of groups structure on G

We construct a new graph of groups structure on $\pi_1(F)$, with a new underlying graph $\mathcal{G}_0$ as follows.

Assume that the collection of drilled fibers is reduced in the sense of Definition 4.18. The vertex groups of $\mathcal{G}_0$ are all isomorphic to $H = \pi_1(S)$, and consist of the following:

(1) For every singular fiber S_v (see Definition 1.2), necessarily undrilled by construction, we have a vertex group G_v isomorphic to H.

(2) For a drilled edge e, let $S_{x_1}, \dots, S_{x_p}$ denote the drilled fibers where x_i are points in order along e, where e has initial and final vertices v_i and v_f. Interpolate vertices $v_i = v_0, v_1, \dots, v_p = v_f$ where v_i lies between x_i and x_{i+1} for $1 \le i \le p-1$. The points $x_1, \dots, x_p$ will be referred to as *special drilled points*, and $v_1, \dots, v_{p-1}$ will be referred to as *interpolating points*.

The vertices of $\mathcal{G}_0$ consist of the vertices of $\mathcal{G}$ *along with interpolating points*.

We now define the edge spaces. The intervals $[v_i, v_{i+1}]$ will be termed *subdivision intervals*. Note that each subdivision interval contains a unique special drilled point x_{i+1}. Let $\eta = [c, d]$ be a subdivision interval, so that the vertex groups G_c and G_v are isomorphic to H. Now, restrict the drilled bundle F to η to obtain F_η. Then F_η is homeomorphic to $S \times [0, 1]$ after removing ϵ-neighborhoods of a nonempty

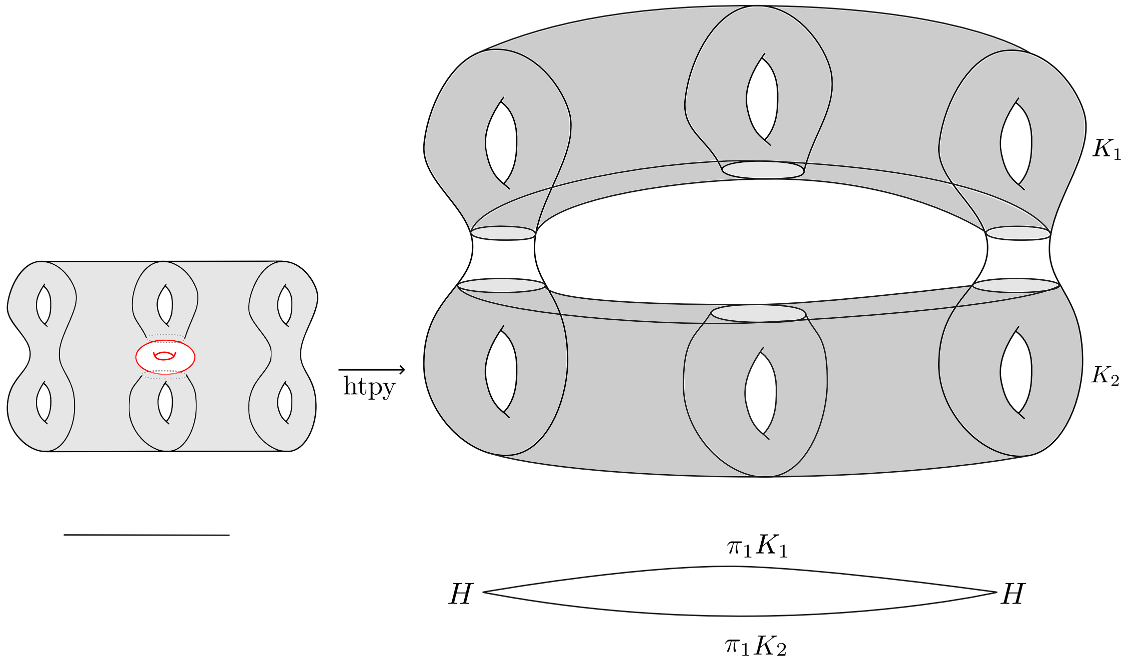

Figure 13: The graph of groups structure for a drilled atom.

