BECK, JENNIFER. Ph.D. Asynchronous Automaticity and Products of Hyperbolic Spaces. (2024)
Directed by Dr. Talia Fernós. 82 pp.

The main result of this thesis is a general procedure for constructing an asynchronous automatic structure for some finitely generated groups quasi-isometric to products of non-elementary hyperbolic spaces. An asynchronous automatic structure, in turn, can be used to represent the group computationally, by now-classical means which we describe in some detail. We refer to the structures at the heart of this procedure as factor-language systems, and give certain criteria which guarantee their existence. The particular criteria we describe enjoy an intriguing analogy with certain criteria of discreteness and reducibility in the theory of lattices in products of trees. Along the way, we explore the geometry of path systems, finite-state automata, regular languages, automatic relations, hyperbolic geometry, quasi-isometries, and HNN-extensions.

ASYNCHRONOUS AUTOMATICITY AND PRODUCTS OF HYPERBOLIC SPACES

by

Jennifer Beck

A Dissertation Submitted to
the Faculty of The Graduate School at
The University of North Carolina at Greensboro
in Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy

Greensboro

2024

Approved by

Dr. Talia Fernós,
Committee Chair

APPROVAL PAGE

This dissertation written by Jennifer Beck has been approved by the following committee of the Faculty of The Graduate School at The University of North Carolina at Greensboro.

Committee Chair ______________________________

Talia Fernós

Committee Members ______________________________

Michael Hull

____

Ian Leary

____

Clifford Smyth

____

Dan Yasaki

____

Date of Acceptance by Committee

____

Date of Final Oral Examination

ACKNOWLEDGMENTS

I would like to thank my advisor Talia Fernós, deeply and from the bottom of my heart, for their mentorship, guidance, encouragement, and seemingly boundless patience throughout my graduate studies. I would likely never have come this far in my mathematical journey without them, and for that I owe them a debt I can never truly repay.

It would require an inordinate amount of space to properly thank all the UNCG faculty whose personal impacts on my life propelled me through the long years of my higher education. I owe thanks to them all, and I hope these few sentences are enough to at least hint at the depth of my gratitude.

The author was partially supported by NSF grant DMS–2005640.

Table of Contents

List of Figures

Chapter 1: Introduction

1.1 Notation

In the sequel, the following notational conventions will be observed.

- $E : A^* \to G$ is the natural evaluation map to a group generated by the finite set A from the free monoid A^*.

- If $v \in A^*$ is a word in the free monoid A^*, $|v|$ denotes the word length of v. If $\vec{\gamma}$ is a rectifiable curve, $\ell(\vec{\gamma})$ denotes the arclength of $\vec{\gamma}$

- If $w \in A^*$ is a word in the free monoid A^*, we denote by $\vec{w}$ the length-parametrized path in the 'Cayley Graph' of A^* (i.e. a tree whose edges are labeled by elements of A, viewed as a metric space) which interpolates the prefixes of w. If $G = \langle A \rangle$ is a group, we let $E\vec{w}$ be the length-parametrized path in $\mathrm{Cay}(G, A)$ which interpolates the images of the prefixes of $\vec{w}$ under the evaluation map. That is, we extend the evaluation map to the edges of the Cayley Graph.

- Given a group $G = \langle A \rangle$, we denote by $|g|_A$ the length of a geodesic word over A representing g, and denote by $d_{G,A}$ the word metric on G induced by the generating set A (i.e. $d_{G,A}(g,h) = |h^{-1}g|_A$). We extend $d_{G,A}$ to the entire Cayley graph $\mathrm{Cay}(G,A)$ by the customary expedient of identifying each edge with a copy of the unit interval.

- We will occasionally assign a common constant C to multiple independent parameters whose only use is to be 'sufficiently large' (e.g. the constants of a quasi-isometry, of a quasi-geodesic, or of a fellow-traveling path system).

- For a product $X = X_1 \times \cdots \times X_n$ and element $x \in X$, we denote by $\pi_i(x)$ the i-th coordinate projection of x, and by σ_i^x the i-th 'coordinate factor at x', i.e. the set of all points in X which can be obtained from x by changing only its i-th coordinate. If a group is embedded in X, we omit the superscript in the case that x is the image of the identity.

- $N(r,x)$ denotes the **closed** neighborhood of radius r about a point x in a metric space. This notation is used in preference to the standard notation for closed balls, on account of the preponderance of other uses in this paper of the capital Latin letter 'B'.

- In Chapter 6, we use the notation $(G,A) \curvearrowright (X,o)$ to refer to the action of a group G *with a particular generating set* A on a metric space X *with distinguished basepoint* o (following convention, we refer to X as a 'pointed metric space' in this context). This is to reinforce the fact that we wish to envision a *specific quasi-isometry* from a *specific Cayley graph* of G into X, allowing us to refer to the orbit map explicitly and unambiguously.

1.2 Background

In recent decades, the family of **hyperbolic metric spaces** has seen successful application in many domains, particularly the theory of finitely generated groups. Developed by Mikhail Gromov in the 1980s [10, 11], the theory of hyperbolic metric spaces can be understood as generalizing the large-scale metric properties of metric trees and hyperbolic n-space. Key to the success of the theory is the naturalness of hyperbolic spaces when considering metric structures *coarsely*, that is, up to **quasi-isometry**. This is because the category of hyperbolic metric spaces is closed under quasi-isometry, and quasi-isometries of hyperbolic spaces induce true homeomorphisms of their boundaries - objects which in some sense encode the 'horizon' of the space, or, less figuratively, the long-term behavior of infinitely extended geodesics. Thus the large-scale properties of hyperbolic spaces are highly robust under quite loosely controlled perturbations, and this permits analysis of such spaces even when the details of their internal geometry are not fully understood.

It is natural to wish to extend this theory to include more complex spaces which are constructed on hyperbolic spaces in some way. This is a subject of ongoing research, with some recent attention being given to *products* of hyperbolic spaces (e.g. the hierarchically hyperbolic spaces of [1], the preservation of coarse median structures in products noted in [2] and exploited in [16], or the powerful Theorem K in [20] which we shall reference presently). While nontrivial products of unbounded hyperbolic spaces do lose many of the convenient properties of hyperbolicity, we will still witness hyperbolic behavior within the **coordinate factors**, i.e. the isometrically embedded copies of the factor spaces obtained by restricting the coordinates in all but one factor to a given point. If we can guarantee that these coordinate factors will be (coarsely) preserved by quasi-isometry, then we have a robust way of considering products of hyperbolic spaces in their 'natural environment' (that is, as coarse geometric objects), while preserving their product structures.

Quasi-isometries, however, do not always preserve the product structure in a product of hyperbolic spaces: an example of this is given by rotation in $\mathbb{R}^n$ when $n \geq 2$, since almost all rotations will map coordinate factors to subspaces at infinite Hausdorff distance from any coordinate factor. However, this issue turns out to be unique to products including an 'elementary' factor (i.e. a factor which is either bounded or quasi-isometric to $\mathbb{R}$), at least when considering spaces sufficiently homogeneous to be quasi-isometric to groups. It was shown by Bowditch [3], who credited the essence of the result to previous work of Kapovich, Kleiner, and Leeb [17], that quasi-isometries between products of so-called 'bushy' hyperbolic spaces must coarsely preserve the factors. Notably, the metric homogeneity of groups allows us to conclude that the factors in a product of hyperbolic spaces *quasi-isometric to a group* must be either elementary or bushy, and so the hypothesis in this case can be weakened to the stipulation that the factors be non-elementary. This result is strengthened significantly in Margolis's *Discretisable Quasi-Actions I: Topological Completions and Hyperbolicity* [20], by the following theorem:

Theorem 1.1 (Margolis)**.** *Let Γ be a finitely generated group quasi-isometric to $\prod_{i=1}^n X_i$, where each X_i is a cocompact proper non-elementary hyperbolic metric space. Then Γ acts geometrically on $\prod_{i=1}^n Y_i$, preserving the product structure, where each Y_i is quasi-isometric to X_i and is either a rank one symmetric space of non-compact type or a locally finite graph.*

This is an extremely useful result for many reasons, not least because it gives us permission to jump immediately from a quasi-isometry between a group and a hyperbolic product to a true geometric action of that group on a product with substantially simplified factors. As Margolis points out in a later theorem, recent work on **Helly groups** [7] allows us to conclude that, when the factors Y_i are graphs, the group Γ is **biautomatic**, a notion from algorithmic group theory which lies at the heart of the computational aspect of the present work.

The theory of automatic groups originates in the early 1990s with the monograph *Word Processing in Groups* by Epstein, Cannon, Holt, Levy, Paterson, and Thurston [8], and a more modern introduction can be found in *Groups, Languages, and Automata* by Holt, Rees, and Röver [13]. In terms of practical application, the purpose of an **automatic structure** on a group is to provide a computable *representation* of the group, where elements are represented by certain strings of characters from a finite alphabet of symbols, and the group operation is realized by an algorithm which takes those strings as input. The algorithm is required to be especially simple - in computational terms, it is based on **finite-state automata**. A finite-state automaton is a sort of stripped-down Turing machine (a formal model of computation of which we give an informal description in Chapter 4) which takes a string as input, reads it character by character with no memory other than its finite set of internal 'states', and outputs only whether the string is 'accepted' or not. For an automatic structure, the set of group elements will be represented by a set of such accepted strings called a **regular language**, and the group operation will also be realized by a collection of regular languages over a specially constructed alphabet described in Chapter 4. In that same chapter, we also give an overview of the classical result that joins the theory of automatic groups to geometric group theory: **the fellow-traveling condition**.

In Chapter 2, we will describe the properties of **path systems**, by which we mean functions which assign, to each pair of points in a geodesic space, some path connecting them. A fellow-traveling condition on a path system is a bound on the 'distance' between two paths (we may make this notion precise in multiple ways, as will be described in Chapter 2), as a function of the distance between the endpoints. When a path system satisfies a fellow-traveling condition, we will say it is **bounded**. It turns out, for reasons that will be described Chapter 4, that a finitely generated group G is automatic *if and only if* there exists a regular language $\mathcal{L}$ with the property that the path system which joins elements of G by paths which 'spell out' words in $\mathcal{L}$ is what we will call **synchronously bounded**. A parallel notion, using a different metric on paths, similarly identifies **asynchronous automaticity** (a generalization of an automatic structure which sacrifices the time complexity of the algorithm computing the group operation in favor of embracing a strictly larger class of groups) with the **asynchronous boundedness** of the path system induced by some regular language.

In Chapter 3, we give a brief introduction to quasi-isometry and the concept of a hyperbolic metric space. Among their many other useful features, hyperbolic metric spaces satisfy the property that any path system consisting of uniform-quality quasi-geodesics will be *asynchronously bounded* (a fact which follows easily from the famous **Morse Lemma**, in combination with a useful lemma proved in Chapter 2). Furthermore, we will prove in Lemma 3.3 that asynchronously bounded path systems can be 'transported across' quasi-isometries in a consistent manner. Combining these notions, we can conclude that any group quasi-isometric to a product of hyperbolic spaces must admit an asynchronously bounded path system: we begin with the path system in the product which 'hugs' coordinate factors one by one in some fixed order, then move this path system across the quasi-isometry to the group. At this point, the only obstacle between us an an asynchronously automatic structure is assuring that such a path system can be given by a **regular language**. We describe such an asynchronously automatic structure in detail in Chapter 5, where we introduce the notion of a **factor-language system**. A factor-language system consists of a group G with generating set A, a space X which decomposes as a product of non-elementary hyperbolic spaces, a quasi-isometry $\phi : G \to X$, and regular languages over the alphabet A whose images under ϕE 'shadow' coordinate factors in X. These components are shown in Chapter 5 to yield an asynchronously automatic structure on G.

The remainder of the discussion concerns conditions which guarantee the existence of 'factor languages' which correspond to the coordinate factors of X. One way to accomplish this is hinted at in the definition of the coordinate factor σ_i^x: it is a subspace of a product obtained by restricting all coordinates save one to a point, which suggests that we may gain some traction by looking at an appropriately path-conscious coarsening of point stabilizers for the action of the group on the factors. In Chapter 6, we show that a factor-language system can always be obtained when the quasi-isometry ϕ maps G into X in such a way that the paths with 'small' projections to any given factor are given by a regular language. This allows us to show the existence of factor languages when ϕ is the orbit map of a geometric action satisfying a condition we call **level-determinism**, which can be thought of as a special kind of *uniform discreteness* (in the topology of pointwise convergence) of the group G in the space of maps from A^* to the factors of X.

We go on in Chapter 7 to show the methods described in Chapters 5 and 6 in action, by using them to obtain asynchronous automatic structures on HNN-extensions of non-elementary hyperbolic groups with finite-index associated subgroups. Since hyperbolic groups are examples of automatic groups, this can be thought of as a geometric generalization of Theorem 2.2 in [12] in the case in which the base group is hyperbolic.

1.3 Notes on Theory and Computation

This thesis is intended to be of both theoretical and computational interest, and these two approaches are intertwined throughout the chapters to come. All of the material we discuss will come into play when we prove the main theorems of chapters 5 and 6, but it would be mostly accurate to say that chapters 2 and 3 are of principally theoretical interest, while chapters 4 and 7 are mostly computational in nature.

The main computational result in the following is the construction of an explicit *asynchronous automatic* structure for a large family of groups in Chapter 5, along with a condition (see Corollary 6.5 in Chapter 6) which can be used to find the components of this structure for certain groups. *Synchronously* automatic groups are already well-implemented in Magma, with most of the computational machinery being dedicated to *finding* an automatic structure for a group given only its presentation. The main obstruction to similarly implementing *asynchronously* automatic structures in a useful way is that these quick methods for finding the structure cannot be used. Once such a structure is found, however, the process of using it to perform calculations in the group is well-known and described in [8]. In presenting a blueprint for converting geometric data about certain groups into asynchronous automatic structures explicitly, we provide the means to skip the search step and begin using the existing machinery to perform computations. That said, algorithms using asynchronous as opposed to synchronous automata are computationally costly, and will require much better optimization to be of practical use. The word problem for asynchronously automatic groups, for example, requires exponential time to solve using current methods, and it is not known whether a polynomial-time algorithm exists at all.

The main theoretical results are found in Chapter 6 and concern the above-mentioned observation about obtaining a factor-language system when the action of a group on the hyperbolic factors of a product is level-deterministic, which we noted to be a kind of discreteness condition. This bears an interesting analogy to a well-known fact about lattices in products of trees, namely that such lattices are **reducible** (i.e. virtually split as a direct product of lattices) when their projections to the factors are **discrete** (see e.g. [6]). In the topology of pointwise convergence on the space of maps between two metric spaces, a set of maps is discrete precisely when its elements can be identified by their behavior on some ball of finite radius. If the same radius can be used for all maps in the set, we might call such a set *uniformly* discrete. Level-determinism is a condition akin to uniform discreteness, but weakens the notion to require only that the image of a map on a ball of fixed radius determine the image of that map on a *slightly larger bounded set* in the domain. Then, instead of concluding that the group in question virtually splits as a direct product of *groups*, we get instead that the group decomposes (in a way that respects the underlying geometric product) as a product of *rational subsets*, which are themselves images (under the evaluation map) of *regular languages* in the free monoid generated by the group generators. This is curious, because it allows us to extract a kind of reducibility from groups which may not admit a nontrivial product structure in the group-theoretic sense.

Chapter 2: Bounded Path Systems

In this chapter, we introduce the properties of **path systems**, a notion which we adapt from [21]. A path system on a path-connected space is simply a function which chooses a single path connecting pairs of points in the space. We begin our discussion by defining three distinct metrics on space of all paths in a geodesic space, and then go on to classify certain path systems by their behavior under these respective metrics. Most of the technical work in this chapter occurs in Section 2.2, where we prove a useful lemma that will allow us to pass freely between two of these metrics when the paths in the system under examination satisfy a coarse 'no large loops' property.

2.1 Metrics on the Space of Paths

Let X be a geodesic metric space, and denote by $\vec{X}$ the space of rectifiable curves on X. For ease of notation, we shall think of curves $\vec{\gamma} \in \vec{X}$ as being parametrized by arclength on $[0, \ell(\vec{\gamma})]$ (where $\ell(\vec{\gamma})$ is the length of the curve $\vec{\gamma}$) and satisfying $\vec{\gamma}(t) = \vec{\gamma}(\ell(\vec{\gamma}))$ for $t \geq \ell(\vec{\gamma})$ or $t = \infty$. We may metrize $\vec{X}$ in three relevant ways:

Definition 2.1 (Metrics on $\vec{X}$)**.** *Denote the metric on X by d, and let $\vec{\gamma}_1, \vec{\gamma}_2 \in \vec{X}$.*

- *The* ***synchronous distance*** *between* $\vec{\gamma}_1$ *and* $\vec{\gamma}_2$*, denoted* d_{sync}*, is*

$$d_{sync}(\vec{\gamma}_1, \vec{\gamma}_2) = \sup_t d(\vec{\gamma}_1(t), \vec{\gamma}_2(t))$$

- *The* ***asynchronous distance*** *between* $\vec{\gamma}_1$ *and* $\vec{\gamma}_2$*, denoted* d_{Async}*, is*

$$d_{Async}(\vec{\gamma}_1, \vec{\gamma}_2) = \inf_{\rho,\rho'} \sup_t d(\vec{\gamma}_1(\rho(t)), \vec{\gamma}_2(\rho'(t)))$$

  where ρ *and* ρ' *range over monotone surjections* $[0, \infty) \to [0, \infty)$*.*

- *The* ***Hausdorff (pseudo-) distance*** *between* $\vec{\gamma}_1$ *and* $\vec{\gamma}_2$*, denoted* d_{Haus}*, is the usual Hausdorff distance between the images* $\vec{\gamma}_1([0, \infty))$*,* $\vec{\gamma}_2([0, \infty))$ *in* X*.*

The Hausdorff distance is a well-known metric on compact subsets of metric spaces, though it fails to be a metric on the space of all *paths* since it can return a distance of 0 between two distinct paths (e.g. two copies of the same path, but with initial and terminal points reversed). If we choose one orientation for each path, and consider only paths which have finitely many self-intersections, then the Hausdorff distance does yield a metric. The other two distances can be easily verified to give true metrics on $\vec{X}$. The synchronous distance can be thought of as the maximal separation between two travelers moving along two given paths at the same speed, while the asynchronous distance is the optimal maximum separation between two such travelers when they are not permitted to backtrack. Note that these definitions are meaningful for any maps $f : [0, \infty) \to X$, and so we can talk about the 'asynchronous distance' between any pair of functions $f_1, f_2 : \mathbb{R} \to X$, but we do need further assumptions (e.g. that the maps be arclength-parametrized curves as above) to conclude that d_{sync} and d_{Async} give metrics and not pseudometrics.

