CUBULATING HYPERBOLIC FREE-BY-CYCLIC GROUPS: THE IRREDUCIBLE CASE

MARK F. HAGEN AND DANIEL T. WISE

Abstract. Let V be a finite graph and let $\phi : V \to V$ be an irreducible train track map whose mapping torus has word-hyperbolic fundamental group G. Then G acts freely and cocompactly on a CAT(0) cube complex. Hence, if F is a finite-rank free group and $\Phi : F \to F$ an irreducible monomorphism so that $G = F*_\Phi$ is word-hyperbolic, then G acts freely and cocompactly on a CAT(0) cube complex. This holds in particular if Φ is an irreducible automorphism with $G = F \rtimes_\Phi \mathbb{Z}$ word-hyperbolic.

Contents

Date: August 18, 2015.
2010 *Mathematics Subject Classification.* Primary: 20F65; Secondary: 57M20.
Key words and phrases. free-by-cyclic group, CAT(0) cube complex, train track map.

Table of notation

Since there is a great deal of recurring notation, the reader may find the following table helpful:

$\mathbf{A}:\overline{W}\to\widetilde{X}$	Approximation map
$\mathbf{A}(\overline{W})$	Approximation of $\overline{W}$
d_i, d'_i	Primary busts in V, E (respectively)
d_{ij}, d'_{ij}	Secondary busts in V, E (respectively)
$\mathsf{d}_{\widetilde{Y}}$	Metric inside subspace $\widetilde{Y}\subseteq\widetilde{X}$ (e.g., $\widetilde{Y}=\widetilde{V}_n$)
d_L	Weighted metric on $\widetilde{X}^\bullet_L$
d_L	Metric on $\mathbb{R}$-tree $\mathcal{Y}$
E	Image in X_L of $V\times\{\frac{1}{2}\}$
$\widetilde{E}_n$	Preimage of E in $\widetilde{X}$ at position $n+\frac{1}{2}$
$\widetilde{E}_{nL}$	$q_L^{-1}(nL+\frac{1}{2})$
$\lvert e\rvert$	Weight of edge e
$\widetilde{K}$	Knockout
κ_1, κ_2	Quasi-isometry constants for piecewise geodesics (depends on L)
λ, λ_i	Quasi-isometry constants for forward ladders
$\mathfrak{M}$	Transition matrix for train track map ϕ
μ_1, μ_2	Quasi-isometry constants for uniform sub-quasiconvexity
$\mathbf{N}$	Nucleus
$N(\sigma)$	Forward ladder of forward path σ
$N(Y)$	For $Y\subseteq\widetilde{X}$ or $Y\subseteq\widetilde{X}_L$, smallest subcomplex containing Y
$\Phi:F\to F$	Injection on finite-rank free group F
$\phi:V\to V$	Combinatorial map (later train track map) inducing Φ
ϖ	Expansion constant of train track map ϕ
$q_L:\widetilde{X}_L\to\mathbf{R}_L$	Coordinate map onto combinatorial line; $q=q_1$
$\tilde{\phi}^L:\widetilde{X}_L\to\widetilde{X}_L$	Forward flow map
$\rho:\widetilde{X}\to\mathcal{Y}$	Natural map from $\widetilde{X}$ to $\mathbb{R}$-tree
$\widetilde{\varrho}_L:\widetilde{X}^\bullet_L\to\widetilde{X}$	Map "folding" subdivided star levels to rooted tree levels
$S_i^\pm$	Slopes
σ, $\sigma_M(x)$	Forward path, forward path of length M determined by x)
$T^o_L(x)$	Level (union of forward paths ending at x) in $\widetilde{X}$
$T^{o\pm}_i$	Level-part of a tunnel
$T^\pm_i$	Tunnel
$\widehat{W}^\bullet$	Graph obtained by attaching tunnels to nucleus
$W^\bullet$	Space obtained from $\widehat{W}^\bullet$ by folding levels according to ϱ_L
W	Immersed wall (component of $W^\bullet$ in X)
$\overline{W}$	Wall (image of $\widetilde{W}\to\widetilde{X}$)
X_L	Mapping torus of ϕ^L ($L\geq 1$); $X=X_1$
$\widetilde{X}^\bullet_L$	Subdivision of $\widetilde{X}_L$ from pulling back cell structure of $\widetilde{X}$ via $\widetilde{\varrho}_L$
V	Base of mapping torus (a finite graph)
$\widetilde{V}_n$	Preimage of V in $\widetilde{X}$ at q-coordinate n

Introduction

The goal of this paper is to prove the following theorem:

Theorem A. *Let F be a finite-rank free group and let $\Phi : F \to F$ be an irreducible automorphism, and suppose that $G = F \rtimes_\Phi \mathbb{Z}$ is word-hyperbolic. Then G acts freely and cocompactly on a CAT(0) cube complex.*

This result is a special case of Corollary 6.21, which handles the more general case of a hyperbolic ascending HNN extension of a free group by an irreducible endomorphism.

Theorem A provides a widely-studied class of hyperbolic groups for which Gromov's question (see [Gro87]) of whether hyperbolic groups are CAT(0) has a positive answer, but goes further, since nonpositively-curved cube complexes enjoy numerous useful properties beyond having universal covers that admit a CAT(0) metric. For example, combining Theorem A with a result of [Ago12] shows that groups G of the type described in Theorem A are virtually special in the sense of [HW08] and therefore virtually embed in a right-angled Artin group. This implies that G has several nice structural features, including $\mathbb{Z}$-linearity.

A group $G \cong F \rtimes_\Phi \mathbb{Z}$ is word-hyperbolic exactly when Φ is atoroidal [BF92, Bri00], so that Theorem A applies to all mapping tori of irreducible, atoroidal automorphisms of free groups. More generally, ascending HNN extensions are hyperbolic precisely if they have no Baumslag-Solitar subgroups [Kap00].

We actually prove the following more general statement:

Theorem B. *Let $\phi : V \to V$ be a train track map of a finite graph V. Suppose that ϕ is π_1-injective and that each edge of V is expanding. Moreover, suppose that the transition matrix $\mathfrak{M}$ of ϕ is irreducible and that the mapping torus X of ϕ has word-hyperbolic fundamental group G. Then G acts freely and cocompactly on a CAT(0) cube complex.*

Our CAT(0) cube complex arises by applying Sageev's construction [Sag95] to a family of walls in the universal cover $\widetilde{X}$ of X. To ensure that the resulting action of G on the dual cube complex is proper and cocompact, we show that there is a quasiconvex wall separating any two points in ∂G, thus verifying the cubulation criterion in [BW13]. As train track maps are central to the proof that there are many walls in this sense, our results build upon the work of Bestvina, Feighn, and Handel in [BH92, BFH97].

It appears likely that in the case where ϕ is π_1-surjective, the hypothesis that ϕ is irreducible can be removed, and we are currently working on developing the methodology in this paper to generalize Theorem A to all hyperbolic mapping tori of free group automorphisms[1].

Moreover, for the construction of immersed walls in X, hyperbolicity of G plays a minor role. It is therefore natural to wonder which free-by-cyclic groups admit free actions on CAT(0) cube complexes arising from immersed walls constructed essentially as in Section 3. If Φ is fully irreducible and G is not hyperbolic, then Φ is represented by a homeomorphism of a surface, by [BH92, Thm. 4.1]. Consequently, in this case G acts freely on a locally finite, finite-dimensional CAT(0) cube complex [PW]. It is reasonable to conjecture that in general, if $G = F \rtimes_\Phi \mathbb{Z}$ is hyperbolic relative to virtually abelian subgroups, then G acts freely on a locally finite, finite-dimensional CAT(0) cube complex. The techniques in this paper are largely portable to that context. However, one cannot expect to obtain cocompact cubulations for general free-by-cyclic groups. Indeed, Gersten's group $\langle a, b, c, t \mid a^t = a,\, b^t = ba,\, c^t = ca^2\rangle$ is free-by-cyclic but does not act metrically properly by semisimple isometries

[1]We posted a tortuous generalization eight months after submitting this paper; see [HW14].

on a CAT(0) space [Ger94], and hence Gersten's group cannot act freely on a locally finite, finite-dimensional CAT(0) cube complex. Nevertheless, Gersten's group does act freely on an infinite-dimensional CAT(0) cube complex [Wis], so there is still much work to do in this direction.

Summary of the paper. In Sections 1 and 2, we describe the mapping torus X and introduce some features – *levels* and *forward ladders* – that play a role in the construction of immersed walls in X.

In Section 3, we describe *immersed walls* $W \to X$ when $\phi : V \to V$ is an arbitrary π_1-injective map sending vertices to vertices and edges to combinatorial paths, under the additional assumptions that no power of ϕ maps an edge to itself and $\pi_1 X$ is hyperbolic. The immersed wall W is homeomorphic to a graph and has two parts, the *nucleus* and the *tunnels*, and is determined by a positive integer L and a collection of sufficiently small intervals $d_i \subset V$, each contained in the interior of an edge. The nucleus is obtained by removing from V each *primary bust* d_i, along with its ϕ^L-preimage. The tunnels are "horizontal" immersed trees joining endpoints of d_i to endpoints of its preimage. Let $\widetilde{W} \to \widetilde{X}$ be a lift of the universal cover of W and let $\overline{W} \subset \widetilde{X}$ be its image. Since $W \to X$ is not in general π_1-injective, $\widetilde{W} \to \overline{W}$ is not in general an isomorphism. However, under suitable conditions described in Section 4, $\overline{W}$ is a wall in $\widetilde{X}$ whose stabilizer is a quasiconvex free subgroup of G. The immersed walls in X are analogous to the "cross-cut surfaces" introduced in [CLR94], and Dufour used these to cubulate hyperbolic mapping tori of self-homeomorphisms of surfaces [Duf12].

Remark 1 (Wall-approximations)**.** Although the goal of the paper is to produce and understand walls, we study these walls by means of a contrived object of which we now warn the reader. Specifically, to prove that the stabilizer of $\overline{W}$ is a quasiconvex subgroup, we introduce a "combinatorial approximation" $\mathbf{A}(\overline{W})$ of $\overline{W}$. This is a subspace of $\widetilde{X}$ obtained from $\overline{W}$ by, roughly, applying the "forward flow" $\widetilde{X} \to \widetilde{X}$ arising from ϕ. The reason for doing this is that it is difficult to show that $\overline{W}$ is quasiconvex, since the tunnels are not uniformly quasiconvex; they are rooted trees whose branches are paths that can fellow-travel in an uncontrolled fashion. Passing to the approximation $\mathbf{A}(\overline{W})$ folds each tunnel into a single, uniformly quasiconvex path. This allows us to establish uniform quasiconvexity of $\mathbf{A}(\overline{W})$, whose stabilizer coincides with that of the original wall $\overline{W}$. We advise the reader to be alert to this distinction, which we view as the main technical difficulty in the proof.

Section 5 and 6 are devoted to the proof of Theorem B. We use a continuous surjection $\widetilde{X} \to \mathcal{Y}$ to an $\mathbb{R}$-tree that arises in the case where ϕ is a train track representative of an irreducible automorphism (see [BFH97]).

Acknowledgements. We thank the referees for extremely useful and detailed reports containing many helpful comments and corrections that significantly improved this text. We also thank an anonymous referee for creating the table of notation. This is based upon work supported by the National Science Foundation under Grant Number NSF 1045119 and by NSERC.

1. Mapping tori

Let V be a finite connected graph based at a vertex v, and let $\phi : V \to V$ be a continuous, basepoint-preserving map such that $\phi(w)$ is a vertex for each vertex w of V, and such that $\phi(e)$ is a *combinatorial path* in V for each edge e of V. This means that there is a subdivision of e such that vertices of the subdivision map to vertices and whose open edges map homeomorphically to open edges. We also assume that ϕ is parametrized so that these homeomorphisms are linear. Moreover, we require that the map $\Phi : F \to F$ induced by ϕ is injective, where $F \cong \pi_1 V$ is a finite-rank free group. We note that any injective $\Phi : F \to F$ is represented by such a map ϕ.

The reader should have in mind the case where Φ is an irreducible automorphism of F and ϕ is a train track map representing Φ, in the sense of [BH92] (we refer the reader to Section 6.2 for more on train track maps and how we use them):

Definition 1.1 (Train track map)**.** $\phi : V \to V$ is a *train track map* if for all edges e of V and all $n \geq 0$, the path $\phi^n(e) \to X$ is immersed.

For an integer $L \geq 1$, let X_L be obtained from $V \times [0, L]$ by identifying (x, L) with $(\phi^L(x), 0)$ for each $x \in V$, so that X_L is the mapping torus of ϕ^L, and let $X = X_1$. See Figure 1. Let $G = \pi_1 X$ and let $G_L = \pi_1 X_L$ for each $L \geq 1$. Note that if Φ is surjective then $G_L \cong F \rtimes_{\Phi^L} \mathbb{Z}$.

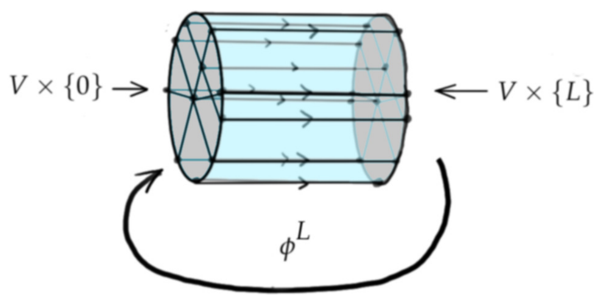

FIGURE 1. The mapping torus X_L.

We regard V as a subspace of X_L, and we denote by E the image in X_L of $V \times \{\frac{1}{2}\}$; the space E plays a role in Section 3.

We now describe a cell structure on X_L. Let $V \times [0, L]$ have the product cell structure: its vertices are $V^0 \times \{0, L\}$, its *vertical edges* are the edges of $V \times \{0, L\}$, and its *horizontal edges* are of the form $\{w\} \times [0, L]$, where $w \in V^0$. We direct each horizontal edge $\{w\} \times [0, L]$ from $\{w\} \times \{0\}$ to $\{w\} \times \{L\}$, and horizontal edges of X_L are directed accordingly. The 2-cells of X_L are images of the 2-cells of $V \times [0, L]$, which have the form $e \times [0, L]$, where e is an edge of V.

For each vertex $w \in V^0 \subset X^0_L$, we let t_w denote the unique horizontal edge outgoing from w. When $L = 1$, let $z \in G$ be the element represented by the loop t_v, where v is the ϕ-invariant basepoint of V. Note that conjugation by z induces the monomorphism $\Phi : F \to F$. For each vertical edge e, joining vertices a, b, there is a 2-cell R_e with attaching map $t_b^{-1} e^{-1} t_a \phi^L(e)$, where t_a, t_b are horizontal edges and $\phi^L(e)$ is a combinatorial path in V.

Define a map $\varrho_L : X_L \to X$ as follows. First, ϱ_L restricts to the identity on V. Each horizontal edge t_w of X_L, joining w to $\phi^L(w) \in V^0$, maps to the concatenation of L horizontal edges of X beginning at w. This determines $\varrho_L : X^1_L \to X$. This map extends to the 2-skeleton by mapping each 2-cell R_e of X_L to a disc diagram $D_e \to X$. Specifically, for $0 \leq i \leq L$, the i^{th} component P_i of the vertical 1-skeleton of D_e is the path $\phi^i(e)$, and there are strips of 2-cells between consecutive vertical components. The boundary path of D_e consists of P_0, P_L, and the horizontal edges of X joining the initial [terminal] point of P_i to the initial [terminal] point of P_{i+1} for $0 \leq i \leq L-1$. Such a diagram D_e is a *long 2-cell* and is depicted in Figure 2. Note that $D_e \to X$ is an immersion when ϕ is a train track map. Otherwise, the paths $P_i \to X$ are not necessarily immersed and the map $D_e \to X$ need not be locally injective. The map $G_L \to G$ induced by ϱ_L embeds G_L as an index-L subgroup of G.

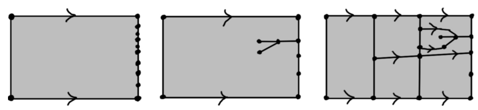

FIGURE 2. A 2-cell of X_L is shown at left, its image in X_L is at the center, and its image in X, which is the image of a long 2-cell in X, is shown at right. Here $L = 3$. Horizontal edges have arrows and all non-arrowed edges are vertical.

The universal cover $\widetilde{X}_L \to X_L$ inherits a cell structure from X_L. Let $\tilde{v} \in \widetilde{X}^0_L$ be a lift of the basepoint $v \in X^0_L$, and let $\widetilde{V}_0$ denote the smallest F-invariant subgraph containing the $\tilde{v}$ component of the preimage of V. Let $\widetilde{V}_{nL} = z^{nL}\widetilde{V}_0$ for $n \in \mathbb{Z}$.

There is a *forward flow* map $\tilde{\phi}_L : \widetilde{X}_L \to \widetilde{X}_L$ defined as follows. For each $p \in V \times \{0\}$, let S_p be the path $\{p\} \times [0, L] \to X_L$. The *horizontal ray* $m_p \to X_L$ at p is the concatenation $S_p S_{\phi^L(p)} S_{\phi^{2L}(p)} \cdots$. For $\tilde{p} \in \widetilde{V}_{nL}$ mapping to p, let $\tilde{m}_{\tilde{p}}$ be the lift of m_p at $\tilde{p}$. For any $\tilde{a} \in \tilde{m}_{\tilde{p}}$, the point $\tilde{\phi}_L(\tilde{a})$ is defined by translating $\tilde{a}$ a positive distance L along $\tilde{m}_{\tilde{p}}$. When $L = 1$, we denote $\tilde{\phi}_L$ by $\tilde{\phi}$.

Let $\mathbf{R}_L$ denote the *combinatorial line* with a vertex for each $nL \in L\mathbb{Z}$ and an edge for each $[nL, nL + L]$ and let $\mathbf{S}_L$ be a circle with a single vertex and a single edge of length L. We define a map $q_L : \widetilde{X}_L \to \mathbf{R}_L$ as follows. There is a map $\bar{q}_L : X_L \to \mathbf{S}_L$ induced by the projection $V \times [0, L] \to [0, L]$. The map $\bar{q}_L$ lifts to the desired map q_L. Note that q_L sends vertical edges to vertices and horizontal edges and 2-cells to edges of $\mathbf{R}_L$. We let $\mathbf{R} = \mathbf{R}_1$, $\mathbf{S} = \mathbf{S}_1$, and $q = q_1$. The map q is the *(horizontal) coordinate projection.*

Let $\widetilde{E}_{nL} = q_L^{-1}(nL + \frac{1}{2})$. Each horizontal edge $t_w \cong \{w\} \times [0, L] \subset \widetilde{X}_L$ intersects $\widetilde{E}_{nL}$ at the point $\{w\} \times \{\frac{1}{2}\}$ for a unique $n \in \mathbb{Z}$.

1.1. **Metrics and subdivisions.** For each edge e of X, let $|e|$ be a positive real number, with $|t_w| = 1$ for each horizontal edge t_w. The assignment $e \mapsto |e|$ is a *weighting* of X^1, and pulls back to a G-equivariant weighting of $\widetilde{X}^1$, with all horizontal edges having unit weight. Regarding e as a copy of $[0, 1]$, the subinterval $d \cong [a, b] \subset e$ has weight $|d| = (b - a)|e|$. Consider an embedded path $P \to \widetilde{X}^1$ (not necessarily combinatorial). The *length* $|P|$ of P is the sum of the weights of $P \cap e$, where e varies over all edges. This yields a geodesic metric d on $\widetilde{X}^1$ such that $(\widetilde{X}^1, \mathsf{d})$ is quasi-isometric to $\widetilde{X}^1$ with the usual combinatorial path-metric in which edges have unit length.

For each $L \geq 1$, let $\widetilde{X}^\bullet_L$ be the subdivision of $\widetilde{X}_L$ such that the lift $\tilde{\varrho}_L : \widetilde{X}_L \to \widetilde{X}$ of ϱ_L sends open cells homeomorphically to open cells. The resulting map $\widetilde{X}^\bullet_L \to \widetilde{X}$ is an isomorphism on subspaces $\widetilde{V}_{nL}$ and sends 2-cells to long 2-cells. Note that 2-cells of $\widetilde{X}_L$ do not immerse in $\widetilde{X}$ unless ϕ is a train track map. Pulling back weights of edges in $\widetilde{X}$ to $\widetilde{X}^\bullet_L$ yields a metric d_L on $(\widetilde{X}^\bullet_L)^1$ with respect to which $(\widetilde{X}^\bullet_L)^1 \to \widetilde{X}^1$ is a distance-nonincreasing quasi-isometry. We shall work mainly in $\widetilde{X}$, except in Section 5, where it is essential to consider $\widetilde{X}^\bullet_L$. We refer the reader to Figure 5 to see the differences between $\widetilde{X}_L, \widetilde{X}^\bullet_L$, and $\widetilde{X}$.

Beginning in Section 3, we shall assume that G is word-hyperbolic, so that there exists $\delta \geq 0$ such that $(\widetilde{X}^1, \mathsf{d})$ is δ-hyperbolic.

2. Forward ladders and levels

In this section, we define various subspaces of $\widetilde{X}$ needed in the construction and analysis of quasiconvex walls in $\widetilde{X}$ and $\widetilde{X}^\bullet_L$.

Definition 2.1 (Midsegment)**.** Let $R_e \to \widetilde{X}$ be a 2-cell with boundary path $t_b^{-1}e^{-1}t_a\tilde{\phi}(e)$, where e is a vertical edge joining vertices a, b. Regarding R_e as a Euclidean trapezoid with parallel sides of length $|e|$ and $|\tilde{\phi}(e)|$, the *midsegment in* R_e determined by $x \in e$ is the line segment joining x to $\tilde{\phi}(x)$. The *midsegment* in $\widetilde{X}$ determined by x is the image of the midsegment in R_e determined by x under the map $R_e \to \widetilde{X}$, and is denoted m_x. Midsegments are directed so that x is initial and $\tilde{\phi}(x)$ is terminal. The midsegment m_x is *singular* if $\tilde{\phi}(x) \in \widetilde{X}^0$ and *regular* otherwise. In general, $R_e \to \widetilde{X}$ is not an embedding, and there may be distinct $x, y \in e$ with the property that the terminal points of m_x and m_y coincide. Note, however, that the intersection of two midsegments contains at most one point. See Figure 3.

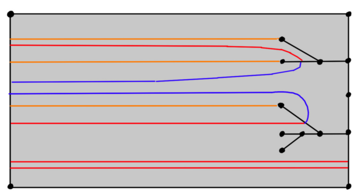

FIGURE 3. Some midsegments in the image of a 2-cell in $\widetilde{X}$.

Definition 2.2 (Forward path, forward ladder)**.** Let $x \in \widetilde{V}_n$ for some $n \in \mathbb{Z}$ and let $M \in \mathbb{Z}$. The *forward path* $\sigma_M(x)$ of length M determined by x is the embedded path that is the concatenation of midsegments starting at x and ending at $\tilde{\phi}^M(x)$. In other words, $\sigma_M(x)$ is isomorphic to the combinatorial interval $[0, M]$, whose vertices are the points $\tilde{\phi}^i(x)$, $0 \leq i \leq M$ and whose edges are the midsegments joining $\tilde{\phi}^i(x)$ to $\tilde{\phi}^{i+1}(x)$. Any path σ of this form is a *forward path.* Note that σ is a directed path with respect to the directions of midsegments in the sense that each internal point in which σ intersects the vertical 1-skeleton of $\widetilde{X}$ has exactly one incoming and one outgoing midsegment.

The forward path σ is *singular* if it contains a vertex and *regular* otherwise.

The *forward ladder* $N(\sigma)$ associated to σ is the smallest subcomplex of $\widetilde{X}$ containing σ. The 1-skeleton $N(\sigma)^1$ plays an important role in many arguments. See Figure 4.

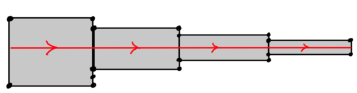

FIGURE 4. A forward ladder for a regular forward path. The forward path is labeled with arrows.

A subgraph Y of $\widetilde{X}^1$ is λ-*quasiconvex* if every geodesic of $\widetilde{X}^1$ starting and ending on Y lies in $\mathcal{N}_\lambda(Y)$. We use the notation $\mathcal{N}_r(Y)$ to denote the closed r-neighborhood of Y.

Proposition 2.3 (Quasiconvexity of forward ladders)**.** *There exist constants $\lambda_1 \geq 1, \lambda_2 \geq 0$ such that for each regular forward path $\sigma \to \widetilde{X}$, the inclusion $N(\sigma)^1 \hookrightarrow \widetilde{X}^1$ is a (λ_1, λ_2)-quasi-isometric embedding. Hence, if $\widetilde{X}^1$ is δ-hyperbolic, there exists $\lambda \geq 0$ such that each $N(\sigma)^1$ is λ-quasiconvex.*

Proof. Let σ join x to $\tilde{\phi}^M(x)$, so that $\sigma = m_x m_{\tilde{\phi}(x)} \cdots m_{\tilde{\phi}^{M-1}(x)}$. Let R_i be the 2-cell containing $m_{\tilde{\phi}^i(x)}$. Then a geodesic P of $N(\sigma)^1$ joining x to $\tilde{\phi}^M(x)$ has the form $P = Q_0 t_0 Q_1 t_1 Q_2 \cdots t_{M-1} Q_M$, where each t_i is a horizontal edge in R_i and each $|Q_i| \leq \max_e\{|\phi(e)|\}$. Since each m_i is a midsegment, $R_i \neq R_j$ for $i \neq j$, whence the coordinate projection $q(P)$ is a combinatorial interval of length M, and the preimage in $N(\sigma)^1$ of each point in $q(P)$ is uniformly bounded. Hence P is a uniform quasigeodesic in $\widetilde{X}^1$. □

In the case that $\widetilde{X}^1$ is δ-hyperbolic, we denote by λ the resulting quasiconvexity constant of the 1-skeleton of a forward ladder. The forward ladder for a singular forward path is also uniformly quasi-isometrically embedded, by an argument very similar to the proof of Proposition 2.3, but we do not require this fact.

Definition 2.4 (Level)**.** Let $x \in \widetilde{V}_n$ and let $L \geq 0$. Note that the preimage $(\tilde{\phi}_L)^{-1}(x)$ is a finite set $\{x_i\}$ in $\widetilde{V}_{n-L}$. Let $\sigma_L(x_i)$ be the forward path beginning at x_i and ending at x. The *level* $T^o_L(x)$ is the subspace $\cup_i \sigma_L(x_i)$. The point x is the *root* of $T^o_L(x)$ and L is the *length*. The *carrier* $N(T^o_L(x))$ is the smallest subcomplex of $\widetilde{X}$ containing $T^o_L(x)$. Note that $N(T^o_L(x)) = \cup_i N(\sigma_i)$, where σ_i varies over the finitely many maximal forward paths in $T^o_L(x)$. Note that each level has a natural directed graph structure in which edges are midsegments.