family $\{\sigma_i\}$ of disjoint, homotopically distinct, essential simple closed curves on a unique drilled fiber S_{x_i}. Let $\Sigma_{\eta,i}$ for $i = 1, \dots, m$ denote the components of $S \setminus \bigcup_i \sigma_i$. Then $\pi_1(F_\eta)$ admits a graph of groups description with two vertex groups G_c and G_d, each isomorphic to H, and m edge groups, isomorphic to $\pi_1(\Sigma_{\eta,i})$ for $i = 1, \dots, m$. Let $\{e_{\eta,i}, i = 1, \dots, m\}$ denote the resulting edges "living over" the subdivision on interval η. We call these *subdivision edges*. The groups $\pi_1(\Sigma_{\eta,i})$ for $i = 1, \dots, m$ are called the *subdivision edge groups*.

The edges of $\mathcal{G}_0$ are obtained as follows:

(1) If there is an edge of $\mathcal{G}$ that has no special drilled point, leave it as it is in $\mathcal{G}_0$. We refer to these as *undrilled edges* of $\mathcal{G}_0$.

(2) Next, replace each interval with at least one special drilled point by subdivision intervals, and finally each subdivision interval by the subdivision edges that live over it.

Note that the above discussion goes through even when the initial and final vertices of the drilled edge e coincide, so that there is a monodromy map ϕ for the bundle E restricted to the (closed) e. We explicate the inclusion maps for $\pi_1(\Sigma_{\eta,i})$ for $i = 1, \dots, m$ into the vertex groups H_i and H_f corresponding to the vertices v_i and v_f (the initial and final vertices of e). We identify $\pi_1(\Sigma_{\eta,i})$ for $i = 1, \dots, m$ with subgroups of H_i via the product structure on $S \times [c, d)$. Then, modulo this identification, the edge-to-vertex group maps for $\pi_1(\Sigma_{\eta,i})$ with $i = 1, \dots, m$ into G_c is the inclusion. The same holds for $d \neq v_f$. For $v_f = d$, the edge-to-vertex group maps for $\pi_1(\Sigma_{\eta,i})$ with $i = 1, \dots, m$ into G_d is given by inclusion followed by ϕ_*, the map induced by the monodromy ϕ.

It remains to identify the edges and edge groups for the underlying graph $\mathcal{G}_0$. The edges are of exactly two kinds: undrilled edges, and subdivision edges. The edge groups for undrilled edges correspond to H. The edge groups for subdivision edges correspond to the subdivision edge groups $\pi_1(\Sigma_{\eta,i})$ for $i = 1, \dots, m$. Finally, the edge-to-vertex group maps are given as above.

Definition 5.1 A maximal connected undrilled subgraph $\mathcal{K}$ of $\mathcal{G}$ will be called an *undrilled component of* $\mathcal{G}$.

We finally modify $\mathcal{G}_0$ to another graph $\mathcal{G}_{\mathbb{K}}$ by collapsing each undrilled component (see Definition 5.1) $\mathcal{K}$ of $\mathcal{G}$ to a single vertex, ie $\mathcal{G}_{\mathbb{K}}$ is the quotient space obtained from $\mathcal{G}_0$ under the equivalence relation $x \sim y$ if and only if x and y belong to the same undrilled component (see Definition 5.1) $\mathcal{K}$ of $\mathcal{G}$. Let $v_{\mathcal{K}}$ be the resulting vertex of $\mathcal{G}_{\mathcal{K}}$.

The vertex space associated to $v_{\mathcal{K}}$ is then declared to be $\mathcal{F}_{\mathcal{K}}$. The edge to vertex inclusions are given by the composition of

(1) edge to vertex inclusion maps over $\mathcal{G}_0$, composed with
(2) inclusions of vertex spaces over $v \in \mathcal{G}_0$ to the vertex spaces over $v_{\mathcal{K}} \in \mathcal{G}_{\mathbb{K}}$, where $\mathcal{K}$ is the undrilled component (see Definition 5.1) of $\mathcal{G}$ containing v.

