

Coarse differentiation of quasi-isometries II: Rigidity for Sol and Lamplighter groups

Alex Eskin, David Fisher and Kevin Whyte

1 Introduction and statements of results

This paper continues the work announced in [EFW1] and begun in [EFW2]. For a more detailed introduction, we refer the reader to those papers. As discussed in those papers, all our theorems stated above are proved using a new technique, which we call *coarse differentiation*. Even though quasi-isometries have no local structure and conventional derivatives do not make sense, we essentially construct a “coarse derivative” that models the large scale behavior of the quasi-isometry. From this point of view, the coarse derivatives of maps studied here are constructed in [EFW2] and this paper consists entirely of a coarse analysis of coarsely differentiable maps.

We now state the main results whose proofs are begun in [EFW2] and finished here. The group $\text{Sol} \cong \mathbb{R} \rtimes \mathbb{R}^2$ with \mathbb{R} acting on \mathbb{R}^2 via the diagonal matrix with entries $e^{z/2}$ and $e^{-z/2}$. As matrices, Sol can be written as :

$$\text{Sol} = \left\{ \left(\begin{array}{ccc} e^{z/2} & x & 0 \\ 0 & 1 & 0 \\ 0 & y & e^{-z/2} \end{array} \right) \middle| (x, y, z) \in \mathbb{R}^3 \right\}$$

The metric $e^{-z}dx^2 + e^zdy^2 + dz^2$ is a left invariant metric on Sol. Any group of the form $\mathbb{Z} \rtimes_T \mathbb{Z}^2$ for $T \in SL(2, \mathbb{Z})$ with $|tr(T)| > 2$ is a cocompact lattice in Sol.

The following theorem proves a conjecture of Farb and Mosher:

Theorem 1.1. *Let Γ be a finitely generated group quasi-isometric to Sol. Then Γ is virtually a lattice in Sol.*

We also prove rigidity results for wreath products $\mathbb{Z} \wr F$ where F is a finite group. The name lamplighter comes from the description $\mathbb{Z} \wr F = F^{\mathbb{Z}} \rtimes \mathbb{Z}$ where the \mathbb{Z} action

is by a shift. The subgroup $F^{\mathbb{Z}}$ is thought of as the states of a line of lamps, each of which has $|F|$ states. The "lamplighter" moves along this line of lamps (the \mathbb{Z} action) and can change the state of the lamp at her current position. The Cayley graphs for the generating sets $F \cup \{\pm 1\}$ depend only on $|F|$, not the structure of F . Furthermore, $\mathbb{Z}\wr F_1$ and $\mathbb{Z}\wr F_2$ are quasi-isometric whenever there is a d so that $|F_1| = d^s$ and $|F_2| = d^t$ for some s, t in \mathbb{Z} . The problem of classifying these groups up to quasi-isometry, and in particular, the question of whether the 2 and 3 state lamplighter groups are quasi-isometric, were well known open problems in the field, see [dlH].

Theorem 1.2. *The lamplighter groups $\mathbb{Z}\wr F$ and $\mathbb{Z}\wr F'$ are quasi-isometric if and only if there exist positive integers d, s, r such that $|F| = d^s$ and $|F'| = d^r$.*

For a rigidity theorem for lamplighter groups, see Theorem 1.3 below.

To state Theorem 1.3 we need to describe a class of graphs. These are the Diestel-Leader graphs, $DL(m, n)$, which can be defined as follows: let T_1 and T_2 be regular trees of valence $m+1$ and $n+1$. Choose orientations on the edges of T_1 and T_2 so each vertex has n (resp. m) edges pointing away from it. This is equivalent to choosing ends on these trees. We can view these orientations as defining height functions f_1 and f_2 on the trees (the Busemann functions for the chosen ends). If one places the point at infinity determining f_1 at the top of the page and the point at infinity determining f_2 at the bottom of the page, then the trees can be drawn as:




Figure 1: The trees for $DL(3, 2)$. Figure borrowed from [PPS].

The graph $DL(m, n)$ is the subset of the product $T_1 \times T_2$ defined by $f_1 + f_2 = 0$. The analogy with the geometry of Sol is made clear in [EFW2, Section 3]. For $n = m$ the Diestel-Leader graphs arise as Cayley graphs of lamplighter groups $\mathbb{Z}\wr F$ for $|F| = n$. This observation was apparently first made by R.Moeller and P.Neumann [MN] and is described explicitly, from two slightly different points of view, in [Wo] and [W]. We prove the following:

Theorem 1.3. *Let Γ be a finitely generated group quasi-isometric to the lamplighter group $\mathbb{Z} \wr F$. Then there exists positive integers d, s, r such that $d^s = |F|^r$ and an isometric, proper, cocompact action of a finite index subgroup of Γ on the Diestel-Leader graph $DL(d, d)$.*

Remark: The theorem can be reinterpreted as saying that any group quasi-isometric to $DL(|F|, |F|)$ is virtually a cocompact lattice in the isometry group of $DL(d, d)$ where d is as above.

Acknowledgements. First author partially supported by NSF grant DMS-0244542. Second author partially supported by NSF grants DMS-0226121 and DMS-0541917. Third author partially supported by NSF grant DMS-0349290 and a Sloan Foundation Fellowship. The second author thanks the University of Chicago for hospitality while this paper was being written. The authors also thank Irine Peng for comments on an early version of the manuscript.

2 Results from [EFW2] and what remains to be done

Remark: All terminology in the following theorems is defined in [EFW2]. Most of it is recalled in §3 below. In particular, whenever we wish to make a statement that refers to either Sol or $DL(m, m)$ we will use the notation $X(m)$ and refer to the space as the *model space*. As in [EFW2], $\text{Sol}(m)$ denotes Sol with the dilated metric

$$ds^2 = dz^2 + e^{-2mz} dx^2 + e^{2mz} dy^2.$$

The main result of this paper is the following. The analogue of this theorem for $X(m, n)$ is proved in [EFW2, Section 5].

Theorem 2.1. *For every $\delta > 0$, $\kappa > 1$ and $C > 0$ there exists a constant $L_0 > 0$ (depending on δ, κ, C) such that the following holds: Suppose $\phi : X(n) \rightarrow X(n')$ is a (κ, C) quasi-isometry. Then for every $L > L_0$ and every box $B(L)$, there exists a subset $U \subset B(L)$ with $|U| \geq (1 - \delta)|B(L)|$ and a height-respecting map $\hat{\phi}(x, y, z) = (\psi(x, y, z), q(z))$ such that*

(i)

$$d(\phi|_U, \hat{\phi}) = O(\delta L).$$

(ii) For z_1, z_2 heights of two points in $B(L)$, we have

$$\frac{1}{2\kappa}|z_1 - z_2| - O(\delta L) < |q(z_1) - q(z_2)| \leq 2\kappa|z_1 - z_2| + O(\delta L). \quad (1)$$

(iii) For all $x \in U$, at least $(1 - \delta)$ fraction of the vertical geodesics passing within $O(1)$ of x are $(\eta, O(\delta L))$ -weakly monotone.

This theorem, combined with results in [EFW2, Section 6] proves that any quasi-isometry $\phi : X(m) \rightarrow X(m')$ is within bounded distance of a height respecting quasi-isometry. This is done in two steps there, the first stated as [EFW2, Theorem 6.1] roughly shows that ϕ respects height difference to sublinear error. Then in [EFW2, Section 6.2] we give an argument that shows this implies ϕ is at bounded distance from height respecting. The deduction of Theorem 1.1 from this fact is already given explicitly in [EFW2, Section 7].

The proof of Theorem 2.1 uses the following consequence of [EFW2, Theorem 4.1]:

Theorem 2.2. *Suppose $\epsilon, \theta > 0$. Let $\phi : \text{Sol} \rightarrow \text{Sol}$ be a (κ, C) quasi-isometry. Then for any L' sufficiently large (depending on κ, C, θ), there exists constants R and L with $C \ll R \ll L \ll L'$ and $e^{\epsilon R} \gg L'$ such that for any box $B(L')$ there exist a collection of disjoint boxes $\{B_i(R)\}_{i \in I}$, a subset I_g of I , and for each $i \in I_g$ a subset $U_i \subset B_i(R)$ with $|U_i| \geq (1 - \theta)|B_i(R)|$ such that the following hold:*

(i)

$$\bigsqcup_{i \in I_g} B_i(R) \subset \phi^{-1}(B(L')) \subset \bigsqcup_{i \in I} B_i(R),$$

(ii)

$$|\bigsqcup_{i \in I_g} U_i| \geq (1 - \theta)|\phi^{-1}(B(L'))| \text{ and } \left| \phi \left(\bigsqcup_{i \in I_g} U_i \right) \right| \geq (1 - \theta)|B(L')|$$

(iii) For each $i \in I_g$ there exists a product map $\hat{\phi}_i : B_i(R) \rightarrow \text{Sol}$ such that

$$d(\phi|_{U_i}, \hat{\phi}_i) = O(\epsilon R).$$

Proof. Choose L large enough that [EFW2, Theorem 4.1] holds with the given ϵ and some $\theta_0 < \theta$ for any box of size L . We cover $\phi^{-1}(B(L'))$ by boxes of size L in the domain. Because ϕ is a quasi-isometry, $\phi^{-1}(B(L'))$ is a Følner set which allows us to

cover $\phi^{-1}(B(L'))$ by $\cup_{k \in K} B_k(L)$ such that the measure of $\cup_{k \in K} B_k(L) - \phi^{-1}(B(L'))$ is small provided $L' \gg L$. We apply [EFW2, Theorem 4.1] to the finite family of boxes $\{B_k(L) | k \in L\}$ and let I_g be the good boxes which we index without reference to k . By choosing θ_0 small enough and using the Følner condition on $\phi^{-1}(B(L'))$, it is easy to see that the conclusions of the theorem are satisfied. \square

Recommendations to the reader: We strongly recommend that the reader study [EFW2] before this paper. In reading this paper, we recommend that the reader first assume that the map ϕ restricted to each $U_i \subset B_i(R)$ for $i \in I_g$ is within $O(\epsilon R)$ of b -standard map, or better yet, the identity. (Replacing a b -standard map with the identity amounts to composing with a quasi-isometry of controlled constants and so has no real effect on our arguments.) This allows the reader to become familiar with the general outline of our arguments without becoming too caught up on technical issues.

The reader familiar with [EFW2] can then read §3 and essentially all of §5, skipping §4 entirely. In first reading §3, the reader might initially read §3.1 through §3.4 and skip §3.5. This last subsection is only required in the case of solvable groups and then only at the very end of §5.4. As remarked there, some of the definitions in §3.3 may also be omitted on first reading.

Remarks on the proof: It is possible to rewrite the arguments here and first prove that ϕ restricted to $U_i \subset B_i(R)$ for $i \in I_g$ is within $O(\epsilon R)$ of a b -standard map. However, the arguments needed to prove this, while not so different in flavor from the arguments in §4, are extremely intricate and technical. The proof given here, while slightly more difficult in some later arguments, is essentially the same proof one would give after proving that fact. See §5 for more discussion.

3 Geometric preliminaries

In this section, we describe some key elements of the spaces we consider. There is some duplication with [EFW2], but the emphasis here is different.

3.1 Boxes, product maps, almost product maps

We recall the notion of a box from [EFW2], first in $\text{Sol}(m)$. Let $B(L, \vec{0}) = [-\frac{e^{2mL}}{2}, \frac{e^{2nL}}{2}] \times [-\frac{e^{2mL}}{2}, \frac{e^{2nL}}{2}] \times [-\frac{L}{2}, \frac{L}{2}]$. In our current setting, $|B(L, \vec{0})| \approx Le^{2L}$ and $\text{Area}(\partial B(L, \vec{0})) \approx e^{2L}$, so $B(L)$ is a Følner set.

To define the analogous object in $DL(m, m)$, we look at the set of points in $DL(m, m)$ we fix a basepoint $(\vec{0})$ and a height function h with $h(\vec{0}) = 0$. Let L be an even integer and let $DL(m, m)_L$ be the $h^{-1}([-\frac{L+1}{2}, \frac{L+1}{2}])$. Then $B(L, \vec{0})$ is the connected component of $\vec{0}$ in $DL(m, m)_L$. We are assuming that the top and bottom of the box are midpoints of edges, to guarantee that they have zero measure.

We call $B(L, \vec{0})$ a box of size L centered at the identity. In Sol, we define the box of size L centered at a point p by $B(L, p) = T_p B(L, \vec{0})$ where T_p is left translation by p . We frequently omit the center of a box in our notation and write $B(L)$. For the case of $DL(m, m)$ it is easiest to define the box $B(L, p)$ directly. That is let $DL(m, m)_{[h(p)-\frac{L+1}{2}, h(p)+\frac{L+1}{2}]} = h^{-1}([h(p)-\frac{L+1}{2}, h(p)+\frac{L+1}{2}])$ and let $B(L, p)$ be the connected component of p in $DL(m, m)_{[h(p)-\frac{L+1}{2}, h(p)+\frac{L+1}{2}]}$. It is easy to see that isometries of $DL(m, m)$ carry boxes to boxes.

For Sol, we write $B(R) = S_X \times S_Y \times S_Z$. We think of S_X as a subset of the lower boundary, S_Y as a subset of the upper boundary, and S_Z as a subset of \mathbb{R} . In the $DL(n, n)$ case, by $S_X \times S_Y \times S_Z$ we mean the set $\{p \in DL(n, n) : h(p) \in S_Z\}$ intersected with the union of all vertical geodesics connecting points of S_X to points of S_Y . We also write $S_Z = [h_{bot}, h_{top}]$. We will use the notation $\partial^+ X$ for the upper boundary and $\partial_- X$ for the lower boundary.

Definition 3.1 (Product Map, Standard Map). A map $\hat{\phi} : \text{Sol} \rightarrow \text{Sol}(n')$ is called a *product map* if it is of the form $(x, y, z) \rightarrow (f(x), g(y), q(z))$ or $(x, y, z) \rightarrow (g(y), f(x), q(z))$, where f, g and q are functions from $\mathbb{R} \rightarrow \mathbb{R}$. A product map $\hat{\phi}$ is called *b-standard* if it is the composition of an isometry with a map of the form $(x, y, z) \rightarrow (f(x), g(y), z)$, where f and g is Bilipshitz with the Bilipshitz constant bounded by b .

The discussion of standard and product maps in the setting of $DL(m, m)$ is slightly more complicated. We let \mathbb{Q}_m be the m -adic rationals. The complement of a point in the boundary at infinity of T_{m+1} is easily seen to be \mathbb{Q}_m . Let x be a point in \mathbb{Q}_m viewed as the lower boundary, and y a point in \mathbb{Q}_m (viewed as the upper boundary). There is a unique vertical geodesic in $DL(m, m)$ connecting x to y . To specify a point in $DL(m, m)$ it suffices to specify x, y and a height z . We will frequently abuse notation by referring to the (x, y, z) coordinate of a point in $DL(m, m)$ even though this representation is highly non-unique.

We need to define product and standard maps as in the case of solvable groups, but there is an additional difficulty introduced by the non-uniqueness of our coordinates. This is that maps of the form $(x, y, z) \rightarrow (f(x), g(y), q(z))$, even when one assumes they are quasi-isometries, are not well-defined, different coordinates for the same points will give rise to different images. We will say a quasi-isometry

ψ is at bounded distance from a map of the form $(x, y, z) \rightarrow (f(x), g(y), q(z))$ if $d(\psi(p), (f(x), g(y), q(z)))$ is uniformly bounded for all points and all choices $p = (x, y, z)$ of coordinates representing each point. It is easy to check that $(x, y, z) \rightarrow (f(x), g(y), q(z))$ is defined up to bounded distance if we assume that the resulting map is a quasi-isometry. The bound depends on κ, C, m, n, m' and n' .

Definition 3.2 (Product Map, Standard Map). A map $\hat{\phi} : \text{DL}(m, m) \rightarrow \text{DL}(m', m')$ is called a *product map* if it is within bounded distance of the form $(x, y, z) \rightarrow (f(x), g(y), q(z))$ or $(x, y, z) \rightarrow (g(y), f(x), q(z))$, where $f : \mathbb{Q}_m \rightarrow \mathbb{Q}_{m'}$ (or $\mathbb{Q}_{n'}$), $g : \mathbb{Q}_n \rightarrow \mathbb{Q}_{n'}$ (or $\mathbb{Q}_{m'}$) and $q : \mathbb{R} \rightarrow \mathbb{R}$. A product map $\hat{\phi}$ is called *b-standard* if it is the composition of an isometry with a map within bounded distance of one of the form $(x, y, z) \rightarrow (f(x), g(y), z)$, where f and g are Bilipshitz with the Bilipshitz constant bounded by b .

Definition 3.3. Given a quasi-isometric embedding $\phi : B(R) \rightarrow X(n')$, we say ϕ is an (α, θ) almost a product map if there exist subsets $U \subset B(R), E_1 \subset S_X$ and $E_2 \subset S_Y$ of relative measure $1 - \theta$ such that $U = \{(x, y, z) : x \in E_1, y \in E_2, z \in S_Z\}$ and all geodesics connecting points in E_1 to points in E_2 have ϵ monotone images under ϕ .