Proposition 2.5 (Properties of levels)**.** *Let $T^o_L(x)$ be a level. Then:*

(1) $T^o_L(x)$ is a directed tree in which each vertex has at most one outgoing edge.
(2) If $x \notin \widetilde{X}^0$, then there exists a topological embedding $T^o_L(x) \times [-1, 1] \to \widetilde{X}$ such that $T^o_L(x) \times \{0\}$ maps isomorphically to $T^o_L(x)$.
(3) If $L' \geq L$, then $T^o_L(x) \subseteq T^o_{L'}(x)$.

Proof. $T^o_L(x)$ is connected since it is the union of a collection of paths, each of which terminates at x. Each vertex of $T^o_L(x)$ has at most one outgoing edge. Hence any cycle in $T^o_L(x)$ is directed. The map $q : T^o_L(x) \to \mathbf{R}$ thus shows that there are no cycles in $T^o_L(x)$. This establishes assertion (1).

Let $x_i \in (\tilde{\phi}_L)^{-1}(x)$ and let $\sigma_i \subset T^o_L(x)$ be the forward path joining x_i to x. Then σ_i is disjoint from $\widetilde{X}^0$, since $T^o_L(x)$ is regular. Hence there exists $\epsilon_i > 0$ such that $N(\sigma_i)$ contains an embedded copy of $\sigma_i \times [-\epsilon_i, \epsilon_i]$ with $\sigma_i \times \{0\} = \sigma_i$, which we denote by F_i. Let $\epsilon = \min_i \epsilon_i$. For each i, let $F'_i \subset F_i$ be $\sigma_i \times [-\epsilon, \epsilon] \subseteq \sigma_i \times [-\epsilon_i, \epsilon_i]$, and let $F = \cup_i F'_i$. Since $\sigma_i \cap \sigma_j$ is a forward path for all i, j, the subspace $F \cong T^o_L(x) \times [-\epsilon, \epsilon]$. See the right side of Figure 5.

Assertion (3) follows from the fact that $\tilde{\phi}^{L'} : \widetilde{X} \to \widetilde{X}$ factors as $\widetilde{X} \xrightarrow{\tilde{\phi}^L} \widetilde{X} \xrightarrow{\tilde{\phi}^{L'-L}} \widetilde{X}$. □

For each $L \geq 1$, forward paths and levels are defined in precisely the same way in $\widetilde{X}_L$ and $\widetilde{X}^\bullet_L$. A level of $\widetilde{X}_L$ is subdivided when we formed $\widetilde{X}^\bullet_L$ in Section 1.1. Accordingly, each length-L level in $\widetilde{X}^\bullet_L$ is isomorphic to a star whose edges are subdivided into length-L paths. The map $\tilde{\varrho}_L : \widetilde{X}^\bullet_L \to \widetilde{X}$ sends each length-n level of $\widetilde{X}_L$, each of whose maximal forward paths contains nL midsegments of $\widetilde{X}^\bullet_L$, to a length-nL level in $\widetilde{X}$. Thus $\tilde{\varrho}_L$ maps subdivided stars to rooted trees, as shown in Figure 5.

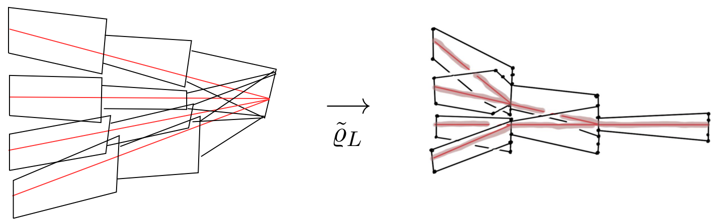

FIGURE 5. The product neighborhood of a regular level in $\widetilde{X}$ is shown at right; the corresponding level in $\widetilde{X}_L$ appears at left. In general, the product neighborhood may contain several subintervals of each vertical edge since ϕ is not in general an immersion on edges.

The image in X_L of a level from $\widetilde{X}_L^\bullet$ is also referred to as a level; it will be clear from the context whether we are working in the base space or the universal cover.

The following observation about forward ladders is required in several places in Section 4 and Section 5.

Lemma 2.6. *Let σ be a forward path. Then for each $R \geq 0$, there exists $\Theta_R \geq 0$, independent of σ and n, such that* $\operatorname{diam}(\mathcal{N}_R(N(\sigma)^1) \cap \mathcal{N}_R(\widetilde{V}_n)) \leq \Theta_R$ *for all $n \in \mathbb{Z}$.*

Proof. This follows from the fact that the coordinate projection $q(\widetilde{V}_n) = n$, while the image of $q|_{N(\sigma)}$ is an interval, each of whose points has uniformly bounded preimage in $N(\sigma)^1$. $\square$

3. Immersed walls, walls, and approximations

In this section, we will describe immersed walls $W \to X$, which are determined by two parameters. The first parameter is a collection $\{d_i\}$ of subintervals of edges in V, called *primary busts*. The second parameter is an integer $L \geq 1$ called the *tunnel length*. The graph W consists of $V - \cup_i d_i - \cup_i (\phi^L)^{-1}(d_i)$ together with a collection of rooted trees called *tunnels*, and is immersed in X_L. We shall show that when L is sufficiently large, $W \to X$ corresponds to a quasiconvex codimension-1 subgroup of G.

3.1. **Primary busts.** Let $\{e_1, \ldots, e_k\}$ be edges of $V \subset X$. For each i, let e_i' be the image of e_i under the isomorphism $V \to E$ given by $(x, 0) \mapsto (x, \frac{1}{2})$. The subspaces e_i', regarded as edges of E, are *primary busted edges*. We will choose closed nontrivial intervals $d_i' \subset \operatorname{Int}(e_i')$, whose distinct endpoints we denote by $p_i^\pm$. The corresponding subinterval of e_i is denoted d_i, and its endpoints $q_i^\pm$ correspond to $p_i^\pm$. Let $E^\flat = E - \cup_{i=1}^k \operatorname{Int}(d_i')$ and let $V^\flat$ denote its preimage in V under the above isomorphism $V \to E$. The subspace $E^\flat$ is the *primary busted space*, and each d_i (or d_i') is a *primary bust*; $\operatorname{Int}(d_i)$ (or $\operatorname{Int}(d_i')$) is an *open primary bust*.

Let C be a component of $V^\flat$ and let $\widetilde{C}$ be a lift of its universal cover to some $\widetilde{V}_n$. Since $C \hookrightarrow V \hookrightarrow X$ is π_1-injective, $\widetilde{C}$ embeds in $\widetilde{V}_n$. Its parallel copy $\widetilde{C}' \subset \widetilde{E}_n$ is a *primary nucleus*, and likewise, each component of $E^\flat$ is a *primary nucleus* in X.

Remark 3.1 (Quasiconvexity of $\widetilde{C}$ under various conditions)**.** In our applications, we will require $\widetilde{C}$ to be quasiconvex in $\widetilde{X}^1$. This is achievable in several ways. Clearly, if $\{e_i\}$ contains enough edges that $E^\flat$ is a forest, then the subspaces $\widetilde{C} \subset \widetilde{X}^1$ are finite trees and therefore quasiconvex.

Quasiconvexity of $\widetilde{C}$ occurs under other circumstances. For example, suppose that $\Phi : F \to F$ is an automorphism and ϕ is a train track map that is *aperiodic* in the sense of [Mit99], i.e.

$\phi^n(e)$ traverses f for all edges e, f and all sufficiently large n. Then, provided $\{e_i\}$ contains at least one edge corresponding to a nontrivial splitting of F, the following theorem of Mitra (see [Mit99, Prop. 3.4]), which is an analog of a result of Scott and Swarup [SS90], ensures that each $\widetilde{C}$ is quasiconvex:

Theorem 3.2. *Let $\Phi : F \to F$ be an aperiodic automorphism of the finite-rank free group F. If $H \leq F$ is a finitely generated, infinite-index subgroup, then H is quasiconvex in $F \rtimes_\Phi \mathbb{Z}$.*

3.2. **Constructing immersed walls.** We now assume that $\widetilde{X}^1$ is δ-hyperbolic. Let $L \geq 1$ be an integer, called the *tunnel length*. For any set $\{d_i\}$ of nontrivial primary busts, the spaces $E^\flat$ and $V^\flat$ embed in X_L by maps factoring through $E \hookrightarrow X_L$ and $V \hookrightarrow X_L$ respectively. For each i, let $\{d_{ij}\}_j$ denote the finite set of components of $(\phi^L)^{-1}(d_i)$. For each i, j, let d'_{ij} be the parallel copy of d_{ij} in E. Each d_{ij} or d'_{ij} is a *secondary bust*. In order to choose busts, we will assume that each edge e of V is *expanding* in the sense that $\phi^k(e) \neq e$ for all $k > 0$. This assumption is justified by Lemma 3.4 below (see also [BH92]).

Definition 3.3. We say $x \in V$ is *periodic* if $\phi^n(x) = x$ for some $n \geq 1$. A point $x \in V$ has *period* m if $\phi^m(x) = x$ and $\phi^k(x) \neq x$ for $0 < k < m$. We then refer to x as being *m-periodic*.

A forward path $\sigma \to \widetilde{X}$ is *periodic* if it is a subpath of a bi-infinite forward path whose stabilizer in G is nontrivial. Note that this holds exactly when each point of $\widetilde{X}^1 \cap \sigma$ projects to a periodic point of V.

Recall that the map $\phi : V \to V$ is *irreducible* if for all edges e, f, there exists $n \geq 0$ such that $\phi^n(e)$ traverses f.

Lemma 3.4. *Let $F \rtimes_\Phi \mathbb{Z}$ be hyperbolic. Then $\Phi : F \to F$ can be represented by a map $\phi : V \to V$ with respect to which each edge of V is expanding and no edge is mapped to a point. Moreover, if Φ has an irreducible train track representative, then $\phi : V \to V$ can be chosen to be an irreducible train track map with respect to which each edge is expanding.*

Proof. We begin with a representative $\phi : V \to V$, which we will adjust by contracting subtrees of V. Let $U \subset V$ be the union of all vertices and all closed edges e such that $|\phi^k(e)|$ is bounded as $k \to \infty$. First, note that $\phi(U) \subseteq U$. Second, each component of U is contractible, since otherwise either ϕ is not π_1-injective or X would contain an immersed torus, contradicting hyperbolicity. We now collapse the ϕ-invariant forest U as in [BH92, Page 7], resulting in a graph $\overline{V}$ and a map $\bar{\phi} : \overline{V} \to \overline{V}$ representing Φ (by reparametrizing, we can assume that the restriction of $\bar{\phi}$ to each edge is a combinatorial path). Note that either U contained no edges (so all edges were expanding and did not map to points), or $\overline{V}$ has strictly fewer edges than V. We repeat the above procedure finitely many times to obtain a graph $\overline{V}$ and a map $\bar{\phi} : \overline{V} \to \overline{V}$ such that edges map to nontrivial paths and all edges are expanding.

The collapse of U preserves the property of being a train track map. Indeed, let $\bar{e}$ be an edge of $\overline{V} = V/U$ that is the image of an edge e of V. Let $n > 0$, and consider the restriction of $\bar{\phi}^n$ to $\bar{e}$. The path $\bar{\phi}^n(\bar{e})$ is obtained from the immersed path $\phi^n(e)$ by collapsing each edge that maps to U. Let $\bar{u}, \bar{v}$ be consecutive edges of $\bar{\phi}^n(\bar{e})$ that fold. Then there is a subpath $u^{-1}fv \subset \phi^n(e)$, where $u \mapsto \bar{u}, v \mapsto \bar{v}$ and f is an immersed path in U. Observe that f is a closed path since u, v have the same initial point. This contradicts the fact that U is a forest.

Finally, the property of irreducibility is preserved by collapsing invariant forests. Indeed, let $\bar{e}, \bar{f}$ be edges of $\overline{V}$ that are images of edges e, f of V. Then by irreducibility of ϕ, there exists $m > 0$ such that $\phi^m(e)$ passes through f, and hence $\bar{\phi}^m(\bar{e})$ passes through $\bar{f}$. $\square$

A point $y \in V$ is *singular* if $\phi^k(y) \in V^0$ for some k.

Lemma 3.5. *Let $L \in \mathbb{N}$. Let $\{e_i\}_{i=1}^k$ be a set of expanding edges in V, let $x_i \in \mathrm{Int}(e_i)$ for each i, and let $\epsilon > 0$. Then there exists a collection $\{d_i\}_{i=1}^k$ of closed subintervals (primary busts), with each $d_i \subset \mathrm{Int}(e_i)$, such that:*

(1) $\cup_i d_i$ is disjoint from the associated secondary busts $\cup_{ij} d_{ij}$.
(2) $\cup_j d_{ij}$ lies in the ϵ-neighborhood of $(\phi^L)^{-1}(x_i)$ for each i.
(3) The endpoints $p_i^\pm$ of d_i are nonsingular.
(4) If x_i is nonsingular and $\phi^L(x_i) \neq x_i$ then we can choose d_i such that x_i is an endpoint of d_i, i.e. $x_i \in \{p_i^\pm\}$.
(5) ϕ^L restricts to an embedding on d_i, for each i.
(6) Suppose that $\phi^L(x_i) \neq \phi^L(x_j)$ for all $i \neq j$. Then $\phi^L(d_i) \cap \phi^L(d_j) = \emptyset$.

Proof. We first establish the finiteness of the set $\mathcal{S}$ consisting of points $s \in e_i$ such that $\phi^L(s) = s$. Each component b of $e_i \cap (\phi^L)^{-1}(e_i)$ is the concatenation of one or more subintervals of e_i, each of which maps homeomorphically to e_i. Since e_i is expanding, Brouwer's fixed point theorem implies that each such subinterval contains a unique point s with $\phi^L(s) = s$. As there are finitely many such b, we conclude that $\mathcal{S}$ is finite.

Let $z_i \in \mathrm{Int}(e_i) - \mathcal{S}$. There exists a nonempty closed interval h_i containing z_i such that $h_i \cap (\phi^L)^{-1}(h_i) = \emptyset$. Indeed, if $h_i \cap (\phi^L)^{-1}(h_i) \neq \emptyset$ for each closed interval containing z_i then there would be a sequence of points converging to z_i whose ϕ^L-images also converge to z_i, and so $z_i \in S$. Property (1) holds whenever $d_i \subset h_i$.

By continuity of ϕ^L, there exists $\delta > 0$ such that $\mathcal{N}_\delta((\phi^L)^{-1}(x_i)) \subset (\phi^L)^{-1}(\mathcal{N}_\epsilon(x_i))$. Property (2) holds by choosing $z_i \in \mathcal{N}_\delta((\phi^L)^{-1}(x_i))$ and letting d_i be a nontrivial component of $h_i \cap Cl\left(\mathcal{N}_\delta((\phi^L)^{-1}(x_i))\right)$. As there are countably many singular points, Property (3) holds since we can assume that neither endpoint of h_i is singular. Property (4) holds by letting $z_i = x_i$, and then choosing h_i above so that it has x_i as an endpoint.

To prove (5), note that e_i has a subdivision where the vertices are points of $(\phi^L)^{-1}(V^0)$. If d_i is properly contained in a single closed edge in this subdivision, then ϕ^L restricts to an embedding on d_i. This can be arranged by choosing d_i sufficiently small (fixing x_i).

We prove (6) by induction on $|\{e_i\}|$. The base case, where $k = 0$, is vacuous. Suppose that $d_1, \ldots, d_{k-1}$ have been chosen to satisfy (1)-(6), with each d_i satisfying $x_k \notin \phi^L(d_i)$ for $1 \leq i < k$. Choose d_k with properties (1)-(5) small enough to avoid the finitely many $\phi^L(d_i)$, $1 \leq i < k$. □

Clearly $d_i \cap d_{i'} = \emptyset$ for $i \neq i'$, since $d_i, d_{i'}$ are contained in distinct open edges. Consequently $d_{ij} \cap d_{i'j'} = \emptyset$ unless $i = i'$ and $j = j'$.

The subspace $\mathbf{N}$ of $E^\flat$ obtained by removing the image under $V \xrightarrow{\sim} E$ of each open secondary bust is the *nucleus*. Observe that $\mathbf{N}$ need not be connected. For each i, j, let $q_{ij}^\pm$ be the endpoints of d_{ij}, which map to $q_i^\pm \in V$, and let $p_{ij}^\pm$ be the corresponding points of $d'_{ij} \subset E$. See Figure 6.

For each i, let $\check{T}_i^{o\pm}$ be the image in X_L of the level $T_L^o(\tilde{q}_i^\pm) \subset \widetilde{X}_L$, where $\tilde{q}_i^\pm$ is an arbitrary lift of $q_i^\pm \in d_i$. Recall that $\check{T}_i^{o\pm}$ is an embedded star of length L rooted at $q_i^\pm$ with leaves at the various $q_{ij}^\pm$ (Proposition 2.5.). Let $S_i^\pm$ be a segment in the 2-cell R_{e_i} of X_L that joins $q_i^\pm$ to $p_i^\mp \in E$. (Note that S_i^+ joins p_i^+ to q_i^- and S_i^- joins p_i^- to q_i^+.) The arcs $S_i^\pm$ are *slopes*. The *level-part* $T_i^{o\pm}$ is the rooted subtree of $\check{T}_i^{o\pm}$ with leaves at $p_{ij}^\pm$. The subspace $T_i^\pm = T_i^{o\pm} \cup S_i^\pm$ obtained by joining the level-part $T_i^\pm$ and the slope along the common point

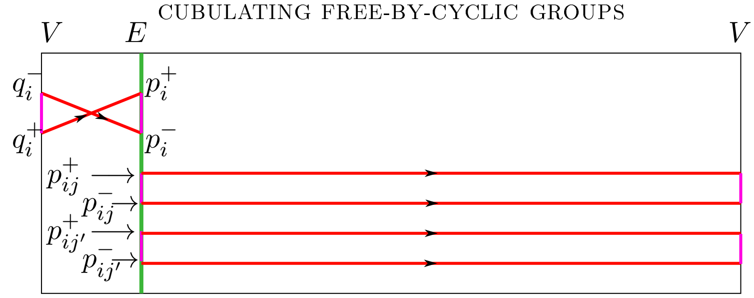

FIGURE 6. Constructing a wall in X_L.

$q_i^\pm$ is a *tunnel*. The space $\widehat{W}^\bullet$ determined by the primary busts $\{d_i\}$ and the tunnel length L is the graph obtained by joining each tunnel $T_i^\pm$ to $\mathbf{N}$ along $\{p_{ij}^\pm\} \cup \{p_i^\mp\}$. The inclusion $\mathbf{N} \hookrightarrow X_L$ and the inclusions $T_i^\pm \hookrightarrow X_L$ induce a (non-combinatorial) immersion $\widehat{W}^\bullet \to X_L$. Note that $T_i^\pm \cap T_j^\pm = \emptyset$ when $i \neq j$ since $d_i \cap d_j = \emptyset$. Note that for each i, the tunnels T_i^+ and T_i^- intersect in the single point $S_i^+ \cap S_i^-$. Composing with the map $X_L \to X$ gives an immersion $\widehat{W}^\bullet \to X$. This extends to a local homeomorphism $\widehat{W}^\bullet \times [-1,1] \to X$ with $\widehat{W}^\bullet$ identified with $\widehat{W}^\bullet \times \{0\}$. Indeed, we described a map $T_i^\pm \times [-1,1] \to X$ earlier, and $\mathbf{N} \times [-1,1] \to X$ is an embedding since $\mathbf{N} \subset E$, and each $S_i^\pm$ lies in a 2-cell. Appropriately chosen neighborhoods $T_i^+ \times [-1,1], T_i^- \times [-1,1]$, and $S_i^+ \times [-1,1], S_i^- \times [-1,1]$, and $\mathbf{N} \times [-1,1]$ can be glued to form $\widehat{W}^\bullet \times [-1,1] \to X_L$. These gluings can be chosen to preserve a "normal vector" at each point of the tunnel, and hence the result is a trivial $[-1,1]$ bundle. The map $\widehat{W}^\bullet \to X$ factors through an immersion $W^\bullet \to X$, where $W^\bullet$ is obtained from $\widehat{W}^\bullet$ by folding the levels according to the map $\varrho_L : X_L \to X$ illustrated in Figure 5. The spaces $\widehat{W}^\bullet$ and $W^\bullet$ are shown in Figure 7.

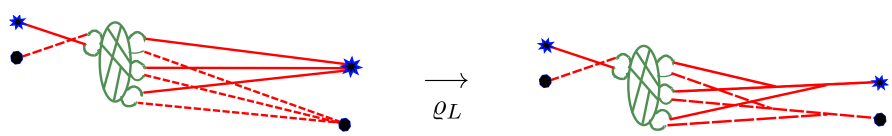

FIGURE 7. The above figure shows $\varrho_L : \widehat{W}^\bullet \to W^\bullet$. In each of the domain and the target, identify the two starred points and the two dotted points; these are the points in V where the slopes, shown at the left of each picture, intersect the levels, shown at right.

Definition 3.6. A component W of $W^\bullet$ is an *immersed wall.*

3.3. **Description of $\overline{W}$.** We now define a wall $\overline{W}$ in $\widetilde{X}$. The map $W \to X$ lifts to a map $\widetilde{W} \to \widetilde{X}$ of universal covers. For each component C of $\mathbf{N}$, the universal cover $\widetilde{C}$ of C lifts to $\widetilde{W}$, and the restriction of $\widetilde{W} \to \widetilde{X}$ to each such $\widetilde{C}$ is an embedding. Moreover, each tunnel lifts to $\widetilde{W}$, and the map $\widetilde{W} \to \widetilde{X}$ restricts to an embedding on each tunnel $T_i \subset \widetilde{W}$. We define $\overline{W}$ to be $\operatorname{im}(\widetilde{W} \to \widetilde{X})$. We conclude that:

Remark 3.7. Let $H_W = \operatorname{Stab}_G(\overline{W})$. When $\overline{W}$ is locally isomorphic to W, the trivial $[-1,1]$-bundle discussed above ensures that there are exactly two components of $\widetilde{X} - \overline{W}$, each of which is H_W-invariant.

Remark 3.8 (Future shape of $\overline{W}$)**.** We now describe the structure of $\overline{W}$ in the situation in which distinct tunnels are disjoint. Note that tunnels $T_i^{\pm}$ and $T_j^{\pm}$ in $\overline{W}$ are disjoint when $i \neq j$, since they map to disjoint tunnels in $\mathrm{im}(W \to X)$. Moreover, we shall show below that, under certain conditions, tunnels $T, T' \subset \overline{W}$, mapping to T_i^+, T_i^- respectively, are disjoint when L is large. In this situation, $\overline{W}$ will be shown to have the structure of a tree of spaces, whose underlying vertices are equipped with a 2-coloring (call the colors "black" and "white). Black vertices correspond to slopes, while white vertices correspond to subspaces that are maximal connected unions of universal covers of nuclei and lifts of level-parts. Note that $\overline{W}$ may still fail to be simply connected – i.e. $\widetilde{W}$ may still fail to embed – since subspaces corresponding to green vertices may not be simply-connected. If $\overline{W}$ contains a nucleus in $\widetilde{E}_n$, then all nuclei lie in $\widetilde{E}_{n+kL}$, $k \in \mathbb{Z}$, and any two nuclei contained in a common vertex space lie in the same space $\widetilde{E}_{n+kL}$. A heuristic picture of $\overline{W}$ is shown in Figure 8, and Figure 9 shows a part of $\overline{W}$ inside $\widetilde{X}$.

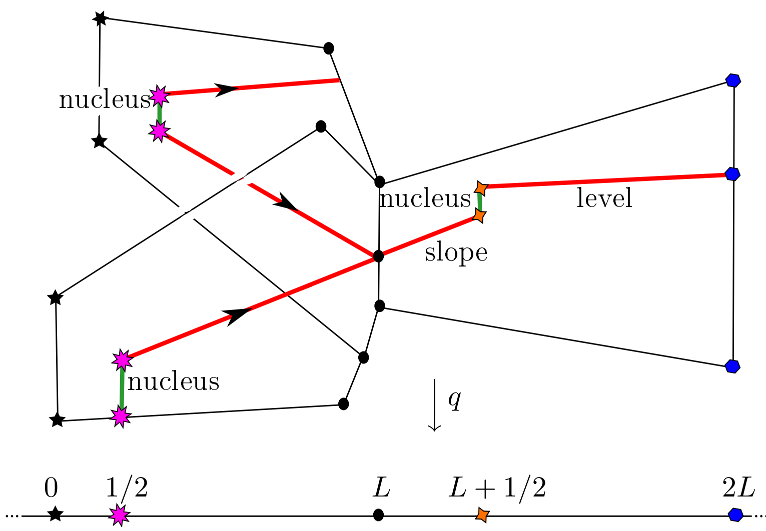

FIGURE 8. A heuristic picture of part of a wall in $\widetilde{X}$ and the effect of the coordinate projection q on the various parts of the wall. Points in the same fiber of q are decorated according to their q–images. The two nuclei at left, and the arrowed levels, belong to the same knockout. This knockout does not contain the slope or the nucleus and level at right.

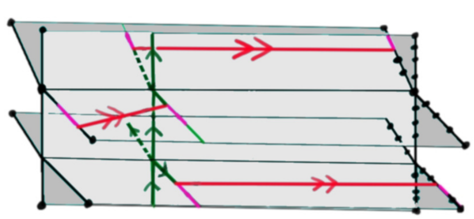

FIGURE 9. Part of a wall $\widetilde{W} \to \widetilde{X}$. The single-arrowed segments belong to a nucleus, while the double-arrowed segments are tunnels.

3.4. **The approximation.** Let $N(\overline{W})$ denote the union of all closed 2-cells of $\widetilde{X}$ that intersect $\overline{W}$. We will show that $N(\overline{W})^1$ is quasiconvex in $\widetilde{X}^1$ under certain conditions, notably sufficiently large tunnel length. However, the quasiconvexity constant will depend on the tunnel length. This is partly because levels are not uniformly quasiconvex and partly because distinct levels emanating from very close secondary busts may contain long forward paths that closely fellow-travel. To achieve uniform quasiconvexity we define the *approximation* of $\overline{W}$, which also has the key feature that it lifts to a geometric wall in $\widetilde{X}^{\bullet}_L$.