Definition 5.2 $\mathcal{G}_{\mathcal{K}}$ will be called the *canonical reduced form* of $\mathcal{G}$ for F.

We conclude this subsection with the following observation, which follows from the above construction.

Lemma 5.3 *The edge spaces of $\mathcal{G}_{\mathcal{K}}$ are precisely the components of $S_d \setminus \bigcup_i \sigma_i$, where S_d ranges over a reduced collection of drilled fibers* (see Definition 4.18).

5.2 Cubulating drilled bundles

We shall need the following theorem due to Wise:

Theorem 5.4 [34, Theorem 15.1] *Let G be a group satisfying the following*:

(1) *G is hyperbolic relative to virtually abelian subgroups.*
(2) *G splits as a graph of groups Γ where each edge group is relatively quasiconvex in G.*
(3) *Each vertex group is virtually special.*
(4) *For each edge e, the edge group G_e has trivial intersection with each Z^2 in the fundamental group of the graph of groups $\Gamma - e$.*

Then G is the fundamental group of a virtually special cube complex.

For us, $G = \pi_1(F)$, and let $\Pi_0 \colon F \to \mathcal{G}$ denote the natural projection.

Theorem 5.5 *Suppose that the graph $\mathcal{G}_{\mathbb{K}}$ is the canonical reduced form of $\mathcal{G}$. Suppose further that for every undrilled component $\mathcal{K}$* (see Definition 5.1) *of $\mathcal{G}$, $\pi_1(\mathcal{F}_{\mathcal{K}})$ is virtually special cubulable. Then the group $G = \pi_1(F)$ is virtually special cubulable.*

Proof Theorem 3.1 proves that $\widetilde{F}$ is strongly hyperbolic relative to $\mathcal{P}$. Since the stabilizer of each $P \in \mathcal{P}$ is $\mathbb{Z} + \mathbb{Z}$, Theorem 5.4(1) is satisfied.

By Lemma 5.3, the edge spaces over $\mathcal{G}_{\mathcal{K}}$ are given by the components of $S_d \setminus \bigcup_i \sigma_i$, where S_d ranges over a reduced collection of drilled fibers (see Definition 4.18). The fundamental groups of these components are relatively quasiconvex in G by Proposition 4.21. Hence Theorem 5.4(2) is satisfied.

Theorem 5.4(3) follows from the hypothesis that $\pi_1(\mathcal{F}_{\mathcal{K}})$ is virtually special cubulable.

Theorem 5.4(4) follows from the hypothesis that $\mathcal{G}_{\mathbb{K}}$ is the canonical reduced form (see Definition 5.2) of $\mathcal{G}$. Indeed, this hypothesis guarantees that there are no accidental parabolics in the components Σ of $S_d \setminus \bigcup_i \sigma_i$, where S_d ranges over a reduced collection of drilled fibers (see Definition 4.18), ie no nonperipheral essential curve in any Σ is freely homotopic to a drilled curve.

Hence, by Theorem 5.4, G is virtually special cubulable. □

5.3 Examples

We now give examples of surface bundles E over graphs $\mathcal{G}$ (see Definition 1.2) and drilled curves such that the hypotheses of Theorem 5.5 are satisfied:

Example 5.6 Each edge of $\mathcal{G}$ contains a drilled point. In this case, the undrilled components $\mathcal{K}$ are precisely the vertices of $\mathcal{G}$, and the associated spaces are given by the fiber S. Since $\pi_1(S)$ is special cubulable, the hypotheses of Theorem 5.5 are satisfied.

Example 5.7 Undrilled components $\mathcal{K}$ are either non-self-intersecting loops in $\mathcal{G}$ or vertices. The associated vertex spaces are either hyperbolic 3-manifolds M fibering over the circle, or the fiber S. By [1], $\pi_1(M)$ is virtually special and hence the hypotheses of Theorem 5.5 are satisfied. More generally, undrilled components $\mathcal{K}$ could be homotopy equivalent to circles or contractible.