For our purposes, we will consider subsets of $\overrightarrow{X}$ which assign a unique path to a pair of endpoints in X. In particular, we will often wish to fix a basepoint and consider collections of paths having this basepoint as their initial point. The following definition generalizes the notion of a path system on a graph introduced in [21] to arbitrary geodesic spaces:

Definition 2.2 (Path System)**.** *A* ***path system*** *on X is a partial map $P : X^2 \to \vec{X}$ which is either a total map or has domain of the form $\{b\} \times X$ for some $b \in X$, with the property that $\vec{\gamma} = P(x, y)$ satisfies $\vec{\gamma}(0) = x$ and $\vec{\gamma}(\infty) = y$ for all (x, y) in the domain of P. We say P is* ***two-sided*** *if P is total, and we say P is* ***one-sided*** *if the domain of P is of the form $\{b\} \times X$.*

We will not worry too much about the distinction between one-sidedness and two-sidedness, since a one-sided path system in a group can always be extended to a two-sided path system by translation, and a two-sided path system can always be restricted to a one-sided path system by a choice of basepoint.

In the theory of automatic groups, certain algorithmic objects associated to a finitely generated group G (namely, word-difference automata) are shown to exist if and only if an associated path system in the Cayley Graph of G satisfies a *fellow-traveling* condition [8]. We will initially distinguish fellow-traveling conditions both by the metrics used to characterize them and the functions which bound them. Then we will show that, for geodesic spaces, we may assume the bounding function is linear.

Definition 2.3 (f-Bounded Path System)**.** *A path system P is* ***synchronously f-bounded*** *if $f : \mathbb{R} \to \mathbb{R}$ is a nondecreasing function such that $f(x) \geq x$ for all $x \in \mathbb{R}$,* *and*

$$d_{sync}(\vec{\gamma}_1, \vec{\gamma}_2) \leq f(D)$$

for all $\vec{\gamma}_1, \vec{\gamma}_2$ *in the image of* P*, where*

$$D = \max\{d(\vec{\gamma}_1(0), \vec{\gamma}_2(0)), d(\vec{\gamma}_1(\infty), \vec{\gamma}_2(\infty))\}$$

Substituting the corresponding metric for d_{sync}*, we similarly define* ***asynchronously*** f***-bounded*** *and* ***Hausdorff*** f***-bounded****.*

Note that the condition that $f(x) \geq x$ for all x is satisfied automatically if f bounds a path system synchronously or asynchronously, since, defining D as above, we have

$$f(D) \geq d_{sync}(\vec{\gamma}_1, \vec{\gamma}_2) \geq d_{Async}(\vec{\gamma}_1, \vec{\gamma}_2) \geq D$$

as a consequence of the definitions of d_{sync} and d_{Async}. We include this condition in the definition to exclude pathological choices of f for Hausdorff-bounded path systems.

We will typically care about bounding functions only up to coarse Lipschitz equivalence. Consequently, if a path system has a linear bounding function, we will omit the function and simply say the path system is **bounded**:

Definition 2.4. *A path system* P *is synchronously (asynchronously, Hausdorff) bounded if it is synchronously (asynchronously, Hausdorff)* f*-bounded for some linear function* f*.*

As was mentioned previously (and will be proved shortly), we can always assume f is linear if X is a geodesic space. Given $P(x, y)$ and $P(x', y')$ for an f-bounded path system P, we examine the sequence of paths

$$P_n := P(\overrightarrow{\gamma}_{xx'}(nC), \overrightarrow{\gamma}_{yy'}(nC))$$

where $C > 0$, n ranges from 0 to $N = \lceil \frac{\max\{d(x,x'),d(y,y')\}}{C} \rceil$, and $\overrightarrow{\gamma}_{xx'}$ and $\overrightarrow{\gamma}_{yy'}$ are geodesics which join x to x' and y to y', respectively. We have

$$\begin{aligned} d_{sync}(P_{i-1}, P_i) &\leq f(\max\{d(\overrightarrow{\gamma}_{xx'}((i-1)C), \overrightarrow{\gamma}_{xx'}(iC)), d(\overrightarrow{\gamma}_{yy'}((i-1)C), \overrightarrow{\gamma}_{yy'}(iC))\}) \\ &\leq f(C) \end{aligned}$$

for all $i = 1, \cdots, N$, since P is f-bounded. Hence,

$$\begin{aligned} d_{sync}(P(x,y), P(y,y')) &\leq f(C)N \\ &\leq f(C)(\frac{1}{C}\max\{d(x,x'), d(y,y')\} + 1) \end{aligned}$$

Defining the linear function $\bar{f}(d) = \frac{f(C)}{C}d + f(C)$, this shows P is in fact $\bar{f}$-bounded, and the same argument holds for asynchronously- or Hausdorff-bounded path systems. In other words,

Lemma 2.5. *Let X be a geodesic space with a path system P. Then P is bounded if and only if it is f-bounded for some function f.*

2.2 Hausdorff-Bounded Path Systems are Asynchronously Bounded

In this section, we prove a lemma that will greatly expedite the proofs in Section 5. In that section, we use the Morse lemma for quasigeodesics in hyperbolic spaces to show that a path system is asynchronously bounded. The conclusion of the Morse lemma, however, only bounds the Hausdorff distance between quasigeodesic paths. It is shown in [8] that, in finitely-generated groups, path systems given by regular normal forms asynchronously fellow-travel if and only if they are Hausdorff-bounded and admit what the authors refer to as a **departure function**. Their proof uses the automata recognizing the normal form in a fairly essential way - here, we prove the more general fact that *any* Hausdorff-bounded path system on a geodesic space is asynchronously bounded if it admits a departure function (our definition of a departure function is generalized to this broader case).

To begin with, we define an object associated to a pair of paths in a geodesic space which, motivated by visual metaphor, we call a **ladder**:

Definition 2.6 (Ladder)**.** *Let X be a geodesic space, let $\alpha, m, n > 0$ be given, and let $\vec{\gamma_1}, \vec{\gamma_2} \in \vec{X}$. An (α, m, n)-**ladder** for $(\vec{\gamma_1}, \vec{\gamma_2})$ is a finite sequence $(s_i, t_i)_{i=0}^{K} \subset \mathbb{R}^2$ such that*

1. $s_0 = t_0 = 0$
2. $s_K \leq \ell(\vec{\gamma_1})$ *and* $t_K \leq \ell(\vec{\gamma_2})$
3. $s_{i+1} - s_i \in [0, m]$ *and* $t_{i+1} - t_i \in [0, n]$*, and either* $s_{i+1} \neq s_i$ *or* $t_{i+1} \neq t_i$*, for all* $i \in [0, K]$

4. $d(\overrightarrow{\gamma_1}(s_i), \overrightarrow{\gamma_2}(t_i)) \leq \alpha$ *for all* $i \in [0, K]$

We will use the existence of certain ladders to prove Lemma 2.11. The following lemma shows that the existence of a ladder bounds the asynchronous distance between the associated paths:

Lemma 2.7. *Let* X *be a geodesic space, and let* $\overrightarrow{\gamma_1}, \overrightarrow{\gamma_2} \in \vec{X}$*. If there exists an* (α, m, n)*-ladder* $(s_i, t_i)_{i=0}^K$ *for* $(\overrightarrow{\gamma_1}, \overrightarrow{\gamma_2})$ *with* $s_K = \ell(\overrightarrow{\gamma_1})$ *and* $t_K = \ell(\overrightarrow{\gamma_2})$*, then*

$$d_{Async}(\overrightarrow{\gamma_1}, \overrightarrow{\gamma_2}) \leq \frac{1}{2}\max\{m, n\} + \min\{m, n\} + \alpha$$

Proof:

Choose any ρ, ρ' such that $\rho(i) = s_i$ and $\rho'(i) = t_i$ for all i (this is always possible since the sequences $(s_i)_{i=0}^K$ and $(t_i)_{i=0}^K$ are monotone by the third condition of Definition 2.6). Let $x \in (0, K)$ be given. Since $(s_i, t_i)_{i=0}^K$ is an (α, m, n)-ladder, we know that $\overrightarrow{\gamma_1}(\rho(x))$ lies on an arc of $\overrightarrow{\gamma_1}$ between $\overrightarrow{\gamma_1}(s_{\lfloor x \rfloor})$ and $\overrightarrow{\gamma_1}(s_{\lfloor x \rfloor + 1})$ whose length does not exceed m. Similarly, we see that $\overrightarrow{\gamma_2}(\rho'(x))$ lies on an arc of $\overrightarrow{\gamma_2}$ between $\overrightarrow{\gamma_2}(t_{\lfloor x \rfloor})$ and $\overrightarrow{\gamma_2}(t_{\lfloor x \rfloor + 1})$ whose length does not exceed n. Assume without loss of generality that $m \geq n$, and that the arc of $\overrightarrow{\gamma_1}$ from $\overrightarrow{\gamma_1}(\rho(x))$ to $\overrightarrow{\gamma_1}(s_{\lfloor x \rfloor})$ has length $\leq \frac{1}{2}m$. By assumption, there is some geodesic of length $\leq \alpha$ from $\overrightarrow{\gamma_1}(s_{\lfloor x \rfloor})$ to $\overrightarrow{\gamma_2}(t_{\lfloor x \rfloor})$, and the arc of $\overrightarrow{\gamma_2}$ connecting $\overrightarrow{\gamma_1}(s_{\lfloor x \rfloor})$ to $\overrightarrow{\gamma_2}(\rho'(x))$ has length $\leq n$. Concatenating these three curves gives a path of length $\leq \frac{1}{2}m + n + \alpha = \frac{1}{2}\max\{m, n\} + \min\{m, n\} + \alpha$ between $\overrightarrow{\gamma_1}(\rho(x))$ and $\overrightarrow{\gamma_2}(\rho'(x))$, and so

$$d(\overrightarrow{\gamma_1}(\rho(x)), \overrightarrow{\gamma_2}(\rho'(x))) \leq \frac{1}{2}\max\{m, n\} + \min\{m, n\} + \alpha$$

for all $x \in (0, K)$. If $x = 0$ or $x \geq K$, then $d(\overrightarrow{\gamma_1}(\rho(x)), \overrightarrow{\gamma_2}(\rho'(x))) \leq \alpha$ since $(\rho(x), \rho'(x))$ is a point on the ladder. The lemma is thus proved in the case that $\max\{m, n\} = m$ and $\rho(x) - s_{\lfloor x \rfloor} \leq s_{\lfloor x \rfloor + 1} - \rho(x)$, and this proof may be applied similarly to the remaining cases. □

A converse to this lemma is easy to show: if two paths asynchronously fellow-travel, then any evenly-spaced sequence of points along them can be used as the 'rungs' of a ladder:

Lemma 2.8. *Let X be a geodesic space with $\overrightarrow{\gamma_1}, \overrightarrow{\gamma_2} \in \vec{X}$, and let $\alpha, \beta > 0$. If $d_{Async}(\overrightarrow{\gamma_1}, \overrightarrow{\gamma_2}) \leq \alpha$, then $(\overrightarrow{\gamma_1}, \overrightarrow{\gamma_2})$ admits a $(\beta, \beta, \alpha + \beta)$-ladder $(s_i, t_i)_{i=0}^{K}$ such that $s_i, t_i \in \beta\mathbb{Z}$ for all $i \in \{0, \cdots, K\}$, $s_K = \ell(\overrightarrow{\gamma_1})$, and $t_K = \ell(\overrightarrow{\gamma_2})$.*

Proof: Choose monotone surjections ρ_1 and ρ_2 such that $d(\overrightarrow{\gamma_1}(\rho_1(t)), \overrightarrow{\gamma_2}(\rho_2(t))) < \beta$ for all $t \geq 0$, and let $K = \lceil \frac{\max\{\ell(\overrightarrow{\gamma_1}), \ell(\overrightarrow{\gamma_2})\}}{\beta} \rceil$. Then for all $i \in \{0, \cdots, K\}$, we let $s_i = i\beta$. There must exist $T_i \geq 0$ so that $\rho_1(T_i) = s_i$, and so

$$d(\overrightarrow{\gamma_1}(s_i), \overrightarrow{\gamma_2}(\rho_2(T_i))) < \alpha$$

Furthermore, there must be $J_i \in \{0, \cdots, K\}$ such that $|J_i\beta - \rho_2(T_i)| < \beta$, and hence $d(\overrightarrow{\gamma_2}(\rho_2(T_i)), \overrightarrow{\gamma_2}(J_i\beta)) < \beta$. Setting $t_i = J_i\beta$, the triangle inequality yields

$$d(\overrightarrow{\gamma_1}(s_i), \overrightarrow{\gamma_2}(t_i)) < \alpha + \beta$$

for all $i \in \{0, \cdots, K\}$. A simple induction shows that we can choose the t_i to be monotone, and we can collapse any subsequences with $(s_i, t_i) = (s_{i+1}, t_{i+1})$ to obtain the desired $(\beta, \beta, \alpha + \beta)$-ladder. □

We will prove Lemma 2.11 by appealing to a particular partial order on the set of ladders for a pair of paths. The following lemma shows that this partially ordered set contains maximal elements, which we need for our proof:

Lemma 2.9. *Let X be a geodesic space, and take $\vec{\gamma}_1, \vec{\gamma}_2 \in \vec{X}$ and $\alpha, m, n \geq 0$. Denote by H the set of all (α, m, n)-ladders on $(\vec{\gamma}_1, \vec{\gamma}_2)$, and define a partial order $\leq_H$ on H so that $(s_i, t_i)_{i=0}^K \leq_H (s'_i, t'_i)_{i=0}^{K'}$ if and only if $s_K \leq s'_{K'}$ and $t_K \leq t'_{K'}$. Then $(H, \leq_H)$ has a maximal element.*

Proof:

We show that every chain in $(H, \leq_H)$ has an upper bound, and the conclusion follows by Zorn's lemma. Let C be a chain in $(H, \leq_H)$. If C is a finite set it has a maximum, so suppose C is infinite. We may project the elements of H to the poset $(\mathbb{R}^2, \leq_{\mathbb{R}^2})$ (where $(a, b) \leq_{\mathbb{R}^2} (c, d) \Leftrightarrow a \leq_{\mathbb{R}} c$ and $b \leq_{\mathbb{R}} d$) via $\pi : (s_i, t_i)_{i=0}^K \mapsto (s_K, t_K)$, and we see that $\pi(C)$ is also a chain and has a least upper bound $(x_C, y_C) \leq_{\mathbb{R}^2} (\ell(\vec{\gamma}_1), \ell(\vec{\gamma}_2))$. We now construct an (α, m, n)-ladder $(s_i, t_i)_{i=0}^K$ with $(s_K, t_K) = (x_C, y_C)$:

Let $((s_i^j, t_i^j)_{i=0}^{K_j})_{j=0}^\infty$ be a sequence in C whose image under π converges in $\mathbb{R}^2$ to (x_C, y_C), i.e. $\lim_{j\to\infty}(s_{K_j}^j, t_{K_j}^j) = (x_C, y_C)$. Since $d(\vec{\gamma}_1(s_i^j), \vec{\gamma}_2(t_i^j)) \leq \alpha$ for all $i, j \geq 0$, we have $d(\vec{\gamma}_1(x_C), \vec{\gamma}_2(y_C)) \leq \alpha$ by continuity. Now choose j so that $x_C - s_{K_j} \in [0, m]$ and $y_C - t_{K_j} \in [0, n]$, and define a new ladder $(s_i, t_i)_{i=0}^{K_j+1}$ so that $(s_i, t_i) = (s_i^j, t_i^j)$ for $0 \leq i \leq K_j$ and $(s_{K_j+1}, t_{K_j+1}) = (x_C, y_C)$. This satisfies the definition of an (α, m, n)-ladder, and is an upper bound for C. $\square$

Finally, we define what we mean by a 'departure function' for a path system on a geodesic space:

Definition 2.10 (Departure Function)**.** *Let $P : X^2 \to \vec{X}$ be a path system on X. We say a function $D : \mathbb{R} \to \mathbb{R}$ is a **departure function** for P if, for all $\vec{\gamma}$ in the image of P and all $M, s, t > 0$, we have $d(\vec{\gamma}(s), \vec{\gamma}(s+t)) > M$ whenever $t > D(M)$ and $s + D(M) \leq \ell(\vec{\gamma})$*

More succinctly, D is a departure function for a path system if no path in the system contains a subpath of length greater than $D(M)$ whose endpoints are within distance M of each other. In other words, D controls the extent to which paths in the system can 'loop back' on themselves. We shall assume, without loss of generality, that D is strictly increasing and continuous.

We can now prove the desired result: Hausdorff-bounded path systems are asynchronously bounded.

Lemma 2.11. *Let X be a geodesic space, and let P be a path system on X. If P is Hausdorff f-bounded and has departure function D, then P is asynchronously $(\frac{1}{2}D(2f + D(2f)) + D(2f) + f)$-bounded.*

Proof:

Let $\vec{\gamma}_1, \vec{\gamma}_2 \in P$, and let $\alpha = f(\max\{d(\vec{\gamma}_1(0), \vec{\gamma}_2(0)), d(\vec{\gamma}_1(\infty), \vec{\gamma}_2(\infty))\})$, so that $d_{Haus}(\vec{\gamma}_1, \vec{\gamma}_2) \leq \alpha$ by assumption. We wish to show that $d_{Async}(\vec{\gamma}_1, \vec{\gamma}_2) \leq (\frac{1}{2}D(2\alpha + D(2\alpha)) + D(2\alpha) + \alpha)$. By Lemma 2.7, this will follow if we can show that $(\vec{\gamma}_1, \vec{\gamma}_2)$ admits an $(\alpha, D(2\alpha + D(2\alpha)), D(2\alpha))$-ladder $(s_i, t_i)_{i=0}^K$ such that $s_K = \ell(\vec{\gamma}_1)$ and $t_K = \ell(\vec{\gamma}_2)$.