Remark. We think of f and g as defined only on E_i . So by $f(I)$ we mean $f(I \cap E_1)$.

Lemma 3.4. *Given a (α, θ) -almost product map ϕ there exists a subset $U^* \subset U$ with relative measure $1 - 128\theta^{\frac{1}{2}}$ and a (partially defined) product map $\hat{\phi} : U^* \rightarrow X(m')$ such that*

$$d(\phi|_U(p), \hat{\phi}(p)) \leq \alpha R$$

for all $p \in U$.

Proof. This is the content of [EFW2, Lemma 4.11 and Proposition 4.12] □

Remark. With an appropriate choice of constants, the converse of Lemma 3.4 is also true.

Notation. Using Lemma 3.4, we write an (almost) product map $\hat{\phi} : B(R) \subset X(n) \rightarrow X(n')$ as $(x, y, z) \rightarrow (f(x), g(y), q(z))$, so the domain of f is S_X etc. We will always work with (almost) product maps of this kind, the arguments for those of the form $(x, y, z) \rightarrow (f(y), g(x), q(z))$ are almost identical. One can also formally deduce any result we need about almost product maps of the form $(x, y, z) \rightarrow (f(y), g(x), q(z))$ from the analogous fact about those of the form $(x, y, z) \rightarrow (f(x), g(y), q(z))$ by noting that these two forms of almost product map differ by either pre- or post-composition with an isometry.

3.2 Discretizing Sol

In this subsection, we introduce a discrete model for $\text{Sol}(n)$ which has some technical advantages at some points in the argument. We will often make arguments for the discrete model instead of for $\text{Sol}(n)$ itself. The discrete model is quasi-isometric to $\text{Sol}(n)$ and in fact $(1, \rho_1)$ quasi-isometric for a parameter ρ_1 which we will choose so that $C \ll \rho_1 \ll \epsilon R$.

The basic idea is to take a ρ_1 net in $\text{Sol}(n)$ and replace $\text{Sol}(n)$ by a graph on this net. It is possible to do this by taking an arithmetic lattice in Sol , taking a deep enough congruence subgroup, and taking the Cayley graph. More concretely, we write $\text{Sol}(n)$ as $\mathbb{R} \times \mathbb{R}^2$, and consider $\rho_1 \mathbb{Z} \subset \mathbb{R}$ and $\rho_1 \mathbb{Z}^2 \subset \mathbb{R}^2$. Here we view $\mathbb{R}^2 \subset \text{Sol}(n)$ as the plane at height zero. We then form a ρ_1 net in $\text{Sol}(n)$ by taking the union

$$\mathcal{G} = \bigcup_{a \in \rho_1 \mathbb{Z}} a \cdot \rho_1 \mathbb{Z}^2.$$

To make this a graph, we connect by an edge any pair of points in \mathcal{G} whose heights differ by ρ_1 and which are within $10\rho_1$ of one another. We metrize this graph by letting lengths of edges be the distance between the corresponding points in $\text{Sol}(n)$, so all edges have length $O(\rho_1)$.

We can also replace $DL(m, m)$ with a graph whose edges have length $O(\rho_1)$. For this we assume $\rho_1 \in \mathbb{N}$. Consider only vertices in $DL(m, m)$ in $h^{-1}(\rho_1 \mathbb{Z})$. Join two vertices by an edge of length ρ_1 if there is a monotone vertical path between them. The resulting graph is clearly quasi-isometric to $DL(m, m)$ and is in fact $DL(m^{\rho_1}, m^{\rho_1})$ but with the edge length fixed as ρ_1 instead of 1.

We remark here that constants that are said to depend only on K, C and the model geometries often also depend on the discretization scale. This is because the discretization process effectively replaces the model space with a graph.

3.3 Shadows, slabs and coarsenings

Shadows and projections: Let H be a subset of an y -horocycle, and suppose $\rho > 1$. By the ρ -shadow of H , denoted $\text{Sh}(H, \rho)$, we mean the union of the vertical geodesic rays which start within distance ρ of H and go down. If H is a x -horocycle, then we use the same definition except that the geodesic rays are going up.

Given a subset of a y -horocycle H , we let $\pi_-(H) = \partial_- X \cap \text{Sh}(H, \rho_1)$. We define $\pi^+(H)$ for a subset of an x -horocycle similarly. Note that we are suppressing ρ_1 in the notation. In any context where π^+ or π_- are used, ρ_1 will be fixed in advance.

The number $\Delta(H)$. For a horocycle H in a box $B(R)$, let $\Delta(H) = \min(h_{top} - h(H), h(H) - h_{bot})$. Thus, $\Delta(H)$ measures how far is $H \cap B(R)$ from the boundary of $B(R)$.

The branching numbers B_X and B'_X . We define B_X to be the branching constant of X . For solvable Lie groups $B_{X(n)} = n$, for Diestel-Leader graphs, $B_{X(n)} = \log(n)$. We use the shorthand B'_X for $B_{X(n')}$.

Measures on the boundary at infinity. Note that the boundaries $\partial_-(X)$ and $\partial^+(X)$ are homogeneous spaces, and thus have a natural Haar measure. (This measure is Lebesgue measure on \mathbb{R} if $X = \text{Sol}$ and the natural measure on the Cantor set if $X = DL(n, n)$). We normalize the measures by requiring that the shadow of a point at height 0 has measure 1. These measures are all denoted by the symbol $|\cdot|$. Note that for any point $p \in X$,

$$|\pi_-(\{p\})|e^{-B_X h(p)} = 1 \tag{2}$$

The parameter β'' . We choose an arbitrary β'' with $\beta'' \ll 1$, with the understanding that ϵ and θ will be chosen so that $\epsilon \ll \beta''$ and $\theta \ll \beta''$. The parameter β'' will be fixed until §5.5.

Slabs: The objects we refer to as *slabs* will always be subsets of the part of the box $B(R)$ which is at least $4\kappa^2\beta''R$ from the boundary of $B(R)$, will always be defined in reference to a horocycle H in $B(R)$, and are always contained in $Sh(H, \rho)$. We give definitions only for slabs in shadows of y horocycles, those for x horocycles are analogous and can be obtained by applying an appropriate flip. If we choose $h_2 < h_1 < h(H)$, a *slab* in $B(R)$ below H is the subset $Sl_2^1(H)$ which is defined to be the subset of the shadow of H which is between heights h_2 and h_1 .

Recommendation to the reader: The remainder of this subsection might be omitted on first reading.

Coarsening: In order to work with more regular sets, we define an operation to *coarsen* subsets of either boundary.

Let a_1, a_2 be two points in a (log model) hyperbolic plane (which we think of as the xz plane in Sol). Let $h^+(a_1, a_2)$ be the height at which vertical geodesics leaving a_1 and a_2 are one unit apart. This function clearly extends to the lower boundary of the hyperbolic plane. We further extend the function to Sol by letting $h^+(p_1, p_2) = h^+(\pi_{xz}(p_1), \pi_{xz}(p_2))$ where $\pi_{xz}(x, y, z) = (x, z)$. If $I = [a, b]$ is an interval, we write $h^+(I)$ for $h^+(a, b)$. Note that we can define h_- similarly in a yz plane. For $DL(n, n)$ we define $h^+(a_1, a_2)$ as the height in T_n at which vertical geodesics leaving a_1 and a_2 meet. Again h_- is defined similarly.

The operation of coarsening replaces any set F by a set $\mathcal{C}_z(F)$ which is a union of open intervals of a certain size depending on z . For $F \subset \partial_- X$ and $z \in \mathbb{R}$, let $\mathcal{C}_z(F)$ denote the set of $x \in \partial_- X$ such that there exists $x' \in F$ with $h^+(x, x') < z$. Similarly, for $F \subset \partial^+ X$ and $z \in \mathbb{R}$, let $\mathcal{C}_z(F)$ denote the set of $y \in \partial^+(X)$ such that there exists $y' \in F$ with $h^-(y, y') > z$.

Generalized Slabs: Given two sets $E^+ \subset \partial^+ X$ and $E_- \subset \partial_- X$, and two heights $h_2 < h_1$, we define a set

$$S(E_-, E^+, h_2, h_1) = \{(x, y, z) \text{ such that } h_2 < z < h_1 \text{ and } x \in E_-, y \in E^+\}.$$

In words $S(E_-, E^+, h_2, h_1)$ is the set of points on geodesics joining E^+ to E_- with height between h_1 and h_2 . We refer to these sets as *generalized slabs*, though in general there geometry can be very bad, depending on the geometry of E^+ and E_- . Generalized slabs will always be subsets of the part of the box $B(R)$ which is at least $4\kappa^2\beta''R$ from the boundary of $B(R)$, even if this is not explicit in our specification of E^+ and E_- . In particular, slabs as defined above are special cases of generalized slabs, with $Sl_2^1(H) = S(\pi_-(H), S_Y, h_2, h_1)$ where $h_2 < h_1 < h(H)$.

Clearly boxes are very special generalized slabs, and we prefer to work in general with generalized slabs that are unions of boxes. One can obtain a generalized slab that is a union of boxes by coarsening E^+ and E_- . Let h_3 and h_4 be two additional heights, and consider $S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1)$. Observe that as long as $h_3 \leq h_2$ and $h_4 \geq h_1$ we have

$$S(E_-, E^+, h_2, h_1) = S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1).$$

We will need some information concerning the geometry of coarse enough generalized slabs.

Lemma 3.5. *Choose $h_3 \geq h_1$ and $h_4 \leq h_2$. Then any generalized slab of the form $S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1)$ is a union of boxes of size $h_1 - h_2$. In the DL(m, m) case, $S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1)$ is a disjoint union of boxes of size $h_1 - h_2$. In the Sol case, $S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1)$ is not a disjoint union of boxes, but any such set contains a disjoint union of boxes of height $h_1 - h_2$ that contain $\frac{1}{25}$ of the volume of $S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1)$. Furthermore, the number of vertical geodesics in $S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1)$ is comparable to:*

$$\frac{\text{Vol}((\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1))}{h_1 - h_2} e^{B_X(h_1 - h_2)}$$

i.e. it is comparable to the area of the cross-section times $e^{B_X(h_1 - h_2)}$

Proof. That $S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1)$ is a union of boxes is clear from the definition of coarsening. In the DL(m, m) case, the set between h_1 and h_2 is a disjoint union of boxes of size $h_1 - h_2$, so the result follows. For Sol one proves the result by considering the set $W = S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1) \cap h^{-1}(z)$ for any $z \in (h_2, h_1)$. It is clear that W is covered by its intersection with boxes of size $h_1 - h_2$, all of which are rectangles of the same size and shape. Using the Vitali covering lemma, one finds a subset of the boxes which cover a fixed fraction of the measure of W . Since the volume of $S(\mathcal{C}_{h_3}(E_-), \mathcal{C}_{h_4}(E^+), h_2, h_1)$ is the area of the cross section times $h_1 - h_2$, we are done.

The claim concerning numbers of vertical geodesics is obvious for a box. The proof in general can be reduced to that case using the earlier parts of this lemma. \square

3.4 The trapping lemma

In this subsection we state some results relating to areas, lengths and shadows. These are used in the proof of Theorem 5.24. Some similar statements are contained in [EFW2, §5.2].

Lemma 3.6. *Let Q be a subset of an x -horocycle H . Then $\pi_-(Q) = \pi_-(H)$ and*

$$\ell(Q) \approx |\pi_-(H)| |\pi_+(Q)|,$$

where by $\ell(Q)$ we mean the length of the intersection of the 3ρ neighborhood of Q with H , and the implied constants depend on ρ .

Proof. This follows from (2). \square

Lemma 3.7. *Suppose $\gamma \subset B(R)$ is a path. Let L be a euclidean plane intersecting $B(R)$, and suppose $U \subset L \cap B(R)$. Suppose also that any vertical geodesic segment from the bottom of $B(R)$ to the top of $B(R)$ which intersects U also intersects the ρ -neighborhood of γ . Then,*

$$\ell(\gamma) \geq c(\rho) \text{Area}(U)$$

(in the above, $c(\rho)$ is a constant, and both the length and the area are measured using the $X(n)$ metric).

Proof. First note that if L' is another Euclidean plane, and U' is the (vertical) projection of U on L' , then $\text{Area}(U') = \text{Area}(U)$.

Now subdivide γ into k segments of length ρ . Let x_i be the midpoints of such a segment. Let R_i be a rectangle at the same height as i , such that x_i is in the center of R_i , and the sides of the rectangle have length 2ρ . Then the $X(n)$ -area of R_i is independent of i , and the projection of the union of the R_i to L must cover U . Therefore $k > c_2(\rho) \text{Area}(U)$, and hence $\ell(\gamma) > c_1(\rho) \text{Area}(U)$. \square

Lemma 3.8 (Trapping Lemma). *Suppose $\rho_1 \gg C$, $\phi : X(n) \rightarrow X(n)$ is a (κ, C) quasi-isometry, and H is a subset (not necessarily connected) of an x -horocycle in $X(n)$.*

Suppose Q is a subset of a finite union of horocycles in $X(n)$, such that such that the $\kappa\rho_1$ -neighborhood of $\phi(Q)$ intersects every vertical geodesic starting from the ρ_1 -neighborhood of H and going down. Then,

$$\ell(Q) \geq c_1 \ell(H)$$

where $c_1 = c_1(\rho_1)$.

Proof. Discretize H on the scale ρ_1 , and apply Lemma 3.7. \square

Lemma 3.8 is sufficient for applications to $DL(n, n)$. For applications to Sol, we will need a generalization that is stated in the next subsection.

3.5 Tangling and generalized trapping

The following (obvious) result about DL graphs is used implicitly in the proof of Theorem 5.24.

Lemma 3.9. *Suppose $\rho > 1$ and p and q are two points in $DL(n, n)$. Suppose also $p \in \text{Sh}(H, \rho)$, $q \in \text{Sh}(H, \rho)^c$. Then any path connecting p to q passes within ρ of H .*

Proof: The point is simply that if π_T is the projection to the tree T_{n+1} transverse to H , then $\pi_T(\text{Sh}(H, \rho))$ is exactly the set directly below the unique point x which is ρ units above the projection of $\pi_T(H)$. And removing x disconnects T_{n+1} . \square

The lemma above is false for the case of Sol. We will need the following variant: Fix an integer $\rho > 100$ for the remainder of this section.

Definition 3.10 (Tangle). Let H be a horocycles. We say that a path $\bar{\gamma}$ *tangles with H within distance D* if either $\bar{\gamma}$ intersects the ρ neighborhood of H or

$$\tau(\bar{\gamma}, H) = \sum_{j=1}^{\frac{D}{\rho}} \frac{\ell(\bar{\gamma} \cap \{p : ja \leq d(p, H) \leq (j+1)3\rho\})}{\nu(ja)} > 100.$$

Here $\nu(r)$ is the volume of the ball of radius r in the hyperbolic plane.

We say $\bar{\gamma}$ tangles with a finite union of horocycles \mathcal{H} within distance D if

$$\sum_{H \in \mathcal{H}} \tau(\bar{\gamma}, H) > 100$$

where D is implicit in our definition of τ .

We first state an easy lemma to illustrate situations in which paths can be forced to tangle with a horocycle.

Lemma 3.11. *Let ρ be as above and let H be horocycle in Sol. Suppose p and q are two points in Sol such that $p \in \text{Sh}(H, \rho/3)$ and $q \in \text{Sh}(H, \rho)^c$. Then any path from p to q of length less than L tangles with H at distance $\log(L)$.*

Proof: This is an easy hyperbolic geometry argument applied to the projection of the path a hyperbolic plane transverse to H . \square

For our applications, we require a more technical variant of Lemma 3.11. In our arguments, we deal with $\text{Sh}(H, \rho_1)$ where ρ_1 is the discretization scale. For this reason, $\text{Sh}(H, \frac{\rho_1}{3})$ is not a good notion and we need to specify the set we consider differently. Given an horocycle H and constant D' , we say a point p is D' -deep in $\text{Sh}(H, \rho)$ if p is more than D' below H and more than $\frac{D'}{9}$ from the edges of the shadow.

Lemma 3.12. *Let ρ be as above, and choose constants $\rho \ll D_1 \ll D_2$. Let H be a horocycle and suppose p and q are two points in Sol such that p is D_2 -deep in $\text{Sh}(H, \rho)$ and $q \in \text{Sh}(H, \rho)^c$. Then any path from p to q of length less than e^{D_1} tangles with \mathcal{H} within distance D_2 .*

Proof: This is an easy hyperbolic geometry argument applied to the projection of the path a hyperbolic plane transverse to H . \square

For a family \mathcal{F} of vertical geodesic segments, we let $\|\mathcal{F}\|$ denote the area of $\mathcal{F} \cap P$, where P is a Euclidean plane intersecting all the segments in \mathcal{F} . (If there is no such plane we break up \mathcal{F} into disjoint subsets \mathcal{F}_i for which such planes exist, and define $\|\mathcal{F}\| = \sum_i \|\mathcal{F}_i\|$).