Definition 3.9 (Approximation)**.** Let $W \to X$ be an immersed wall with tunnel length L and primary busted edges $\{e_i\}$. Let $\overline{W}$ be the image of a lift $\widetilde{W} \to \widetilde{X}$ of the universal cover of W to $\widetilde{X}$. We define a map $\mathbf{A} : \overline{W} \to \widetilde{X}$ as follows. First, suppose that $\widetilde{C} \subset \overline{W}$ is the universal cover of a component of the nucleus of W. Let $n \in \mathbb{Z}$ be such that $\widetilde{C} \subset \widetilde{E}_n$, and let $\widetilde{C}' \subset \widetilde{V}_n$ be the parallel copy of $\widetilde{C}$. For each $c \in \widetilde{C}$, let c' denote the corresponding point of $\widetilde{C}'$. Then $\mathbf{A} : \widetilde{C} \to \widetilde{X}$ is defined by $\mathbf{A}(c) = \tilde{\phi}^L(c')$. For each level-part T^o of $\overline{W}$, let q be the root of T^o. Then $\mathbf{A}(t) = q$ for each $t \in T_o$. Finally, let $S \subset \overline{W}$ be a slope, beginning at q and ending on a point p in a nucleus component $\widetilde{C}$. Then p is an endpoint of a primary bust $d_i \subset \widetilde{E}_n$. The map $\mathbf{A}$ sends the slope S homeomorphically to the path $d_i'P$, where d_i' is the parallel copy of d_i in $\widetilde{V}_n$ that joins q to p' and P is the forward path joining p' to $\tilde{\phi}^L(p')$. See Figure 10.

The *approximation* $\mathbf{A}(\overline{W})$ of $\overline{W}$ is the image of $\overline{W}$ under the map $\mathbf{A}$. Note that $\mathbf{A}(\overline{W})$ is the union of length-L forward paths together with subspaces of $\widetilde{V}_{nL}$ for each $n \in \mathbb{Z}$. Let $N(\mathbf{A}(\overline{W}))^1$ be the 1-skeleton of the smallest subcomplex of $\widetilde{X}$ containing $\mathbf{A}(\overline{W})$.

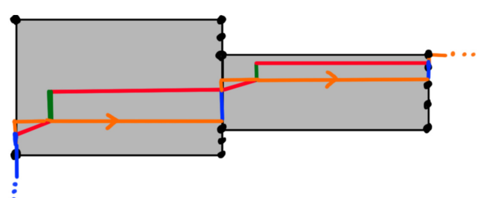

FIGURE 10. Part of a wall and its approximation. The arrowed paths are the approximations of the slopes intersecting them.

Remark 3.10. Observe that the use of slopes ensures that $\mathbf{A}(\overline{W})$ passes through each primary busted edge intersecting $\overline{W}$, which is crucial in the proof Proposition 4.6, which says that $\overline{W}$ is actually a wall.

Remark 3.11. If $\widetilde{C}_1, \widetilde{C}_2$ are nuclei of $\overline{W}$ intersecting a level-part of a tunnel of $\overline{W}$, then $\mathbf{A}(\widetilde{C}_1) \cap \mathbf{A}(\widetilde{C}_2) \neq \emptyset$. For each $n \in \mathbb{Z}$, each component of $\mathbf{A}(\overline{W}) \cap \widetilde{V}_n$ is formed as follows. A *knockout* $\widetilde{K}$ is a maximal connected subspace of $\overline{W}$ that does not contain an interior point of a slope. The knockout $\widetilde{K}$ is *at position* n if it is the union of nuclei in $\overline{W} \cap \widetilde{E}_{n-L}$ together with level-parts traveling from $\widetilde{E}_{n-L}$ to $\widetilde{V}_n$. The knockout in Figure 8 is at position L.

To each position-n knockout $\widetilde{K}$, we associate a component of $\mathbf{A}(\overline{W}) \cap \widetilde{V}_n$, namely the one obtained from the connected subspace $\mathbf{A}(\widetilde{K}) \subset \widetilde{V}_n$ by adding all (closed) primary bust intervals that intersect $\mathbf{A}(\widetilde{K})$. See Figure 11.

Remark 3.12. If S_1, S_2 are distinct slopes, rooted at primary busts d_1, d_2, then $\mathbf{A}(S_1) \cap \mathbf{A}(S_2) = \emptyset$ by Lemma 3.5.(5)-(6).

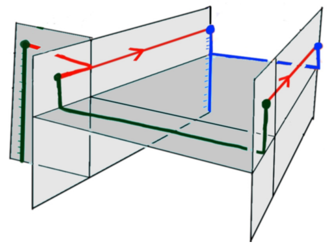

FIGURE 11. Two nuclei intersecting a common level-part have intersecting approximations whose union avoids primary busts.

3.5. **Hypotheses on $\mathbf{A}(\overline{W})$ enabling quasiconvexity.** We recall that we are assuming, for the rest of the paper, that there exists $\delta \geq 0$ such that $\widetilde{X}^1$ is δ-hyperbolic. The following statements are instrumental in proving that, provided L is sufficiently large, $N(\mathbf{A}(\overline{W}))^1 \hookrightarrow \widetilde{X}$ is quasi-isometrically embedded, and the quasi-isometry constants are independent of L.

Definition 3.13 (Bust-quasiconvex). W is *bust-quasiconvex* if there exist constants μ_1', μ_2' such that each component in $\widetilde{X}$ of the preimage of $V - \cup_i \mathrm{Int}(d_i)$ is (μ_1', μ_2')-quasi-isometrically embedded. For example, as noted in Remark 3.1, W is bust-quasiconvex if $V^\flat$ is a forest or if Φ is an aperiodic isomorphism and there is at least one bust.

Lemma 3.14 (Quasiconvexity of approximations of nuclei). *Approximations have the following properties when W is bust-quasiconvex:*

1. *For each nucleus $\widetilde{C}$ and each open primary bust $\tilde{d}$, we have $\mathbf{A}(\widetilde{C}) \cap \tilde{d} = \emptyset$.*
2. *Let $\widetilde{K}$ be a knockout. Then $\mathbf{A}(\widetilde{K}) \cap \tilde{d} = \emptyset$ for each open primary bust $\tilde{d}$.*
3. *Hence there exist $\mu_1 \geq 1, \mu_2 \geq 0$ such that, for each $n \in \mathbb{Z}$ and each component $\mathbf{K}$ of $\widetilde{V}_n \cap \mathbf{A}(\overline{W})$, the inclusion $\mathbf{K} \to \widetilde{X}^1$ is a (μ_1, μ_2)-quasi-isometric embedding. Moreover, μ_1 and μ_2 are independent of $\{d_i\}$ and L.*

Proof.

1. $\widetilde{C}$ has empty intersection with the set of open secondary busts in $\widetilde{V}_n$, and hence maps to the complement of the set of open primary busts in $\widetilde{V}_{n+L}$.
2. This follows immediately from (1) because level-parts map to points.
3. Since W is bust-quasiconvex, statement (2) implies that $\mathbf{K}$ is a subtree of a uniform neighborhood in $\widetilde{V}_n$ of some $\mathbf{A}(\widetilde{K})$, and the first claim follows. Since there are finitely many possible sets of primary busted edges, the constants μ_1, μ_2 can be chosen independently not only of L and $\{d_i\}$, but also of $\{e_i\}$. Indeed, each set $\{e_i\}$ of edges yielding a bust-quasiconvex immersed wall gives rise to a pair of quasi-isometry constants, and we take μ_1, μ_2 to be the maximal such constants. □

Lemma 3.14 provides uniform quasiconvexity of nucleus approximations, and Proposition 2.3 provides uniform quasiconvexity of forward ladders. Lemma 2.6 provides a bound on the diameters of coarse intersections of nucleus approximations and carriers of approximations of slopes. To prove quasiconvexity of $N(\mathbf{A}(\overline{W}))^1$ requires the following additional property.

Definition 3.15 (Ladder overlap property). A family of immersed walls $\{W_i \to X\}$ has the *ladder overlap property* if there exists $B \geq 0$ such that for all i and all distinct tunnels $T_1, T_2 \subset \overline{W}_i$ intersecting a common nucleus,

$$\mathrm{diam}\left(\mathcal{N}_{3\delta+2\lambda}(N(\mathbf{A}(T_1))) \cap \mathcal{N}_{3\delta+2\lambda}(N(\mathbf{A}(T_2)))\right) \leq B,$$

where λ is the constant from Proposition 2.3.

Remark 3.16. The purpose of the ladder overlap property is to guarantee that, when L is large and W is bust-quasiconvex, paths of the form $\beta\alpha\beta'$ are uniform quasigeodesics, where β, β' are geodesics of carriers of slope-approximations and α is a vertical geodesic in $\mathbf{A}(\overline{W})$.

If the interiors of β, β' have disjoint images in $\mathbf{R}$, Lemma 2.6 ensures that $\beta\alpha\beta'$ is a uniform quasigeodesic. The interesting situations are those in which β, β' are both incoming or both outgoing with respect to the vertical part of $\mathbf{A}(\overline{W})$ containing α. A thin quadrilateral argument shows that in either case, the ladder overlap property ensures that β, β' have uniformly bounded coarse intersection, from which one concludes that $\beta\alpha\beta'$ is a uniform quasigeodesic (see Lemma 4.3 below).

4. Quasiconvex codimension-1 subgroups from immersed walls

In this section, we determine conditions ensuring that $N(\mathbf{A}(\overline{W}))^1$ is quasiconvex and $\overline{W}$ is a wall in $\widetilde{X}$, continuing to assume that $\widetilde{X}^1$ is δ–hyperbolic.

4.1. **Uniform quasiconvexity.** A collection $\{W \to X\}$ of immersed walls is *uniformly bust-quasiconvex* if there exist constants μ_1, μ_2 such that $\mathbf{A}(\widetilde{K}) \hookrightarrow \widetilde{X}^1$ is a (μ_1, μ_2)-quasi-isometric embedding for each W and each knockout $\widetilde{K}$ of $\overline{W}$. The first goal of this section is to prove:

Proposition 4.1. *Let $\mathbb{W} = \{W \to X\}$ be a uniformly bust-quasiconvex set of immersed walls with the ladder-overlap property. Then there exists L_0, κ_1, κ_2 such that for all $W \in \mathbb{W}$ with tunnel length at least L_0, the inclusion $N(\mathbf{A}(\overline{W}))^1 \hookrightarrow \widetilde{X}^1$ is a (κ_1, κ_2)-quasi-isometric embedding.*

The constants are $\kappa_1 = 4\lambda_1\mu_1$ and $\kappa_2 = \frac{\mu_2}{2} + 2L_0(1 + \frac{1}{4\lambda_1\mu_1})$. Here μ_1, μ_2 are the quasi-isometry constants from uniform bust-quasiconvexity, and λ_1 is the multiplicative quasi-isometry constant for 1-skeleta of forward ladders. We emphasize that these are independent of L and of the collection of primary busts. We postpone the proof of Proposition 4.1 until after the following necessary lemmas.

Lemma 4.2. *Let Z be δ-hyperbolic, and let $P = \alpha_0\beta_1\alpha_1 \cdots \beta_k\alpha_k$ be a path in Z with all α_i and β_i geodesic. Suppose there exists $B \geq 0$ such that for all i, each intersection below has diameter $\leq B$:*

$$\mathcal{N}_{3\delta}(\beta_i) \cap \beta_{i+1}, \quad \mathcal{N}_{3\delta}(\beta_i) \cap \alpha_i, \quad \mathcal{N}_{3\delta}(\beta_i) \cap \alpha_{i-1}.$$

Then if $|\beta_i| \geq 12(B + \delta)$ for each i, then $\|P\| \geq \frac{1}{2}\left(\sum_{i=0}^{k} |\alpha_i| + \sum_{i=1}^{k} |\beta_i|\right)$.

Proof. This is a standard argument. We refer, for instance, to [HW12, Thm 2.3]. □

We now promote Lemma 4.2 to a statement about piecewise-quasigeodesics.

Lemma 4.3. *Let Z be δ-hyperbolic and let $P = \alpha_0\beta_1\alpha_1 \cdots \beta_k\alpha_k$ be a path in Z such that each β_i is a (λ_1, λ_2)-quasigeodesic and each α_i is a (μ_1, μ_2)-quasigeodesic. Suppose that for each $R \geq 0$ there exists $B_R \geq 0$ such that for all i, each intersection below has diameter $\leq B_R$:*

$$\mathcal{N}_{3\delta+R}(\beta_i) \cap \beta_{i+1}, \quad \mathcal{N}_{3\delta+R}(\beta_i) \cap \alpha_i, \quad \mathcal{N}_{3\delta+R}(\beta_i) \cap \alpha_{i-1}.$$

Then there exists L_0 such that, if $|\beta_i| \geq L_0$ for each i, then $\|P\| \geq \frac{1}{4\lambda_1\mu_1}|P| - \frac{\mu_2}{2}$.

Proof. For each i, let $\bar{\alpha}_i$ [respectively, $\bar{\beta}_i$] be a geodesic with the same endpoints as α_i [respectively, β_i], and let $\overline{P} = \bar{\alpha}_0\bar{\beta}_1\bar{\alpha}_1 \cdots \bar{\beta}_k\bar{\alpha}_k$ be a piecewise-geodesic with the same endpoints as P. Since Z is δ-hyperbolic, there exists $\mu = \mu(\mu_1, \mu_2, \delta)$ such that α_i and $\bar{\alpha}_i$ lie at Hausdorff distance at most μ, and there exists $\lambda = \lambda(\lambda_1, \lambda_2, \delta)$ such that $\bar{\beta}_i$ and β_i lie at Hausdorff distance at most λ.

Note that if $R_1 \leq R_2$, then we may assume $B_{R_1} \leq B_{R_2}$. By hypothesis, $\mathcal{N}_{3\delta+2\lambda}(\beta_i) \cap \beta_{i+1}$ has diameter $\leq B_{2\lambda}$. Moreover, if $\bar{\beta}' \subset \bar{\beta}_i$ is a subpath that 3δ-fellowtravels with a subpath $\bar{\alpha}'$ of $\bar{\alpha}_i$ or $\bar{\alpha}_{i-1}$, then β' fellowtravels at distance $3\delta + \mu + \lambda$ with a subpath α'' of α_i or α_{i-1}, whence $|\beta'| \leq B_{\mu+\lambda}$ by hypothesis. Letting $L_0 \geq 12(\delta + B_{\mu+2\lambda})$ and applying Lemma 4.2 shows that $\overline{P}$ is a $(\frac{1}{2}, 0)$-quasigeodesic, and we have:

$$\|P\| = \|\overline{P}\| \geq \frac{1}{2}\left|\overline{P}\right|. \tag{1}$$

Since $\mu_1, \lambda_1 \geq 1$, we can bound $|\overline{P}|$ as follows:

$$\begin{aligned}
\left|\overline{P}\right| = \sum_{i=1}^{k} |\bar{\beta}_i| + \sum_{i=0}^{k} |\bar{\alpha}_i| \quad &\geq \quad \sum_{i=1}^{k} (\lambda_1^{-1}|\beta_i| - \lambda_2) + \sum_{i=0}^{k} (\mu_1^{-1}|\alpha_i| - \mu_2) \\
&= \quad \left[\lambda_1^{-1} \sum_{i=1}^{k} (|\beta_i| - \lambda_1(\lambda_2 + \mu_2)) + \mu_1^{-1} \sum_{i=0}^{k} |\alpha_i|\right] - \mu_2 \\
&\geq \quad (\lambda_1\mu_1)^{-1} \left[\sum_{i=1}^{k} (|\beta_i| - \lambda_1(\lambda_2 + \mu_2)) + \sum_{i=0}^{k} |\alpha_i|\right] - \mu_2.
\end{aligned}$$

If $L_0 \geq 2\lambda_1(\lambda_2 + \mu_2) + 1$, then, provided that $|\beta_i| \geq L \geq L_0$, we have:

$$\left|\overline{P}\right| \geq \frac{1}{2\lambda_1\mu_1}\left[\sum_{i=1}^{k} |\beta_i| + \sum_{i=0}^{k} |\alpha_i|\right] - \mu_2.$$

Combining this with Equation (1) yields $\|P\| \geq \frac{1}{4\lambda_1\mu_1}|P| - \frac{\mu_2}{2}$. $\square$

Proof of Proposition 4.1. For a path P in $\widetilde{X}^1$, as usual $\|P\|$ denotes the distance in $\widetilde{X}^1$ between the endpoints of P. If P is a geodesic of $N(\mathbf{A}(\overline{W}))^1$, then its edge-length $|P|$ equals the distance in $N(\mathbf{A}(\overline{W}))^1$ between the endpoints of P. We will show that when $L \geq L_0$, there are constants κ_1, κ_2 such that $\|P\| \geq \kappa_1^{-1}|P| - \kappa_2$.

Alternating geodesics: Let P' be a geodesic in the graph $N(\mathbf{A}(\overline{W}))^1$. Suppose P' *alternates*, in the sense that $P' = \alpha_0\beta_1\alpha_1 \cdots \beta_k\alpha_k$, where each α_i is a vertical geodesic path, and each β_i is a geodesic of the 1-skeleton of a length-L forward ladder (and thus a (λ_1, λ_2)-quasigeodesic). We allow the possibility that α_0 or α_k has length 0.

Each α_i is a (μ_1, μ_2)-quasigeodesic by our hypothesis that knockout-approximations are quasi-isometrically embedded. Since W has the ladder overlap property, $\operatorname{diam}(\mathcal{N}_{3\delta+2\lambda}(\beta_i) \cap \mathcal{N}_{3\delta+2\lambda}(\beta_{i+1})) \leq B$. Let $B_0 = \max(B, \Theta_{3\delta+2\lambda})$, where $\Theta_{3\delta+2\lambda}$ is as in Lemma 2.6. Applying Lemma 4.3 yields a constant L_0 such that, if $L \geq L_0$, then $\|P'\| \geq \frac{1}{4\lambda_1\mu_1}|P'| - \frac{\mu_2}{2}$.

$\mathbf{A}(\overline{W})$ quasi-isometrically embeds: Let P be a geodesic of $N(\mathbf{A}(\overline{W}))^1$. By construction $P = \beta_0'P'\beta_{k+1}'$ where P' is alternating and β_0', β_{k+1}' are (possibly trivial) paths in forward ladders. If $|\beta_0'|, |\beta_{k+1}'| \geq L_0$, then $\|P\| \geq \frac{1}{4\lambda_1\mu_1}|P| - \frac{\mu_2}{2}$ by Lemma 4.3. If $|\beta_0'|, |\beta_{k+1}'| \leq L_0$, then since P' is alternating,

$$\|P\| \geq \frac{1}{4\lambda_1\mu_1}|P'| - \frac{\mu_2}{2} - 2L_0 \geq \frac{1}{4\lambda_1\mu_1}|P| - \frac{\mu_2}{2} - 2L_0(1 + \frac{1}{4\lambda_1\mu_1}).$$

In the remaining case, without loss of generality, $P = \beta_0' P''$, where $|\beta_0'| \leq L_0$ and P'' satisfies $\|P''\| \geq \frac{1}{4\lambda_1\mu_1}|P''| - \frac{\mu_2}{2}$ by Lemma 4.3. The proof is thus complete with $\kappa_1 = 4\lambda_1\mu_1$ and $\kappa_2 = \frac{\mu_2}{2} + 2L_0(1 + \frac{1}{4\lambda_1\mu_1})$. □

4.2. $\overline{W}$ **is a wall when tunnels are long.** A subspace $Y \subset \widetilde{X}$ is a *wall* if $\widetilde{X} - Y$ has exactly two components, each of which is stabilized by $\mathrm{Stab}(Y)$. Note that this definition is stricter than usual. For more about wallspaces and the various definitions, background, and references, see [HW]. Our goal is now to show that if $W \to X$ is an immersed wall with sufficiently long tunnels, then $\overline{W}$ is a wall. We need the following useful consequence of quasiconvexity.

Proposition 4.4. *Let $\mathbb{W}$ satisfy the hypotheses of Proposition 4.1. There exists $L_1 \geq L_0$ such that $\mathbf{A}(\overline{W})$ is a tree for each $W \in \mathbb{W}$ with tunnel length $L \geq L_1$.*

Proof. Let Q be an immersed path in $\mathbf{A}(\overline{W})$, and let Q' be a geodesic of $N(\mathbf{A}(\overline{W}))^1$ with the same endpoints as Q. Proposition 4.1 implies that $\|Q'\| \geq \kappa_1^{-1}|Q| - \kappa_2$. Hence if $|Q'| \geq L_1 = \max(L_0, \kappa_1\kappa_2 + \kappa_1)$, then Q is not closed. Any immersed path Q in $\mathbf{A}(\overline{W})$ either lies in a single vertex space and is thus not closed, or contains a slope approximation and thus Q' has length at least $L \geq L_1$. □

Remark 4.5 (Tree of spaces structure on $\overline{W}$)**.** Proposition 4.4 justifies our claim in Remark 3.8 that $\overline{W}$ is a tree of spaces when L is sufficiently large, assuming that W is bust-quasiconvex and has the ladder overlap property. Indeed, any cycle in $\overline{W}$ that is not contained in a knockout will map to a cycle in $\mathbf{A}(\overline{W})$, contradicting Proposition 4.4.

Proposition 4.6. *Let $W \to X$ be an immersed wall in a collection $\mathbb{W}$ satisfying the hypotheses of Proposition 4.1. The image $\overline{W} \subset \widetilde{X}$ of $\widetilde{W} \to \widetilde{X}$ is a wall provided that W has tunnel length $L \geq L_1$, where L_1 is the constant from Proposition 4.4.*

Proof. Since $\mathsf{H}^1(\widetilde{X}) = 0$, it suffices to show that $\overline{W}$ has an open neighborhood homeomorphic to $\overline{W} \times [-1,1]$ with $\overline{W}$ identified with $\overline{W} \times \{0\}$. The local homeomorphism $W \times [-1,1] \to X$ lifts to a map $\widetilde{W} \times [-1,1] \to \widetilde{X}$. The image of $\widetilde{W} \times [-1,1] \to \widetilde{X}$ would provide the desired neighborhood $\overline{W} \times [-1,1]$ provided that this map is a covering map onto its image. By choosing the image of $W \times [-1,1]$ to be sufficiently narrow, the only place where this could fail is where distinct slopes of $\overline{W}$ intersect. To exclude this possibility, we will show that distinct tunnels T_0, T_1 of $\overline{W}$ are disjoint.

Suppose that $T_0 \neq T_1$ and $T_0 \cap T_1 \neq \emptyset$. Let e be the primary busted edge dual to T_0 and T_1. Since $T_0, T_1 \subset \overline{W}$, there exists a path $P \to \overline{W}$ that starts on T_0, ends on T_1, and which is disjoint from the interiors of T_0 and T_1. Indeed, let $\widetilde{P} \to \widetilde{W}$ be a path joining lifts of $\widetilde{T}_0, \widetilde{T}_1$ and let P be the image of $\widetilde{P}$ in $\overline{W}$. Moreover, we assume that $\widetilde{P}$ is minimal in the sense that it is disjoint from intervening lifts of T_0, T_1. The minimality of $\widetilde{P}$ ensures that P has the desired property.

There are three cases. The first is where P starts and ends on the levels of T_0, T_1. The second is where P starts and ends on the slopes of T_0, T_1. The third case is where P starts on the level of (say) T_0 and ends on the slope of T_1. See Figure 12.

Observe that $e \subset \mathbf{A}(\overline{W})$, as shown in Figure 12, by the definition of $\mathbf{A}(\overline{W})$. Indeed, $e = \alpha_0 d \alpha_1$, where d is the primary bust and α_0, α_1 lie in the approximations of the nuclei attached to the levels of T_0, T_1. (The approximations of T_0, T_1 overlap along d.) The primary bust d is included in both of the associated slope-approximations (see Definition 3.9).

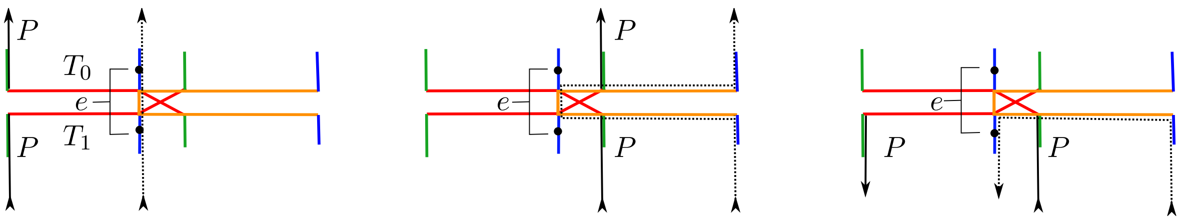

FIGURE 12. The three cases in the proof of Proposition 4.6. The solid arrowed lines are parts of the path P. The dotted lines are parts of the closed paths in $\mathbf{A}(\overline{W})$ constructed by joining endpoints of $\mathbf{A}(P)$.

In the first case, at left in Figure 12, the approximation $\mathbf{A}(P)$ of the image of P is a connected subspace of $\mathbf{A}(\overline{W})$ that contains the endpoints of e but does not contain the entire edge e. Hence $\mathbf{A}(P) \cup e$ contains a cycle. Since $\mathbf{A}(P) \subset \mathbf{A}(\overline{W})$ and $e \subset \mathbf{A}(\overline{W})$,

Similarly, in the second case, shown in the center of Figure 12, $\mathbf{A}(P)$ is disjoint from e, so that $\mathbf{A}(P) \cup e \cup \mathbf{A}(T_0) \cup \mathbf{A}(T_1)$ is a subspace of $\mathbf{A}(\overline{W})$ that contains a cycle. In the third case, shown at right, the contradictory subspace is $\mathbf{A}(P) \cup \mathbf{A}(T_1) \cup e$. $\square$

We note the following corollary:

Corollary 4.7. *Let $\mathbb{W}$ be a set of bust-quasiconvex immersed walls such that $\mathbb{W}$ has the ladder overlap property. Then there exists L_1 such that for all $W \in \mathbb{W}$ with tunnel length $L \geq L_1$, the stabilizer $H_W \leq G$ of $\overline{W}$ is a quasiconvex, codimension-1 free subgroup.*

5. Cutting geodesics

In this section, we recall the criterion for cocompact cubulation of hyperbolic groups given in [BW13] and describe how a sufficiently rich collection of quasiconvex walls in $\widetilde{X}$ ensures that this criterion is satisfied.

5.1. **Separating points on $\partial\widetilde{X}$.** Let $\partial\widetilde{X}$ denote the Gromov boundary of $\widetilde{X}^1$. Let $W \to X$ be an immersed wall with the property that $N(\mathbf{A}(\overline{W}))^1$ is quasiconvex in $\widetilde{X}^1$ and $\overline{W}$ is a wall. Let $\overleftarrow{W}$ and $\overrightarrow{W}$ be the components of $\widetilde{X} - \overline{W}$, and let $N(\overleftarrow{W}), N(\overrightarrow{W})$ be the smallest subcomplexes containing $\overleftarrow{W}, \overrightarrow{W}$ respectively. Then $N(\overleftarrow{W})^1 \cap N(\overrightarrow{W})^1 = N(\overline{W})^1$, which is coarsely equal to $N(\mathbf{A}(\overline{W}))^1$. Let $\partial\overline{W}$ denote $\partial N(\overline{W})^1 = \partial N(\mathbf{A}(\overline{W}))^1$, which is a closed subset of $\partial\widetilde{X}$ since $N(\mathbf{A}(\overline{W}))^1$ is quasiconvex in $\widetilde{X}^1$. Let $\partial\overleftarrow{W} = \partial N(\overleftarrow{W})^1 - \partial\overline{W}$ and let $\partial\overrightarrow{W} = \partial N(\overrightarrow{W})^1 - \partial\overline{W}$, so that $\partial\overleftarrow{W}$ and $\partial\overrightarrow{W}$ are disjoint open subsets of $\partial\widetilde{X}$. Note that $\partial\overleftarrow{W}$ and $\partial\overrightarrow{W}$ are H_W-invariant, since $N(\overleftarrow{W})$ and $N(\overrightarrow{W})$ are H_W-invariant by Remark 3.7.

