Geometric properties of the golden ratio Thompson's group

Denys Svetelik

Submitted in fulfilment of the requirements for the
Degree of Master of Science

School of Mathematics and Statistics
College of Science and Engineering
University of Glasgow

January 2025

Abstract

We show that all three golden ratio Thompson's groups F_τ, T_τ and V_τ embed in the asynchronous rational group. We prove properties of the Cayley graph of the monoid $M = \langle L, R : LR^2 = RL^2 \rangle$, whose topological full group is V_τ. Particularly we compute a distance function for the Cayley graph of the monoid M. Additionally, we prove that this Cayley graph is hyperbolic in the sense of Gromov. Our analysis reveals that the horofunction boundary of this graph is homeomorphic to a space resembling a Cantor-like set, with additional isolated points situated between each pair of breakpoints.

Contents

List of Figures

Declaration

With the exception of Chapters 1 and 2, which contain introductory material, all work in this thesis was carried out by the author unless otherwise explicitly stated.

Chapter 1

Introduction

1.1 Overview

In 1965 R. Thompson introduced three groups F, T, V, which were used in [27] for construction of finitely-presented groups with unsolvable word problems. In [34] Thompson showed that the groups T and V are finitely-presented, infinite simple groups and used the group V to prove that a group with finitely many generators has a solvable word problem if and only if it can be embedded into a finitely generated simple subgroup of a finitely presented group. The group F was used in a number of works related to homotopy idempotents [16, 19]. In [20] it was proven that F is the first known example of a torsion-free infinite-dimensional FP_∞ group. Later on it was proven in [29] that F is simply connected at infinity, hence showing that the group does not have any homotopy at infinity.

Thompson's group F can be defined in several equivalent ways:

- a group of piecewise linear homeomorphisms of the unit interval [0,1], where each homeomorphism has finitely many breakpoints at dyadic rational points and slopes that are powers of 2, formally $F = G([0,1]; \mathbb{Z}[\frac{1}{2}], \langle 2 \rangle)$;
- a group of tree pair diagrams, where elements are represented by pairs of finite binary trees with the same number of leaves;
- a finitely presented group with two generators and two relations.

The Thompson's groups F, T, and V can be seen as groups of piecewise linear homeomorphisms of the unit interval, the unit circle and the Cantor set respectively, that map dyadic rational numbers to dyadic rational numbers, and are differentiable everywhere except for a finite number of dyadic numbers, and the derivative of the interval of differentiability is always a power of 2. For a detailed introduction on Thompson's groups we refer the reader to [12] and [9].

In 2000 S. Cleary in his paper [13] introduced irrational slope Thompson's groups. These irrational slope Thompson's groups are a variation of the classic Thompson's group, with the breakpoints being now in $\mathbb{Z}[\tau]$ and slopes being power of τ, where $\tau = \frac{\sqrt{5}-1}{2}$ is the positive square root of the equation $x^2 + x = 1$

and is called the *small golden ratio*. In [8] and [10] J. Burillo, B. Nucinkis, and L. Reeves proved that the commutator subgroup of F_τ is simple but T_τ and V_τ have index-2 normal subgroups; they also gave a finite presentation of these groups and represented them in terms of binary trees.

The approach we adopt for *rationality* builds on the work of R. I. Grigorchuk, V. V. Nekrashevych, and V. I. Sushchanskii [21], who employed finite state machines to describe sets, relations, and functions. A central objective of this framework is to define homeomorphisms on the Cantor set $\{0,1\}^\omega$ using asynchronous machines, where a single symbol is read as input and a finite string, composed of elements of $\{0,1\}$, is produced as output during each step of computation. These homeomorphisms correspond to *rational functions*.

A homeomorphism of the Cantor set $\{0,1\}^\omega$ is called *rational* if it can be realized by an asynchronous transducer (or asynchronous Mealy machine) that acts on infinite binary strings. In [21], R. I. Grigorchuk, V. V. Nekrashevych, and V. I. Sushchanskii observed that the collection of all rational homeomorphisms of $\{0,1\}^\omega$ form a group $\mathcal{R}$ under the operation of composition, which they called the *rational group*. Additionally, they point out that the group of rational homeomorphisms of A^ω, for any finite alphabet A with at least two symbols, is isomorphic to $\mathcal{R}$.

Over the last few decades relatively low focus has been directed towards the class of groups generated by asynchronous transducers, particularly the full asynchronous rational group $\mathcal{R}$. It is established that the group $\mathcal{R}$ is simple but not finitely generated [6]. Furthermore, finitely generated subgroups of $\mathcal{R}$ have a solvable word problem, however there is no solution to the periodicity problem for elements of $\mathcal{R}$ [3, 21]. Additionally, $\mathcal{R}$ contains several well-known groups: the Thompson groups F, T, and V, the Brin-Thompson groups nV and groups like the Röver group V_Γ [3, 21, 33]. Furthermore, any group generated by synchronous automata can be embedded into $\mathcal{R}$.

In 2023 J. Belk, C. Bleak, F. Matucci and M. Zaremsky proved that all hyperbolic groups satisfy the Boone–Higman conjecture in [4]. The Boone–Higman conjecture was proposed in 1973 and states that a finitely generated group with a solvable word problem can be embedded in a finitely presented simple group. The authors proved that hyperbolic groups meet this criterion by demonstrating that each hyperbolic group embeds within a finitely presented simple group. In their proof they introduced a new class of groups, which they called *rational similarity groups*. Rational similarity groups generalize self-similar groups by allowing for rational, "canonical similarity" transformations within subshifts of finite type. Specifically, the authors showed that each hyperbolic group embeds in a full, contracting rational similarity group, which embeds in a finitely presented simple group. Additionally, rational similarity groups emerged as fundamental objects in the study of asynchronous group actions, suggesting their potential broader applicability in embedding problems within geometric and combinatorial group theory.

In [5], J. Belk, C. Bleak and F. Matucci proved that every Gromov hyperbolic group G embeds in the rational group $\mathcal{R}$, where by Gromov hyperbolic group we mean a group with a Cayley graph Γ that is hyperbolic in Gromov sense. This embedding uses a framework where elements of G act on a boundary space (horofunction boundary $\partial_h G$) of binary sequences by asynchronous transducers. In the same paper the authors proved that for any hyperbolic group G, the action of G on its horofunction

boundary $\partial_h G$ is rational. In the construction, the authors introduce a specific tree structure within the hyperbolic graph Γ, which they call the *tree of atoms*. They observed that when a group G acts properly and cocompactly on Γ then the tree of atoms has a self-similar structure and in addition is naturally homeomorphic to the horofunction boundary $\partial_h\Gamma$ (also known as the metric boundary). It was shown in [36] that the horofunction boundary $\partial_h\Gamma$ is compact, totally disconnected and has the Gromov boundary $\partial\Gamma$ as a quotient.

1.2 Summary of main results

In section 3 we prove the following theorem. Let $\{0,1\}^\omega$ be the set of all binary sequences. Let X_i, Y_i, C_{i+1} and Π_i, for $i = \{0,1\}$ be rational homeomorphisms of the Cantor set $\{0,1\}^\omega$ defined by the automata shown in Figure 3.1 (which differ only in their initial states).

Theorem 1.2.1. *The group G of homeomorphisms of $\{0,1\}^\omega$ generated by X_0, X_1, Y_0, Y_1, C_1, C_2, Π_0 and Π_1 is isomorphic to the golden ratio Thompson group V_τ. Moreover V_τ is a rational similarity group.*

The first part of the theorem is proven directly in Theorem 3.1.1. First, we identify a quotient map from the space of binary sequences $\{0,1\}^\omega$ to the unit interval $[0,1]$. We then establish that an isomorphism exists between the two groups. The second part is proven in Theorem 3.2.5. The proof heavily relies on the ideas presented by J. Belk, C. Bleak, F. Matucci and M. Zaremsky in [4]. They stated the conditions for a set of maps to be a nucleus of injections (Definition 3.2.3) and the conditions for a group to be a rational similarity group with a given nucleus (Theorem 3.2.4). We adapt these ideas to find the nucleus $\mathcal{N}$ for G, this lays a background in proving that V_τ is a rational similarity group.

From this, we derive a corollary: F_τ and T_τ are isomorphic to subgroups of G formed by the generators X_i, Y_i and X_i, Y_i, C_{i+1} respectively. This provides insight into the interaction of irrational slope Thompson groups within the context of automata and geometric representations.

Section 4 is heavily based on studying the properties of the Cayley graph of the monoid $M = \langle L, R : LR^2 = RL^2\rangle$, whose topological full group is V_τ. Let $\mathcal{M}$ be the respective Cayley graph, see Figure 4.1. Observe that the monoid M acts as the vertex set in $\mathcal{M}$, on which we impose with the path metric. Let x, m be vertices in $\mathcal{M}$, we say that $x \in Cone(m)$, if there is a word for x that starts with m. The main result of the section is the following theorem:

Theorem 1.2.2. *Let x, y be any given pair of vertices in the Cayley graph $\mathcal{M}$ of the monoid $M = \langle L, R : LR^2 = RL^2\rangle$. Then the distance between x and y is given by:*

$$d(x,y) = \begin{cases} d(x',y'), & \text{when } x = mx', y = my' \text{ for some nontrivial } m \in M; \\ |x| + |y| - 2, & \text{when } x \in Cone(LR), y \in Cone(RL) \text{ and } x, y \notin Cone(LR^2); \\ |x| + |y|, & \text{else.} \end{cases}$$

The proof is presented in Theorem 4.3.1. The first step we make is to observe that every cone in $\mathcal{M}$ is strongly geodesically convex, where by strongly geodesically convex we mean that if $x, y \in Cone(m)$ then

any geodesic between x and y is fully contained in $Cone(m)$. The remaining proof involves manipulating possible paths that a geodesic between x and y can take. The key idea here is to observe that if x starts with L and y starts with R then a geodesic either visits the root 1 or passes though the $Cone(LR^2)$, which is self-similar to the whole graph itself.

This theorem provides a useful distance function in a Cayley graph, which has a complicated structure. It is not only a profound tool in Sections 5 and 6 of this paper, but also it lays a solid background in further study of the monoid M.

Section 5 explores the horofunction boundary of the graph $\mathcal{M}$, introducing a Cantor-like set $\mathcal{D}_\tau$. This set includes additional isolated points between each pair of breakpoints, as shown below.

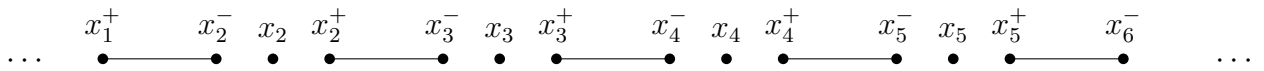

where every breakpoint x_i belongs to the ring $\mathbb{Z}[\tau] \cap (0, 1)$, and each interval above represents a similar Cantor set with extra isolated points.

Theorem 1.2.3. *The horofunction boundary $\partial_h\mathcal{M}$ of the Cayley graph $\mathcal{M}$ of the monoid $M = \langle L, R : LR^2 = RL^2\rangle$ is naturally homeomorphic to $\mathcal{D}_\tau$.*

The proof is presented in Theorem 5.3.2. Our proof heavily relies on the ideas presented by J. Belk, C. Bleak, and F. Matucci in [5]. The main idea of the proof is to construct a tree of atoms. This tree's boundary is shown to be homeomorphic to the horofunction boundary (Theorem 5.1.8). We start by considering vector fields, where each vector corresponds to a distance function between the two vertices, on the set of edges contained in balls of radius $n \geq 0$ around the root of $\mathcal{M}$. Since $\mathcal{M}$ is a locally finite graph, for every n, there is a finite number of possible vector fields. For each distinct vector field we associate a set of vertices from the whole graph, for which the vector field coincides with the distance function. Whenever the set is infinite, we call it an *atom*. With each increment of n, the previous atoms decompose into new ones. Eventually we observe a pattern in atoms and construct a tree of atoms. This tree has a self-similar structure and can be represented as a directed multi-graph. We show that the path space of this graph is $\mathcal{D}_\tau$. Then, following [5], we conclude that the horofunction boundary of the graph $\mathcal{M}$ is homeomorphic to $\mathcal{D}_\tau$.

This result establishes a solid understanding of the horofunction boundary of the Cayley graph of the monoid $M = \langle L, R : LR^2 = RL^2\rangle$ potentially leading to discovering new properties of the monoid M. This case raises an interesting question: could the horofunction boundary of all monoids of type $\langle L, R : LR^n = RL^n\rangle$ for $n > 1$, also be a Cantor-like space with additional isolated points between each pair of breakpoints? Exploring this could unveil a broader, underlying structure in the horofunction boundaries of monoids.

In Section 6 we embed the graph $\mathcal{M}$ into $\mathcal{M}'$, which is the same graph but with additional edges between every pair of adjacent vertices. We prove that these two graphs are quasi-isometric. Relying on the work of S. Kong, K. Lau and X. Wang [25] we prove certain proprieties (Theorem 6.1.4) of $\mathcal{M}'$ that

are equivalent of it being a hyperbolic graph, thus resulting in the following theorem, which is proved in Theorem 6.3.1.

Theorem 1.2.4. *The Cayley graph $\mathcal{M}$ of the monoid $M = \langle L, R : LR^2 = RL^2 \rangle$ is δ-hyperbolic.*

Chapter 2

Background

2.1 Word spaces

In this section, we recall the fundamental concepts of alphabets and the languages they generate. We rely on the definitions stated in [21, Section 2.1]. Our goal is to construct a background to produce infinite sequences from finite sets, which will be used throughout this paper.

Let X be a finite set with at least two elements, we refer to it as an *alphabet*. From X, we construct the set X^*, known as the *free monoid* generated by X. Elements of X^* are finite sequences of symbols from X, which are called *words*. The set also includes the *empty word* ϵ. If $v = a_1 a_2 \dots a_n \in X^*$, then $|v| = n$ represents the *length* of v. The length of the empty word ϵ is 0.

In addition to finite words, we also consider infinite sequences $a_1 a_2 a_3 \dots$, where $a_i \in X$. The set of all such infinite sequences is denoted by X^ω. For any $v \in X^*$ and $u \in X^* \cup X^\omega$, we can define the concatenation (product) $vu \in X^\omega$ in a natural way. A word $v \in X^*$ is a prefix of another word $u \in X^* \cup X^\omega$ if $u = vy$ for some $y \in X^* \cup X^\omega$. For any given set of words $X \subseteq X^* \cup X^\omega$, there exists a unique longest common prefix of all words in X, this prefix is infinite if and only if X consists of a single infinite word.

We endow the set X with the discrete topology, and X^ω with the product topology. This space is homeomorphic to the Cantor set, which implies that its topological type does not depend on the choice of X. For any finite word $v \in X^*$, the set $c(v) = \{vu : u \in X^\omega\}$ is both open and closed in this topology. The family $\{c(v) : v \in X^*\}$ forms a basis for the topology on X^ω. The sets $c(v_1)$ and $c(v_2)$ have a nonempty intersection if and only if one of v_1 or v_2 is a prefix of the other. In such case, the set corresponding to the longer word is a subset of the other.

2.2 Automaton and rational groups

In this section we recall the concepts of automata theory formulated in [21, Section 2]. We state the definitions for automata, rational groups and how they are connected from [5, Section 1.1]. Our goal is to build up the language that connects graph theory with a special variation of Thompson's groups.

Definition 2.2.1. [5, Definition 1.1] An *automaton* has the following components:

1. Two finite sets X_{in} and X_{out}, which are called the input and output alphabet.
2. A finite set of states Q, where each element represents a distinct state.
3. An initial state q_0 that belongs to the set Q.
4. A transition function t that maps a state $q \in Q$ and an input symbol $a \in X_{in}$ to a new state $q' \in Q$, represented as $t : Q \times X_{in} \to Q$.
5. An output function o that maps a state $q \in Q$ and an input symbol $a \in X_{in}$ to a string of output symbols $s \in X_{out}$, represented as $o : Q \times X_{in} \to X_{out}^*$.

The automaton is classified as *synchronous* if the output mapping o generates a single symbol from the output alphabet X_{out} for each state $q \in Q$ and input symbol $a \in X_{in}$. If this condition is not met, the automaton is called *asynchronous*.

An automaton can be graphically represented as a finite directed graph. Each state is depicted as a node in the graph. The transitions and outputs are illustrated by directed edges. Specifically, for each state $q \in Q$ and input symbol $a \in X_{in}$, there is a directed edge from the node representing q to the node representing $t(q, a)$, labeled with $a/o(q, a)$.

If the reader is unfamiliar with directed graphs, a formal definition appears in Section 5.

Given an automaton $T = (X_{in}, X_{out}, Q, q_0, t, o)$, an *input word* for T is an infinite sequence $a_1a_2a_3... \in X_{in}^{\omega}$. The corresponding *output word* is the concatenation of the outputs produced by the output mapping o for each state transition, expressed as:

$$o(q_0, a_1)o(q_1, a_2)o(q_2, a_3)...,$$

where $\{q_n\}$ is the sequence of states starting from the initial state q_0, defined recursively as $q_n = t(q_{n-1}, a_n)$.

It's worth noting that the output word may be finite if the output function $o(q, a)$ is the empty word for all but finitely many input symbols a in a given state q. However, our focus is on automata whose output words are always infinite. Such automata are called *nondegenerate*, otherwise we call them *degenerate*. A nondegenerate automaton defines a mapping from infinite input words to infinite output words over the respective alphabets.

Example 2.2.2. Observe the automaton $\mathcal{X}$ shown in Figure 2.1, where f is the initial state. When an infinite input word consisting of 1s (represented as $111\ldots$) is processed through the automaton starting from state f, the output is 01. In other words, $f(111\ldots) = 01$. Note that if an automaton contains a cycle that produces empty output words, it is degenerate.

Definition 2.2.3. [5, Definition 1.2]A function $f : X_{\text{in}}^{\omega} \to X_{\text{out}}^{\omega}$ is called *rational* if there exists a non-degenerate automaton $\mathcal{X}$ over the alphabets X_{in} and X_{out} such that $f(\psi) = \mathcal{X}(\psi)$ for all $\psi \in X_{\text{in}}^{\omega}$.

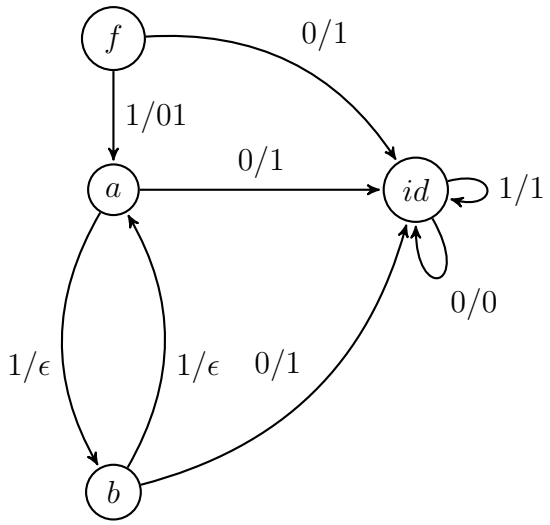

Figure 2.1: An automaton that is degenerate.

Rational functions possess several key properties [5, Proposition 1.3]:

1. Rational functions are continuous with respect to the product topologies on X^ω_{in} and X^ω_{out}.
2. The composition of two rational functions is again a rational function.
3. If $f : X^\omega_{\text{in}} \to X^\omega_{\text{out}}$ is a rational bijection, then the inverse function $f^{-1} : X^\omega_{\text{out}} \to X^\omega_{\text{in}}$ is also rational.

Definition 2.2.4. [5, Definition 1.4] Let X be a finite alphabet, consisting of at least two symbols, we define the *rational group* $\mathcal{R}_X$. This group consists of all rational homeomorphisms of X^ω. All rational groups $\mathcal{R}_X$ are isomorphic to each other, we refer to a single rational group $\mathcal{R}$ without specifying the alphabet.

When G is a group, any injective homomorphism $G \to \mathcal{R}_X$ is a rational representation of this group. The rational group $\mathcal{R}$ can be studied through an infinite rooted tree X^* of finite strings over X. The boundary ∂X^* of this tree is homeomorphic to X^ω, and the action of $\mathcal{R}$ on X^ω can be understood by analyzing the restrictions of rational functions to the subtrees of X^*. See section 5.1 for the definitions of trees and boundaries.

It has been shown in [5] that every finitely generated group, whose Cayley graph is delta-hyperbolic, embeds in the rational group $\mathcal{R}$. This result creates a bridge between geometric group theory and the theory of automata. This opens questions about what properties of hyperbolic groups are preserved by this embedding.

Definition 2.2.5. [21, Definition 2.3] Let $f : X^\omega_{in} \to X^\omega_{out}$ be a continuous non-constant function and $w \in X^*_{in}$ be a finite word. The *local action* of f at the word w is denoted by $f|_w$, is the mapping $f|_w : X^\omega_{in} \to X^\omega_{out}$ defined as:

$$\forall u \in X^\omega_{in},\ f(wu) = vf|_w(u),$$

where v is the longest common prefix of the sets of words $\{f(wu) : u \in X_{in}^{\omega}\}$. If this set has only a single word, then the largest common prefix is infinite, in which case $f|_w$ is undefined.

In other words, the local action $f|_w$ describes how f acts on infinite words that begin with the finite prefix w. If the output of f on words with prefix w is independent of the infinite suffix after w, then the local action at w is undefined. It is not allowed for f to be a constant function, because a constant function may map two distinct input words to the same output word, which is not allowed.

Theorem 2.2.6. *[21, Theorem 2.5] A continuous mapping $f : X_{in}^{\omega} \to X_{out}^{\omega}$ is rational if and only if it has a finite number of local actions.*

Definition 2.2.7. [5, Definition 2.1] Let Γ be a finite directed graph, then the *subshift of finite type* is assosiated to Γ is the set Σ_Γ of all infinite directed paths in Γ.

Let α be a finite path in a directed graph Γ, then a *cone* is the set $\mathcal{C}_\alpha \subseteq \Sigma_\Gamma$ containing all infinite paths that start with the prefix α. When v is a vertex in Γ, then cone $\mathcal{C}_v$ is the of all infinite directed paths that start from v. We denote $\mathcal{C}_\emptyset$ as the set of all infinite paths in Γ. Every cone $\mathcal{C}$ is a clopen subset of Σ_Γ with the product topology, in addition they for a basis for a topology. [5, Section 2.1]

Definition 2.2.8. [4, Definition 2.19] A subshift Σ_Γ has an *irreducible core* if there exists an induced subgraph Γ_0 of Γ such that the following conditions hold:

1. the graph Γ_0 is irreducible;
2. for every vertex $v \in \Gamma_0$, there is a directed path in Γ from v to a vertex in Γ_0;
3. there exists $n \geq 0$ such that every directed path in Γ of length n terminates in Γ_0.

Proposition 2.2.9. *[5, Proposition 2.12] Let Σ_Γ be a subshift of finite type with no isolated points or empty cones, and $E \subseteq \Sigma_\Gamma$ be a nonempty clopen set. Then the set $\mathcal{R}_{\Gamma,E}$ of rational homeomorphisms $E \to E$ forms a group under composition.*

Definition 2.2.10. [5, Section 2.1] A *canonical similarity*

$$L_\alpha : \mathcal{C}_{t(\alpha)} \to \mathcal{C}_\alpha$$

is a homeomorphism defined by the formula $L_\alpha(\omega) = \alpha \cdot \omega$, where α is a finite path. Generally when α and β are finite paths with $t(\alpha) = t(\beta)$, the composition $L_\beta \circ L_\alpha^{-1}$ defines the canonical similarity $\mathcal{C}_\alpha \to \mathcal{C}_\beta$, which maps $\alpha \cdot \omega$ to $\beta \cdot \omega$, for all $\omega \in \mathcal{C}_\alpha$.

Definition 2.2.11. [5, Definition 2.32] Let Σ_Γ be a subshift of finite type with no isolated points or empty cones. Let $E \subseteq \Sigma_\Gamma$ be a nonempty clopen set, and let $\mathcal{R}_{\Gamma,E}$ be the associated rational group. A subgroup $G \leq \mathcal{R}_{\Gamma,E}$ is called a *rational similarity group* if, for every pair of cones $\mathcal{C}_\alpha, \mathcal{C}_\beta$ neither of which are contained within E with $t(\alpha) = t(\beta)$, there exists an element $g \in G$ that maps $\mathcal{C}_\alpha$ to $\mathcal{C}_\beta$ via the canonical similarity.

2.3 Irrational slope Thompson groups

In this section, we introduce the Thompson group F_τ, which is an interesting variant of the Thompson group F that was introduced by S. Clearly in 2005 in [13]. The group F_τ shares many structural properties with the original Thompson group F, but with a key difference: the slopes and breakpoints of the group elements are tied to the golden ratio $\frac{\sqrt{5}-1}{2}$, also known as the *small golden ratio*. This group has been extensively studied by J. Burillo, B. Nucinkis, and L. Reeves, who provided a finite presentation for F_τ and explored its combinatorial structure in terms of binary trees in [8]. In [30] L. Molyneux, B. Nucinkis, and Y. Rego calculated the BNSR-invariants for F_τ.

Let $\tau = \frac{\sqrt{5}-1}{2} \approx 0.6180339887...$ be the *small golden ratio*, which is a root of the polynomial x^2+x-1. This value of τ is a unit in the ring $\mathbb{Z}[\tau]$, which consists of elements of the form $a + b\tau$, where a and b are integers.

One crucial property of τ is that $\tau + \tau^2 = 1$. This allows us to subdivide the unit interval into two subintervals of length τ and τ^2 in two different ways: $[0, 1] = [0, \tau^2] \cup [\tau^2, 1]$ and $[0, 1] = [0, \tau] \cup [\tau, 1]$ respectively. Observe that since $\tau^n = \tau^{n+1} + \tau^{n+2}$, every subdivision of the unit interval can be subdivided into two subintervals of lengths of some powers of τ. Therefore every end-point of a sequence of subdivisions will belong to $\mathbb{Z}[\tau]$, see [13] for a detailed proof. Taking this into account, we define the following group:

Definition 2.3.1. F_τ is the group of homeomorphism from the unit interval to itself such that:

- they are piecewise linear and orientation-preserving;
- is non-differentiable at a finite number of points that belong to $\mathbb{Z}[\tau] \cap [0, 1]$;
- where differentiable, the slope is a power of τ.

Formally:

$$F_\tau = G([0, 1]; \mathbb{Z}[\tau], \langle\tau\rangle)$$

Here we adopt the standard notation for representing the Thompson's group F, where $[0, 1]$ represents the domain and range of the functions, $\mathbb{Z}[\tau]$ represents the set of non-differentiable points, and $\langle\tau\rangle$ represents the set of possible slops for the differentiable segments.

Like Thompson's group F, elements of F_τ can be represented by pairs of binary trees. Each binary tree encodes a subdivision of the interval into subintervals, and the leaves of the trees correspond to these subintervals. However, unlike in F, we must distinguish between two types of subdivisions in F_τ one of length τ and one of length τ^2. This is represented in the tree diagram by carets with edges of different lengths, where the longer edge corresponds to the shorter interval and vice versa, to encode the length as a power of τ. Following [8], we define two types of carets in F_τ:

Definition 2.3.2. [8, Definition 1.1] For the group F_τ, we define two types of carets:

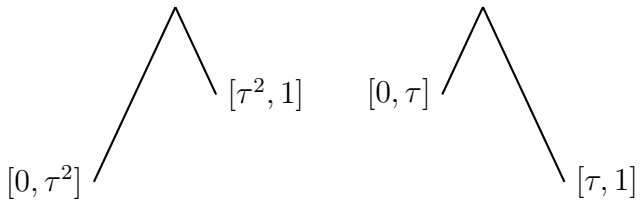

Figure 2.2: A pair of tree diagrams representing the domain and range of a function in F_τ, where the domain caret is of x-type and range caret is of y-type

- an *x-type caret* has a long left edge and a short right edge. It subdivides an interval $[x, y]$ into $[x, x+\tau^2(y-x)] \cup [x+\tau^2(y-x), y]$.
- a *y-type caret* has a short left edge and a long right edge. It subdivides an interval $[x, y]$ into $[x, x+\tau(y-x)] \cup [x+\tau(y-x), y]$.