Let H be the set of $(\alpha, D(2\alpha + D(2\alpha)), D(2\alpha))$-ladders on $(\vec{\gamma}_1, \vec{\gamma}_2)$. Since we have assumed $f(x) \geq x$ for all x, we know $\alpha \geq d(\vec{\gamma}_1(0), \vec{\gamma}_2(0))$. This means the single-element sequence $(0,0)$ is an $(\alpha, D(2\alpha + D(2\alpha), D(2\alpha))$-ladder, and H is nonempty. Since H is nonempty, it contains a $\leq_H$-maximal element by Lemma 2.9. Let $(s_i, t_i)_{i=0}^K$ be such a maximal ladder.

We wish to show that $(s_i, t_i)_{i=0}^K$ satisfies $(s_K, t_K) = (\ell(\vec{\gamma}_1), \ell(\vec{\gamma}_2))$. By way of contradiction, suppose that this is false - i.e. we have either $s_K < \ell(\vec{\gamma}_1)$ or $t_K < \ell(\vec{\gamma}_2)$. We will treat the case in which $\ell(\vec{\gamma}_1) - s_K > D(2\alpha + D(2\alpha))$ in detail, then briefly address the small differences in the proof of the complementary case, which proceeds along nearly the same lines.

We will refer to a pair $(x, y) \in [0, \ell(\vec{\gamma}_1)] \times [0, \ell(\vec{\gamma}_2)]$ as **Hausdorff partners** if $d(\vec{\gamma}_1(x), \vec{\gamma}_2(y)) \leq \alpha$. Note that the pairs (s_i, t_i) in our maximal ladder are all Hausdorff partners. For ease of notation, we will label the following intervals (keeping track of which arcs of $\vec{\gamma}_1$ and $\vec{\gamma}_2$ they correspond to):

- $I = (s_K, s_K + D(2\alpha + D(2\alpha))]$ (a subpath of $\vec{\gamma}_1$)
- $A = [0, t_K)$ (a subpath of $\vec{\gamma}_2$)
- $B = [t_K, t_K + D(2\alpha))$ (a subpath of $\vec{\gamma}_2$)
- $C = [t_K + D(2\alpha), \ell(\vec{\gamma}_2)]$ (a subpath of $\vec{\gamma}_2$)

These intervals are motivated by the following illustration, which may be used as a visual aid for the ensuing proof:

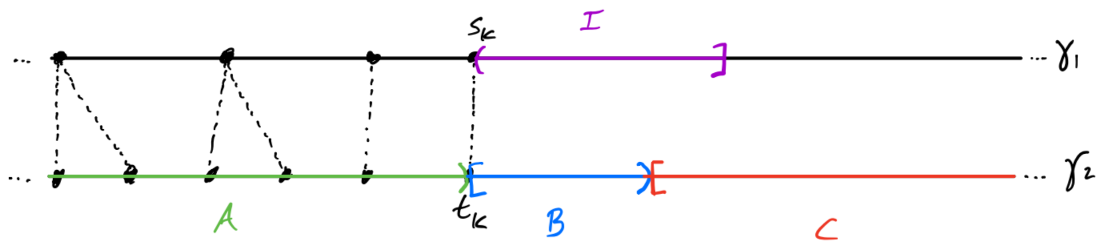

Figure 2.1. Two possible paths, with a maximal ladder and labeled intervals. Hausdorff partners are joined by dotted lines.

We obtain the desired contradiction if we can produce a pair of Hausdorff partners $(s,t) \in \overline{I \times B} \setminus \{(s_K, t_K)\}$, since we can append this pair to our maximal ladder $(s_i, t_i)_{i=0}^K$ and obtain a strictly greater ladder under the order $\leq_H$.

Since $\vec{\gamma}_1$ and $\vec{\gamma}_2$ are Hausdorff α-close, we know that every point on $\vec{\gamma}_1$ has some Hausdorff partner on $\vec{\gamma}_2$, and vice-versa. Thus, the following four cases are exhaustive:

1. Some point in I has a Hausdorff partner in B, or
2. Every Hausdorff partner of every point in I lies in C, or
3. Every Hausdorff partner of every point in I lies in A, or
4. There exists both a point of I with a Hausdorff partner in A and a point of I with a Hausdorff partner in C.

We now show that, in each of these three cases, we can 'advance' our purportedly maximal ladder to create a strictly greater one.

- **Case 1:** There exists $s \in I$ such that s has a Hausdorff partner $t \in B$. As mentioned above, this produces the desired contradiction immediately.

- **Case 2:** For every $s \in I$, every Hausdorff partner t of s lies in C. Writing $s = s_K + \epsilon$, we have $d(\vec{\gamma}_1(s), \vec{\gamma}_1(s_K)) \leq \epsilon$, and so

$$\begin{aligned} d(\vec{\gamma}_2(t), \vec{\gamma}_2(t_K)) &\leq d(\vec{\gamma}_2(t), \vec{\gamma}_1(s)) + d(\vec{\gamma}_1(s), \vec{\gamma}_1(s_K)) + d(\vec{\gamma}_1(s_K), \vec{\gamma}_2(t_K)) \\ &\leq 2\alpha + \epsilon \end{aligned}$$

Since D is a departure function for $\vec{\gamma}_2$, this means that $t - t_K \leq D(2\alpha + \epsilon)$. Furthermore, since we have assumed that $t \in C$, we also have $t - t_K \geq D(2\alpha)$. Letting $\epsilon \to 0$, we conclude that $(s_K, t_K + D(2\alpha))$ must be Hausdorff partners. We may append $(s_K, t_K + D(2\alpha))$ to $(s_i, t_i)_{i=0}^K$ to obtain a ladder in H strictly greater than our maximal ladder: a contradiction.

- **Case 3:** This case is impossible, since $s_K + D(2\alpha + D(2\alpha))$ must have a Hausdorff partner in C. To see this, let $s > s_K + D(2\alpha + D(2\alpha))$ be given. If s has a Hausdorff partner $t \in A$, then by the definition of a ladder there exists a point (s_i, t_i) on our maximal ladder such that $d(\vec{\gamma}_2(t), \vec{\gamma}_2(t_i)) \leq d(2\alpha)$. Then we have

$$\begin{aligned} d(\vec{\gamma}_1(s), \vec{\gamma}_1(s_i)) &\leq d(\vec{\gamma}_1(s), \vec{\gamma}_2(t)) + d(\vec{\gamma}_2(t), \vec{\gamma}_2(t_i)) + d(\vec{\gamma}_2(t_i), \vec{\gamma}_2(s_i)) \\ &\leq 2\alpha + D(2\alpha) \end{aligned}$$

By the definition of a departure function, we must have $s - s_i \leq D(2\alpha + D(2\alpha))$. However, we have assumed that $s > s_K + D(2\alpha + D(2\alpha))$, and so $s - s_i > D(2\alpha + D(2\alpha))$, since $s_i \leq s_K$. We conclude that s cannot have a Hausdorff partner in A.

Now we take any sequence $(s_j)_{j=0}^{\infty}$ converging to $s_K + D(2\alpha + D(2\alpha))$ from above.

By assumption, there is a corresponding sequence

$$(t_j)_{j=0}^{\infty} \subseteq B \cup C$$

such that (s_j, t_j) are Hausdorff partners for all j. Passing to a converging subsequence of $(t_j)_{j=0}^{\infty}$, we obtain a Hausdorff partner for $s_K + D(2\alpha + D(2\alpha))$ which also lies outside A, since $B \cup C$ is closed.

- **Case 4:** There exist Hausdorff partners $(s, t) \in I \times A$ and $(s', t') \in I \times C$. Let s'' be the supremum of points in I with Hausdorff partners in A. By approaching s'' from below and passing to a subsequence as in the preceding case, we conclude that s'' has a Hausdorff partner $t_A \in A$. Furthermore, since every $s > s''$ has a Hausdorff partner in C, we obtain a Hausdorff partner $t_C \in C$ for s''. Thus s'' has a Hausdorff partner $t_A < t_K$ and a Hausdorff partner $t_C \geq t_K + D(2\alpha)$. Now we have

$$\begin{aligned} d(\vec{\gamma}_2(t_A), \vec{\gamma}_2(t_C)) &\leq d(\vec{\gamma}_2(t_A), \vec{\gamma}_1(s'')) + d(\vec{\gamma}_1(s''), \vec{\gamma}_2(t_C)) \\ &\leq 2\alpha \end{aligned}$$

  Since $t_C - t_A > D(2\alpha)$, this is a contradiction.

Having exhausted all possible cases, we conclude that the maximal ladder $(s_i, t_i)_{i=0}^{K}$ must satisfy $\ell(\vec{\gamma}_1) - s_K \in [0, D(2\alpha + D(2\alpha))]$. If $\ell(\vec{\gamma}_2) - t_K > D(2\alpha)$, we may repeat the cases above almost identically to obtain the same contradiction: either s_K and t_K can be advanced by a small amount (Case 1), $d(\vec{\gamma}_1(s_K), \vec{\gamma}_2(t_K + D(2\alpha))) \leq \alpha$ (Case 2), or we can apply the fact that $(\ell(\vec{\gamma}_1), \ell(\vec{\gamma}_2))$ is a pair of Haudsorff partners to find $s \in (s_K, \ell(\vec{\gamma}_1)]$ with two Hausdorff partners t, t' such that $t \leq t_K \leq t'$, $t' - t_K \leq D(2\alpha)$, and the ladder can be extended to (s, t'). We conclude that $\ell(\overrightarrow{\gamma}_2) - t_K \leq D(2\alpha)$, and note that our ladder can now be extended to include $(\ell(\overrightarrow{\gamma}_1), \ell(\overrightarrow{\gamma}_2))$. Thus a maximal ladder in H must satisfy $s_K = \ell(\overrightarrow{\gamma}_1)$ and $t_K = \ell(\overrightarrow{\gamma}_2)$. $\square$

Chapter 3: Quasi-Isometry and Hyperbolicity

This section provides a minimal introduction to quasi-isometry, hyperbolicity, and metric products. Most of the following material is elementary in the theory of hyperbolic geodesic spaces, a comprehensive introduction to which may be found in [4], with the exception of Sections 3.2 (which contains observations about the behavior of bounded path systems under quasi-isometry) and 3.5 (which contains elementary facts about metric products and quotes a theorem of Bowditch).

3.1 Quasi-Isometry

We begin by defining the notion of a **quasi-isometry**, a coarse generalization of isometry which permits the metric to be perturbed within fixed linear bounds:

Definition 3.1. *Let (X, d_X) and (Y, d_Y) be metric spaces, and let $\lambda, c \geq 0$ be given. We say a map $\phi : X \to Y$ is a $\boldsymbol{(\lambda, c)}$ **quasi-isometry** if the following hold:*

1. $\frac{1}{\lambda} d_Y(\phi(x), \phi(x')) - c \leq d_X(x, x') \leq \lambda d_Y(\phi(x), \phi_Y(x')) + c$ *for all* $x, x' \in X$.

2. *For all* $y \in Y$*, there exists* $x \in X$ *such that* $d_Y(\phi(x), y) < c$

Note that we are adopting the convention that the 'coarse surjectivity constant' in condition 2 is identical with the additive constant of quasi-isometry. This is not a universal convention, but it is commonly used and conveniently keeps the number of distinctly named constants to a minimum.

Essentially, a quasi-isometry is an isometry modulo some controlled distortion and bounded perturbation. In particular, while a quasi-isometry $\phi : X \to Y$ does not necessarily need to be a bijection, it will always admit a **quasi-inverse** - that is, a quasi-isometry $\bar{\phi} : Y \to X$ such that $\phi\bar{\phi}$ and $\bar{\phi}\phi$ are at bounded distance from the identity maps on Y and X, respectively. A quasi-inverse is not necessarily unique, and can be constructed by letting $\bar{\phi}(y)$ be any x satisfying condition 2 in the definition above for all $y \in Y$.

As is the case with isometries, the composition of two quasi-isometries is again a quasi-isometry, though its constants λ and c may be different. One consequence of this is that quasi-isometries induce an equivalence relation on metric spaces, so that it is meaningful to talk about the 'quasi-isometry class' of a given metric space, or to deal with such equivalence classes instead of the particular metric spaces they comprise. We run into trouble, however, when we try to fix the *quality* of the quasi-isometries under consideration; it is not generally the case that the composition of two (λ, c) quasi-isometries is again a (λ, c) quasi-isometry.

Thus, when we speak in later sections of a group acting on a space by uniform quasi-isometries, or more generally 'quasi-acting' on a metric space, we mean much more than that the group is *generated* by uniform-quality quasi-isometries. We mean additionally that all these quasi-isometries conspire among themselves to compose without disturbing the constants λ and c. This means, for example, that a group of uniform-quality quasi-isometries cannot contain elements that uniformly dilate or contract the original metric (e.g. non-isometric homotheties in affine n-space), even though such maps are prototypical examples of quasi-isometries in general.

The following famous theorem concerning group actions and quasi-isometries is central to geometric group theory, and will be needed in later sections. Recall that an action of a group G on a metric space X is said to be **geometric** if the following hold:

1. G acts on X by isometries
2. For all compact sets $K \subseteq X$, we have $\#\{g \in G \mid g \cdot K \cap K \neq \emptyset\} < \infty$
3. There exists $K \subseteq X$ compact such that $G \cdot K = X$

In fact, the following result can be generalized fairly substantially while preserving the same conclusion. In particular, we can let K be bounded instead of compact in condition 3, and can even relax the isometric action in condition 1 to a quasi-action. We present the theorem here in its most well-known form, and will address any generalizations if they become relevant.

Theorem 3.2 (Švarc-Milnor Lemma)**.** *Let X be a proper geodesic space, and let G be a group acting on X geometrically. Then G is finitely generated, and the orbit map $g \mapsto g \cdot x$, for any basepoint $x \in X$, is a quasi-isometry from G (metrized by the word metric with respect to any finite generating set) to X.*

The interested reader may find a proof for this famous theorem in Part I, Proposition 8.19 of [4].

3.2 Quasi-Isometries and Path Systems

If two geodesic spaces X and Y are quasi-isometric, it is natural to wonder how the systems of paths $\vec{X}$ and $\vec{Y}$ are related. In particular, what can we say about bounded path systems on X, if we know Y admits bounded path systems? It turns out that the answer varies depending on whether you are interested in synchronously or asynchronously bounded path systems. Asynchronously bounded path systems are well-behaved under quasi-isometry, as illustrated by the following lemma:

Lemma 3.3. *Suppose X and Y are geodesic spaces, $\phi : X \to Y$ is a (λ, c)-quasi-isometry, and $\vec{\gamma}_1, \vec{\gamma}_2$ are curves in X such that $\phi\vec{\gamma}_1$ and $\phi\vec{\gamma}_2$ are (monotone reparametrizations of) curves in Y. If $d_{async}(\phi\vec{\gamma}_1, \phi\vec{\gamma}_2) < B$, then $d_{async}(\vec{\gamma}_1, \vec{\gamma}_2) \leq \lambda B + 3c$.*

Proof: We can choose a quasi-inverse $\bar{\phi}$ for ϕ so that $d(\bar{\phi}\phi x, x) \leq c$ for all $x \in X$. Now we choose reparametrizations ρ, ρ' so that $d_Y(\phi\vec{\gamma}_1\rho(t), \phi\vec{\gamma}_2\rho'(t)) \leq B$ for all t, and we see that

$$\begin{aligned} d_X(\vec{\gamma}_1\rho(t), \vec{\gamma}_2\rho'(t)) &\leq d_X(\bar{\phi}\phi\vec{\gamma}_1\rho(t), \bar{\phi}\phi\vec{\gamma}_2\rho'(t)) + d_X(\vec{\gamma}_1\rho(t), \bar{\phi}\phi\vec{\gamma}_1\rho(t)) \\ &\quad + d_X(\vec{\gamma}_2\rho'(t), \vec{\gamma}_2\rho'(t)) \\ &\leq \lambda B + 3c \end{aligned}$$

□

Note that it is not generally the case that the image of a rectifiable curve under a quasi-isometry is again a curve. This does not prove to be an obstruction, however, and we lose no generality by assuming that ϕ maps curves to curves. This is because the asynchronous boundedness of a path system is equivalent to the existence of uniform-quality ladders for nearby pairs of paths (see Lemmas 2.7 and 2.8). If we have a path $\vec{\gamma} \in \vec{X}$ and $\phi : X \to Y$ is a quasi-isometry, then we can take any sequence $(t_i)_{i=0}^K$ that gives a uniform discretization of $\vec{\gamma}$ (as in the proof of Lemma 2.8) and associate to $\vec{\gamma}$ any piecewise-geodesic interpolation $\vec{\gamma_\phi}$ of the discrete path given by the sequence $(\phi(t_i))_{i=0}^K$. Thus while $\phi\vec{\gamma}$ may not be a curve itself, it does asynchronously fellow-travel the curve $\vec{\gamma_\phi}$ with bound depending only on the quasi-isometry ϕ and the coarseness of the discretization $(\phi(t_i))_{i=0}^K$. This allows us to conclude

Proposition 3.4. *Let X and Y be geodesic spaces, $\phi : X \to Y$ a quasi-isometry, and P an asynchronously bounded path system on X. Then there exists an asynchronously bounded path system P_ϕ on Y, such that $\phi P(x, y)$ and $P_\phi(\phi(x), \phi(y))$ are at uniformly bounded asynchronous distance for all $x, y \in X$.*

Finally, it should be pointed out that synchronously bounded path systems do *not* exhibit such nice behavior under quasi-isometry. This is one of several compelling reasons to think of the asynchronous distance as the 'natural' notion for applications of path geometry in geometric group theory.

By way of illustration, consider the following two spaces, with their respective intrinsic metrics (i.e. the distance between two points is the length of the shortest path connecting them):

- $A := [0, 1] \times [0, \infty) \subset \mathbb{R}^2$
- $B := A \setminus \bigcup_{n\in\mathbb{N}} N(1, (2n-1, 0))^\circ$

We see that A has a synchronously bounded path system P given by geodesics, and that the identity map gives a quasi-isometry $B \to A$. However, if we take a path in P and follow the procedure from Proposition 3.4 - discretizing paths with steps of uniform size, chasing them through the quasi-isometry, and joining them up by geodesic segments - we often get a path that does not synchronously fellow travel the original.