Let $p, q \in \partial\widetilde{X}$ be the endpoints of a bi-infinite geodesic $\gamma : \mathbf{R} \to \widetilde{X}^1$. Then γ is *cut* by $\overline{W}$ if $p \in \partial\overleftarrow{W}$ and $q \in \partial\overrightarrow{W}$ or vice versa.

The following holds by [BW13, Thm 1.4]:

Proposition 5.1. *Suppose that for every geodesic $\gamma : \mathbf{R} \to \widetilde{X}^1$, there exists a wall $\overline{W}$ of the type described in Section 3, such that $N(\overline{W})$ is quasiconvex and such that $\overline{W}$ cuts γ. Then* *there exists a G-finite collection $\{\overline{W}\}$ of walls in $\widetilde{X}$ such that G acts freely and cocompactly on the dual CAT(0) cube complex.*

The utility of Proposition 5.1 is that we can build an enormous (in particular, G–infinite) collection of quasiconvex walls in $\widetilde{X}$, and the proposition will provide a suitable G–finite collection of walls.

5.2. **A method for cutting the two types of geodesics.**

Definition 5.2 (Ladderlike, deviating)**.** Let $M \geq 0$ and let $\gamma \subset \widetilde{X}^1$ be an embedded infinite or bi-infinite path whose image is ξ-quasiconvex, for some $\xi \geq 0$. Then γ is *M-ladderlike* if there exists a forward ladder $N(\sigma)$, where σ is a forward path of length M, such that a geodesic of $N(\sigma)$ joining the endpoints of σ fellow-travels with a subpath of γ at distance $2\delta + \lambda + \xi$. Here, $\widetilde{X}^1$ is δ-hyperbolic and λ is the constant from Proposition 2.3. Otherwise, γ is *M-deviating.*

Note that if γ is M-deviating, for each $R \geq 0$ there exists M_R depending only on ξ, M, R such that for all forward paths σ, we have $\mathrm{diam}(\gamma \cap \mathcal{N}_R(N(\sigma)^1)) \leq M_R$. Moreover, if γ is $2M$-deviating, the same bound holds with σ replaced by any level, since any geodesic in a level decomposes as the concatenation of two (possibly trivial) forward paths.

Definition 5.3 (Many effective walls)**.** A set $\mathbb{W}$ of immersed walls in X is *spreading* if:

- For arbitrarily large L, there exists $W \in \mathbb{W}$ with tunnel length L.
- $\mathbb{W}$ has the ladder overlap property of Definition 3.15.

$\widetilde{X}$ has *many effective walls* if Conditions (1) and (2) below hold.

(1) For each regular $y \in V$, there exists a spreading set $\mathbb{W}_y$ such that for each $\epsilon > 0$ and each $m \in \mathbb{N}$, there exists $L > m$ and $W \in \mathbb{W}_y$ with tunnel length L, a primary bust in each edge of V, and a primary bust in the ϵ-neighborhood of y.

(2) For each $a \in \widetilde{V}_0$ whose image in V is periodic and whose corresponding point in $\widetilde{E}_0$ is denoted by a', there exists $k = k(a) \geq 0$ and a spreading set $\mathbb{W}_a$ such that all of the following hold:
   - $\mathbb{W}_a$ is uniformly bust-quasiconvex;
   - for each $W \in \mathbb{W}_a$, each edge of V contains a primary bust of W;
   - for each primary bust d' of $\overline{W}$ in $\widetilde{E}_0$ (corresponding to an interval $d \subset \widetilde{V}_0$) that is joined to $a' \in \widetilde{E}_0$ by a path in a knockout of $\overline{W}$, we have $\mathsf{d}_{\widetilde{X}^1}(\tilde{\phi}^n(a), \tilde{\phi}^n(d)) \geq 3\delta + 2\lambda$ for all $n \geq k$. See Figure 13.

   (Observe that $\mathbb{W}_a$ need not be a maximal set of immersed walls satisfying the above conditions, and indeed our applications we choose a very specific $\mathbb{W}_a$.)

Remark 5.4. Firstly, assuming that the third part of Definition 5.3.(2) holds, the constant k can be chosen independently of the point a. For each $a \in \widetilde{V}_0$ whose image $\bar{a}$ in V is periodic, let $k'(a)$ be chosen so that for each bust d and each $n \geq k'(a)$, we have $\mathsf{d}_{\widetilde{X}^1}(\tilde{\phi}^n(a), \tilde{\phi}^n(d)) \geq 3\delta+2\lambda+1$. This is possible since the existence of $k(a)$ implies that the forward rays emanating from a and any point of d diverge. By translating, we note that if a_1 is another lift of $\bar{a}$ to $\widetilde{V}_0$, then we can take $k'(a_1) = k'(a)$.

Fix $\epsilon \in (0,1)$ and consider the neighborhood of a given by $U_a = (\tilde{\phi}^{k'(a)})^{-1}(\mathcal{N}_\epsilon(\tilde{\phi}^{k'(a)}(a)))$. Then for each $b \in U_a$, we have $\mathsf{d}_{\widetilde{X}^1}(\tilde{\phi}^n(b), \tilde{\phi}^n(d)) \geq 3\delta+2\lambda+1-\epsilon > 3\delta+2\lambda$ for each $n \geq k'(a)$. Hence we may choose $k(b) \leq k'(a)$.

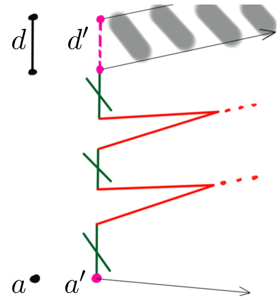

FIGURE 13. Definition 5.3.(2).

Let $V^{\circlearrowleft} \subseteq V$ be the set of periodic points, and let $Cl\,(V^{\circlearrowleft})$ denote its closure, which is compact since V is compact. Let $\overline{U}_{\bar{a}}$ be the image of U_a in V, which is open since U_a is open. Note that $\{\overline{U}_{\bar{a}} : \bar{a} \in V^{\circlearrowleft}\}$ is an open covering of $Cl\,(V^{\circlearrowleft})$, since $\bar{a} \in \overline{U}_{\bar{a}}$ for each $\bar{a} \in V^{\circlearrowleft}$. Hence the claim follows since the open covering has a finite subcovering consisting of finitely many sets $\overline{U}_{\bar{a}}$, and we can take k to be the maximum of the associated $k(a')$.

Secondly, since the collection $\mathbb{W}_a$ has the ladder overlap property, we claim there is likewise a constant B_a such that for all $W \in \mathbb{W}_a$, any two tunnels T, T' of $\overline{W}$ that are joined by a path in $\overline{W}$ not traversing a slope have the property that the $(3\delta + 2\lambda)$-neighborhoods of $\mathbf{A}(T), \mathbf{A}(T')$ intersect in a set of diameter at most B_a. Let V_a be the set of images in V of points $b \in \widetilde{V}$ such that for all $W \in \mathbb{W}_a$, the point b lies in a nucleus of some $\overline{W}$ and for all primary busts d of $\overline{W}$, we have $\mathsf{d}_{\widetilde{X}^1}(\tilde{\phi}^n(a), \tilde{\phi}^n(d)) \geq 3\delta + 2\lambda$ for all $n \geq k$. The previous argument showed that, with k chosen appropriately, the set V_a is open. It follows that if $\widetilde{X}$ has many effective walls, the ladder overlap constant B_a can be chosen independently of a. This is used in the proof of Proposition 5.19.

Definition 5.5 (Separating level)**.** $\widetilde{X}$ is (M,K)-*separated* if for each M-deviating geodesic γ there exists $y \in \widetilde{X}$ such that the following holds for all sufficiently large n: the set $\gamma \cap T^o_n(\tilde{\phi}^n(y))$ has odd cardinality, and the distance in $T^o_n(\tilde{\phi}^n(y))$ from $\gamma \cap T^o_n(\tilde{\phi}^n(y))$ to the root or to any leaf of $T^0_n(\tilde{\phi}^n(y))$ exceeds $M+K$. See Figure 14. We say $\widetilde{X}$ is *level-separated* if it is (M,K)-separated for all $M > 0, K \geq 0$.

Remark 5.6. If the level $T^o_n(\tilde{\phi}^n(y))$ separates γ in the above sense, then we can choose y so that the image $\bar{y} \in V$ of y is not periodic. Indeed, if y', y are sufficiently close, then $T^o_n(\tilde{\phi}^n(y))$ and $T^o_n(\tilde{\phi}^n(y'))$ both separate γ. There are points y' arbitrarily close to y whose images in V are not periodic since there are only countably many periodic points.

Definition 5.7 (Bounded level-intersection)**.** $\widetilde{X}$ has *bounded level-intersection* if for each $z \in \widetilde{X}^1$ and each vertical edge $e \subset \widetilde{X}^1$, there exists $K = K(z,e)$ such that for every level T with a leaf at z, we have $|T \cap e| \leq K$.

Remark 5.8. In the case of greatest interest, where X is the mapping torus of a train track map, each level intersects each vertical edge in at most a single point (Lemma 6.3), and hence $\widetilde{X}$ has bounded level-intersection. This holds in particular for the complexes $\widetilde{X}$ considered in Theorem 6.16. More generally, this holds whenever there is a continuous map from $\widetilde{X}$ to an $\mathbb{R}$-tree that is constant on levels and sends edges to concatenations of finitely many arcs.

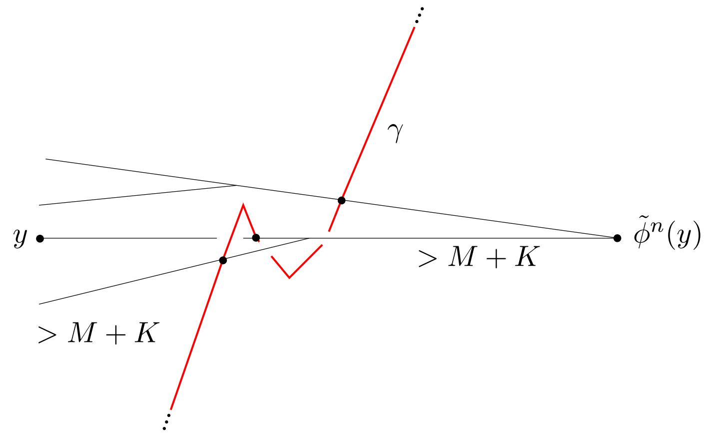

FIGURE 14. The M–deviating geodesic γ is separated by $T_n^o(\tilde{\phi}^n(y))$: their intersection consists of 3 points, all far from the root and leaves of $T_n^o(\tilde{\phi}^n(y))$.

Definition 5.9 (Exponentially expanding)**.** The train track map $\phi : V \to V$ is *exponentially expanding* if there exists an *expansion constant* $\varpi > 1$ such that for all edges e of V and all arcs $\alpha \subset e$, and all $L \geq 0$, we have $|\phi^L(\alpha)| \geq \varpi^L|\alpha|$. Note that if ϕ is an irreducible train track map and edges are expanding, then ϕ is exponentially expanding, as can be seen by taking ϖ to be the Perron-Frobenius eigenvalue of the transition matrix of ϕ. See Section 6.2 for more on the eigenvalues of the transition matrix.

The main result of this section is:

Proposition 5.10. *Suppose that $\phi : V \to V$ is a π_1-injective train track map. Let X be the mapping torus of ϕ. Suppose that $\pi_1 X$ is word-hyperbolic and that $\widetilde{X}$ satisfies:*

(1) $\widetilde{X}$ is level-separated.
(2) $\widetilde{X}$ has many effective walls.
(3) Every finite regular forward path fellow-travels at uniformly bounded distance with a periodic regular forward path.
(4) ϕ is exponentially expanding.

Then G acts freely and cocompactly on a CAT(0) cube complex.

Proof. Proposition 5.19 shows that there exists M such that every M-ladderlike geodesic is cut by a wall. Proposition 5.18 shows that each M-deviating geodesic is cut by a wall; Proposition 5.18 requires $\widetilde{X}$ to have bounded level-intersection, which is the case since ϕ is a train track map. The claim then follows from Proposition 5.1 since each geodesic that is not M-ladderlike is by definition M-deviating. □

Convention 5.11. In the remainder of this section, $\phi : V \to V$ is assumed to satisfy the initial hypotheses of Proposition 5.10, except that the enumerated hypotheses will be invoked as needed.

5.3. **Walls in $\widetilde{X}_L$.** Let $W \to X$ be an immersed wall with tunnel length $L \geq 1$, and suppose that $\overline{W}$ is a wall and $N(\mathbf{A}(\overline{W}))^1$ is (κ_1, κ_2)-quasi-isometrically embedded and κ-quasiconvex. Each primary bust has regular endpoints, by Lemma 3.5.(3), so that each level-part of $\overline{W}$ is disjoint from $\widetilde{X}^0$. Similarly, $\widetilde{X}^0$ is disjoint from $\mathbf{A}(S)$ for each slope S of $\overline{W}$.

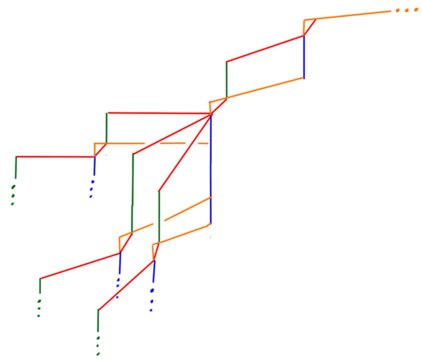

FIGURE 15. $\overline{W}_L$ and $\mathbf{A}(\overline{W})$ inside $\widetilde{X}^\bullet_L$.

Recall that $\widetilde{X}^\bullet_L$ denotes the subdivision of $\widetilde{X}_L$ obtained by pulling back the 1-skeleton of $\widetilde{X}$. For each $n \in \mathbb{Z}$, the inclusion $\widetilde{V}_{nL} \hookrightarrow \widetilde{X}$ lifts to an embedding $\widetilde{V}_{nL} \hookrightarrow (\widetilde{X}^\bullet_L)^1$, and we continue to use the notation $\widetilde{V}_{nL}$ for this subspace. We make the same observation and convention about $\widetilde{E}_{nL}$. By translating, we can assume that $\overline{W}$ has a primary bust in $\widetilde{V}_0$, and hence all primary busts in $\overline{W}$ lie in the various $\widetilde{V}_{nL}$ and the map $\widetilde{W} \to \widetilde{X}$ lifts to $\widetilde{W} \to \widetilde{X}^\bullet_L$. Let $\overline{W}_L$ be the image of $\widetilde{W} \to \widetilde{X}^\bullet_L$, so that we have the commutative diagram:

$$\begin{array}{ccc} \overline{W}_L & \to & \widetilde{X}^\bullet_L \\ \downarrow & & \downarrow \\ \overline{W} & \longrightarrow & \widetilde{X} \end{array}$$

Note that $\overline{W}_L$ and $\overline{W}$ are very similar: each tunnel $\overline{T}_L$ of $\overline{W}_L$ consists of a slope and a level-part that is a (subdivided) star, and $\overline{W}$ is obtained from $\overline{W}_L$ by folding each such subdivided star into a tree (see Figure 5 and Figure 7). The halfspaces $\overleftarrow{W}_L, \overrightarrow{W}_L$ in $\widetilde{X}^\bullet_L$ associated to $\overline{W}_L$ respectively map to the halfspaces $\overleftarrow{W}, \overrightarrow{W}$ in $\widetilde{X}$.

The approximation map $\mathbf{A}$ is defined in $\widetilde{X}_L$ just as it is in $\widetilde{X} = \widetilde{X}_1$. Consider $\mathbf{A} : \overline{W}_L \to \widetilde{X}_L$, which is a lift of $\mathbf{A} : \overline{W} \to \widetilde{X}$. There is a corresponding commutative diagram:

$$\begin{array}{ccc} \mathbf{A}(\overline{W}_L) & \to & \widetilde{X}^\bullet_L \\ \downarrow & & \downarrow \\ \mathbf{A}(\overline{W}) & \longrightarrow & \widetilde{X} \end{array}$$

in which the map $\mathbf{A}(\overline{W}_L) \to \mathbf{A}(\overline{W})$ is an isomorphism. Thus $\mathbf{A}(\overline{W}) \to \widetilde{X}$ lifts to an embedding $\mathbf{A}(\overline{W}) \to \widetilde{X}^\bullet_L$ whose image is $\mathbf{A}(\overline{W}_L)$. Figure 15 depicts $\overline{W}_L$ and $\mathbf{A}(\overline{W}_L)$.

There is also a lift $N(\mathbf{A}(\overline{W}))^1 \hookrightarrow \widetilde{X}^\bullet_L$. Since $N(\mathbf{A}(\overline{W}))^1 \hookrightarrow \widetilde{X}^1$ factors as $N(\mathbf{A}(\overline{W}))^1 \hookrightarrow (\widetilde{X}^\bullet_L)^1 \to \widetilde{X}^1$ and since $(\widetilde{X}^\bullet_L)^1 \to \widetilde{X}^1$ is distance nonincreasing, $N(\mathbf{A}(\overline{W}))^1 \to (\widetilde{X}^\bullet_L)^1$ is a (κ_1, κ_2)-quasi-isometric embedding. Thus $\partial N(\mathbf{A}(\overline{W}))^1$ embeds in $\partial \widetilde{X}^\bullet_L$ as a closed subset.

The following proposition explains that the tree $\mathbf{A}(\overline{W}_L)$ determines a wall in $\widetilde{X}^\bullet_L$, and therefore determines a coarse wall in $\widetilde{X}$ that coarsely agrees with $\overline{W}$.

Proposition 5.12. *For each lifted wall $\overline{W}_L$, the space $\widetilde{X}^\bullet_L$ contains subspaces $\overleftarrow{A}, \overrightarrow{A}$ such that $\overleftarrow{A} \cup \overrightarrow{A} = \widetilde{X}^\bullet_L$ and $\overleftarrow{A} \cap \overrightarrow{A} = \mathbf{A}(\overline{W}_L)$. Both $\overleftarrow{A} - \mathbf{A}(\overline{W}_L)$ and $\overrightarrow{A} - \mathbf{A}(\overline{W}_L)$ are connected. Moreover, the images of $\overleftarrow{A}$ and $\overrightarrow{A}$ under the map $\widetilde{X}^\bullet_L \to \widetilde{X}$ are coarsely equal to $\overleftarrow{W}_L$ and $\overrightarrow{W}_L$.*

Proof. It suffices to produce the subspaces $\overleftarrow{A}, \overrightarrow{A}$ so that each is coarsely equal to a component of $\widetilde{X}^\bullet_L - \overline{W}_L$. Let $\overleftarrow{W}_L, \overrightarrow{W}_L$ be the closures of the components of $\widetilde{X}^\bullet_L - \overline{W}_L$. The halfspaces $\overleftarrow{A}$ and $\overrightarrow{A}$ will be obtained from $\overleftarrow{W}_L$ and $\overrightarrow{W}_L$ by adding and subtracting "discrepancy zones", which are subspaces between $\overline{W}_L$ and $\mathbf{A}(\overline{W}_L)$ suggested by Figure 15.

Discrepancy zones: Let $e \subset \widetilde{V}_{nL}$ be a primary busted edge with outgoing long 2-cell $R_e \subset \widetilde{X}^\bullet_L$. Let $d \subset e$ be the closed primary bust with endpoints p, q. Let p', q' be the points at distance $\frac{1}{2}$ to the right of p, q within R_e. The slope S travels from p to q', as shown in Figure 16. Let $Z^\uparrow$ be the 2-simplex in R_e bounded by S and the part of $\mathbf{A}(S)$ between p and q'. The disc $Z^\uparrow$ is an *upward discrepancy zone*.

Let $\widetilde{C} \subset \widetilde{E}_{nL}$ be a nucleus in $\overline{W}_L$ and let $\mathbf{A}(\widetilde{C}) \subset \widetilde{V}_{nL+L}$ be its approximation. Consider the map $\widetilde{C} \times [\frac{1}{2}, L] \to \widetilde{X}^\bullet_L$ that restricts to the inclusion $\widetilde{C} \times \{t\} \hookrightarrow \widetilde{V}_{nL} \times \{t\} \subset \widetilde{X}^\bullet_L$ for $t < L$ and acts as the map $\tilde{\phi}^L : \widetilde{C} \to \widetilde{V}_{nL+L}$ on $\widetilde{C} \times \{L\}$. The image of this map is a *downward discrepancy zone* $Z^\downarrow$. In other words, $Z^\downarrow$ is the closure of $\widetilde{C} \times [\frac{1}{2}, L)$ in $\widetilde{X}^\bullet_L$. See Figure 17.

FIGURE 16. An upward discrepancy zone.

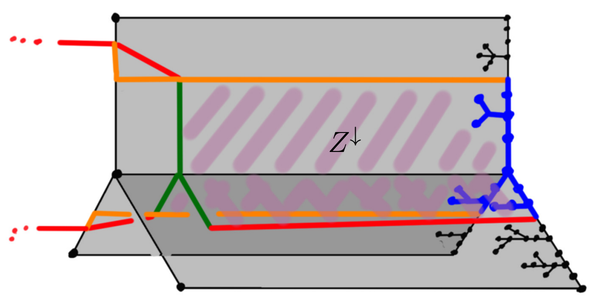

FIGURE 17. A downward discrepancy zone is shaded.

The halfspaces $\overleftarrow{A}$ and $\overrightarrow{A}$: Let $\mathfrak{Z}^\uparrow$ be the union of all upward discrepancy zones associated to $\mathbf{A}(\overline{W}_L)$, and likewise let $\mathfrak{Z}^\downarrow$ be the union of all downward discrepancy zones. Let

$$\overleftarrow{A} = Cl\left(\left(\overleftarrow{W}_L - \mathfrak{Z}^\uparrow\right) \cup \mathfrak{Z}^\downarrow\right) \quad \text{and} \quad \overrightarrow{A} = Cl\left(\left(\overrightarrow{W}_L - \mathfrak{Z}^\downarrow\right) \cup \mathfrak{Z}^\uparrow\right).$$

Since each discrepancy zone lies at distance less than L from $\overline{W}_L$, we see that $\overleftarrow{A}$ coarsely equals $\overleftarrow{W}_L$. By construction, $\overleftarrow{A} \cup \overrightarrow{A} = \widetilde{X}^\bullet_L$. Finally, suppose that $x \in \overleftarrow{A} \cap \overrightarrow{A}$. Then x must lie on the boundary of a discrepancy zone. If $x \in \overline{W}_L$, and $x \in \mathfrak{z}^\uparrow$, then $x \notin \overleftarrow{A}$ unless $x \in \mathbf{A}(\overline{W}_L) \cap \overline{W}_L$. Similarly, if $x \in \overline{W}_L$ and $x \in \mathfrak{z}^\downarrow$, then $x \notin \overrightarrow{A}$ unless $x \in \mathbf{A}(\overline{W}_L) \cap \overline{W}_L$. Hence $\overleftarrow{A} \cap \overrightarrow{A} \subseteq \mathbf{A}(\overline{W}_L)$. On the other hand, every point in $\mathbf{A}(\overline{W}_L)$ lies in the boundary of a discrepancy zone, and thus $\mathbf{A}(\overline{W}_L) \subseteq \overleftarrow{A} \cap \overrightarrow{A}$.

Observe that $\overleftarrow{A} - \mathbf{A}(\overline{W}_L)$ is homeomorphic to $\overleftarrow{W}_L - \overline{W}_L$, which is connected. Likewise $\left(\overrightarrow{A} - \mathbf{A}(\overline{W}_L)\right) \cong \left(\overrightarrow{W}_L - \overline{W}_L\right)$. Hence $\overleftarrow{A} - \mathbf{A}(\overline{W}_L)$ and $\overrightarrow{A} - \mathbf{A}(\overline{W}_L)$ are connected. □

5.4. **Lifting and cutting geodesics in $\widetilde{X}^\bullet_L$.** We now describe a criterion ensuring that a given geodesic in $\widetilde{X}$ is cut by a wall, in terms of quasigeodesics and walls $(\overleftarrow{A}, \overrightarrow{A})$ in $\widetilde{X}^\bullet_L$.

5.4.1. *Lifted augmentations of geodesics.* The following construction adjusts a bi-infinite quasigeodesic $\gamma \to \widetilde{X}^1$ so that it can be lifted to a bi-infinite quasigeodesic $\widehat{\gamma_\succ} \to \widetilde{X}^\bullet_L$ such that γ and $\widehat{\gamma_\succ}$ determine the same pair of points in $\partial\widetilde{X} \cong \partial\widetilde{X}^\bullet_L$.

Construction 5.13 (Lifted augmentations of quasigeodesics)**.** Let $\gamma : \mathbf{R} \to \widetilde{X}^1$ be an embedded quasigeodesic. The *augmentation* $\gamma_\succ$ of γ is defined as follows. For each (possibly trivial) bounded maximal horizontal subpath $P \subset \gamma$, with endpoints $p, p' \in \widetilde{V}_n, \widetilde{V}_{n'}$, let n'' be the smallest multiple of L greater than or equal to $\max\{n, n'\}$ and let $p'' = \tilde{\phi}^{n''-n}(p) = \tilde{\phi}^{n''-n'}(p')$. Let Q' be the horizontal path $pp''p'$, and replace P by Q'. Performing this replacement for each such P yields $\gamma_\succ$. Note that $\gamma_\succ$ is a quasigeodesic that L-fellowtravels with γ, so that $\partial\gamma_\succ = \partial\gamma$. We use the following notation. First, $P = P_1P_2$, where P_1 and P_2^{-1} are forward horizontal paths, one of which is trivial. Then $Q' = P_1QP_2$, where $Q = Q_1Q_1^{-1}$, with Q_1 a forward path. The terminal point p'' of Q_1 is the *apex* of Q, and $Q = Q_1Q_1^{-1}$ is an *augmentation* of γ.