Lemma 3.13 (Generalized Trapping Lemma). *Let $\rho \ll D_2$ be constants as above. Suppose \mathcal{F} is a family of vertical geodesic segments, suppose Q is a subset of a finite union \mathcal{H} of horocycles. Suppose also that for each $\gamma \in \mathcal{F}$, γ tangles with \mathcal{H} within distance D_2 and that γ is contained in $N(Q^c, D_2)$. Then, $\ell(Q) \geq \omega \|\mathcal{F}\|$, where ω depends only on κ, C, n and the constants in the definition of tangle.*

Proof. We assume P is a Euclidean plane intersecting all the geodesics in \mathcal{F} , the general case is not much harder. Let $S(r) = \{p : r \leq d(p, Q) \leq r + a\}$. Then, $|S(r)| = c \ell(Q) \nu(r)$, where c depends only on a . Then, we have by [EFW2, Proposition 5.4],

$$\ell(Q) = \frac{|S(r)|}{c\nu(r)} \geq \omega_1 \frac{|(S(r))|}{\nu(r)} \geq \omega_1 \int_{\mathcal{F} \cap P} \frac{\ell(\gamma \cap (S(r)))}{\nu(r)} d\gamma \geq \omega_2 \int_{\mathcal{F} \cap P} \frac{\ell(\gamma) \cap S(r)}{\nu(r)} d\gamma$$

where we have identified the space of vertical geodesics with P and ω_1 and ω_2 depend only on (κ, C, a) . After writing $r = ja$, summing the above equation over j and using the assumption that γ tangles with \mathcal{H} within distance D_2 and is contained in $N(Q^c, D_2)$ for all $\gamma \in \mathcal{F}$, we obtain that $\ell(Q) \geq \omega |\mathcal{F} \cap P|$ as required. \square

4 Improving almost product maps

In this section, we make some arguments that improve the information available concerning $\phi|_{B_i(R)}$ where $i \in I_g$. More or less, by throwing away another set of small measure, we show that ϕ maps many slabs to particular nice generalized slabs. We also show that the map q can be taken to be a linear map.

Recommendation to the reader: The reader may wish to skip this section on first reading and continue reading assuming that $\phi|_{U_i}$ is b -standard or within $O(\epsilon R)$ of a b -standard map. All the results in this section are somewhat technical in nature.

4.1 Bilipschitz in measure bounds

It is clear that the image of a slab under a product map is a generalized slab and that the image of a slab under a b -standard map is a slab. We need to work instead with images of slabs under almost product maps. Given an almost product map $\phi : B(R) \rightarrow \text{Sol}$ one wants to understand the image of $Sl_2^1(H)$. In general, there is not an obvious relation between $\phi(Sl_2^1(H))$ and $S(f(\pi_-(H)), g(S_Y), q(h_2), q(h_1))$. We will show that this is true, at least up to sets of small measure, for appropriately chosen slabs, once we coarsen the image of the slab. To this end we let $h = h(H)$ and fix a height $h_1 < h$ and define:

$$\hat{Sl}_2^1(H) = S(\mathcal{C}_{q(h_1)}(f(\pi_-(H))), \mathcal{C}_{g(h_2)}(g(S_Y)), q(h_2), q(h_1)). \quad (3)$$

Note that $\hat{Sl}_2^1(H)$ is a union of boxes of size $q(h_1) - q(h_2)$. When the choice of H is clear, we suppress reference to H and consider h_1, h_2 and also write h for $h(H)$.

In this section, we prove two lemmas which show that we can restrict attention to $Sl_2^1(H)$ which are almost entirely in U^* and whose (coarsened) image is mostly a collection of boxes contained in (a small neighborhood) of the image of U^* .

Terminology: In order to discuss properties of $\hat{Sl}_2^1(H)$ without fixing either the orientation of H or the almost product map on $B_i(R)$, we introduce some terminology. This terminology is justified by comparison with the case where $\hat{Sl}_2^1(H)$ is a slab. We refer to the direction in z that goes from $q(h_2(H))$ to $q(h_1(H))$ as *towards the horocycle* and the opposite direction as *away from the horocycle*. Similarly, there is a direction, either x or y that one can think of as being *along the horocycle* where the other direction is *transverse to the horocycle*. If H is an x horocycle and our product map is of the form $(x, y, z) \rightarrow (f(x), g(y), q(z))$ then x is along the horocycle and y is transverse to the horocycle.

Let ϕ be an (ϵ, R) almost product map and $\hat{\phi}$ the corresponding (partially defined) product map. The following equation follows from the definitions. It says that the image of the intersection of certain slabs with the good set is essentially contained in a corresponding slab.

$$\phi(Sl_2^1(H) \cap U_*) \subset N_{O(\epsilon R)} \hat{\phi}(Sl_2^1(H) \cap U_*) \subset N_{O(\epsilon R)}(\hat{Sl}_2^1(H)). \quad (4)$$

The following two lemmas yield a strengthening of the equation above. The first lemma provides a lower bound on the measure of $Sl_2^1(H) \cap U_*$ and so on the measure of $N_{O(\epsilon R)}(\hat{Sl}_2^1(H))$ for most choices of H . The second lemma provides an upper bound on $N_{O(\epsilon R)}(\hat{Sl}_2^1(H))$ and even $N_{O(\rho_1)}(\hat{Sl}_2^1(H))$ for a more restricted set of choices of H . To do this, we actually need to modify $\hat{Sl}_2^1(H)$ in a way that we describe in Lemma 4.2.

Given any subset $A \subset B(R)$ and any constant $d < 1$, we denote by A^d the intersection of A with the points in $B(R)$ more than dR of the $\partial B(R)$.

Lemma 4.1. *Given $\beta' \gg \beta \gg \alpha \gg 1$, there exist constants c_1, c_2 depending on ϵ, θ and β' and a subset E_{**} of S_X with $|S_X \setminus E_{**}| \leq c_1(\theta, \epsilon, \beta')|S_X|$ with the following properties. Given a y -horocycle H intersecting $B(R)$ more than $2\kappa\beta'R$ away from $\partial B(R)$ and with $\pi_-(H)$ containing a point of E_{**} and any slab $Sl_2^1(H)$ such that $\beta'R > |h_1(H) - h_2(H)| > \beta R$, $4\beta R > |h(H) - h_1(H)| > 2\alpha R$, we have*

$$|Sl_2^1(H) \cap U_*| \geq (1 - c_2(\theta, \epsilon))|Sl_2^1(H)|. \quad (5)$$

Our current notion of $\hat{Sl}_2^1(H)$ is a bit too coarse. In particular, there can be points in this set that are $O(R)$ away from $\phi(Sl_2^1(H))$. We introduce some notations

needed to describe a subset of $\hat{S}l_2^1(H)$ which can be controlled more easily. Given a set $D \subset B(R)$, we denote by $S_Y \cap D$ the set in S_Y consisting of y coordinates of points in D . We then define

$$\tilde{S}l_2^1(H, D) = S(\mathcal{C}_{q(h_2)}(f(\pi_-(H))), \mathcal{C}_{g(h_2)}(g(S_Y \cap D)), q(h_2), q(h_1)).$$

The fact that we only intersect the y coordinate with D is not an accident, it is due to the fact that we consider sets which are “large” in the y direction and “small” in the x direction.

Lemma 4.2. *Given $\beta'' \gg \beta' \gg \beta \gg \alpha \gg 1$, there exist constants c_3, c_4 depending on ϵ, θ and β' and a subset a subset E_* of S_X with $|S_X \setminus E_*| \leq c_3(\theta, \epsilon)|S_X|$ with the following properties. For any y -horocycle H_0 intersecting $B(R)$ more than $4\kappa^2\beta''R$ away from $\partial B(R)$ with $\pi_-(H_0)$ containing a point of E_* , consider all horocycles H in $S = S(\pi_-(H_0), S_Y, h(H_0), h(H_0) + \beta''R) \cap U_*$ with $\pi_-(H)$ containing a point of E_* and any constants $\beta'R > |h_1(H) - h_2(H)| > \beta R$, $4\beta R > |h(H) - h_1(H)| > \alpha R$ such that the slab $Sl_2^1(H)$ is also contained in S , letting $\tilde{S}l_2^1(H) = \tilde{S}l_2^1(H, \phi(S))$, we have:*

$$|\tilde{S}l_2^1(H) \cap N_{\rho_1}(\phi(U_* \cap Sl_2^1(H)))| \geq (1 - c_4(\theta, \epsilon))|\tilde{S}l_2^1(H)|. \quad (6)$$

For $i = 3, 4$, we have $c_i(\theta, \epsilon) = c_i(\theta, \epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$ and $\theta \rightarrow 0$.

Saying H_0 intersects $B(R)$ more than $4\beta''^2R$ from $\partial B(R)$ is the same as saying H_0 intersects the box $B((1 - 2\beta'')R)$ with the same center as $B(R)$. The point is to stay away from the edge of the box. See the remarks in the definition of slabs and generalized slabs.

Before proving the lemma, we state and prove a corollary concerning measures of cross sections. We note that by the definitions of the measures on the boundary, for a generalized slab $S(E_-, E^+, h_2, h_1)$, and $h_1 < z < h_2$, the area (or equivalently volume) of the $O(1)$ neighborhood of the cross section at height z (i.e. of $S(E_-, E^+, h_2, h_1) \cap h^{-1}(z)$) is $|\mathcal{C}_z(E_+)| |\mathcal{C}_z(E_-)|$.

Corollary 4.3. *Assume H satisfies the hypotheses of Lemmas 4.1 and 4.2. Let w_1, w_2 be such that $2\kappa\alpha R < |h(H) - w_1| < \frac{2}{\kappa}\beta R$ and $2\kappa\beta R < |w_2 - w_1| < \frac{1}{2\kappa}\beta'R$. Then:*

$$|\mathcal{C}_{w_1}(f(\pi_-(H)))| |\mathcal{C}_{w_2}(g(S_Y \cap S))| \geq \omega |\pi_-(H)| |S_Y|, \quad (7)$$

and

$$|\mathcal{C}_{w_1}(f(\pi_-(H)))| |\mathcal{C}_{w_2}(g(S_Y \cap S))| \leq b |\pi_-(H)| |S_Y|, \quad (8)$$

where ω and b depend only on κ and C .

Proof of Corollary. Note that from the structure of U and the fact that ϕ is a quasi-isometry, it follows that for $z_1, z_2 \in S_Z$, we have

$$\frac{1}{2\kappa}|z_1 - z_2| - O(\epsilon R) < |q(z_1) - q(z_2)| \leq 2\kappa|z_1 - z_2| + O(\epsilon R). \quad (9)$$

In particular, q is essentially monotone (up to $O(\epsilon R)$ error).

Given w_1, w_2 as in the Corollary, there exist heights $h_1(H)$ and $h_2(H)$ satisfying the hypotheses of Lemmas 4.1 and 4.2 such that $q(h_1(H)) = w_1, q(h_2(H)) = w_2$. We apply those lemmas to the resulting $Sl_2^1(H)$ and $\tilde{Sl}_2^1(H)$.

Let $\text{Vol}'(X) = |N_{\rho_1}(X)|$ with $\epsilon R \gg \rho_1 \gg C$. Recall that $|h_1 - h_2| > \beta R$ for some $\beta \gg \epsilon$. By Lemma 3.5 and the fact that the measure of the $O(\epsilon R)$ neighborhood of a box of size βR is comparable to the measure of a box of size βR , we have

$$(1 - c) \text{Vol}'(N_{O(\epsilon R)}(\tilde{Sl}_2^1(H))) \leq \text{Vol}'(\tilde{Sl}_2^1(H)), \quad (10)$$

where c is a constant that depends only on $\frac{\epsilon}{\beta}$ and which goes to 0 as ϵ goes to zero.

Note that (4) continues to hold when we replace $\hat{Sl}_2^1(H)$ by $\tilde{Sl}_2^1(H)$. Therefore, by (4) and (6) we have

$$(1 - c_3)(1 - c) \text{Vol}'(\tilde{Sl}_2^1(H)) \leq (1 - c) \text{Vol}'(\phi(Sl_2^1(H) \cap U_*)) \leq \text{Vol}'(\tilde{Sl}_2^1(H)), \quad (11)$$

But by [EFW2, Lemma 5.4] and (5),

$$\omega_1^{-1} \text{Vol}'(Sl_2^1(H)) \leq \text{Vol}'(\phi(Sl_2^1(H) \cap U_*)) \leq \omega_1 \text{Vol}'(Sl_2^1(H)), \quad (12)$$

where ω_1 depends only on (κ, C) . Now (7) and (8) follow from (12), (11) (9) and the fact that the volume of a sufficiently coarsened generalized slab is the area of the cross section times the difference in height. \square

4.2 Proof of Lemmas 4.1 and 4.2

We first prove a preliminary estimate:

Lemma 4.4. *Given p_1, p_2 in U^* , then $h^+(\phi(p_1), \phi(p_2)) = q(h^+(p_1, p_2)) + O(\epsilon R)$.*

Proof. By the definition of U^* we can find \tilde{p}_i in U^* with $\pi_{xz}(\tilde{p}_i) = \pi_{xz}(p_i)$ and vertical geodesic segments $\gamma_i \subset U$ going up from \tilde{p}_i which come within $O(1)$ at $h^+(p_1, p_2)$. Since each γ_i is in U , each $\phi(\gamma_i)$ is within $O(\epsilon R)$ of a vertical geodesic $\tilde{\gamma}_i$ and the

$\tilde{\gamma}_i$ come within $O(\epsilon R)$ of one another only at $h^+(\phi(p_1), \phi(p_2)) + O(\epsilon R)$. But by the definition of the product map, $\tilde{\gamma}_1$ is within $O(\epsilon R)$ of $\tilde{\gamma}_2$ at $q(h^+(p_1, p_2))$. \square

Proof of Lemma 4.1 Let $c_2 = c_2(\epsilon, \theta)$ be a constant to be chosen later. Fix $i < j$. Let $E_1 \subset S_X$ be such that for $x \in E_1$ there exists a horocycle H_x such that $x \in I_x \equiv \pi_-(H_x)$ and (5) fails for some slab $Sl_2^1(H_k)$ as in the statement of the lemma. Note that by assumption $Sl_2^1(H_k) \subset B(R)$. Thus we have a cover of E_1 by the intervals I_x . Then, by the Vitali covering lemma there are intervals $I_k = \pi_-(H_k)$ such that the inequality opposite to (5) holds for H_k , $\sum_k |I_k| \geq (1/5)|E_1|$, and also the I_k are strongly disjoint, i.e. for $j \neq k$, $d(I_j, I_k) \geq (1/2) \max(|I_j|, |I_k|)$. Then the sets $Sl_2^1(H_k)$ are also disjoint. By construction, $|SL_2^1(H_k) \cap U_*^c| \geq c_2 |Sl_2^1(H_k)|$. Summing this over k , we get that

$$|B(R) \cap U_*^c| \geq c_2 \sum_k |Sl_2^1(H_k)| \geq (c_2/2) \sum_k |h_1(H_k) - h_2(H_k)| |I_k| |S_Y|.$$

Since $|B(R) \cap U_*^c| \leq \theta R |S_X| |S_Y|$, we get

$$|E_1| \leq 5 \sum_k |I_k| \leq \frac{10\theta}{\beta' c_2} |S_X|.$$

If $c_1 c_2 \beta' = 20\theta$ this implies that $|E_1| < \frac{\epsilon_1}{2} |S_X|$. So letting $E_{**} = S_X \setminus E_1$, we are done. \square

Proof of Lemma 4.2 We construct E_* as a subset of E_{**} from Lemma 4.1, so any H satisfying the hypotheses of Lemma 4.2 satisfies the conclusions of Lemma 4.1.

We now show that $\tilde{Sl}_2^1(H) \subset \phi(B(R))$. Recall that H_0 is more than $4\kappa^2 \beta'' R$ from the edge of $B(R)$. By definition

$$\tilde{Sl}_2^1(H) = S(\mathcal{C}_{q(h_c)}(f(\pi_-(H))), \mathcal{C}_{g(h_2)}(g(S_Y \cap S)), q(h_2), q(h_1)).$$

Since $S \subset U_*$, for any $y \in S_Y \cap S$ there is a point $p = (x, y, z) \in S$ such that ϕ maps p to within $O(\epsilon R)$ of $(f(x), g(y), q(z))$ with x in $\pi_-(H)$. The point p is at most $\beta'' R$ from H_0 and $\phi(p)$ is within $O(\epsilon R)$ of a vertical geodesic γ which stays within $O(\epsilon R)$ of the image of a vertical geodesic through p . Note that any point q in $S(f(\pi_-(H)), g(S_Y \cap S), q(h_2), q(h_1))$ is on a vertical geodesic γ' which stays within ϵR of the image of a geodesic which passes through S and therefore through H_0 . The point $\phi(p)$ is within $\kappa \beta'' R$ of where the geodesics γ and γ' come within $O(\epsilon R)$ since p is within $\beta'' R$ of the point where the corresponding geodesics in the domain come close. This implies that

q is within $3.1\kappa\beta''R$ of $\phi(p)$. By the definition of coarsening, this implies that any point in $S(\mathcal{C}_{q(h_c)}(f(\pi_-(H))), \mathcal{C}_{g(h_2)}(g(S_Y \cap S)), q(h_2), q(h_1))$ is within $4\kappa\beta''R$ of $\phi(p)$. By our assumptions on S and p , this shows that $\tilde{S}l_2^1(H) \subset \phi(B(R))$.