See Figure 2.3 for a visual representation of these caret types.

A pair of tree diagrams has a non-trivial application. Whenever both trees have the same number of leaves they can be used to represent the domain and range. Let's view a short example.

Example 2.3.3. Take the function $f : [0, 1] \to [0, 1]$ defined by:

$$f(x) = \begin{cases} x\tau \text{ when } 0 \leq x < \tau^2; \\ x\tau^{-1} - \tau \text{ when } \tau^2 \leq x \leq 1. \end{cases}$$

Let's verify that f belongs to F_τ.

Proof. The function f is piecewise-linear and since $0 = f(0)$ and $1 = f(1)$, it is orientation preserving. The derivative of f is:

$$f'(x) = \begin{cases} \tau \text{ when } 0 \leq x < \tau^2; \\ \tau^{-1} \text{ when } \tau^2 < x \leq 1. \end{cases}$$

Hence, the slopes of f are powers of τ and f is not differentiable only at a finite number of points that belong to $\mathbb{Z}[\tau]$. Hence, f belongs to F_τ. □

We can express a function that belongs to F_τ as a pair of tree diagrams, where the left tree represents the domain and the right tree represents the range. See Figure 2.2 for the representation of f. The general rule for reading a tree diagram is that the n^{th} interval in the domain tree, counting from the left, is mapped to the n^{th} interval in the range tree, unless specified otherwise. The length of intervals in a tree diagram corresponds to τ^n, where n is the distance from the root to the end of the caret of the respective interval. Short carets have a length of 1, while long carets have a length of 2.

The distinction between the two types of carets reflects the relationship between the depth of a node in the tree and the length of the associated interval. The tree depth corresponds to powers of τ, with

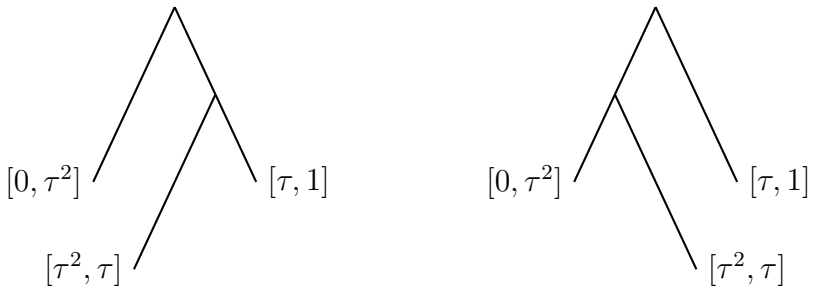

Figure 2.3: Interchangeable tree patterns

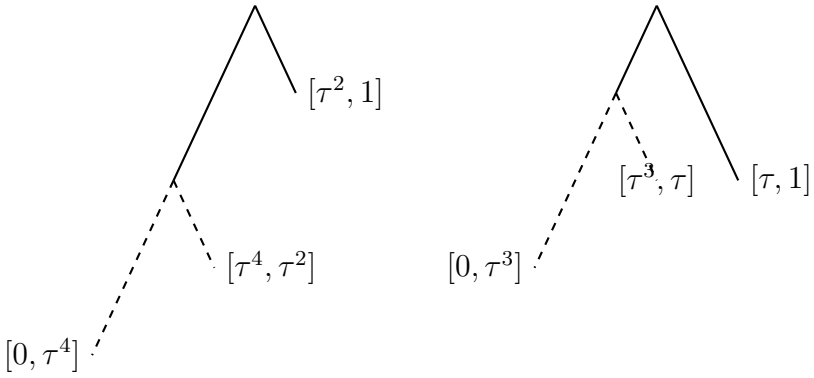

Figure 2.4: Adding carets to the tree diagram without changing the element of F_τ

shorter intervals being represented by carets with longer edges, as this encodes the exponential scaling of interval lengths.

One interesting property of F_τ is that there exist subdivisons that correspond to multiple tree diagrams. Observe, a subdivision of $[0, 1]$ into three subintervals of lengths τ^2, τ^3, and τ^2 can be represented by two distinct trees, as shown in Figure 2.3. This process of interchanging subtrees, while preserving the overall structure, is called a *basic move*. Note that the first tree is constructed by 2 x-type carets, whereas the second is made of 2 y-types. Such patterns are always interchangeable even if they appear as subtrees, while preserving the overall structure of the interval subdivision.

This property leads to interesting algebraic consequences, such as the fact that two different reduced diagrams may represent the same element in F_τ, unlike in the classical Thompson group F, where such diagrams can always be reduced further. This flexibility plays an important role in defining operations within the group, such as composition, which is performed by combining tree pair diagrams after expanding them to a common subdivision.

Since there are an infinite number of binary tree representations for the same element of F_τ, one can further subdivide the n-th subinterval of the domain and range tree by adding carets of the same type to corresponding leaves in both trees. See Figure 2.4 for an example of how a caret can be added to the tree diagrams without altering the element of F_τ. Therefore carets of the same pair can be canceled without the risk of changing the element. This leads to the conclusion of [8, Theorem 7.3], which states that every element of F_τ has a unique representation without cancellative carets.

Theorem 2.3.4. *[8, Theorem 4.4] The group F_τ is generated by four key elements: x_0, x_1 and y_0, y_1*

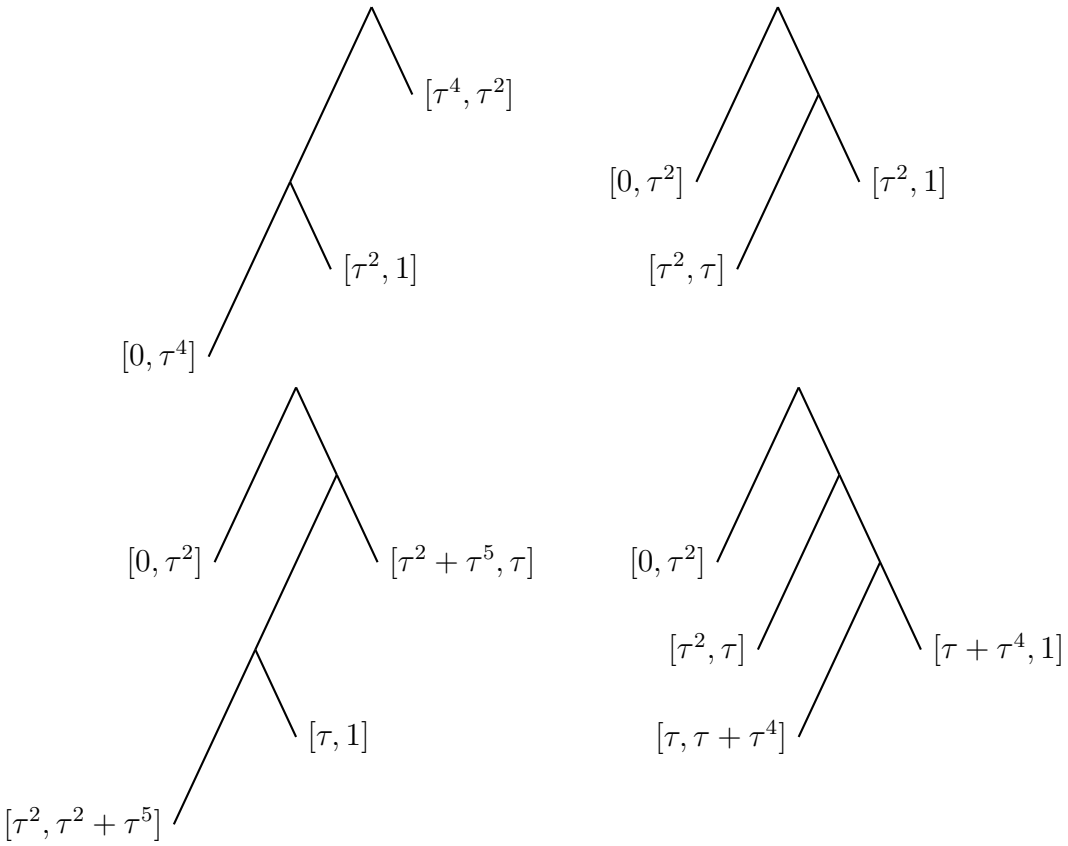

Figure 2.5: The generators x_0 and x_1 as tree diagrams.

(see Figures 2.5, 2.6 for the generators).

These generators are analogous to the generators x_0 and x_1 of Thompson's group F, but the presence of both x and y type carets introduces additional complexity. These generators satisfy relations similar to those in F:

$$x_j x_i = x_i x_{j+1}, \quad x_j y_i = y_i x_{j+1}, \quad y_j x_i = x_i y_{j+1}, \quad y_j y_i = y_i y_{j+1}, \quad y_i^2 = x_i x_{i+1}.$$

This finite presentation provides a framework for constructing any element of F_τ as a product of these generators. Figure 2.7 represents the infinite generators of the group F_τ.

The group T_τ is the group of piecewise-linear, orientation-preserving homeomorphisms of the circle, such that it is not differentiable only at a finite number of points in $\mathbb{Z}[\tau]$, and having derivatives of the differentiable intervals as powers of τ.

Thompson's group T_τ can be constructed by considering maps on the unit circle. We identify the two endpoints of the unit interval $[0, 1]$, hence we can consider the maps of the interval in such a way that the images of 0 and 1 are equal. For a more detailed introduction to T_τ, we refer the reader to [10].

Theorem 2.3.5. *[10, Theorem 2.2] The Thompson group T_τ is generated by the generators of the Thompson group F_τ and, in addition, c_1 and c_2. See Figure 2.8 for these generators.*

The rules for T_τ are identical to those for F_τ with the only difference being that we allow permutation

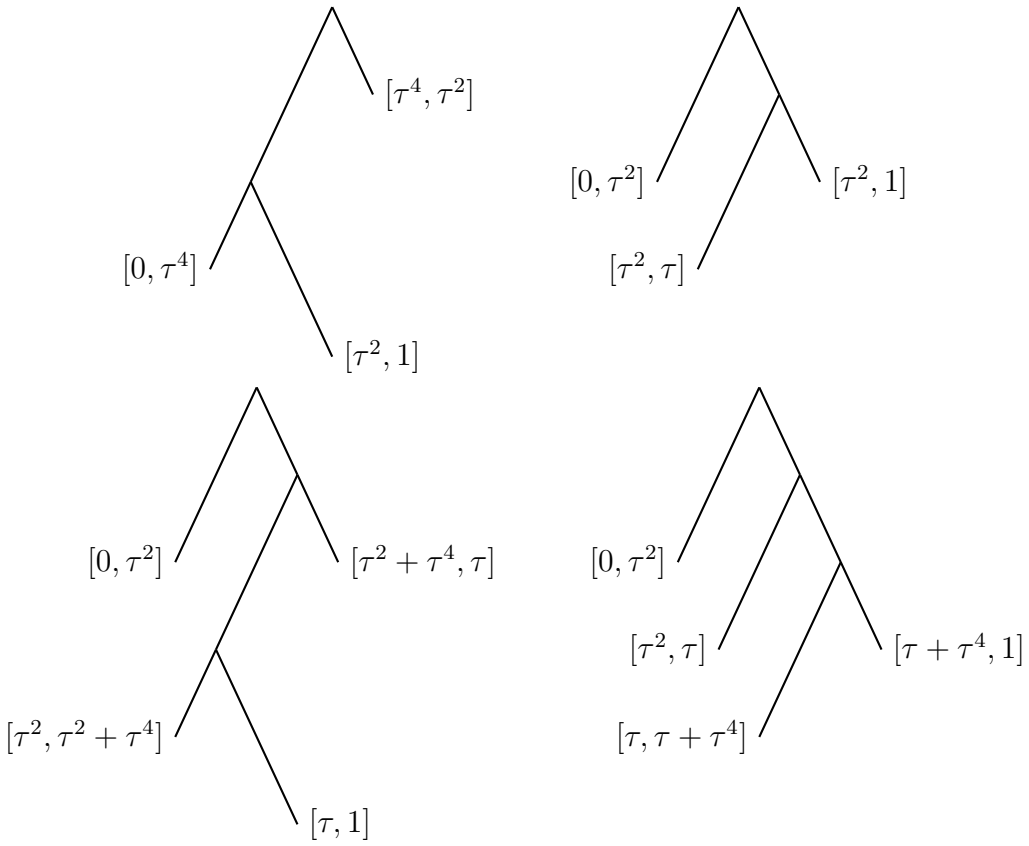

Figure 2.6: The generators y_0 and y_1 as tree diagrams.

of the order of the range intervals. For example, the range intervals of the c_n generators are shifted once anticlockwise.

The group V_τ is the group of piecewise-linear, orientation-preserving bijections of the interval, such that it is not differantiable only at a finite number of points in $\mathbb{Z}[\tau]$, and having derivatives of the differential intervals as powers of τ. We consider two elements of V_τ to be the same if they agree everywhere but on a finite set of points. For a more detailed introduction to V_τ, we refer the reader to [10].

Theorem 2.3.6. *[10, Section 4] The Thompson group V_τ is generated by the generators of the Thompson group T_τ and, in addition, π_0 and π_1. See Figure 2.8 for these generators.*

The rules for V_τ are identical to those for F_τ. The only difference is that V_τ has additional generator that permute the penultimate and ultimate intervals in the range tree.

We will now present a construction of the Cantor set on which V_τ acts. For a detailed introduction on the Cantor set we refer the reader to [37, Section 30].

Definition 2.3.7. Let E be a countable dense subset of $(0, 1)$. *The blowup* of $[0, 1]$ along E is the set

$$([0, 1] \setminus E) \cup \{x^- : x \in E\} \cup \{x^+ : x \in E\}.$$

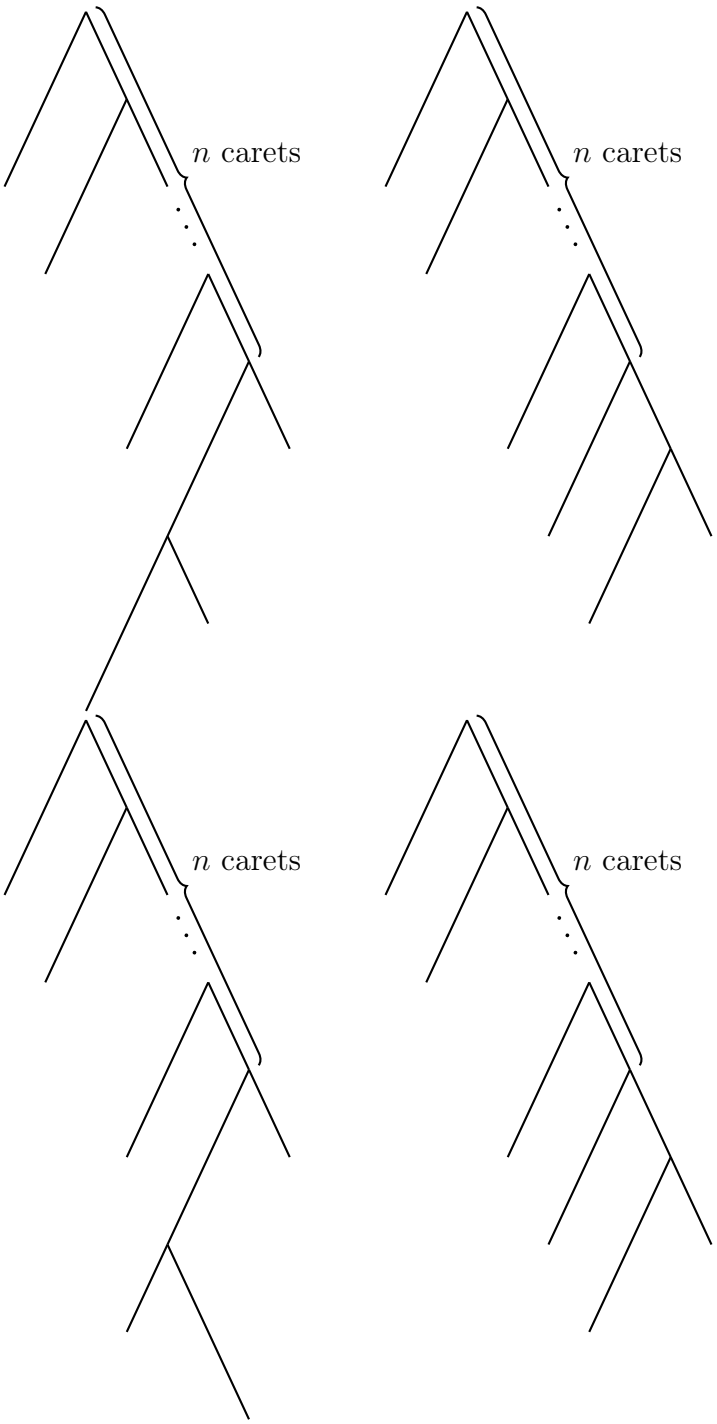

Figure 2.7: The generators x_n and y_n as tree diagrams.

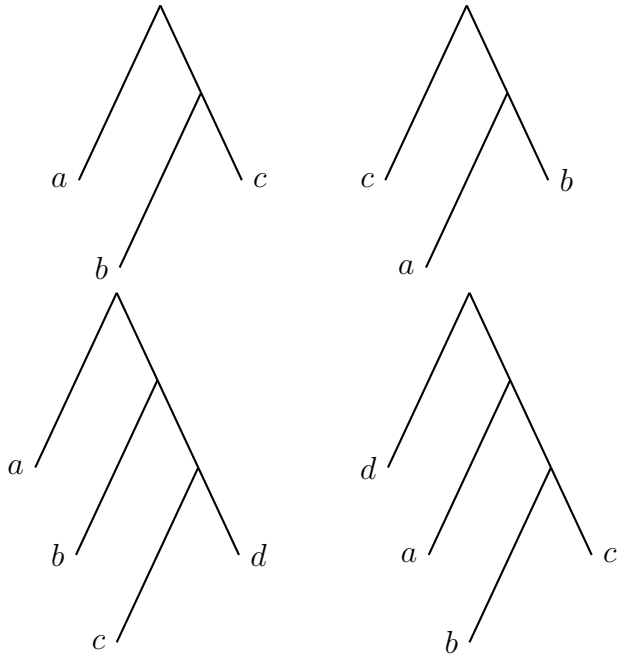

Figure 2.8: The generators c_1 and c_2 as tree diagrams. The labels on the leaves determine which domain intervals map to which range intervals.

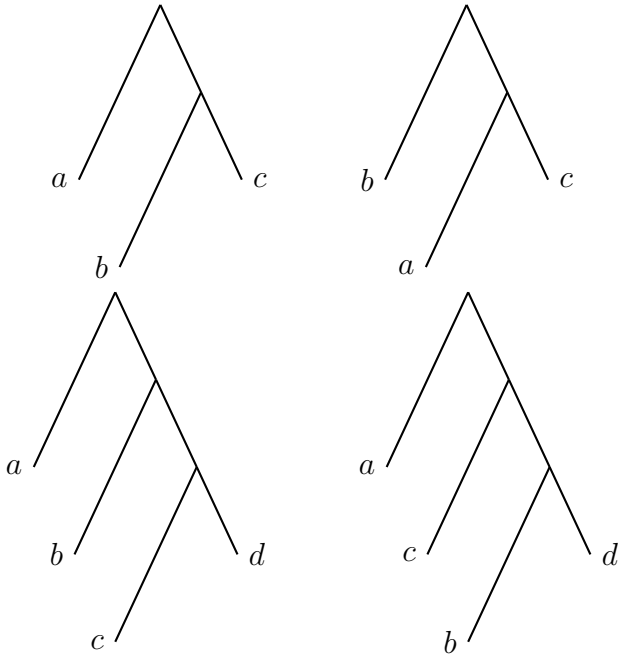

Figure 2.9: The generators π_0 and π_1 as tree diagrams.

We endow this blowup with the order topology, where $a < x^- < x^+ < b$ for any $x \in E$ and $a, b \in [0,1] \setminus E$ such that $a < x < b$.

Theorem 2.3.8. *[24, Cantor's isomorphism theorem. Theorem 4.3] Any two countable, dense, unbounded linear orders are order-isomorphic*

We can now prove that a blowup along $[0,1]$ is an efficient tool in construction of a Cantor set.

Theorem 2.3.9. *Any blowup of* $[0,1]$ *along a countable dense subset* E *of* $(0,1)$ *is homeomorphic to the Cantor set.*

Proof. By Cantor's theorem we have an order isomorphism $f\colon E \to D$, where $D = \mathbb{Z}\left[\frac{1}{2}\right] \cap (0,1)$. This extends continuously to a homeomorphism $f\colon [0,1] \to [0,1]$ that maps E to D, and it follows that the blowup of $[0,1]$ along E is homeomorphic to the blowup of $[0,1]$ along D. But the latter is obviously homeomorphic to the usual Cantor set. □

Definition 2.3.10. Let $I_\tau = \mathbb{Z}[\tau] \cap (0,1)$, where $\mathbb{Z}[\tau] = \{a + b\tau : a, b \in \mathbb{Z}\}$. Then $\mathcal{C}_\tau$ is a Cantor set defined by a blowup along I_τ.

Let $x_i \in I_\tau$ then the Cantor set $\mathcal{C}_\tau$ is depicted as:

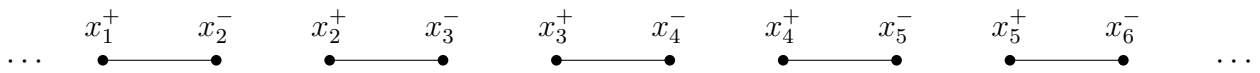

Where each interval above represents a similar Cantor set. The Thompson's group V_τ acts on this Cantor set. Specifically, if $f \in V_\tau$ maps an interval (a,b) to (c,d) for some $a, b, c, d \in \mathbb{Z}[\tau]$, then f maps $[a^+, b^-]$ homeomorphically to $[c^+, d^-]$. It can be seen as the group of all possible permutations of non-dividable intervals contained in $\mathcal{C}_\tau$,

Chapter 3

Irrational slope Thompson groups as an asynchronous automaton

In this section, we explore the Thompson groups F_τ, T_τ, and V_τ and their representation as asynchronous automata. This connection between group theory and automaton theory provides valuable insights into the structure and properties of V_τ. We begin by recalling the generators of V_τ and then proceed to construct an automaton that captures its behavior.

3.1 Thompson groups F_τ, T_τ and V_τ as group homeomorphisms of the Cantor set

In this section we use the automaton theory to explore the dynamical properties of the Thompson groups F_τ, T_τ and V_τ. Particularly we prove that the Thompson groups F_τ, T_τ and V_τ are isomorphic to a certain groups of rational homeomorphisms of the Cantor set. Let $X_0, X_1, Y_0, Y_1, C_1, C_2, \Pi_0$ and Π_1 be rational homeomorphisms of the Cantor set $\{0,1\}^\omega$ defined by the automata shown in Figure 3.1 (which differ only in their initial states). We prove the following theorem.

Theorem 3.1.1. *The group G of rational homeomorphisms of $\{0,1\}^\omega$ generated by X_0, X_1, Y_0, Y_1, C_1, C_2, Π_0 and Π_1 is isomorphic to V_τ. Formally, $G = \langle X_0, X_1, Y_0, Y_1, C_1, C_2, \Pi_0, \Pi_1\rangle \simeq V_\tau$. See Figure 3.1 for the automata that generate the group G.*

Note that G is a subgroup of $\mathrm{Homeo}(\{0,1\}^\omega)$. The rest of the section is dedicated to the proving this theorem.

Let's recall the generators of the Thompson group V_τ. It has been shown in [10, Section 4] that V_τ has an infinite generating set $\{x_i, y_i, c_{i+1}, \pi_i\}$, where $i \in \mathbb{Z}$, and a finite generating set $\{x_0, x_1, y_0, y_1, c_1, c_2, \pi_0, \pi_1\}$, see Figures 2.5, 2.6, 2.8 and 2.9 for the generating elements. Our goal is to demonstrate that there exists an asynchronous automaton with initial states corresponding to the generating set $\{x_0, x_1, y_0, y_1, c_1, c_2, \pi_0, \pi_1\}$.

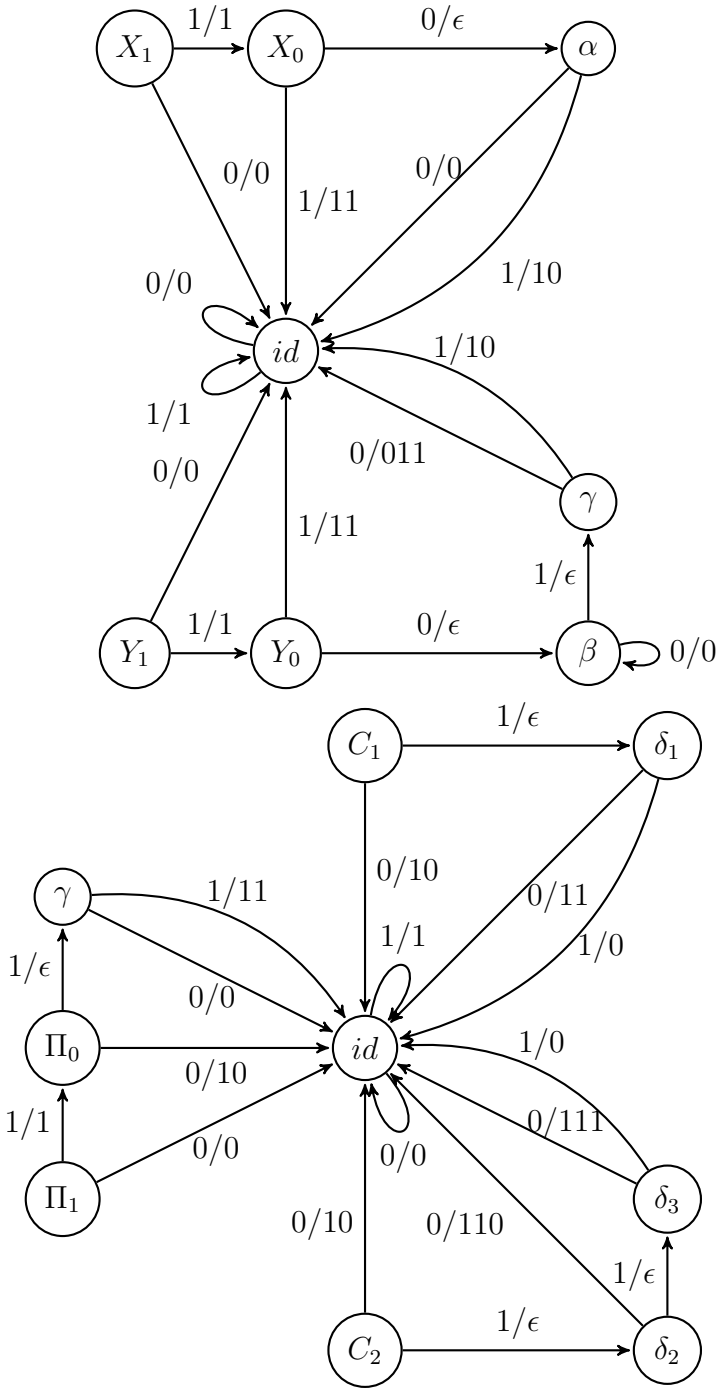

Figure 3.1: Automata that generate the Thompson group V_τ.