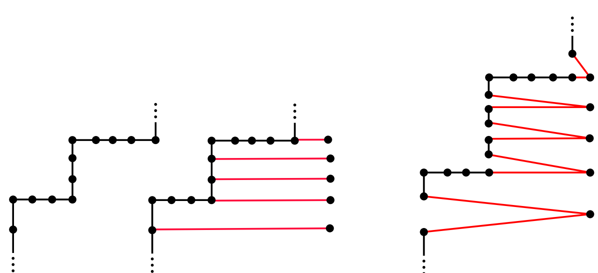

Figure 18. At left is part of a geodesic in $\widetilde{X}$. In the middle is the image of its augmentation. At right is the lifted augmentation in $\widetilde{X}^\bullet_L$.

The path $\gamma_\succ$ lifts to a quasigeodesic $\widehat{\gamma_\succ} \to \widetilde{X}^\bullet_L$. More specifically, each lift of the union of the vertical edges of $\gamma_\succ$ determines a unique lift of $\gamma_\succ$ to a quasigeodesic. Indeed, we can write $\gamma_\succ$ in one of the following four forms, where each e_i is a vertical edge and each A_i, B_i is a horizontal path, with A_i starting at an apex and each B_i ending at an apex:

(1) $\cdots A_{-1}e_{-1}B_{-1}A_0e_0B_0A_1e_1B_1A_2e_2B_2\cdots$, where $e_{\pm i}$ are present for all $i \in \mathbb{N}$;
(2) $A_0e_0B_0A_1\cdots$, where A_0 is unbounded;
(3) $\cdots A_0e_0B_0$, where B_0 is unbounded;

(4) $A_s e_s B_s \cdots A_t e_t B_t$ with A_s, B_t unbounded. (This includes the case $B_0 A_1$ in which $\gamma = \gamma_\succ$ is horizontal.)

Observe that each lift of e_i determines lifts of A_i and B_i to $\widetilde{X}_L^\bullet$. Since the apexes lift uniquely, any lift of B_i is concatenable with any lift of A_{i+1}, and we conclude that a lift of $\{e_i\}$ induces a lift of $\gamma_\succ$. In the case where γ is horizontal, $\gamma = \gamma_\succ$ lifts uniquely since any horizontal path starting and ending in $\cup_k \widetilde{V}_{kL}$ lifts uniquely. Under the quasi-isometry $(\widetilde{X}_L^\bullet)^1 \to \widetilde{X}^1$, the quasigeodesic $\widehat{\gamma_\succ}$ is sent to $\gamma_\succ$, and thus $\partial\widehat{\gamma_\succ} = \partial\gamma$. Finally, if some augmentation of γ has a subpath that lifts to a backtrack in $\widehat{\gamma}_\succ$, then we truncate $\gamma_\succ$ accordingly and define $\widehat{\gamma}_\succ$ to be the lift of the truncated augmentation. An augmentation where this truncation is nontrivial is a *truncated augmentation*. We call $\widehat{\gamma_\succ}$ a *lifted augmentation* of γ. See Figure 18.

5.4.2. *Cutting in $\widetilde{X}_L^\bullet$.* We now establish a criterion, in terms of lifted augmentations, ensuring that a wall $\overline{W}$ cuts a given quasigeodesic in $\widetilde{X}^1$.

Proposition 5.14. *Let $\gamma : \mathbf{R} \to \widetilde{X}^1$ be an embedded quasigeodesic, and let $\widehat{\gamma_\succ} \to \widetilde{X}_L^\bullet$ be a lifted augmentation. Let C_o be a bounded subset of $\widehat{\gamma_\succ} \cap \mathbf{A}(\overline{W}_L)$ and let C be the smallest subgraph of $\widehat{\gamma_\succ}$ containing C_o. Let $\widehat{\gamma_\succ} \vee_C N(\mathbf{A}(\overline{W}_L))^1 \to (\widetilde{X}_L^\bullet)^1$ be the graph obtained by wedging $\widehat{\gamma_\succ} \to \widetilde{X}$ and $N(\mathbf{A}(\overline{W}_L))^1 \to \widetilde{X}$ along the common subgraph C. Suppose that:*

(1) $\widehat{\gamma_\succ} \vee_C N(\mathbf{A}(\overline{W}_L))^1 \to (\widetilde{X}_L^\bullet)^1$ is a quasi-isometric embedding.

(2) There are nontrivial intervals $f, f' \subset \widehat{\gamma_\succ}$, immediately preceding and succeeding C_o within $\widehat{\gamma_\succ}$, that lie in $\overleftarrow{A}$ and $\overrightarrow{A}$ respectively.

(3) For every component D of $\widehat{\gamma_\succ} \cap \mathbf{A}(\overline{W}_L)$ disjoint from C_o, the 1-neighborhood in $\widehat{\gamma_\succ}$ of D lies entirely in $\overleftarrow{A}$ or $\overrightarrow{A}$.

Then $\overline{W}$ cuts γ.

Proof. Hypotheses (2) and (3) together imply that $\widehat{\gamma_\succ}$ decomposes as a concatenation $\overleftarrow{\gamma}\bar{\gamma}\overrightarrow{\gamma}$, where $\bar{\gamma}$ is a bounded path containing C_o and $\overleftarrow{\gamma}, \overrightarrow{\gamma}$ are rays contained in $\overleftarrow{A}, \overrightarrow{A}$ respectively. The image of $\widehat{\gamma_\succ} \vee_C N(\mathbf{A}(\overline{W}_L))^1 \to \widetilde{X}_L^\bullet$ is $\widehat{\gamma_\succ} \cup N(\mathbf{A}(\overline{W}_L))^1$, which is quasi-isometrically embedded in $(\widetilde{X}_L^\bullet)^1$ by hypothesis (1). The inclusion $\widehat{\gamma_\succ} \cup N(\mathbf{A}(\overline{W}_L))^1 \hookrightarrow (\widetilde{X}_L^\bullet)^1$ thus induces an embedding $\partial\widehat{\gamma_\succ} \sqcup \partial N(\mathbf{A}(\overline{W}_L))^1 \to \partial\widetilde{X}_L^\bullet$: the two points of $\partial\widehat{\gamma_\succ}$ are $\partial\overleftarrow{\gamma} \in \partial\overleftarrow{A}$ and $\partial\overrightarrow{\gamma} \in \partial\overrightarrow{A}$, and neither of these points lies in $\partial N(\mathbf{A}(\overline{W}_L))^1$ since hypothesis (1) implies that no sub-ray of $\widehat{\gamma_\succ}$ lies in a bounded neighborhood of $N(\mathbf{A}(\overline{W}_L))^1$. Applying the quasi-isometry $\widetilde{X}_L^\bullet \to \widetilde{X}$ shows that the points of $\partial\gamma \subset \partial\widetilde{X}$ lie in $\partial N(\overleftarrow{W}) - \partial\overline{W}$ and $\partial N(\overrightarrow{W}) - \partial\overline{W}$, whence $\overline{W}$ cuts γ. □

5.4.3. *Narrow discrepancy zones.* We now analyze discrepancy zones. Specifically, we need the following notion of "narrow exceptional discrepancy zones", and the ability to construct immersed walls with narrow exceptional zones, in order to use Proposition 5.14 to cut geodesics.

Definition 5.15 (Exceptional zone, narrow exceptional zones)**.** Let $W \to X$ be an immersed wall with tunnel-length L. An *exceptional zone* is a downward discrepancy zone in $\widetilde{X}_L^\bullet$ whose boundary path intersects the interior of a slope approximation. The downward discrepancy zone shown in Figure 17 is exceptional.

We say that W has *narrow exceptional zones* if for each exceptional zone $Z \subset \widetilde{X}_L^\bullet$ associated to a nucleus $\widetilde{C}$ of $\overline{W}_L$, the image in $Z \subset \widetilde{X}_L^\bullet$ of $\widetilde{C} \times [\frac{1}{2}, \frac{3L}{4}]$ does not contain a vertex. See Figure 19, in which the exceptional zone at left is not narrow and the exceptional zone at right is narrow. (There is also a non-exceptional downward discrepancy zone at right.)

Lemma 5.16. *Suppose that $\phi : V \to V$ is a train track map with expanding edges. For each edge e_i of V, let y'_i be a periodic regular point of e_i, and let $\epsilon' \in (0, \min_i \mathsf{d}_V(y'_i, V^0))$. There exists $L' = L'(\epsilon', \{y'_i\})$ so that the following holds. Let $\{d'_i\}$ be a set of primary busts with each $d'_i \subset \mathcal{N}'_\epsilon(y'_i)$, and let $L \geq L'$ be chosen so that L and $\{d'_i\}$ satisfy the conclusion of Lemma 3.5.*

Then for any choice of nontrivial subintervals $\{d_i \subset d'_i\}$, there is an immersed wall $W \to X$ with tunnel length L and primary busts $\{d_i\}$, with the following property: each exceptional discrepancy zone Z lies in the interior of a single long 2-cell of $\widetilde{X}^\bullet_L$, and hence Z intersects a single slope-approximation.

Proof. Let α be a component of $e_i - \mathcal{N}'_\epsilon(y'_i)$. Then there exists L'_i such that the path $\phi^L(\alpha)$ traverses an entire edge provided $L \geq L'_i$. Hence, for any subinterval d of $\mathcal{N}'_\epsilon(y'_i)$, the same is true for any component α' of $e_i - d$, since $\alpha \subset \alpha'$. Let $L' = \max_i L'_i$. Given $L \geq L'$ and the primary busts $\{d'_i \subset \mathcal{N}'_\epsilon(y'_i)\}$ satisfying the conclusions of Lemma 3.5, observe that the conclusions of Lemma 3.5 continue to hold (with the same L) when each d'_i is replaced by a nontrivial subinterval d_i. Thus, for each $L, \{d_i\}$ as above, we have an immersed wall $W \to X$ with tunnel-length L and primary busts $\{d_i\}$ so that, for all i and all components α of $e_i - d_i$, the path $\phi^L(\alpha)$ traverses an entire edge.

Let $\mathbf{A}(S)$ be a slope-approximation, associated to an exceptional zone and lying in a long 2-cell R based at the vertical edge $e \subset \widetilde{V}_n$. Then the nucleus $\widetilde{C}$ incident to S is the copy in $\widetilde{E}_n$ of a subinterval of the interior of e. Let α be a component of $e - \mathrm{Int}(d)$, where d is the primary bust associated to S. Then for all sufficiently large L, the path $\phi^L(\alpha)$ traverses an entire edge, and therefore contains a primary bust, and the claim follows. ☐

Lemma 5.17. *Suppose that $\phi : V \to V$ is a train track map with exponentially expanding edges. Let $y_1, \ldots, y_s \in V$ be regular points such that each edge of V contains exactly one y_i, and let $\epsilon > 0$. Then for all sufficiently large L, there exists an immersed wall $W \to X$ with tunnel-length L, such that each primary bust is in the ϵ-neighborhood of some y_i, and W has narrow exceptional zones.*

Proof. Let $\varpi > 1$ be the expansion constant of ϕ (see Definition 5.9). For each i, let $y'_i \in V$ be a periodic regular point in the edge e_i containing y_i with $\mathsf{d}_{e_i}(y'_i, y_i) < \frac{\epsilon}{2}$; such a y'_i exists since periodic points are dense in each edge by Lemma 6.19. Let $\chi_i = \min\{\mathsf{d}_V(\phi^k(y'_i), V^0) \colon k \geq 0\}$, which is positive since y'_i is periodic and regular. Let $\chi = \min_i \chi_i$. Let

$$L_0 = \max\left\{4 \log_\varpi \left(\frac{2 \max_i |e_i|}{\chi}\right),\ \log_\varpi \frac{\epsilon}{\chi}\right\}.$$

We now apply Lemma 5.16 to the collection of points $\{y'_i\}$, choosing primary busts d_i, each in the $\epsilon' = \frac{\epsilon}{2\varpi^{L_0}}$–neighborhood of y'_i, and tunnel-length $L \geq \max\{L_0 + 4, L'\}$, where L' is the constant from Lemma 5.16.

Let Z be the image in $\widetilde{X}$ of an exceptional zone between $\overline{W}$ and $\mathbf{A}(\overline{W})$. By Lemma 5.16, there is a unique slope S such that the forward part of $\mathbf{A}(S)$ forms part of the boundary path of Z. See Figure 19. If $v \in Z$ is a vertex of some $\widetilde{V}_n$ at horizontal distance $h > \lfloor \frac{L}{4} \rfloor$ from the nucleus-approximation $\mathbf{A}(\widetilde{C})$ on the right of Z, then since $L \geq L_0$, the right boundary path of Z contains a complete edge e', and thus a primary bust, which is impossible.

To see this, let ζ be the vertical geodesic arc from v to the forward part of $\mathbf{A}(S)$. Then the right boundary path of Z contains $\phi^h(\zeta)$, which satisfies

$$|\phi^h(\zeta)| > |\phi^{\lfloor \frac{L}{4} \rfloor}(\zeta)| \geq \varpi^{\lfloor \frac{L}{4} \rfloor}|\zeta| > \varpi^{\frac{L_0}{4}}|\zeta| \geq \frac{2 \max_i |e_i|}{\chi}|\zeta|.$$

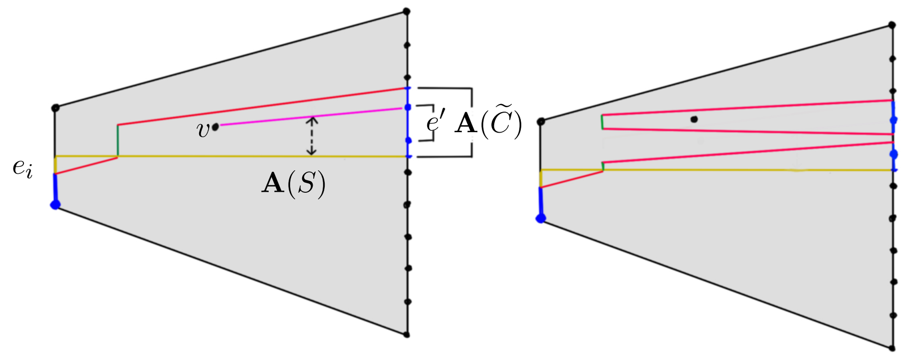

FIGURE 19. The exceptional zone corresponding to $\mathbf{A}(S)$ cannot contain the vertex v when L is sufficiently large. Such a vertex v could only be contained in a non-exceptional downward discrepancy zone, as shown at right.

Now $|\zeta| \geq \chi - \frac{\epsilon}{2\varpi^{L_0}}$, so our choice of L_0 ensures $|\zeta| \geq \frac{\chi}{2}$, whence $|\phi^h(\zeta)| > \max_i |e_i|$. Hence $\phi^h(\zeta)$ traverses a complete edge as claimed. □

5.5. **Cutting deviating geodesics.**

Proposition 5.18. *Let $\widetilde{X}$ satisfy the hypotheses of Proposition 5.10, let $M \geq 0$, and let $\gamma : \mathbf{R} \to \widetilde{X}^1$ be an M-deviating geodesic. Then there exists a wall $\overline{W}$ such that $N(\mathbf{A}(\overline{W}))^1$ is quasiconvex in $\widetilde{X}^1$ and $\overline{W}$ cuts γ.*

Proof. We will find a wall $\overline{W} \to \widetilde{X}$ satisfying the hypotheses of Proposition 5.14.

An oddly-intersecting forward path: Since $\widetilde{X}$ is level-separated, there exists $z \in \widetilde{X}$ such that for all sufficiently large n, there is a regular level $\mathcal{T}_n = T^o_n(\tilde{\phi}^n(z))$ with a leaf at z, such that $\mathcal{T}_n$ has odd intersection with γ and the distance in $\mathcal{T}_n$ from $\gamma \cap \mathcal{T}_n$ to the root or to any leaf of $\mathcal{T}_n$ exceeds $12(M + \delta)$.

The fact that $\widetilde{X}$ has bounded level intersection and γ is M-deviating implies that there exists N and a finite, odd-cardinality set $C'_o \subset \gamma$ such that $\mathcal{T}_n \cap \gamma = C'_o$ for all $n \geq N$. Each $\mathcal{T}_n$ is the union of finitely many maximal forward paths emanating from leaves. For each $n \geq N$, we wish to choose a leaf y of $\mathcal{T}_n$ such that the maximal forward path $\sigma_n \subset \mathcal{T}_n$ emanating from y has the property that $\sigma_n \cap \gamma$ is a fixed odd-cardinality subset $C_o \subseteq C'_o$. However, to achieve this, we shall replace γ with an embedded deviating uniform quasigeodesic that coincides with the original γ outside a diameter-$2M$ subset, as follows.

We now describe the modification of γ. Let $e_1, \ldots, e_{|C'_o|}$ be the edges of γ intersecting $\mathcal{T}_n$ for $n \geq N$. Index these so that e_i precedes e_j in the geodesic γ if and only if $i < j$. The set $\{e_1, \ldots, e_{|C'_o|}\}$ is partially ordered as follows: $e_i \preceq e_j$ if there exists a forward path in $\mathcal{T}_n$ originating at e_i and passing through e_j. The edges e_i, e_j are *confluent* if there exists k such that $e_i, e_j \preceq e_k$. We claim that confluence is an equivalence relation; it suffices to check transitivity. If e_i, e_j are confluent (witnessed by forward paths from e_i, e_j to some e_k), and $e_j, e_{i'}$ are confluent (witnessed by forward paths from $e_{i'}, e_j$ to some $e_{k'}$), then since ϕ is a train track map, the forward paths from e_j to e_k and $e_{k'}$ have the same initial point, whence $e_i \preceq e_{k'}$ or $e_{i'} \preceq e_k$, i.e. $e_i, e_{i'}$ are confluent. Observe that there is exactly one confluence class for each $\preceq$-maximal edge. Since $|C'_o|$ is odd, there exists an odd-cardinality confluence class, and we let e_k be its $\preceq$-maximal element. Let e_i, e_j be the elements of the confluence class of e_k such that the indices i, j are respectively minimal and maximal. Let α_i, α_j be forward paths in $\mathcal{T}_n$ joining e_i, e_j to e_k. Let A_i be an embedded combinatorial path in the forward ladder $N(\alpha_i)$ that joins the terminal vertex v_i of e_i to a vertex v_k of e_k and does not intersect α_i. The edge e_j contains a vertex v_j on the same side of $\mathcal{T}_n$ as the terminal vertex of e_i. Let A_j be an embedded combinatorial path in $N(\alpha_j)$ joining v_j to v_k and not intersecting α_j. Since γ is deviating, $\mathsf{d}(e_i, e_k)$ and $\mathsf{d}(e_j, e_k)$ are uniformly bounded. Hence, since $N(\alpha_i)^1, N(\alpha_j)^1$ are uniformly quasiconvex, the paths A_i, A_j have uniformly bounded length. Let A be the path obtained from $A_i A_j^{-1}$ by removing backtracks. We replace the subpath of γ between v_i and v_j by A, and thus replace γ by a bi-infinite embedded quasigeodesic γ'. By construction, for all $n \geq N$, there is a forward path σ_n of $\mathcal{T}_n$, that intersects the modified path exactly once, namely in a point of e_i. The argument proceeds using the new γ', which is an embedded quasigeodesic that is M'-deviating, with M' a new, larger constant. However, since the quasi-isometry constants of γ' play no essential role in the argument, we will continue, for simplicity, with γ and M as before.

By the above construction, there exists $\epsilon > 0$ such that for all $x \in \mathcal{N}_\epsilon(y)$, any forward path σ_x of length $n \geq N$ emanating from x intersects γ in a set C_o^x of interior points of edges that has the same cardinality as C_o and has the property that the smallest subcomplex C containing C_o is exactly the smallest subcomplex containing C_o^x. The wall we will choose will contain a slope S such that $\mathbf{A}(S)$ contains such a σ_x as its forward part.

Quasi-isometric embedding of $\gamma \vee_C N(\mathbf{A}(\overline{W}))^1 \to \widetilde{X}^1$: Let $W \to X$ be an immersed wall such that every edge of V contains a primary bust, and suppose $\overline{W} \subset \widetilde{X}$ is the image of a lift $\widetilde{W} \to \widetilde{X}$ such that $\overline{W}$ contains a slope S with the forward part of $\mathbf{A}(S)$ equal to a path σ_x, emanating from some $x \in \mathcal{N}_\epsilon(y)$, as above. Suppose moreover that W was drawn from a set of immersed walls with uniformly bounded ladder-overlap.

Since every edge contains a primary bust, Proposition 4.1 provides constants L_0, κ_1, κ_2, depending only on $\widetilde{X}$, such that if the tunnel length of W is at least L_0, then $N(\mathbf{A}(\overline{W}))^1$ is (κ_1, κ_2) quasi-isometrically embedded. Recall also that $\overline{W}$ is a genuine wall if the tunnel-length exceeds a uniform constant L_1, by Proposition 4.6.

There exist constants $L_2 \geq L_1, \kappa_1', \kappa_2'$, depending on $\widetilde{X}$ and M such that if W has tunnel-length at least L_2, then $\gamma \vee_C N(\mathbf{A}(\overline{W}))^1 \to \widetilde{X}^1$ is a (κ_1', κ_2')-quasi-isometric embedding. Indeed, this follows from an application of Lemma 4.3, since γ is M-deviating and hence has uniformly bounded $(3\delta + 2\lambda)$-overlap with $\mathbf{A}(S)$.

Verification that $\gamma \vee_{C_o} \mathbf{A}(\overline{W})$ embeds: By construction, γ does not intersect any point of $\mathbf{A}(S)$ outside of C_o. Hence suppose that $\tau\beta_1\alpha_1 \cdots \beta_k\alpha_k\beta_{k+1}$ is a path in $N(\mathbf{A}(\overline{W}))^1 \cup \gamma$ that begins and ends in C_o, such that: τ is a subpath of γ, and each β_i lies in the carrier of a slope-approximation, and each α_i lies in a nucleus approximation, and $|\beta_i| \geq L$ except when $i = k+1$. If L is sufficiently large and $|\beta_2| \geq L$, then the existence of such a closed path contradicts the above conclusion that $N(\mathbf{A}(\overline{W}))^1 \vee_C \gamma$ uniformly quasi-isometrically embeds. The remaining possibility is that a path of the form $\tau\beta_1\alpha_1\beta_2$ or $\tau\beta_1\alpha_1$ is closed in $\widetilde{X}$. In either case, when L is sufficiently large, a thin quadrilateral argument shows that γ is forced to $(2\delta + \lambda)$-fellow-travel with β_1 or β_2 for distance exceeding M. Indeed, the fellow-traveling between α_1 and β_i is controlled by Lemma 2.6, since α_1 is vertical, while the fellow-traveling between β_1, β_2 (if the latter exists) is controlled by construction. Thinness of the quadrilateral thus forces fellow-traveling between β_1 and γ. Hence $\gamma \vee_{C_o} \mathbf{A}(\overline{W})$ embeds in $\widetilde{X}$.

Preventing short augmentations from crossing $\mathbf{A}(\overline{W}_L)$ at an apex: We now compute the tunnel-length $L_3 \geq L_2$ necessary to ensure that each augmentation QQ^{-1} in $\gamma_\succ$ either fails to intersect $\mathbf{A}(\overline{W})$ or has length at least $\frac{L}{4}$, where $L \geq L_3$ is the tunnel-length of

W. Note that if QQ^{-1} is a truncated augmentation in the sense of Construction 5.13, then the apex lies in some $\widetilde{V}_n$ with $n \notin L\mathbb{Z}$, and hence $QQ^{-1} \cap \mathbf{A}(\overline{W}) = \emptyset$, so we only need consider non-truncated augmentations.

Let W have tunnel-length $L \geq L_2$ and let QQ^{-1} be an augmentation whose apex p lies in $\mathbf{A}(\overline{W})$, and hence in a nucleus-approximation. Suppose that $|Q| \leq \frac{L}{4}$. Let γ' be the subpath of γ between C and the initial point of Q, let β be a geodesic of $N(\mathbf{A}(\overline{W}))^1$ joining p to the terminal point of $\mathbf{A}(S)$, and let τ be a geodesic of $N(\mathbf{A}(S))^1$ joining the initial point of γ' to the terminal point of β. Since γ is deviating, the path $\gamma' Q$ is a quasigeodesic with constants depending only on M and λ. Meanwhile, since $\overline{W}$ has uniform ladder-overlap and $L \geq L_2 \geq L_0$, the path $\beta\tau^{-1}$ is a (κ_1, κ_2)-quasigeodesic. Hence $\gamma' Q$ fellow-travels with $\tau\beta^{-1}$ at distance depending only on M and $\widetilde{X}$. This is impossible for sufficiently large L, since γ', τ have $(2\delta + 2\lambda)$-overlap of length at most M. (Note that fellow-traveling between subpaths of β and γ' would force impossible fellow-traveling between β or γ' and τ.)

Choosing $\overline{W}$: Since $\widetilde{X}$ has many effective walls, there exists an immersed wall $W \to X$ with tunnel length $L \geq L_3$, involving a primary bust in every edge of V, such that the image $\overline{W}$ of a lift $\widetilde{W} \to \widetilde{X}$ satisfies the following:

(1) $\overline{W}$ is a wall (since $L_3 \geq L_2$).
(2) $N(\mathbf{A}(\overline{W}))^1$ is (κ_1, κ_2)-quasiconvex.
(3) $\gamma \cap N(\mathbf{A}(\overline{W}))^1 = C$, which is contained in the carrier of a slope-approximation $\mathbf{A}(S)$.
(4) $N(\mathbf{A}(\overline{W}))^1 \vee_C \gamma \to \widetilde{X}^1$ is a quasi-isometric embedding.
(5) Any augmentation QQ^{-1} of γ that intersects $\mathbf{A}(\overline{W})$ has the property that $|Q| > \frac{L}{4}$ (since $L \geq L_3$).

W is chosen from the spreading set $\mathbb{W}$ given in Definition 5.3.(1).

An arbitrary lifted augmentation: Let $\widehat{\gamma_{\succ}} \to \widetilde{X}^\bullet_L$ be a lifted augmentation of γ. Since the map $\widetilde{X}^\bullet_L \to \widetilde{X}$ is a quasi-isometry and restricts to the identity on $\mathbf{A}(\overline{W}_L)$ and sends $\widehat{\gamma_{\succ}}$ to $\gamma_{\succ}$, the intersection $\widehat{C} = \widehat{\gamma_{\succ}} \cap N(\mathbf{A}(\overline{W}_L))^1$ is bounded and $\widehat{\gamma_{\succ}} \vee_{\widehat{C}} N(\mathbf{A}(\overline{W}_L))^1 \to \widetilde{X}^\bullet_L$ is a quasi-isometric embedding. (We could have chosen a specific lifted augmentation to make $\widehat{C} \neq \emptyset$, but it follows from the discussion below that this holds for any lifted augmentation.) Thus any $\widehat{\gamma_{\succ}}$, together with $\mathbf{A}(\overline{W}_L)$, satisfies Hypothesis (1) of Proposition 5.14.