A simple example for Example 5.7 is given by a graph $\mathcal{G}$ with two vertices v_1 and v_2, an edge $[v_1, v_2]$, and a loop at each vertex. Further, the edge $[v_1, v_2]$ has a single drilled fiber S_w with one simple closed curve $\sigma \subset S_w$ drilled.

Example 5.8 Undrilled components $\mathcal{K}$ are homotopy equivalent to a wedge of circles, and the restriction of E over any such $\mathcal{K}$ are examples from [24]. The main theorem of [24] then guarantees that the hypotheses of Theorem 5.5 are satisfied.

More generally, components $\mathcal{K}$ could be a mixture of these cases, ie they could be

(1) contractible, in which case the associated vertex space is homotopy equivalent to S,

(2) homotopy equivalent to a circle, in which case the associated vertex space is homotopy equivalent to hyperbolic 3-manifolds M fibering over the circle,

(3) homotopy equivalent to a wedge of circles, with the associated vertex space homotopy equivalent to one of the examples from [24].

6 Virtual algebraic fibering

Notation to be used in this section:

- K_η denotes a homotoped copy of an elementary drilled atom F_η, being the union of one torus T_i glued to S along σ_i each for every drilled curve σ_i of F_η.

Definition 6.1 A finitely generated group G is said to *virtually algebraically fiber* if there exists a finite-index subgroup G_1 of G such that G_1 admits a surjective homomorphism to $\mathbb{Z}$ with finitely generated kernel.

A theorem of Kielak [20] gives the following criterion for virtual algebraic fibering.

Theorem 6.2 [20] *Let G admit a geometric action on a* $\mathrm{CAT}(0)$ *cube complex. Then the following are equivalent*:

(1) *G virtually algebraically fibers.*

(2) *the first l^2 Betti number $\beta_1^{(2)}(G)$ equals zero.*

6.1 Vanishing first l^2 Betti number

The aim of this subsection is to show:

Proposition 6.3 *Let F be a drilled surface bundle over a finite graph $\mathcal{G}$* (see Definition 1.2). *Let $G = \pi_1(F)$. Then $\beta_1^{(2)}(G) = 0$.*

As an immediate consequence of Proposition 6.3 and Theorem 6.2, we have the following:

Theorem 6.4 *Let F be a drilled surface bundle over a finite graph $\mathcal{G}$* (see Definition 1.2) *satisfying the hypotheses of Theorem 5.5. Then $G = \pi_1(F)$ virtually algebraically fibers.*

To prove Proposition 6.3, we shall need a couple of results. A fundamental theorem of Lott and Lück gives the following:

Theorem 6.5 [21, Theorem 0.1] *Let M_r be a drilled atom of F* (see Definition 3.2), *and $H_r = \pi_1(M_r)$. Then $\beta_1^{(2)}(H_r) = -\chi(M_r)$.*

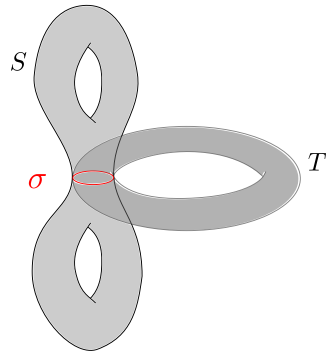

Figure 14: A homotoped elementary drilled atom.

We shall also use the following theorem of Fernós and Valette:

Theorem 6.6 [12, Theorem 1.1] *Let G be a graph of groups with at least one edge, such that all vertex groups satisfy $\beta_1^{(2)}(G_v) = 0$. If every edge group is infinite, and for every edge-to-vertex group inclusion the edge group is of infinite index in the vertex group, then $\beta_1^{(2)}(G) = 0$.*

Recall from Section 5.1 that if η is a subdivision interval, the restriction F_η of F to η has a simple topology: F_η is homeomorphic to $S \times [0,1]$ after removing ϵ-neighborhoods of a nonempty family $\{\sigma_i\}$ of disjoint homotopically distinct essential simple closed curves on a unique drilled fiber S_x. In this subsection, we shall refer to such an F_η as an *elementary drilled atom*.