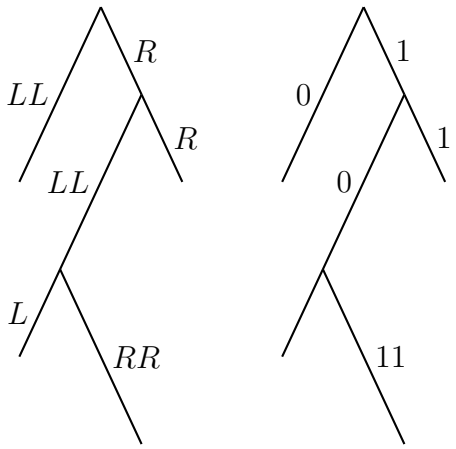

Figure 3.2: Carets of the y_1 generator addressed by pairs L's,R's and 0's, 1's

To accurately interpret the domain and ranges in trees using binary sequences, we need to establish a standardized notation for caret addresses. Let a left caret be denoted as L and a right caret as R to provide clarity in representation. In this context, two letters of the same kind represent a long caret, while one letter symbolizes a short caret. Consistently, we assign values 0 and 1 to LL and R respectively, ensuring a precise and coherent portrayal of caret addresses in binary sequences.

Example 3.1.2. Figure 3.2 demonstrates the intuition behind the notation we use to address carets in a tree representation. Note that a small left caret denoted by L is not defined in the binary $\{0,1\}$ representation. However this is not a problem as we are allowed to add additional carets and perform basic moves.

Remark 3.1.3. Recall that the tree diagrams for the x_n and y_n generators are read by following a sequence of carets that start from the root and, at each step, slide down the tree until reaching a caret that does not have any children. Note that a finite tree has at most m such sequences, where m is the total number of carets in the tree.

Our proof that the generators of G are rational homeomorphisms will use the golden ratio base for numbers, where a base τ expansion of $x \in \mathbb{R}$ is $0.\alpha_1\alpha_2\cdots$, where $\alpha_i \in \{0,1\}$ and $\sum_{i=1}^{\infty} \alpha_i \tau^i = x$. See [7] for a general introduction to irrational bases. Every $x \in \mathbb{R}$ has many different base τ expansions. For example, since $\tau = \tau^2 + \tau^3$, it follows that τ has expansions $0.1000\cdots$ and $0.011000\cdots$, additionally $\tau = \tau^2 + \tau^4 + \tau^5$, so τ has also expansion $0.01011000\cdots$. However, we can make the base τ expansion unique for most numbers in $[0,1]$ if we impose the rules that $\alpha_1 = 0$ and all finite strings of 0's in $\alpha_2, \alpha_3, \ldots$ have even length. If we use these rules, then any number is $\mathbb{Z}[\tau] \cap (0,1)$ has 2 expansions and any other number in $[0,1]$ has a unique expansion. By applying these rules we obtain that two base τ expansions represent the same x if they have forms $0.\beta00111\cdots$ and $0.\beta1000\cdots$ for some finite $\beta \in \{0,1\}^*$. For example, imposing these rules limits the possible expansions of τ to $0.1000\cdots$ and $0.00111\cdots$.

The following propositions follows immediately from the golden-ratio base.

Proposition 3.1.4. *Let $q : \{0,1\}^\omega \to [0,1]$ be defined as $q = q_2 q_1(\omega)$, where:*

$$q_1(\omega_1) = \begin{cases} 00q_1(\omega_2), & \text{when } \omega_1 = 0\omega_2 \\ 1q_1(\omega_2), & \text{when } \omega_1 = 1\omega_2 \end{cases};$$
$$q_2(\omega) = 1 - \sum_{n=1}^{\infty} \tau^{n+1}(1-\omega_n) = \sum_{n=1}^{\infty} \tau^{n+1}\omega_n.$$

Note that q is an order preserving surjection, where each point in $\mathbb{Z}[\tau] \cap (0,1)$ has 2 preimages and every other point in $[0,1]$ has 1 preimage; that is, q is an almost one-to-one surjection. Moreover, q commutes with the rational map generating V_τ. We check this for the X_0 and x_0 generators in Proposition 3.1.6

The function q_1 performs a simple operation. It replaces every occurrence of 0 with 00. Whereas q_2 mimics the behavior of the left L or right R operations. Observe that L only changes the upper bound of the interval by a power of τ, whereas R only changes the lower bound. Take a finite sequence $RRR...$ At each sequential R the interval is shrunk from the left by τ^{n+1}. This representation forces the lower bound to converge with the upper bound when the sequence is infinite. As a result, there are exactly two formulas that can convert the infinite sequence to a point in the interval. Either by summing up powers of τ at the positions of 1's or by subtracting from 1 the powers of τ at the 0's positions.

To illustrate how the formula works let's consider an example. Let $\omega_{ex} = 010100\ldots$

$$q(\omega_{ex}) = q(010100\ldots) = q_2 q_1(010100\ldots) = q_2(0010010000\ldots)$$
$$= \sum_{n=1}^{\infty} \tau^{n+1}(\omega_n) = \tau^4 + \tau^7.$$

The following corollary follows from Proposition 3.1.4.

Corollary 3.1.5. *There exists a homeomorphism h such that the following diagram commutes.*

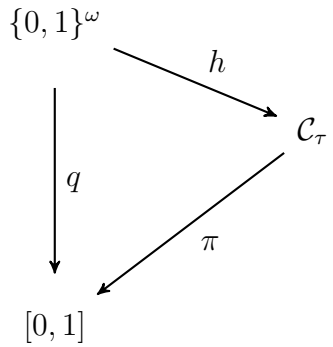

Since q is order-preserving, the fiber over each x in $[0,1]$ consist of 2 points if the x is in $\mathbb{Z}[\tau]$ and one point otherwise.

Proposition 3.1.6. *The map q makes the following diagram commute, where X_0 is a generator of G*

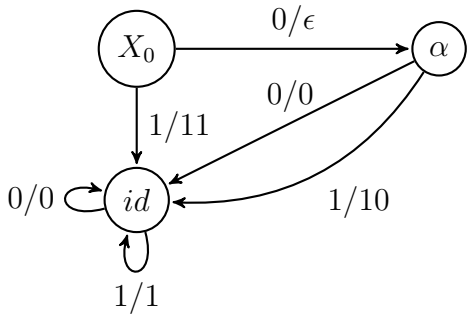

Figure 3.3: The X_0 generator

and x_0 is a generator of V_τ.

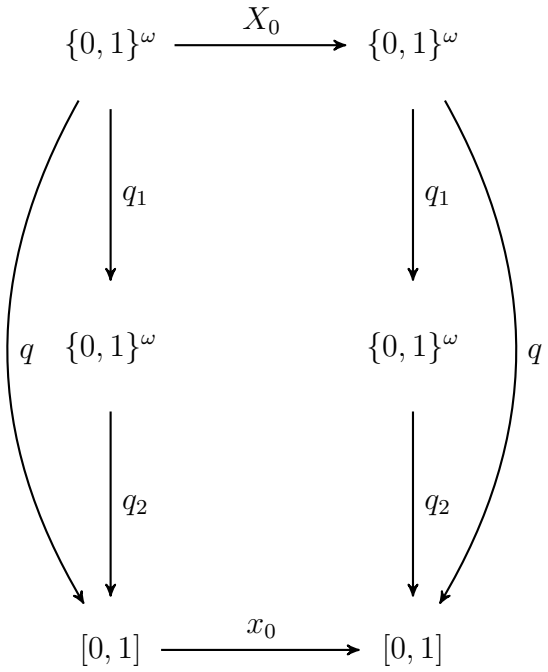

Proof. We will verify that $x_0 q = q X_0$. Let's recall the x_0 generator and the X_0 generator, see Figures 2.5 and 3.3. From reading the diagram that describes the generator x_0, we introduce two linear functions f_L, $f_R : [0,1] \to [0,1]$, defined by:

$$f_L(t) = \tau t;$$
$$f_R(t) = (1-\tau) + \tau t.$$

For simplicity we denote $f_x \circ f_y = f_{xy}$. Due to being linear these functions satisfy the following equations:

$$x_0 f_{L^4}(\omega_1) = f_{L^2}(\omega_1);$$
$$x_0 f_{L^2R}(\omega_1) = f_{RL^2}(\omega_1);$$
$$x_0 f_R = f_{R^2}(\omega_1);$$
$$f_L q_2(\omega_1) = q_2(0\omega_1);$$
$$f_R q_2(\omega_1) = q_2(1\omega_1).$$

We distinguish 3 cases when the input word starts with 00, 01 and 1, corresponding to all possible inputs.

$$\begin{aligned} x_0 q_2 q_1(00\omega) &= x_0 q_2(0000\omega_2) = x_0 f_{L^4} q_2(\omega_2) = f_{L^2} q_2(\omega_2) \\ &= q_2(00\omega_2) = q_2 q_1(0\omega) = q_2 q_1 X_0(00\omega); \\ x_0 q_2 q_1(01\omega) &= x_0 q_2(001\omega_2) = x_0 f_{L^2R} q_2(\omega_2) = f_{RL^2} q_2(\omega_2) \\ &= q_2(100\omega_2) = q_2 q_1(10\omega) = q_2 q_1 X_0(01\omega); \\ x_0 q_2 q_1(1\omega) &= x_0 q_2(1\omega_2) = x_0 f_R q_2(\omega_2) = f_{R^2} q_2(\omega_2) \\ &= q_2(11\omega_2) = q_2 q_1(11\omega) = q_2 q_1 X_0(1\omega). \end{aligned}$$

Hence, q is in fact maps X_0 to x_0. □

It is similarly straightforward to check that this rational map acts the same on other pairs of generators.

We now have enough material to prove Theorem 3.1.1.

Proof of Theorem 3.1.1. For each $g \in G$, let $\phi(g)$ be the unique element of V_τ that makes the following diagram commute:

$$\begin{array}{ccc} \{0,1\}^\omega & \xrightarrow{\quad g \quad} & \{0,1\}^\omega \\ \Big\downarrow{\scriptstyle h} & & \Big\downarrow{\scriptstyle h} \\ \mathcal{C}_\tau & \xrightarrow{\quad \phi(g) \quad} & \mathcal{C}_\tau \end{array}$$

Let $v \in \{0,1\}^\omega$, then by construction it will be mapped to $[0,1]$ in two different ways $q \circ g(v)$ and $\phi(g) \circ q(v)$, thus ϕ is well defined. We will now show that ϕ is a homomorphism. Let $g_1, g_2 \in G$, then $g_2 \circ g_1$ is also an element of G. Hence, by the definition of ϕ the following holds:

$$q \circ g_1 = \phi(g_1) \circ q;$$
$$q \circ g_2 = \phi(g_2) \circ q;$$
$$q \circ g_2 \circ g_1 = \phi(g_2 \circ g_1) \circ q.$$

By manipulating the first two equations we obtain:

$$q \circ g_2 \circ g_1 = \phi(g_2) \circ \phi(g_1) \circ q.$$

And thus:

$$\phi(g_2 \circ g_1) \circ q = \phi(g_2) \circ \phi(g_1) \circ q.$$

Therefore, ϕ is a homomorphism. Figure 3.4 represents a diagram of this homomorphism for an arbitrary pair $g_1, g_2 \in G$. Combining this with Lemma 3.1.5 makes ϕ an isomorphism from G to V_τ. □

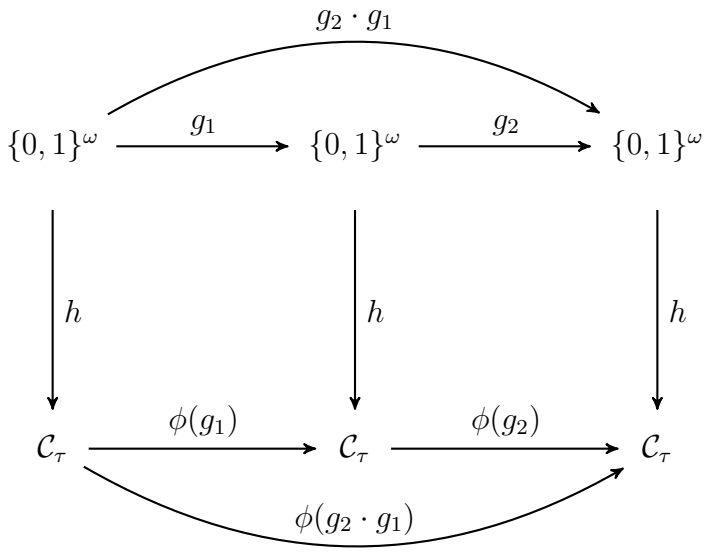

Figure 3.4: A diagram representing the homomorphism of ϕ

Since by Theorems 2.3.5 and 2.3.6 the generators of both T_τ and F_τ are contained in V_τ we obtained the following two corollaries.

Corollary 3.1.7. *The group of homeomorphisms of $\{0,1\}^\omega$ generated by X_0, X_1, Y_0 and Y_1 is isomorphic to F_τ. See Figure 3.1 for the automata.*

Corollary 3.1.8. *The group of homeomorphisms of $\{0,1\}^\omega$ generated by X_0, X_1, Y_0, Y_1, C_1 and C_2 is isomorphic to T_τ. See Figure 3.1 for the automata.*

3.2 Nucleus of the automata that generates V_τ

Having established the automaton representation of V_τ, we now turn our attention to the nucleus, which is a crucial concept in automata theory. The nucleus of an automaton provides a compact representation of its essential states and transitions. In this section, we will identify the nucleus for the automaton generating V_τ, which will offer deeper insights into the group's structure.

Definition 3.2.1. [4, Section 2.6] Let $f \in \mathcal{R}_{X,E}$ be a non-degenerate map and α, β be finite paths in X, then f has only a finite number of local actions $f|_\alpha$. The set of all local actions that occur infinitely many times is called the *nucleus* and is denoted by $\mathcal{N}_f$. Formally $\mathcal{N}_f = \{f|_\alpha : f|_\alpha = f|_\beta$ for infinitely many $\beta\}$. As a consequence $f|_\alpha \notin \mathcal{N}_f$ only for a finite number of different α.

Definition 3.2.2. [4, Definition 2.41] Let $G \leq \mathcal{R}_X$, then the *nucleus* $\mathcal{N}_G$ of G is defined as the union of all $\mathcal{N}_g$ for $g \in G$. Hence, $\mathcal{N}_G$ is the smallest set of functions such that for any $g \in G$, $g|_\alpha \in \mathcal{N}_G$ holds for all but for a finite number of $\alpha \in X^*$. We call $G \leq \mathcal{R}_X$ *contracting* whenever $\mathcal{N}_G$ is finite and Σ_X has an irreducible core.

Definition 3.2.3. [4, Definition 2.43] A set of maps $\mathcal{N}$ is the *nucleus of injections* if it satisfies the following conditions:

1. The identity map belongs to $\mathcal{N}$;
2. For every map $x \in \mathcal{N}$, $x|_\alpha \in \mathcal{N}$.
3. For every map $x \in \mathcal{N}$ there exists $f \in \mathcal{N}$ such that $x \in \mathcal{N}_f$;
4. For every map $x \in \mathcal{N}$, $\mathcal{N}_{x^{-1}} \subseteq \mathcal{N}$;
5. For every pair of maps x and $y \in \mathcal{N}$, $\mathcal{N}_{xy} \subseteq \mathcal{N}$.

Theorem 3.2.4. *[4, Theorem 2.46] Let $\mathcal{N}$ be a nucleus of injections over Σ_X and $E \subseteq \Sigma_X$ be an nonempty clopen set, then*

$$G = \{f \in \mathcal{R}_{X,E} : \mathcal{N}_f \subseteq \mathcal{N}\}$$

is a rational similarity group with nucleus $\mathcal{N}$.

Theorem 3.2.5. *The nucleus of the automaton that generates the Thompson group V_τ is given by the states $\{\beta, \gamma, id\}$.*

Proof. To identify the nucleus of the automaton that generates the Thompson group V_τ, we make an educated guess based on the properties of the nucleus. It is known that every map in the nucleus must occur for infinitely many words. The only maps that could potentially satisfy this condition are β, γ, id since nucleus elements must follow a cycle in the automaton diagram, see Figure 3.1. Let us denote the set of these maps as $\mathcal{N} = \{\beta, \gamma, id\}$.

We now proceed to verify that the set $\mathcal{N}$ satisfies the required conditions for being the nucleus of injections.

The first condition follows since the identity map id belongs to $\mathcal{N}$.

The second condition is satisfied, as for every map $x \in \mathcal{N}$, $x|_\alpha \in \mathcal{N}$ for all α, by observation of Figure 3.1.

The third condition is immediate since $\{\beta, \gamma, id\} \subseteq \mathcal{N}_\beta$ by observation of Figure 3.1.

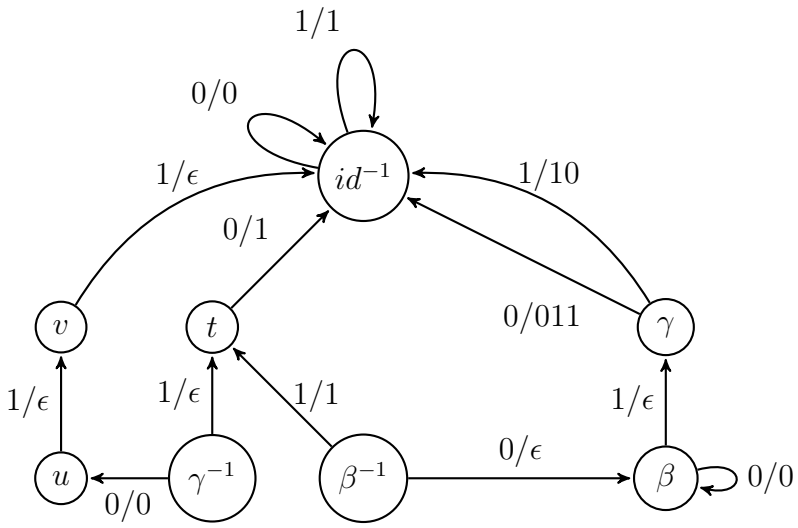

Figure 3.5: Automata of the the inverse maps of β, γ and id of the Thompson group V_τ

The fourth condition is verified through direct computation of the inverse maps of β, γ, and id, as shown in Figure 3.5. Observe that β^{-1} is defined for words that start with 10 or 0, and $\beta^{-1}|_{10} = id$, and $\beta^{-1}|_0 = \beta$. Whereas γ^{-1} is defined only for words that start with 011 or 10, and $\gamma^{-1}|_{011} = \gamma^{-1}|_{10} = id$. Therefore, $\mathcal{N}_{\beta^{-1}} = \{\beta, \gamma, id\}$, $\mathcal{N}_{\gamma^{-1}} = \{id\}$ and $\mathcal{N}_{id^{-1}} = \mathcal{N}_{id}$. Thus, the fourth condition is satisfied.

The fifth condition is the most complex, as it requires us to consider all possible compositions of maps in the nucleus. We will analyze these compositions case by case, showing that they always result in maps within our proposed nucleus. Verifying that the nucleus of a product with the identity state id belongs to the set $\mathcal{N}$ is trivial. Therefore, we only need to confirm that the nuclei of γ^2, $\gamma\beta$, $\beta\gamma$, and β^2 are contained in the set $\mathcal{N}$. We can verify this by considering the following eight possible scenarios:

$$\begin{aligned}
&\gamma^2(0\omega) = \gamma(1\omega) = 0\omega;\\
&\gamma^2(1\omega) = \gamma(0\omega) = 1\omega;\\
&\gamma\beta(0\omega) = \gamma(0)\beta(\omega) = 1\beta(\omega);\\
&\gamma\beta(1\omega) = \gamma^2(\omega);\\
&\beta\gamma(0\omega) = \beta(1\omega) = 1\gamma(\omega);\\
&\beta\gamma(1\omega) = \beta(0\omega) = 0\beta(\omega);\\
&\beta^2(0\omega) = \beta(0)\beta(\omega) = 0\beta^2(\omega);\\
&\beta^2(1\omega) = \beta(1)\gamma(\omega) = 1\gamma^2(\omega).
\end{aligned}$$

See Figure 3.6 for the automaton diagram of these maps. Observe that the maps γ^2 and β^2 act as the identity. Whereas the maps $\gamma\beta$ and $\beta\gamma$ map all possible sequences to the maps γ, β and $\gamma^2 = id$. Therefore, since every possible sequence gets mapped back to a state in the set $\mathcal{N}$, the fifth condition is also satisfied.

Therefore, the set $\mathcal{N} = \{\beta, \gamma, id\}$ satisfies all five conditions required for being the nucleus of the

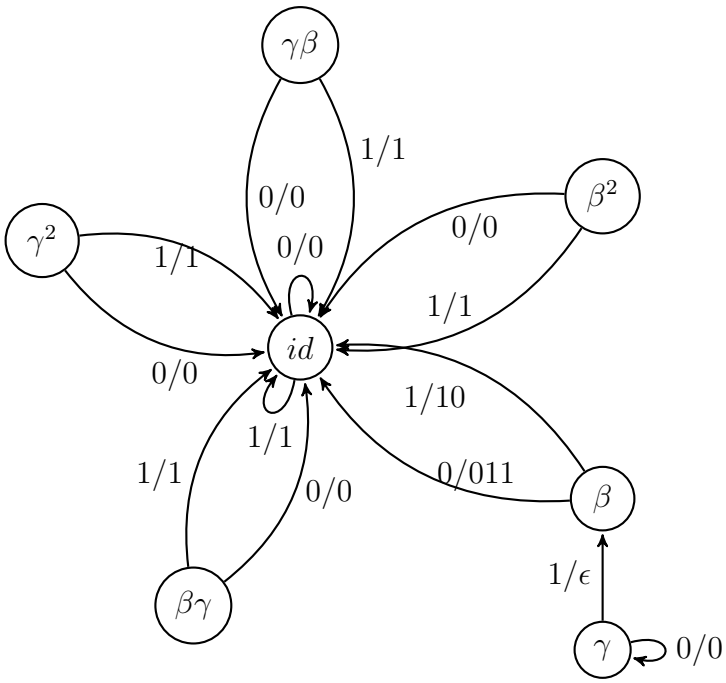

Figure 3.6: The composition of states of $\mathcal{N}$

automaton that generates the Thompson group V_τ.

Furthermore, given that the binary alphabet is irreducible and the nucleus $\mathcal{N}$ is finite, we conclude that the nucleus is contracting. □

Remark 3.2.6. It is not hard to see that V_τ is precisely this group G, in particular V_τ is a rational similarity group.

Chapter 4

The monoid M

This section introduces the reader to the main object of study in this thesis: the monoid $M = \langle L, R : LR^2 = RL^2 \rangle$. We start by building up the necessary language to talk about graphs, after which we introduce the monoid M and its associated Cayley graph. We end the section by presenting the distance formula for the Cayley graph of M and by introducing an intriguing representation of M in terms of the real line.

4.1 Preliminary on graphs

In this section we restate some graph background definitions stated in [25, Section 2]. We will define what a graph is, impose a partial order relation on it and induce a metric and a topology on it.

An *undirected graph* $\Gamma = (V, E)$ consists of a countable set V, elements of which are called *vertices*, and a symmetric subset $E \subseteq \big\{\{x, y\} : x, y \in V \text{ and } x \neq y\big\}$, whose elements are unordered pairs and are called *edges*. We represent edges as pairs of vertices, we do not allow loops, and we allow at most one edge between any given pair of vertices. The graph is said to be *locally finite* if for every vertex $x \in V$, $\deg(x) := \#\{y : \{x, y\} \in E\} < \infty$, in other words, if every vertex emits a finite number of edges. A path of length n is a sequence of edges $\{x_0, x_1\}, \{x_1, x_2\} \dots \{x_{n-1}, x_n\}$, where paths of length 0 are vertices.

A *geodesic* is a shortest path between two vertices. For $x, y \in V$, a geodesic from x to y is denoted by $\pi(x, y)$. The distance between x and y is the length of the geodesic and is denoted by $d(x, y)$. In case no geodesic exists then $d(x, y) = \infty$. If for all $x, y \in V$, $d(x, y)$ is finite then the graph $\Gamma = (V, E)$ is *connected*; in such case d is a word metric on V.

A *root* of a connected graph is a fixed vertex from which there exists a least one geodesic to every other vertex in the graph. When a graph (V, E) is a locally finite connected graph and there exists a vertex $r \in V$ that acts as a fixed root then the triple (V, E, r) is called a *rooted graph*. From now on we reserve the variable r for the root of a given graph. Note that for a given connected graph any vertex can be chosen as its root. For $x, r \in V$ we denote $|x| := d(v, r)$ and let $V_n := \{x \in V : |x| = n\}$. Then $V = \bigcup_{n=0}^{\infty} V_n$. We introduce a partial order relation $\preceq$ on V with $y \preceq x$ if and only if y belongs

to some $\pi(r, x)$.

Let's briefly check that $\preceq$ is indeed a partial order relation.

Let $x \in V$, then $x \in \pi(r, x)$ then $x \preceq x$, hence, the relation $\preceq$ in reflexive.

Let $x, y \in V$ be such a pair that $x \preceq y$ and $y \preceq x$. Then $x \in \pi(r, y)$ implies that $|x| \leq d(r, y) = |\pi(r, y)| = |y|$. And $y \in \pi(v, x)$ implies that $|y| \leq d(v, x) = |\pi(r, x)| = |x|$. Hence, $|x| = |y|$, but $x \in \pi(r, y)$, therefore $x = y$. This shows that the relation $\preceq$ is antisymmetric.

Let $x, y, z \in V$ such that $x \preceq y$ and $y \preceq z$. Then $y \in \pi(r, x)$ and $z \in \pi(r, y)$. Since $y \in \pi(r, x)$ then a geodesic $\pi(r, x)$ can be decomposed into $\pi(r, y)$ and $\pi(y, x)$. But $z \in \pi(r, y)$, hence, $z \in \pi(r, x)$. This shows that $\preceq$ is transitive.

Hence, $\preceq$ is indeed a partial order relation.

Let $m \geq 0$ and $x \in V$, then

$$J_m(x) := \{y \in V : x \preceq y, |y| = |x| + m\}$$

is called the *m-th descendant set*, and

$$J_{-m}(x) := \{z \in V : x \in J_m(z)\}$$

is called the *m-th predecessor set* of x respectively. In principle we allow $J_m(x)$ to be empty.

4.2 The monoid M and its Cayley graph $\mathcal{M}$

In this section we define the monoid M, properties of which we will analyse in the remaining sections of this paper. A major goal of this paper is to study the geometric properties of M, such as horofunction boundary and hyperbolicity. Therefore we will also present the Cayley graph $\mathcal{M}$.

Let $M = \langle L, R : LR^2 = RL^2 \rangle$ be the monoid generated by 2 elements L and R with the action being multiplication by the right with the relation $LR^2 = RL^2$. As like any other monoid, M satisfies the following properties:

- closure, i.e., if $x, y \in M$ then $x \cdot y \in M$;
- associativity, i.e., if $x, y, z \in M$ then $(x \cdot y) \cdot z = x \cdot (y \cdot z)$;
- presence of the identity element, i.e., there exists $1 \in M$ such that for all $x \in M$, $1 \cdot x = x \cdot 1 = x$.

Note that unlike some other monoids, M lacks right cancellation, except for the identity element.