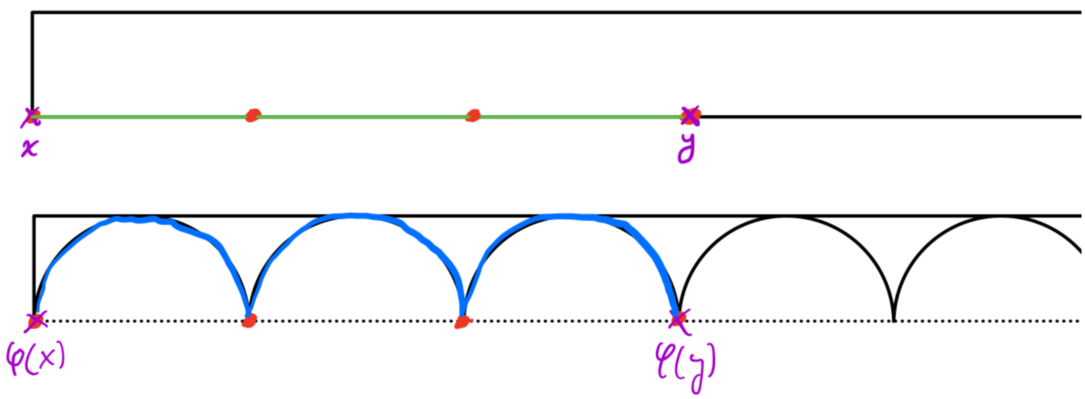

Figure 3.1. Synchronously fellow-traveling paths in A with quasi-isometric images in B, constructed as in the proof of Proposition 3.4 using steps of length 2

3.3 Hyperbolic Spaces

The notion of hyperbolicity for general metric spaces was introduced by Mikhail Gromov in the 1980s [10, 11], and is defined in its greatest generality in terms of the **Gromov product**. Since we are interested in geodesic spaces, we will use an equivalent characterization in terms of geodesic triangles. There are several characterizations of hyperbolicity for geodesic spaces, but the following definition has become standard:

Definition 3.5. *Let X be a geodesic space, and let $\delta > 0$ be given. We say X is **δ-hyperbolic** if each side of every geodesic triangle in X is contained in the δ-neighborhood of the other two sides.*

This 'thin triangle condition' means that triangles in hyperbolic spaces approximate tripods at large scales. One consequence of this is that every asymptotic cone (a construction that formalizes the idea of 'zooming out to infinity' when looking at a metric space) of a hyperbolic space is an $\mathbb{R}$-tree.

Hyperbolic spaces are fantastically well-behaved in many ways. Immediately relevant to us are the two classic results in the field: the Morse lemma, and the local-to-global property for quasi-geodesics. We begin with the Morse Lemma ([4], Part III, Theorem 1.7):

Lemma 3.6 (Morse Lemma)**.** *Let X be a δ-hyperbolic geodesic space, and let $\lambda, c \geq 0$. Then there exists a constant $M > 0$ such that, for any geodesic $\overrightarrow{\gamma_1}$ and (λ, c) quasi-geodesic $\overrightarrow{\gamma_2}$ such that $\overrightarrow{\gamma_1}(0) = \overrightarrow{\gamma_2}(0)$ and $\overrightarrow{\gamma_1}(\infty) = \overrightarrow{\gamma_2}(\infty)$, we have*

$$d_{Haus}(\overrightarrow{\gamma_1}, \overrightarrow{\gamma_2}) < M$$

In other words, quasi-geodesics with common endpoints have images contained within an M-neighborhood of each other., where M only depends on the quality of the quasi-geodesics. In light of the discussion in Section 2, we can restate this lemma both quantitatively and in terms of the **asynchronous** distance between quasi-geodesics:

Lemma 3.7. *In a δ-hyperbolic geodesic space X, the asynchronous distance between any pair of (λ, c) quasi-geodesics $\overrightarrow{\gamma_1}, \overrightarrow{\gamma_2}$ is bounded by a linear function of the distance between their endpoints. In other words, any path system on X consisting of (λ, c) quasi-geodesics is asynchronously bounded.*

Proof: Because the departure function for a path system of fixed-quality quasi-geodesics is linear, Lemma 2.11 shows that the asynchronous distance between $\overrightarrow{\gamma_1}$ and $\overrightarrow{\gamma_2}$ is a linear function of the Hausdorff distance between their images. We may apply the Morse lemma to $\overrightarrow{\gamma_1}$ and the concatenated path $\overrightarrow{\gamma_L}\overrightarrow{\gamma_2}\overrightarrow{\gamma_R}$, where $\overrightarrow{\gamma_L}$ and $\overrightarrow{\gamma_R}$ are geodesics joining the endpoints of $\overrightarrow{\gamma_1}$ and $\overrightarrow{\gamma_2}$. Note that $\overrightarrow{\gamma_L}\overrightarrow{\gamma_2}\overrightarrow{\gamma_R}$ is a $(\lambda, c+d)$-quasigeodesic, where d is the maximal distance between pairs of corresponding endpoints. Referencing [9], we find that the Hausdorff distance between $\overrightarrow{\gamma_1}$ and $\overrightarrow{\gamma_2}$ is bounded above by $184\lambda^2(c+d+\delta)$, which is linear in d. □

The second foundational result we need to cite is the local-to-global property of quasi-geodesics in hyperbolic spaces. A ***k*-local geodesic** is a curve $\vec{\gamma}$ such that $d(\vec{\gamma}(s), \vec{\gamma}(t)) = |s-t|$ whenever $|s-t| < k$ and $s, t \in [0, \ell(\vec{\gamma})]$. In hyperbolic spaces, k-local geodesics for sufficiently large k are guaranteed to be quasi-geodesics whose quality depends on δ and k ([4], Part III, Theorem 1.13):

Theorem 3.8. *Let X be a δ-hyperbolic geodesic space, and let $k > 8\delta$. Then every k-local geodesic $\vec{\gamma}$ is a global $(\frac{k+4\delta}{k-4\delta}, 2\delta)$ quasi-geodesic.*

3.4 The Gromov Boundary

Mostly, these nice properties of hyperbolic spaces can be derived from the behavior of geodesics and **geodesic rays** - that is, maps $\vec{\gamma} : [0, \infty) \to X$ such that the restriction of $\vec{\gamma}$ to any closed, bounded interval is a geodesic. Significantly for the present discussion, geodesic rays in hyperbolic spaces allow us to define a nice **boundary**, whose topology is intimately related to the internal geometry of X:

Definition 3.9. *Let X be a hyperbolic space. The **Gromov boundary** ∂X of X is the set of equivalence classes of geodesic rays in X with a fixed basepoint, where two rays are equivalent if they are at bounded synchronous distance from each other.*

The Gromov boundary of a proper hyperbolic geodesic space X can be equipped with a topology that makes it a compact, metrizable space. A basis for this topology is given by the set of 'cones', or sets of equivalence classes of geodesics which have representatives identical on some interval $[0, r]$. Furthermore, quasi-isometries of X induce homeomorphisms of ∂X, and this correspondence can be exploited to control the behavior of quasi-isometries on X.

We will only need the Gromov boundary to define the notion of a 'bushy' hyperbolic space, which supports a critical theorem used in Section 5. Before we can give this definition, we need to introduce one more concept: that of a **coarse median**, as defined by Bowditch in [2]. The true definition is a coarsening of the notion of a median function on a so-called 'median space' and applies to a broad class of metric spaces. For simplicity, we only define coarse median maps in the context of hyperbolic spaces.

Definition 3.10. *Let X be a δ-hyperbolic geodesic space. A* ***coarse median map*** *on X is a map $\mu : X^3 \to X$ such that the image of any triple $(x, y, z) \in X^3$ lies in the intersection of the δ-neighborhoods of the sides of the geodesic triangle xyz.*

Recall that hyperbolic geodesic spaces are spaces in which triangles are 'coarsely tripods'. A coarse median map, then, maps the corners of such a coarse tripod to its 'coarse center'. If we allow the corners of a geodesic triangle to escape 'to infinity' along geodesic rays, we can extend the coarse median to the Gromov boundary. With this fact in hand, we introduce the following definition from [3] :

Definition 3.11. *A δ-hyperbolic geodesic space X is said to be* ***bushy*** *if the coarse median map on the boundary, $\mu|_{\partial X} : (\partial X)^3 \to X \cup \partial X$, is coarsely surjective onto X (i.e. has image at bounded Hausdorff distance from X).*

We will use this definition in the next section.

3.5 Products of Hyperbolic Spaces

The main results in this paper concern asynchronous automatic structures on groups quasi-isometric to products of hyperbolic spaces. In this section, we briefly go over the properties of metric products, with emphasis on the special properties enjoyed by products of hyperbolic spaces.

3.5.1 Equivalent Product Metrics

Given a direct product of metric spaces $(X_1, d_1), (X_2, d_2), \cdots (X_N, d_N)$ and $p \in \mathbb{N}$, we define the **p-product metric**

$$d^p((x_1, \cdots, x_N), (y_1, \cdots, y_N)) = (\sum_{n=1}^{N} d_n(x_n, y_n)^p)^{\frac{1}{p}}$$

for any two points $(x_1, \cdots, x_N), (y_1, \cdots, y_N) \in \prod_{n=1}^{N} X_n$. Evaluating this expression in the limit as $p \to \infty$ gives the **sup metric**

$$d^\infty((x_1, \cdots, x_N), (y_1, \cdots, y_N)) = \max\{d(x_n, y_n) \mid n \in \{1, \cdots, N\}\}.$$

Note that the metric topology arising from any of these metrics coincides with the product topology, as we would expect for a 'product metric'.

We say two metrics d and d' on a common space X are **equivalent metrics** if they are Lipschitz equivalent - that is, there exists $C > 0$ so that

$$\frac{1}{C}d(x,y) \leq d'(x,y) \leq Cd(x,y)$$

for all $x, y \in X$. Since we clearly have

$$\begin{aligned} d^\infty((x_1,\cdots,x_N),(y_1,\cdots,y_N)) &\leq d^p((x_1,\cdots,x_N),(y_1,\cdots,y_N)) \\ &\leq Nd^\infty((x_1,\cdots,x_N),(y_1,\cdots,y_N)) \end{aligned}$$

for all $(x_1, \cdots, x_N), (y_1, \cdots, y_N) \in \prod_{n=1}^N X_n$ and all $p \in \mathbb{N}$, we can conclude that all of the product metrics defined above are equivalent. Importantly, this means that the quasi-isometry class of the product $\prod_{n=1}^N X_n$ is independent of which of the above product metrics we choose to assign it. We will largely be concerned with properties of quasi-isometry classes of groups, so we will simply refer to 'the' product metric in later sections.

3.5.2 The Geometry of Coordinate Factors

If $X = \prod_{i=1}^N X_i$ is a product of metric spaces, then for any $x \in X$ and $i \in \{1, \cdots, N\}$, there is a 'canonical' copy of X_i isometrically embedded in X which contains x. We call this copy, denoted σ_i^x, the **i-th coordinate factor at x** and obtain it by setting

$$\sigma_i^x = \{y \in X \ : \ \pi_j(y) = \pi_j(x) \ \forall j \neq i\}$$

If a factor X_i is hyperbolic, it is within the corresponding coordinate factor that hyperbolic behavior will appear in the product. This is because, as mentioned, the coordinate factors are all isometrically embedded. In the context of a product of geodesic spaces, this means that the induced metric on a coordinate factor as a subspace coincides with its intrinsic metric. Hence, the geodesics and quasi-geodesics of X which remain within a single, hyperbolic coordinate factor will be bound by the hyperbolic geometry of that factor.

Furthermore, this is true of any fixed neighborhood of a coordinate factor as well, up to possibly inflating the value of δ. This is because points in coordinate factors are joined by geodesics lying in those coordinate factors, and these geodesics remain geodesic after embedding the coordinate factor into its neighborhood (with its intrinsic metric). This yields the following theorem:

Theorem 3.12. *Let X be a geodesic space, H a δ-hyperbolic geodesic space, $x \in X \times H$ a point in their product, and $\eta > 0$ some fixed constant. Then the η-neighborhood of σ_H^x is a δ'-hyperbolic space, where δ' depends only on δ and η.*

Proof: The inclusion $\iota : H \to N(\eta, \sigma_H^x)$ is an isometric embedding by the observation above. This is in fact a quasi-isometry, since $d_{Haus}(\iota(H), N(\eta, \sigma_H^x)) = \eta$, and the constant of hyperbolicity for the target space depends only on the hyperbolicity constant of H and the coarse-surjectivity constant δ (since the 'quasi-isometry constants' of the isometry ι are $(1, 0)$). □

3.5.3 Products of Hyperbolic Spaces

We now have the terminology to discuss the main result that we will apply to control the behavior of a group quasi-isometric to a product of hyperbolic spaces. This essence of this theorem is originally due to a publication of Kapovich, Kleiner, and Leeb [17], but we will use a refinement of the original result by Bowditch [3] as the latter's hypotheses are slightly more general and require less exposition.

The essence of the following theorem is that quasi-isometries from a product of non-elementary hyperbolic spaces to itself must coarsely decompose as products of quasi-isometries in the factors, up to permutation of quasi-isometric factors. In the language we have developed so far in the present paper, this means that self quasi-isometries of hyperbolic products, composed with an appropriate permutation of coordinates, coarsely preserve coordinate factors. First, we repeat Bowditch's theorem in its entirety:

Theorem 3.13 (Bowditch)**.** *Suppose that for $i = 1, ..., n$, we have hyperbolic spaces Λ_i and Λ'_i. Suppose that $1 \leq q \leq p \leq n$, such that for all $i \leq p$, Λ_i is bushy and for all $j > q$, Λ'_j is quasi-isometric to the real line. Let $L = \prod_{i=1}^n \Lambda_i$ and $L' = \prod_{i=1}^n \Lambda'_i$. Suppose that $\phi : L \to L'$ is a quasi-isometric embedding. Then $p = q$. Moreover, after permuting the indices $1, ..., p$, there are quasi-isometric embeddings $\phi_i : \Lambda_i \to \Lambda'_i$ for $i \leq p$, and a quasi-isometry, $\phi_u : \prod_{i=p+1}^n \Lambda_i \to \prod_{i=p+1}^n \Lambda'_i$, for each $u \in \prod_{i=1}^p \Lambda_i$ such that for all $x \in L$, $\phi(x)$ is a bounded distance from $\psi(u), \phi_u(v)$, where u, v are respectively projections of x to the first p and last $n - p$ coordinates, and where ψ is a direct product of the maps ϕ_i for $i \leq p$, and the bound depends only on the parameters of the hypotheses.*

The key application of this theorem is the following observation:

Corollary 3.14. *Let $X = \prod_{i=1}^N H_i$ be a product of hyperbolic geodesic spaces, and let $\phi : X \to X$ be a quasi-isometry whose induced permutation on the factors of X is trivial. Then ϕ splits as a product of self-quasi-isometries of the factors H_i, up to some bounded error depending only on the space X and the quality of ϕ. In particular, there exists $\alpha > 0$ (depending only on the constants of hyperbolicity, quasi-isometry, and bushiness) such that* $\operatorname{diam}(\pi_j \phi(\sigma_i^x)) < \alpha$ *for all $x \in X$ and all $i, j \in \{1, \cdots, N\}$*

such that $i \neq j$.

Given any group of G of quasi-isometries of X, we can always pass to a finite-index subgroup of G whose elements induce only the trivial permutation on the factors. The finite-index subgroups of G are quasi-isometric to G, so we will assume going forward that our quasi-isometries do not permute factors, and thus satisfy Corollary 3.12.

Chapter 4: Regular Languages and Automatic Groups

Let G be a group generated by a finite set of group elements $A \subseteq G$. For the sake of simplicity, we shall always assume that $A^{-1} = A$. Associated to the group G is a graph, the **Cayley Graph** of G, whose edge relation effectively encodes the entire structure of the group.

Definition 4.1 (Cayley Graph)**.** *Let G be a finitely generated group with generating set A. The **Cayley Graph** of G with respect to A, denoted $Cay(G, A)$ or simply $Cay(G)$ if the generating set is understood, is the labeled directed graph with vertex set $\mathcal{V} = G$ and labeled edge set $\mathcal{E} = \{(g, ga, a) \mid a \in A\}$ (we understand the edge (g, ga, a) to be directed from g to ga and labeled by a).*

The set of finite ordered tuples of elements of A, called **strings** or **words** over the alphabet A, is denoted A^*, and the set of strings of length n will be denoted A^n. We shall use the simplified notation $a_1...a_m$ to indicate the ordered tuple $(a_1, ..., a_m)$. The set A^* admits a natural binary operation in the form of **concatenation**, i.e. the operation $(a_1...a_m, a'_1...a'_n) \mapsto a_1...a_m a'_1...a'_n$. Under this operation, the set A^* has the structure of a **monoid** with the addition of an identity element $\epsilon \notin A$, called the

empty word or the **empty string**. This monoid A^* is also called the **free monoid** generated by the set A.

There is a natural monoid homomorphism $E : A^* \to G$ which maps each letter $a \in A$ to the group element $a \in A \subseteq G$ and maps the empty string to the identity of G. We refer to E as the **evaluation homomorphism**, **evaluation map**, or simply **evaluation**. We will represent the vertices of $Cay(G, A)$ (i.e. the elements of G) by choosing a **normal form** for G in A^*. For us, a normal form will always be a subset $\mathcal{L} \subset A^*$ such that

- $E(\mathcal{L}) = G$, and
- There exists $n \in \mathbb{N}$ such that $\#(E^{-1}(g) \cap \mathcal{L}) \leq n$ for all $g \in G$

It is not uncommon for authors to drop the second condition in this definition, in which case the normal forms we treat here would be specified as *(uniformly) finite-to-one*. Assuming from the outset that all of our normal forms are finite-to-one will save us a lot of technical headaches later on, and will not in any way limit the use of this concept for our purposes. In most cases, it is safe to assume that a normal form consists of *unique* representatives - in the particular case of **automatic groups**, defined in a later section, the generality of this assumption is in fact a theorem [8].

4.1 Regular Languages and Finite-State Automata

A subset $\mathcal{L}$ of the free monoid A^* is also called a **language** over the alphabet A. We shall principally concern ourselves with **regular languages**, as they are the most restrictive and, consequently, computationally simple. There are many equivalent characterizations of regular languages. The most compact, and hence most valuable to computer scientists, is the **regular expression**. A regular expression is a finite string whose characters consist of the alphabet A over which a language is to be defined, parentheses and commas for bracketing and listing, the **alternation** symbol $\vee$ (understood as indicating set union or logical disjunction), and the **Kleene star** *. We have already seen the Kleene star applied to the alphabet A to yield the set A^* of all finite words over A - in general, the Kleene star of a set of strings S is the set of all finite strings which can be obtained by concatenating elements of S.