We now verify that $\widehat{\gamma_{\succ}}$ satisfies the remaining two hypotheses of Proposition 5.14. To this end, let $\hat{\eta}$ be an embedded quasigeodesic in $\widetilde{X}^\bullet_L$ obtained from $\tilde{\phi}_L \circ \widehat{\gamma_{\succ}}$ by removing backtracks, and let η be the image in $\widetilde{X}$ of $\hat{\eta}$. Note that $\hat{\eta}$ is independent of the choice of lifted augmentation of γ.

Intersection of $\hat{\eta}$ with $\mathbf{A}(\overline{W})$: We first work in $\widetilde{X}$. Recall that $\gamma \cap \mathbf{A}(\overline{W})$ is the odd-cardinality set C_o of points in $\mathbf{A}(S)$ for some slope S of $\overline{W}$. Consider the nucleus $\widetilde{M} \subset \overline{W}$ that intersects S and $\mathbf{A}(S)$. Then, since we can assume that L is sufficiently large to ensure that each primary bust is separated from each vertex by a secondary bust, $\widetilde{M}$ corresponds to a subinterval containing no vertex, and hence $\mathbf{A}(\widetilde{M})$ maps to a subspace of the star of a vertex in V. By Lemma 5.17 and our above choice of L, the exceptional zone determined by $\widetilde{M}$ and $\mathbf{A}(\widetilde{M})$ contains no vertex of a vertical edge containing a point of C_o. It follows that the tunnel T' attached to the unique secondary bust of $\widetilde{M}$ intersects γ in a set of points corresponding bijectively to C_o. Let S' be the slope of T'. Then, since the primary busts can be chosen arbitrarily small, we can assume that $\mathbf{A}(S') \cap \eta$ is an odd-cardinality set E'_o. Hence $\mathbf{A}(S') \cap \hat{\eta} \subset \widetilde{X}^\bullet_L$ is an odd-cardinality set E_o mapping bijectively to E'_o.

Since the endpoints of primary busts are regular, the path $\hat{\eta}$ contains a nontrivial interval I ending at a point of E_o and lying in the image of S' under the forward flow; hence $I \subset \overrightarrow{W}_L$ and, since I does not lie in a discrepancy zone, $I \subset \overrightarrow{A}$. There is likewise a nontrivial interval I' in $\hat{\eta}$ beginning on E_o and lying in an exceptional zone determined by the nucleus intersecting $\mathbf{A}(S')$. Moreover, I' and I can be chosen to be separated in $\hat{\eta}$ by E_o. See Figure 20.

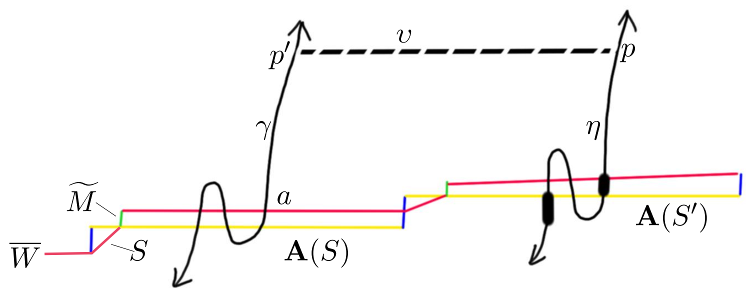

FIGURE 20. The relationship between $\gamma, \eta, \overline{W}, \mathbf{A}(\overline{W})$ in $\widetilde{X}$. The intervals I, I' in $\widetilde{X}^\bullet_L$ map to the bold intervals. The path θ contains the terminal part of $\mathbf{A}(S)$ and the initial part of $\mathbf{A}(S')$ (except for the last case considered in the proof).

We claim that $\mathbf{A}(\overline{W}) \cap \hat{\eta} = E_o$. Otherwise, applying the map $\widetilde{X}^\bullet_L \to \widetilde{X}$ would show that $\mathbf{A}(\overline{W}) \cap \eta$ contains some point $p \notin E'_o$, since $\widetilde{X}^\bullet_L \to \widetilde{X}$ restricts to a bijection on $\mathbf{A}(\overline{W})$. Then there is a forward path υ of length L emanating from a point $p' \in \gamma$ and terminating at p. Let θ be a geodesic of $\mathbf{A}(\overline{W})$ joining p to a closest point a of C_o, and let γ_1 be the subpath of γ joining a to p'. Then $\gamma_1\upsilon$ is a quasigeodesic with quasi-isometry constants depending only on the deviation constant of γ, while θ is a (κ_1, κ_2)-quasigeodesic. Hence $\gamma_1\upsilon$ fellowtravels with θ at distance depending only on $\delta, M, \kappa_1, \kappa_2$ and not on L. It follows that there is a uniform upper bound on $|\gamma_1|$ that is independent of L. Hence, if L is sufficiently large, then since C_o lies at distance at least $\frac{L}{4}$ from all nuclei, $\min_n |q(p) - nL| \geq \frac{L}{4}$, so that p lies at horizontal distance at least $\frac{L}{4}$ from any nucleus approximation. Suppose that $\mathbf{A}(S')$ lies in θ. Then θ contains a point at distance at least $\frac{L}{4}$ from $\gamma_1\upsilon$, and hence $\gamma_1\upsilon$ and θ cannot uniformly fellow-travel when L is sufficiently large. Similarly, if θ enters some other slope-approximation attached to $\mathbf{A}(\widetilde{M})$, we find that θ and $\gamma_1\upsilon$ cannot fellow-travel. The remaining possibility is that there is a path in $\mathbf{A}(\widetilde{M})$ joining the endpoint of $\mathbf{A}(S)$ to a point of υ. This is impossible since, as established above, the level part of T' has odd-cardinality intersection with γ and $p \notin E'_o$.

Conclusion: It follows from the above discussion that $\hat{\eta}$ contains two quasigeodesic rays, one in each of the halfspaces of $\widetilde{X}^\bullet_L$ associated to $\mathbf{A}(\overline{W})$. Since $\hat{\eta}$ fellow-travels with $\widehat{\gamma_\succ}$, we see that $\widehat{\gamma_\succ}$ satisfies all hypotheses of Proposition 5.14, whence $\overline{W}$ cuts γ. $\square$

5.6. **Cutting ladderlike geodesics.**

Proposition 5.19. *Suppose that $\widetilde{X}$ has many effective walls and for each bounded forward path α there exists a periodic regular forward path α' such that $N(\alpha) = N(\alpha')$.*

Then for each geodesic $\gamma : \mathbf{R} \to \widetilde{X}^1$ that is not M-deviating for any M, there exists an immersed wall $W \to X$ such that $\overline{W}$ is a wall, $N(\mathbf{A}(\overline{W}))^1$ is quasiconvex, and $\overline{W}$ cuts γ.

Proof. Suppose that γ contains a path γ' such that for some regular $x \in \widetilde{X}^1$ and some M to be determined, the path γ' fellowtravels at distance $(2\delta + 2\lambda)$ with the sequence $x, \tilde{\phi}(x), \ldots, \tilde{\phi}^M(x)$, where x is a periodic regular point. Such γ' exists for arbitrarily large M by combining the fact that γ is M-ladderlike for arbitrarily large M with the first hypothesis. We shall show that if M is sufficiently large, then there exists a wall $\overline{W}$ that has the desired properties and cuts γ and separates x and $\tilde{\phi}^M(x)$.

Choosing W using many effective walls: Without loss of generality, M is an even integer, and we let $a = \tilde{\phi}^{M/2}(x)$. Note that a is periodic and regular. Let $\{e_i\}$ be the collection of edges of V, and let $W \to X$ be an immersed wall busting each e_i, with tunnel-length L to be determined. Let e_1 be the edge whose interior contains a. By Remark 5.4 and the fact that $\widetilde{X}$ has many effective walls, there exist κ_1, κ_2, L_1 depending only on $\widetilde{X}$ such that we can choose W with tunnel length $L \geq L_1$ so that $\overline{W}$ is a wall and $N(\mathbf{A}(\overline{W}))^1$ is (κ_1, κ_2)-quasi-isometrically embedded. Moreover, we choose W from the the collection $\mathbb{W}_a$ of Definition 5.3.(2), which guarantees that $\overline{W}$ can be chosen with the following properties:

(1) There exists $k \geq 0$ such that for each primary bust d with an endpoint in $\overline{W}$ in the same knockout as a, we have $\mathsf{d}(\tilde{\phi}^n(a), \tilde{\phi}^n(d)) \geq 3\delta + 2\lambda$ for all $n \geq k$.
(2) W has tunnel length $L > \max\{12(\delta + k), L_1\}$, independent of M.
(3) The image of a in V lies in the interior of a nucleus of W and so $\mathbf{A}(\overline{W})$ contains $\tilde{\phi}^L(a)$.

We assume that $M > JL$, where $J \geq 4$ will be chosen below. Let σ be the uniform quasi-geodesic in $\widetilde{X}^1$ obtained from γ by removing γ' and replacing it by the sequence $x, \tilde{\phi}(x), \ldots, \tilde{\phi}^M(x)$ (see Figure 21).

Verifying that $\sigma \vee_{\phi^L(a)} N(\mathbf{A}(\overline{W}))^1$ quasi-isometrically embeds: Consider paths of the form $\alpha_0\beta_0 \cdots \beta_{s-1}\alpha_s\tau$, where β_i is a geodesic of the carrier of a slope-approximation, α_i is a vertical geodesic of $N(\mathbf{A}(\overline{W}))^1$, and α_s terminates at $\tilde{\phi}^L(a)$, and τ is a subpath of σ beginning at the endpoint $\tilde{\phi}^L(a) = \tilde{\phi}^{L+\frac{M}{2}}(x)$ of α_s. (We remind the reader that, because of how σ was defined, the initial part of τ is a subsequence of $x, \tilde{\phi}(x), \ldots \tilde{\phi}^{\frac{M}{2}+L}(x)$ or of $\tilde{\phi}^{\frac{M}{2}+L}(x), \ldots \tilde{\phi}^M(x)$.) The $(3\delta + 2\lambda)$-overlap between α_s and τ and between α_i and β_i and between α_i and β_{i-1} is controlled by Lemma 2.6, and Condition (1) on $\overline{W}$ ensures that the $(3\delta + 2\lambda)$–overlap between τ and β_{s-1} has length at most k. The choice of L now allows us to invoke Lemma 4.3 to conclude that $\sigma \vee_{\tilde{\phi}^L(a)} N(\mathbf{A}(\overline{W}))^1$ is quasi-isometrically embedded in $\widetilde{X}^1$, with constants (κ_1', κ_2') depending only on $\kappa_1, \kappa_2, \lambda$.

Verifying that $\sigma \vee_{\tilde{\phi}^L(a)} \mathbf{A}(\overline{W})$ embeds: We will show that there is no path $\tau \subset \sigma$ beginning at $\tilde{\phi}^L(a)$ and joining the endpoints of a path $\alpha_0\beta_0 \cdots \alpha_m$ or $\alpha_0\beta_0 \cdots \alpha_m\beta_m$ in $N(\mathbf{A}(\overline{W}))^1$ with each α_i vertical, and each β_i a path in the carrier of a slope approximation. Each β_i has length L except for the β_m in the path of the second form. Since $\sigma \vee_{\tilde{\phi}^L(a)} N(\mathbf{A}(\overline{W}))^1$ is quasi-isometrically embedded, it suffices to examine the case where $m \leq 1$. The quadrilateral $\alpha_0\beta_0\alpha_1\tau^{-1}$ is approximated by a quasigeodesic quadrilateral $\bar{\alpha}_0\beta_0\bar{\alpha}_1\tau^{-1}$, where each $\bar{\alpha}_i$ is a geodesic of length exceeding $3\delta + 2\lambda$. This quadrilateral is $(2\delta + 2\lambda)$–thin, and β_0, τ fellow travel at distance $2\delta + 2\lambda$ for length at most k. Moreover, since α_0, α_1 are vertical and $|\beta_0| = L$, there is no $(2\delta+2\lambda)$–fellow traveling between α_0 and α_1. Hence, without loss of generality, $\bar{\alpha}_0$ must fellow-travel with β_0 at distance $2\delta + \lambda$ for distance at least $\frac{11L}{24}$, whence α_0 must $(2\delta + 2\lambda + \mu)$-fellow-travel with β_0 for distance at least $\frac{11L}{24\mu_1} - \mu_2$, which contradicts Lemma 2.6 when L is sufficiently large. (Recall that the quasiconvexity constant

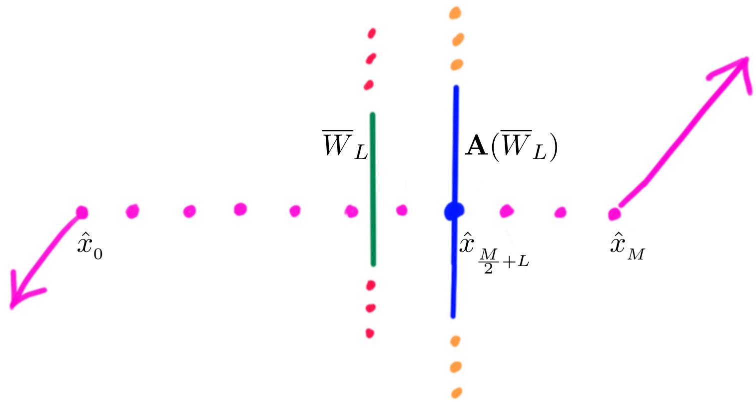

FIGURE 21. Notation in the proof of Proposition 5.19.

μ and the quasi-isometry constants (μ_1, μ_2) of the nucleus approximations are independent of L.) Hence $\sigma \vee_{\tilde{\phi}^L(a)} \mathbf{A}(\overline{W})$ embeds.

Applying Proposition 5.14: Let $\widehat{\sigma_\succ}$ be a lifted augmentation of σ induced by a lift of the forward path joining x to $\phi^M(x)$. Let $\hat{x}_i$ denote the lift of $\tilde{\phi}^i(x)$, so that $\hat{x}_{M/2}$ is a lift of a and $\hat{x}_{M/2+L}$ is a lift of $\tilde{\phi}^L(a)$. The wall $\overline{W}$ is the image of a wall $\overline{W}_L$ such that a nucleus of $\overline{W}_L$ separates $\hat{x}_{M/2}$ from $\hat{x}_{M/2+1}$ and thus $\hat{x}_{M/2+L}$ lies in a nucleus approximation of $\mathbf{A}(\overline{W}_L)$. (Recall that $\mathbf{A}(\overline{W}_L)$ maps isomorphically to $\mathbf{A}(\overline{W})$.) Thus the points $\hat{x}_{M/2+L\pm1}$ lie in distinct halfspaces associated to $\mathbf{A}(\overline{W}_L)$. Indeed, $\hat{x}_{M/2+L-1}$ lies in a downward discrepancy zone and hence in $\overleftarrow{A}$, while $\hat{x}_{M/2+L+1} \in \overrightarrow{A}$. See Figure 21. This verifies Hypothesis (2) of Proposition 5.14.

As in the proof of Proposition 5.18, the fact that $\sigma \vee_{\tilde{\phi}^L(a)} N(\mathbf{A}(\overline{W}))^1$ is quasi-isometrically embedded, together with the fact that $\widetilde{X}_L^\bullet \to \widetilde{X}$ is a quasi-isometry, shows that $\widehat{\sigma_\succ} \vee_{\hat{x}_{M/2+L}} N(\mathbf{A}(\overline{W}_L))^1 \to \widetilde{X}_L^\bullet$ is a quasi-isometric embedding. This verifies Hypothesis (1) of Proposition 5.14.

Let $y \in \widehat{\sigma_\succ} \cap \mathbf{A}(\overline{W}_L)$. Then either y maps to a point of $\sigma \cap \mathbf{A}(\overline{W})$, in which case $y = \tilde{\phi}^L(a)$ since $\sigma \vee_{\tilde{\phi}^L(a)} \mathbf{A}(\overline{W})$ embeds in $\widetilde{X}$, or y is an apex of $\widehat{\sigma_\succ}$. The latter is impossible provided J is sufficiently large compared to κ_1', κ_2'. Indeed, suppose QQ^{-1} is an augmentation beginning on σ and having an apex $p \in \mathbf{A}(\overline{W})$. Let the geodesic $\beta \to N(\mathbf{A}(\overline{W}))^1$ join p to $\tilde{\phi}^L(a)$, let $\tau \to N(\sigma)$ join $\tilde{\phi}^L(a)$ to x, and let P be the subpath of σ joining the initial point of Q to x. Then the concatenation $P\tau^{-1}\beta^{-1}$ is a (κ_1', κ_2')-quasigeodesic. Indeed, $P\tau^{-1}\beta^{-1}$ can be chosen to be a geodesic of the tree $\sigma \vee_{\tilde{\phi}^L(a)} \mathbf{A}(\overline{W})$, which is (κ_1', κ_2')–quasi-isometrically embedded. Moreover, $P\tau^{-1}\beta^{-1}$ contains a subpath of length at least $(J/2-1)L$, namely τ. Hence if $J > 2(\kappa_1'(L+\kappa_2')L^{-1}+1)$, then the offending apex p cannot exist since $|Q| \leq L$. This verifies Hypothesis (3) of Proposition 5.14, and the proof is complete. □

6. Leaf-separation and many effective walls in the irreducible case

In this section, we describe conditions on ϕ ensuring that $\widetilde{X}^1$ satisfies the hypotheses of Proposition 5.10.

6.1. **Leaves.**

Definition 6.1 (Leaf)**.** Let $x, y \in \widetilde{X}$. Then x, y are *equivalent* if there exist forward paths σ_x, σ_y such that $x \in \sigma_x, y \in \sigma_y$ and $\sigma_x \cap \sigma_y \neq \emptyset$. An equivalence class is a *leaf.* We denote the leaf containing x by $\mathcal{L}_x$. The leaf $\mathcal{L}_x$ is *singular* if it contains a 0-cell, and otherwise $\mathcal{L}_x$ is *regular*.

Observe that $\mathcal{L}_x$ is $\tilde{\phi}$-invariant. Moreover, observe that $\mathcal{L}_x$ has a natural directed graph structure: vertices are points of $\mathcal{L}_x \cap \widetilde{X}^1$, and edges are midsegments. From Proposition 6.2.(1) and Proposition 2.5, it follows that this subdivision makes $\mathcal{L}_x$ a directed tree in which each vertex has exactly one outgoing edge and finitely many incoming edges.

Proposition 6.2 (Properties of leaves)**.** *Leaves have the following properties:*

(1) If $\mathcal{L}_x$ is a regular leaf and ϕ is a train track map, then $\mathcal{L}_x$ has a neighborhood homeomorphic to $\mathcal{L}_x \times [-1, 1]$ with $\mathcal{L}_x$ identified with $\mathcal{L}_x \times \{0\}$.

(2) Each level is contained in a unique leaf, and $\mathcal{L}_x$ is an increasing union of levels.

Proof. **Proof of** (1)**:** This uses Lemma 6.3 below. For each vertex $v_{\tilde{e}} = \mathcal{L}_x \cap \tilde{e}$ of $\mathcal{L}_x$, let $U_{\tilde{e}}$ be an open interval in $\tilde{e}$ about $v_{\tilde{e}}$. For each edge $f_{\tilde{c}} = \mathcal{L}_x \cap R_{\tilde{c}}$ of $\mathcal{L}_x$, with vertices at $v_{\tilde{c}}$ and $v_{\tilde{d}}$, let $U(f_{\tilde{c}})$ be the open trapezoid in $R_{\tilde{c}}$ joining $U_{\tilde{c}}$ to $U_{\tilde{d}}$. The desired open neighborhood of $\mathcal{L}_x$ is $\bigcup_{f_{\tilde{c}} \in \mathrm{Edges}(\mathcal{L}_x)} U(f_{\tilde{c}})$, as shown in Figure 22.

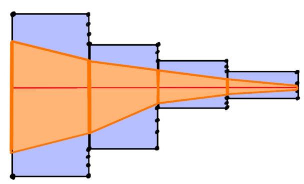

FIGURE 22. A product neighborhood of a regular leaf.

Proof of (2)**:** This follows immediately from the definitions of levels and leaves. □

We denote by $\mathcal{Y}_0$ the set of leaves of $\widetilde{X}$ and define a surjection $\rho_0 : \widetilde{X} \to \mathcal{Y}_0$ by $\rho_0(x) = \mathcal{L}_x$.

Lemma 6.3. *Let $\tilde{e}$ be a vertical edge of $\widetilde{X}$ and let $\mathcal{L}_x$ be a leaf. If ϕ is a train track map, then $|\tilde{e} \cap \mathcal{L}_x| \leq 1$.*

Proof. When ϕ is a train track map, distinct points in each vertical edge e lie on distinct leaves, i.e. the map $\rho_0 : e \to \mathcal{Y}_0$ is injective. Note that in this case, each 2-cell of $\widetilde{X}$ is foliated by a family of distinct fibers of ρ, each of which is a midsegment. □

6.2. **Forward space in the train track case.** Suppose that ϕ is a train track map. We now describe an $\mathbb{R}$-tree $\mathcal{Y}$ whose points are equivalence classes of leaves, and a G-action on $\mathcal{Y}$, and use this to establish that $\widetilde{X}$ is level-separated. This construction mimics the stable tree discussed in [BFH97], although the underlying set is defined differently. The referee explains that it is a special case of a construction in [GJLL98]. Let $\mathcal{E} = \mathbb{R}[\mathrm{Edges}(V)]$ and denote by $\vec{e}_i$ the basis element of $\mathcal{E}$ corresponding to e_i. Let $\mathfrak{M} : \mathcal{E} \to \mathcal{E}$ be the linear map whose matrix with respect to the basis $\{\vec{e}_i\}$ has ij-entry the number of times the path $\phi(e_i)$ traverses e_j, ignoring orientation. Note that this *transition matrix*, which we also denote by $\mathfrak{M}$, is a nonnegative matrix. We further assume that $\mathfrak{M}$ is irreducible.

Let ϖ be the Perron-Frobenius eigenvalue of $\mathfrak{M}$. As shown in [BH92], $\varpi > 1$ since ϕ is irreducible and has infinite order. Let $\mathbf{v}$ be a ϖ-eigenvector, all of whose entries are positive.

For each i, let c_i be the magnitude of the Perron projection of $\vec{e}_i$ onto $\mathbb{R}[\mathbf{v}]$. As is made precise in Definition 5.9, the map ϕ expands edges of V by a factor of ϖ.

We now choose an equivariant weighting of vertical edges in $\widetilde{X}$ by letting $|e_i| = c_i$ for each edge e_i of V, letting each horizontal edge of X have unit weight, and pulling back these weights to $\widetilde{X}$. This determines the metric d on $\widetilde{X}^1$. For each e_i and each $n \in \mathbb{Z}$, we define the *scaled length* of a lift $\tilde{e}_i$ of e_i to $\widetilde{V}_n$ to be $\varpi^{-n}|\tilde{e}_i| = \varpi^{-n}c_i$. Let $\mathsf{d}_{\widetilde{V}_n} : \widetilde{V}_n \times \widetilde{V}_n \to [0,\infty]$ be the resulting path-metric.

Given leaves $\mathcal{L}_x, \mathcal{L}_y$, with $x, y \in \widetilde{V}_k$ for some k, let

$$\mathsf{d}_\infty(\mathcal{L}_x, \mathcal{L}_y) = \lim_{n\to\infty} \mathsf{d}_{\widetilde{V}_n}(\tilde{\phi}^n(x), \tilde{\phi}^n(y)).$$

This limit exists and is finite because $\mathsf{d}_{\widetilde{V}_n}(\tilde{\phi}^n(x), \tilde{\phi}^n(y))$ is non-increasing and bounded. Moreover, $\mathsf{d}_\infty(\mathcal{L}_x, \mathcal{L}_y)$ is well-defined since for other choices $x' \in \mathcal{L}_x \cap \widetilde{V}_{k'}$ and $y' \in \mathcal{L}_y \cap \widetilde{V}_{k'}$, for all but finitely many n, we have $\phi^{n'}(x') = \phi^n(x)$ and $\phi^{n'}(y') = \phi^n(y)$ for some n'.

Lemma 6.4. *Let $\mathcal{Y}$ be the quotient of $\mathcal{Y}_0$ obtained by identifying points $\rho_0(x), \rho_0(y)$ for which $\mathsf{d}_\infty(\rho_0(x), \rho_0(y)) = 0$. Then the induced pseudometric $\mathsf{d}_\infty : \mathcal{Y} \to [0,\infty)$ is a metric. Let $\rho : \widetilde{X} \to \mathcal{Y}$ be the composition $\widetilde{X} \xrightarrow{\rho_0} \mathcal{Y}_0 \to \mathcal{Y}$. Then the restriction of ρ to each vertical edge is an isometric embedding.*

Proof. Note that d_∞ is symmetric and satisfies the triangle inequality. Hence $\mathsf{d}_\infty : \mathcal{Y} \to [0,\infty)$ is a metric. Let e_i be a vertical edge with endpoints x, y. Then the distance in $\widetilde{V}_n$ between the endpoints of $\tilde{\phi}^n(e_i)$ is $\varpi^n|e_i|$, whence $\mathsf{d}_\infty(\rho(x), \rho(y)) = c_i$. Our assumption that $\tilde{\phi}$ has a constant-speed parametrization on each edge implies that the same equality holds for any subinterval of e_i. $\square$

Proposition 6.5. *Suppose that every edge of V is expanding with respect to ϕ. Then:*

1. *The map $\rho : \widetilde{X} \to \mathcal{Y}$ is continuous.*
2. *$(\mathcal{Y}, \mathsf{d}_\infty)$ is a 0-hyperbolic geodesic metric space, i.e. $\mathcal{Y}$ is an $\mathbb{R}$-tree.*
3. *$\mathcal{Y}$ admits a G-action by homeomorphisms with respect to which ρ is G-equivariant.*
4. *The restriction of the G-action on $\mathcal{Y}$ to F is an action by isometries.*
5. *The stabilizer in F of $\rho(\tilde{x})$ is trivial whenever $\tilde{x}$ is a lift to $\widetilde{X}$ of a periodic point in V.*

Proof. **Continuity of ρ:** The restriction of ρ to each vertical edge e is continuous since it is an isometric embedding, and ρ is continuous on each closed 2-cell since ρ is constant on each midsegment and each 2-cell is therefore foliated by fibers of ρ since ϕ is a train track map. Since $\widetilde{X}$ is locally finite, the pasting lemma implies that ρ is continuous on $\widetilde{X}$.

$\mathbb{R}$-tree: Let $x, y \in \widetilde{V}_n$ and let $P \to \widetilde{V}_0$ be a path joining x to y. Then since ρ is continuous, $\rho(P)$ is a path joining $\rho(x)$ to $\rho(y)$, whence $\mathcal{Y}$ is path-connected. Since $\mathcal{Y}$ is a path-connected subspace of an asymptotic cone of the simplicial tree $(\widetilde{V}, \mathsf{d}_0)$, the space $\mathcal{Y}$ is an $\mathbb{R}$-tree [KL95, Prop. 3.6]. (The asymptotic cone in question is built using any non-principal ultrafilter on $\mathbb{N}$, the observation point $(\tilde{v}, \tilde{\phi}(\tilde{v}), \ldots)$, and the scaled metrics $\mathsf{d}_{\widetilde{V}_n}$ on $\widetilde{V}_0$.)