The Cayley graph of M is an undirected locally finite rooted graph denoted by $\mathcal{M} = (M, E, 1)$, where:

- the vertex set M is in a bijection with the set of elements of the monoid M, i.e., every vertex in $\mathcal{M}$ represents a unique element in M;

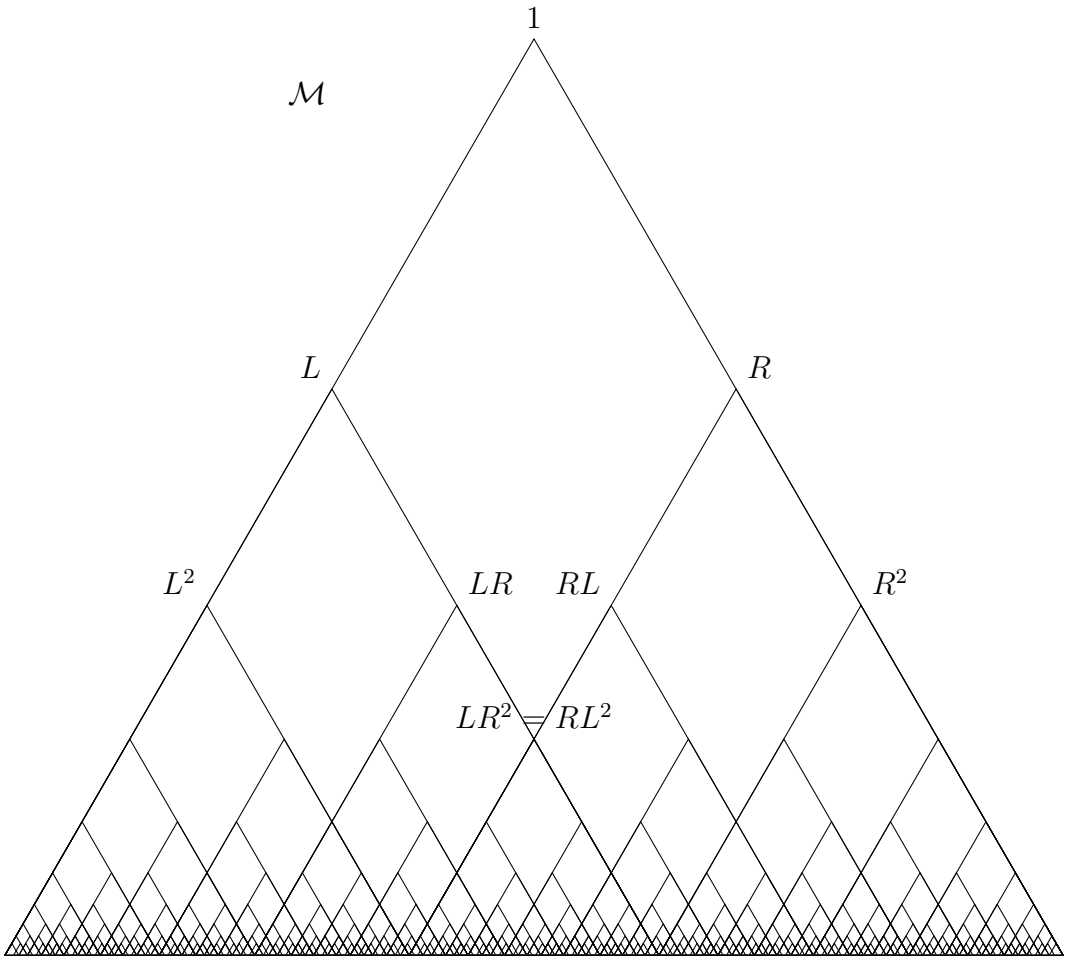

Figure 4.1: The Cayley graph $\mathcal{M}$ of the monoid M

- the edge set E consists of all pairs of vertices $x, y \in M$ such that $x = yL$ or $x = yR$, i.e., there exists an edge between two vertices if one of their corresponding elements in M can be obtained from the second one by right multiplication by L or R;
- the root of the graph is the identity element 1 of the monoid M.

In the case of the monoid M, the defining relation $LR^2 = RL^2$ imposes specific self-similarity constrains on the structure of the Cayley graph $\mathcal{M}$. For a visual representation of the Cayley graph $\mathcal{M}$ refer to Figure 4.1. Note that the picture does not accurately represent the bottom layer, we expect the reader to imagine a fractal structure at the boundary of the graph.

Definition 4.2.1. For any pair of vertices $x, y \in M$ in the graph $\mathcal{M} = (M, E, 1)$, a *code* for a path between x and y is a word $w = w_1 w_2 \cdots w_n$, where each $w_i \in L, R, L^{-1}, R^{-1}$, with L and R being the generators of the monoid M, and L^{-1} and R^{-1} being their formal inverses. The code for a path w from x to y satisfies the equation $xw = xw_1w_2 \cdots w_n = y$ in the free group generated by L and R. The length of w is defined as n corresponding to the number of elements in the word w. We denote the code of a geodesic between x and y by $c(x, y)$.

While L^{-1} and R^{-1} are used in the representation, they do not exist in the monoid M itself, but rather represent backtracking in the Cayley graph. In the actual graph $\mathcal{M}$, each step of the path corresponds to an edge, which is always a right multiplication by either L, R L^{-1} or R^{-1} with inverses indicating moving "backwards" along an L or R edge, respectively.

The *cone* of an element $x \in M$, denoted by $Cone(x)$, is the set of all elements in M that can be obtained by multiplying x on the right by any element of the monoid M, formally,

$$Cone(x) = \{x \cdot m : m \in M\},$$

in other words $m \in M$ is in the cone of x when x is a prefix of m.

Proposition 4.2.2. *If $x, y \in M$ then $Cone(x) \subseteq Cone(y)$ if and only if $x \in Cone(y)$.*

Proof. We will start from the forward condition, let $x, y \in M$ such that $Cone(x) \subseteq Cone(y)$. Then $x \in Cone(x) \subseteq Cone(y)$.

For the opposite direction let $x \in Cone(y)$. Then there exists $m \in M$ such that $x = ym$, then $Cone(x) = Cone(ym) \subseteq Cone(y)$. □

We end this section with a small lemma, which is clear from Figure 4.1.

Lemma 4.2.3. $Cone(L) \cap Cone(R) = Cone(LR^2) = Cone(RL^2)$.

4.3 The distance formula for the graph $\mathcal{M}$

In this section we introduce the distance formula for the graph $\mathcal{M}$. We start by proving that every cone in $\mathcal{M}$ is geodesically convex and then through a number of lemmas we introduce a geodesic formula for any given pair of elements in $\mathcal{M}$. The main purpose of this section is to prove the following theorem, which will be used later in multiple sections.

Theorem 4.3.1. *Let x, y be any given pair of vertices in the graph $\mathcal{M}$. Then the distance between x and y is given by:*

$$d(x,y) = \begin{cases} d(x',y'), \text{ when } x = mx', y = my' \text{ for some nontrivial } m \in M; \\ |x| + |y| - 2, \text{ when } x \in Cone(LR), y \in Cone(RL) \text{ and } x, y \notin Cone(LR^2); \\ |x| + |y|, \text{ else.} \end{cases}$$

This formula offers a straightforward approach to calculating the distance between any given pair of vertices x, y in M. There is a total of three steps to follow. The first step is to remove the common prefix, resulting in obtaining x' and y'. The next step is to identify the first pair of prefixes of x' and y'. Finally, if $x' \in Cone(LR)$ and $y' \in Cone(RL)$, then the distance is $|x'| + |y'| - 2$; if not, it is simply $|x'| + |y'|$.

We will prove this theorem by a number of lemmas throughout this section.

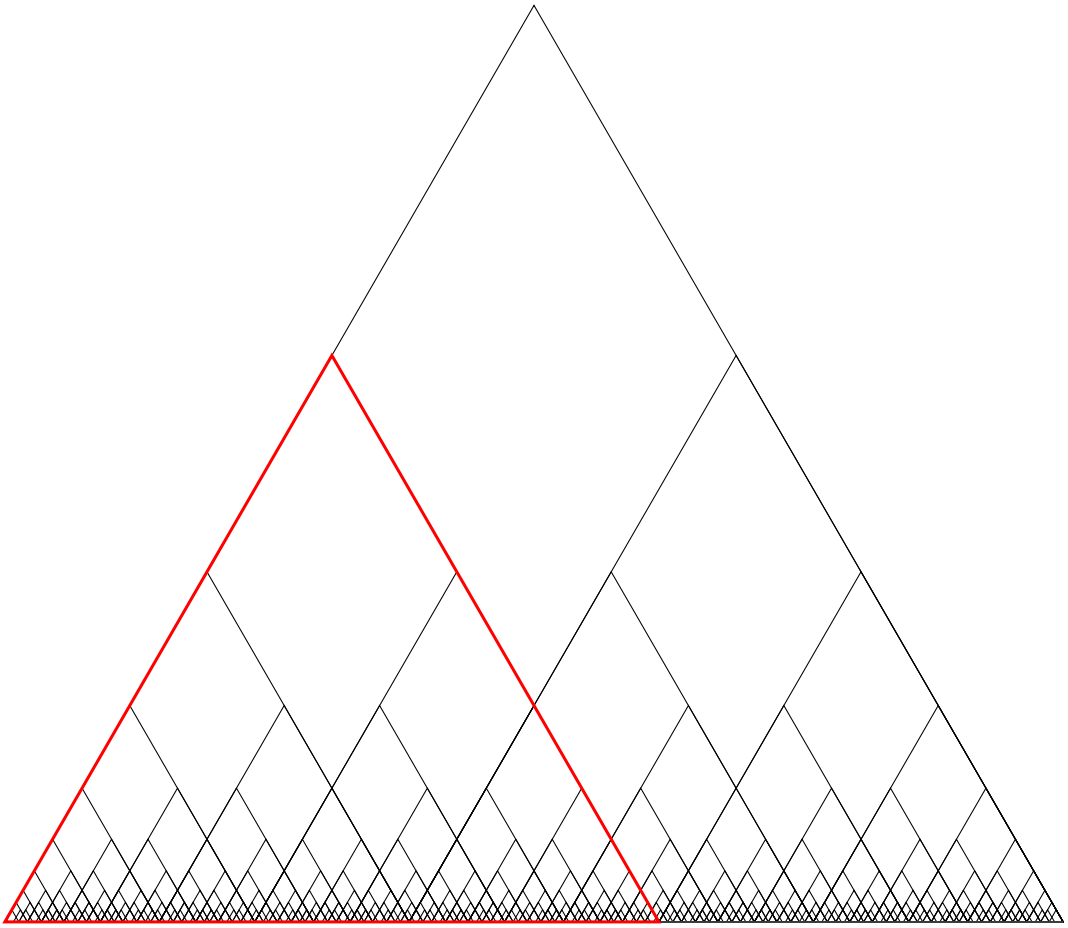

Figure 4.2: $Cone(L)$ of the graph $\mathcal{M}$

Definition 4.3.2. Let $S \subseteq X$ be a subset of a metric space X. The set S is said to be *geodesically convex* if for any two elements $x, y \in S$, there exists at least one geodesic joining x and y that is entirely contained in S.

In addition when every geodesic joining any two elements $x, y \in S$ is entirely contained in S, the set is called *strongly geodesically convex*.

Proposition 4.3.3. *Each cone of $\mathcal{M}$ is strongly geodesically convex.*

Proof. Let $x, y \in Cone(L)$ (represented by the red triangle in Figure 4.2). Let's assume, for the sake of contradiction, that a geodesic $\pi(x, y)$ at some point leaves $Cone(L)$. Then $\pi(x, y)$ can be decomposed into 3 subgeodesics $\pi(x, y) = \pi(x, LR^m) \cdot \pi(LR^m, LR^n) \cdot \pi(LR^n, y)$. We will focus on the second subgeodesic $\pi(LR^m, LR^n)$, as it represents the portion of the path that supposedly leaves $Cone(L)$.

We will now show that unique code $c(LR^m, LR^n)$ is R^{m-n} (without loss of generality let $m > n$). Any deviation from the direct ascending path, whose code is R^{m-n}, would introduce additional steps, making the path longer. Therefore, R^{m-n} is the unique code of the geodesic from LR^m to LR^n.

This shortest path from LR^m to LR^n lies entirely within $Cone(L)$. This contradicts the assumption that a geodesic $\pi(LR^m, LR^n)$ leaves $Cone(L)$. Hence, $Cone(L)$ is strongy geodesically convex.

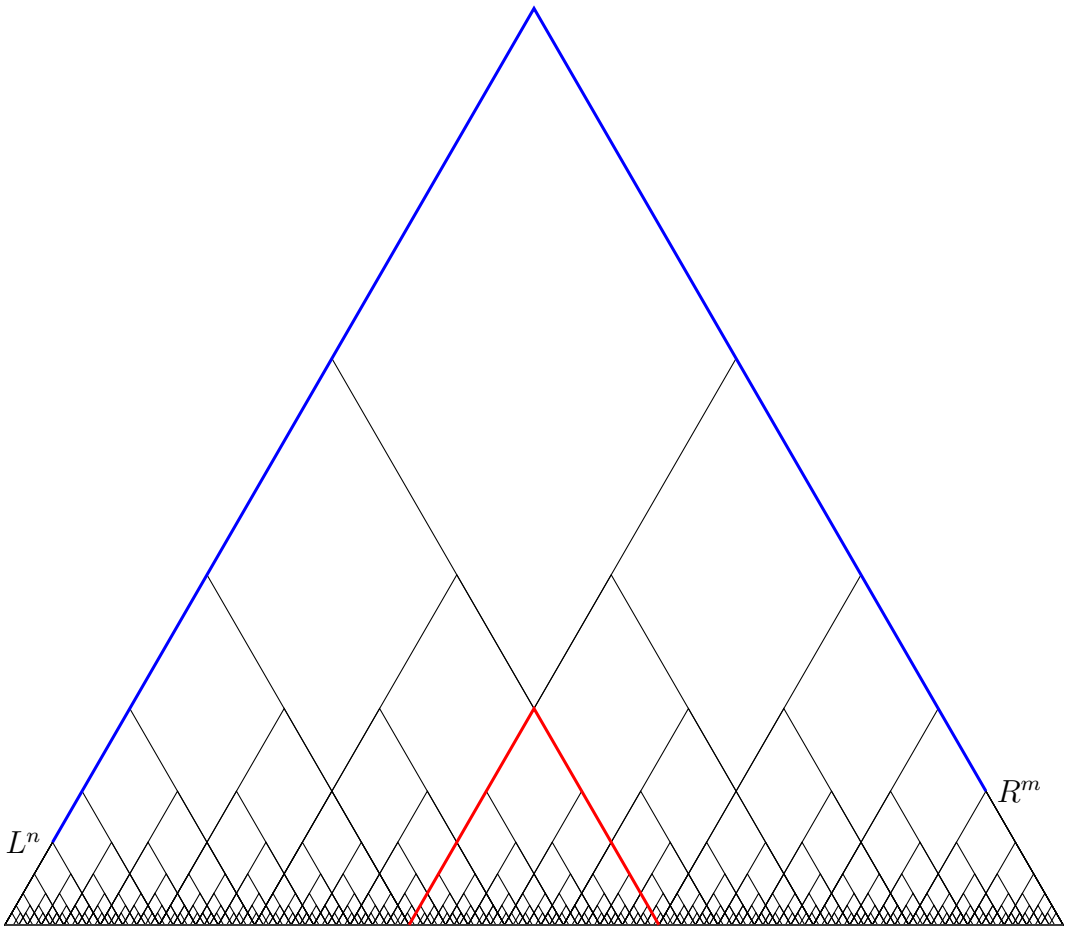

Figure 4.3: The graph $\mathcal{M}$ with the $L^{-n}R^m$ code and $Cone(LR^2)$ highlighted

Due to symmetry, we apply the same argument to $Cone(R)$ by interchanging the roles of L and R.

Observe that $\mathcal{M}$ is isometric to $Cone(L)$ and $Cone(R)$. Now since $Cone(LL)$ and $Cone(LR)$ are strongly geodesically convex in $Cone(L)$ they are also strongly geodesically convex in $\mathcal{M}$. By induction everything is $Cone(L)$ and $Cone(R)$ is strongly geodesically convex.

Suppose $Cone(m)$ is strongly geodesically convex, then $Cone(m)$ is isometric to $\mathcal{M}$. Since $Cone(L)$ and $Cone(R)$ are strongly geodesically convex in $\mathcal{M}$, then $Cone(mL)$ and $Cone(mR)$ are strongly geodesically convex in $Cone(m)$. Thus, $Cone(mR)$ are strongly geodesically convex in $\mathcal{M}$. Hence, every $Cone(m) \in \mathcal{M}$ is strongly geodesically convex. □

Proposition 4.3.4. *Let $n, m \geq 0$, then the code of the unique geodesic $\pi(L^n, R^m)$ is $L^{-n}R^m$.*

See Figure 4.3 for an illustration of the unique geodesic $\pi(L^n, R^m)$ in blue and $Cone(LR^2)$ in red. This visualization will aid in understanding the subsequent proof by contradiction.

Proof. We will prove this by induction on k. For $k \geq 0$, let $n, m \geq 0$, then $n + m \leq k$ implies that $d(L^n, R^m) = n + m$. Let $k \leq 1$, then the statement holds trivially. Now let's assume that the statement holds for $k - 1$. We will now prove it for k.

Suppose, for contradiction, that $L^{-n}R^m$ is not the code of a geodesic $c(L^n, R^m)$, and does not visit the root 1. Then this geodesic must pass through $Cone(LR^2)$. In this case, a code $c(L^n, R^m)$ can be decomposed into 2 subcodes:

$$c(L^n, R^m) = c(L^n, RL^2L^a) \cdot c(RL^2L^a, R^m),$$

where $a \leq n-2$. If $a \geq n-2$, we will end-up with a code longer than $L^{-n}R^m$, contradicting the assumption that $c(L^n, R^m)$ is not the code of a geodesic.

Let's look at the second subcode $c(RL^2L^a, R^m)$. Due to the convex property of the monoid M, a code $c(RL^2L^a, R^m)$ cannot leave $Cone(R)$, hence $c(RL^2L^a, R^m) = c(L^2L^a, R^{m-1})$.

Note that $(a+2)+(m-1) \leq n+m-1 = k-1$. Then by our induction hypothesis, the length of a code $c(RL^2L^a, R^m)$ is $a+m+1$. The code $L^{-2}L^{-a}R^{m-1}$ is of such length, this implies that a code $c(L^n, R^m)$ visits LR^2.

Given that a code $c(L^n, R^m)$ visits LR^2, we observe that it can be decomposed as:

$$c(L^n, R^m) = c(L^n, LR^2) \cdot c(LR^2, R^m).$$

Consider the subcode $c(L^n, LR^2)$. From the convex property of $\mathcal{M}$, we observe that $c(L^n, LR^2) = c(L^{n-1}, R^2)$. Since $n-1+2 < k$, by our inductive hypothesis, $d(L^{n-1}, R^2) = n+1$.

This gives $d(L^n, LR^2) = n+1$ and similarly, $d(LR^2, R^m) = m+1$. Hence,

$$d(L^n, R^m) = (n+1)+(m+1) = n+m+2.$$

But this is greater than $n+m$, which is the length of the code $L^{-n}R^m$. This contradicts that assumption that $c(L^n, R^m)$ does not pass through the root 1.

It is left to prove uniqueness. Observe that $c(L^n, R^m) = c(L^n, 1) \cdot c(1, R^m)$, where both subgcodes have unique ascending and descending respective codes. Therefore, the unique code $c(L^n, R^m)$ must be $L^{-n}R^m$. □

Proposition 4.3.5. *Let $x \in \big(Cone(L) \setminus Cone(LR)\big)$, then $d(x, R^n) = |x| + n$.*

Proof. The distance $d(x, R^n)$ depends on the geodesic $\pi(x, R^n)$, which on its half has two possibilities: either it visits the root 1 or it does not. If it does we are trivially done. If $n = 0$ we are also trivially done, let $n > 0$.

For the purpose of contradiction let's assume that there exists a geodesic $\pi(x, R)$ that does not visit the root and has length at most $|x| + n - 1$. If such a geodesic exists, then at some point it must enter and leave $Cone(LR)$. Then Proposition 4.3.4 implies that it must visit LR.

We can split the assumed geodesic into two subgeodesics in the following way:

$$\pi(x, R) = \pi(x, LR) \cdot \pi(LR, R^n).$$

Following Proposition 4.3.4, we observe that $d(LR, R^n) = n + 2$ and has two possible geodesics, one that travels through the root and another one that goes through LR^2. Then $d(x, R^n) = d(x, LR) + d(LR, R^n) = d(x, LR) + n + 2 \leq |x| + n - 1$ implies that $d(x, LR) \leq |x| - 3$.

But this means that

$$d(x, 1) = d(x, LR) + d(LR, 1) \leq |x| - 3 + 2 = |x| - 1.$$

This is a contradiction as there is no path from x to the root of length $|x|-1$. Therefore, if $x \in \big(Cone(L) \setminus Cone(LR)\big)$, then every geodesic $\pi(x, R^n)$ must pass through the root, consequently, $d(x, R^n) = |x| + n$. □

We now have enough tools to introduce a geodesic code formula for the graph $\mathcal{M}$. Before we formally state it, we will explain how it works.

This geodesic code formula offers a straightforward approach to finding a code of a geodesic between any pair of vertices x and y. This is done in a number of steps:

1. identify the smallest $Cone(m)$ such that $x, y \in Cone(m)$, note that m can be the empty word;

2. remove the common prefix m from x and y, obtaining elements x' and y' respectively. Without loss of generality let x' start with L and y' start with R;

3. if $x' \in Cone(LR)$ and $y' \in Cone(RL)$, then one of the geodesics from x to y starts from x, takes the ascending path to mLR, then through mLR^2 travels to mRL, and finally takes the descending path to y. We denote the code of such geodesic as $c(x, y) = (x'^{-1}RL)(R^{-2}L^{-1}y)$;

4. if $x' \notin Cone(LR)$ or $y' \notin Cone(RL)$, then there may exist multiple geodesics from x to y. One of these takes the ascending path from x to the root of $Cone(m)$ (the vertex m) and then takes the descending path to y. We denote the code of such geodesic as $c(x, y) = x'^{-1}y'$.

Proposition 4.3.6. *Let x, y be any given pair of vertices in the graph $\mathcal{M}$, and let x' and y' be elements obtained from x and y by removing their common prefix m. Then the code of a geodesic between x and y is given by the following equation:*

$$c(x, y) = \begin{cases} c(x', y'), \text{ if } x, y \in Cone(m) \text{ for some nontrivial } m \in M; \\ (x^{-1}RL^2)(R^{-2}L^{-1}y), \text{ if } x \in \big(Cone(LR) \setminus Cone(LR^2)\big), y \in \big(Cone(RL) \setminus Cone(LR^2)\big); \\ x^{-1}y, \text{ else.} \end{cases}$$

Before proving Proposition 4.3.6 we would like to invite the reader to view Figure 4.4. This figure

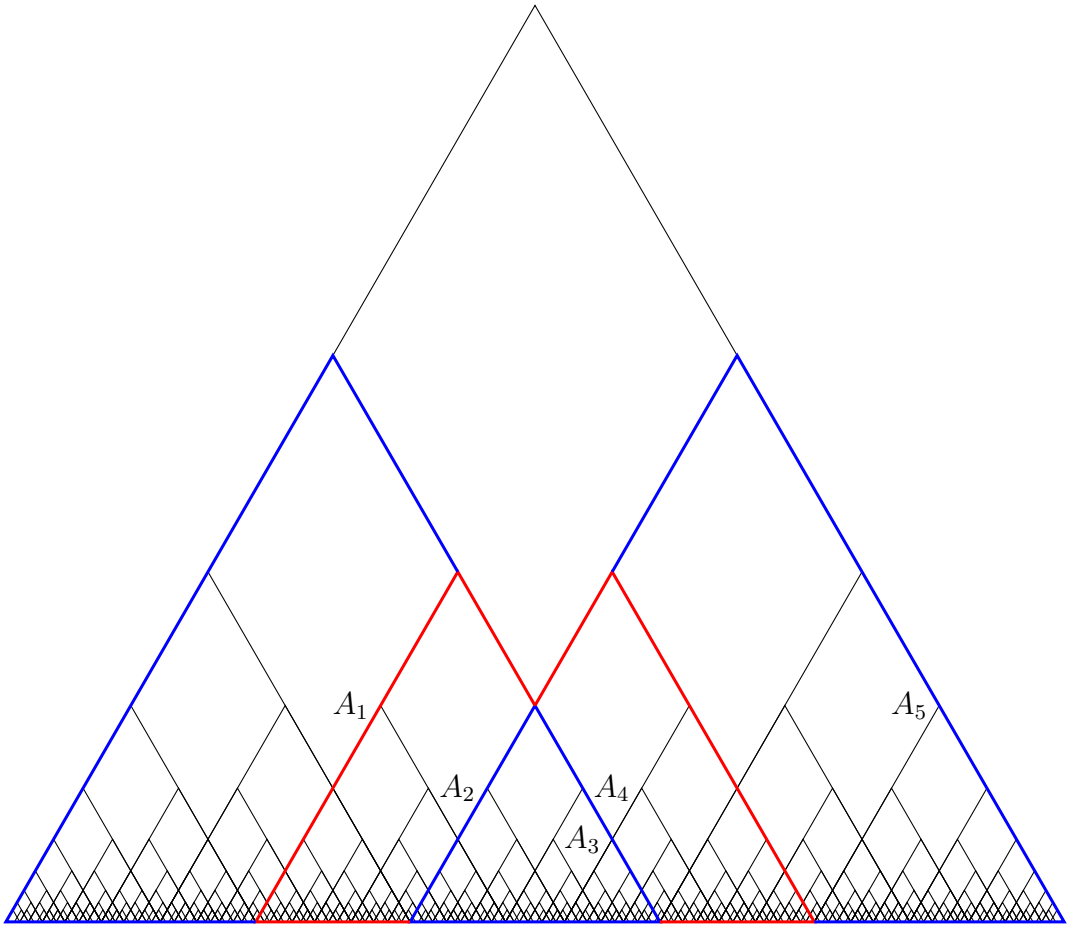

Figure 4.4: The areas A_1 - A_5 on the Cayley graph M

demonstrates 5 subsets of vertices of the graph $\mathcal{M}$.

$$A_1 = Cone(L) \setminus Cone(LR);$$
$$A_2 = Cone(LR) \setminus Cone(LR^2);$$
$$A_3 = Cone(LR^2);$$
$$A_4 = Cone(RL) \setminus Cone(LR^2);$$
$$A_5 = Cone(R) \setminus Cone(RL).$$

Essentially if x' and y' are in the $Cone(m)$, then a geodesic will not pass through the root of $Cone(m)$ only if both x' and y' belong to the areas A_2 and A_4 respectively. If only one of them belongs to its respective area then a geodesic has a choice of passing through the root or not. In case none of them are, then every geodesic passes through the root m.

Proof. Suppose x, y are in $Cone(m)$, where m is not the root, then the first case follows immediately from Proposition 4.3.3.

Now suppose x and y are not in any nontrivial $Cone(m)$. Clearly x and y cannot belong to

$Cone(LR^2)$. Without loss of generality, let x start with L and y start with R. Let's assume that a geodesic $\pi(x, y)$ passes through $Cone(LR^2)$. From Proposition 4.3.4 it is know that such a geodesic must travel along the boundary of $Cone(LR^2)$. Hence, if $\pi(x, y)$ passes though $Cone(LR^2)$ it must visit LR^2.