Example 4.2. *The following regular expression describes the language $\mathcal{L}$ of all strings over the alphabet $\{a, b, c\}$ which contain an even number of occurrences of the letter a:*

$$(\mathtt{b} \vee \mathtt{c})^*(\mathtt{a}(\mathtt{b} \vee \mathtt{c})^*\mathtt{a})^*(\mathtt{b} \vee \mathtt{c})^*$$

We may read this expression as beginning with any string over $\{b, c\}$, followed by arbitrarily many pairs of a's with any string over $\{b, c\}$ between them, followed by any string over $\{b, c\}$.

A more comprehensive perspective on regular languages classifies them by the type of Turing machine that recognizes them. These machines are called **finite-state automata**, and can in fact be thought of as forming an essential component of *every* Turing machine.

Definition 4.3. *A* ***finite-state automaton*** *$\mathcal{M}$ is a tuple $\mathcal{M} = (A, \mathcal{Q}, S, C, \tau)$ where A is a finite alphabet, $\mathcal{Q}$ is a finite set called the set of* ***states****, $S \subseteq \mathcal{Q}$ is the set of* ***start*** *states, $C \subseteq \mathcal{Q}$ is the set of* ***accept*** *states, and $\tau : \mathcal{Q} \times (A \cup \{\epsilon\}) \to 2^{\mathcal{Q}}$ (where ϵ is the empty word of A^*) is the* ***transition function****. We say the language $\mathcal{L} \subseteq A^*$ is* ***accepted*** *or* ***recognized*** *by the finite-state automaton $\mathcal{M}$ if and only if $\mathcal{L}$ consists* *precisely of strings $a_1...a_n$ such that there exists a sequence of states $Q_1, ..., Q_m$, plus a subset $V \subseteq \{1, \cdots, m\}$ of cardinality n, satisfying*

1. $Q_1 \in S$

2. $Q_m \in C$

3. $Q_{i+1} \in \tau(a_i, Q_i)$ *for all* $i \in V$, *and* $Q_{i+1} \in \tau(\epsilon, Q_i)$ *otherwise.*

It is usually more convenient to describe a finite-state automaton as a finite directed graph with vertex set $\mathcal{Q}$, having a set of distinguished 'start' vertices and a set of distinguished 'accept' vertices. We have an edge labeled a from Q_1 to Q_2 if and only if $Q_2 \in \tau(a, Q_1)$. The accepted language is then the language of labels of directed paths beginning at the start vertex and ending at an accept vertex. The illustration below is the state graph for an automaton recognizing the language above:

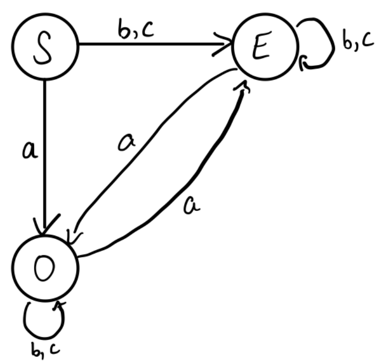

Figure 4.1. The state graph for an automaton recognizing the regular expression in Example 4.2

It is this more visual understanding of finite-state automata that we shall prefer in subsequent discussion, as it allows for intuitive visual proofs.

4.1.1 Properties of Regular Languages

Regular languages satisfy many useful closure properties. We will by no means exhaust all of them here, but we will go over a few that are relevant to the preceding and following discussion.

Alternation If $\mathcal{L}$ and $\mathcal{L}'$ are regular languages, then their union $\mathcal{L}'' = \mathcal{L} \cup \mathcal{L}'$ is also a regular language (in the parlance of regular expressions, we referred to this as *alternation* and notated it $\mathcal{L} \vee \mathcal{L}'$). If $\mathcal{M}$ and $\mathcal{M}'$ are finite state automata recognizing $\mathcal{L}$ and $\mathcal{L}'$, respectively, then we can get an automaton $\mathcal{M}''$ recognizing $\mathcal{L}''$ by simply taking as a state graph the disjoint union of the state graphs of $\mathcal{M}$ and $\mathcal{M}'$, keeping start and accept states as well.

Complementation Given a language $\mathcal{L}$ over an alphabet A, its **complement** $A^* \setminus \mathcal{L}$ is its set-theoretic complement in A^*. Given a finite-state automaton $\mathcal{M}$ recognizing $A^* \setminus \mathcal{L}$, the automaton recognizing $A^* \setminus \mathcal{L}$ is constructed in two steps. First, we complete the automaton $\mathcal{M}$ by adding a new state F, often called a **fail state**, and setting $\tau(a, Q) = \{F\}$ wherever we had $\tau(a, Q) = \emptyset$ before. In terms of state graphs, this corresponds to adding a new vertex with incident loops labeled by each letter in A, and an incoming edge labeled a from every state which previously lacked an outgoing edge labeled a. Once we complete $\mathcal{M}$ to an automaton $\mathcal{M}'$ which recognizes $\mathcal{L}$ and has nonempty transitions for all nonempty inputs, we finish constructing the automaton recognizing $A^* \setminus \mathcal{L}$ by changing the accept set C to $\mathcal{Q} \setminus C$; that is, by reversing the roles of accept and non-accept states.

With procedures for handling alternation (i.e. disjunction) and complementation (i.e. negation), we have an effective procedure for building a finite-state automaton

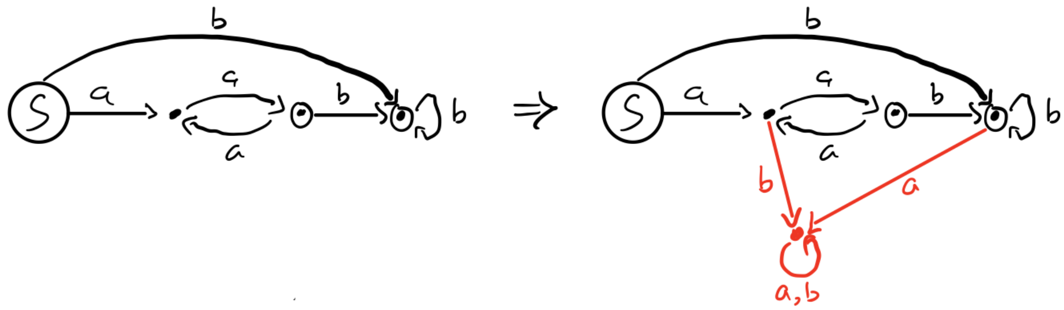

Figure 4.2. An automaton recognizing $\mathcal{L} = \{a^{2m}b^n | m, n \in \mathbb{N}\}$ is completed by adding a fail state.

which recognizes the language of words w satisfying any propositional sentence whose atomic formulas are containment of w in some fixed regular language.

Concatenation If $\mathcal{L}$ and $\mathcal{L}'$ are regular languages, their **concatenation** $\mathcal{L}\mathcal{L}'$ is the language $\{vw | v \in \mathcal{L}, w \in \mathcal{L}'\}$. Given deterministic finite-state automata $\mathcal{M}$ and $\mathcal{M}'$ recognizing $\mathcal{L}$ and $\mathcal{L}'$, respectively, we obtain an automaton recognizing $\mathcal{L}\mathcal{L}'$ by taking the disjoint union of their state graphs and adding an edge labeled ϵ from every accept state of $\mathcal{M}$ to the start state of $\mathcal{M}'$.

Kleene Star Similar to the construction for concatenation, we can construct an automaton recognizing the Kleene star of a finite collection of regular languages, $\{\mathcal{L}_1, ..., \mathcal{L}_n\}^*$, by taking the disjoint union of the state graphs of deterministic finite-state automata $\mathcal{M}_1, ..., \mathcal{M}_n$ recognizing $\mathcal{L}_1, ..., \mathcal{L}_n$, respectively, and adding to the resulting state graph an edge labeled ϵ from each accept state of $\mathcal{M}_i$ to the start state of $\mathcal{M}_j$, for all $i \neq j$.

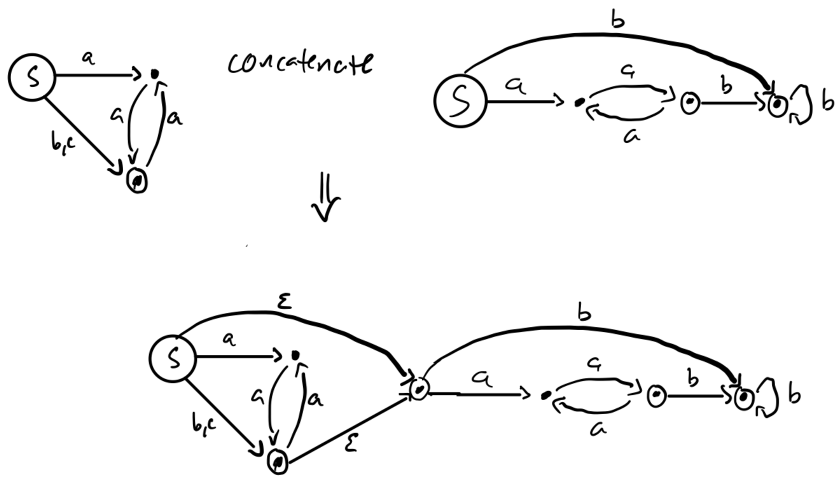

Figure 4.3. Concatenation of automata

4.1.2 Regular Languages and Local Recognition

A subclass of regular languages that will be of interest to us is what we shall refer to as **locally recognized** languages. By 'locally recognized languages', we refer to those languages whose membership can be checked by looking at a candidate word only through a sliding 'window' of finite length. More precisely,

Definition 4.4. *Let A be a finite alphabet, $\mathcal{L} \subseteq A^*$ a language over A, and $k > 0$. We say $\mathcal{L}$ is k-**locally recognized** if and only if there exist finite sets $S, C \subseteq A^*$ and $R \subseteq A^k$ (called, respectively, the **prefixes**, **suffixes**, and **infixes** of $\mathcal{L}$) such that $\mathcal{L}$ is the set of words $w \in A^*$ satisfying the following:*

1. $w \in SA^*C$

2. *For every $x \in A^k$ such that $w = uxv$ for some $u, v \in A^*$, we have $x \in R$*

*If there exists $k > 0$ such that $\mathcal{L}$ is k-locally recognized, we say $\mathcal{L}$ is **locally recognized.***

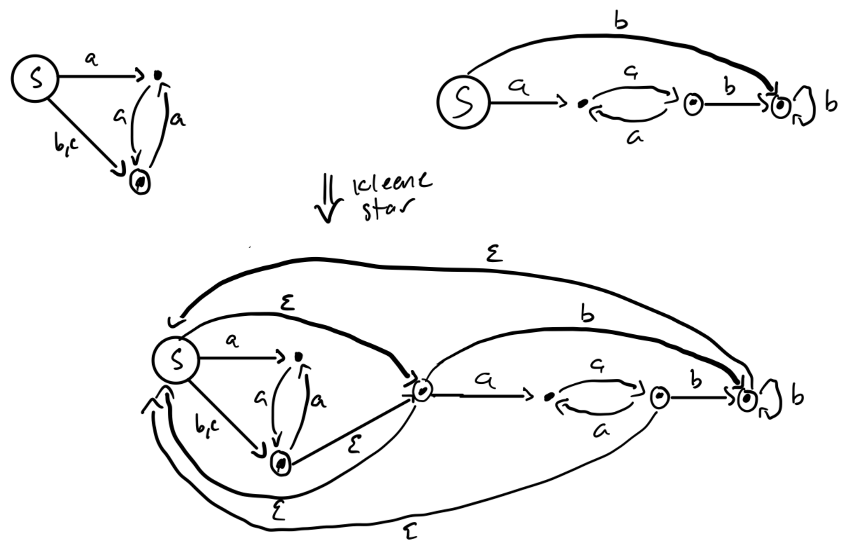

Figure 4.4. Automaton recognizing the Kleene star of a pair of languages

Sometimes this notion is defined so that R is a set of *forbidden* infixes rather than allowed ones. This definition is equivalent, and we can translate between them by complementation in A^k.

Proposition 4.5. *All locally recognized languages are regular.*

Proof: Let $\mathcal{L}$ satisfy definition 4.4 with respect to $k > 0, S, C \subset A^*$, and $R \subseteq A^k$. We will first show that the language

$$\mathcal{L}' = \{w \in A^* | \text{ every length } k \text{ infix of } w \text{ is in } R, \text{ or } w \text{ is a prefix of an element of } R\}$$

is regular. The result will then follow, since $\mathcal{L} = \mathcal{L}' \cap SA^*C$, and the intersection of two regular languages is regular.

We will build a finite-state automaton $\mathcal{M}'$ recognizing $\mathcal{L}'$. First, we create a state for each word $w \in R$. Then we create a start state together with enough new states to draw a path labeled by w from the start state to the state w, for each $w \in R$. Then, for each $a \in A$ and $w = w_1...w_n \in R$, we set $\tau(a, w) = \{w_2...w_n a\} \cap R$. All states are accept states.

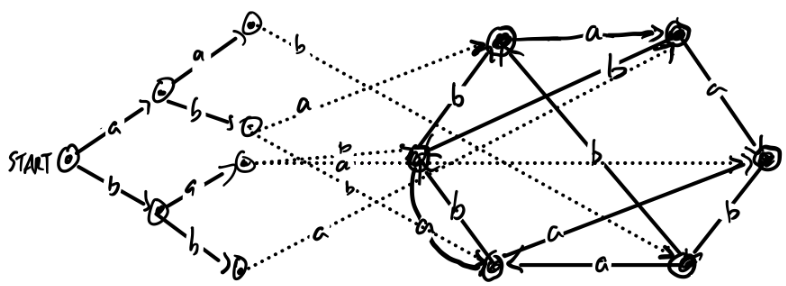

Figure 4.5. $\mathcal{M}'$ when $A = \{a, b\}$ and $R = A^3 \setminus \{aaa, bbb\}$

It is now easy to check that $\mathcal{M}'$ recognizes $\mathcal{L}'$. $\square$

Locally recognized languages occur frequently in normal forms that arise from applications in geometric group theory. Of particular note is the 2006 work of Świątkowski [21] which helpfully generalized earlier results on the automaticity of certain groups; this result has seen much recent reference in works by Chalopin, Osajda and others [7, 14] on injective spaces and Helly graphs.

4.2 Multi-Tape Automata and Automatic Relations

The notion of an **automatic group** was defined and elaborated on in a seminal monograph of Cannon, Epstein, Holt, Levy, Paterson, and Thurston [8], and proofs for the theorems in this section can be found there (we also provide our own proofs here, using the language of path systems developed in the preceding chapters). Automatic groups satisfy an impressive list of useful properties related to complexity, and admit efficient algorithms for resolving group-theoretic questions. Given a group G generated by a finite set A, an **automatic structure** on G is a regular language in A^* which gives a finite-to-one normal form for G, and such that the relation

$$R_a = \{(w, v) \in A^* \ : \ w, v \in \mathcal{L} \text{ and } E(w) = E(v)a\}$$

is what we will call an **automatic relation** for each generator a. Informally, a relation on A^* is called automatic if it is 'given by' a regular language. Our first objective will be to define what we mean by a regular language of *pairs* of words (or, more generally, *tuples* of words). Our definition will agree with the slightly more general notion of an *automatic structure* as introduced by Khoussainov and Nerode [18].

4.2.1 Synchronous and Asynchronous Languages of Tuples

The relation described above is a subset of $A^* \times A^*$, in the same way that regular languages are subsets of A^*. The ambiguity in describing a subset of $A^* \times A^*$ as a language itself lies in the alphabet. We may naively propose A^2 as a candidate alphabet, but this fails since A^2 does not even generate $A^* \times A^*$ as a monoid: it instead generates the submonoid $\{(w, v) \ : \ |w| = |v|\}$.

To account for words of differing lengths, we introduce a new symbol not originally in the alphabet A, called the **padding symbol** and typically denoted \$. Now for any natural number n, we define the **n-letter padded alphabet**

$$A_{(n)} := (A \cup \{\$\})^n \setminus \{(\$, \cdots, \$)\}$$

From here, we can form the free monoid $A^*_{(n)}$ and see that it does represent $(A^*)^n$ via the natural surjection $U : A^*_{(n)} \to (A^*)^n$ which sends \$ to ϵ in each coordinate.

If we stop here and consider a subset $\mathcal{R} \subseteq (A^*)^n$ such that there is a regular language $\mathcal{L} \subseteq A^*_{(n)}$ with $U(A^*_{(n)}) = \mathcal{R}$, then we have what we shall call an **asynchronously automatic relation**. The synchronous counterpart of this notion, as well as the synchronous/asynchronous naming convention itself, will benefit from some motivation.

The classic mechanical metaphor for the Turing Machine envisions a physical device consisting of three parts: an (in principle) infinite tape divided into cells which contain the letters of an input, a read/write head positioned over the tape that processes one cell at a time, and an internal mechanism connected to the read/write head that tracks the machine's state. When the read/write head is only allowed to read and the input tape must proceed sequentially from left to right, this contraption acts as a finite-state automaton and recognizes regular languages. We can now imagine a similar device which has *multiple* read/write heads and admits a tape for each one, advancing at least one tape to the right by one cell at each step in its computation.

An asynchronously automatic relation as described above corresponds to a collection of tuples of input words that *can* be accepted by the machine if their tapes are processed in the correct order. Since the tapes do not have to be read simultaneously, we call the corresponding multi-tape automaton *asynchronous*. The downside is that, a priori, we do not know what sequence of tape advancements will allow an accepted tuple to be successfully processed by the machine. In principle, our best decision algorithm in the general case would require us to try every possible configuration for a given set of inputs before deciding to reject them. This results in exponential-or-worse time complexity for algorithms naively based on asynchronously automatic structures.

Returning to the multi-tape machine, we might imagine imposing a further restriction on its operation: the two input tapes must be read simultaneously until one of them reaches its end, and from there the other tape is read alone. The result is a mechanical realization of a **synchronous multi-tape automaton**.