Let $c_3 = c_3(\epsilon, \theta, \beta')$ be a constant to be chosen later. Let $E_2 \subset S_X \setminus E_1$ be such that for $x \in E_2$ there exists a horocycle H_x such that $x \in I_x \equiv \pi_-(H_x)$ and (6) fails. Thus we have a cover of E_1 by the intervals I_x . Then, by the Vitali covering lemma there are intervals $I_k = \pi_-(H_k)$, such that the inequality opposite to (6) holds for H_k instead of H , $\sum_k |I_k| \geq (1/5)|E_2|$, and also the I_k are strongly disjoint, i.e. for $l \neq k$, $d(I_l, I_k) \geq (1/2) \max(|I_l|, |I_k|)$.

We now claim that

$$\phi(\text{Sh}(H_k, O(1))^c \cap U_*) \cap \tilde{S}l_2^1(H_k) = \emptyset. \quad (13)$$

Indeed suppose $p \in \text{Sh}(H_k, O(1))^c \cap U_*$, and $\phi(p) \in \tilde{S}l_2^1(H_k)$. By the definition of $\hat{S}l_2^1(H_k)$, $\pi_-(\phi(p)) \subset \mathcal{C}_{q(h_c(H))}(f(\pi_-(H_k)))$. Hence, by the definition of coarsening, there exists $p' \in \text{Sh}(H_k, O(1)) \cap U_*$ such that $h^+(\phi(p), \phi(p')) = q(h_c(H)) + O(1)$. Since $p_1 \in \text{Sh}(H_k, O(1))^c$ and $p_2 \in \text{Sh}(H_k, O(1))$, we have $h^+(p_1, p_2) > h(H_k) + O(1)$. This contradicts Lemma 4.4, and thus (13) holds. The same argument shows that the sets $\tilde{S}l_2^1(H_k)$ are disjoint.

Suppose $p \in \tilde{S}l_2^1(H_k)$, $q(h_2(H_k)) + O(\epsilon R) < h(p) < q(h_1(H_k)) - O(\epsilon R)$, and $p \notin N_{O(\epsilon R)}(\phi(Sl_2^1(H_k) \cap B(R)))$. We claim that $p \notin \phi(U_*)$. Indeed, if $p = \phi(p')$ where $p' \in U_*$, then by (13), $p' \notin \text{Sh}(H_k, O(1))^c$. But since $h_2(H_k) < h(p') < h_1(H_k)$, we have $p' \in Sl_2^1(H_k) \cap B(R)$. This is a contradiction, and hence $p \notin \phi(U_*)$. This implies that $\phi(U_* \cap Sl_2^1(H_k)^c) \cap \tilde{S}l_2^1(H_k)$ contributes negligibly to the measure of $\hat{S}l_2^1(H_k)$, i.e. the contribution goes to zero as ϵ goes to zero. So to complete the proof, we need only control $\text{Vol}'(\tilde{S}l_2^1(H_k) \cap \phi(B(R) \cap U_*^c))$.

Thus, since we are assuming the opposite inequality to (6), we have $\text{Vol}'(\tilde{S}l_2^1(H_k) \cap \phi(B(R) \cap U_*^c)) \geq c_3 |\tilde{S}l_2^1(H_k)|$. But then, using the disjointness of the $\tilde{S}l_2^1(H_k)$ we get

$$\begin{aligned} \text{Vol}'(\phi(B(R) \cap U_*^c)) &\geq c_3 \sum_k |\hat{S}l_2^1(H_k)| \geq (c_3)(1-c) \text{Vol}'(\phi(Sl_2^1(H_k) \cap U_*)) \geq \\ &\geq \omega_3 c_3 \sum_k |Sl_2^1(H_k) \cap U_*| \geq \omega_4 c_3 \beta R \sum_k |I_k| |S_Y|. \end{aligned} \quad (14)$$

The first inequality is our assumption, the second uses equation (11). The third is [EFW2, Proposition 5.4] and also uses the fact that each I_k contains a point of $S_X \setminus E_1$ to conclude that $|Sl_2^1(H_k) \cap U_*| \geq (1/2)|Sl_2^1(H_k)|$.

Since by [EFW2, Proposition 5.4], $\text{Vol}'(\phi(B(R) \cap U_*^c)) \leq \omega_5 \theta R |S_X| |S_Y|$, we get

$$|E_2| \leq 5 \sum_k |I_k| \leq \frac{\omega_6 \theta}{c_3 \beta} |S_X|.$$

And so $|E_2| < \frac{c_1}{2}$, provided $c_3 c_1 \beta = 2 \omega_6 \theta$. So after letting $E_* = S_X \setminus E_1 \cup E_2$, the proof is complete. \square

4.3 The map on heights

Suppose $B(R) \subset X(n)$ is a box, and suppose $\phi : B(R) \rightarrow X(n')$ is an (ϵ, θ) almost-product map. Then by definition, there exists a partially defined product map $\hat{\phi} = (f, g, q)$ and a subset $U \subset B(R)$ with $|U| \geq (1 - \theta)|B(R)|$ such that

$$d(\phi|_U, \hat{\phi}) = O(\epsilon R) \tag{15}$$

Proposition 4.5 (Map on heights). *Let $\beta \ll \beta' \ll \beta'' \ll 1$ be as in §4.1. Write $B(R) = S_X \times S_Y \times [h_{\text{bot}}, h_{\text{top}}]$. Suppose $h_{\text{bot}} < z_{\text{bot}} < z_{\text{top}} < h_{\text{top}}$, with $4\beta R \leq |z_{\text{top}} - z_{\text{bot}}| \leq \beta' R$, and also $|h_{\text{top}} - z_{\text{top}}| > 4\kappa^2 \beta'' R$, $|z_{\text{bot}} - h_{\text{bot}}| > 4\kappa^2 \beta'' R$. Then there exists a set $S \subset B(R)$ as in Lemma 4.2 and a function $\epsilon' = \epsilon'(\epsilon, \theta)$ with $\epsilon' \rightarrow 0$ as $\epsilon \rightarrow 0$ and $\theta \rightarrow 0$ such that for all $z \in [z_{\text{bot}}, z_{\text{top}}]$,*

$$q(z) = Az - \frac{1}{B'_X} \log \frac{|\mathcal{C}_{q(z_{\text{bot}})}(g(S_Y \cap S))|}{|S_Y|} + O(\epsilon' R),$$

where $A = B_{X(n)}/B_{X(n')} = B_X/B'_X$ is the ratio of branching constants. In particular, if $n = n'$, $A = 1$.

Remark. In all applications of Proposition 4.5, we change q by $O(\epsilon' R)$ in order to have (4.5) hold with no error term.

Remark. For any n, n' there exists a standard map $\hat{\phi} = (f, g, q) : X(n) \rightarrow X(n')$ with $q(z) = Az$. For solvable groups $\hat{\phi}$ is simply a homothety, for Diestel-Leader graphs it is given by collapsing levels.

The rest of this subsection will consist of the proof of Proposition 4.5. Apply Lemmas 4.1 and 4.2 to get a set $E_* \subset S_X$. Let H, H_0 be y horocycles that satisfy the conditions of Lemmas 4.1 and 4.2 with $h(H_0) > h(H) > z_{\text{top}}$. In particular $\pi_-(H)$ contains a point of E_* . Choose an arbitrary $z \in [z_{\text{bot}}, z_{\text{top}}]$, let $h_1 = z$, $h_2 = z_{\text{bot}}$. For the remainder of this subsection we simplify notation by writing $g(S_Y)$ for $g(S_Y \cap S)$.

Lemma 4.6. *There exists a function $\epsilon' = \epsilon'(\epsilon, \theta)$ with $\epsilon' \rightarrow 0$ as $\epsilon \rightarrow 0$ and $\theta \rightarrow 0$ such that the following holds: Let \mathcal{F} denote the set of vertical geodesic segments in $Sl_2^1(H)$, and let $\hat{\mathcal{F}}$ denote the set of vertical geodesic segments in $\tilde{Sl}_2^1(H)$. Then exists a subset $\mathcal{F}' \subset \mathcal{F}$ with $|\mathcal{F}'| \geq (1/2)|\mathcal{F}|$ and a map $\psi : \mathcal{F}' \rightarrow \hat{\mathcal{F}}$ which is at most $e^{\epsilon'R}$ to one. Also there exists a subset $\hat{\mathcal{F}}' \subset \hat{\mathcal{F}}$ with $|\hat{\mathcal{F}}'| \geq (1/2)|\hat{\mathcal{F}}|$ and a map $\hat{\psi} : \hat{\mathcal{F}}' \rightarrow \mathcal{F}$ with is at most $e^{\epsilon'R}$ to one. Hence,*

$$\log |\mathcal{F}| = \log |\hat{\mathcal{F}}| + O(\epsilon'R). \quad (16)$$

Proof. Let c_2 be as in (5). We let \mathcal{F}' to be the set of vertical geodesics in $Sl_2^1(H)$ more than $O(\epsilon R)$ from the edges and which spend at least $1 - \sqrt{c_2}$ fraction of their length in U_* . Then, by (5), $|\mathcal{F}'| \geq (1/2)|\mathcal{F}|$. Now since ϕ is an almost-product map, for each $\gamma \in \mathcal{F}'$ there exists a geodesic $\hat{\gamma} \in \hat{\mathcal{F}}$ such that $\phi(\gamma \cap U_*)$ is within $O(\epsilon R)$ of $\hat{\gamma}$. We define $\psi(\gamma) = \hat{\gamma}$. The map ψ is at most $e^{O(\epsilon R + \sqrt{c_2}R)}$ to one since two geodesics with the same image must be within ϵR of each other whenever they are in U_* , and by assumption there exist points in U_* on each geodesic within $O(\sqrt{c_2}R)$ of h_{top} and h_{bot} .

The construction of the “inverse” map $\hat{\psi}$ is virtually identical, except that one uses (6) instead of (5) and c_3 instead of c_2 . In the end, we can choose $\epsilon' = O(\epsilon + \sqrt{c_2} + \sqrt{c_3})$. \square

Lemma 4.7. *For all $z \in [z_{bot}, z_{top}]$,*

$$q(z) - q(z_{bot}) = A(z - z_{bot}) + O(\epsilon'R)$$

Proof. We count vertical geodesics using Lemma 4.6. Note that $|\mathcal{F}| \sim |\pi_-(H)||S_Y|e^{B_X(h_1 - h_2)}$, and by Lemma 3.5, $|\hat{\mathcal{F}}|$ is comparable to

$$|\mathcal{C}_{h_1}(f(\pi_-(H)))||\mathcal{C}_{h_2}(g(S_Y))|e^{B_{X'}(q(h_1) - q(h_2))},$$

where as above $h_1 = z$, $h_2 = z_{bot}$. Then, by (16),

$$q(h_1) - q(h_2) = A(h_1 - h_2) + \frac{1}{B'_X} \log \frac{|\mathcal{C}_{q(h_1)}(f(\pi_-(H)))||\mathcal{C}_{q(h_2)}(g(S_Y))|}{|\pi_-(H)||S_Y|} + O(\epsilon'R).$$

Now by Corollary 4.3 the logarithm is bounded between two constants which depend only on κ and C . \square

Proof of Proposition 4.5. Choose $h_1 = (z_{top} + z_{bot})/2$, $h_2 = z_{bot}$. By Lemma 4.1 there exists a horocycle H' with $h(H') = z_{top}$ so that (7) and (8) hold for H' . Then

$$\log \frac{|\mathcal{C}_{q(h_1)}(f(\pi_-(H')))| |\mathcal{C}_{q(z_{bot})}(g(S_Y))|}{|\pi_-(H')| |S_Y|} = O(1). \quad (17)$$

By Lemma 4.4, equation (2) and the fact that we coarsen below the horocycle, we see that,

$$\log \frac{|\mathcal{C}_{q(h_1)}(f(\pi_-(H')))| e^{-B'_X q(h(H'))}}{|\pi_-(H')| e^{-B_X h(H')}} = O(\epsilon R).$$

Since $h(H') = z_{top}$, after rearranging we get,

$$\frac{1}{B'_X} \log \frac{|\mathcal{C}_{q(h_1)}(f(\pi_-(H')))|}{|\pi_-(H')|} = q(z_{top}) - Az_{top} + O(\epsilon R).$$

Substituting into (17) we get,

$$\frac{1}{B'_X} \log \frac{|\mathcal{C}_{q(z_{bot})}(g(S_Y))|}{|S_Y|} = Az_{top} - q(z_{top}) = Az_{bot} - q(z_{bot}), \quad (18)$$

where we have used Lemma 4.7 for the last equality. Now Proposition 4.5 follows from (18) and Lemma 4.7. \square

5 Proof of Theorem 2.1

In this section we prove Theorem 2.1. The basic strategy is to show that for most horocycles H intersecting $\phi^{-1}(B(L'))$, the image $\phi(H)$ is within ϵR of a horocycle, at least for most of its measure. This argument occupies the first four subsections. Subsection §5.5 completes the proof in a manner analogous to [EFW2, Section 5.4].

A key ingredient in our proofs is Lemma 5.19. The reader should think of this “illegal circuit lemma” as a generalization or strengthening of the “quadrilaterals lemma” [EFW2, Lemma 3.1]. The greater generality comes from making weaker assumptions on the paths forming the “legs” of the “quadrilateral”. Lemma 5.19 is used much like [EFW2, Lemma 3.1] to show that points along a horocycle must map by ϕ to points approximately along a horocycle.

Recommendation to the reader: We recommend that the reader read this section first assuming that, for each $i \in I_g$, the map ϕ restricted to U_i in $B_i(R)$ is within $O(\epsilon R)$

of a b -standard map. Under this assumption, the construction of the \hat{S} -graph can be omitted since it suffices to consider only the S -graph. The reader will find that proofs in §5.2 and §5.3 simplify somewhat under this hypothesis, but the main arguments in §5.4 remain essentially the same.

The primary difficulty that occurs here in dropping the assumption that $\phi_i|_{U_i}$ is within $O(\epsilon R)$ of a b -standard map is in guaranteeing that the map preserves the “divergence conditions” on pairs of vertical geodesics required to control paths by the methods of §5.3.

5.1 Constructing the \hat{S} graph and the H -graph

Given a “good enough” horocycle H mostly contained in $\phi^{-1}(B(L'))$, in this section we construct a graph which we use to control $\phi(H)$. To begin, we choose constants and make precise the notion of a “good enough” horocycle.

Choosing Constants: Let $\phi : X(n) \rightarrow X(n')$ be a (κ, C) quasi-isometry. Choose $\rho_1 \gg C$, and discretize on scale ρ_1 as described in subsection 3.2. Let B_X (resp. $B_{X'}$) be the branching constant of the resulting graph and let $B = \max\{B_X, B_{X'}\}$. Let $\epsilon > 0$ and $\theta > 0$ be constants to be specified below, and let L' be sufficiently large so that Theorem 2.2 applies, and fix a box $B(L')$. We call the graph that is the discretization of $B(L')$ the S -graph.

We now apply Theorem 2.2 to $B(L')$. We fix $\epsilon \ll \alpha \ll \beta \ll \beta' \ll \beta''$ and apply the arguments described in Section 4 to each box $B_i(R)$ for $i \in I_g$ as in the conclusion of Theorem 2.2, to obtain sets $(E_*^+)_i \subset \partial^+ X$ and $(E_*^-)_i \subset \partial^- X$. After replacing the set U_i from Theorem 2.2 with a slightly smaller set, we can make sure that for all $(x, y, z) \in U_i$, $x \in (E_*^+)_i$, $y \in (E_*^-)_i$. We still have $|U_i| \geq (1 - \delta_0)|B_i(R)|$, where $\delta_0 \rightarrow 0$ as $\epsilon \rightarrow 0$ and $\theta \rightarrow 0$. As remarked following Proposition 4.5, we further modify q_i so that it satisfies (4.5) with no error term. This makes $\hat{\phi}_i$ within $O(\epsilon' R)$ of ϕ where ϵ' goes to zero as $\epsilon \rightarrow 0$ and $\theta \rightarrow 0$.

We then choose $0 < \eta \ll 1$ such that $\rho_1 \ll 1/\eta$ (We mean that for any function $f(\rho_1)$ and any quantity u which is labeled $O(\eta)$ in the argument, $f(\rho_1)$ is much less than 1.)

We then choose $\rho_2 \gg \rho_1$ so that $f(\rho_1)/B^{\rho_2} \ll \eta$, where $f(\rho_1)$ is any function of ρ_1 which arises during the proof. Now pick ρ_3, ρ_4, ρ_5 so that $\rho_2 \ll \rho_3 \ll \rho_4 \ll \rho_5$.

Choose $0 < \delta_0 \ll 1$, so that $\rho_5 \ll 1/\delta_0$ (The last inequality means that for any function $f(\rho_5)$ and any function $g(\delta_0)$ going to 0 as $\delta_0 \rightarrow 0$ which arise during the argument, $f(\rho_5)g(\delta_0) \ll 1$. We also assume that $1/\eta \ll 1/\delta_0$ (i.e. for any quantity u labeled $O(\eta)$ and any function of $g(\delta_0)$ going to 0 as $\delta_0 \rightarrow 0$ which arises during the

proof, we have $f(\delta_0) \ll u$.