Therefore to pass from $Cone(L)$ to $Cone(R)$ a geodesic $\pi(x, y)$ must visit 1 or LR^2. This implies that

$$\pi(x, y) = min\{\pi(x, 1) \cdot \pi(1, y),\ \pi(x, LR^2) \cdot \pi(LR^2, y)\}.$$

We only need to find the codes for $\pi(x, 1)$ and $\pi(x, LR^2)$, as codes for y will be analogous cases.

Let's start with the geodesic that passes though the root:

$$\pi(x, 1) = x^{-1} \text{ and } d(x, 1) = |x|.$$

We now analyse the second possible geodesic $\pi(x, LR^2)$. We will split this into 2 cases, when $x \in \big(Cone(LR) \setminus Cone(LR^2)\big)$ (area A_2 of Figure 4.4) and when $x \in \big(Cone(L) \setminus Cone(LR)\big)$ (area A_1). Then by Proposition 4.3.5:

$$\pi(x, LR^2) = \begin{cases} |x| - 1, \text{ when } x \in Cone(LR); \\ |x| + 1, \text{ else.} \end{cases}$$

Due to symmetry we obtain the following:

$$\pi(y, LR^2) = \begin{cases} |y| - 1, \text{ when } y \in Cone(RL); \\ |y| + 1, \text{ else.} \end{cases}$$

This is enough to conclude the following three statements:

1. If $x \in \big(Cone(LR) \setminus Cone(LR^2)\big)$ and $y \in \big(Cone(RL) \setminus Cone(LR^2)\big)$, then a geodesic $\pi(x, y)$ passes though LR^2 and has length $|x| + |y| - 2$. One of the possible codes of a geodesic is $(x^{-1}RL^2)(R^{-2}L^{-1}y)$;

2. If only one of the elements is in $Cone(LR) \setminus Cone(LR^2)$ or $Cone(RL) \setminus Cone(LR^2)$ respectively, then a geodesic has a choice of visiting either the root or LR^2, both resulting in the same length. One of the possible codes for such geodesic is $(x^{-1}y)$

3. When neither of the elements is in these regions, a geodesic $\pi(x, y)$ will always pass thorough the root.

Recall from Proposition 4.3.3 that every cone in $\mathcal{M}$ is strongly geodesically convex. Hence combining this with the initial reduction to the smallest common $Cone(m)$ concludes the proof. □

We can now prove Theorem 4.3.1.

Proof of Theorem 4.3.1. Following Proposition 4.3.6 we take the length of a code between x and y. □

Let x, y be a pair of vertices in $\mathcal{M}$ and let $w \in \{L, R\}^*$. Observe that multiplying x and y from the left by w does not alter their distance. This means that the left multiplication is an isometry. We end this section with a small but important consequence of Theorem 4.3.1.

Corollary 4.3.7. *The monoid M is left cancellative.*

4.4 A representation of the monoid M in terms of the real line

This section introduces a set of subintervals I_M that will allow us to study the properties of M in terms of the real line. Particularly, for every element of M we associate a closed subinterval from the unit interval of the real line.

Definition 4.4.1. Let I_c be the set of all closed intervals $[x, y] \subseteq [0, 1]$, such that $x, y \in \mathbb{R}$ and $x < y$.

If $[x, y] \subseteq I_C$ then we define the subinterval as follows:

$$\begin{aligned} [x, y]_1 &\coloneqq [x, y]; \\ [x, y]_L &\coloneqq [x, y - (y - x)\tau^2]; \\ [x, y]_R &\coloneqq [x + (y - x)\tau^2, y]. \end{aligned} \tag{4.1}$$

Recall that τ is the positive root of the equation $x + x^2 = 1$, approximately 0.618034.

If $w = w_1 w_2 \cdots w_n$ is a word in $\{L, R\}^*$, we define $[x, y]_w$ recursively by $[x, y]_w = ([x, y]_{w_1 w_2 \cdots w_{n-1}})_{w_n}$.

Definition 4.4.2. Let $w = w_1 w_2 \cdots w_n$ be a word in $\{L, R\}^*$, then we define $I_M \subseteq I_c$ to be the collection of all subinterval $[0, 1]_w$.

Lemma 4.4.3. *Let $[x, y] \in I_M$, then the following equality holds:*

$$[x, y]_{LR^2} = [x, y]_{RL^2}.$$

Proof. We will verify this rule by direct computation:

$$\begin{aligned} [x, y]_{LR^2} &= [x, x\tau^2 + y - y\tau^2]_{R^2} \\ &= [x - x\tau^2 + x\tau^4 + y\tau^2 - y\tau^4, x\tau^2 + y - y\tau^2]_R \\ &= [x - 2x\tau^2 + 3x\tau^4 - x\tau^6 + 2y\tau^2 - 3y\tau^4 + y\tau^6, x\tau^2 + y - y\tau^2] \\ &= [x - x\tau^2 + y\tau^2, x\tau^2 + y - y\tau^2]; \end{aligned}$$

$$\begin{aligned} [x, y]_{RL^2} &= [x - x\tau^2 + y\tau^2, y]_{L^2} \\ &= [x - x\tau^2 + y\tau^2, x\tau^2 - x\tau^4 + y - y\tau^2 + y\tau^4]_L \\ &= [x - x\tau^2 + y\tau^2, 2x\tau^2 - 3x\tau^4 + x\tau^6 + y - 2y\tau^2 + 3y\tau^4 - y\tau^6] \\ &= [x - x\tau^2 + y\tau^2, x\tau^2 + y - y\tau^2]. \end{aligned}$$

The final simplification in both cases relies on the property $\tau = \tau^2 + \tau^3$. Both computations give the same result, thus proving the lemma. □

Lemma 4.4.4. *If $x, y \in M$ then $Cone(x) \subseteq Cone(y)$ if and only if $[0,1]_x \subseteq [0,1]_y$. See Figure 4.5 for an illustration.*

Proof. We will start by proving the forward direction. Let $x, y \in M$ such that $Cone(x) \subseteq Cone(y)$. Then there exists $m \in M$ such that $x = ym$. Therefore,

$$[0,1]_x = [0,1]_{ym} = ([0,1]_y)_m \subseteq [0,1]_y.$$

We will prove the opposite direction by induction on the length of x. Let $x, y \in M$ such that $|y| \leq |x| \leq n$, then $[0,1]_x \subseteq [0,1]_y$ implies $Cone(x) \in Cone(y)$.

Let's verify the base case for $|x| = 1$. Then $Cone(y)$ is either equal to $Cone(x)$ or to the whole monoid $Cone(1)$. Hence, for $x, y \in M$ such that $|y| \leq |x| = 1$, $Cone(x) \subseteq Cone(y)$ implies y being a prefix of x or the empty word. This forces $[0,1]_x \subseteq [0,1]_y$.

Assume the statement holds for all x such that $|x| \leq n$. Now let $|x| = n + 1$. We point out the key property of I_M:

$$[0,1]_L \cap [0,1]_R = [0,1]_{LR^2} = [0,1]_{RL^2}.$$

Without loss of generality, let x and y have no common prefix, let x start with R, and y start with L. Then $[0,1]_{Rx'} \subseteq [0,1]_{Ly'}$ for some x', y'. This implies $[0,1]_{Rx'} \subseteq [0,1]_L$, which is only possible if x' starts with L^2. Therefore,

$$[0,1]_{Rx'} \subseteq [0,1]_{RL^2} = [0,1]_{LR^2} \subseteq [0,1]_L.$$

Hence, $[0,1]_{x'} \subseteq [0,1]_{R^2} \subseteq [0,1]_R$. By the induction hypothesis, $Cone(x') \subseteq Cone(R)$, this forces $Cone(x) \subseteq Cone(RL^2) = Cone(LR^2) \subseteq Cone(L)$. But this means that x and y have a common prefix that can be canceled. Then by the induction hypothesis $Cone(x_2x_3 \cdots x_n) \subseteq Cone(y_2y_3 \cdots y_n)$, which forces $Cone(x) \subseteq Cone(y)$. □

Lemma 4.4.5. *Let $x, y, z \in M$. Then $[0,1]_x \cap [0,1]_y = [0,1]_z$ if and only if $Cone(x) \cap Cone(y) = Cone(z)$.*

Proof. We will first prove the forward condition. Let $[0,1]_x \cap [0,1]_y = [0,1]_z$, then there exists $m, n \in M$ such that $[0,1]_{xm} = [0,1]_z = [0,1]_{yn}$. This implies:

$$[0,1]_z = [0,1]_{xm} \subseteq [0,1]_x \quad \text{and} \quad [0,1]_z = [0,1]_{yn} \subseteq [0,1]_y$$

By Lemma 4.4.4, this is equivalent to:

$$Cone(z) \subseteq Cone(x) \quad \text{and} \quad Cone(z) \subseteq Cone(y)$$

Therefore, $Cone(z) \subseteq Cone(x) \cap Cone(y)$.

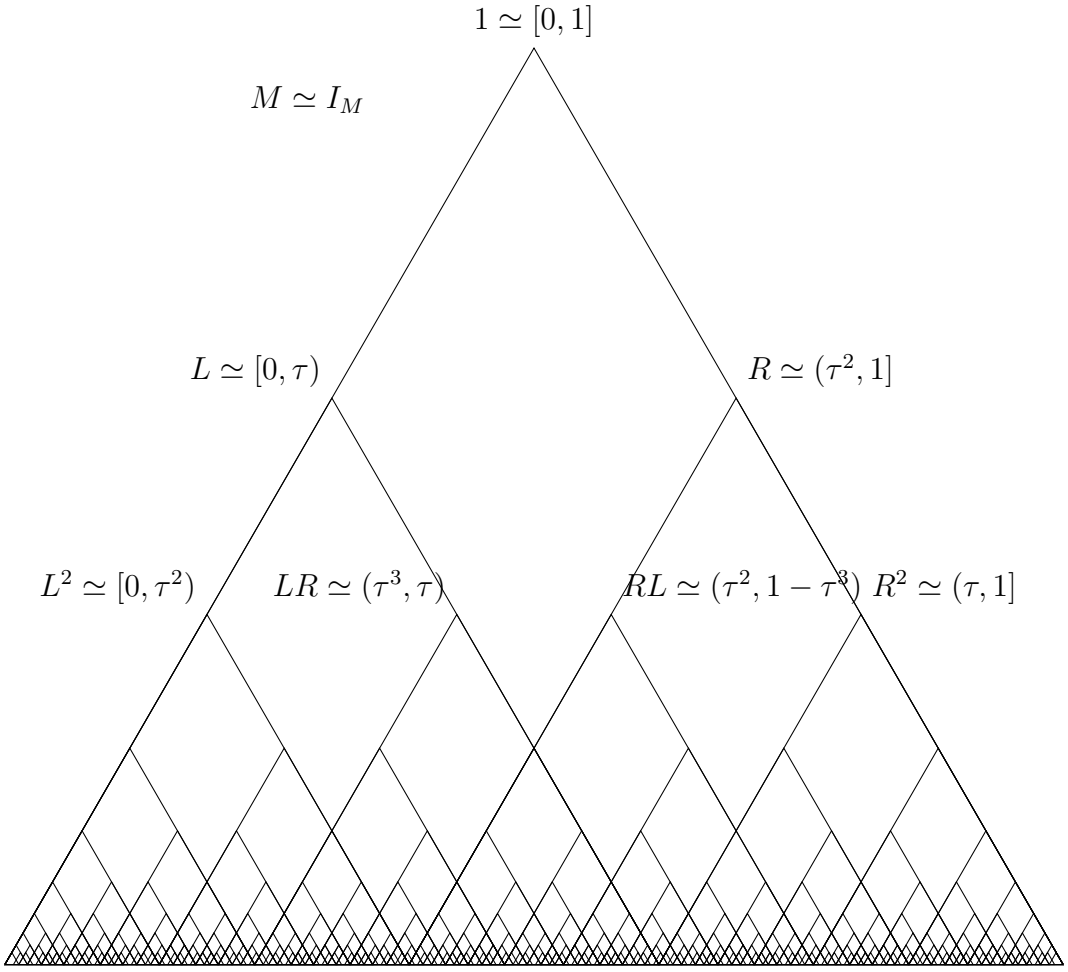

Figure 4.5: The Cayley graph $\mathcal{M}$ with labeled elements of I_M

To show the reverse inclusion, let $w \in Cone(x) \cap Cone(y)$. Then $[0,1]_w \subseteq [0,1]_x \cap [0,1]_y = [0,1]_z$, which implies $w \in Cone(z)$. Thus, $Cone(x) \cap Cone(y) \subseteq Cone(z)$. Now by combining both inclusions, we conclude $Cone(x) \cap Cone(y) = Cone(z)$.

For the opposite direction, let $Cone(x) \cap Cone(y) = Cone(z)$. Then there exist $m, n \in M$ such that $Cone(xm) = Cone(z) = Cone(yn)$. This implies:

$$Cone(z) \subseteq Cone(x) \quad \text{and} \quad Cone(z) \subseteq Cone(y)$$

By Lemma 4.4.4, this is equivalent to:

$$[0,1]_z \subseteq [0,1]_x \quad \text{and} \quad [0,1]_z \subseteq [0,1]_y$$

Therefore, $[0,1]_z \subseteq [0,1]_x \cap [0,1]_y$.

It is left to check the reverse inclusion, let $t \in [0,1]_x \cap [0,1]_y$. Then by Lemma 4.4.4 the element of M corresponding to t is in $Cone(x) \cap Cone(y) = Cone(z)$, which implies $t \in [0,1]_z$. Thus, $[0,1]_x \cap [0,1]_y \subseteq [0,1]_z$. By combining both inclusions, we conclude $[0,1]_x \cap [0,1]_y = [0,1]_z$. □

Chapter 5

The horofunction boundary of $\mathcal{M}$

This section presents one of the main results of this paper: the horofunction boundary of $\mathcal{M}$. We will start by briefly introducing the reader to the geometric concept of horofunction boundaries, introduced by M. Gromov in [22]. This will build a solid background for finding new properties of the monoid M.

5.1 Preliminary on horofunctions

This section introduces key definitions and properties of horofunctions in the context of graph theory. We briefly restate some modified definitions stated in [5, Section 1] and introduce an alternative approach in finding the horofunction boundary using equivalence classes relying on J. Belk, C. Bleak, and F. Matucci paper [5, Section 1.3].

Let $\Gamma = (V, E)$ be a locally finite connected graph. We impose the path metric on V. We define $F(V, \mathbb{Z})$ as the abelian group of all integer-valued functions on V. Let $\overline{F}(V, \mathbb{Z})$ be the quotient of $F(V, \mathbb{Z})$ by the subgroup of constant functions. This means that within $\overline{F}(V, \mathbb{Z})$ two functions f and g from $F(V, \mathbb{Z})$ are equivalent if their difference $f - g$ is a constant function. This construction allows us to focus on the variations of functions over V, where $\overline{F}$ denotes the equivalence classes of F. If $f \in F(V, \mathbb{Z})$ we let $\overline{f}$ denote its image on $\overline{F}(V, \mathbb{Z})$.

Observe that $F(V, \mathbb{Z}) = \mathbb{Z}^V$ is a topological space under the product topology. Consequently, $\overline{F}(V, \mathbb{Z})$, being a quotient of $F(V, \mathbb{Z})$, naturally inherits a quotient topology.

Definition 5.1.1. [5, Definition 1.22] Let $\Gamma = (V, E)$ a locally finite connected graph. Let $x \in V$ then for all $y \in V$ the corresponding *distance function* $d_x : V \to \mathbb{Z}$ is defined by:

$$d_x(y) = d(x, y).$$

The function $i : V \to \overline{F}(V, \mathbb{Z})$ defined by

$$i(x) = \overline{d}_x$$

for all $x \in V$, is called the *canonical embedding*.

Definition 5.1.2. [5, Definition 1.23] Let $\partial_h V$ be the set of accumulation points of $i(V)$ in $\overline{F}(V,\mathbb{Z})$, this set is called the *horofunction boundary* of V.

We call a function $f : V \to \mathbb{Z}$ a *horofunction* when $\overline{f} \in \partial_h V$.

The horofunction boundary $\partial_h V$ is equipped with the topology of pointwise convergence. Specifically, for a sequence of functions $\overline{f}_n \in \overline{F}(V,\mathbb{Z})$ and a function $f \in \overline{F}(V,\mathbb{Z})$, we say $f_n \to f$ if and only if for every $v \in V$, $f_n(v) \to f(v)$ as $n \to \infty$. This topology can be alternatively described as the subspace topology inherited from the product topology on $\mathbb{Z}^V$, where $\mathbb{Z}$ is given the discrete topology, additionally this give us the quotient topology on $\overline{F}(V,\mathbb{Z})$.

The definition of the horofunction boundary given above, while mathematically precise, may not provide intuitive insight into its structure. Following [5], we will adopt an alternative characterization of the horofunctions introduced by Belk, Bleak, and Matucci. This approach uses combinatorial methods that are more aligned with the graph-theoretic context presented in this paper and provides a less complex realization of horofunctions.

A *vector field* on a locally connected graph $\Gamma = (V, E)$ with respect to a chosen distance function d is a function $F : E \to E \cup E_{\pm}$, where $E_{\pm} = \big\{(x, y) : \{x, y\} \in E\big\}$, such that for all pairs $x, y \in E$:

$$F(\{x,y\}) = \begin{cases} \{x,y\} & \text{if } d(x) = d(y); \\ (x,y) & \text{if } d(y) < d(x); \\ (y,x) & \text{if } d(y) > d(x); \end{cases}$$

where $\{x, y\}$ denotes an unordered pair, and (x, y) denotes an ordered pair of vertices.

Essentially the graph $\Gamma_F = \big(V, F(E)\big)$ has the same set of vertices as Γ, however now we allow some or all edges to be directed.

For any vertex $x \in V$, we can define a *principal vector field* F_x such that for all pair or vertexes $y, z \in E$ the following holds:

$$F_x(\{y,z\}) = \begin{cases} \{y,z\} & \text{if } d_x(y) = d_x(z); \\ (y,z) & \text{if } d_x(z) < d_x(y); \\ (z,y) & \text{if } d_x(z) > d_x(y). \end{cases}$$

In a principal vector field F_x all oriented edges point along geodesics towards the given vertex x. Essentially F_x depicts the distance function d_x from all vertices to x. For simplicity we refer to $\Gamma_F = \big(V, F_x(E)\big)$ as the principal vector field F_x, as it is rather simple to depict vectors on a graph than to formally express all relations.

Let's consider a simple example to illustrate how the principal vector field are constructed.

Example 5.1.3. Let $\Gamma = (V, E)$ be the following graph with the a fixed vertex x:

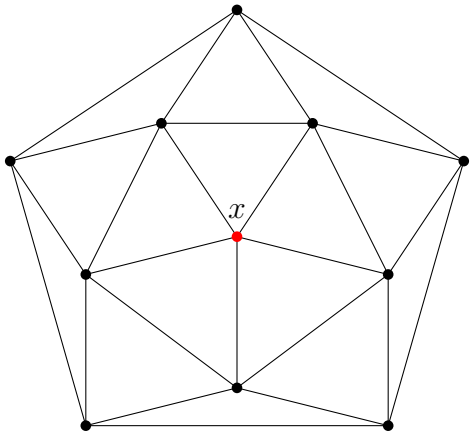

Then principal vector field F_x corresponding to the vertex x, is the following:

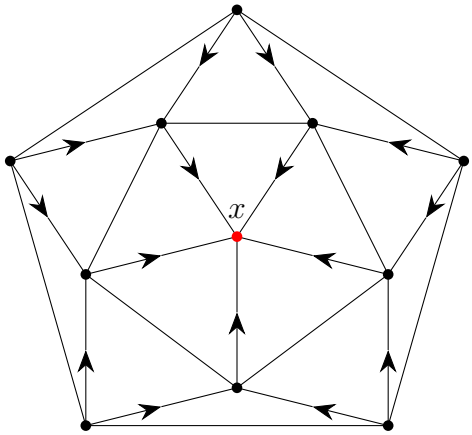

Definition 5.1.4. Fix $B \subseteq E$ a subset of edges. We say $x \underset{B}{\sim} y$ if $F_x|_B = F_y|_B$. Note $\underset{B}{\sim}$ is an equivalence relation. We denote this equivalence class by $[x]_B$.

A set $B_n = \{x \in V : d_r(x) \leq n\}$ is called a *ball*, where r denotes the root of the graph. The corresponding equivalence class is denoted by $\mathcal{A}(B_n)$ and is called an *atom* [5, Definition 3.1].

Example 5.1.5. Let $\Gamma = (V, E, r)$ be a locally finite connected rooted graph and $B_0 \subseteq E$. Then its 0-level atom $\mathcal{A}(B_0)$ is the whole graph V itself.

We call a graph $\Gamma = (V, E)$ a *tree* whenever there exists a unique geodesic for every pair of vertices in V, consecutively $\Gamma = (V, E, r)$ is called a *rooted tree* when this condition is met.

A *geodesic ray* is an infinite sequence of edges in which every finite subsequence of edges acts as a geodesic for the corresponding pair of vertices, in other words, it is an infinite sequence of geodesics $\pi(x_0, x_1), \pi(x_1, x_2), \ldots$ such that for any $n, m \geq 0$ the sequence of geodesics

$$\pi(x_n, x_{n+1}), \pi(x_{n+1}, x_{n+2}), \ldots, \pi(x_{m-2}, x_{m-1}), \pi(x_{m-1}x_m)$$

is a geodesic from x_n to x_m.

Let $\Gamma = (V, E, r)$ be a rooted tree, then the set of equivalence classes of infinite geodesic rays is called the *boundary* of a tree.

Definition 5.1.6. [5, Definition 3.4] For every $n \geq 0$, we denote $\mathcal{A}_n(V)$ as the set of *infinite atoms* in $\mathcal{A}(B_n)$, i.e, an *atom* is in $\mathcal{A}_n(V)$ if and only if its cardinality is infinite. The disjoint union

$$\mathcal{A}(V) = \bigsqcup_{n=0}^{\infty} \mathcal{A}_n(V)$$

is called the *tree of atoms* of V.

The tree of atoms $\mathcal{A}(V)$ is equipped with the following topologies:

- for each $n \geq 0$, $\mathcal{A}_n(V)$ has the discrete topology;
- the topology on $\mathcal{A}(V) = \bigsqcup_{n=0}^{\infty} \mathcal{A}_n(V)$ is the disjoint union topology;
- the boundary of the tree of atoms $\partial\mathcal{A}(V)$ inherits the subspace topology of the product topology on $\mathcal{A}(V)^{\mathbb{N}}$, where elements of the boundary are identified with infinite descending paths in the tree.

Note that, since $\mathcal{A}(V)$ is as a disjoint union, each element in $\mathcal{A}(V)$ can be represented as an ordered pair (n, A), where $n \geq 0$ and $A \in \mathcal{A}_n(V)$. This distinction is important, as it is possible for the same atom A to be in $\mathcal{A}_n(V)$ for multiple values of n. Observe that in the tree of atoms there exists an edge from (n, A) to $(n+1, A')$ if and only if $A' \subseteq A$.

Definition 5.1.7. [5, Definition 3.7] A *morphism* from $[x]_{B_m}$ to $[y]_{B_n}$ is a bijection $\phi : [x]_{B_m} \to [y]_{B_n}$ such that:

1. if $u, v \in [x]_{B_m}$ and $p \in \{L, R\}^*$, then $u = pv$ if and only if $\phi(u) = p\phi(v)$;
2. for all $z \in [x]_{B_m}$ and all $k \geq 0$, $\phi([z]_{B_{m+k}}) = [\phi(z)]_{B_{n+k}}$.

This definition of morphism is weaker modification than the one presented in [5], as we have to use arbitrary graph isomorphisms since there aren't group elements. Last comment was here.

The corresponding morphism is a pair of isomorphisms $\phi_1 : [x]_{B_m} \to [y]_{B_n}$ and $\phi_2 : [y]_{B_n} \to [x]_{B_n}$. We say that $[x]_{B_m} \underset{T}{\sim} [y]_{B_n}$ have the *same equivalence type* if such a morphism exists. We will briefly check that $\underset{T}{\sim}$ is an equivalence relation.

Proof. We will first check reflexivity. Let $[x]_{B_m} \underset{T}{\sim} [x]_{B_m}$, we choose $\phi : [x]_{B_m} \to [x]_{B_m}$ to be the identity map:

- let $u, v \in [x]_{B_m}$ and $p \in \{L, R\}^*$, then $u = pv \iff \phi(u) = u = pv = p\phi(v)$;
- for all $z \in [x]_{B_m}$ and $k \geq 0$, $\phi([z]_{B_{m+k}}) = [z]_{B_{m+k}} = [\phi(z)]_{B_{m+k}}$.

Hence, $\underset{T}{\sim}$ is reflexive.

We will now check symmetry. Let $[x]_{B_m} \underset{T}{\sim} [y]_{B_n}$, then $\phi : [x]_{B_m} \to [y]_{B_n}$ is a morphism, we will show that $\phi^{-1} : [y]_{B_n} \to [x]_{B_m}$ is also a morphism:

- let $u, v \in [y]_{B_n}$ and $p \in E$, then $u = pv \iff \phi\big(\phi^{-1}(u)\big) = p\phi\big(\phi^{-1}(v)\big) \iff \phi^{-1}(u) = p\phi^{-1}(v)$;
- for all $z \in [y]_{B_n}$ and $k \geq 0$, $\phi^{-1}([z]B_{n+k}) = \phi^{-1}\Big([\phi\big(\phi^{-1}(z)\big)]_{B_{n+k}}\Big) = [\phi^{-1}(z)]_{B_{m+k}}$.

Hence, is $\underset{T}{\sim}$ is symmetric.

It is left to check transitivity. Let $[x]_{B_m} \underset{T}{\sim} [y]_{B_n}$ and $[y]_{B_n} \underset{T}{\sim} [z]_{B_r}$, then $\phi : [x]_{B_m} \to [y]_{B_n}$ and $\psi : [y]_{B_n} \to [z]_{B_r}$ are morphisms. We will now show that $\psi \cdot \phi : [x]_{B_m} \to [z]_{B_r}$ is also a morphism:

- let $u, v \in [x]_{B_m}$ and $p \in \{L, R\}^*$, then $u = pv \iff \phi(u) = p\phi(v) \iff \psi\big(\phi(u)\big) = p\psi\big(\phi(v)\big) \iff (\psi \cdot \phi)(u) = p(\psi \cdot \phi)(v)$;
- for all $z \in [x]_{B_m}$ and $k \geq 0$, $(\psi \cdot \phi)([z]_{B_{m+k}}) = \psi\big(\phi([z]_{B_{m+k}})\big) = \psi\big([\phi(z)]_{B_{n+k}}\big) = [\psi\big(\phi(z)\big)]_{B_{r+k}} = [(\psi \cdot \phi)(z)]_{B_{r+k}}$.

Hence, $\underset{T}{\sim}$ is transitive. Since all there conditions are met, we conclude that $\underset{T}{\sim}$ is indeed an equivalence relation. □

Theorem 5.1.8. *[5, Theorem 3.6] The boundary of the tree of atoms $\mathcal{A}(V)$ is homeomorphic to the horofunction boundary $\partial_h V$ of V.*

This theorem plays a key role in this section. It not only provides the necessary tools to find the horofunction boundary using a combinatorial approach but also allows us to visualize the horofunction boundary.

A finite directed *multigraph* is a graph $\Gamma = (V, E)$, satisfying the conditions of a finite directed graph with the only difference of allowing to have multiple edges between the same pair of vertices. The corresponding *path language* $\mathcal{L}(\gamma, r) \subseteq E^*$ is the set of all finite paths $e_1 e_2 \cdots e_n \in \Gamma$ that begin with a fixed vertex r, since the graph is directed we do not allow a path a vertex v to move along edges that enter v. The set $\mathcal{L}(\gamma, r)$ has a natural structure of a locally finite tree, where the root acts as the empty path ϵ, and the boundary $\partial\mathcal{L}(\gamma, r)$ is the set of all infinite paths in Γ starting from the root r [5, Section 2.1].