Let $\sigma_1, \dots, \sigma_k \subset S_x$ denote the drilled curves. Let $\boldsymbol{T}_1, \dots, \boldsymbol{T}_k$ denote k-copies of the standard torus $S^1 \times S^1$. Also, let K_η denote the 2-complex obtained as a quotient space of $S_x \sqcup \bigcup_i \boldsymbol{T}_i$ by identifying $S^1 \times \{0\} \subset \boldsymbol{T}_i$ with σ_i via a homeomorphism (see Figure 14). Then F_η is homotopy equivalent to K_η. Further, $\pi_1(F_\eta) = \pi_1(K_\eta)$ has a graph of groups description, where the underlying graph $\mathcal{G}$ has $k+1$ vertices, $0, \dots, k$ say, and

(1) $\mathcal{G}$ is a tree with one root vertex 0, and all other vertices $1, \dots, k$ are connected to 0 by an edge each,
(2) the vertex group G_0 equals $\pi_1(S_x)$,
(3) for $i = 1, \dots, k$, each vertex group G_i equals $\mathbb{Z} \oplus \mathbb{Z}$,
(4) each edge group is $\mathbb{Z}$.

Lemma 6.7 *Let L be a group with $\beta_1^{(2)}(L) = 0$. Let $H \subset L$ be a subgroup isomorphic to $\pi_1(S_x)$. Let $G = L *_H \pi_1(K_\eta)$, where H is identified via an automorphism with $G_0 = \pi_1(S_x)$. Then $\beta_1^{(2)}(G) = 0$.*

Proof G admits a graph of groups description, where the underlying graph $\mathcal{G}$ has $k+1$ vertices, $0, \dots, k$ say, and

(1) $\mathcal{G}$ is a tree with one root vertex 0, and all other vertices $1, \dots, k$ are connected to 0 by an edge each,

(2) the vertex group G_0 equals L,

(3) for $i = 1, \dots, k$, each vertex group G_i equals $\mathbb{Z} \oplus \mathbb{Z}$,

(4) each edge group is $\mathbb{Z}$.

Then G satisfies the hypotheses of Theorem 6.6. Hence, $\beta_1^{(2)}(G) = 0$. □

Lemma 6.8 *Let F be a drilled surface bundle* (see Definition 1.2) *over a finite graph $\mathcal{G}$ such that $\mathcal{G}$ is homotopy equivalent to a circle. Let $G = \pi_1(F)$. Then $\beta_1^{(2)}(G) = 0$.*

Proof Let $\mathcal{C} \subset \mathcal{G}$ be a cycle with no repeated vertices. Then $\mathcal{G} = \mathcal{C} \cup \bigcup_{i=1,\dots,k} \mathcal{T}_i$, where

(1) each $\mathcal{T}_i$ is a finite tree

(2) $\mathcal{T}_i$ intersects $\mathcal{C}$ at a single point p_i,

(3) $\mathcal{T}_i \setminus \{p_i\}$ is disjoint from $\bigcup_{j \neq i} \mathcal{T}_j$.

Each edge of $\mathcal{T}_i$ for $i = 1, \dots, k$, can be subdivided as in Section 5.1, so that each edge is a subdivision edge. In particular, after such subdivision, for every edge η of $\mathcal{T}_i$ for $i = 1, \dots, k$, F_η is an elementary drilled atom.

Let $F_{\mathcal{C}}$ denote the restriction of F to the cycle $\mathcal{C}$. Then $F_{\mathcal{C}}$ is a 3-manifold whose boundary components are all tori. Hence $\chi(F_{\mathcal{C}}) = 0$. By Theorem 6.5, $\beta_1^{(2)}(\pi_1(F_{\mathcal{C}})) = 0$.