Recap. We have

$$C \ll \rho_1 \ll \rho_2 \ll \rho_3 \ll \rho_4 \ll \rho_5 \ll (1/\delta_0).$$

also

$$\rho_1 \ll 1/\eta \ll 1/\delta_0$$

and

$$\rho_5 \ll \epsilon'R \ll R \ll L'.$$

We do not assume that e.g. $e^{\epsilon'R}\delta_0$ is small.

Note. We assume $e^{\epsilon'R} \gg L'$. Both of these are consequences of the proof of Theorem 2.2. We always assume that any path we consider has length $O(L')$ which is much smaller than $e^{\epsilon'R}$.

The sets U' and U . Let $U_i, i \in I_g$ be as in the second paragraph of this subsection. Let $U' = \bigcup_i U_i$. Then $|U'^c \cap \phi^{-1}(B(L'))| \leq 2\delta_0|\phi^{-1}(B(L'))|$.

Let U'' denote the subset of $\phi^{-1}(B(L'))$ which is distance at most ρ_1 from U' . Then $|U''| \geq |U'| \geq (1 - \delta'_0)|\phi^{-1}(B(L'))|$. Also note that since $\rho_1 \ll \epsilon'R$, for $i \in I_g$, the restriction of ϕ to $U \cap B_i(R)$ is an $(\epsilon'R + \rho_1, \theta)$ -almost product map. We define a set U by $U^c = N_{\rho_5 + \rho_1}(U''^c)$. An elementary covering lemma argument shows that $|U| \geq (1 - \delta''_0)|\phi^{-1}(B(L'))|$ where δ''_0 goes to zero with δ'_0 .

Favorable Horocycles. We define a horocycle H to be *favorable* if H does not stay within $\beta''R$ of the walls of the $B_i(R)$, and also

$$|H \cap U'| \geq (1 - \delta'''_0)|H \cap \phi^{-1}(B(L'))|.$$

We call H *very favorable* if the same holds with U in place of U' .

Remark: If a horocycle is very favorable, any horocycle within ρ_5 of it is favorable.

Lemma 5.1. *There exists $\hat{\theta} > 0$ such that the fraction of $B(L')$ which is contained in the image of a very favorable x -horocycle and a very favorable y -horocycle is at least $(1 - \hat{\theta})$. Here $\hat{\theta}$ is a function of δ_0 and β'' which goes to 0 as $\delta_0 \rightarrow 0$ and $\beta'' \rightarrow 0$.*

Proof. This is immediate from the construction. □

Notation. For most of the argument, we fix a very favorable horocycle H , whose image $\phi(H)$ intersects $B(L')$. For notational simplicity, we assume that H is an y -horocycle. We also fix a favorable horocycle H_0 so that $\rho_5/2 < d(H, H_0) < \rho_5$ and $H \subset \text{Sh}(H_0, \rho_1)$.

The sets $I_g(H)$, \tilde{B} and U_* . Let $I_g(H)$ denote the set of indices $i \in I_g$ such that

$$|H \cap U' \cap B_i(R)| \geq (1 - \delta_0'')|H \cap B_i(R)| > 0. \quad (19)$$

Now let $\tilde{B} = \bigcup_{i \in I_g(H)} B_i(R)$, and let $U_* = U' \cap \tilde{B}$.

Good and bad boxes. We refer to boxes $B_i(R)$ with $i \in I_g(H)$ as “good boxes”, and to boxes $B_i(R)$ intersecting H with $i \in I \setminus I_g(H)$ as “bad boxes”.

Shadows of H and $\phi(H)$: Let $h_1 = h(H) - (\alpha + \beta)R$ and $h_2 = h(H) - (\alpha + \beta + \frac{\beta'}{2})R$. For each $i \in I_g$, we let h_0^i to be specified below be such that $(\alpha + \frac{\beta}{2})R < |h(H) - h_0^i| < (\alpha + \beta)R$. For each $B(R)_i$ intersecting H , we denote $W(H)_i = h^{-1}(h_0) \cap \text{Sh}(H_0, \rho_1)$. For all bad boxes, we fix $h_0^i = h(H) - (\alpha + \beta)R$, for good boxes h_0^i will be fixed during the proof of Lemma 5.2 below. For each $B_i(R)$ intersecting H with $i \in I_g$ we let

$$\hat{W}(H)_i = \{(x, y, z) | x \in \mathcal{C}_{q(h_1)}(f(\pi_-(H_0))), y \in \mathcal{C}_{q(h_2)}(g(S_Y \cap Y)), z = q(h_0)\}.$$

Let $W(H) = \bigcup_{i \in I} W(H)_i$, and $\hat{W}(H) = \bigcup_{i \in I_g} \hat{W}(H)_i$. We define these sets in terms of H_0 not H so as to be able to consider points above H in certain arguments below. Recall q is fixed so that the $O(\epsilon'R)$ term in Proposition 4.5 is 0. We let $R'_i = h(H) - h_0^i$. We frequently suppress reference to i in our notation for R' and h_0 . **FIX H VS H_0 BELOW.**

Shadow vertices. We now define a set of *shadow vertices* in the discretization of $X(n')$. By shifting the discretization, we can assume that $\hat{W}(H)$ contains a ρ_1 net of S -vertices. Every S -vertex in $\hat{W}(H)$ is a *shadow vertex*. If some vertical geodesic going down $\beta'R$ from s contains a point of $\phi(U')$ below h_1 and s is not within $10\kappa\epsilon'R$ of an edge of $\hat{W}(H)$ then we call s a *good shadow vertex*. Any S -vertex in $\hat{W}(H)$ which is not a good shadow vertex is a *bad shadow vertex*. We now add additional shadow vertices, not necessarily in $\hat{W}(H)$. We also make any S vertex in $N_{\rho_1}\phi(U_*^c \cap W(H))$ a *bad shadow vertex*, even if it is a good shadow vertex by our previous definition. The bad shadow vertices in $N_{\rho_1}\phi(U_*^c \cap W(H))$ are not necessarily close to $\hat{W}(H)$, even if they come from good boxes. While these bad shadow vertices are not well controlled, they make up a small proportion of all shadow vertices and so do not interfere with our arguments, see Lemma 5.2 below.

For either good boxes or bad boxes, the number of shadow vertices coming from B_i is proportional to the length of $H \cap B_i$. The proportionality constant depends only on κ, C and the geometry of the model spaces.

Lemma 5.2. *There is a constant $c_4(\delta_0, \epsilon')$ such that, for appropriate choices of h_0^i , only c_4 fraction of all shadow vertices are bad and c_4 goes to zero as δ_0, ϵ' go to zero.*

Proof Bad shadow vertices are defined in two stages. First we have the set S_1 of vertices in $\hat{W}(H)_i$ not within $10\kappa\epsilon'R$ of an edge and not within βR of a point in $\phi(U')$ below h_0 . That this set has small measure in $\hat{W}(H)_i$ follows from two facts. First, the subset within $10\kappa\epsilon'R$ of the boundary has measure going to zero with ϵ' . Second, if S_1 contains θ fraction of the vertices in $\hat{W}(H)_i$, then the set of points on geodesics going down $\beta'R$ from S_1 contains $\frac{\theta}{2}$ fraction of the measure of $\tilde{Sl}_2^1(H_0)$. But $\phi(U')$ contains $1 - c_4$ of the measure in $\tilde{Sl}_2^1(H_0)$ by Lemma 4.1, so this implies that $\theta < 2c_4$. Since c_4 goes to zero with ϵ' and δ_0 , this implies $|S_1|$ goes to zero with them as well.

In the second stage, we enlarge the set of bad vertices in $\hat{W}(H)$ by adding the set $N_{\rho_1}\phi(U'^c \cap W(H))$ to the set of bad vertices. The fraction of shadow vertices coming from bad boxes, being proportional to the fraction of the length of H in bad boxes, clearly goes to zero as δ_0 goes to zero. So it suffices to control the size of $N_{\rho_1}\phi(U'_* \cap W(H)_i)$. To show that $W(H)_i \cap U'_*$ contains a small fraction of the measure of $W(H)_i$, we use the flexibility in our choice of h_0^i . Here we make this flexibility explicit by letting $W(H)_i(h_0^i)$ be the set of possible $W(H)_i$'s, parametrized by choices of h_0^i . Let $\rho(h_0^i)$ be the fraction of $W(H)_i(h_0^i)$ contained in $W(H)_i \cap U'_*$. Consider the slab $Sl_2^1(H_0)$ with $h_1 = (\alpha + \frac{\beta}{2})R$ and $h_2 = h(H) - (\alpha + 2\beta)R$. Since all $W(H)_i(h_0^i)$ are contained in $Sl_2^1(H)$, using Lemma 4.1, we have that

$$\sum_{h_0=(\alpha+\frac{\beta}{2})R}^{(\alpha+\beta)R} \rho(h_0) \leq 2c_3(\epsilon', \theta)$$

which implies that for some h_0^i , we have $\rho(h_0) < 2\sqrt{c_3}$. We fix some h_0^i with this property.

Lastly we need to see that this contribution remains small relative to the number of good shadow vertices coming from $B_i(R)$. To see this, we use Corollary 4.3 which implies that $|W(H)_i| \sim |\hat{W}(H)_i|$ for constants depending only on κ and C . Combined with [EFW2, Proposition 5.4], this implies that the ratio of $|N_{\rho_1}\phi(U'_* \cap W(H)_i)|$ to the number of good vertices in $\hat{W}(H)_i$ goes to zero with ϵ' and δ_0 . □

The \hat{S} -graph. It is convenient to modify the S -graph near the image of H . For $x \in \partial_- X$, $y \in \partial^+ X$ and $t \in [q(h_0), q(h(H_0)) - \frac{\rho_5}{4}]$, let $\gamma_{x,y}(t) = (x, y, t)$, so that $\gamma_{x,y}$ is a vertical geodesic segment of length $q(h(H_0)) - \frac{\rho_5}{4} - q(h_0)$. Let K_i be the union of $\gamma_{x,y}$ where $x, y, q(h_0) \in \hat{W}(H)_i$. We begin by replacing K_i as a subset of the S graph by the disjoint union of the $\gamma_{x,y}$. We then define the \hat{S} graph by defining a new set of vertices and a new incidence relation on K_i . For $1 \leq j \leq q(h(H_0)) - \frac{\rho_5}{4} - q(h_0)/\rho_1$

let $t_j = q_i(h_0) + j\rho_1$. We introduce *pre-vertices* along each $\gamma_{x,y}$ at each t_j . An *irregular \hat{S} -vertex* will be an equivalence class of pre-vertices. Each pre-vertex has coordinates $\{x, y, t_j\}$. At each height level t_j in $X'(n)$, we tile the y -horocycle by disjoint segments T_y of length $10\rho_1$. At each height level $q^{-1}(t_j)$ in $X(n)$ we tile each x horocycle by disjoint segments T_x of length $10\kappa^2\rho$. (These tilings are best thought of as tilings of horocycles in the corresponding trees or hyperbolic planes.) We identify two pre-vertices if:

1. their projections to the yt plane are in the same T_y and
2. the points $(f_i^{-1}(x), q_i^{-1}(t_j))$ and $(f_i^{-1}(x'), q_i^{-1}(t'_j))$ are in the same T_x .
3. $\pi_-(T_x) \cap f((E_* *^-)_i)$ contains at least half the measure in $\pi_-(T_x)$.

Any segment ending at a bad shadow vertex is removed. The \hat{S} -vertices which are S -vertices outside of K_i are called *regular*.

The cloud of an \hat{S} -vertex. Note that for any \hat{S} -vertex v , $h(v)$ and the y -coordinate of v are well defined. For an irregular \hat{S} -vertex the x coordinate is “fuzzy”. More precisely, the *cloud* of an \hat{S} -vertex v is the set of points at height $h(v)$ which are on the vertical segments incident to v . Then for a regular \hat{S} -vertex, the cloud is essentially a point (it has size $O(\rho_1)$), whereas for an irregular \hat{S} -vertex the cloud can have size $D\epsilon'R$ where D is a constant depending only on κ, C and the model geometry.

The set $\hat{\phi}(H')$. Note that if H' is within ρ_4 of H , then the set $\hat{\phi}_i(H')$ consisting of the \hat{S} -vertices v with $q_i(h(H')) = h(v) + O(\rho_1)$ and the x -coordinate of H' is $f_i^{-1}(v) + O(\rho_1)$ is well defined. (The notation is explained by the fact that for any $v \in \hat{\phi}_i(H')$, $\hat{\phi}_i^{-1}(v)$ is within $O(\rho_1)$ of H' . We then define $\hat{\phi}(H')$ to be $\bigcup_{i \in I_g} \hat{\phi}_i(H')$).

Lemma 5.3. *There exist constants M_l and M_u depending only on κ, C such that for any two \hat{S} -vertices v_1 and v_2 in $B(L')$, the ratio of the number of vertical geodesics in $B(L')$ passing through v_1 to the number of vertical geodesics in $B(L')$ passing through v_2 is bounded between M_l and M_u .*

Proof. The proof is mainly a computation of the valence of (i.e. the number of vertical paths incident to) an irregular vertex. We give the proof in the DL case first. In the DL case the valence of a regular vertex is clearly $e^{B'_x L'}$ and we will see that irregular vertices have the same valence. For Sol, the valence of regular vertices can vary by a factor of 2 due to edge effects. This same factor of 2 occurs in the first step of the computation below.

Let h_{top} denote the height of the top of $B(L')$, and $h_{bot} = h_{top} - L'$ denote the height of the bottom of $B(L')$. Suppose v is an irregular vertex in K_i . We can choose a horocycle H' so that $v \in \hat{\phi}_i(H')$, hence $q_i(h(H')) = h(v)$. Note that by definition, $\pi_-(H')$ contains a point in E_{**} . Then the number of paths going up from v to the height h_{top} is $\approx e^{B'_X(h_{top}-h(v))}$. Now the number of paths going down from v to $\hat{W}(H)$ (at height h_0) is

$$\begin{aligned}
&\approx |\mathcal{C}_{q_i(h_0)}(f_i(\pi_-(H')))| e^{-B'_X q_i(h_0)} && \text{by (2)} \\
&\approx |\mathcal{C}_{q_i(h_0)}(f_i(\pi_-(H')))| e^{-B_X h_0} \frac{|\mathcal{C}_{q_i(h_2)}(g_i(S_Y \cap S))|}{|S_Y|} && \text{by Proposition 4.5} \\
&\approx |\pi_-(H')| e^{-B_X h_0} && \text{by Corollary 4.3} \\
&\approx e^{B_X(h(H')-h_0)} && \text{by (2)} \\
&= e^{B'_X(h(v)-q_i(h_0))} && \text{by Proposition 4.5}
\end{aligned}$$

Thus the total number of paths going down from v to h_{bot} is $\approx e^{B'_X(h(v)-h_{bot})}$, and thus the total number of paths incident to v is $\approx e^{B'_X L}$ as required. \square

The H -graph. An irregular \hat{S} -vertex $v \in K_i$ is an H -vertex if and only if $q_i(h(H)) = h(v) + O(\rho_1)$ and the x -coordinate of H is $f_i^{-1}(v) + O(\rho_1)$. These vertices are called “good”. (Note that for any good H -vertex $v \in K_i$, $\hat{\phi}_i^{-1}(v)$ is within $O(\rho_1)$ of H).

We also declare the “bad” H -vertices to be the bad shadow vertices (these are always regular \hat{S} -vertices. The “good” and “bad” vertices thus defined comprise all the vertices of the H -graph. An edge of the H -graph is a vertical path in the \hat{S} -graph which either connects two H -vertices, or connects an H -vertex to the top or bottom of the box $B(L')$. An edge with one endpoint at the top or bottom of the box is called a *leaf edge*.

We will count edges with multiplicity. An edge has multiplicity equal to the number of vertical paths in the \hat{S} -graph which contain it.

Notation. We denote the H -graph by $\mathcal{G}(H)$. Let \mathcal{V} denote the set of vertices of $\mathcal{G}(H)$, and let \mathcal{E} denote the set of edges. Let $\mathcal{V}_1 \subset \mathcal{V}$ denote the set of “good” vertices as defined above. We call an H vertex y oriented (resp. x oriented) if the horocycle segment containing it is a y horocycle (resp. x horocycle). We also refer to an orientation for \hat{S} vertices, which is just the orientation of H vertices in the same box.

Lemma 5.4. *The valence of H vertices is bounded between two constants M_l and M_u depending only on κ, C and the model geometries. Furthermore $|\mathcal{V}_1| \geq (1 - c_5)|\mathcal{V}|$ where $c_5 = c_5(\epsilon', \delta_0)$ goes to zero with $\delta_0 \rightarrow 0$ and $\epsilon' \rightarrow 0$.*

Proof. The first statement of the lemma is immediate from Lemma 5.3. To show the final claim, let F denote the set of vertical paths passing through the good shadow vertices. By definition, every such path is incident to a good H -vertex, and also every vertical path incident to a good H -vertex belongs to F . Thus F is also equal to the set of vertical paths incident to good shadow vertices. Let A denote the set of good shadow vertices. Since the valence of each H -vertex is between M_l and M_u times the valence of each good shadow vertex, we have $M_l|A| \leq |F| \leq M_u|A|$, and $M_l|\mathcal{V}_1| \leq |F| \leq M_u|\mathcal{V}_1|$. Thus, $|\mathcal{V}_1| \geq (M_l/M_u)|A|$. But by Lemma 5.2, $|\mathcal{V} \setminus \mathcal{V}_1| \leq c_4|A|$, where $c_4(\epsilon', \delta_0) \rightarrow 0$ as $\epsilon' \rightarrow 0$ and $\delta_0 \rightarrow 0$. Thus $|\mathcal{V} \setminus \mathcal{V}_1| \leq (M_l/M_u)^2 c_4|\mathcal{V}_1|$. \square

5.2 Averaging over the H -graph

Choose $0 < \theta_3 < \theta_4 \ll 1$. The θ 's will be functions of δ_0 which go to 0 as $\delta_0 \rightarrow 0$.