At the algebraic level, we begin defining a 'synchronizing map' $S : (A^*)^n \to A^*_{(n)}$, which we illustrate in the case of two-tape automata: for words $w = w_1 \cdots w_n$ and $v = v_1 \cdots v_m$ (assuming without loss of generality that $m \leq n$), we set

$$S(w, v) = (w_1, v_1) \cdots (w_m, v_m)(w_{m+1}, \$) \cdots (w_n, \$)$$

The extension of this map to automata with more than two tapes is similar. Now we will declare a subset $\mathcal{R} \subseteq (A^*)^n$ to be a **synchronously automatic relation** if and only if $S(\mathcal{R})$ is a regular language.

Before we move on to define synchronously and asynchronously automatic *groups*, two additional observations may be of interest. Firstly, there is no loss of generality in assuming that a multi-tape automaton reads from only one tape in any time step. This corresponds to replacing $A_{(n)}$ with the alphabet

$$A_{[n]} := \{(a_1, \cdots, a_n) \in (A \cup \{\$\})^n \ : \ a_i \neq \$ \text{ for exactly one } i\}$$

Given an automaton recognizing an asynchronously automatic relation, we can obtain an equivalent automaton over $A^*_{[n]}$ by replacing every transition labeled $(a_1, \cdots, a_n)$ with a sequence of transitions (and new states between them) labeled $(a_1, \$, \cdots, \$)$, $(\$, a_2, \$, \cdots, \$)$, etc.

Secondly, there is nothing sacred about the synchronizing function S. The performance advantage enjoyed by synchronous automata is shared by any asynchronous automaton equipped with a similar 'conditioning' function; e.g. an asynchronous automaton whose accepted pairs are always read at a rate of 2 to 1 can be treated like a synchronous automaton with a modified synchronizing function S' which encodes inputs in a way that reflects this.

4.2.2 Automatic Groups

The notions of automatic and asynchronously automatic groups were first introduced in [8]; our definitions below and in the preceding section are equivalent to theirs and differ only by a little terminology, with one exception. The familiar reader may recall that their definition of an asynchronous multi-tape automaton differs from ours significantly. In particular, their asynchronous multi-tape automaton recognizes 'shuffles' of strings over a single alphabet, while ours recognizes padded languages of tuples as in the synchronous case. That these notions yield identical characterizations of automatic groups is nontrivial, and was first established by Derek Holt [8].

An automatic group is a finitely generated group $G = \langle A \rangle$ with a regular normal form $\mathcal{L} \subset A^*$ such that equality in G and differing on the right by a given generator are both automatic relations. Automatic groups satisfy a number of useful geometric properties, and also admit efficient algorithms for certain group-theoretic problems like the word problem. We will call such a group **asynchronously automatic** if the associated relations are asynchronously automatic relations.

Definition 4.6. *Let G be a group generated by a finite set A, and $\mathcal{L} \subset A^*$ a regular language. We say $\mathcal{L}$ furnishes an **automation** for G, $(G, \mathcal{L})$ is an **automatic structure**, or simply that G is an **automatic group** with normal form $\mathcal{L}$, if the following conditions hold:*

- *$\mathcal{L}$ is a normal form for G*

- *The sets*

$$\mathcal{R}_\epsilon := \{(v,w) \in \mathcal{L} \times \mathcal{L} \ : \ E(v) = E(w)\}$$

and

$$\mathcal{R}_a := \{(v,w) \in \mathcal{L} \times \mathcal{L} \ : \ E(v) = E(w)a\}$$

(for each $a \in A$) are synchronously automatic relations.

If the relations in condition 2 are instead asynchronously automatic, then we say G is an ***asynchronously automatic group****, that $\mathcal{L}$ furnishes an* ***asynchronous automation*** *for G, or that $(G, \mathcal{L})$ is an* ***asynchronously automatic structure****.*

Once it is known that a group is automatic with respect to some normal form, the problem of constructing the automata for an automatic structure can be reduced to determining a single constant associated with the geometry of the normal form $\mathcal{L}$ as a path system in $\mathrm{Cay}(G, A)$. It is this geometric association, often called a *fellow-traveling property*, that makes automatic groups of such interest to geometric group theorists.

Consider an automatic structure $(G, \mathcal{L})$ with generating set A. Choose $a \in A \cup \{\epsilon\}$ and let $\mathcal{M}$ be a two-tape automaton recognizing $\mathcal{R}_a$. Now take any $(v,w) \in \mathcal{R}_a$, and let W be a word in $A^*_{[2]}$ which is accepted by $\mathcal{M}$ and satisfies $U(W) = (v,w)$. Letting W_i denote the i-th prefix of W, we similarly define $v_i = \pi_1 U(W_i)$ and $w_i = \pi_2 U(W_i)$, where π_i is the canonical projection to the i-th factor in $A^* \times A^*$.

Since W_i is a prefix of an accepted word, there is a directed path in the state graph of $\mathcal{M}$ from the state after reading W_i to an accept state. We can always choose this path so that it crosses fewer than N edges, where N is the number of states in $\mathcal{M}$.

Let W_i' be the label of this path, and define $v_i' = \pi_1 U(W_i')$ and $w_i' = \pi_2 U(W_i')$. Since $W_i W_i'$ is accepted by $\mathcal{M}$, we have

$$E(v_i v_i') = E(w_i w_i')a$$

or, equivalently,

$$E(w_i)^{-1}E(v_i) = E(w_i')aE(v_i')^{-1}.$$

This allows us to conclude

$$d_{G,A}(E(v_i), E(w_i)) \leq N + 1.$$

In other words, $\vec{v}$ and $\vec{w}$ admit a $(1, 1, N+1)$-ladder. By Lemma 2.7, we can now conclude that every pair of words $v, w \in \mathcal{L}$ such that $d_{G,A}(\vec{v}(\infty), \vec{w}(\infty)) \leq 1$ satisfies

$$d_{Async}(\vec{v}, \vec{w}) \leq \frac{2N+5}{2}$$

In other words, the path system on $\mathrm{Cay}(G, A)$ induced by $\mathcal{L}$ is bounded; synchronously if the automatic structure is synchronous and asynchronously otherwise.

Now suppose that $G = \langle A \rangle$ is a finitely-generated group and $\mathcal{L}$ is a regular normal form for G which induces a bounded path system on $\mathrm{Cay}(G, A)$. Then for every $a \in A \cup \{\epsilon\}$, the set $\mathcal{R}_a$ consists of pairs of words which fellow-travel with some uniform constant k. We can now define an automaton $\mathcal{M}_a$ recognizing $\mathcal{R}_a$. The states of $\mathcal{M}_a$ are the vertices of the ball of radius k in $\mathrm{Cay}(G, A)$. We add a transition with label $(a', \$)$ from g to ga' for each state g and each $a' \in A$ such that ga' is a state, and we add a transition with label $(\$, a')$ from g to $(a')^{-1}g$ for each state g and each letter

$a' \in A$ such that $(a')^{-1}g$ is a state. Finally, we make the vertex a the unique accept state. We see now that, for any word $W \in A^*_{[2]}$ which can be read by $\mathcal{M}_a$, the state of $\mathcal{M}_a$ upon reading W is precisely the word-difference

$$(E\pi_2 U(W))^{-1} E\pi_1 U(W)$$

That is, if we write $U(W) = (v, w)$, then the state of this automaton after reading W is precisely the element $g \in G$ such that $E(v) = E(w)g$. Since the only accept state is a, this automaton recognizes those pairs of words which evaluate to elements of G differing on the right by a. Intersecting this automatic relation with $\mathcal{L} \times \mathcal{L}$ shows that $\mathcal{R}_a$ is itself an automatic relation (or asynchronously automatic, if the path system induced by $\mathcal{L}$ is asynchronously bounded).

This proves the following foundational theorem in the theory of automatic groups, which was originally shown in [8]:

Theorem 4.7. *A group G generated by a finite set A is automatic if and only if there is a regular normal form $\mathcal{L} \subseteq A^*$ for G which induces a synchronously bounded path system on* $\mathrm{Cay}(G, A)$*. G is asynchronously automatic if and only if there is a regular language $\mathcal{L} \subseteq A^*$ which induces an asynchronously bounded path system on* $\mathrm{Cay}(G, A)$*.*

Consequently, the standard way to show that a group is automatic (or asynchronously automatic) is to find a bounded path system on the group and show that it is given by a regular language. The 'multiplier automata' recognizing the relations $\mathcal{R}_a$ are rarely specified in practice.

It should be noted here that one can show that the image of a regular language $\mathcal{L}$ in a group (sometimes called a **rational subset**) has prefixes at bounded distance from $E(\mathcal{L})$, using an approach very similar to that used to show that automatic relations fellow-travel in groups. We will use this fact to prove a crucial lemma in Section 5.

Lemma 4.8. *Let G be a group generated by a finite set A, and let $\mathcal{L} \subseteq A^*$ be a regular language. Let $v \in \mathcal{L}$ be a word in $\mathcal{L}$, and denote by v_i the product in A^* of the first i letters of the word v. Then for any prefix v_i of V, we have $d_{G,A}(E(v_i), E(\mathcal{L})) \leq N$, where N is the number of states in an automaton recognizing $\mathcal{L}$.*

Proof: Since v_i is a prefix of an accepted word, there is a directed path in the state graph of an automaton recognizing $\mathcal{L}$ from the state upon reading v_i to an accept state. This path can be chosen to cross no more than N edges. If w is the label of this path, then we have $|w| \leq N$ and $v_i w \in \mathcal{L}$ - hence $d_{G,A}(E(v_i), E(\mathcal{L})) \leq N$, as desired. □

Before we move on from automatic groups, we will first note a minor generalization of a well-established fact about automatic subgroups. It is well-known that automaticity is preserved by finite-index supergroups; that is, G is automatic (or asynchronously automatic) if it has a finite-index subgroup which is automatic (or asynchronously automatic). This is easy to prove, and the proof does not rely on the fact that the automatic subgroup is a *group*. Rather, it easily generalized to the case in which a regular language gives a *partial* path system on $\mathrm{Cay}(G, A)$ which has coarsely dense image in G (here, we denote by A_k the set of words over A of length at most k):

Theorem 4.9. *Let G be a group with finite generating set A, and let $\mathcal{L} \subset A^*$ be a regular language whose image in G under the evaluation map gives a synchronously (asynchronously) bounded, finite-to-one path system. If $E(\mathcal{L})$ is coarsely dense in G, then there exists $k > 0$ so that $\mathcal{L}A_k$ gives a synchronous (asynchronous) automation for G.*

Proof: Choose k so that $E(\mathcal{L}A_k) = G$, and let $w_1, w_2 \in \mathcal{L}A_k$ be words in A^* such that $E(w_1)$ and $E(w_2)$ differ by a generator (i.e. $d_{G,A}(E(w_1), E(w_2)) = 1$). We wish to show that the paths $E\overrightarrow{w_1}$ and $E\overrightarrow{w_2}$ synchronously (asynchronously) fellow-travel. Since w_1 and w_2 differ from words in $\mathcal{L}$ by at most k letters, we conclude that w_1 and w_2 have prefixes w_1' and w_2' such that $w_1', w_2' \in \mathcal{L}$, and

$$d_{G,A}(E(w_1'), E(w_1)), d_{G,A}(E(w_2'), E(w_2)) \leq k$$

Then $d_{G,A}(E(w_1'), E(w_2')) \leq 2k+1$, and so the synchronous (or asynchronous) distance between $\overrightarrow{w_1'}$ and $\overrightarrow{w_2'}$ must be bounded above by $f(2k+1)$. The final arcs of $E\overrightarrow{w_1}$ and $E\overrightarrow{w_2}$, which join $E(w_1')$ to $E(w_1)$ and $E(w_2')$ to $E(w_2)$, have length bounded above by k and endpoints at distance bounded above by $2k+1$, and so their images are contained in a common neighborhood of radius $4k+1$. Thus the paths $E\overrightarrow{w_1}$ and $E\overrightarrow{w_2}$ synchronously (or asynchronously) fellow-travel with constant $\max\{f(2k+1), 4k+1\}$. The evaluation $E : \mathcal{L}A^k \to G$ is clearly still finite-to-one, and so $\mathcal{L}A^k$ gives a synchronous (or asynchronous) automatic structure on G. $\square$

Finally, we record the following classic result from [8] on hyperbolic groups, for later reference:

Theorem 4.10. *Let G be a finitely generated hyperbolic group, with finite generating set A. Then the set $\mathcal{L} \subset A^*$ of geodesic words is a regular language. Furthermore, G is biautomatic with respect to $\mathcal{L}$.*

Note that the asynchronous automaticity of hyperbolic groups is already guaranteed by combining Lemma 3.7, Theorem 3.8, Proposition 4.5, and Theorem 4.9. This will result in an automation furnished by a regular language of uniform-quality quasi-geodesics. The additional detail asserted by the theorem above is that the language of *geodesic* words (rather than just the language of *locally* geodesic words) in a hyperbolic group is always regular.

Chapter 5: Factor Languages

In this chapter, we define the central object which we shall use to describe asynchronous automatic structures for groups quasi-isometric to products of non-elementary hyperbolic spaces - the **factor-language system** - and show that it does yield an asynchronous automation. This definition is recursive, with the idea being that we will introduce one hyperbolic factor H to the product at a time, always concatenating a language which embeds along the coordinate factor σ_H quasi-isometrically. The language corresponding to the product of the previous factors is allowed to be embedded in this larger product more loosely, being required only to map quasi-isometrically to the corresponding factors *through the canonical projection*. This generality is inspired by the embedding of the factor languages in the asynchronous automatic structure studied in Chapter 7, where one factor language embeds quasi-isometrically to an $\mathbb{H}^2$ coordinate factor while the other (corresponding to the tree factor T) is 'spread out' in the sense that its projection to the $\mathbb{H}^2$ factor is not bounded.

We begin with the definition of a factor-language system. This definition has a recursive flavor and specifies a slightly more general object than the one we will care about in our main theorem, but this somewhat contrived construction has been chosen to simplify the subsequent proofs.

Definition 5.1 (Factor Language). *Let G be a group with finite generating set $A \subseteq G$,* *let* $\hat{X} = \prod_{i=1}^{N} H_i$ *be a metric space which splits as a finite product of non-elementary hyperbolic geodesic spaces, let* $\hat{\phi} : G \to \hat{X}$ *be a quasi-isometry. Now suppose that* $\mathcal{L}_1, \mathcal{L}_2 \subset A^*$ *are regular languages, and* $\phi : E(\mathcal{L}_1\mathcal{L}_2) \to X$*, where* $X = \prod_{i=1}^{M} H_i$ *with* $M \leq N$*, is a quasi-isometry which agrees with* $\pi_X\hat{\phi}$ *on* $E(\mathcal{L}_1\mathcal{L}_2)$*. We say* $(G, A, X, \mathcal{L}_1, \mathcal{L}_2, \phi)$ *is a* ***factor-language system*** *if the following conditions hold:*

1. *The Hausdorff distance between* $\phi E(\mathcal{L}_2)$ *and the coordinate factor* $\sigma_M \cong H_M$ *is finite, and the paths given by words in* $E(\mathcal{L}_2)$ *are uniform-quality quasi-geodesics in* G*.*

2. *One of the following holds for* $X' = \prod_{i=1}^{M-1} H_i$*:*

   (a) $\mathcal{L}_1 = \{\epsilon\}$ *and* X' *is hyperbolic*

   (b) *There exist regular languages* $\mathcal{L}_1', \mathcal{L}_2' \subseteq A^*$ *such that* $\mathcal{L}_1 = \mathcal{L}_1'\mathcal{L}_2'$*, and* $(G, A, X', \mathcal{L}_1', \mathcal{L}_2', \pi_1\phi)$ *is a factor-language system*

If the other elements are understood, we say $(\mathcal{L}_1, \mathcal{L}_2)$ *is a* ***factor-language pair****.*

In the following, it will be convenient if ϕ takes edges in $\mathrm{Cay}(G, A)$ to geodesics in X, so that the image of a rectifiable curve under ϕ is a rectifiable curve with bounded length. We shall assume this without loss of generality.

Lemma 5.2. *Let* $(G, A, X, \mathcal{L}_1, \mathcal{L}_2, \phi)$ *be a factor-language system. Then there exists* $\alpha > 0$ *such that* $d_{H_1}(\pi_1\phi E(wv)), \pi_1\phi E(w)) < \alpha$ *for all* $w \in \mathcal{L}_1$ *and all* $v \in \mathcal{L}_2$*.*

Proof: By assumption, $\phi(E(v))$ is at finite Hausdorff distance from σ_2 and thus we may take $C > 0$ so that $\phi(E(v))$ lies in $N_C(\sigma_2)$, the C-neighborhood of σ_2. Hence,

$$d_{H_1}(\pi_1\phi E(v), \pi_1\phi E(e)) \leq C$$

Furthermore, the elements $E(w) \in E(\mathcal{L}_1)$ induce uniform-quality quasi-isometries on $\hat{X}$ by composing the left-action of G on itself with $\hat{\phi}$ on the left and (a choice of) its quasi-inverse on the right. Letting $X = X' \times H_M$ as in the definition, Corollary 3.14 shows that each of these $E(w)$ induces a quasi-isometry on X which splits as a product of quasi-isometries $E(w)_1 \times E(w)_2 : X' \times H_M \to X' \times H_M$, up to bounded error. Without loss of generality, we will take these all to be (C, C)-quasi-isometries, and we will also take the error bound to be C. Thus,

$$\begin{aligned} d_{H_1}(\pi_1\phi E(wv), \pi_1\phi E(w)) &= d_{H_1}(\pi_1\phi(E(w)E(v)), \pi_1\phi(E(w)E(e))) \\ &= d_{H_1}(\pi_1(E(w) \cdot \phi E(v)), \pi_1(E(w) \cdot \phi E(e))) \\ &\leq d_{H_1}(E(w)_1 \cdot \pi_1\phi E(v), E(w)_1 \cdot \pi_1\phi E(e)) + 2C \\ &\leq C d_{H_1}(\pi_1\phi E(v), \pi_1\phi E(e)) + 3C \\ &\leq C^2 + 3C \end{aligned}$$

Taking $\alpha > C^2 + 3C$ completes the proof. □

Lemma 5.3. *Let $(\mathcal{L}_1, \mathcal{L}_2)$ be the factor-language pair in the system described above. Then there exists $\eta > 0$ such that, for any $w \in \mathcal{L}_1$ and $v \in \mathcal{L}_2$, the subpath of $\phi E\overrightarrow{wv}$ joining $\phi E(w)$ to $\phi E(wv)$ has image contained in an η-neighborhood of $\sigma_2^{\phi E(w)}$.*

Proof: Let $\overrightarrow{\gamma}$ denote the subpath in question, let γ denote its image in X, and let v_i denote the i-th prefix of v for all $0 \leq i \leq |v|$. By assumption, $\overrightarrow{\gamma}$ is a piecewise-geodesic interpolation of the discrete path $(\gamma_0, ..., \gamma_{|v|})$, where $\gamma_i = \phi E(wv_i)$ for $0 \leq i \leq |v|$. From Lemma 5.2, we know that there is a uniform constant $\alpha > 0$ such that

$$d_{X'}(\pi_1\phi E(w), \pi_1\phi E(wv')) < \alpha.$$

for all $v' \in \mathcal{L}_2$. Assuming without loss of generality that C is greater than the number of states in an automaton recognizing $\mathcal{L}_2$, we see by Lemma 4.8 that every point γ_i is joined by a piecewise-geodesic arc to an element of $\phi E(w\mathcal{L}_2)$, and this arc can be chosen to consist of no more than C geodesic pieces. These pieces are the images under ϕ of edges in the Cayley graph $\text{Cay}(G, A)$, which have length 1, and consequently have length not exceeding $2C$. Thus the entire arc has length bounded above by $2C^2$, and every point on it is at distance at most C^2 from an endpoint lying in the α-neighborhood of $\sigma_2^{\phi E(w)}$. Letting $\eta = \alpha + C^2$, every point of γ is contained in the η-neighborhood of $\sigma_2^{\phi E(w)}$, as desired. $\square$

Theorem 5.4. *Let $(G, A, X, \mathcal{L}_1, \mathcal{L}_2, \phi)$ be a factor-language system such that $E(\mathcal{L}_1\mathcal{L}_2)$ is at bounded Hausdorff distance from* $\text{Cay}(G, A)$*. Then there exists $k > 0$ such that $\mathcal{L} = \mathcal{L}_1\mathcal{L}_2A^k$ gives an asynchronous automatic structure for G.*

Proof: We will show that the paths given by words in $\mathcal{L}_1\mathcal{L}_2$ give an asynchronously bounded path system in G through the evaluation map E, and the desired result will follow from Theorem 4.9.