Proceeding by induction on the number of elementary drilled atoms in F, and applying Theorem 6.6 inductively, we conclude that $\beta_1^{(2)}(G) = 0$. □

Proof of Proposition 6.3 The first part of the argument in Section 5.1 allows us to modify $\mathcal{G}$ to a graph where each edge group is isomorphic to $\pi_1(S)$. Thus, assume without loss of generality that each edge in $\mathcal{G}$ has edge group isomorphic to $\pi_1(S)$.

Let $\mathcal{G}_0$ denote a maximal subgraph of $\mathcal{G}$ such that $\mathcal{G}_0$ is homotopy equivalent to a circle. Let F_0 denote the restriction of F to $\mathcal{G}_0$. Let $L = \pi_1(F_0)$. Let n denote the number of edges of $(\mathcal{G} \setminus \mathcal{G}_0)$. Then $G = \pi_1(F)$ admits a new graph of groups decomposition, where the base graph has one vertex w and n loops, with $G_w = L$, and each edge group isomorphic to $\pi_1(S)$.

By Lemma 6.8, $\beta_1^{(2)}(L) = 0$. Since $\mathcal{G}_0$ is homotopy equivalent to a circle, each edge group is of infinite index in the vertex group L. Hence, by Lemma 6.7, $\beta_1^{(2)}(G) = 0$. □

6.2 Questions

In this paper, we have only drilled simple closed curves σ in fibers S_x. A similar drilling operation could be carried out even when a realization of σ as a geodesic in the fiber S has self-intersections. This can be done by homotoping σ slightly in a product neighborhood of S_x to convert σ into a knot in E.

Note that this is a noncanonical operation, as we have a choice of over- and under-crossings at every self-intersection point. At any rate, after such a homotopy, the resulting knot can be drilled. We assume, however, that σ represents a primitive element of $H = \pi_1(S)$. We call this process *generalized drilling*, and the resulting F a *generalized drilled bundle*. Relative hyperbolicity of the generalized drilled G follows by essentially the same argument as in Theorem 3.1.

Question 6.9 *Let $\Gamma = \pi_1(E)$ be hyperbolic. Let F be a generalized drilled bundle obtained by the above generalized drilling operation applied to E. Is $G(:= \pi_1(F))$ cubulable?*

Acknowledgments

Mj would like to thank Daniel Groves and Jason Manning for helpful conversations. We also thank Daniel Groves for comments and corrections to an earlier draft. Finally, we thank the referee for many constructive suggestions and a careful reading of an earlier draft.

Both authors were supported by the Department of Atomic Energy, Government of India, under project 12-R&D-TFR-5.01-0500 and also by an endowment of the Infosys Foundation. Mj was also supported in part by a DST JC Bose Fellowship. Mj was supported in part by the Institut Henri Poincaré (UAR 839 CNRS-Sorbonne Universite), LabEx CARMIN, ANR-10-LABX-59-01, during his participation in the trimester program *Groups acting on fractals*, *hyperbolicity and self-similarity*, April–June 2022.

References

[1] **I Agol**, *The virtual Haken conjecture*, Doc. Math. 18 (2013) 1045–1087 MR

[2] **I Belegradek**, *Rigidity and relative hyperbolicity of real hyperbolic hyperplane complements*, Pure Appl. Math. Q. 8 (2012) 15–51 MR

[3] **I Belegradek**, **G C Hruska**, *Hyperplane arrangements in negatively curved manifolds and relative hyperbolicity*, Groups Geom. Dyn. 7 (2013) 13–38 MR

[4] **N Bergeron**, **D T Wise**, *A boundary criterion for cubulation*, Amer. J. Math. 134 (2012) 843–859 MR

[5] **M Bestvina**, **M Feighn**, *A combination theorem for negatively curved groups*, J. Differential Geom. 35 (1992) 85–101 MR Correction in 43 (1996) 783–788

[6] **B H Bowditch**, *Intersection numbers and the hyperbolicity of the curve complex*, J. Reine Angew. Math. 598 (2006) 105–129 MR

[7] **B H Bowditch**, *Relatively hyperbolic groups*, Int. J. Algebra Comput. 22 (2012) art. id. 1250016 MR