Definition 5.5 (Good Edges). The following defines sets of “good” edges. See also Definition 5.6.

\mathcal{E}_1 : Either connects two vertices in \mathcal{V}_1 or is a leaf edge based on a vertex of \mathcal{V}_1 .

\mathcal{E}_3 : An \mathcal{E}_1 edge e such that for all \hat{S} -vertices $x \in e$, $1 - \theta_3$ fraction of the edges (forward) branching at x are in \mathcal{E}_1 . (note that x is not supposed to be a vertex of the H -graph).

\mathcal{E}_4 : An \mathcal{E}_3 edge such that for any \hat{S} -vertex $x \in e$, $1 - \theta_4$ fraction of the edges reverse branching from x are in \mathcal{E}_1 .

Remark: There \mathcal{E}_2 edges, they will be defined below in §5.3.

Choose $1 \gg \nu_3 > \nu_2 > 0$. The ν 's will be functions of δ_0 which tend to 0 as $\delta_0 \rightarrow 0$.

Definition 5.6 (Good Vertices). The following defines sets of “good” vertices. See also Definition 5.5.

\mathcal{V}_1 : The set of “good” vertices as defined in the previous section.

\mathcal{V}_2 : In \mathcal{V}_1 and $1 - \nu_2$ fraction of the outgoing edges are in \mathcal{E}_1 .

\mathcal{V}_3 : In \mathcal{V}_2 and $1 - \nu_3$ fraction of the outgoing edges are in \mathcal{E}_4 .

\mathcal{V}_4 : In \mathcal{V}_3 and is not a strange vertex (see Definition 5.13 below).

Lemma 5.7. *If $L' \gg L$, we can choose a horocycle H such that for the H -graph $\mathcal{G}(H)$, $1 - \delta_1$ fraction of vertices are in \mathcal{V}_1 . Here, δ_1 is a function of δ_0 which tends to 0 as $\delta_0 \rightarrow 0$.*

Proof. Note that $\phi^{-1}(B(L'))$ has small boundary area (compared to the volume). Now tile $\phi^{-1}(B(L'))$ by boxes $B(L)$. Since $L' \gg L$, most boxes are completely in the interior of $\phi^{-1}(B(L'))$.

Let \mathcal{U} denote the set where we know the map is locally standard (but could be right side up or upside down). Note that for every box $B(L)$, $|\mathcal{U} \cap B(L)| \geq 0.999|B(L)|$.

This implies that for most H ,

$$|H \cap \phi^{-1}(B(L')) \cap \mathcal{U}| \geq 0.99|H \cap \phi^{-1}(B(L'))|$$

Then for such H , \mathcal{V}_1 , which consists of vertices on $\phi(H) \cap B(L') \cap \phi(\mathcal{U})$, satisfies the conditions of the lemma. \square

We now fix H such that Lemma 5.7 holds.

Lemma 5.8. *At least $1 - \epsilon_1$ fraction of the edges of $\mathcal{G}(H)$ are in \mathcal{E}_1 . Here, ϵ_1 is a function of δ_0 which tends to 0 as $\delta_0 \rightarrow 0$.*

Proof. Recall that $m \leq M(v) \leq M$ where $M(v)$ is the degree of v . This implies that

$$\frac{1}{m} \leq \frac{|\mathcal{V}(H)|}{|\mathcal{E}(H)|} \leq \frac{1}{M}.$$

Since each edge not in \mathcal{E}_1 must be quasi-incident on a vertex not in \mathcal{V}_1 and each vertex is incident to at most M edges, we have:

$$|\mathcal{E}_1^c| \leq 2M|\mathcal{V}_1^c|.$$

Combined with equation (5.2) this implies

$$\frac{|\mathcal{E}_1^c|}{|\mathcal{E}(H)|} \leq 2 \frac{M}{m} \frac{|\mathcal{V}_1^c|}{|\mathcal{V}(H)|}$$

Thus the lemma follows from Lemma 5.7. \square

Lemma 5.9. *At least $1 - \delta_2$ fraction of the vertices of $\mathcal{G}(H)$ are in \mathcal{V}_2 . Here, δ_2 is a function of δ_0 which tends to 0 as $\delta_0 \rightarrow 0$.*

Proof. This follows immediately from Lemma 5.8. \square

Lemma 5.10. *At least $1 - \epsilon_3$ fraction of the edges of $\mathcal{G}(H)$ are in \mathcal{E}_3 . Here, ϵ_3 is a function of δ_0 which tends to 0 as $\delta_0 \rightarrow 0$.*

Proof. In view of Lemma 5.9, it enough to prove that for any $v \in \mathcal{V}_2$, almost all the edges outgoing from v belong to \mathcal{E}_3 .

Suppose $v \in \mathcal{V}_2$. Let $\mathcal{E}(v)$ denote all the edges which are incident to v . We know that most edges in $\mathcal{E}(v)$ belong to \mathcal{E}_2 , i.e.

$$|\mathcal{E}_2^c \cap \mathcal{E}(v)| \leq \delta_2 |\mathcal{E}(v)| \quad (20)$$

Let $A_v = \mathcal{E}(v) \cap \mathcal{E}_3^c$ denote the edges outgoing from v which are not in \mathcal{E}_3 . We know that for any $e \in A_v$, there exists $x \in e$ such that at least θ_3 of the edges branching from e at x are not in \mathcal{E}_1 . Thus there exists a neighborhood U of e such that

$$|\mathcal{E}_2^c \cap U \cap \mathcal{E}(v)| \geq \theta_3 |U \cap \mathcal{E}(v)|$$

We thus get a cover of A_v by U 's. Then by Vitali's covering lemma, there exists disjoint U_j such that

$$\sum_{j=1} |U_j| \geq \frac{1}{2} |A_v|.$$

Thus,

$$|A_v| \leq 2 \sum |U_j| \leq \frac{2}{\theta_3} \sum |U_j \cap \mathcal{E}_2^c \cap \mathcal{E}(v)| \leq \frac{2}{\theta_3} |\mathcal{E}_2^c \cap \mathcal{E}(v)|$$

Then, by (20),

$$|A_v| \leq \frac{2\delta_2}{\theta_3} |\mathcal{E}(v)|.$$

We now choose $\theta_3 = \epsilon_3 = \sqrt{2\delta_2}$. \square

Lemma 5.11. *At least $1 - \epsilon_4$ fraction of the edges of $\mathcal{G}(H)$ are in \mathcal{E}_4 . Here, ϵ_4 is a function of δ_0 which tends to 0 as $\delta_0 \rightarrow 0$.*

Proof. It follows immediately from Lemma 5.10 that $1 - 2\epsilon_3$ proportion of the for non-leaf edges have the reverse branching property.

Let $\alpha = \epsilon_1^{1/6}$. If the proportion of the leaf edges is at most α , we are already done (with $\epsilon_4 = 2\epsilon_3 + \alpha$). Thus we may assume that the proportion of leaf edges is at least α .

Let Y be the set of all vertical paths in $B(L)$ going from top to bottom, and $Y' \subset Y$ is the subset consisting of paths which pass through a vertex not in \mathcal{V}_1 . Let $D(\gamma) = 1$ if $\gamma \in Y'$ and $D(\gamma) = 0$ otherwise. By lemma 5.8, we have:

$$\sum_{\gamma \in Y} D(\gamma) \leq 2\epsilon_1 |\mathcal{E}(H)|. \quad (21)$$

From (21),

$$\sum_{\gamma \in Y} D(\gamma) \leq \epsilon_1 |\mathcal{E}(H)| \leq \frac{\epsilon_1}{\alpha} |\mathcal{E}_{leaf}|,$$

where $\mathcal{E}_{leaf} \subset \mathcal{E}(H)$ denotes the set of leaf edges. For a point $v \in \partial B(R)$, let Y_v denote the set of geodesics emanating from v . We get

$$\sum_{v \in \partial B(R)} \sum_{\gamma \in Y_v} D(\gamma) \leq \sum_{v \in \partial B(R)} \frac{\epsilon_1}{\alpha} |\mathcal{E}_{leaf}(v)|,$$

where $\mathcal{E}_{leaf}(v)$ denotes the set of leaf edges emanating from v . Let $\theta' = \epsilon_1^{2/3}$, and let

$$P = \left\{ v \in \partial B(R) : \sum_{\gamma \in Y_v} D(\gamma) > \theta' |\mathcal{E}_{leaf}(v)| \right\}$$

Note that

$$\sum_{v \in P} |\mathcal{E}_{leaf}(v)| \leq \sum_{v \in P} \frac{1}{\theta'} \sum_{\gamma \in Y_v} D(\gamma) \leq \frac{1}{\theta'} \sum_{v \in Y} D(\gamma) \leq \frac{\epsilon_1}{\alpha \theta'} |\mathcal{E}_{leaf}|$$

Thus, since we choose α and θ' so that $\frac{\epsilon_1}{\alpha \theta'} \ll 1$, it is enough to prove that for $v \notin P$, most of the edges in $\mathcal{E}_{leaf}(v)$ are in \mathcal{E}_4 .

Now assume $v \notin P$. Thus we have

$$\sum_{\gamma \in Y_v} D(\gamma) < \theta' |\mathcal{E}_{leaf}(v)|$$

Choose $\theta_4 = \epsilon_1^{1/12}$. Let $A_v = \mathcal{E}_{leaf}(v) \cap \mathcal{E}_4^c$ denote the leaf edges outgoing from v which are not in \mathcal{E}_4 . We know that for any $e \in A_v$, there exists $x \in e$ such that

at least θ_4 of the edges branching from e at x are not in \mathcal{E}_2 . Thus there exists a neighborhood $U \subset Y_v$ with $e \in U$ such that

$$|\mathcal{E}_2^c \cap U| \geq \theta_4 |U|$$

hence using the definition of \mathcal{E}_2 ,

$$\sum_{\gamma \in U} D(\gamma) \geq \theta_2 \theta_4 |U|.$$

We thus get a cover of A_v by U 's. Then by Vitali, there exists disjoint U_j such that

$$\sum_{j=1} |U_j| \geq \frac{1}{2} |A_v|.$$

Thus,

$$|A_v| \leq 2 \sum_j |U_j| \leq \frac{2}{\theta_2 \theta_4} \sum_j \sum_{\gamma \in U_j} D(\gamma) \leq \frac{2}{\theta_2 \theta_4} \sum_{\gamma \in Y_v} D(\gamma) \leq \frac{2\theta'}{\theta_2 \theta_4} |\mathcal{E}_{leaf}(v)|.$$

Since $\theta_2 = \epsilon_1^{1/2}$, $\frac{2\theta'}{\theta_2 \theta_4} = \epsilon_1^{\frac{1}{12}}$ and the lemma follows. \square

Lemma 5.12. *At least $1 - \delta_4$ fraction of the vertices of $\mathcal{G}(H)$ are in \mathcal{V}_3 . Here, δ_4 is a function of δ_0 which tends to 0 as $\delta_0 \rightarrow 0$.*

Proof. This follows from Lemma 5.11. \square

Let H_* be an horocycle intersecting $B(L')$. We say that an S -vertex on w on H_* is *marked* by a \mathcal{V}_1 H -vertex v if the cloud of v contains a point of H_* , and also $h(v) = h(H_*) + O(\rho_2)$, and also the coordinates of v and w along H_* must agree up to $O(\rho_2)$. (In particular the orientation of v must be such that the coordinate of v along H_* is not “fuzzy”).

Definition 5.13 (Strange Vertex). An H -vertex $v \in \mathcal{V}_3$ is called *strange* if there is an horizontal segment (i.e. piece of horocycle) K marked by v such that more than $1 - \nu_4$ fraction of the S -vertices on K are marked by H -vertices which are \mathcal{V}_1 but not in \mathcal{V}_3 .

Lemma 5.14. *At least $1 - \delta_6$ fraction of the vertices of $\mathcal{G}(H)$ are in \mathcal{V}_4 (i.e are in \mathcal{V}_3 and not strange). Here, δ_6 is a function of δ_0 which tends to 0 as $\delta_0 \rightarrow 0$.*

Proof. Let v_1, \dots, v_m be the strange vertices, and let K_1, \dots, K_m be horocycle segments marked by the strange vertices. The K_i are not quite uniquely defined, but we address this issue below.

Note that the number of H -vertices which can mark a given S -vertex is $O(\rho_2)$. Indeed, any two such vertices must be within $O(\epsilon'R)$ of each other, which means that they must have come from the same good box, which implies that heights and their transverse coordinates must agree. (Recall that the vertices which come from near the edges of a good box are automatically not in \mathcal{V}_1).

The same argument shows that one can choose the K_i so that for $i \neq j$, $d(K_i, K_j) > 3D\epsilon'R$. Now we can apply the Vitali covering lemma to the K_i . This lemma applies since each K_i is one-dimensional and the different K_i do not interact with each other. Also the density of the \mathcal{V}_1 vertices which are not in \mathcal{V}_3 is small by Lemma 5.12. This implies that the strange vertices are a small fraction of all the vertices. \square

5.3 Circuits

The projection π_H and the function $\rho_H(\cdot, \cdot)$. Let H be a horocycle. Let $\pi_H : \text{Sol} \rightarrow \mathbb{H}^2$ denote the orthogonal projection to the hyperplane orthogonal to H . We let $\rho_H(p, q) = (\pi_H(p) | \pi_H(q))_{\pi_H(H)}$ be the Gromov product of $\pi_H(p)$ and $\pi_H(q)$ with respect to $\pi_H(H)$ in \mathbb{H}^2 . Recall that for three points x, y, z in a metric space X , the Gromov product is defined as:

$$(y|z)_x = \frac{1}{2} \{d_X(x, y) + d_X(x, z) - d_X(z, y)\}.$$

Let γ_{yz} be the geodesic joining y to z . In a δ -hyperbolic space X satisfies

$$d_X(\gamma_{yz}, x) - \delta \leq (y|z)_x \leq d_X(\gamma_{yz}, x)$$

see e.g. [GhdIH, Lemma 2.17]. We note the following properties of ρ_H :

Lemma 5.15. (i) *Suppose $d(p', p) \ll d(p, H)$, $d(q', q) \ll d(q, H)$, and $\rho_H(p, q) \ll \min(d(p, H), d(q, H))$. Then,*

$$\rho_H(p, q) \approx \rho_H(p', q').$$

(ii) *Suppose $h(p') < h(p)$, $h(q') < h(q)$, the points p and p' can be connected by a vertical geodesic, and the same for the points q and q' . Suppose also $d(p, H) \gg \rho_H(p, q)$ and $d(q, H) \gg \rho_H(p, q)$. Then,*

$$\rho_H(p, q) \approx \rho_H(p', q').$$

(iii) If $\rho_H(p, q) > s$ and $\rho_H(q, q') > s$, then $\rho_H(p, q') > s$ (up to a small error).

Proof. The statements (i), (ii) and (iii) are standard hyperbolic geometry. In particular (iii) follows immediately from the “thin triangle” property. \square

In the following lemma, the horocycle H is assumed to be an y horocycle. An analogous lemma, with a few sign changes, holds for x horocycles.

Lemma 5.16. *Suppose $p, q \in X(n)$ are connected by a path $\hat{\gamma}$ in such that*

$$h(x) \leq h(H) - \rho_4 \text{ for all } x \in \gamma, \quad (22)$$

Further assume the initial segments of γ at both p and q are vertical geodesics going down for length at least $\epsilon'R$, that γ stays below $h(H) - R'$ except on these initial segments and that the length of γ is less than $e^{\epsilon'R}$. Then, $\rho_H(p, q) > \Omega(\rho_4)$.

Proof. This is just a standard hyperbolic lemma applied to $\pi_H(\gamma)$. \square

Notation. An \mathcal{E}_2 edge is a monotone vertical path in the \hat{S} -graph which is a subset of an \mathcal{E}_1 edge (or possibly a subset the extension of an \mathcal{E}_1 edge by at most ρ_4 at each end).