Let's see a short example, which demonstrates the intuition behind the path language.

Example 5.1.9. Figure 5.1 demonstrates such a tree.

A *self-similar* structure [5, Section 2.1] on $\mathcal{L}(\Gamma, r)$ can be defined by the following statements:

- if two paths $p, q \in \mathcal{L}(\Gamma, r)$ end up in the same vertex of Γ, then we say that they have the same type;

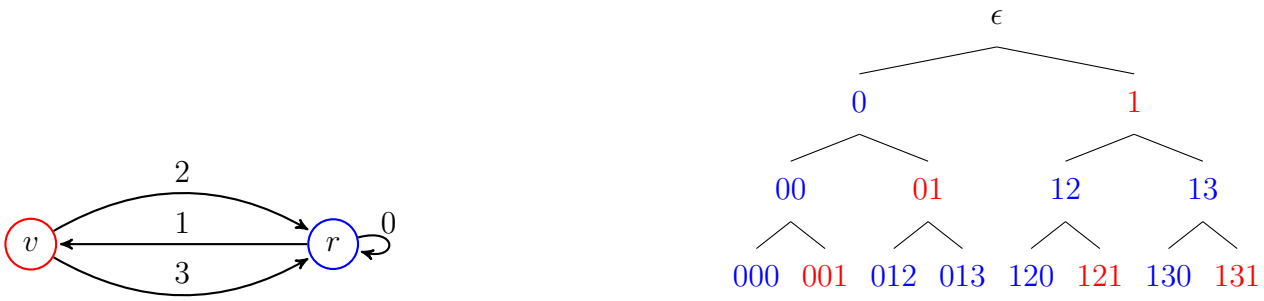

Figure 5.1: A directed multigraph Γ and its corresponding path language tree $\mathcal{L}(\Gamma, r)$

- let two paths $p, q \in \mathcal{L}(\Gamma, r)$ end up in the same vertex v, then for every finite directed path r staring from v, we define single morphism $\gamma_{pq} : \mathcal{L}(\Gamma, r)_p \to \mathcal{L}(\Gamma, r)_q$ by

$$\gamma_{pq}(pr) = qr.$$

  call γ_{pq} a *prefix replacement morphism.*

According to [5, Proposition 2.21], every self-similar tree T is isomorphic to a path language. We define define the *type graph* Γ of a self-similar tree T as follows:

- for every vertex type in T there is exactly one vertex in Γ;
- for every pair of vertices $x, y \in T$, the number of edges in Γ between x and y corresponds to the number of children of type y that every vertex of type x has in T.

By [5, Proposition 2.21], the tree of atoms $\mathcal{A}(V)$ is isomorphic to the set all finite directed path starting form the root. Hence, the horofunction boundary $\partial_h V$ is naturally homeomorphic to the space of all infinite paths that start from the root in its corresponding type graph.

5.2 The relationship between cones and distances in the graph $\mathcal{M}$

This section builds up the required tools to study atoms of $\mathcal{M}$. In particular we will introduce the connection between the cones of $\mathcal{M}$ and the distance functions. Although intuitively the structure of the cones of $\mathcal{M}$ seem complicated, by the end of the section we will see that the structure is understood through 4 statements.

Lemma 5.2.1. *Let $x \in M$, then the following hold:*

1. $x \in Cone(L) \Leftrightarrow d(x, L) < d(x, 1)$;
2. $x \in Cone(R) \Leftrightarrow d(x, R) < d(x, 1)$;
3. $x \notin Cone(L) \Leftrightarrow d(x, L) > d(x, 1)$;

4. $x \notin Cone(R) \Leftrightarrow d(x,R) > d(x,1)$;
5. $x \in Cone(LR^2) \Leftrightarrow d(x,R), d(x,L) < d(x,1)$;
6. $x \in \big(Cone(LR) \cup Cone(RL)\big) \Leftrightarrow d(x,LR) < d(x,L) \Leftrightarrow d(x,RL) < d(x,R)$.

Proof. We begin by proving the first equivalent statement.

1. We will start from the forward direction. Let $x \in Cone(L)$. By definition of the graph $\mathcal{M}$, $d(x,1) = |x|$. Since $x \in Cone(L)$, there exists $x' \in M$ such that $x = Lx'$. For $d(x,L)$, we apply Proposition 4.3.3 that states that every cone of $\mathcal{M}$ is convex and Theorem 4.3.1 that tells us that $d(x,L) = d(x',1) = |x'|$. Note that $|x'| < |x|$ because $x = Lx'$. Hence, $x \in Cone(L) \Rightarrow d(x,L) < d(x,1)$.

   For the opposite direction let $x \in M$ such that $d(x,L) < d(x,1)$. Then $d(x,1) = |x|$, which implies that $d(x,L) \leq |x| - 1$. By Theorem 4.3.1 this possible only when $x \in Cone(L)$.

2. The second equivalent statement is proven analogously to the first by interchanging the roles of L and R.

3. We will show that for all $x \in M$ $d(x,1) \neq d(x,L)$. Observe that $d(x,1) = |x|$. Now by Theorem 4.3.1 $d(x,L)$ is either $|x| - 1$ or $|x| + 1$ depending on x.

   Hence the statement $x \notin Cone(L) \Leftrightarrow d(x,L) > d(x,1)$ is the contrapositive of the first equivalence, and thus follows directly from it.

4. The fourth equivalent statement is a symmetric case of the third one.

5. The fifth equivalent statement is a direct consequence of the first two.

   For the forward direction, let $x \in M$ such that $x \in Cone(LR^2)$. Then $x \in Cone(LR^2) \subseteq Cone(L)$ and $x \in Cone(LR^2) = Cone(RL^2) \subseteq Cone(R)$. This implies that $x \in Cone(L)$ and $x \in Cone(R)$. Now we apply equivalences 1 and 2 to conclude that $d(x,L) < d(x,1)$ and $d(x,R) < d(x,1)$.

   For the opposite direction, let $x \in M$ such that $d(x,L) < d(x,1)$ and $d(x,R) < d(x,1)$. Then by equivalences 1 and 2, $x \in Cone(L)$ and $x \in Cone(R)$. Hence, $x \in Cone(L) \cap Cone(R) = Cone(LR^2)$, thus proving the equivalence.

6. The last equivalent statement follows from Theorem 4.3.1. We will prove $x \in \big(Cone(LR) \cup Cone(RL)\big) \Leftrightarrow d(x,LR) < d(x,L)$. As the last equivalence is an analogous case of the second one.

   We will start with the forward direction. Let $x \in \big(Cone(LR) \cup Cone(RL)\big)$. We have two cases corresponding to the first symbol of x. If x begins with L, we are trivially done. Let x start with R, then by Theorem 4.3.1 $d(x,LR) = |x| - 2$, whereas $d(x,L) = |x| - 1$.

   We will now prove the opposite direction. Let $d(x,LR) < d(x,L)$, then following the proof of Proposition 4.3.6, we observe that the set of points for which this inequality holds is precisely $Cone(LR) \cup Cone(RL)$.

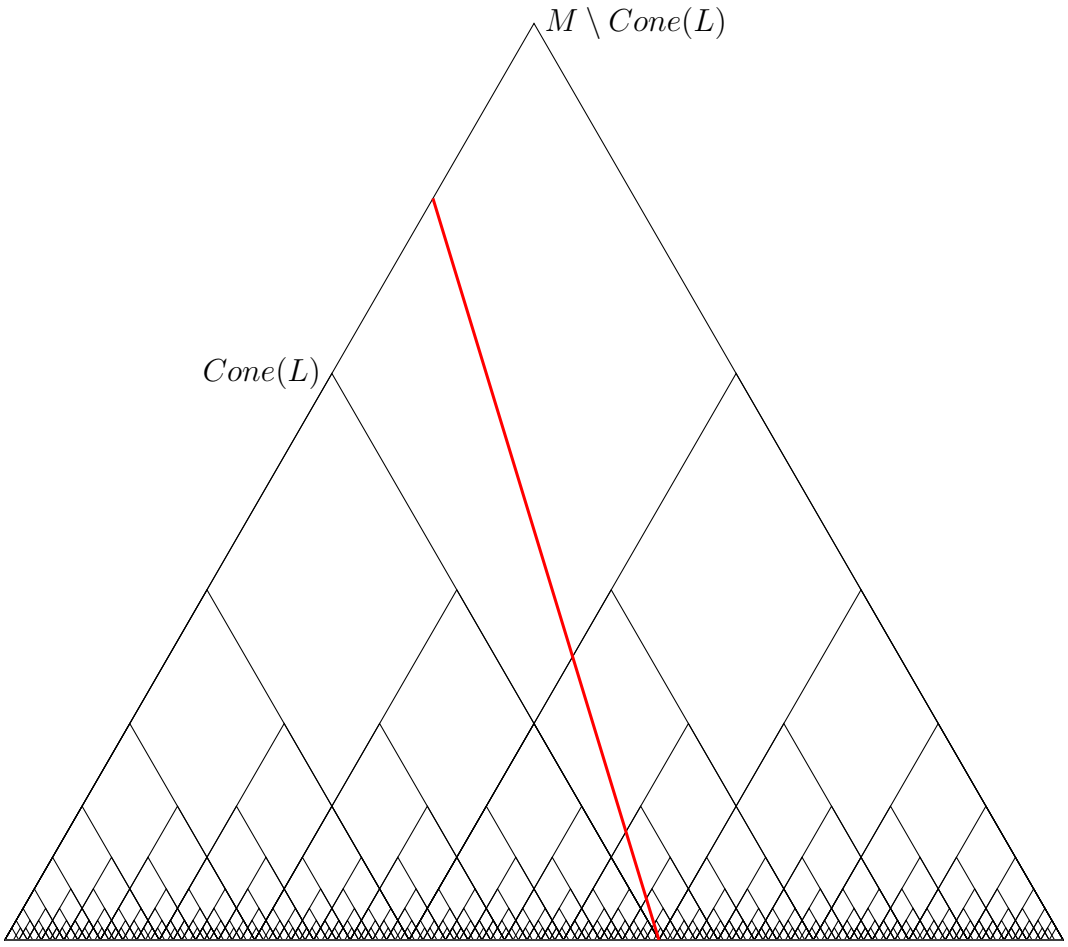

Figure 5.2: A hyperedge in red, splitting the set of vertices

□

Definition 5.2.2. A *hyperedge* is a line that splits a space into 2 disjoint *half-planes*.

In the context of the graph $\mathcal{M}$, where the space is the set of vertices M, a hyperedge selects a specific subset of vertices. To visualize this concept, see to Figure 5.2, where the hyperedge, depicted in red, divides the set of vertices into two parts:

- the left half satisfies the first equivalent statement of Lemma 5.2.1;
- the right half satisfies the negation of this statement, i.e., the third statement of Lemma 5.2.1.

The orientation of a vector field from L to 1 indicates the direction towards a half-plane where the vertices are closer to 1 than they are to vertex L. Note, every hyperedge crosses at least 2 edges. Whenever a vector field is imposed on one edge that is crossed by a hyperedge it automatically predetermines the orientations of all other edges that this hyperedge crosses. See Figure 5.3 for a visualisation of a vector-field imposed on a hyperedge that satisfy statement 2 of Lemma 5.2.1.

Lemma 5.2.3. *Let $x, m \in M$. Then $x \in \big(Cone(mLR) \cup Cone(mRL)\big) \Leftrightarrow d(x, mLR) < d(x, mL) \Leftrightarrow d(x, mRL) < d(x, mR)$. See Figure 5.4 for visualisation.*

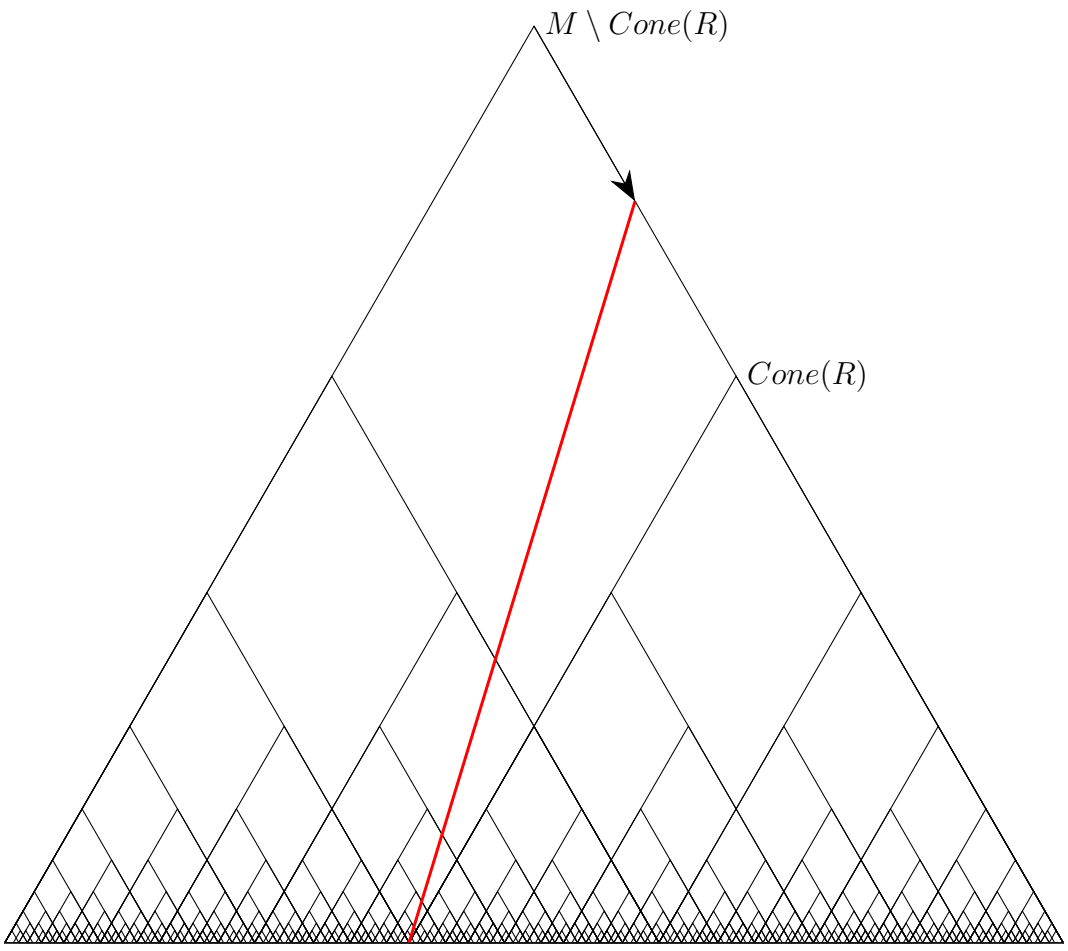

Figure 5.3: A vector field pointing towards the half-plane, whose elements are closer to R than to 1

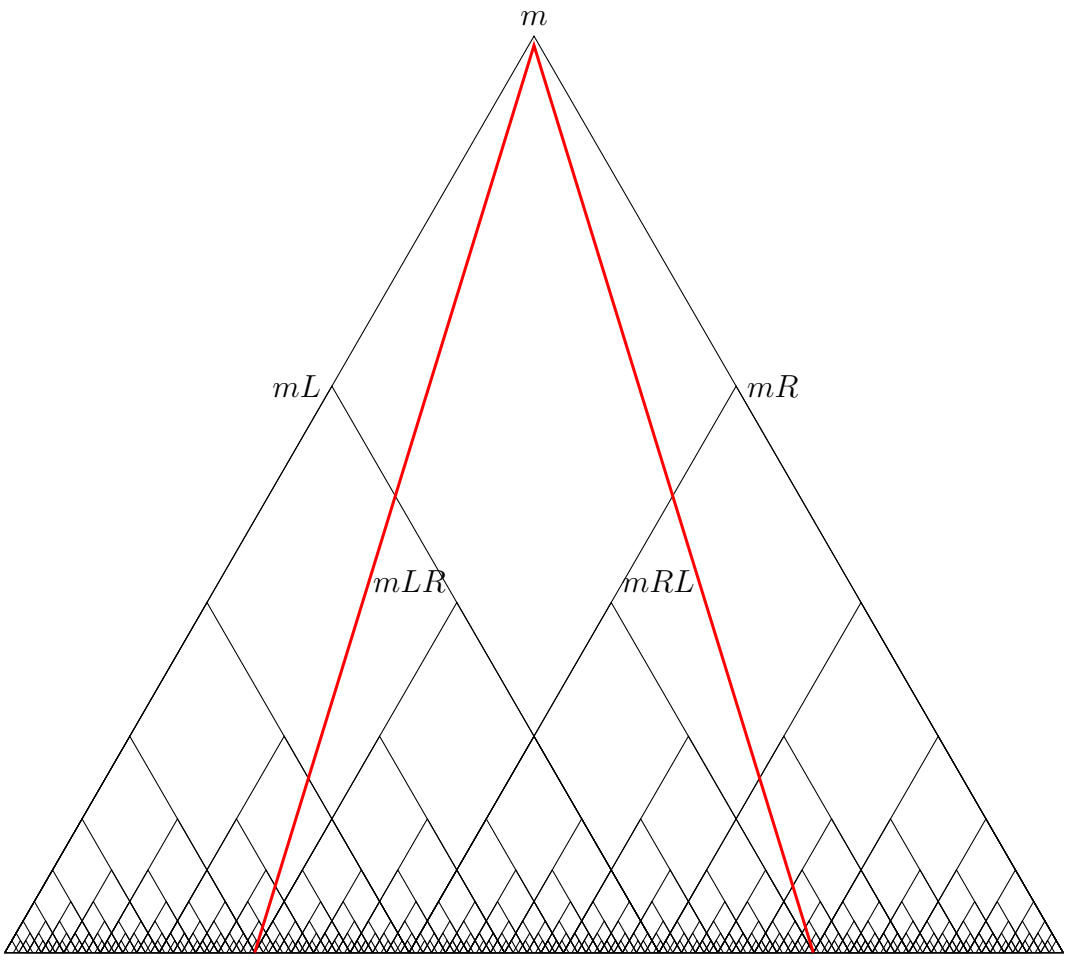

Figure 5.4: A hyperedge showing $Cone(mLR) \cup Cone(mRL)$

Proof. We will start from the forward direction. Let $x \in \big(Cone(mLR) \cup Cone(mRL)\big)$, then by Proposition 4.3.3 we know that $d(x, mLR) = d(x', LR)$ and $d(x, mL) = d(x', L)$, where $x = mx'$. Now by Theorem 4.3.1 $d(x', L) < d(x', LR)$. An analogous argument is applied to prove $d(x', R) < d(x', RL)$.

Now let $d(x, mLR) < d(x, mL)$, now assume that there exist $x \notin Cone(m)$ that satisfies this condition. Then there exists a geodesic from mL to x that passes through mR but by Proposition 4.3.6 this is not possible as a geodesic first takes the ascending path until it reaches the root n, nLR or nRL of the smallest $Cone(n)$ containing both mL and x.

Hence, the condition $d(x, mLR) < d(x, mL)$ only makes sense when $x \in Cone(m)$. Now we remove the common prefix m and apply part 6 of Lemma 5.2.1. Hence, $d(x, mLR) < d(x, mL)$ implies $x \in \big(Cone(mLR) \cup Cone(mRL)\big)$.

To prove that $d(x, mRL) < d(x, mR)$ implies $x \in \big(Cone(mLR) \cup Cone(mRL)\big)$, we apply an analogous argument by interchanging the roles of L and R.

Hence, the statement $x \in \big(Cone(mLR) \cup Cone(mRL)\big) \Leftrightarrow d(x, mLR) < d(x, mL) \Leftrightarrow d(x, mRL) < d(x, mR)$ is indeed valid. □

Corollary 5.2.4. *Let $x, m \in M$, such that there exists $m' \in M$ that satisfies $m = m'L$ then $d(x, mR) < d(x, m) \Leftrightarrow x \in \big(Cone(mR) \cup Cone(m'RL)\big)$.*

Proof. The proof requires studying two cases, when m' exists and not. When $m' \notin M$, the conclusion follows trivially.

Let's look at the first case when $m' \in M$. We apply Lemma 5.2.3 to $Cone(m')$, see Figure 5.5. Hence, $d(x, mR) < d(x, m) \Leftrightarrow x \in \big(Cone(mR) \cup Cone(m'RL)\big)$ follows immediately. □

Corollary 5.2.5. *Let $x, m \in M$, then $d(x, mLR^2) < d(x, mLR) \Leftrightarrow d(x, mR) < d(x, m)$.*

Proof. This is another consequence of Theorem 4.3.1. Observe that the element mR satisfies the relation $d(x, mLR^2) < d(x, mLR)$. This implies that every element of $Cone(mR)$ satisfies the inequality, which implies $d(x, mR) < d(x, m)$. □

These corollaries have non-trivial results that play key roles in understanding the atoms of $\mathcal{M}$. Observe that whenever $m = m'L$, $d(x, mR^2) < d(x, mR)$ triggers a recursive effect as $d(x, m'LR^2) < d(x, m'LR)$ implies $d(x, m'R) < d(x, m')$. This will continue to recurse until it hits the boundary of the graph or a state $d(x, nLR) < d(x, nL)$. This means that the set of points corresponding to the inequality $d(x, mR) < d(x, m)$ will be the greatest cone containing mR but not m, and the cone $m'RL$.

Bringing it together gives us the following equivalent statements.

Proposition 5.2.6. *Let $x, m \in M$, then there exists $n \in M$ such that $mL = nR^i$, where i is as large as possible. There exists $n' \in M$ such that $nL = n'R^j$, where $j = max\{0, 1\}$.*

Then the following equivalent statements are valid:

1. $d(x, mL) < d(x, m) \Leftrightarrow x \in \big(Cone(n) \cup Cone(n')\big)$;

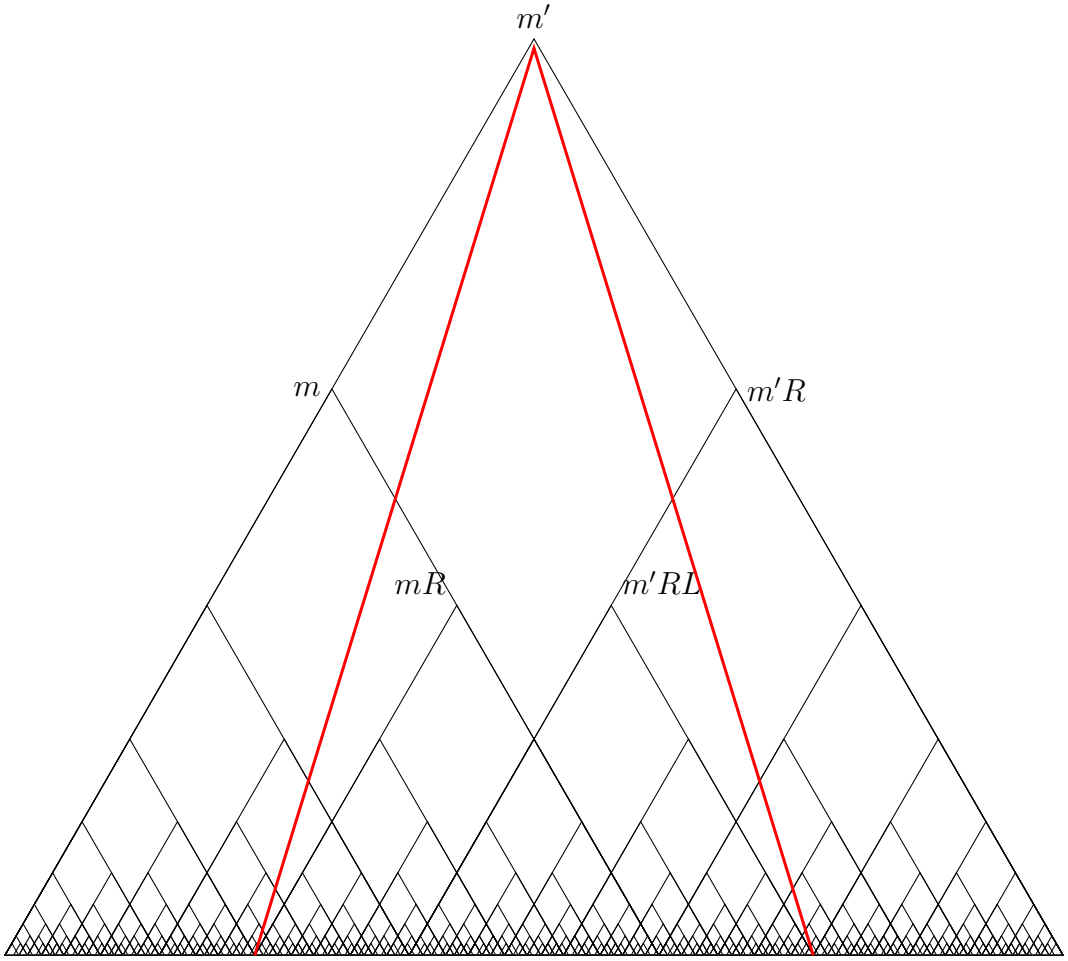

Figure 5.5: A hyperedge picking out elements that satisfy the inequality $d(x, m) < d(x, mR)$

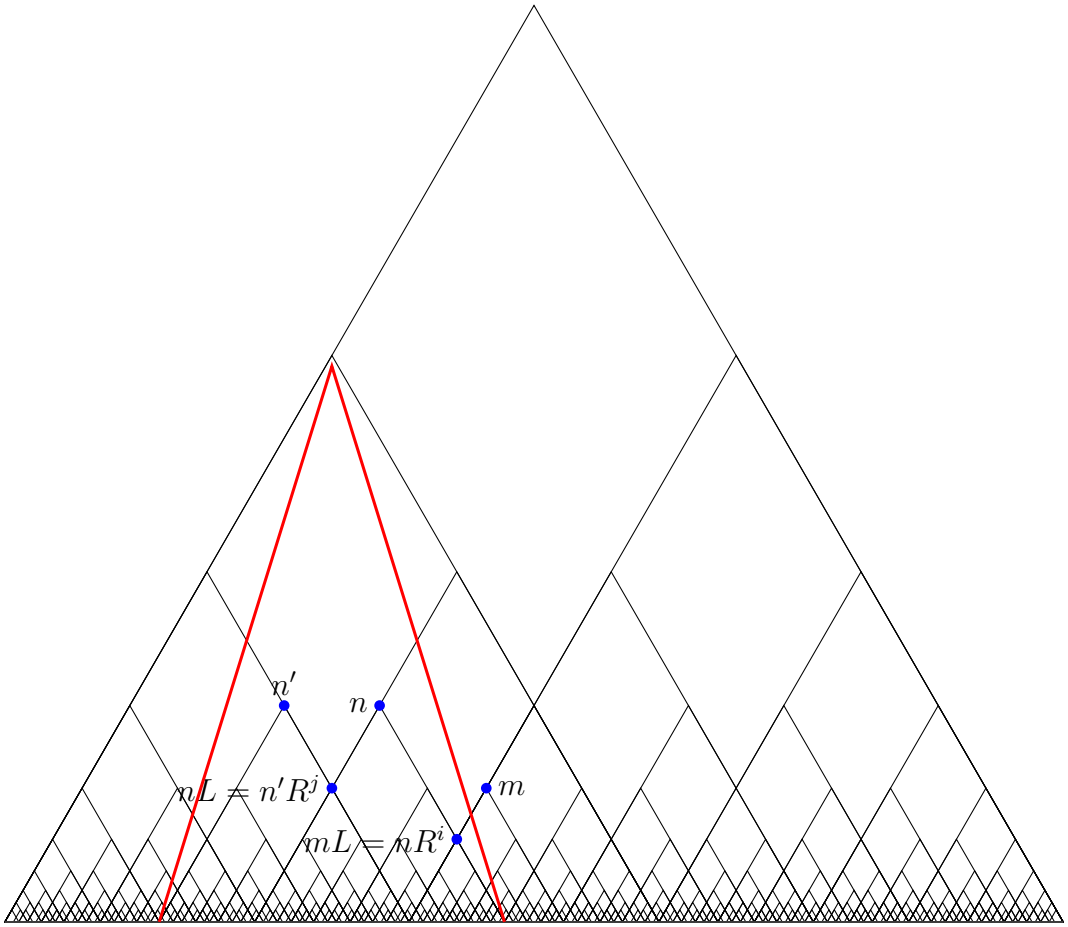

Figure 5.6: Visualisation of statements 1 and 2 of Proposition 5.2.6

2. $d(x, mL) > d(x, m) \Leftrightarrow x \notin \big(Cone(n) \cup Cone(n')\big)$.