The G-action: For $g \in G$ and $x \in \widetilde{X}$, let $g\rho(x) = \rho(gx)$. This defines an action since G takes leaves in $\widetilde{X}$ to leaves. The action is by homeomorphisms since ρ is continuous and G acts by homeomorphisms on $\widetilde{X}$.

The F-action is isometric: Let $x, y \in \widetilde{X}$. Since F acts by isometries on each $\widetilde{V}_n$, for each $f \in F$, we have

$$\begin{aligned}
\mathsf{d}_\infty(f\rho(x), f\rho(y)) &= \lim_n \mathsf{d}_{\widetilde{V}_n}(\phi^n(fx), \phi^n(fy)) \\
&= \lim_n \mathsf{d}_{\widetilde{V}_n}(\Phi^n(f)\phi^n(x), \Phi^n(f)\phi^n(y)) = \lim_n \mathsf{d}_{\widetilde{V}_n}(\phi^n(x), \phi^n(y)) = \mathsf{d}_\infty(\rho(x), \rho(y)).
\end{aligned}$$

The F-action is free on periodic points: Let $x \in V$ be a periodic point and let $\tilde{x}$ be a lift of x to $\widetilde{X}$. For $f \in F$, by Corollary 6.12, either $\rho(\tilde{x}) \neq f\rho(\tilde{x})$, and we are done, or the forward rays $\sigma_{\tilde{x}}$ and $f\sigma_{\tilde{x}}$ emanating from $\tilde{x}$ and $f\tilde{x}$ respectively lie at finite Hausdorff distance. It follows that the immersed vertical path $\widetilde{P}$ joining $\tilde{x}$ to $f\tilde{x}$ projects to an essential closed path $P \to V$, based at x, such that $\phi^k(P)$ is a periodic Nielsen path for some $k \geq 0$. This contradicts the hyperbolicity of G. □

Remark 6.6. When ϕ is a π-isomorphism, and G is hyperbolic, the action of F on $\mathcal{Y}$ can be shown to be free using Lemma 6.11 and the fact that there are no nontrivial periodic Nielsen paths. We expect that this is true for a general hyperbolic monomorphism, but a free action on the set of periodic points suffices for our purposes.

6.3. **Level-separation in the train track case.** The purpose of this subsection is to prove Lemma 6.15.

Definition 6.7 (Transverse)**.** Let $\mathcal{T}$ be an $\mathbb{R}$-tree. The map $\theta : \mathbf{R} \to \mathcal{T}$ is *transverse* to $y \in \mathcal{T}$ if for each $p \in \theta^{-1}(y)$, there exists $\epsilon > 0$ such that $\theta((p-\epsilon, p))$ and $\theta((p, p+\epsilon))$ lie in distinct components of $\mathcal{T} - \{y\}$. Note that if θ is transverse to y, then $\theta^{-1}(y)$ is a discrete set.

We denote by $\mathbf{R}^+$ a combinatorial sub-ray of the combinatorial line $\mathbf{R}$.

Lemma 6.8. *Let $\mathcal{T}$ be an $\mathbb{R}$-tree. Let $\mathcal{T}_0 \subseteq \mathcal{T}$ have the property that $\mathcal{T} - \{y\}$ has two components for each $y \in \mathcal{T}_0$ and each open arc of $\mathcal{T}$ contains a point of $\mathcal{T}_0$. Let $\theta : \mathbf{R} \to \mathcal{T}$ or $\theta : \mathbf{R}^+ \to \mathcal{T}$ be a continuous map. Suppose θ is transverse to every point in $\mathcal{T}_0$. Moreover, suppose that each edge e of the domain of θ has connected intersection with the preimage of each point in $\mathcal{T}$. Then one of the following holds:*

1. *There exists a nontrivial arc $\alpha \subset \mathcal{T}$ such that $|\theta^{-1}(y)|$ is odd for all $y \in \alpha \cap \mathcal{T}_0$.*
2. *There exists $y \in \mathcal{T}$ with $\theta^{-1}(y)$ having infinitely many components.*
3. *For each $r \geq 0$, there exists $y_r \in \mathcal{T}$ such that $\theta^{-1}(y_r)$ has diameter at least r.*

Proof. For each $p \in \mathbf{R}$, we denote by $\bar{p}$ its image in $\mathcal{T}$ and by $|\theta^{-1}(x)|$ the number of components of the preimage of $x \in \mathcal{T}$ in $\mathbf{R}$.

We now show that either (3) holds or $\operatorname{im}(\theta)$ is locally compact since each point of $\theta(\mathbf{R})$ has a neighborhood intersecting the images of only finitely many edges. We first claim that either (3) holds, or for each edge e of $\mathbf{R}$, there are (uniformly) finitely many edges f such that $\theta(f) \cap \theta(e) \neq \emptyset$. Indeed, if there are arbitrarily many such f, then for each $r \geq 0$, we can choose f such that $\mathsf{d}_{\mathbf{R}}(e, f) > r$ but $\theta(e) \cap \theta(f) \neq \emptyset$, yielding (3). Second, choose a point $p \in \mathcal{T}$. Our first claim shows that either (3) holds or the set $\{e_j\}_{j \in J}$ of edges with $p \in \theta(e_i)$ is finite. Assume the latter. Then for each $i \in J$ we can choose $\epsilon_i > 0$ such that the ϵ_i-neighborhood of p in $\theta(e_i)$ is disjoint from the image of each edge not in $\{e_j\}_{j \in J}$. Let $\epsilon = \min_i \epsilon_i$. Then the ϵ-neighborhood of p in $\operatorname{im}(\theta)$ lies in $\cup_i \theta(e_i)$ and thus and thus intersects the images of only finitely many edges.

There exist sequences $\{a_i\}$ and $\{b_i\}$ in $\mathbf{R} = (-\infty, \infty)$ converging to ∞ and $-\infty$ respectively, whose images are sequences $\{\bar{a}_i\}$ and $\{\bar{b}_i\}$ that converge to points $\bar{a}_\infty$ and $\bar{b}_\infty$ in $\operatorname{im}(\theta) \cup$

$\partial \operatorname{im}(\theta)$. Indeed, since $\operatorname{im}(\theta)$ is a locally compact $\mathbb{R}$-tree, $\operatorname{im}(\theta) \cup \partial \operatorname{im}(\theta)$ is compact by [BH99, Exmp. II.8.11.(5)].

Suppose $\bar{a}_\infty \neq \bar{b}_\infty$. Let α be a nontrivial arc in the geodesic between $\bar{a}_\infty$ and $\bar{b}_\infty$, and note that $\alpha \cap \mathcal{T}_0 \subset \operatorname{im}\theta$. Note that $\theta^{-1}(\bar{c})$ has either odd or infinite cardinality for each $\bar{c} \in \alpha \cap \mathcal{T}_0$, since $\bar{c}$ must separate a_i from b_i for all but finitely many i. Hence either conclusion (1) or (2) holds.

Suppose $\bar{a}_\infty$ and $\bar{b}_\infty$ are equal to the same point $\bar{p}_\infty$. We can assume that $\bar{p}_\infty \in \partial\theta(\mathbf{R})$ since, as above, either (3) holds or each point of $\theta(\mathbf{R})$ has a neighborhood intersecting the images of finitely many edges. Let $\bar{o}$ denote the image of the basepoint o of $\mathbf{R}$. The intersections of the geodesic segments $\bar{o}\bar{a}_i \cap \bar{o}\bar{p}_\infty$ converge to the segment $\bar{o}\bar{p}_\infty$. The same holds for $\bar{o}\bar{b}_i$. We use this to choose a new pair of sequences $\{a_i'\}$ and $\{b_i'\}$ that still converge to $\pm\infty$, and with the additional property that $\bar{a}_i' = \bar{b}_i'$. We do this by choosing the image points far out in $\bar{o}\bar{p}_\infty$. We have thus found arbitrarily distant points in $\mathbf{R}$ with the same images, verifying conclusion (3).

The case of the ray $\mathbf{R}^+$ is similar. □

By Lemma 6.3 and Proposition 6.2, for each regular leaf there is a pair $(\overleftarrow{\mathcal{L}}, \overrightarrow{\mathcal{L}})$ of closed halfspaces in $\widetilde{X}$ such that $\overleftarrow{\mathcal{L}} \cup \overrightarrow{\mathcal{L}} = \widetilde{X}$ and $\overleftarrow{\mathcal{L}} \cap \overrightarrow{\mathcal{L}} = \mathcal{L}$. Points of $\rho(\widetilde{X}^0)$ are *singular* points of $\mathcal{Y}$, and the other points are *regular*. If $\mathcal{L}$ is a regular leaf, then $\mathcal{Y} - \rho(\mathcal{L})$ has two components, namely the interiors of the images of $\overleftarrow{\mathcal{L}}$ and $\overrightarrow{\mathcal{L}}$. Since there are countably many singular points in $\mathcal{Y}$, each open arc in $\mathcal{Y}$ contains a regular point.

Lemma 6.9. *For any geodesic $\gamma : \mathbf{R} \to \widetilde{X}^1$, the map $\theta = \rho \circ \gamma : \mathbf{R} \to \mathcal{Y}$ is transverse to regular points.*

Proof. Let $y \in \mathcal{Y}$ be a regular point, so that each $x \in \rho^{-1}(y)$ lies in the interior of a vertical 1-cell, which in turn embeds in $\mathcal{Y}$ by Lemma 6.4. The image of the vertical 1-cell is separated by $\rho(x) = y$. □

The goal of the rest of this subsection is to prove Corollary 6.10, which depends on Corollary 6.12. We first give a proof of the latter in the case where ϕ is π_1-surjective incorporating the technology of [BFH00], followed by a self-contained proof in the general case.

Corollary 6.10. *Let $\gamma : \mathbf{R} \to \widetilde{X}^1$ be an M-deviating geodesic for some $M \geq 0$. Then there exists a regular leaf $\mathcal{L}$ such that $|\gamma \cap \mathcal{L}|$ is finite and odd.*

Proof. Consider $\theta = \rho \circ \gamma$. By Lemma 6.9, $\rho|_\gamma$ is transverse to regular points. By Lemma 6.8, one of the following holds:

- There exists a regular point $y \in \mathcal{Y}$ such that $\rho^{-1}(y) \cap \gamma$ has finite, odd cardinality.
- For all $r \geq 0$, there exists $y_r \in \mathcal{Y}$ such that $\operatorname{diam}(\rho^{-1}(y_r) \cap \gamma) > r$. (This includes the case in which some point in $\mathcal{Y}$ has infinite preimage.)

In the first case, note that $\rho^{-1}(y)$ is the union of regular leaves, one of which must therefore have odd intersection with γ. We will now show that the second case leads to a contradiction.

In the second case, for each $r \geq 0$, we claim there exists $m \in \mathbb{Z}$ and forward rays σ_1, σ_2 originating at points of γ and traveling through $\widetilde{V}_m$, such that $\rho(\sigma_1) = \rho(\sigma_2)$ and $\mathsf{d}_{\widetilde{V}_m}(\sigma_1 \cap \widetilde{V}_m, \sigma_2 \cap \widetilde{V}_m) > r$. Indeed, let $x_1, x_2 \in \gamma$ be chosen so that $\rho(x_1) = \rho(x_2) = y_r$, and the coordinate projections satisfy $q(x_1) \leq q(x_2) = m$, and $\mathsf{d}_{\widetilde{X}}(x_1, x_2) > r + M + \delta$. For some $k \geq 0$, we have $\tilde{\phi}^k(x_1) = x_1' \in \widetilde{V}_m$. We also have $\rho(x_1') = y_r$. Since γ is M-deviating, considering the δ-thin triangle $x_1 x_2 x_1'$ shows that $\mathsf{d}_{\widetilde{X}}(x_1', x_2) > r$. Hence $\mathsf{d}_{\widetilde{V}_m}(x_1', x_2) > r$.

We now apply Corollary 6.12. The rays σ_1, σ_2 cannot fellowtravel when r is sufficiently large, since the conclusion of a thin quadrilateral argument would then contradict the hypothesis that γ is M-deviating. Hence, by Corollary 6.12, we see that $\rho(\sigma_1) \neq \rho(\sigma_2)$, a contradiction. $\square$

The *tightening* of a path P in a graph is the immersed path that is path-homotopic to P. A *periodic Nielsen path* in V is an essential path P such that the tightening of $\phi^k(P)$ is path-homotopic to P for some $k > 0$. The following is a rephrasing of a special case of [Lev09, Lem. 6.5], which splits into [BFH00, Lem. 4.1.4, Lem. 4.2.6, Lem. 5.5.1].

Lemma 6.11 (Splitting lemma)**.** *Let $\phi : V \to V$ be a π_1-surjective train track map. Let $P \to V$ be a path. Then there exists n_0 such that the tightening of $\phi^{n_0}(P)$ is a concatenation $Q_1 \cdots Q_k$, where each Q_s is of one of the following types:*

(1) a periodic Nielsen path;
(2) an edge of V;
(3) a subinterval of an edge of V, if $s \in \{1, k\}$;

Moreover, for all $n \geq n_0$, the tightening of $\phi^n(P)$ is equal to a concatenation of the tightenings of the paths $\phi^{n-n_0}(Q_s)$.

Corollary 6.12. *Let σ_1, σ_2 be forward rays beginning on $\widetilde{V}_m$. Then either $N(\sigma_1), N(\sigma_2)$ lie at finite Hausdorff distance or $\rho(\sigma_1) \neq \rho(\sigma_2)$.*

Corollary 6.12 means that for each $y \in \mathcal{Y}$, any two forward rays in $\rho^{-1}(y)$ fellowtravel, in the sense that they lie at finite Hausdorff distance.

Proof of Corollary 6.12 when ϕ is π_1-surjective. Let $P \to \widetilde{V}_m$ be a path from σ_1 to σ_2. Lemma 6.11 implies that for some $n \geq 0$, the tightening of $\widetilde{\phi}^n(P)$ splits as the concatenation of periodic Nielsen paths and edges. If $\widetilde{\phi}^n(P)$ is the concatenation of periodic Nielsen paths, then σ_1, σ_2 fellowtravel. Otherwise the splitting contains an edge e and for all $n' \geq n$, we have $\mathsf{d}_{n'}(\sigma_1 \cap \widetilde{V}_{n'}, \sigma_2 \cap \widetilde{V}_{n'}) \geq |e|$ (Lemma 6.4), whence $\rho(\sigma_1) \neq \rho(\sigma_2)$. $\square$

Proof of Corollary 6.12 in the general case. If σ_1, σ_2 do not fellowtravel, then by Lemma 6.13 and Lemma 6.14, the geodesic of $\widetilde{V}_m$ joining the initial points of σ_1, σ_2 contains an open arc $\alpha \subset e$, for some edge e, such that each regular leaf intersecting α separates σ_1, σ_2. For each $n \geq m$, let $a_n = \sigma_1 \cap \widetilde{V}_n$ and $b_n = \sigma_2 \cap \widetilde{V}_n$. Then for each n, the geodesic of $\widetilde{V}_n$ joining a_n, b_n contains $\phi^n(\alpha)$. Regarding e as a copy of $[0,1]$ with weight $|e|$, and $\alpha = (t_1, t_2) \subset [0,1]$, we see that $\mathsf{d}_{\widetilde{V}_n}(a_n, b_n) \geq |e|(t_2 - t_1)$. Hence $\mathsf{d}_\infty(\rho(\sigma_1), \rho(\sigma_2)) > 0$. $\square$

Lemma 6.13. *Let σ_1, σ_2 be forward rays beginning on $\widetilde{V}_m$. Then either $N(\sigma_1)$ and $N(\sigma_2)$ lie at finite Hausdorff distance or there exists a regular leaf separating σ_1 from σ_2.*

Proof. We first claim that is σ_1, σ_2 are not separated by a regular leaf, then $\rho(\sigma_1) = \rho(\sigma_2)$. Indeed, if $\rho(\sigma_1) \neq \rho(\sigma_2)$, then these points are separated by a point $y \in \mathcal{Y}$ whose preimage is the union of regular leaves. Any path joining σ_1, σ_2 must intersect the union of these leaves in an odd-cardinality set, so some regular leaf in the preimage of y separates σ_1, σ_2.

We now verify that $N(\sigma_1)$ and $N(\sigma_2)$ lie at finite Hausdorff distance under the following two assumptions: $\rho(\sigma_1) = \rho(\sigma_2)$ *and* no regular leaf separates σ_1 from σ_2.

Let $z = \rho(\sigma_i) \in \mathcal{Y}$. Let $p \geq m$ be such that there is a vertical geodesic $I_p \to \widetilde{V}_p$ joining $\sigma_1 \cap \widetilde{V}_p, \sigma_2 \cap \widetilde{V}_p$ and with the property that $\rho^{-1}(z) \cap I_p$ has minimal cardinality. For simplicity, having chosen p, we will translate so that $p = 0$.

Having chosen I_0, we now inductively define paths $I_n \to \widetilde{V}_n$ joining σ_1 to σ_2 as follows. For $n \geq 0$, express $I_n = e_1 e_2 \cdots e_k$ as a concatenation of partial edges: e_1, e_k are closed subintervals of edges and the other e_i are entire edges. Let $I_{n+1} \to \widetilde{V}_{n+1}$ be the path $\tilde{\phi}(e_1) \cdots \tilde{\phi}(e_k)$. Let $\bar{I}_n$ be the image of I_n in $\widetilde{X}$ and note that $\bar{I}_n$ is a finite subtree of $\widetilde{V}_n$. Observe that $T = \rho(\bar{I}_0) \subset \mathcal{Y}$ is a finite tree, since it is the union of finitely many closed embedded arcs. Let $\rho_n : I_n \to \mathcal{Y}$ be the composition $I_n \to \widetilde{X} \xrightarrow{\rho} \mathcal{Y}$. Since each $\bar{I}_n \to \bar{I}_{n+1}$ is surjective, $\rho(\bar{I}_n) = T$ for all $n \geq 0$. The maps $e_i \to \tilde{\phi}(e_i)$ induce a map $I_n \to I_{n+1}$ so that the following diagram commutes.

$$\begin{tikzcd}
I_n \arrow[r] \arrow[d] & I_{n+1} \arrow[d] \arrow[dr, "\rho_{n+1}"] & \\
\bar{I}_n \arrow[r, "\tilde{\phi}"] & \bar{I}_{n+1} \arrow[r, "\rho"] & T \subset \mathcal{Y}
\end{tikzcd}$$

Since $\rho(\sigma_1) = \rho(\sigma_2)$, each $\rho_n : I_n \to T$ is a closed path in T beginning and ending at $z \in T$. If $d_{\widetilde{V}_n}(\sigma_1 \cap \bar{I}_n, \sigma_2 \cap \bar{I}_n)$ is uniformly bounded as $n \to \infty$, then σ_1, σ_2 lie at uniformly bounded vertical distance, and so $N(\sigma_1)$ and $N(\sigma_2)$ lie at finite Hausdorff distance.

Since I_n is vertical, $\rho_n^{-1}(z)$ is finite, and $I_n = Q_1 Q_2 \cdots Q_r$, where the interiors of the Q_i are the components of $I_n - \rho_n^{-1}(z)$. Let $\bar{Q}_i$ denote the image of Q_i in $\widetilde{V}_n$. Note that r is independent of n; indeed, this is ensured by the minimality achieved through our choice of p. It follows that no regular leaf intersects $\bar{Q}_i$ and $\bar{Q}_j$ for $i \neq j$, for otherwise we could apply $\tilde{\phi}$ finitely many times and reduce r.

Let a_i and b_i be the endpoints of Q_i, and let $\bar{a}_i, \bar{b}_i$ be their images in $\bar{I}_n$. We will show that there exists M, independent of n, such that $\mathsf{d}_{\widetilde{V}_n}(\bar{a}_i, \bar{b}_i) \leq M$. We conclude that $\mathsf{d}_{\widetilde{V}_n}(N(\sigma_1), N(\sigma_2)) \leq rM$ for all sufficiently large n.

To verify the existence of M, we shall show that there exists a leaf $\mathcal{L}_i$ that intersects the initial and terminal (possibly partial) edges of $\bar{Q}_i$, intersecting these edges in points c_i, d_i respectively. This leaf $\mathcal{L}_i$ must intersect $\bar{I}_0$ in points $\hat{c}_i, \hat{d}_i$ with $\tilde{\phi}^n(\hat{c}_i) = c_i$ and $\tilde{\phi}^n(\hat{d}_i) = d_i$. Hence there are forward paths $\hat{c}_i c_i$ and $\hat{d}_i d_i$ of $N(\mathcal{L}_i)$ whose intersections with $\widetilde{X}^1$ are λ-quasigeodesics lying in forward rays of $N(\mathcal{L}_i)$. The quasigeodesic quadrilateral $\hat{c}_i c_i d_i \hat{d}_i$ shows that $\hat{c}_i c_i$ and $\hat{d}_i d_i$ fellowtravel at distance $M' = M'(\delta, \lambda, |I_0|)$, and hence $\mathsf{d}_{\widetilde{V}_n}(\bar{a}_i, \bar{b}_i) \leq M$, where $M = M' + 2$.

It remains to find the leaf $\mathcal{L}_i$. Note that if $\bar{a}_i, \bar{b}_i$ lie on a common leaf, we are done. We can assume that no regular leaf separates $\bar{a}_i$ from $\bar{b}_i$. Indeed, any such leaf could not separate σ_1, σ_2, since we are assuming that σ_1, σ_2 are not separated (otherwise we would be done). Thus any such leaf must end on $\bar{Q}_j$ for some $i \neq j$, which was ruled out above. Hence each leaf emanating from the image of the initial edge of $\bar{Q}_i$ intersects the images of an even number of edges of $\bar{Q}_i$. Let $\mathcal{L}$ be such a regular leaf, and let $b_i' \in \mathcal{L} \cap \bar{Q}_i$ be a point outside of the image of the initial partial edge of $\bar{Q}_i$. We claim that by choosing $\mathcal{L}$ to intersect the $\bar{Q}_i$ at a point a_i' sufficiently close to $\bar{a}_i$, we can ensure that b_i' lies in the image of the terminal partial edge of $\bar{Q}_i$.

Indeed, choose a sequence $\{a_{ik}'\}_k$ of regular points in the initial edge of $\overline{Q}_i$, with $a_{ik}' \to \bar{a}_i$. For each k, let $\mathcal{L}_{ik}$ be the regular leaf containing a_{ik}'. If $\mathcal{L}_{ik}$ intersects the terminal edge of $\overline{Q}_i$, we are done, so we let b_{ik}' be a point of $\mathcal{L}_{ik} \cap \overline{Q}_{ik}$ that lies in a non-terminal, non-initial edge. By possibly passing to a subsequence, compactness allows us to assume that $\{b_{ik}'\}$ converges to some $b_i' \in \overline{Q}_i$ different from $\bar{b}_i$. Since ρ is continuous and $\rho(b_{ik}') = \rho(a_{ik}') \to z$, we have $\rho(b_i') = z$, contradicting the fact that the interior of $\overline{Q}_i$ contains no point in $\rho^{-1}(z)$. $\square$

Lemma 6.14. *Let σ_1, σ_2 be forward rays beginning on $\widetilde{V}_m$ that do not fellowtravel. Suppose there exists a regular leaf $\mathcal{L}$ separating σ_1, σ_2. Then the geodesic of $\widetilde{V}_m$ joining the initial points of σ_1, σ_2 contains an open arc $\alpha \subset e$, for some edge e, such that each regular leaf intersecting α separates σ_1, σ_2.*

Hence for all $n \geq m$, the geodesic of $\widetilde{V}_n$ joining $\sigma_1 \cap \widetilde{V}_n$ and $\sigma_2 \cap \widetilde{V}_n$ contains $\tilde{\phi}^n(\alpha)$.

Proof. Let $P \to \widetilde{V}_m$ be a vertical geodesic joining σ_1, σ_2. For any $n \geq m$, given a path $U \to \widetilde{V}_n$ joining σ_1, σ_2, a *syllable* of U is a maximal subpath Q that is *legal* in the sense of [BH92]; since ϕ is a train track map, this means that $\tilde{\phi}^k(Q)$ is embedded for all $k \geq 0$. Consider the vertical geodesic T_k joining the endpoints of $\tilde{\phi}^k(P)$ for $k \geq 0$. Each T_k can be expressed as a concatenation of syllables, since ϕ is a train track map, and this decomposition is unique. Choose $k \geq 0$ such that the number of syllables in the decomposition of T_k is equal to the number of syllables in $T_{k'}$ for all $k' \geq k$. Let $T_k = Q_1 \cdots Q_n$ be a decomposition into syllables. Observe that nonconsecutive syllables of T_k cannot intersect a common leaf, for otherwise applying some iterate of $\tilde{\phi}$ would result in a path with fewer syllables.

Since $\mathcal{L}$ intersects each syllable in at most one point, the minimality of T_k guarantees that $|\mathcal{L} \cap T_k| < 3$ and hence, since this cardinality is odd, $|\mathcal{L} \cap T_k| = 1$. Hence there exists a unique Q_i such that $\mathcal{L} \cap T_k$ is contained in $\mathrm{Int}(Q_i)$.

For each $p \in \mathbb{N}$, let $B_i^{\pm}(\frac{1}{p})$ be the two half-open $\frac{1}{p}$-neighborhoods in Q_i bounded at $\mathcal{L} \cap Q_i$. If the lemma does not hold, then for each p there exists a regular leaf $\mathcal{L}^{\pm}(\frac{1}{p})$ intersecting $B_i^{\pm}(\frac{1}{p})$ but failing to separate σ_1 and σ_2. Each $\mathcal{L}^{\pm}(\frac{1}{p})$ has even intersection with T_k and thus also intersects $Q_{i\pm 1}$ in a single point. The sequence $\{\mathcal{L}^{\pm}(\frac{1}{p}) \cap Q_{i\pm 1}\}_p$ has a subsequence that converges to a point $z_{\pm} \in Q_{i\pm 1}$ such that $\rho(z_{\pm}) = \rho(\mathcal{L})$. Observe that no regular leaf separates z_+ (say, the remainder of the argument works analogously for z_-) from $\mathcal{L} \cap Q_i$, since such a separating regular leaf would have to intersect some $\mathcal{L}^+(\frac{1}{p})$, which is impossible since leaves are disjoint. Hence, by Lemma 6.13, the forward rays σ_{z_+} and σ, respectively emanating from z_+ and $\mathcal{L} \cap Q_i$, must fellowtravel.