Lemma 5.17. *Suppose $\gamma = \overline{p_0q_0}$ is an \mathcal{E}_2 edge going up from an x -oriented irregular \hat{S} -vertex (or going down from an y -oriented irregular \hat{S} -vertex). Suppose $p \in \gamma$ is within the same $B_i(R)$ as p_0 , and $q \in \gamma$ is within the same $B_{i'}(R)$ as q_0 and $d(p, p_0)$ and $d(q, q_0)$ is at least $10\epsilon'R$. Then the following hold:*

(i) *Except at near its endpoints, γ never passes through any irregular \hat{S} -vertices. In other words, in its interior, γ never comes within R' of a good H vertex.*

(ii) *We have $\rho_H(\hat{\phi}^{-1}(p), \hat{\phi}^{-1}(q)) > \Omega(\rho_4)$.*

Remark. In the above $\rho_H(\hat{\phi}^{-1}(p), \hat{\phi}^{-1}(q))$ is well defined since for an \hat{S} -vertex v , $\pi_H(\hat{\phi}^{-1}(v))$ is well defined (even though $\hat{\phi}^{-1}(v)$ may not be).

Informal outline of proof. We first outline the proof, and then give the full argument. Consider $\phi^{-1}(\gamma)$. Note that below height $h(H) - R'$, $\phi^{-1}(\gamma)$ cannot move transverse to H because it is of length at most $O(L)$. Because of this, whenever γ attempts to cross above height $h(H) - R'$ it must do so in the image of $W(H)$. Consider the point q' where it does so. Since γ cannot hit a bad shadow vertex, q' must be essentially in $U \cap B_i(R) \cap W(H)$. But then, by the definition of the \hat{S} -graph,

γ must hit an H -vertex. Thus, q' is near the endpoint of γ , and thus (i) holds. Now (ii) follows from Lemma 5.16 since we know that $\phi^{-1}(\gamma)$ has not passed above height $h(H) - R'$ except near the endpoints.

Proof. Let p_1 be the first place where γ hits $\hat{W}(H)$. Then, since γ cannot hit a bad shadow vertex, there exists $p'_1 \in U_* \cap W(H)$ such that $\hat{\phi}(p'_1) = p_1$ and $d(\phi^{-1}(p_1), \hat{\phi}^{-1}(p_1)) = O(\epsilon'R)$. Note that p'_1 and $\phi^{-1}(p_1)$ are both $\Omega(\epsilon'R)$ from the sides of $W(H)$.

Let p'_2 be the next point after p'_1 when $\phi^{-1}(\gamma)$ intersects $\tilde{B} \cap \{x : h(x) = h(H) - R'\}$ at $\phi^{-1}(p_2)$. Since γ is an \mathcal{E}_2 edge and in particular a vertical geodesic, we know $d(p'_1, p'_2)$ is $\Omega(\beta R)$. By the choice of p'_2 and the definition of \mathcal{E}_2 edge, $\phi^{-1}\gamma$ never hits a shadow vertex between p'_1 and p'_2 . This fact and the fact that $|\phi^{-1}(\gamma)| < O(L)$ imply that p'_2 must be in $W(H)$. Since γ is \mathcal{E}_2 , p'_2 is not a bad shadow vertex and in particular is away from the edge of $W(H)$. Together this implies that p'_2 is in $W(H) \cap \tilde{B}$ and that the continuation of γ past $p_2 = \hat{\phi}(p'_2)$ must, by the definition of the \hat{S} and H graphs, hit an H -vertex. Since γ is an \mathcal{E}_2 edge and does not contain good vertices in its interior, this implies that p_2 and q_0 are in the same box and that the segment from p_2 to q_0 contains q . Now by Lemma 5.15,

$$\rho_H(\hat{\phi}^{-1}(p), \hat{\phi}^{-1}(q)) = \rho_H(\hat{\phi}^{-1}(p_1), \hat{\phi}^{-1}(p_2)) \approx \rho_H(\phi^{-1}(p_1), \phi^{-1}(p_2)) \geq \Omega(\rho_4)$$

□

Lemma 5.18. *Suppose $\overline{p_0q_0}$ is an \mathcal{E}_2 edge (which goes up from a x -oriented vertex and down from an y -oriented vertex), $\rho_3 \gg s \gg \rho_1$, and p (resp. q) is on γ distance s away from p_0 (resp. from q_0). Then there exists a horocycle H' such that p and q are within $O(\rho_1)$ of $\hat{\phi}(H')$.*

Proof. Choose points p' and q' on $\overline{p_0q_0}$ close to where $\overline{p_0q_0}$ enters the respective good boxes. Applying Lemma 5.17 we see that $\rho_H(\hat{\phi}^{-1}(p'), \hat{\phi}^{-1}(q')) > \underline{\Omega(\rho_4)}$. By the usual δ -thin triangle properties, this implies that the geodesic segments $\pi_H(\hat{\phi}^{-1}(p))\pi_H(H)$ and $\pi_H(\hat{\phi}^{-1}(q))\pi_H(H)$ stay close till roughly for roughly ρ_4 units from $\pi_H(H)$. Since $d(\hat{\phi}^{-1}(p), H) \ll \rho_3 < \rho_4$ and similarly for $\hat{\phi}^{-1}(q)$ this implies that $\pi_H(\hat{\phi}^{-1}(p))$ and $\pi_H(\hat{\phi}^{-1}(q))$ are within 2δ of the same vertical geodesic through the point $\pi_H(H)$. But since they are at the same height, this implies that $\pi_H(\hat{\phi}^{-1}(p)) = \pi_H(\hat{\phi}^{-1}(q))$. □

Suppose H' is a horocycle obtained by moving up less than ρ_3 from H . Recall that the set $\hat{\phi}(H')$ is a well defined subset of the \hat{S} -graph (see §5.1). We always assume that $\hat{\phi}(H')$ runs along vertices in the \hat{S} -graph (or else project it). By Lemma 5.18,

given any collection of \mathcal{E}_2 edges with (some) endpoints on H , we may replace them with \mathcal{E}_2 edges with (some) endpoints on H' .

Lemma 5.19 (Illegal Circuit). *Suppose n is some finite even integer which is not too large (we will use $n = 4$ and $n = 6$), and for $0 \leq i \leq n - 1$, p_i are \hat{S} -vertices. Also suppose that for $0 \leq i \leq n - 1$, $\overline{p_{i-1}p_i}$ are subsets of \mathcal{E}_2 edges, where $i - 1$ is considered mod n . For $1 \leq i \leq n$, let $r^\pm(p_i)$ denote the maximum distance the geodesic $\overline{p_{i\pm 1}p_i}$ can be continued beyond p_i while remaining a subset of an \mathcal{E}_2 edge, and let $r(p_i) = \max(r^+(p_i), r^-(p_i))$.*

Suppose there is an index k such that $r(p_k) \ll \rho_4$, and for all $i \neq k$, $r(p_i) > r(p_k) + 2\rho_1$. Then $\overline{p_{k-1}p_k}$ and $\overline{p_k p_{k+1}}$ cannot have only the point p_k in common.

Remark: Roughly, the point of the lemma is that one cannot find a loop of length $O(L)$ through a point on the horocycle which begins by going up in two distinct directions unless the loop comes back to the original horocycle.

Proof. Without loss of generality, $k = 0$. Let H' be the horocycle passing through $\hat{\phi}^{-1}(p_0)$. By Lemma 5.18 and the discussion following, we can consider H' in place of H , namely we can replace all H vertices that occur in our arguments with vertices in H' . Let p_{i-1}^+ be the first time when $\overline{p_{i-1}p_i}$ leaves $\hat{\phi}(\tilde{B} \cap W(H))$, and let p_{i-1}^- be the last time when $\overline{p_{i-1}p_i}$ enters $\hat{\phi}(\tilde{B} \cap W(H))$ (so $d(p_{i-1}, p_{i-1}^+) \approx R' \leq R$, and $d(p_{i-1}^-, p_i) \approx R' \leq R$). By applying Lemma 5.17 to each segment $\overline{p_{i-1}^+, p_i^-}$ we see that $\rho_{H'}(\hat{\phi}^{-1}(p_{i-1}^+), \hat{\phi}^{-1}(p_i^-)) \geq \Omega(\rho_4)$.

Now, by assumption, for all $i \in [0, n - 1]$ except $k = 0$, $\rho_{H'}(\hat{\phi}^{-1}(p_i^-), \hat{\phi}^{-1}(p_i^+)) \geq 2\rho_1$, but for $i = 0$, $\rho_{H'}(\hat{\phi}^{-1}(p_0^-), \hat{\phi}^{-1}(p_0^+)) \leq \rho_1$. This is a contradiction to Lemma 5.15 (iii). \square

5.4 Families of geodesics

Let $B[\lambda]$ be a box in $X(n')$ of combinatorial size λ (i.e. the number of edges from the top to the bottom is λ , and the distance from the top to the bottom is $\rho_1\lambda$). Let \mathfrak{b} be the branching number of each vertex (i.e the valence of each vertex counting both up and down branching is $2\mathfrak{b}$). Note that \mathfrak{b} is related to the constant B'_X of §3.3 by $\mathfrak{b}^{2\lambda} = e^{B'_X \rho_1 \lambda}$, so $\log \mathfrak{b} = B'_X \rho_1 / 2$.

Thus the number of \hat{S} -vertices on the top edge of $B[\lambda]$ is \mathfrak{b}^λ , and so is the number of \hat{S} -vertices on the bottom edge. The total number of vertical geodesics in $B[\lambda]$ is $\mathfrak{b}^{2\lambda}$.

Lemma 5.20. *The number of vertices of the H -graph in $B[\lambda]$ is at most $c_9(\rho_1)\mathfrak{b}^\lambda$.*

Proof. Apply Lemma 3.8 with Q the union of the top edge and the bottom edge of the box. \square

Given a box $B(D)$ and a vertical geodesic segment γ of length D in B , we say γ is *through* if γ does not hit any H vertex in B . The following lemma applies to families of geodesics in a box. Note that the geodesics are not assumed to be part of the H -graph. The point of the lemma is that if too many paths through the box are blocked by good vertices, then some good vertex must block many paths. This really only depends on the fact that there are not too many good vertices in the box.

Lemma 5.21. *Let $B[\lambda]$ be a box of combinatorial size λ . Suppose \mathcal{F} is a family of vertical geodesics (actually monotone paths in the modified \hat{S} -graph going from the top of $B[\lambda]$ to the bottom) with the following properties:*

- (a) *Each geodesic in \mathcal{F} does not hit any bad vertices*
- (b) *$|\mathcal{F}|$ (i.e. the number of geodesics in \mathcal{F}) is at least $\sigma \mathfrak{b}^{2\lambda}$, where $0 < \sigma < 1$.*
- (c) *For some $\rho \in \mathbb{N}$ (we will always use $\rho = \rho_2$), fewer than $1 - \frac{c_9(\rho_1)}{\mathfrak{b}^\rho} \mathfrak{b}^{2\lambda}$ of the geodesics in \mathcal{F} are through (i.e. do not contain any H -vertices in their interior).*

Then there exists a vertex $v \in \mathcal{V}_1$ not on the bottom edge or within ρ of the top edge of $B[\lambda]$, and two geodesics in \mathcal{F} which pass through v and stay together for fewer than ρ \hat{S} -edges.

Thus, if $\sigma \gg \frac{c_9(\rho_1)}{\mathfrak{b}^\rho}$, almost all of the geodesics in \mathcal{F} are through unless we have a configuration as described in the conclusion of the lemma.

Proof. Let \mathcal{F}_0 denote the family of all vertical geodesics on the unmodified \hat{S} -graph, passing from the top of $B[\lambda]$ to the bottom. Clearly $|\mathcal{F}_0| = \mathfrak{b}^{2\lambda}$. Note that $B[\lambda]$ has $\lambda \mathfrak{b}^\lambda$ \hat{S} -vertices, and each geodesic contains λ \hat{S} -vertices. This implies that each \hat{S} -vertex lies on $M = \mathfrak{b}^\lambda$ geodesics in \mathcal{F}_0 .

Now suppose v is an H -vertex in \mathcal{V}_1 , and v is not on the bottom edge or within ρ of the top edge. Assuming the conclusion of the lemma fails then v can belong to at most $M \mathfrak{b}^{-\rho}$ geodesics in \mathcal{F} . Thus, using Lemma 5.20, we see that the total number of geodesics in \mathcal{F} which pass through a vertex in \mathcal{V}_1 in B is at most

$$M \mathfrak{b}^{-\rho} c_9(\rho_1) \mathfrak{b}^\lambda = \frac{c_9(\rho_1)}{\mathfrak{b}^\rho} \mathfrak{b}^{2\lambda}$$

which implies that all but $\frac{c_9(\rho_1)}{\mathfrak{b}^\rho} \mathfrak{b}^{2\lambda}$ of the geodesics are through, contradicting (c).

□

Convention For the remainder of this subsection, we assume that we have a x horocycle H whose image (at least in some initial box) is x -oriented. The proof proceeds by extending the set on which the image is horocycle and so all points in the H graph we consider will be x -oriented.

The intervals $I_\lambda(v)$ and $I'_\lambda(v)$. For any \hat{S} -vertex v , let $I_\lambda(v)$ denote the set of vertices on the same x -horocycle as v which are within combinatorial distance 2λ . Let $I'_\lambda(v)$ denote the set of vertices which can be reached from v by a monotone path going up for exactly λ steps (so $I'_\lambda(v)$ is a piece of y -horocycle). Note that for DL graphs, each point of $I_\lambda(v)$ is connected to each point of $I'_\lambda(v)$ by a monotone path of length λ , and for any $w \in I_\lambda(v)$, $I_\lambda(w) = I_\lambda(v)$ and $I'_\lambda(w) = I'_\lambda(v)$. For Sol slightly more complicated variants of these statements hold, for instance for any $w \in I_\lambda(v)$, $I_\lambda(w)$ and $I_\lambda(v)$ intersect in a set that contains more than half the measure of each. And the relative measure of this intersection in each set is close to one, unless w is close to an edge of $I_\lambda(v)$.

Let v be any \hat{S} -vertex. Let $\mathcal{U}(v, \lambda)$ denote the set of distinct monotone geodesic segments in the \hat{S} graph going up from v for distance exactly λ . Then $\mathcal{U}(v, \lambda)$ is the set of geodesics joining v to $I'_\lambda(v)$. Similarly, we let $\mathcal{D}(w, \lambda)$ be the set of distinct monotone geodesics segments in the \hat{S} -graph going down distance λ from w . If $w \in I'_\lambda(v)$ then $\mathcal{D}(w, \lambda)$ is the set of monotone geodesics joining w to points in $I_\lambda(v)$.

Proposition 5.22 (Extension of Horocycles I). *Suppose $v \in \mathcal{V}_3$. Suppose $\sigma \gg \eta \gg c_9(\rho_1)/\mathfrak{b}^{\rho_2}$ and suppose λ is such that at least σ -fraction of the edges going up from v are \mathcal{E}_4 edges of length at least $\lambda + \rho_2$. Then at least $1 - O(\eta)$ fraction of the \hat{S} -vertices in $I_\lambda(v)$ are in fact H -vertices.*

Proof. We assume that $v \in \mathcal{V}_3$ and $\hat{\phi}(H)$ is oriented as a x -horocycle near v . Let E denote the set of \mathcal{E}_4 edges coming out of v which have length at least $\lambda + \rho_2$. Let E_λ be the set of vertices in $I'_\lambda(v)$ which are on of $\lambda + \rho_2$ unobstructed geodesics leaving v . By assumption, we have

$$|E_\lambda| \geq \sigma \mathfrak{b}^\lambda. \quad (23)$$

We now let $\mathcal{F}'_0 = \bigcup_{w \in E_\lambda} \mathcal{D}_\lambda(w)$ and let \mathcal{F}' be all the geodesics segments in \mathcal{F}'_0 which do not contain a bad vertex. Assume for a contradiction that many geodesics in \mathcal{F}' are not through, i.e. that (c) of Lemma 5.21 holds for \mathcal{F}' . We verify that Lemma 5.21(a) and (b) hold for \mathcal{F}' . Since $v \in \mathcal{V}_3$,

$$(1 - \theta_4)|E_\lambda| \mathfrak{b}^\lambda \leq |\mathcal{F}'| \leq |E_\lambda| \mathfrak{b}^\lambda \quad (24)$$

Note that by (23) and (24), we have $|\mathcal{F}'| \geq \sigma \mathfrak{b}^{2(\lambda)}$. Hence Lemma 5.21 (b) holds. Note that all the geodesics in \mathcal{F}' end at points of $I_\lambda(v)$.

Now by Lemma 5.21 there exists $w \in \mathcal{V}_1$ with $h(w) > h(v)$ and an \hat{S} -vertex w_1 with $h(w_1) > h(w)$ and $d(w, w_1) < \rho_2$ so that at least two geodesics in \mathcal{F}' meet at w_1 and continue to w . (See figure 2). Let $x \in I'_\lambda(v)$ and $y \in I'_\lambda(v)$ be the starting points of these two geodesics.




Figure 2. Proof of Proposition 5.22.
The filled boxes denote H -vertices.

Let z be the last common point of the geodesics \overline{vx} and \overline{vy} . We now apply Lemma 5.19 to the points $\langle w_1, x, z, y \rangle$. Note that $r(w_1) < \rho_2$ (because of w). Also by assumption, $r(x) \geq \rho_2 > r(w_1)$ and $r(y) \geq \rho_2 > r(w_1)$. Note that $h(z) = h(w_1)$, hence $r(z) = h(z) - h(v) = h(w_1) - h(v) > h(w_1) - h(w) = r(w_1)$. Hence we get a contradiction by Lemma 5.19. Hence we cannot have condition (c) of Lemma 5.21 therefore all but $O(\eta)$ of the geodesics in \mathcal{F}' are unobstructed. Therefore the number of unobstructed geodesics in \mathcal{F}' is at least

$$(1 - O(\eta))|\mathcal{F}'| \geq (1 - O(\eta))(1 - \theta_4)|E_\lambda|\mathfrak{b}^\lambda \quad (25)$$

where we have used (24) to get the second estimate.