[8] **R D Canary**, *A covering theorem for hyperbolic 3-manifolds and its applications*, Topology 35 (1996) 751–778 MR

[9] **S Dowdall**, **R P Kent, IV**, **C J Leininger**, *Pseudo-Anosov subgroups of fibered 3-manifold groups*, Groups Geom. Dyn. 8 (2014) 1247–1282 MR

[10] **B Farb**, *Relatively hyperbolic groups*, Geom. Funct. Anal. 8 (1998) 810–840 MR

[11] **B Farb**, **L Mosher**, *Convex cocompact subgroups of mapping class groups*, Geom. Topol. 6 (2002) 91–152 MR

[12] **T Fernós**, **A Valette**, *The Mayer–Vietoris sequence for graphs of groups, property (T), and the first ℓ^2-Betti number*, Homology Homotopy Appl. 19 (2017) 251–274 MR

[13] **D Groves**, **P Haïssinsky**, **J F Manning**, **D Osajda**, **A Sisto**, **G S Walsh**, *Drilling hyperbolic groups*, preprint (2024) arXiv 2406.14667

[14] **F Haglund**, **D T Wise**, *Special cube complexes*, Geom. Funct. Anal. 17 (2008) 1551–1620 MR

[15] **U Hamenstädt**, *Word hyperbolic extensions of surface groups*, preprint (2005) arXiv math/0505244

[16] **U Hamenstädt**, *Geometry of the complex of curves and of Teichmüller space*, from "Handbook of Teichmüller theory, I", IRMA Lect. Math. Theor. Phys. 11, Eur. Math. Soc., Zürich (2007) 447–467 MR

[17] **J Kahn**, **V Markovic**, *Immersing almost geodesic surfaces in a closed hyperbolic three manifold*, Ann. of Math. 175 (2012) 1127–1190 MR

[18] **M Kapovich**, *Hyperbolic manifolds and discrete groups*, Progr. Math. 183, Birkhäuser, Boston, MA (2001) MR

[19] **R P Kent, IV**, **C J Leininger**, *Shadows of mapping class groups: capturing convex cocompactness*, Geom. Funct. Anal. 18 (2008) 1270–1325 MR

[20] **D Kielak**, *Residually finite rationally solvable groups and virtual fibring*, J. Amer. Math. Soc. 33 (2020) 451–486 MR

[21] **J Lott**, **W Lück**, *L^2-topological invariants of 3-manifolds*, Invent. Math. 120 (1995) 15–60 MR

[22] **M Mitra**, *Cannon–Thurston maps for trees of hyperbolic metric spaces*, J. Differential Geom. 48 (1998) 135–164 MR

[23] **M Mj**, *Tight trees and model geometries of surface bundles over graphs*, J. Lond. Math. Soc. 102 (2020) 1178–1222 MR

[24] **M Mj**, *Cubulating surface-by-free groups*, J. Topol. 17 (2024) art. id. e70011 MR With an appendix by J Manning, M Mj, and M Sageev

[25] **M Mj**, **A Pal**, *Relative hyperbolicity, trees of spaces and Cannon–Thurston maps*, Geom. Dedicata 151 (2011) 59–78 MR

[26] **M Mj**, **K Rafi**, *Algebraic ending laminations and quasiconvexity*, Algebr. Geom. Topol. 18 (2018) 1883–1916 MR

[27] **M Mj**, **L Reeves**, *A combination theorem for strong relative hyperbolicity*, Geom. Topol. 12 (2008) 1777–1798 MR

[28] **M Mj**, **P Sardar**, *A combination theorem for metric bundles*, Geom. Funct. Anal. 22 (2012) 1636–1707 MR

[29] **G P Scott**, **G A Swarup**, *Geometric finiteness of certain Kleinian groups*, Proc. Amer. Math. Soc. 109 (1990) 765–768 MR

[30] **A Sisto**, *On metric relative hyperbolicity*, preprint (2012) arXiv 1210.8081