Let $w_1, w_2 \in \mathcal{L}_1$ and $v_1, v_2 \in \mathcal{L}_2$. We wish to show that there exists a function f so that

$$d_{Async}(\phi E\overrightarrow{w_1v_2}, \phi E\overrightarrow{w_2v_2}) \leq f(D),$$

where $D = d(\phi E(w_1v_1), \phi E(w_2v_2))$. Letting $X = \prod_{i=1}^{N} H_i$ be the associated decomposition of X as a product of non-elementary hyperbolic spaces, we shall assume by way of induction that factor-language systems give asynchronously bounded path systems if X is a product of $N-1$ or fewer non-elementary hyperbolic geodesic spaces. The basis for this induction is the case in which $N = 1$, X is hyperbolic, and $\mathcal{L}_1\mathcal{L}_2 = \mathcal{L}_1$ gives uniform-quality quasi-geodesics in G, which are asynchronously bounded by the

Morse Lemma.

As a first step, we show that $\phi E\overrightarrow{w_1}$ and $\phi E\overrightarrow{w_2}$ asynchronously fellow-travel with bound linear in D. This will follow from the inductive hypothesis above, if we can show that $d_{G,A}(E(w_1), E(w_2))$ is bounded by a linear function of D. Since $D = d(\phi E(w_1 v_1), \phi E(w_2 v_2))$, we have

$$d_{H_1}(\pi_1 \phi E(w_1 v_1), \pi_1 \phi E(w_2 v_2)) \leq D,$$

and by Lemma 5.2, we have

$$d_{H_1}(\pi_1 \phi E(w_1 v_1), \pi_1 \phi E(w_1)), d_{H_1}(\pi_1 \phi E(w_2 v_2), \pi_1 \phi E(w_2)) \leq \alpha$$

for some fixed constant $\alpha > 0$. We may therefore conclude that

$$d_{H_1}(\pi_1 \phi E(w_1), \pi_1 \phi E(w_2)) \leq D + 2\alpha$$

by the triangle inequality. Since $\pi_1 \phi : E(\mathcal{L}_1) \to H_1$ is a quasi-isometry, this shows that $d_{G,A}(E(w_1), E(w_2))$ is bounded by a function of D, as desired.

Now we need to show that the arcs joining $\phi E(w_1)$ to $\phi E(w_1 v_1)$ and $\phi E(w_2)$ to $\phi E(w_2 v_2)$ asynchronously fellow-travel with bound a function of D. Denoting these arcs by $\overrightarrow{\gamma_1}$ and $\overrightarrow{\gamma_2}$, respectively, and denoting their images by γ_1 and γ_2, we see from Lemma 5.3 that there exists a constant η, independent of $\overrightarrow{\gamma_1}$ and $\overrightarrow{\gamma_2}$, such that γ_i lies in the η-neighborhood of $\sigma_2^{\phi E(w_i)}$, for $i = 1, 2$. Let β be the Hausdorff distance between $\sigma_2^{\phi E(w_1)}$ and $\sigma_2^{\phi E(w_2)}$. By the preceding paragraph, β is bounded by a function of D. Furthermore, both γ_1 and γ_2 are contained in an $2\eta + \beta$-neighborhood of $\sigma_2^{\phi E(w_1)}$. By Theorem 3.12, this neighborhood is δ'-hyperbolic, where δ' is linear in

β and hence linear in D. Since $E\overrightarrow{v_1}$ and $E\overrightarrow{v_2}$ (hence their translates $E(w_1)E\overrightarrow{v_1}$ and $E(w_2)E\overrightarrow{v_2}$) are uniform-quality quasi-geodesics, we see by Lemma 3.7 that $\overrightarrow{\gamma_1}$ and $\overrightarrow{\gamma_2}$ must asynchronously fellow-travel with bound linear in both δ' and the maximal distance between corresponding endpoints - all of which we have shown to be bounded by linear functions of D.

Since $E\overrightarrow{w_1}$ and $E\overrightarrow{w_2}$ asynchronously fellow travel with bound $f(D)$, and $\overrightarrow{\gamma_1}$ and $\overrightarrow{\gamma_2}$ asynchronously fellow-travel with bound $f(D)$, the asynchronous distance between their respective concatenations, $\phi E\overrightarrow{w_1 v_1}$ and $\phi E\overrightarrow{w_2 v_2}$, is also bounded by $f(D)$. $\square$

Chapter 6: Level-Deterministic Actions

In this section, we give a sufficient condition for an isometric action of a finitely generated group $G = \langle A \rangle$ on a pointed geodesic space (X, o) to admit regular languages of paths which do not depart a fixed neighborhood of the basepoint o. First, we shall discuss the application of such languages to the construction of the preceding section.

Definition 6.1. *Let (G, A) be a group and finite generating set, and let (X, o) be a pointed geodesic space on which G acts cocompactly by isometries. Let $r > 0$ be a real number. We say the action $(G, A) \curvearrowright (X, o)$ has an* ***r-coarse regular stabilizer*** *if the language*

$$\mathcal{L}_r = \{w \in A^* \mid d_X(o, E(w_i) \cdot o) \leq r \text{ for all prefixes } w_i \text{ of } w\}$$

is regular.

Note that 'coarse regular stabilizers' are distinct from the idea of coarse stabilizers used in [20]. For one thing, the languages $\mathcal{L}_r$ only give sub*sets* of G under evaluation, rather than subgroups. For another, these languages consist of words whose *entire trajectories* in X remain within a fixed neighborhood, rather than just words which evaluate to small translations of o.

The utility of this property can be seen in the case that X is a hyperbolic factor in a product of hyperbolic spaces as in Section 5, as is shown in the following theorem:

Theorem 6.2. *Let G be a group acting geometrically on a pointed, finite product $(X, o) = (\prod_i^N H_i, o)$ of non-elementary hyperbolic geodesic spaces, preserving the factors, and suppose that the action of G on each factor H_i has an r-coarse regular stabilizer. Then G is asynchronously automatic if $r \geq D$, where D is a constant depending only on the action $(G, A) \curvearrowright (X, o)$.*

Proof: For $i \in \{1, \cdots, N\}$, denote by $\mathcal{L}_r^i$ the coarse regular stabilizer with radius r for the action of G on H_i, and let C be the coarse density constant of the orbit quasi-isometry associated with the geometric action of G on X. We begin by showing that, for large enough r, we have $E(\bigcap_{i \neq j} \mathcal{L}_r^i) \cdot o$ at finite Hausdorff distance from σ_j, the j-th coordinate factor, for all indices j.

Let us now fix j. Since $N(C, \sigma_j)$, the C-neighborhood of σ_j, has intersection with $G \cdot o$ which is coarsely dense in $N(C, \sigma_j)$, we obtain the desired result if we choose r large enough that $g \in \mathcal{L}_r^i$ for every g such that $g \cdot o \in N(C, \sigma_j)$. Fix such an element g. We show that we may choose r so that $N(r, \sigma_j)$ contains the entire image of a path in G from the identity to g, and the desired result follows.

We accomplish this by noting that $g \cdot o$ and o are joined in X by a $(1, C)$ quasi-geodesic $\overrightarrow{\gamma}$ whose image lies in $N(C, \sigma_j)$. We can construct $\overrightarrow{\gamma}$ by taking any geodesic in σ_j connecting o to a closest-point projection of $g \cdot o$ to σ_j, and then joining the endpoint of that geodesic to $g \cdot o$ by a path of length not exceeding C. By Proposition 3.4, there is a path in G which asynchronously fellow-travels the image of $\overrightarrow{\gamma}$ under a quasi-inverse of the orbit map, and we see that this path must be a quasi-geodesic with quality depending only on the action $(G, A) \curvearrowright (X, o)$. Applying the orbit map to this path yields a path in X lying in a B-neighborhood of the image of $\overrightarrow{\gamma}$ (hence in a B-neighborhood of $N(C, \sigma_j)$), where B depends only on $(G, A) \curvearrowright (X, o)$. Choosing $D = C + B$ ensures that $\bigcap_{i \neq j} \mathcal{L}_r^i$ contains a representative for every $g \in G$ such that $g \cdot o \in N(C, \sigma_j)$ whenever $r \geq D$, as desired. Furthermore, we have also shown that $\bigcap_{i \neq j} \mathcal{L}_r^i$ contains a representative for g that evaluates to a uniform-quality quasi-geodesic in G, with quality determined by the action.

Since $\bigcap_{i \neq j} \mathcal{L}_r^i$ is regular (as a finite intersection of regular languages), prefix-closed, and has image in X Hausdorff-close to a coordinate factor, it is almost a (second) factor language in the sense of Definition 5.1. To align with this definition, our language must also consist of words which evaluate to uniform-quality quasi-geodesics. This is easy enough to arrange, since the paths in $\bigcap_{i \neq j} \mathcal{L}_r^i$ have images lying entirely in a fixed neighborhood of σ_j, which is hyperbolic by Theorem 3.12. Consequently, these paths lie entirely in a hyperbolic subspace H' of G - namely, the preimage of the previous neighborhood under the orbit map. Invoking Theorem 4.5, we note that, for any $k \in \mathbb{N}$, the k-local geodesic words in G form a regular language $\mathcal{L}_{LG(k)}$ since they are locally recognized. Local geodesics in G whose images lie in H' are local geodesics in H', and therefore they are in fact global quasi-geodesics in H' by Theorem 3.8. Consequently,

$$\mathcal{L}_j := \mathcal{L}_{LG(k)} \cap \bigcap_{i \neq j} \mathcal{L}_r^i$$

is a regular language consisting of uniform-quality quasi-geodesics. Since we have chosen r large enough to ensure that $\bigcap_{i \neq j} \mathcal{L}_r^i$ contains uniform-quality quasi-geodesic representatives for a coarsely dense subset of a neighborhood of σ_j, this intersection still satisfies the remaining hypotheses for factor languages when k is sufficiently large.

By Theorem 5.4, G is asynchronously automatic, with normal form given virtually by the concatenation of the $\mathcal{L}_j$. $\square$

Now that we have an intermediary condition that gives us the desired factor-language system, we will produce a condition on the action $(G, A) \curvearrowright (X, o)$ which is similar in flavor to the way that reducibility of such an action on a product follows from discreteness of the action on the factors. Inspired by this comparison, we call this property **level-determinism** and will define it after establishing the following notation: we associate to each $g \in G$ a map $g|_* : A^* \to X$ by defining

$$g|_*(w) := gE(w) \cdot o$$

for all $w \in A^*$. Note that the maps $g|_*$ are K-Lipschitz, where K is the maximum distance between o and $a \cdot o$ for any generator of $a \in A$,

Letting $A_n = \bigcup_{i=0}^{n} A^n$, we denote by $g|_n$ the restriction of $g|_*$ to A_n. More generally, we will write $g|_Y$ for the restriction of $g|_*$ to any subset $Y \subset A^*$. For $m \geq n$, denote the restriction to A_n of maps $A_m \to X$ by $R_n^m : X^{A_{n+1}} \to X^{A_n}$. Finally, we will make special use of the following subsets of A^* parametrized by $g \in G$, $r \in \mathbb{R}^{\geq 0}$, and $n \in \mathbb{Z}^{\geq 0}$:

$$S(g, r, n) := A_n \cup \bigcup_{g|_*(a) \in N(r,o)} aA^n$$

In other words, $S(g, r, n)$ is A_n plus those words in A^{n+1} that begin with a letter $a \in A$ such that $g|_*(a)$ is within the closed r-neighborhood of the origin o.

Definition 6.3. *Let G be a group with finite generating set A, and let $(G, A) \curvearrowright (X, o)$ by isometries for some pointed geodesic space (X, o). Let $r \geq 0$ be a nonnegative real* *number, and let $n \geq 0$ be a nonnegative integer. We say the action $(G, A) \curvearrowright (X, o)$ is* $(\mathbf{r}, \mathbf{n})$ ***level-deterministic*** *if the following conditions hold:*

1. $\#(G \cdot o \cap N(s, o)) < \infty$ *for all* $s > 0$

2. *For all $g, h \in G$ such that $g \cdot o$, $h \cdot o \in N(r, o)$, we have* $g|_{S(g,r,n)} = h|_{S(h,r,n)} \Leftrightarrow g|_n = h|_n$

The second condition may seem somewhat contrived, but more or less means that the elements of G, viewed as maps $A^{n+1} \to X$, are fully determined by their behavior on A^n. Isolating the potentially smaller set $S(g, r, n)$ is necessary to allow for a special case that will be discussed in Section 7. In turn, the condition that $g|_n$ determines $g|_{n+1}$ is a weakening of the condition that $g|_n$ determines $g|_*$, which is a kind of uniform discreteness for G as a set of maps $A^* \to X$.

The idea behind this setup is to produce a sufficient condition for the set $\mathcal{L}_r$ in Definition 6.1 to be a regular subset of A^*, by explicitly producing an automaton recognizing it. The states of the automaton are the points in $G \cdot o \cap N(r, o)$, together with a choice of embedding of A_n that sends the empty word to the given point. Condition 1 of Definition 6.3 implies there will only be finitely many states. Condition 2 then ensures that the transition function is well-defined for all $a \in A$ that should be accepted by $\mathcal{L}_r$.

Theorem 6.4. *Let G be a group with finite generating set A, and let $(G, A) \curvearrowright (X, o)$ by isometries for some pointed geodesic space (X, o). If this action is (r, n) level-deterministic for some $r \in \mathbb{R}^{\geq 0}$, $n \in \mathbb{Z}^{\geq 0}$, then the action has an r-coarse regular stabilizer.*

Proof: As mentioned above, we will explicitly construct a finite-state automaton recognizing $\mathcal{L}_r$. Given $r \in \mathbb{R}^{\geq 0}$ and $n \in \mathbb{Z}^{\geq 0}$, the state set for this automaton will be

$$Q := \{g|_n \mid g \cdot o \in N(r, o)\}$$

Since the action is level-deterministic and the maps $g|_n$ are uniformly Lipschitz, there are only finitely many possible images in $G \cdot o$ for each $w \in A_n$, once the image of the empty word is fixed. This shows that Q is finite. The start state is $e|_n$, where e is the identity in G, and all states are accept states.

The transition function for this automaton is defined as follows: for $g|_n \in Q$ and $a \in A$, we pick any representative $g \in G$ which restricts to $g|_n$ as a map $A_n \to X$. We then define the state transition

$$(g|_n, a) \mapsto (gE(a))|_n$$

for all a such that $g|_*(a) \in N(r, o)$. Note that $g|_n = h|_n \Rightarrow g_*(a) = h_*(a)$ as long as $n > 0$, so the $a \in A$ for which this transition is added are independent of our choice of representative (that is, $g|_n = h|_n \Rightarrow S(g, r, n) = S(h, r, n)$ always holds if $n > 0$). If $n = 0$, then this is an assumption of Definition 6.3 which must be verified independently.

We now show this transition is well-defined - that is, it does not depend on our choice of representative g. Let $\hat{a} : A_n \to A_{n+1}$ be the map that sends $w \mapsto aw$ for all

$w \in A_n$. Then we have the following for all $w \in A_n$:

$$\begin{aligned}(gE(a))|_n(w) &= gE(a)E(w) \cdot o \\ &= gE(aw) \cdot o \\ &= gE(\hat{a}(w)) \cdot o \\ &= g|_{n+1}(aw)\end{aligned}$$

Now if g' is any other representative for $g|_n$, we have $g'|_{n+1}(aw) = g|_{n+1}(aw)$ since $aw \in S(g, r, n) = S(g', r, n)$. Thus $(gE(a))|_n = (g'E(a))|_n$, as desired.