By interchanging the roles of L and R we get the following.

Let $x, m \in M$, then there exists $n \in M$ such that $mR = nL^i$, where i is as large as possible. There exists $n' \in M$ such that $nR = n'L^j$, where $j = max\{0, 1\}$.

Then the following equivalent statements are valid:

3. $d(x, mR) < d(x, m) \Leftrightarrow x \in \big(Cone(n) \cup Cone(n')\big)$;

4. $d(x, mR) > d(x, m) \Leftrightarrow x \notin \big(Cone(n) \cup Cone(n')\big)$.

Note that it is possible for $Cone(n') \subset Cone(n)$. See Figure 5.6 for a visualisation of this process. Although this proposition alone may not appear immediately insightful, it has promising consequences. If x is an element in $Cone(m)$ then knowing the signs of the inequalities $d(x, m) \neq d(x, mL)$ and $d(x, m) \neq d(x, mR)$, tells all the distances between every element $n \notin Cone(m)$ and its offsprings. In other words for every $x \in Cone(m)$, the distance vectors $\{m, mL\}$ and $\{m, mR\}$ determine the distance vectors outside $Cone(m)$.

We end this section with a theorem that is a final consequence of these propositions. For any $m \in M$ we call set the $Cone(pLR) \cup Cone(pRL)$ a *triangle* in $\mathcal{M}$. We call the sets $Cone(L^n)$ and $Cone(R^n)$ *side cones* for any $n > 0$.

Theorem 5.2.7. *Let S be a triangle or a side cone in $\mathcal{M}$. Let $p, q \in M$ be adjacent vertices with $p \in S$ and $q \in S^c$. Then the following statements hold:*

1. *if $x \in S$ then $d(x,p) < d(x,q)$;*
2. *if $x \in S^c$ then $d(x,q) < d(x,p)$.*

This theorem implies that for every edge in the Cayley graph $\mathcal{M}$ there is exactly one hyperedge that crosses it. As a consequence we say that two atoms $\mathcal{A}$ and $\mathcal{B}$ are of the same type if there exist $Cone(m) \supseteq \mathcal{A}$ and $Cone(n) \supseteq \mathcal{B}$ such that the following hold:

- for every vertex $a \in \mathcal{A}$ there exists a vertex $b \in \mathcal{B}$ such that $m^{-1}a = n^{-1}b$, i.e., the atoms have same shapes;
- $F_{\mathcal{A}}|_{Cone(m)} = F_{\mathcal{B}}|_{Cone(n)}$, i.e., the cones that contain $\mathcal{A}$ and $\mathcal{B}$ have the same vector fields.

5.3 The horofunction boundary of M

This section presents one of the main results of this thesis, in particular demonstrating the connection between the horofunction boundary $\partial_h M$ of the Cayley graph $\mathcal{M}$ of the monoid M and the small golden ratio $\tau = \frac{\sqrt{5}-1}{2}$. We start by introducing the space $\mathcal{D}_\tau$, followed by a number of lemmas building up the connection between the horofunction boundary $\partial_h M$ and $\mathcal{D}_\tau$.

Let $I_\tau = \mathbb{Z}[\tau] \cap (0,1)$, where $\mathbb{Z}[\tau] = \{a + b\tau : a, b \in \mathbb{Z}\}$ and $\tau = \frac{1+\sqrt{5}}{2} \approx 0.618034$ is the small golden ratio. Recall that by Definition 2.3.10 a blowup of $[0,1]$ along I_τ is a Cantor set denoted by $\mathcal{C}_\tau$.

Definition 5.3.1. Let $\mathcal{D}_\tau$ be the set

$$\mathcal{D}_\tau = ([0,1] \setminus I_\tau) \cup \{x^- : x \in I_\tau\} \cup \{x : x \in I_\tau\} \cup \{x^+ : x \in I_\tau\}$$

equipped with a linear order $\prec$ satisfying the following conditions:

- $\prec$ agrees with the standard order $<$ on $[0,1] \setminus I_\tau$;
- for all $x \in I_\tau$, $x^- \prec x \prec x^+$;
- if $x \in I_\tau$ and $y \in [0,1] \setminus I_\tau$, then $x^+ \prec y \iff x \prec y \iff x^- \prec y$;
- if $x, y \in I_\tau$ and $x < y$, then $x^- \prec x \prec x^+ \prec y^- \prec y \prec y^+$.

The topology on $\mathcal{D}_\tau$ is the order topology induced by $\prec$. The set $\mathcal{D}_\tau$ can be viewed as a "double-blowup" of $[0,1]$ along I_τ, in other words every point in I_τ is replaced by three points instead of two points.

While $\mathcal{D}_\tau$ is not itself a Cantor set, it contains $\mathcal{C}_\tau$ as a subset, which is homeomorphic to the Cantor set. The key difference is that $\mathcal{D}_\tau$ contains additional isolated points (the elements $x \in I_\tau$) between each pair of blowup points in $\mathcal{C}_\tau$.

Let $x_i \in I_\tau$ then the Cantor set $\mathcal{C}_\tau$ is depicted as:

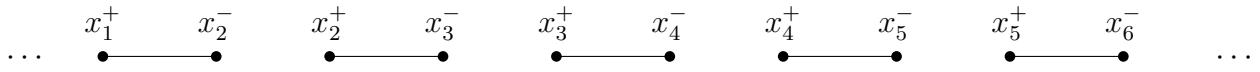

However $\mathcal{D}_\tau$ has additional isolated points in between every pair of non-dividable intervals. The following depicts the Cantor-like set $\mathcal{D}_\tau$.

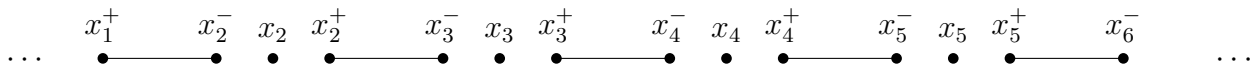

Recall that by Lemma 4.4.4 there is a one-to-one correspondence between the elements of the monoid M and subintervals of the unit interval I_M, where $x \in M$ can be represented by $[0,1]_x \in I_M$. The system of equations 4.1 represent the actions of L and R on the interval.

We claim that these equations must be slightly modified to be able to use them in terms of infinite paths of $\mathcal{M}$:

$$\begin{aligned} [x,y]_1 &:= [x,y]; \\ [x,y]_L &:= [x,(y-(y-x)\tau^2)^-]; \\ [x,y]_R &:= [(x+(y-x)\tau^2)^+,y]. \end{aligned} \tag{5.1}$$

The logic behind this step is motivated by first 2-level decompositions of atoms of $\mathcal{M}$ presented in Propositions 5.3.3 and 5.3.4. In these propositions we will see that the same element $RLLL\cdots = L^2RRR\cdots = LRLRLR\cdots$ of the monoid, depending on the chosen path belongs to distinct atoms. This is not allowed by the definition of atoms. To solve this issue we introduce a double-blowup along elements of $\mathbb{Z}[\tau] \cap (0,1)$, resulting in the introduced modified equations above.

Theorem 5.3.2. *The horofunction boundary $\partial_h\mathcal{M}$ of $\mathcal{M}$ is naturally homeomorphic to $\mathcal{D}_\tau$.*

The rest of the section is dedicated to proving this theorem. To understand the structure of the horofunction boundary $\partial_h\mathcal{M}$, we will decompose infinite atoms in sequential infinite atom levels until no new infinite atom types appear.

Proposition 5.3.3. *Let $\mathcal{A}_1(\mathcal{M})$ be the set of 1-level infinite atoms of $\mathcal{M}$. Then $\mathcal{A}_1(\mathcal{M}) = a_1 \cup b_1 \cup c_1$, where*

$$\begin{aligned} a_1 &= Cone(L) \setminus Cone(LR^2); \\ b_1 &= Cone(LR^2); \\ c_1 &= Cone(R) \setminus Cone(LR^2). \end{aligned}$$

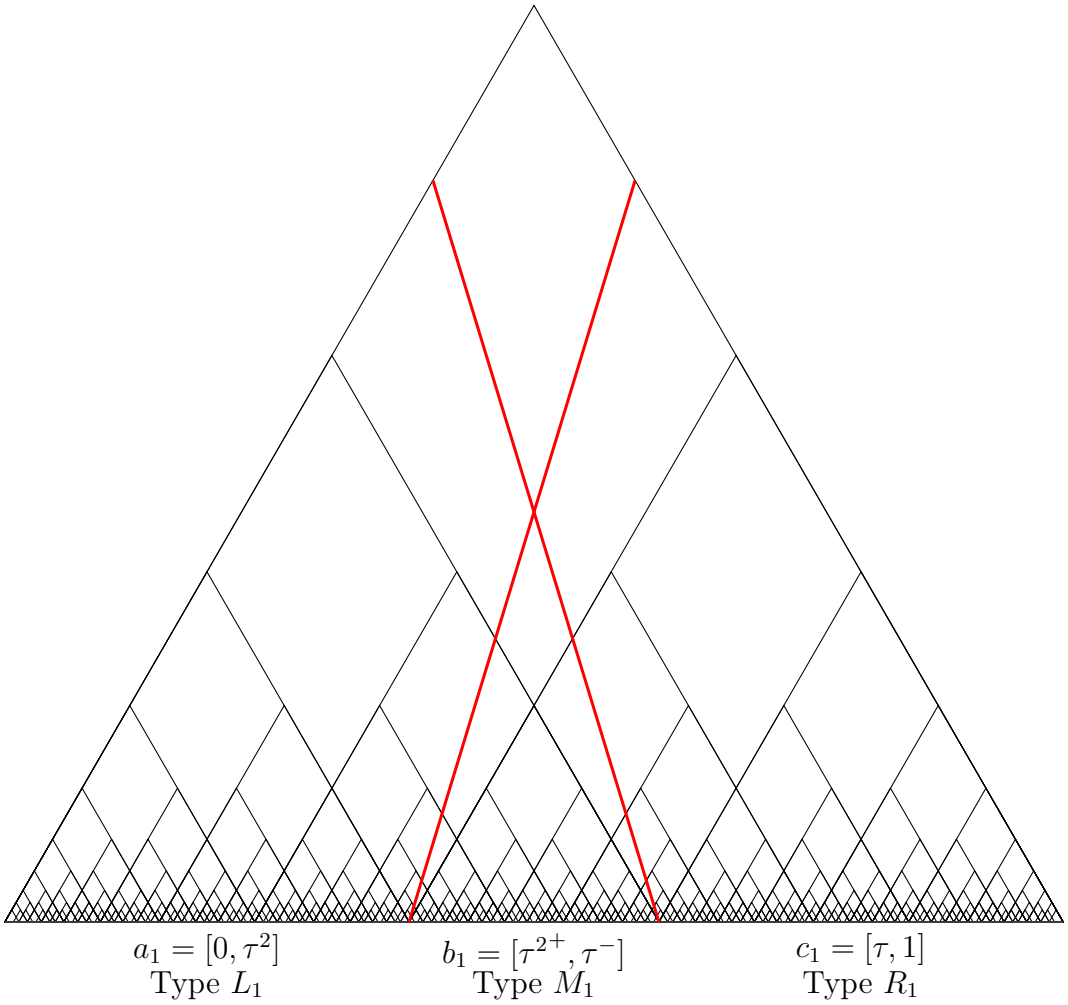

Figure 5.7: Partition of 1-level infinite atoms $\mathcal{A}_1(\mathcal{M})$ in the graph $\mathcal{M}$

In other words, the graph $\mathcal{M}$ is decomposed into 3 infinite sets by different vector fields of the unit ball centered at the root. See Figure 5.7 for a visualisation.

Proof. Let B_1 be a unit ball centered at the root of $\mathcal{M} = (M, E, 1)$. Observe that there is a total of 4 possible principal vectors in B_1. We will denote them as:

1. $F_1(\{1, L\}) = (L, 1)$ and $F_1(\{1, R\}) = (R, 1)$;
2. $F_L(\{1, L\}) = (1, L)$ and $F_L(\{1, R\}) = (R, 1)$;
3. $F_R(\{1, L\}) = (L, 1)$ and $F_R(\{1, R\}) = (1, R)$;
4. $F_{LR^2}(\{1, L\}) = (1, L)$ and $F_{LR^2}(\{1, R\}) = (1, R)$.

To ease the understanding of the complicated representations of the possible vector fields in B_1 we invite

the reader to take a look at their scaled graphic representation:

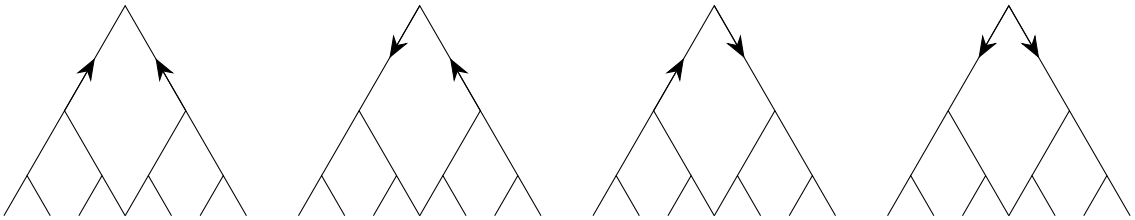

Recall that a vector $F(\{n,m\}) = (n,m)$ is equivalent to the statement $d(x,n) > d(x,m)$.

We will now analyse the four cases corresponding to the possible vector fields in B_1. We start from the first case.

Let $x \in M$ such that $d(x,L) > d(x,1)$ and $d(x,R) > d(x,1)$, following Lemma 5.2.1 we observe that $x \notin Cone(L)$ and $x \notin Cone(R)$. Hence, $x \in \{1\}$. This is a finite set, hence by definition it does not contribute to the horofunction boundary.

We move to the second case. Let $x \in M$ be such that $d(x,L) < d(x,1)$ and $d(x,R) > d(x,1)$. Following Lemma 5.2.1 we observe that $x \in Cone(L)$ and $x \notin Cone(R)$. Hence, $x \in Cone(L) \setminus Cone(LR^2)$. Let this atom be denoted as a_1. This is an infinite set, hence it contributes to the horofunction boundary.

The third case is similar to the second one. Let $x \in M$ such that $d(x,L) > d(x,1)$ and $d(x,R) < d(x,1)$, following Lemma 5.2.1 we observe that $x \notin Cone(L)$ and $x \in Cone(R)$. Hence, $x \in Cone(R) \setminus Cone(LR^2)$. Let this atom be denoted as c_1

We are left with the last case. Let $x \in M$ such that $d(x,L) < d(x,1)$ and $d(x,R) < d(x,1)$, following Lemma 5.2.1 we observe that $x \in Cone(L)$ and $x \in Cone(R)$. Hence, $x \in Cone(LR^2)$. Let this atom be denoted as b_1

The following is a scaled visualisation of these 4 sets, for a full size picture we refer the reader to Figure 5.7.

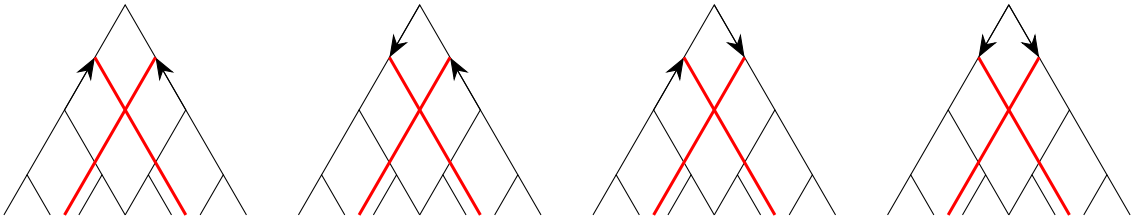

Let atoms a_1, b_1 and c_1 be of types L_1, M_1 and R_1 respectively. □

Following Lemma 4.4.4 we should be able to represent the atoms a_1, b_1 and c_1 in terms of the real line. However we immediately encounter a problem. The atom $a_1 = Cone(L) \setminus Cone(LR^2)$ corresponds to the interval $[0, \tau^2]$ and the atom $b_1 = Cone(LR^2)$ corresponds to the interval $[\tau^2, \tau]$. Observe that the point τ^2 belongs both to a_1 and b_1. This is due to the fact that τ^2 corresponds to multiple elements

of M in particular $L^2RRR\cdots$ and $RLLL\cdots$. When studying atoms we do not allow such constructions as infinite paths of atoms must lead to distinct elements. Hence, there is a blowup of τ^2 in such a way that $\tau^2 \neq \tau^{2'}$, $\tau^2 < \tau^{2'}$, and $\tau^2 \in a_1$, $\tau^{2'} \in b_1$.

However the statement above is not true when studying the atoms of $\mathcal{M}$. In order to not repeat ourselves we would like the reader to believe that the atom a_1 will decompose in such a way that the point τ^2 will blowup again. This will be seen in the next proposition.

Hence, the equations 5.1 can be used to represent atoms. We will now state the atoms a_1, b_1 and c_1 in terms of intervals.

We start from a_1. The condition $a_1 \subseteq Cone(L)$ is equivalent to $a_1 \subseteq [0,1]_L = [0,\tau^-]$ and $a_1 \not\subseteq [0,1]_R = [\tau^+,1]$, hence a_1 is the set $[0,\tau^2]$.

For the atom b_1 the condition $b_1 \subseteq Cone(LR^2)$ is equivalent to $b_1 = [0,1]_{LR^2} = [\tau^{2^+},\tau^-]$.

We invite the reader to verify that $c_1 = [\tau,1]$.

Proposition 5.3.4. *The atom a_1 decomposes into three disjoint infinite atoms in $\mathcal{A}_2(\mathcal{M})$, denoted as a_2, b_2 and c_2. These infinite atoms correspond to the following subsets of $\mathcal{M}$:*

$$\begin{aligned} a_2 &= Cone(L^2) \setminus Cone(L^2R^2);\\ b_2 &= Cone(L^2R^2);\\ c_2 &= Cone(LR) \setminus \big(Cone(LRL^2) \cup Cone(LR^2)\big). \end{aligned}$$

Moreover, $a_1 = \{L\} \cup a_2 \cup b_2 \cup c_2$. See Figure 5.8 for a visualisation.

Proof. Let's consider the a_1 atom, whose vector field in B_1 points towards $Cone(L)$, see case 2 of Proposition 5.3.3 for an illustration. Following Proposition 5.2.6 we know that the vectors in $Cone(L)$ determine the decomposition of the atoms in $Cone(L)$, therefore we ignore the vectors of $Cone(L) \setminus Cone(R)$ as they will be predetermined by the ones in $Cone(L)$. Hence, there will be a total of 4 cases:

1. $F_L(\{L,L^2\}) = (L^2,L)$ and $F_L(\{L,LR\}) = (LR,L)$;
2. $F_{L^2}(\{L,L^2\}) = (L,L^2)$ and $F_{L^2}(\{L,LR\}) = (LR,L)$;
3. $F_{LR}(\{L,L^2\}) = (L^2,L)$ and $F_{LR}(\{L,LR\}) = (L,LR)$;
4. $F_{L^2R^2}(\{L,L^2\}) = (L,L^2)$ and $F_{L^2R^2}(\{L,LR\}) = (L,LR)$.

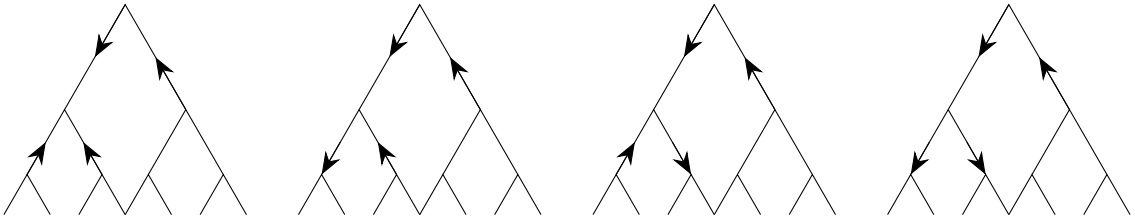

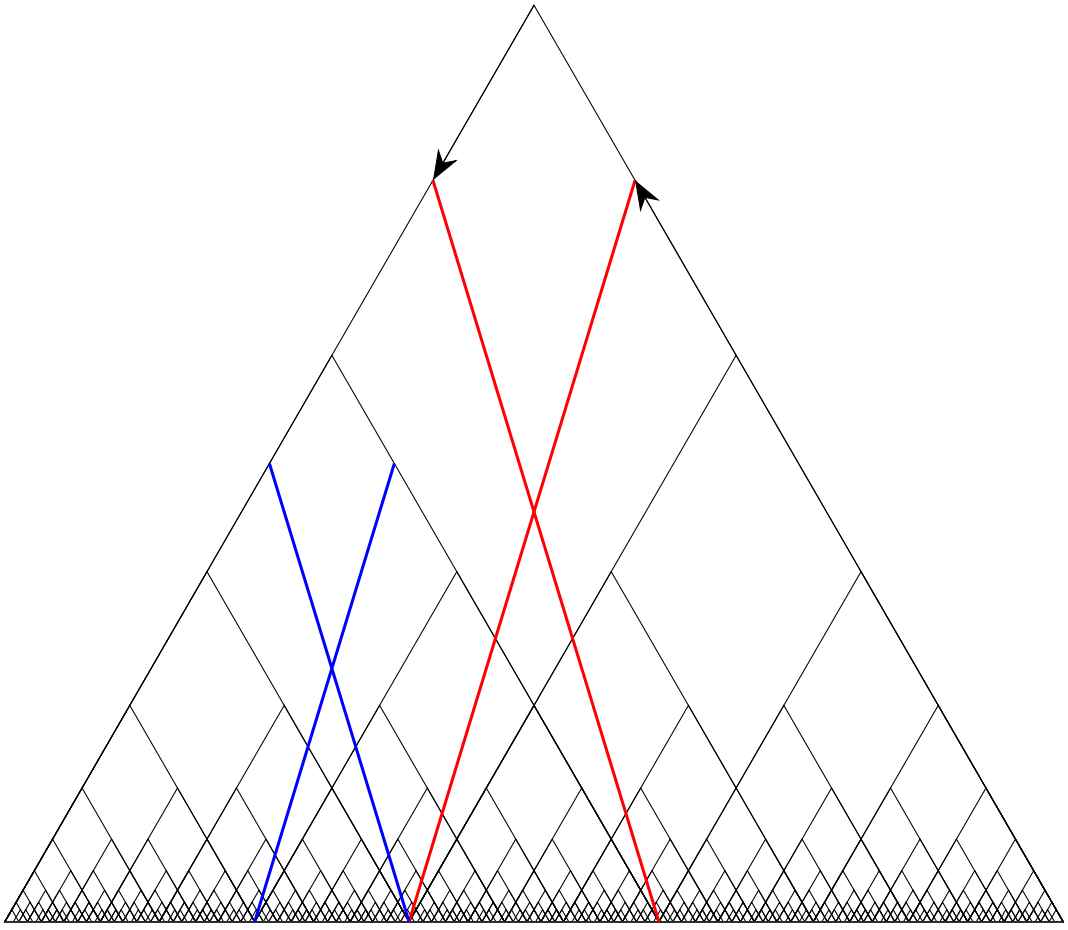

Figure 5.8: Decomposition of the atom a_1 into a_2, b_2, and c_2 in the second sequential level of the $\mathcal{A}_2(\mathcal{M})$

We leave it to the reader to check that the cases 1, 2 and 4 are analogous to the corresponding cases analysed in Propositions 5.3.3 subject to scaling.

However we will check the decomposition of the third case, as it has an nontrivial result. Let $x \in M$ such that $d(x, LR) < d(x, L) < d(x, 1) < d(x, R)$ and $d(x, LR) < d(x, L^2)$. By Lemma 5.2.1 this is equivalent to $x \notin Cone(R)$, $x \in Cone(L)$, $x \notin Cone(L^2)$ and $x \in Cone(LR)$. Hence, $x \in Cone(LR) \setminus \big(Cone(L^2R^2) \cup Cone(LR^2)\big)$. Let this atom be denoted as c_2.

Translating this in terms of intervals $c_2 \subseteq [0,1]_{LR} = [\tau^{3^+}, \tau^-]$, $b_2 \not\subseteq [0,1]_{L^2R^2} = [0, \tau^{2^-}]$ and $b_2 \not\subseteq [0,1]_{LR^2} = [\tau^{2^+}, \tau^-]$. Hence, $c_2 = \{\tau^2\}$ is a point. □

Observe that the pairs of atoms a_1, a_2 and b_1 and b_2 have the same shape and type. Let c_2 be an atom of type P_L.

Proposition 5.3.5. *The 1-level infinite atom b_1 does not decomposes in $\mathcal{A}_2(\mathcal{M})$.*

To prove this propositions one can apply the same method presented above and see that there is only one possibility for the 2-level atom vector field configuration and that it does not decompose b_1. We will refer to this infinite atom as d_2.

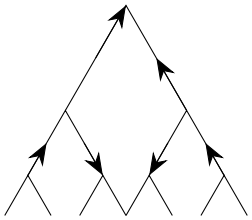

Note that b_1 and d_2 are atoms of different types. Let d_2 be of type M_2.

From now on we will omit the formal proofs of atom decomposition, as essentially they follow the same format presented above. However, we invite the reader to convince themselves that the stated propositions are valid.

Corollary 5.3.6. *The graph $\mathcal{M}$ contains 7 infinite atoms in $\mathcal{A}_2(\mathcal{M})$. See Figure 5.9 for visualisation.*

Proposition 5.3.7. *The 2-level infinite atom d_2 decomposes into three disjoint infinite atoms in $\mathcal{A}_3(\mathcal{M})$, denoted as f_3, g_3 and h_3. These infinite atoms correspond to the following subsets of $\mathcal{M}$:*

$$\begin{aligned} f_3 &= Cone(LR^2L^2); \\ g_3 &= Cone(L^2R^2) \setminus \{Cone(LR^2L^2) \cup Cone(LR^4)\}; \\ h_3 &= Cone(LR^4). \end{aligned}$$

Moreover, $d_2 = f_3 \cup g_3 \cup h_3$. See Figure 5.10 for visualisation.

The proof follows the same logic as for the previous decompositions.

Note that f_3 and h_3 are atoms of type M_1, this can be seen by looking at $Cone(LR)$ and $Cone(RL)$ and observing that the regions of atoms f_3, h_3 and b_1 within their respective cones are the same. Whereas g_3 is an infinite atom of a new type, let it be denoted by M_3.