[31] **P Susskind**, **G A Swarup**, *Limit sets of geometrically finite hyperbolic groups*, Amer. J. Math. 114 (1992) 233–250 MR

[32] **W P Thurston**, *The geometry and topology of three-manifolds*, lecture notes, Princeton University (1979) Available at https://url.msp.org/gt3m

[33] **D T Wise**, *From riches to raags: 3-manifolds, right-angled Artin groups, and cubical geometry*, CBMS Region. Conf. Ser. Math. 117, Amer. Math. Soc., Providence, RI (2012) MR

[34] **D T Wise**, *The structure of groups with a quasiconvex hierarchy*, Ann. of Math. Stud. 209, Princeton Univ. Press (2021) MR

MM, BN: *School of Mathematics, Tata Institute of Fundamental Research Mumbai, India*

mahan@math.tifr.res.in, biswajit@math.tifr.res.in

Received: 11 December 2023 Revised: 15 August 2024

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2025 is US $760/year for the electronic version, and $1110/year (+$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

mathematical sciences publishers

nonprofit scientific publishing

https://msp.org/

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 25 Issue 7 (pages 3789–4436) 2025

Algebras for enriched ∞-operads 3789
RUNE HAUGSENG
An overtwisted convex hypersurface in higher dimensions 3813
RIVER CHIANG and KLAUS NIEDERKRÜGER-EID
Presheaves of groupoids as models for homotopy types 3833
LÉONARD GUETTA
Product and coproduct on fixed point Floer homology of positive Dehn twists 3875
YUAN YAO and ZIWEN ZHAO
Multitwists in big mapping class groups 3921
GEORGE DOMAT, FEDERICA FANONI and SEBASTIAN HENSEL
Cofibrantly generated model structures for functor calculus 3931
LAUREN BANDKLAYDER, JULIA E BERGNER, RHIANNON GRIFFITHS, BRENDA JOHNSON and REKHA SANTHANAM
Endomorphisms of Artin groups of type D 3975
FABRICE CASTEL and LUIS PARIS
Linear bounds of the crosscap number of knots 4009
ROB MCCONKEY
Atiyah–Segal completion for the Hermitian K-theory of symplectic groups 4037
JENS HORNBOSTEL, HERMAN ROHRBACH and MARCUS ZIBROWIUS
A minimality property for knots without Khovanov 2-torsion 4073
ONKAR SINGH GUJRAL and JOSHUA WANG
Local indicability of groups with homology circle presentations 4077
AGUSTÍN NICOLÁS BARRETO and ELÍAS GABRIEL MINIAN
Cubulating drilled bundles over graphs 4095
MAHAN MJ and BISWAJIT NAG
Infinitely many homeomorphic hyperbolic plugs with the same basic sets 4147
FANGFANG CHEN
One-point compactifications of configuration spaces and the self duality of the little disks operad 4163
CONNOR MALIN
Braided multitwists 4185
RODRIGO DE POOL
Positive intermediate Ricci curvature on connected sums 4209
PHILIPP REISER and DAVID J WRAITH
The $\mathbb{S}_n$-equivariant Euler characteristic of the moduli space of graphs 4229
MICHAEL BORINSKY and JOS VERMASEREN
Classification of metric fibrations 4257
YASUHIKO ASAO
Spaces over BO are thickened manifolds 4287
HIRO LEE TANAKA
Representation-graded Bredon homology of elementary abelian 2-groups 4321
MARKUS HAUSMANN and STEFAN SCHWEDE
The $\mathbb{Z}/p$-equivariant cohomology of the genus-zero Deligne–Mumford space with $1+p$ marked points 4341
DAIN KIM and NICHOLAS WILKINS
Geometry of free extensions of free groups via automorphisms with fixed points on the complex of free factors 4357
PRITAM GHOSH and FUNDA GÜLTEPE
Coarse cohomology of configuration space and coarse embedding 4391
ARKA BANERJEE
Correction to the article An algebraic model for finite loop spaces 4427
CARLES BROTO, RAN LEVI and BOB OLIVER