Now, let $U' \subset I_\lambda(v)$ be the set of \hat{S} -vertices (at height $h(v)$) which are the endpoints of at least two geodesics in \mathcal{F}' . Since every vertex can be reached by at most $|E_\lambda|$ geodesics, we have by (25),

$$|U'| \geq (1 - O(\eta))(1 - \theta_4)\mathfrak{b}^\lambda. \quad (26)$$

I.e., then U' has almost full measure in $I_\lambda(v)$.

Now suppose $w \in I_\lambda(v)$ is such that two unobstructed geodesics in \mathcal{F}' end at w . Let us denote these geodesics by \overline{wx} and \overline{wy} where $x, y \in I'_\lambda(v)$. By definition of \mathcal{F}' , \overline{xv} and \overline{yv} are unobstructed. We now apply Lemma 5.19 to the points $\langle w, x, v, y \rangle$. Note that $r(v) = 0$ (since v is an H -vertex), and also $r(x) \geq \rho_2$, and $r(y) \geq \rho_2$. Thus, by Lemma 5.19, we get a contradiction unless $r(w) = 0$, i.e. w is an H -vertex. \square

If $v \in \mathcal{V}_4$ then the conclusion is strengthened automatically to imply that most vertices in $I_\lambda(v)$ are in \mathcal{V}_3 . This is used in the following proposition.

Proposition 5.23 (Zero-One Law). *Suppose $v \in \mathcal{V}_4$. Suppose λ is such that the fraction of the edges in $\mathcal{U}(v, \lambda)$ which are in \mathcal{E}_4 and are unobstructed for at least length $\lambda + \rho_2$ is at least $\sigma \gg c_9(\rho_1)/\mathfrak{b}^{\rho_2}$. Let $\mathcal{F} = \bigcup_{w \in I_\lambda(v)} \mathcal{U}(w)$. Then at least $1 - O(\eta)$ fraction of the edges in \mathcal{F} are unobstructed for length $\lambda + \rho_2$.*

Proof. As in the previous proposition, let E_λ be the set of vertices in $I'_\lambda(v)$ which are on of $\lambda + \rho_2$ unobstructed geodesics leaving v . Also let $U' \subset I_\lambda(v)$ and \mathcal{F}' be as in Proposition 5.22.

Now since v is not a strange vertex, the subset U'' of $I_\lambda(v)$ consisting of \mathcal{V}_3 vertices in U' is of almost full measure in $I_\lambda(v)$. Let

$$\mathcal{F}'' = \bigcup_{w \in U''} \mathcal{U}(w) \cap \mathcal{E}_2.$$

(so \mathcal{F}'' consists of all the \mathcal{E}_2 edges coming out of all the “good” H -vertices on $I_\lambda(v)$). We cut off all the geodesics in \mathcal{F}'' after they cross $I'_\lambda(v)$.

We want to apply Lemma 5.21 to \mathcal{F}'' in the box of size λ , but there are technical difficulties here in verifying Lemma 5.21. To overcome these difficulties, we look at a horocircle H' that is ρ_4 units below H with the same orientation. By a discussion similar to Lemma 5.18 and following and the fact that \mathcal{F}'' consists of edges in \mathcal{E}_2 , we can extend every geodesic segment in \mathcal{F}'' by ρ_4 on top and bottom in all possible ways to obtain a family \mathcal{F}'_{long} . We will apply Lemma 5.21 to \mathcal{F}''_{long} instead. If almost all segments in \mathcal{F}''_{long} are unobstructed by H' , it is immediate that almost all segments in \mathcal{F}'' . We let U'_{long} be the set of H' vertices within ρ_4 of U' .

We have that $|\mathcal{F}''_{long}| \geq (1 - O(\eta))\mathfrak{b}^{2(\lambda + \rho_4)}$, so (b) is satisfied. Also (a) is satisfied since the relevant edges are in \mathcal{E}_2 . If (c) does not hold, we are done, so we assume (c) holds. This implies that the conclusion of the lemma is true, and we show this yields an illegal circuit. (see Figure 3).

By Lemma 5.21 exists an H' -vertex q with $h(q) < h(v) + \lambda + 2\rho_4$ and an \hat{S} -vertex q_* with $h(q) - \rho_2 < h(q_*) \leq h(q)$ such that at least two geodesics in \mathcal{F}''_{long} come




Figure 3. Proof of Proposition 5.23.
The filled boxes denote H' -vertices.

together at q_* . Let these geodesics be $\overline{u_1 q_*}$ and $\overline{u_2 q_*}$ where for $i = 1, 2$, $u_i \in U'_{long}$. Let $w_i = \overline{u_i q_*} \cap U'$ denote the corresponding point in U' . Since $w_i \in U'$, there exists $x_i \in I'_{\lambda + \rho_4}(v)$ such that $\overline{w_i x_i}$ and $\overline{x_i v}$ are both \mathcal{E}_2 and unobstructed. Since v' denote any point on U'_{long} that is ρ_4 units below v . We now apply Lemma 5.19 to the points $\langle q_*, w_1, x_1, v, x_2, w_2 \rangle$. Note that by construction, $r(q_*) < \rho_2 \ll \rho_3$, $r(v) = \rho_4$, and for $i = 1, 2$, $r(w_i) = \rho_4$, $r(x_i) \geq \rho_2$. Thus by Lemma 5.19, $q_* w_1$ and $q_* w_2$ do not diverge at q_* , which is a contradiction. \square

Theorem 5.24 (Extension of Horocycles II). *Suppose $v \in \mathcal{V}_4$ is x -oriented. Let s denote the height difference between v and the top of $B(L')$ and assume $s > 4\kappa^2 \beta'' R$. Then, the density of x -oriented \mathcal{V}_3 H -vertices along $I_s(v)$ is $1 - O(\eta)$.*

Remark: The proof of this Theorem is considerably simpler in the case of DL-graphs as boxes in DL graphs have “no sides”. We give the proof first in this case. The Sol case is complicated by needing to avoid having paths “escape off the sides of the box.”

Proof for DL graphs. For an x -oriented \mathcal{V}_4 vertex w , let where $f(w, \lambda)$ denotes the proportion of edges in $\mathcal{U}(w)$ which are \mathcal{E}_4 and unobstructed for length $\lambda + \rho_2$. Let

$$f^*(v, \lambda) = \sup_{w \in I_\lambda(v) \cap \mathcal{V}_4} f(w, \lambda).$$

In view of Proposition 5.23, for any λ for which $f^*(v, \lambda) \geq O(\eta)$, $f^*(v, \lambda) > 1 - O(\eta)$.

Thus, either for all $1 \leq \lambda \leq s$, $f_j^*(v, \lambda) \geq 1 - O(\eta)$, in which case Theorem 5.24 holds in view of Proposition 5.22 and 5.23, or else there exists minimal λ such that

$f^*(v, \lambda) > 1 - O(\eta)$, and also $f^*(v, \lambda + 1) < O(\eta)$. Note that $\lambda > \Omega(\beta''R)$ by the definition of good vertices and the \hat{S} and H -graphs. Let $w \in I_\lambda(v) \cap \mathcal{V}_4$ be such that the sup in the definition of $f^*(v, \lambda)$ is realized at w . Hence, by Proposition 5.22 (i), all but $O(\eta)$ fraction of the \hat{S} -vertices in $I_\lambda(w) = I_\lambda(v)$ are H -vertices. By the choice of w at least $1 - O(\eta)$ fraction of the geodesics in $\mathcal{U}(w)$ are in \mathcal{E}_4 , unobstructed for length $\lambda + \rho_2$, and hit an H -vertex (in \mathcal{V}_1) at length $\lambda + \rho_2 + 1$. Thus, in particular, the density of H -vertices on $I'_{\lambda+\rho_2+1}(w)$ is at least $1 - O(\eta)$.

Let $\tilde{H} = I_\lambda(w) \cap \mathcal{V}_1$. We consider the family $\mathcal{E}(w)$ of monotone geodesic segments “going up” length L' from points at height h_1 in $\text{Sh}(N(\phi^{-1}(\tilde{H}), O(\epsilon'R)) \cap H, \rho_1)$ and use the behavior of this family to derive a contradiction. We first modify $\mathcal{E}(w)$ by throwing away some bad parts of the set. This modification is unnecessary if we are assuming that $\phi|_{U_i}$ is within $O(\epsilon'R)$ of a b -standard map. We throw out any geodesic γ in $\mathcal{E}(w)$ whose intersection with $SL_2^1(H)$ has more than $100c_2^{\frac{1}{2}}$ of its measure outside $SL_2^1(H) \cap U_*$. By Lemma 4.1, this throws away at most $O(c_2^{\frac{1}{2}})$ of the geodesics in $\mathcal{E}(w)$. After this modifications, it follows that each geodesic in $\mathcal{E}(w)$ has ϵ' -monotone image on an initial segment of length at least $\Omega(\beta'R)$.

Note that $N(\phi^{-1}(\tilde{H}), O(\epsilon'R)) \cap H$ contains a set of large measure in H and that $(I_\lambda(w) \cup I'_{\lambda+\rho_2+1}(w)) \cap \mathcal{V}_1$ is contained in the $O(\epsilon'R)$ neighborhood of $\phi(H)$.

Since every geodesic in $\mathcal{E}(w)$ diverges linearly from H and the initial segments of all $\mathcal{E}(w)$ of length $\Omega(\beta'R) > \Omega(\epsilon'R)$ have ϵ' -monotone image for ϕ , we have that any quasi-geodesic in $\phi(\mathcal{E}(w))$ diverges linearly from $\phi(H)$ and in particular, never comes within $\Omega(\beta'R)$ of $\phi(H)$.

Let $Q_u \subset I'_\lambda(w)$ be the subset of vertices v such that all vertices on $I'_{\lambda+\rho_2+1}$ within $\frac{R}{100\kappa^3} - \rho_2 - 1$ are not in \mathcal{V}_1 . Since $\ell(\mathcal{V}_1^c \cap I'_{\lambda+\rho_2+1}) < O(\eta)\ell(I'_{\lambda+\rho_2+1})$, we have $\ell(Q_u) \ll O(\eta)\ell(I'_\lambda(w)) = O(\eta)\mathfrak{b}^\lambda$. Any quasi-geodesic in $\phi(\mathcal{E}(w))$ crossing $I'_\lambda(w)$ does so on Q_u .

Similarly let $Q_d = I_\lambda(w) \cap \mathcal{V}_1^c$ and note that $\ell(Q_d) \ll O(\eta)\ell(I_\lambda(w)) = O(\eta)\mathfrak{b}^\lambda$ and that any quasi-geodesic in $\phi(\mathcal{E}(w))$ crossing $I_\lambda(w)$ must cross it on Q_d .

Now as all quasi-geodesics in $\phi(\mathcal{E}(w))$ diverge linearly from $\phi(H)$, they must all eventually leave the box of size λ bounded by $I_\lambda(w)$ and $I'_\lambda(w)$. This implies that every quasi-geodesic in $\phi(\mathcal{E}(w))$ eventually crosses $Q_u \cup Q_d$ or that every geodesic in $\mathcal{E}(w)$ eventually crosses $\phi^{-1}(Q_u \cup Q_d)$. This is impossible by Lemma 3.7, since

$$\ell(\phi^{-1}(Q_u \cup Q_d)) \leq O(\eta)\ell(H), \quad (27)$$

and $c(\rho_1)O(\eta) \ll 1$. □

Before reading the proof for Sol, the reader should be sure to read §3.5.

Proof for Sol.

We need to modify the proof given above in two ways in order to avoid “escape off the sides” of the box of size λ . As this is a modification of the previous proof, we only sketch the necessary changes.

We choose w as in the proof for DL graphs. We remark that it is easy to see that w can be chosen away from the edge of $B[\lambda]$. This can be deduced from Proposition 5.23. We will assume that we have chosen such a w . It is also possible to work with w near the edge of the box, but that one use a more complicated definition of points deep in the shadow of horocycles.

As above we consider $\tilde{H} = I_\lambda(w) \cap \mathcal{V}_1$. We consider the family $\mathcal{E}(w)$ of monotone geodesics “going up” length L' from points at height h_1 in $\text{Sh}(N(\phi^{-1}(\tilde{H}), O(\epsilon'R)) \cap H, \rho_1)$ and use the behavior of this family to derive a contradiction. We first modify $\mathcal{E}(w)$ exactly as before. We now further modify $\mathcal{E}(w)$ to only include those geodesics whose images at the end of the initial segment are $\beta'R$ -deep in $B[\lambda]$. By this we mean that they are $\beta'R$ deep in the shadows of the top and bottom of $B(R)$. This subset still contains a large proportion of the original elements of $\mathcal{E}(w)$. Let $Q = Q_u \cup Q_d$. Then as before, we see that paths in $\mathcal{E}(w)$ can only come near the top and bottom of $B[\lambda]$ in $N(Q^c, \frac{\beta'R}{2\kappa})$.

We now apply the results of §3.5 with $\rho = \rho_1, D_1 = \epsilon'R, D_2 = \frac{\beta'R}{2\kappa}$ and $D_3 = \lambda$. By Lemma 3.12 if a path $\gamma \in \phi(\mathcal{E}(w))$ leaves the box, it must tangle with the union of the top and the bottom of the box. Since $\gamma \in N(Q^c, \frac{\beta'R}{2\kappa})$, Lemma 3.13 implies

$$\ell(Q_u \cup Q_d) = \ell(Q) \geq \omega \|\mathcal{E}(w)\|,$$

where ω depends only on κ and C . But we have $\|\mathcal{E}(w)\| \geq \omega' \ell(H)$. where ω' depends only on κ and C . This is a contradiction to (27), if η is sufficiently small. As before η can be made arbitrarily small by taking ϵ' and δ_0 sufficiently small. \square

5.5 Completing the proof of Theorem 2.1

Proof. By Theorem 1.3 ϕ^{-1} of any very favorable horocycle in $B(L')$ is within $O(\epsilon'R)$ error of a horocycle. Given $\hat{\theta} > 0$, Lemma 5.1 implies that, by choosing β'' and δ_0 small enough, that $1 - \hat{\theta}$ of the measure of $B(L')$ consists of points in the image of both a very favorable x -horocycle and a very favorable y -horocycle. By an argument from the proof of [EFW2, Lemma 4.11], this implies that ϕ^{-1} respects level sets of height to within $O(\epsilon'R)$ error.

From this, it is not hard to show that ϕ^{-1} of most vertical geodesics are weakly monotone. This is very similar to the proof of [EFW2, Lemma 6.5]. There are some

additional difficulties due to the fact that we only control the map on most of the measure, but these can be handled in a manner similar to the proofs of [EFW2, Lemma 5.10 and Corollary 5.12].

Once we know ϕ^{-1} of most vertical geodesics are weakly monotone, the conclusion of the theorem follows from as in the proof of [EFW2, Theorem 5.1] at the end of [EFW2, Section 5.4]. \square

References

- [dlH] de la Harpe, Pierre. *Topics in geometric group theory*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 2000.
- [EFW1] Eskin, Alex; Fisher, David; Whyte, Kevin. Quasi-isometries and rigidity of solvable groups, to appear *Pur. Appl. Math. Q.*
- [EFW2] Eskin, Alex; Fisher, David; Whyte, Kevin. Coarse Differentiation of Quasi-isometries I; spaces not quasi-isometric to Cayley graphs, in preparation.
- [FM] Farb, Benson; Mosher, Lee. Problems on the geometry of finitely generated solvable groups. *Crystallographic groups and their generalizations (Kortrijk, 1999)*, 121–134, Contemp. Math., 262, Amer. Math. Soc., Providence, RI, 2000.
- [GhdH] Sur les groupes hyperboliques d’après Mikhael Gromov. (French) [Hyperbolic groups in the theory of Mikhael Gromov] Papers from the Swiss Seminar on Hyperbolic Groups held in Bern, 1988. Edited by É. Ghys and P. de la Harpe. Progress in Mathematics, 83. Birkhuser Boston, Inc., Boston, MA, 1990.
- [MN] Letter from R.Moeller to W.Woess, 2001.
- [PPS] Yuval Peres, Gabor Pete, Ariel Scolnicov. Critical percolation on certain non-unimodular graphs, preprint.
- [Wo] Woess, Wolfgang. Lamplighters, Diestel-Leader graphs, random walks, and harmonic functions, *Combinatorics, Probability & Computing* 14 (2005) 415–433.
- [W] Wortman, Kevin. A finitely presented solvable group with small quasi-isometry group, preprint.

Department of Mathematics, University of Chicago, 5734 S. University Avenue, Chicago, Illinois 60637.

Department of Mathematics, Indiana University, Rawles Hall, Bloomington, IN, 47405.

Department of Mathematics, Statistics, & Computer Science, University of Illinois at Chicago 322 Science Engineering Offices (M/C 249), 851 S. Morgan Street Chicago, IL 60607-7045.