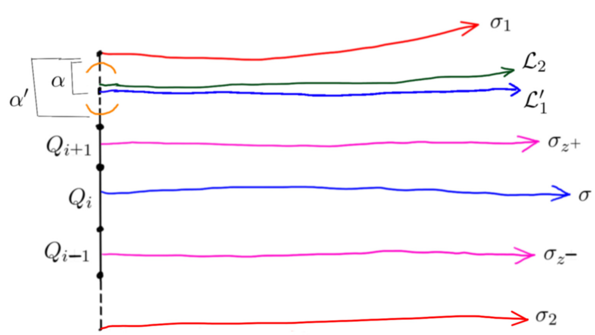

FIGURE 23. The forward rays and leaves in the proof of Lemma 6.14. Only one of σ_{z_+} and σ_{z_-} need exist.

If σ fellowtravels with σ_1 [resp. σ_2] and σ_{z_+} fellowtravels with σ_2 [resp. σ_1], then since σ, σ_{z_+} fellowtravel, we conclude that σ_1, σ_2 fellowtravel, contradicting our hypotheses. See Figure 23. If σ, σ_1 (for example) do not fellowtravel, then σ_1, σ_{z_+} also do not fellowtravel. Lemma 6.13 implies that a regular leaf $\mathcal{L}_1$ separates σ_{z_+}, σ_1. The part of T_k subtended by

σ_{z_+}, σ_1 has strictly fewer syllables than T_k, so by induction, there is an open interval $\alpha' \subset T_k$ with the following properties:

(1) α' is contained in the interior of some edge.
(2) α' intersects a regular leaf $\mathcal{L}_1'$ that separates σ_{z_+} and σ_1.
(3) α' lies on the part of T_k between σ_{z_+} and σ_1.
(4) All regular leaves intersecting α' separate σ_{z_+} from σ_1.

Let $\mathcal{L}_2$ be a regular leaf intersecting α' between $\mathcal{L}_1'$ and σ_1. Then $\mathcal{L}_2$ separates σ_1 from σ_{z_+} by the induction hypothesis, and therefore $\mathcal{L}_2$ separates σ_1 from σ_2. The subinterval of α' between $\mathcal{L}_1'$ and σ_1 (and so containing all such $\mathcal{L}_2$) is the desired interval α. See Figure 23.

In the base case, T_k has a single syllable, and any open subinterval of an edge suffices.

Finally, let $n \geq m$ and let $P_n \to \widetilde{V}_n$ be the geodesic joining σ_1, σ_2. For any $x \in \tilde{\phi}^n(\alpha)$, and any $\epsilon > 0$, there exists a regular point $y \in \tilde{\phi}^n(\alpha)$ at distance less than ϵ from x, since edges are expanding. The regular leaf $\mathcal{L}_y$ separates σ_1, σ_2, so that $y \in P_n$. Since this holds for arbitrarily small ϵ and P_n is closed, $x \in P_n$. □

We have now arrived at the main goal of this subsection:

Lemma 6.15. *Suppose that ϕ is a train track map, that every edge of V is expanding, and that $\mathfrak{M}$ is irreducible. Then $\widetilde{X}$ is level-separated.*

Proof. Let $\gamma : \mathbf{R} \to \widetilde{X}^1$ be an M-deviating geodesic and let $K \geq 0$. By Corollary 6.10, there is a regular leaf $\mathcal{L}$ such that $|\mathcal{L} \cap \gamma|$ is finite and odd. Let $C_0 = \mathcal{L} \cap \gamma$ and choose $y \in \mathcal{L}$ such that the coordinate projections satisfy $q(c) - q(y) > M + K$ for all $c \in C_0$. Then for all sufficiently large n, there is a level $T_o^n(\tilde{\phi}^n(y)) \subset \mathcal{L}$ that contains y as one of its leaves and satisfies $T_o^n(\tilde{\phi}^n(y)) \cap \gamma = C_0$. Hence $\widetilde{X}$ is level-separated. □

6.4. **Proof of Theorem B.**

Theorem 6.16. *Let $\phi : V \to V$ be a train track map of a finite graph V. Suppose that ϕ is π_1-injective and that each edge of V is expanding. Moreover, suppose that the transition matrix $\mathfrak{M}$ of ϕ is irreducible and that the mapping torus X of ϕ has word-hyperbolic fundamental group G. Then G acts freely and cocompactly on a CAT(0) cube complex.*

Proof. Let $\mathcal{Y}$ be the forward space arising from the map $\tilde{\phi} : \widetilde{X} \to \widetilde{X}$. Since ϕ is a train track map, $\widetilde{X}$ has bounded level intersection by Remark 5.8 and is level-separated by Lemma 6.15. By Lemma 6.19, each finite forward path uniformly fellow-travels with a periodic forward path. Hence by Proposition 5.10, it suffices to show that $\widetilde{X}$ has many effective walls by verifying Conditions (1) and (2) of Definition 5.3.

Condition (1)**:** Let $y \in V$ be regular and let $\epsilon > 0$. Let $\mathbf{S}$ be a finite subtree of $\widetilde{V}_0$ such that each contractible subspace of V has one or more lifts to $\mathbf{S}$.

Let $x_0 \in V$ be a periodic point in the interior of the edge e_0 containing y, chosen so that $\mathsf{d}_{e_0}(x_0, y) < \epsilon$. This choice is possible since periodic points are dense in each edge of V, by Lemma 6.19 below.

By Lemma 6.17 below, we have a set $\{x_0, \ldots, x_r\}$ of periodic points in V such that:

(1) Each edge e_i of V contains exactly one point x_i in its interior.
(2) The point x_0 lies in the interior of the edge e_0 containing y and $\mathsf{d}_{e_0}(x_0, y) < \epsilon$.
(3) Let $\tilde{x}_{ip} \neq \tilde{x}_{jq}$ be lifts of x_i, x_j in the closure of a lift of a component of $V - \cup_i \{x_i\}$. Then $\rho(\tilde{x}_{ip}) \neq \rho(\tilde{x}_{jq})$. This holds by construction when $i \neq j$, and holds by Proposition 6.5.(5) when $i = j$ and $p \neq q$.

For each $\tilde{x}_{ip}$, let Λ_{ip} be the bi-infinite periodic forward path containing $\tilde{x}_{ip}$. Let $N(\Lambda_{ip})$ be the 1-skeleton of the smallest subcomplex of $\widetilde{X}$ containing Λ_{ip}, so that $N(\Lambda_{ip})$ is λ-quasiconvex in $\widetilde{X}^1$ by Proposition 2.3.

We now show that for each $R \geq 0$ there exists B_R such that

$$\mathrm{diam}(\mathcal{N}_R(N(\Lambda_{ip})) \cap \mathcal{N}_R(N(\Lambda_{jq}))) \leq B_R$$

whenever $\Lambda_{ip} \neq \Lambda_{jq}$. Since $\Lambda_{ip}, \Lambda_{jq}$ are periodic, they either fellow-travel or have bounded coarse intersection; the following argument precludes the former possibility, whence the claimed B_R exists since there are finitely many pairs $\Lambda_{ip}, \Lambda_{jq}$. Let $\mathsf{d}_{(ip,jq)} = \mathsf{d}_\infty(\rho(\tilde{x}_{ip}), \rho(\tilde{x}_{jq}))$. By definition of d_∞, when $\Lambda_{ip} \neq \Lambda_{jq}$, there exists $n^o_{(ip,jq)} > 0$ such that for all $n \geq n^o_{(ip,jq)}$ we have

$$\mathsf{d}_{\widetilde{V}_n}(\tilde{\phi}^n(\tilde{x}_{ip}), \tilde{\phi}^n(\tilde{x}_{jq})) \geq \frac{\varpi^n \mathsf{d}_{(ip,jq)}}{2}.$$

Let $n_{(ip,jq)} \geq n^o_{(ip,jq)}$ have the property that $\varpi^{n_{(ip,jq)}} \mathsf{d}_{(ip,jq)} \geq 2R$. Let $m = \max\{n_{(ip,jq)}\}$. Then for all $\Lambda_{ip} \neq \Lambda_{jq}$, and all $n \geq m$ we have

$$\mathsf{d}_{\widetilde{V}_n}(\tilde{\phi}^n(\tilde{x}_{ip}), \tilde{\phi}^n(\tilde{x}_{jq})) \geq R.$$

We now construct the uniformly bust-quasiconvex spreading set $\mathbb{W}_y$. Choose J such that $\phi^J(x_s) = x_s$ for all $0 \leq s \leq r$. For each $L \geq 1$ divisible by J, let $\epsilon' = \frac{\epsilon}{\varpi^L}$. By Lemma 3.5, there exist primary busts $d_i \subset e_i$, each disjoint from its ϕ^L-preimage, with $d_i \subset \mathcal{N}_{\epsilon'}(x_i)$. Let $W \to X$ be the immersed wall with tunnel-length L and primary busts d_i. We choose $\mathbb{W}_y$ to be the set of all walls constructed in this way.

It remains to check ladder-overlap. First, $\mathbb{W}_y$ is uniformly bust-quasiconvex since each component of $V - \cup_i \mathrm{Int}(e_i)$ is a finite tree. Let T_i, T_j be distinct tunnels of $\overline{W}$ and suppose that $\mathbf{A}(T_i), \mathbf{A}(T_j)$ intersect a common nucleus approximation $\mathbf{N}$. The forward parts of $\mathbf{A}(T_i), \mathbf{A}(T_j)$ begin at endpoints of primary busts $\tilde{d}_{ip}, \tilde{d}_{jq}$ which are lifts of primary busts d_i, d_j near the periodic points x_i, x_j respectively. Let $\tilde{x}_{ip}, \tilde{x}_{jq}$ be the lifts of x_i, x_j at distance ϵ' from $\tilde{d}_{ip}, \tilde{d}_{jq}$. There are three cases according to whether each of $\mathbf{A}(T_i), \mathbf{A}(T_j)$ is incoming or outgoing at $\mathbf{N}$. In the case where one is incoming and the other outgoing, consideration of the coordinate projection q shows that the diameter of the intersection of the R-neighborhoods of $N(\mathbf{A}(T_i))^1$ and $N(\mathbf{A}(T_j))^1$ is bounded by a function of R.

Suppose that $\mathbf{A}(T_i)$ and $\mathbf{A}(T_j)$ are both outgoing from $\mathbf{N}$. Our choice of ϵ' ensures that $\mathbf{A}(T_i)$ fellow-travels at distance ϵ with the forward path of length L emanating from $\tilde{x}_{ip}$ and similarly for $\mathbf{A}(T_j)$ and $\tilde{x}_{jq}$. (More precisely, each point of $\mathbf{A}(T_i) \cap \widetilde{X}^1$ is at distance at most ϵ from the corresponding point of the forward path emanating from $\tilde{x}_{ip}$.) Hence the coarse intersection of $\mathbf{A}(T_i)$ and $\mathbf{A}(T_j)$ is controlled by the function $R \mapsto B_R$ and the uniform constant ϵ.

Suppose that $\mathbf{A}(T_i)$ and $\mathbf{A}(T_j)$ are both incoming to $\mathbf{N}$. By translating, we can assume that $\mathbf{N} \subset \mathbf{S}$. Because $J \mid L$, we have that $\tilde{\phi}^L(\tilde{x}_{ip})$ and $\tilde{\phi}^L(\tilde{x}_{jq})$ are again lifts of x_i, x_j to $\mathbf{N} \subset \mathbf{S}$ and thus lie on the bi-infinite periodic forward paths $\Lambda_{ip}, \Lambda_{jq}$ that diverge according to the map $R \mapsto B_R$. As before, $\mathbf{A}(T_i)$ and $\mathbf{A}(T_j)$ are (uniformly) coarsely contained in the ϵ-neighborhoods of Λ_{ip} and Λ_{jq}.

Condition (2): Let $a \in \widetilde{V}_0$ and let its image $\bar{a} \in V$ be periodic with period J_a. As before, let $\mathbf{S}$ be a finite subtree of $\widetilde{V}_0$ containing a and having the property that every contractible subspace of V lifts to $\mathbf{S}$ and let $e_0, \ldots, e_r$ be the edges of V, with $\bar{a} \in e_0$. Let $x_{-1} = \bar{a}$. We temporarily subdivide e_0, writing $e_0 = e'_{-1}e'_0$ with $x_{-1} \in e'_{-1}$. We now apply Lemma 6.18 to

V, and then remove the subdivision vertex, yielding periodic points $x_i \in \mathrm{Int}(e_i), 0 \leq i \leq r$ so that: for all $i, j \geq -1$ and all $n \geq 0$, any lifts $\tilde{x}_{ip}, \tilde{x}_{jq}$ of $\phi^n(x_i), \phi^n(x_j)$ to $\mathbf{S}$ satisfy $\rho(\tilde{x}_{ip}) \neq \rho(\tilde{x}_{jq})$. As before, let J be the least common multiple of the periods of the x_i.

Let $L \geq 0$ and $\epsilon > 0$. Applying Lemma 3.5, for each $i \geq 0$ let $d_i \subset \mathrm{Int}(e_i)$ be a primary bust such that $d_i \subset \mathcal{N}_{\frac{\epsilon}{\varpi^L}}(x_i)$ and such that there is an immersed wall $W \to X$ with tunnel length L and primary busts d_i. The collection $\mathbb{W}_a$ of such walls with $J \mid L$ is uniformly bust-quasiconvex since each component of the complement of the primary busts is contractible. Arguing as in the verification of Condition (1), the characteristic property of $\{x_i\}$ ensures that $\mathbb{W}_a$ has uniformly bounded ladder overlap, with bound independent of L. Likewise, there is a uniform bound $k(a)$ on $3\delta + 2\lambda$ fellow-traveling between two forward ladders, one emanating from an endpoint of $\tilde{d}_{ip}$ and one from $a = \tilde{x}_{0q}$, whenever $\tilde{d}_{ip}$ is a lift of some d_i that is joined to a by a path in a knockout of $\overline{W}$. Indeed, in this situation, $\tilde{\phi}^L(a)$ is a lift of $\bar{a}$ to the finite nucleus approximation containing the lift $\tilde{\phi}^L(\tilde{x}_{ip})$ of x_i, whence the forward paths emanating from $\tilde{d}_{ip}$ and a have uniformly bounded coarse intersection. The other case, where a and $\tilde{d}_{ip}$ lie on the same nucleus approximation, is handled as in the analogous case in the verification of Condition (1). □

Lemma 6.17. *Let e_0 be an edge of V and let $y \in \mathrm{Int}(e_0)$. Then there exists a set $\{x_0, \ldots, x_r\}$ of periodic points in V such that:*

(1) Each edge e_i of V contains exactly one point x_i in its interior.
(2) The point x_0 lies in the interior of the edge e_0 containing y and $\mathsf{d}_{e_0}(x_0, y) < \epsilon$.
(3) Let $\tilde{x}_{ip} \neq \tilde{x}_{jq}$ be lifts of x_i, x_j in the closure of a lift of a component of $V - \cup_i \{x_i\}$. Then $\rho(\tilde{x}_{ip}) \neq \rho(\tilde{x}_{jq})$.

Proof. Let $e_1, \ldots, e_r$ be the edges of V, except e_0. Using density of periodic points in e_0 (Lemma 6.19), choose a periodic point $x_0 \in \mathrm{Int}(e_0)$ satisfying assertion (2). Suppose, by induction, that we have chosen periodic points $\{x_i \in \mathrm{Int}(e_i)\}$ for $0 \leq i < s$, for some $s \leq r$, with the property that $\rho(\tilde{x}_i) \neq \rho(\tilde{x}_j)$ for $i \neq j$ and any lifts $\tilde{x}_i, \tilde{x}_j$ of x_i, x_j to $\mathbf{S}$. Let $\tilde{e}_{i1}, \ldots, \tilde{e}_{ip_i}$ be the lifts of e_i to $\mathbf{S}$. Likewise, let $\tilde{x}_{ij}$ be the lift of x_i to $\tilde{e}_{ij}$. Choose $x_s \in \mathrm{Int}(e_s)$ to be a periodic point with the property that no lift of x_s to $\mathbf{S}$ lies in $\cup_{i<s,j} \rho^{-1}(\{\rho(\tilde{x}_{ij})\})$. Iterating this procedure yields the desired set $\{x_0, \ldots, x_r\}$. Indeed, these points are periodic and satisfy assertions (1) and (2) by construction. Denoting by $\tilde{x}_{s1}, \ldots, \tilde{x}_{sp_s}$ the lifts of x_s to $\mathbf{S}$, we see that assertion (3) holds by construction when $i \neq j$, and holds by Proposition 6.5.(5) when $i = j$ and $p \neq q$. □

Lemma 6.18. *Let $x_{-1} \in V$ be a periodic point in an edge e_{-1} and let $e_0, \ldots, e_r$ be a collection of edges in V. Then for $0 \leq i \leq r$, there exist periodic points $x_i \in \mathrm{Int}(e_i)$ such that for all $i, j \geq -1, n \geq 0$ and for all distinct lifts $\tilde{x}_{ip}, \tilde{x}_{jq}$ of $\phi^n(x_i), \phi^n(x_j)$ to $\mathbf{S}$, we have $\rho(\tilde{x}_{ip}) \neq \rho(\tilde{x}_{jq})$.*

Proof. For all i, any two distinct lifts of $\phi^n(x_i)$ to $\mathbf{S}$ have distinct images in $\mathcal{Y}$ by Proposition 6.5.(5). It therefore suffices to verify the claim of the lemma for points $\tilde{x}_{ip}, \tilde{x}_{jq}$ with $i \neq j$.

We argue by induction on r. In the base case where $r = -1$, there is nothing to prove. Supposing that $x_{-1}, \ldots, x_{r-1}$ satisfy the conclusion of the lemma, we will choose x_r. Since ρ is an embedding on each edge and $\mathbf{S}$ is the union of finitely many edges, there exists $K \in \mathbb{N}$ such that for all $y \in \rho(\mathbf{S})$, we have $|\rho^{-1}(y) \cap \mathbf{S}| \leq K$. Let

$$Q = |\{\rho(\tilde{x}_{ip}) : -1 \leq i \leq r-1,\ 1 \leq p \leq p_i\}|,$$

where p_i is the number of lifts of x_i to $\mathbf{S}$.

We claim that there exists $m \in \mathbb{N}$ such that e_r intersects at least $KQ+1$ ϕ-orbits of m-periodic points. It suffices to show that there exists m so that the number of m–periodic points in e_r is at least $(KQ+1)m$. To this end, choose $C > 0$ so that for each edge e_i of V, the immersed path $\phi^k(e_i)$ traverses e_r for some $k \leq C$ (such a C exists by irreducibility). For any $n \in \mathbb{N}$, consider the paths $\phi^n(e_r), \phi^{n+1}(e_r), \ldots, \phi^{n+C}(e_r)$. These paths collectively contain at least $|\phi^n(e_r)| \geq \varpi^n$ occurrences of the edge e_r, since each one contains at least one occurrence of e_r in the image of each edge of $\phi^n(e_r)$. On the other hand, there are $C+1$ such paths. Hence there exists $k \leq C$ so that $\phi^{n+k}(e_r)$ traverses e_r at least $(C+1)^{-1}\varpi^n$ times. As in the proof of Lemma 6.19 below, Brouwer's fixed-point theorem implies that e_r contains at least $(C+1)^{-1}\varpi^n$ points of period dividing $n+k$. Hence there exist arbitrarily large $m \in \mathbb{N}$ so that e_r contains at least $C'\varpi^m$ points of period dividing m, where $C' = \frac{\varpi^{-C}}{2(C+1)}$, and the claim follows by, e.g., choosing m to be a sufficiently large prime.

For each such m-periodic u, a *lifted orbit* of u is the set of all lifts to $\mathbf{S}$ of all points $\phi^k(u)$ with $0 \leq k < m$. Note that if u, u' are m-periodic points with distinct ϕ-orbits, then their lifted orbits are disjoint since their projections to V are distinct ϕ-orbits of the same cardinality and are hence disjoint. By the pigeonhole principle, there exists an m-periodic point $x_r \in e_r$ with the desired property. Indeed, the points $\rho(\tilde{x}_{ip})$ with $i < r$ ruled out at most KQ of the $KQ+1$ lifted orbits. $\square$

Lemma 6.19. *Let ϕ be as in Theorem 6.16. Then for each edge e of V, the set of periodic points of e is dense in e. Consequently, for any finite forward path $\sigma \to \widetilde{X}$, there exists a periodic forward path χ with $\sigma \subset N(\chi)$. If σ is regular, then χ can be chosen so that $N(\sigma) = N(\chi)$.*

Proof. To prove the first assertion, let e be an edge of V and let α be a nontrivial closed subinterval of e. Since ϕ is an exponentially expanding irreducible train track map (see Definition 5.9), for each $L \geq 0$, the path $\phi^L(\alpha)$ is an immersed path in V satisfying $|\phi^L(\alpha)| \geq \varpi^L|\alpha|$, which is unbounded as $L \to \infty$. Hence there exists L so that $\phi^L(\alpha)$ traverses an entire edge of V. Since ϕ is irreducible, there thus exists $L' \geq L$ so that $\phi^{L'}(\alpha)$ traverses e and hence traverses α. It follows from Brouwer's fixed-point theorem that α contains a point x with $\phi^{L'}(x) = x$, verifying the first assertion.

The second assertion follows from the fact that periodic points are dense in V and the fact that distinct forward rays diverge at a rate governed by ϖ. First, if σ starts at a vertex, then observe that since vertices are periodic, σ must actually lie in a periodic forward path. Hence let x be the initial point of σ, suppose that x is not a vertex, and let $n = |\sigma|$. We choose y to be a point at distance at most $\frac{K}{2\varpi^n}$ from x, with the image of y in V periodic regular, where K is the distance from x to a nearest vertex. Then the length-n forward path σ_y fellow-travels with σ at distance $\frac{K}{2}$, and hence the first and last vertical edges in the carriers of σ, χ are equal. $\square$

Remark 6.20. The period of χ is unbounded as the length of σ increases.

We conclude with the following:

Corollary 6.21. *Let $\Phi : F \to F$ be a monomorphism of the finitely generated free group F. Suppose that Φ is irreducible and that the ascending HNN extension $G = F*_{\Phi}$ is word-hyperbolic. Then G acts freely and cocompactly on a CAT(0) cube complex.*

Proof. This follows from the fact that such Φ is represented by a map $\phi : V \to V$ satisfying the hypotheses of Theorem 6.16. Indeed, any irreducible endomorphism has an irreducible train track representative [BH92, Rey10, DV96]. □

References

[Ago12] I. Agol. The virtual Haken conjecture. Preprint, arXiv:1204.2810. With an appendix by Ian Agol, Daniel Groves, and Jason Manning, 2012.

[BF92] M. Bestvina and M. Feighn. A combination theorem for negatively-curved groups. *J. Diff. Geom.*, 35:85–101, 1992.

[BFH97] M. Bestvina, M. Feighn, and M. Handel. Laminations, trees, and irreducible automorphisms of free groups. *Geom. Func. Anal.*, 7(2):215–244, 1997.

[BFH00] Mladen Bestvina, Mark Feighn, and Michael Handel. The Tits alternative for $Out(F_n)$ I: Dynamics of exponentially-growing automorphisms. *Annals of Math.*, 151:517–623, 2000.

[BH92] Mladen Bestvina and Michael Handel. Train tracks and automorphisms of free groups. *Ann. of Math.*, 135:1–51, 1992.

[BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*. Springer-Verlag, Berlin, 1999.

[Bri00] P. Brinkmann. Hyperbolic automorphisms of free groups. *Geom. Func. Anal.*, 10(5):1071–1089, 2000.

[BW13] Nicolas Bergeron and Daniel T. Wise. A boundary criterion for cubulation. *Amer. J. Math.*, 2013.

[CLR94] D. Cooper, D. D. Long, and A. W. Reid. Bundles and finite foliations. *Invent. Math*, 118:255–283, 1994.

[Duf12] Guillaume Dufour. *Cubulations de variétés hyperboliques compactes*. PhD thesis, Université Paris-Sud, 2012.

[DV96] W. Dicks and E. Ventura. *The Group Fixed by a Family of Injective Endomorphisms of a Free Group*. Contemporary mathematics. American Mathematical Society, 1996.

[Ger94] S. M. Gersten. The automorphism group of a free group is not a CAT(0) group. *Proc. Amer. Math. Soc.*, 121:pp. 999–1002, 1994.

[GJLL98] Damien Gaboriau, Andre Jaeger, Gilbert Levitt, and Martin Lustig. An index for counting fixed points of automorphisms of free groups. *Duke Mathematical Journal*, 93(3):425–452, 1998.

[Gro87] M. Gromov. Hyperbolic groups. In *Essays in group theory*, volume 8 of *Math. Sci. Res. Inst. Publ.*, pages 75–263. Springer, New York, 1987.

[HW] G. Christopher Hruska and Daniel T. Wise. Finiteness properties of cubulated groups. *Comp. Math.* pp. 1–58, to appear.

[HW08] Frédéric Haglund and Daniel T. Wise. Special cube complexes. *Geom. Funct. Anal.*, 17(5):1 551–1620, 2008.

[HW12] Tim Hsu and Daniel T. Wise. Cubulating malnormal amalgams. 2012.

[HW14] Mark F Hagen and Daniel T Wise. Cubulating hyperbolic free-by-cyclic groups: the general case. *arXiv preprint arXiv:1406.3292*, 2014.

[Kap00] Ilya Kapovich. Mapping tori of endomorphisms of free groups. *Communications in Algebra*, 28(6):2895–2917, 2000.

[KL95] M. Kapovich and B. Leeb. On asymptotic cones and quasi-isometry classes of fundamental groups of 3-manifolds. *Geom. Func. Anal.*, 5:582–603, 1995.

[Lev09] Gilbert Levitt. Counting growth types of automorphisms of free groups. *Geometric and Functional Analysis*, 19(4):1119–1146, 2009.

[Mit99] Mahan Mitra. On a theorem of Scott and Swarup. *Proc. Amer. Math. Soc.*, 127:1625–1631, 1999.

[PW] Piotr Przytycki and Daniel T. Wise. Mixed 3-manifolds are virtually special. pages 1–24. Available at arXiv:1205.6742.

[Rey10] Patrick Reynolds. Dynamics of irreducible endomorphisms of F_n. 2010. cite arxiv:1008.3659.

[Sag95] Michah Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc. (3)*, 71(3):585–617, 1995.

[SS90] G. P. Scott and G. A. Swarup. Geometric finiteness of certain Kleinian groups. *Proc. Amer. Math. Soc.*, 109(3):765–768, 1990.

[Wis] Daniel T. Wise. Cubular tubular groups. *Trans. Amer. Math. Soc.* To appear.

Dept. of Pure Maths. and Math. Stat., University of Cambridge, Cambridge, UK
E-mail address: markfhagen@gmail.com

Dept. of Math. and Stat., McGill University, Montreal, Quebec, Canada
E-mail address: wise@math.mcgill.ca