It is now straightforward to show that this automaton in fact recognizes $\mathcal{L}_r$, the set of words in A^* whose corresponding paths in X remain within distance r of o. By way of induction on word length, suppose that $w \in A^*$ describes a path in X which does not depart $N(r, o)$, and let $a \in A$. When the automaton finishes reading w, it must be in state $E(w)|_n$ (this can either be shown directly or included in the inductive hypothesis). We have $wa \in \mathcal{L}_r$ if and only if $E(w)|_*(a) \in N(r, o)$, which is true if and only if there is an outgoing transition from the state $E(w)|_n$ labeled a. That is, wa is accepted by our automaton if and only if $wa \in \mathcal{L}_r$. □

Combining Theorems 6.2 and 6.4, we obtain the following corollary:

Corollary 6.5. *Let G be a group acting geometrically on a pointed, finite product $(X, o) = (\prod_i H_i, o)$ of non-elementary hyperbolic geodesic spaces, preserving the factors, and suppose that the action of G on each factor H_i is (r, n) level-deterministic. Then G is asynchronously automatic if $r \geq D$, where D is a constant depending only on the action $(G, A) \curvearrowright (X, o)$.*

The next section covers a special case of this approach, where the factor language in question is given by an entire subgroup. In this case, choosing the right generators allows us to take $D = r = n = 0$.

Chapter 7: Application: HNN Extensions of Hyperbolic Groups

In this chapter, we illustrate the tools developed in the preceding chapters by finding a factor-language system for a broad class of groups: HNN extensions by commensurated subgroups of hyperbolic groups. That these groups are asynchronously automatic was recently observed by Hughes and Valiunas in [15], a note following the authors' groundbreaking work in [16]. The family of languages they provide coincides coarsely with the one obtained here, but it may still be of interest to see how these languages fit within the more general pattern established in the present work.

The groups described in [16], which we shall refer to as the **Hughes-Valiunas groups**, are obtained as commensurated HNN extensions of lattices in $PSL(2, \mathbb{R})$. We omit the details here - the curious reader may consult the detailed treatment in [16] to confirm that these groups satisfy the conditions of Theorem 7.5.

First, we define the HNN-extension:

Definition 7.1. *Let G be a group, and suppose $H, K \leq G$ are subgroups of G with $\tau : H \to K$ an isomorphism. The* ***HNN extension*** *of G with respect to τ, denoted*

$G^\star_\tau$, or $G^\star$ if τ is understood, is the quotient

$$G * \langle t\rangle / \langle\langle tht^{-1}\tau(h)^{-1} \ |h \in H\rangle\rangle$$

where $\langle\langle\cdot\rangle\rangle$ denotes the normal closure of a subset.

Alternatively, we can think of the HNN extension in terms of group presentations. If G is presented by generators A and relators R, then a presentation for the HNN extension $G^\star_\tau$ is obtained by adding t to the generators and adding the relation $h^t = \tau(h)$ for all $h \in H$.

Let T_H and T_K be coset transversals of H and K, respectively. Every element of $G^\star$ can be written as an alternating product

$$g_0 t^{\epsilon_0} g_1 t^{\epsilon_1} \cdots g_n t^{\epsilon_n}$$

where $g_i \in G$ and $\epsilon_i \in \{-1, 1\}$ for all $i = 0, 1, \cdots, n$. Now for any $i \in \{0, 1, \cdots, n\}$, we can use the HNN presentation to rewrite $g_i t^{\epsilon_i} = c_i t^{\epsilon_i} g_i'$, where $g_i' \in G$, $c_i \in T_H$ if $\epsilon_i = -1$, and $c_i \in T_K$ if $\epsilon_i = 1$. Iterating this process from left to right produces an expression of the form

$$c_0 t^{\epsilon_0} c_1 t^{\epsilon_1} \cdots c_n t^{\epsilon_n} g$$

where $g \in G$ and the c_i satisfy the preceding condition in relation to ϵ_i. If we pick representative words for the c_i and write g in a normal form for G, we obtain a normal form for $G^\star$:

Definition 7.2. *Let $G^\star$ be an HNN extension of a group $G = \langle A\rangle$ with isomorphic subgroups H and K, and let T_H and T_K consist of unique representatives in A^* for coset transversals of H and K, respectively. A* ***left HNN normal form*** *for $G^\star$ is* *given by freely reduced words which match the regular expression*

$$(\mathtt{T_H t^{-1}} \vee \mathtt{T_K t})^* \mathcal{L}$$

where $\mathcal{L}$ is a normal form over A^* for G. Given a word in this normal form, the initial string of coset representatives and powers of t is called the ***prefix****.*

Up to choice of representatives for coset transversals and normal form for G, the left HNN normal form gives unique representatives for elements of $G^\star$. This unique normal form can be used to show that the canonical map $G \to G * \langle t \rangle \to G^\star$ is injective - that is, $G \leq G^\star$. This embedded copy of G can be used to produce an action of $G^\star$ on a tree.

HNN extensions play a central role in the development of *Bass-Serre theory*, which characterizes groups acting on simplicial trees by graph automorphisms. In particular, there is associated to each HNN extension $G^\star$ a regular tree T, called the **Bass-Serre tree**, on which $G^\star$ acts by graph automorphisms (i.e. isometries, if T is given the graph metric).

Definition 7.3. *Let $G^\star$ be an HNN extension of a group G with isomorphic subgroups H and K. The* ***Bass-Serre tree*** *associated to $G^\star$ is the graph T whose vertices are the left cosets of G in $G^\star$ and whose edges are of the form (xG, xtG) or $(xG, xt^{-1}G)$.*

If $m = [G : H]$ and $n = [G : K]$, then T is an $m + n$-regular tree.

Proposition 7.4. *Let $G^\star$ be an HNN extension. Then $G^\star$ acts on its Bass-Serre tree by graph automorphisms, with the action on the vertex set given by $g \cdot xG = gxG$.*

Under this action, the vertex G is clearly stabilized by the subgroup $G \leq G^\star$. This shows, among other things, that the prefixes of the left HNN normal form uniquely address the vertices of the Bass-Serre tree. In other words, the prefixes are a coset transversal of G in $G^\star$. More generally, one can show that the vertex stabilizers for the action of $G^\star$ on the Bass-Serre tree are precisely the conjugates of G in $G^\star$.

HNN extensions have been a fruitful source of interesting examples in group theory. They were used in the 1960's by Britton [5] to give a simplified proof of the famous result that that there exist finitely generated groups with unsolvable word problem. More recently, Leary and Minasyan [19], as well as Hughes and Valiunas [16], used carefully constructed HNN extensions to provide negative solutions to a pair of open problems in geometric group theory. The Hughes-Valiunas examples, as was mentioned above, are HNN extensions of certain lattices in the hyperbolic plane, and they furnish examples of hierarchically hyperbolic groups which are not biautomatic. The Leary-Minasyan groups are HNN extensions of finite-rank free abelian groups, some of which can be shown to be CAT(0) but not biautomatic.

Both cases are currently potential candidates to produce a group which is automatic but not biautomatic, and both use a similar strategy. Namely, they form HNN extensions by embedding the base group G in a larger isometry group and then conjugating it by some isometry which commensurates G. Since the stable letter acts by isometries on the same space as G, the result is a group which is quasi-isometric to the product of a finite-valence Bass-Serre tree and a space with nice geometry (Euclidean n-space, or the hyperbolic plane), with a geometric action that assures the desired properties (the CAT(0) condition, or hierarchical hyperbolicity). While the potential automaticity of such groups remains an open question, we can use Theorem 5.4 to show that this kind of construction produces a group which is *asynchronously* automatic, whenever the base group is non-elementary hyperbolic.

Theorem 7.5. *Let X be a non-elementary proper hyperbolic geodesic space, and suppose $G \leq \text{Isom}(X)$ acts geometrically on X. For any element $\tau \in \text{Isom}(X)$ and subgroup $H \leq G$, the HNN-extension*

$$G^{\star} := \langle G, t \mid g^t = g^{\tau}, g \in (H^{\tau} \cap G)^{\tau^{-1}} \rangle$$

acts geometrically on $T \times X$, where T is the Bass-Serre tree.

Furthermore, if $(H^{\tau} \cap G)$ and $(H^{\tau} \cap G)^{\tau^{-1}}$ have finite index in G, then $(G^{\star}, A \cup \{t\}, T, X, \mathcal{L}_1, \mathcal{L}_2, \phi)$ is a factor-language system, where A is a finite set of generators for G, $\phi : G \to T \times X$ is the orbit map, $\mathcal{L}_1 \subseteq (A \cup \{t\})^*$ gives the prefixes of the left HNN normal form for $G^{\star}$, and $\mathcal{L}_2 \subseteq A^*$ is the set of geodesic words in G. In particular, $G^{\star}$ is asynchronously automatic.

Proof: To simplify notation, we will refer to the isomorphic subgroups of G as H and K.

$G^{\star}$ acts on the Bass-Serre tree T by isometries, and also acts by isometries on X by letting the stable letter t act as the isometry τ. This action is clearly cocompact (it is cocompact on X and transitive on T), and it is also properly discontinuous, since T is discrete and the vertex stabilizers of the action of $G^{\star}$ all act geometrically on X.

By the Švarc-Milnor lemma, the orbit map $\phi : G \to T \times X$, with respect to any basepoint (v_0, x_0), is a quasi-isometry when $T \times X$ is proper. This will occur precisely when T is of finite valence - that is, when H and K both have finite index in G, as in the hypothesis.

We wish to verify that $(\mathcal{L}_1, \mathcal{L}_2)$ is a factor-language pair. We see immediately that $\mathcal{L}_2$ meets the criteria: $\mathcal{L}_2$ is regular by Theorem 4.10, consists of uniform quasi-geodesics, and is prefix-closed; furthermore, $E(\mathcal{L}_2) = G$, the stabilizer of the basepoint

v_0 in T, and so $\phi E(\mathcal{L}_2) = \{v_0\} \times G \cdot x_0$, which is at bounded Hausdorff distance from $\sigma_2 = \sigma_2^{(v_0,x_0)}$.

$\mathcal{L}_1$ consists of prefixes for the left HNN normal form, so we know $\mathcal{L}_1$ is $(|w|+1)$-prefix closed, where w is a maximal-length word in $T_H \cup T_K$. Since $\mathcal{L}_1$ is a regular language (it is given by a regular expression), it remains only to show that $\pi_1\phi : E(\mathcal{L}_1) \to H_1$ is a quasi-isometry.

Let T_H and T_K be the representatives of coset transversals that determine $\mathcal{L}_1$, and let $g, h \in E(\mathcal{L}_1)$. Since ϕ is a quasi-isometry and π_1 is 1-Lipschitz, we only need to show that $d_{G,A}(g,h) \leq K d_T(\pi_1\phi(g), \pi_1\phi(h)) + C$ for some uniform K, C.

Since $g, h \in E(\mathcal{L}_1)$, we can write g and h as words in $\mathcal{L}_1$:

$$g = c_1 t^{\epsilon_1} \cdots c_m t^{\epsilon_m}$$

$$h = d_1 t^{\epsilon'_1} \cdots d_n t^{\epsilon'_n}.$$

Assume without loss of generality that $c_m \neq d_n$. Then

$$\begin{aligned} d_{G,A}(g,h) &= |gh^{-1}| \\ &\leq |g| + |h| \\ &\leq (\sum_{i=0}^{m} |c_i| + m) + (\sum_{i=0}^{n} |d_n| + n) \\ &\leq M(m+n) \end{aligned}$$

where $M = \max\{|w| \mid w \in T_H \cup T_K\}$. On the other hand, we have $\pi_1\phi(g) = gG$ and $\pi_1\phi(h) = hG$ (recall that vertices of T are cosets of G in $G^\star$). Using the same expression for g, we see

$$(c_1 t^{\epsilon_1} \cdots c_i t^{\epsilon_i} c_{i+1} t^{\epsilon_{i+1}}) t^{-\epsilon_{i+1}} G = c_1 t^{\epsilon_1} \cdots c_i t^{\epsilon_i} c_{i+1} G$$
$$= c_1 t^{\epsilon_1} \cdots c_i t^{\epsilon_i} G$$

Hence $\pi_1 \phi(c_1 t^{\epsilon_1} \cdots c_{i+1} t^{\epsilon_{i+1}})$ is adjacent to $\pi_1 \phi(c_1 t^{\epsilon_1} \cdots c_i t^{\epsilon_i})$ for all $i \in \{1, \cdots, m-1\}$, and the same is true for the above expansion of h. These adjacencies give paths from gG and hG to the root vertex G, and these paths are geodesic since the prefix of the left HNN normal form is freely reduced (i.e. does not backtrack in T). Since we assumed $c_m \neq d_n$, gG and hH are on different branches of T from the root vertex, the concatenation of these paths is also geodesic: $d_T(\pi_1\phi(g), \pi_1\phi(h)) = m + n$. Thus,

$$d_{G,A}(g, h) \leq M d_T(\pi_1\phi(g), \pi_1\phi(h))$$

as desired. □

Corollary 7.6. *The Hughes-Valiunas groups are asynchronously automatic.*

It should be pointed out that this asynchronous automatic structure is very nearly synchronous. If $g \in G^\star$, we can write g in left HNN normal form:

$$g = E(c_1 t^{\epsilon_1} \cdots c_n t^{\epsilon_n} w)$$

Now if h differs from g by a generator, either $h = ga$ for some $a \in A$, in which case

$$h = E(c_1 t^{\epsilon_1} \cdots c_n t^{\epsilon_n} wa)$$

and a geodesic representative for $E(wa)$ must synchronously fellow-travel $E\overrightarrow{w}$. The only possible complication occurs when $h = gt$, for then

$$\begin{aligned} h &= E(c_1 t^{\epsilon_1} \cdots c_n t^{\epsilon_n} w t) \\ &= E(c_1 t^{\epsilon_1} \cdots c_n t^{\epsilon_n} c_{n+1} t^{\epsilon_{n+1}} w') \end{aligned}$$

and it may be the case that $E\overrightarrow{w'}$ is of very different length than $E\vec{w}'$. Synchronous automaticity can be assured if A, the generating set for G, can be chosen so that $|\tau(h)| = |h|$ for all $h \in H$, up to uniformly bounded error.

Finally, we can show that the machinery of Section 6 applies here as well. Since the action on the Bass-Serre tree is transitive and the images of elements of G are all at distance 0 from the second coordinate factor σ_2, we have $D = 0$ in the conditions of Corollary 6.5. We see that the action on the Bass-Serre tree is $(0, 0)$ level-deterministic by the following argument:

Condition 1 of Definition 6.3 is clearly satisfied, so we turn to condition 2. If $g \cdot o$ and $h \cdot o$ lie in $N(0, o)$, then $g, h \in G$. Since $g, h \in G$, we have $ga \cdot o = G$ and $ha \cdot o = G$ if and only if a is one of the generators of G. Thus the automaton constructed in Theorem 6.4 is a single state $e|_0$ with loops labeled by the generators of G and no transitions labeled by the stable letter or its inverse. This shows what we already established above: that the entire language A^*, where A is a generating set for G, gives the 0-coarse regular stabilizer. Intersecting this with the language of geodesics in G gives $\mathcal{L}_2$, but we could also have used local geodesics as in Theorem 6.2 to obtain an alternative factor language (which would incidentally contain our original $\mathcal{L}_2$).

References

[1] Jason Behrstock, Mark F. Hagen, and Alessandro Sisto. Hierarchically hyperbolic spaces, I: Curve complexes for cubical groups. *Geom. Topol.*, 21(3):1731–1804, 2017.

[2] Brian H. Bowditch. Coarse median spaces and groups. *Pacific J. Math.*, 261(1):53–93, 2013.

[3] Brian H. Bowditch. Quasi-isometric maps between direct products of hyperbolic spaces. *Internat. J. Algebra Comput.*, 26(4):619–633, 2016.

[4] Martin Bridson and André Haefliger. *Metric Spaces of Non-Positive Curvature*, volume 319. 01 2009.

[5] John L. Britton. The word problem. *Ann. of Math. (2)*, 77:16–32, 1963.

[6] Marc Burger and Shahar Mozes. Lattices in product of trees. *Inst. Hautes Études Sci. Publ. Math.*, (92):151–194, 2000.

[7] Jérémie Chalopin, Victor Chepoi, Anthony Genevois, Hiroshi Hirai, and Damian Osajda. Helly groups (arxiv preprint), 2023.

[8] David B. A. Epstein, M. S. Paterson, J. W. Cannon, D. F. Holt, S. V. Levy, and W. P. Thurston. *Word Processing in Groups.* A. K. Peters, Ltd., USA, 1992.

[9] Sébastien Gouëzel and Vladimir Shchur. Corrigendum: A corrected quantitative version of the Morse lemma [MR3003738]. *J. Funct. Anal.*, 277(4):1258–1268, 2019.

[10] M. Gromov. Hyperbolic manifolds, groups and actions. In *Riemann surfaces and related topics: Proceedings of the 1978 Stony Brook Conference (State Univ. New York, Stony Brook, N.Y., 1978)*, volume No. 97 of *Ann. of Math. Stud.*, pages 183–213. Princeton Univ. Press, Princeton, NJ, 1981.

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

[12] Derek F. Holt. Automatic groups, subgroups and cosets. In *The Epstein birthday schrift*, volume 1 of *Geom. Topol. Monogr.*, pages 249–260. Geom. Topol. Publ., Coventry, 1998.

[13] Derek F. Holt, Sarah Rees, and Claas E. Röver. *Groups, Languages and Automata.* London Mathematical Society Student Texts. Cambridge University Press, 2017.

[14] Jingyin Huang and Damian Osajda. Helly meets Garside and Artin. *Invent. Math.*, 225(2):395–426, 2021.

[15] Sam Hughes and Motiejus Valiunas. A note on asynchronously automatic groups and notions of non-positive curvature (arxiv preprint), 2023.

[16] Sam Hughes and Motiejus Valiunas. Commensurating HNN-extensions: Hierarchical hyperbolicity and biautomaticity. *Comment. Math. Helv.*, 99(2):397–436, 2024.

[17] Michael Kapovich, Bruce Kleiner, and Bernhard Leeb. Quasi-isometries and the de Rham decomposition. *Topology*, 37(6):1193–1211, 1998.

[18] Bakhadyr Khoussainov and Anil Nerode. Automatic presentations of structures. In *Logic and computational complexity (Indianapolis, IN, 1994)*, volume 960 of *Lecture Notes in Comput. Sci.*, pages 367–392. Springer, Berlin, 1995.

[19] Ian J. Leary and Ashot Minasyan. Commensurating HNN extensions: nonpositive curvature and biautomaticity. *Geom. Topol.*, 25(4):1819–1860, 2021.

[20] Alex Margolis. Discretisable quasi-actions i: Topological completions and hyperbolicity (arxiv preprint), 2022.

[21] Jacek Świątkowski. Regular path systems and (bi)automatic groups. *Geom. Dedicata*, 118:23–48, 2006.