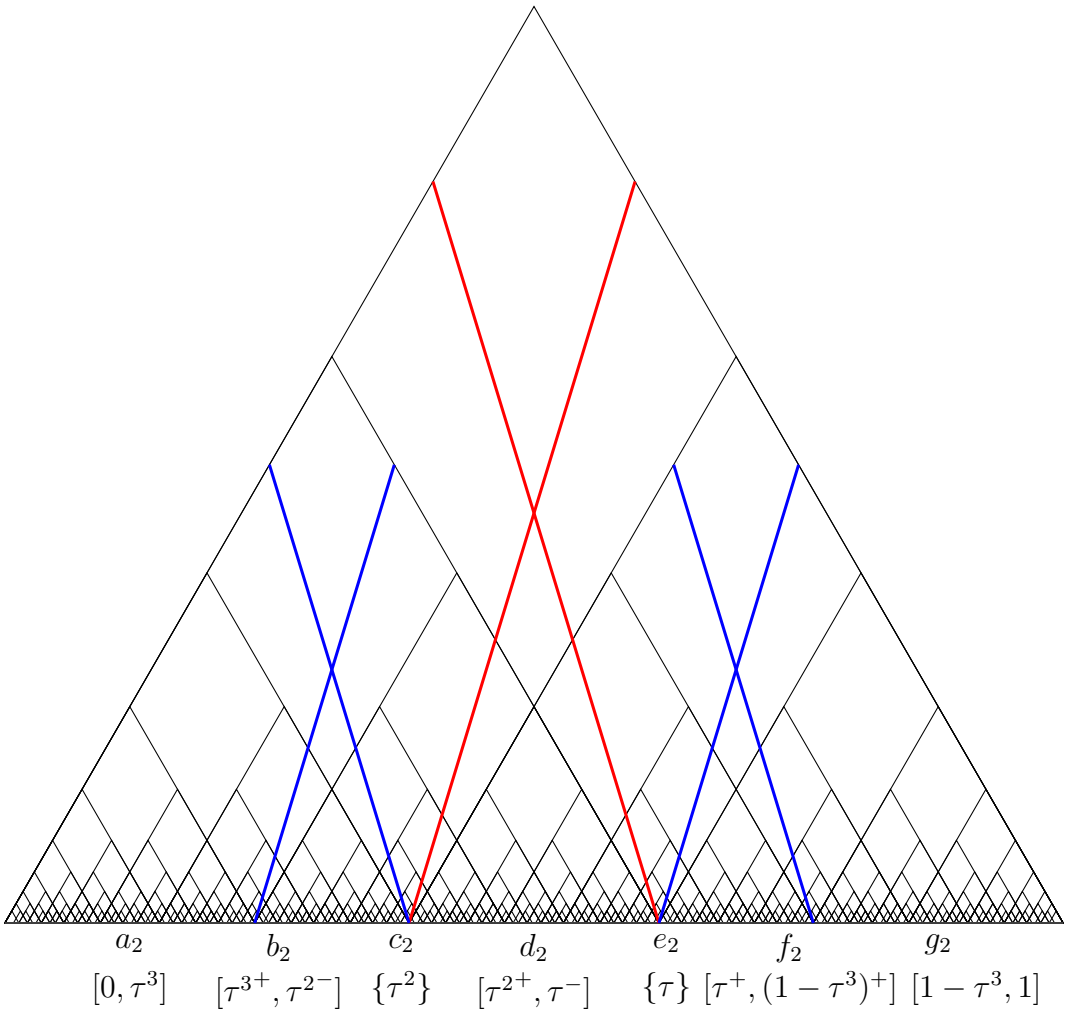

Figure 5.9: Partition of 2-level infinite atoms $\mathcal{A}_2(\mathcal{M})$ in the graph $\mathcal{M}$

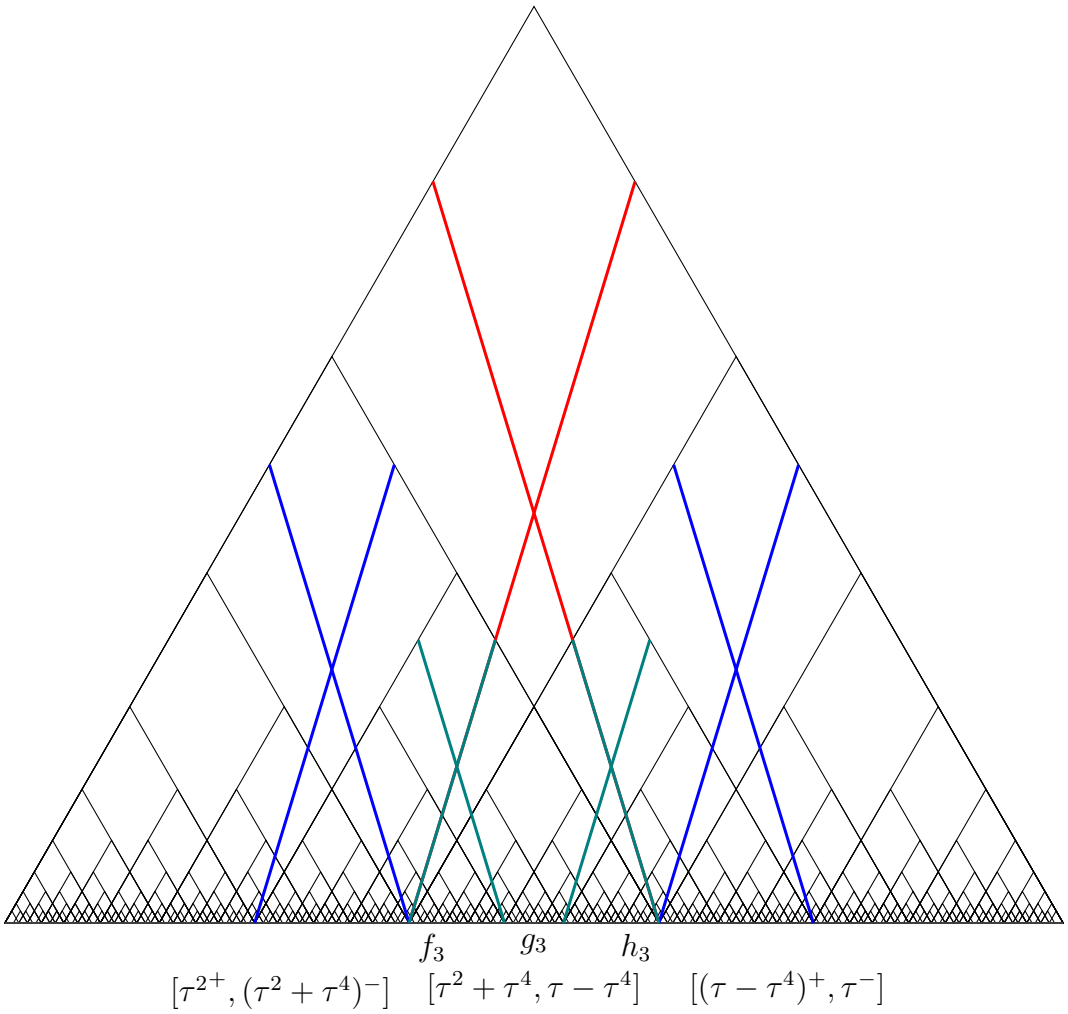

Figure 5.10: Decomposition of the atom d_2 into f_3, g_3, and h_3 in the second sequential level of the $\mathcal{A}_3(\mathcal{M})$

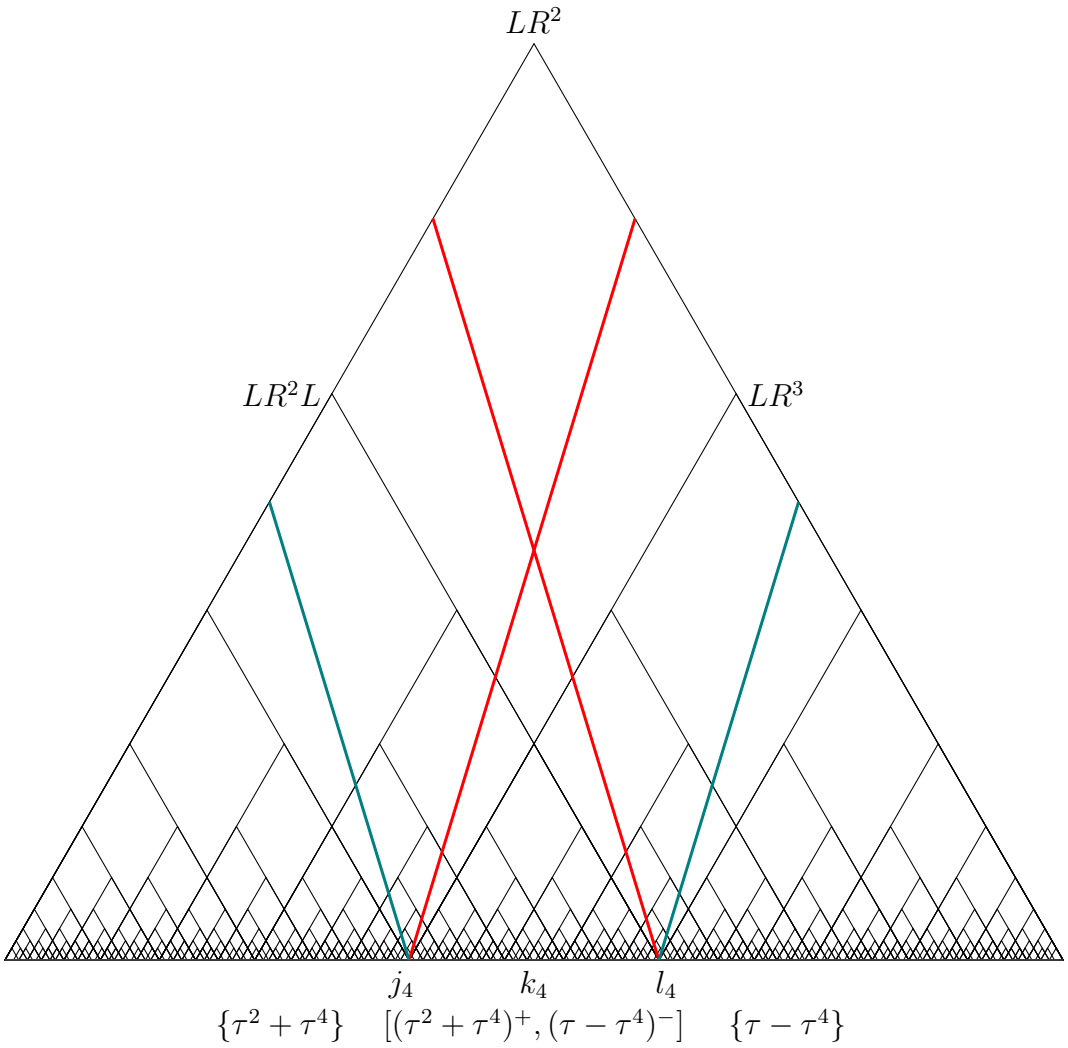

Figure 5.11: Representation of the decomposition of the atom f_3 into j_4, k_4, and l_4 in the second sequential level of the $\mathcal{A}_4(\mathcal{M})$

Proposition 5.3.8. *The 3-level infinite atom g_3 decomposes into three disjoint infinite atoms in $\mathcal{A}_4(\mathcal{M})$, denoted as j_4, k_4 and l_4. These infinite atoms correspond to the following subsets of $\mathcal{M}$:*

$$\begin{aligned} j_4 &= Cone(LR^2L) \setminus \big(Cone(LR^2L^2) \cup Cone(LR^2LR^2)\big);\\ k_4 &= Cone(LR^2LR^2);\\ l_4 &= Cone(RL^2R) \setminus \big(Cone(RL^2R^2) \cup Cone(RL^2RL^2)\big). \end{aligned}$$

Moreover, $f_3 = j_4 \cup k_4 \cup l_4 \cup \{LR^2\}$. See Figure 5.11 for visualisation.

The proof follows the same logic as before.

Note that j_4 is an atom of type P_R, k_4 is of type M_1 and l_4 is of type P_L. Since no new infinite atom types have appeared, we obtain the following corollary.

Corollary 5.3.9. *The tree of infinite atoms $\mathcal{A}(\mathcal{M})$ has only finite number of atom types. See Figure 5.12 for the type graph $\mathcal{T}$ and Figure 5.13 for the tree of infinite atoms $\mathcal{A}(\mathcal{M})$.*

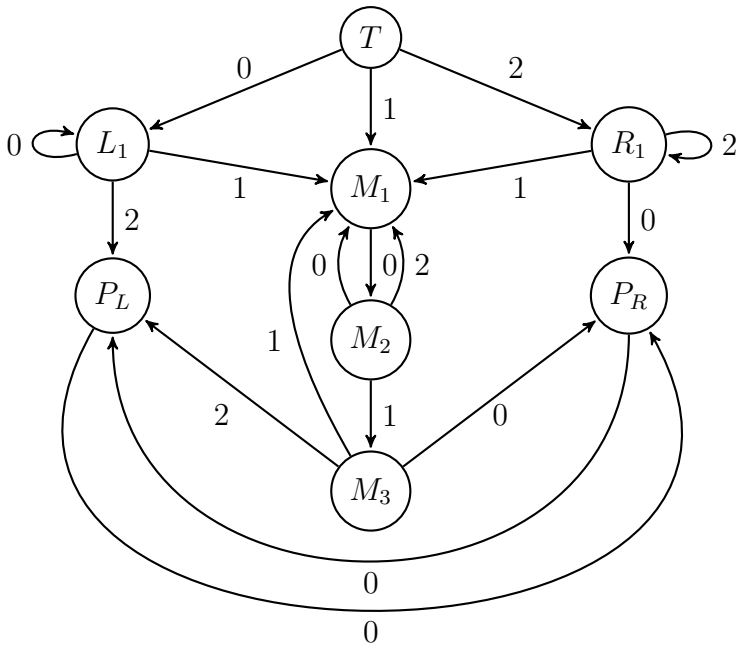

Figure 5.12: The type graph $\mathcal{T}$ of the infinite atoms in $\mathcal{M}$ with labeled edges

The path space $P(\mathcal{T})$ of the graph $\mathcal{T}$ (Figure 5.12) is the set of all infinite sequences of labels of infinite directional paths starting from the vertex T. The topology on $P(\mathcal{T})$ is generated by cylinder sets of the form $C_\alpha = \{w \in P(\mathcal{T}) : w \text{ starts with } \alpha\}$ for finite paths α.

Proposition 5.3.10. *The path space of the graph $\mathcal{T}$ depicted by Figure 5.12 is naturally homeomorphic to $\mathcal{D}_\tau$.*

Proof. We will construct a homeomorphism $f : P(\mathcal{T}) \to \mathcal{D}_\tau$ recursively as follows:

$$f(w) = [0,1]^T_w,$$

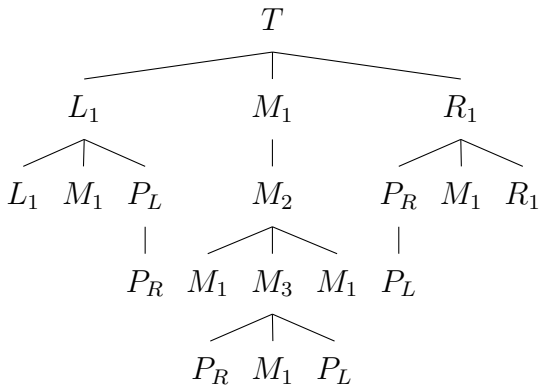

Figure 5.13: The tree of atoms $\mathcal{A}(\mathcal{M})$

where $[x,y]_w^A$ is defined for atoms $A \in \{T, L_1, M_1, R_1, M_2, M_3, P_L, P_R\}$ by:

$$\begin{aligned}
[0,1]_{0\omega}^{T} &= [0,\tau^2]_{\omega}^{L_1}; & [x,y]_{1\omega}^{R_1} &= [x^+, \big(y-\tau(y-x)\big)^-]_{\omega}^{M_1};\\
[0,1]_{1\omega}^{T} &= [\tau^{2^+},\tau^-]_{\omega}^{M_1}; & [x,y]_{2\omega}^{R_1} &= [y-\tau(y-x), y]_{\omega}^{R_1};\\
[0,1]_{2\omega}^{T} &= [\tau,1]_{\omega}^{R_1}; & [x^+,y^-]_{0\omega}^{M_2} &= [x^+, (x+\tau^2(y-x))^-]_{\omega}^{M_1};\\
[x,y]_{0\omega}^{L_1} &= [0,\tau y]_{\omega}^{L_1}; & [x^+,y^-]_{1\omega}^{M_2} &= [x+\tau^2(y-x), y-\tau^2(y-x)]_{\omega}^{M_3};\\
[x,y]_{1\omega}^{L_1} &= [(\tau y)^+, y^-]_{\omega}^{M_1}; & [x^+,y^-]_{2\omega}^{M_2} &= [\big(y-\tau^2(y-x)\big)^+, y^-]_{\omega}^{M_1};\\
[x,y]_{2\omega}^{L_1} &= \{y\}; & [x,y]_{0\omega}^{M_3} &= \{x\};\\
[x^+,y^-]_{0\omega}^{M_1} &= [x^+,y^-]_{\omega}^{M_2}; & [x,y]_{1\omega}^{M_3} &= [x^+,y^-]_{\omega}^{M_1};\\
[x,y]_{0\omega}^{R_1} &= \{x\}; & [x,y]_{2\omega}^{M_3} &= \{y\}.
\end{aligned} \tag{5.2}$$

Although this function is written in a complicated form it is not difficult to understand, as it just represents recursive atom decompositions. This definition of f means that it takes an infinite path $w = w_1w_2\cdots \in P(\mathcal{T})$ and produces a sequence of nested intervals in $[0,1]$, each corresponding to a subdivision of an atom according to the choices of $w_1, w_2, \cdots$. At each step, the function f applies a transformation to the interval $[0,1]$ based on the current value of w_i, as specified by the recursive rules. Formally, $\bigcap_{n=1}^{\infty}[0,1]_{w_1w_2\cdots w_n}^{T} = \{f(w_1w_2\cdots)\}$. Note the width of $[0,1]_{w_1\cdots w_n}^{T}$ is less than $\tau^{\frac{n}{2}}$, we add a factor of $\frac{1}{2}$ to compensate the $[x^+,y^-]_{0\omega}^{M_1}$ action that does not alter the interval.

By construction, for any $w \in P(\mathcal{T})$, $f(w) \in \mathcal{D}_\tau$. Therefore, f is well-defined.

The inverse map $g : \mathcal{D}_\tau \to P(\mathcal{T})$ is defined as follows. Given $x \in \mathcal{D}_\tau$, we express x as a sequence of nested intervals, and recover the sequence $w = (w_1, w_2, w_3, \dots)$ in $P(\mathcal{T})$ such that:

$$w_n = \begin{cases} 0 & \text{when } x \in [0,1]_{w_1w_2\cdots w_{n-1}0}^{T};\\ 1 & \text{when } x \in [0,1]_{w_1w_2\cdots w_{n-1}1}^{T};\\ 2 & \text{when } x \in [0,1]_{w_1w_2\cdots w_{n-1}2}^{T}. \end{cases}$$

This defines a sequence $w = (w_1, w_2, w_3, \dots)$ that corresponds to x. Since g is the inverse of f, we have $g\big(f(w)\big) = w$ for all $w \in P(\mathcal{T})$. Hence, f is a bijection.

We will now show that open sets $U \in P(\mathcal{T})$ are mapped to open sets $f(U) \in \mathcal{D}_\tau$. By construction every open set in $P(\mathcal{T})$ is a cylinder set. Recall that a cylinder set is defined by $C_\alpha = \{w \in P(\mathcal{T}) : w \text{ starts with } \alpha\}$ for finite paths α. Hence, we only need to show that every finite path in $P(\mathcal{T})$ is mapped by f to an open set in $D_\mathcal{T}$. Due to the structure of $D_\mathcal{T}$, every point and closed interval contained in $D_\mathcal{T}$ can be expressed as an open interval. We will change the representation of equations 5.2 to show that

every image of a finite path $\alpha \in P(\mathcal{T})$ under f is an open set:

$$
\begin{aligned}
[0,1]^{T}_{0\omega} &= \left(-\infty, (\tau^2)^+\right)^{L_1}_{\omega}; & [x,y]^{R_1}_{1\omega} &= \left(x, y-\tau(y-x)\right)^{M_1}_{\omega};\\
[0,1]^{T}_{1\omega} &= (\tau^2,\tau)^{M_1}_{\omega}; & [x,y]^{R_1}_{2\omega} &= \left(\left(y-\tau(y-x)\right)^-, \infty\right)^{R_1}_{\omega};\\
[0,1]^{T}_{2\omega} &= (\tau^-,+\infty)^{R_1}_{\omega}; & [x^+,y^-]^{M_2}_{0\omega} &= \left(x, x+\tau^2(y-x)\right)^{M_1}_{\omega};\\
[x,y]^{L_1}_{0\omega} &= \left(-\infty, (\tau y)^+\right)^{L_1}_{\omega}; & [x^+,y^-]^{M_2}_{1\omega} &= \left(\left(x+\tau^2(y-x)\right)^-, \left(y-\tau^2(y-x)\right)^+\right)^{M_3}_{\omega};\\
[x,y]^{L_1}_{1\omega} &= (\tau y, y)^{M_1}_{\omega}; & [x^+,y^-]^{M_2}_{2\omega} &= \left(y-\tau^2(y-x), y\right)^{M_1}_{\omega};\\
[x,y]^{L_1}_{2\omega} &= (y^-,y^+); & [x,y]^{M_3}_{0\omega} &= (x^-,x^+);\\
[x^+,y^-]^{M_1}_{0\omega} &= (x,y)^{M_2}_{\omega}; & [x,y]^{M_3}_{1\omega} &= (x,y)^{M_1}_{\omega};\\
[x,y]^{R_1}_{0\omega} &= (x^-,x^+); & [x,y]^{M_3}_{2\omega} &= (y^-,y^+).
\end{aligned}
\tag{5.3}
$$

Formally, if α is a finite path in $P(\mathcal{T})$ and $f(\alpha)$ transforms the interval $[0,1]$ n times, then we use 5.2 for first $n-1$ transformations and 5.3 for the last n^{th} transformation. This approach guarantees that the function remains well-defined and the resulting interval is open. Hence, the inverse g is a continuous function.

By [35, Theorem 2.1], that states that the path space of any directed graph is a locally compact Hausdorff space, we conclude that $P(\mathcal{T})$ is a Hausdorff space.

Since the Cantor-like space $D_{\mathcal{T}}$ is a union of a Cantor set and isolated points, it is closed and bounded, and hence compact.

Finally, we use the fact that a bijective continuous function from a compact space to a Hausdorff space is a homeomorphism. The map $f : P(\mathcal{T}) \to D_{\mathcal{T}}$ is such a function, hence f is indeed a homeomorphism. □

We now have all the necessary tools to prove Theorem 5.3.2. This proof brings together our analysis of the tree infinite atom $\mathcal{A}(M)$, the limit behavior of infinite sequences in $\mathcal{A}(M)$, and the unique properties of the small golden ratio.

Proof of Theorem 5.3.2. By Proposition 5.3.10 the space of the path space $P(\mathcal{T})$ of the type graph $\mathcal{T}$ (see Figure 5.12) is homeomorphic to $\mathcal{D}_{\tau}$. According to [5, Proposition 2.21] the boundary of the tree of atoms $\mathcal{A}(M)$ (see Figure 5.13) is isomorphic to the path space of the type graph $\mathcal{T}$. Finally, by Theorem 5.1.8 the boundary of the tree of atoms $\mathcal{A}(M)$ is homeomorphic to the horofunction boundary $\partial_h\mathcal{M}$ of $\mathcal{M}$.

Hence, horofunction boundary $\partial_h\mathcal{M}$ of $\mathcal{M}$ is homeomoprhic to $\mathcal{D}_{\tau}$. □

This result reveals intricate large-scale geometry properties of the monoid M. The horofunction boundary of its Cayley graph is homeomorphic to a Cantor-like space union a countable set of isolated points. This duality suggests that, while the space exhibits self-similar properties at large scales, local

variations near the isolated points influence the geometric behavior at infinity. The resulting horofunction boundary captures this complexity, offering new insights into the asymptotic structure of M and its associated Cayley graph.

Remark 5.3.11. It was an open question if any non-elementary hyperbolic groups have isolated points in their horofunction boundaries. In the next chapter we show that M is hyperbolic, so this gives an example of a hyperbolic monoid whose horofunction boundary has isolated points.

Chapter 6

Gromov hyperbolicity

This section provides proof of one of the main results of this paper, Theorem 6.3.1, which states that the Cayley graph $\mathcal{M}$ of the monoid $M = \langle L, R : LR^2 = RL^2 \rangle$ is hyperbolic in the Gromov sense. We start by recalling some background on graphs and quasi-isometries. After this is done, we begin the proof by introducing a modified graph and proving that it is a quasi-isometry of the graph M. Relying on Theorem 6.1.4 from the results of Kong, Lau and Wang [25], we prove that the newly introduced graph is hyperbolic.

6.1 Preliminary on Horizontal Graphs

In this section, we restate some background definitions and properties following [25, Sections 1-2]. We introduce the reader with the notion of horizontal edges, which play a key role in the subsequent sections.

Let $\Gamma = (V, E, r)$ be a rooted graph. The *vertical edge set* $E_v \subseteq E$ is defined as:

$$E_v = \{\{x, y\} \in E : |x| - |y| = \pm 1\}.$$

The *horizontal edge set* $E_h \subset E$ is defined as:

$$E_h = \{\{x, y\} \in E : |x| = |y|\}.$$

Therefore, the edge set E can be partitioned as $E = E_v \cup E_h$. Note that in a connected graph with $E \neq \emptyset$, E_h may be empty, but E_v is always non-empty.

For $\Gamma = (V, E, r)$, we define *horizontal distance* $d_h(\cdot, \cdot)$ as the graph distance on the subgraph induced by (V, E_h). For any pair of vertices $x, y \in V$, $d_h(x, y) = \infty$ whenever $|x| \neq |y|$. Moreover, the inequality $d(x, y) \leq d_h(x, y)$ holds for all $x, y \in V$. When the equality occurs, i.e., $d(x, y) = d_h(x, y)$, there exists a geodesic $\pi(x, y)$ that is entirely contained in (V, E_h). We refer to it as the *horizontal geodesic* of Γ denoted by $\pi_h(x, y)$.

Let's briefly recall from Section 4.1 the notation for descendant and predecessor sets. Let $m \geq 0$

and $x \in V$, then

$$J_m(x) := \{y \in V : x \preceq y, |y| = |x| + m\};$$
$$J_{-m}(x) := \{z \in V : x \in J_m(z)\};$$

are the m-th descendant and m-th predecessor sets of x respectively.

Definition 6.1.1. [25, Definition 2.1] A rooted graph $\Gamma = (V, E, r)$ is called *expansive* if it satisfies the following condition:

$$\forall x, y \in V,\ \forall u \in J_1(x),\ \forall v \in J_1(y),\ d_h(x, y) > 1 \implies d_h(u, v) > 1.$$

Definition 6.1.2. [25, Definition 2.5] A rooted graph $\Gamma = (V, E, r)$ is called (m, k)-*departing* if there exist $m, k \in \mathbb{N}$ such that:

$$\forall x, y \in V,\ d_h(x, y) > k \implies \forall u \in J_m(x), v \in J_m(y),\ d_h(u, v) > 2k.$$

Let's consider a short example of an expansive graph.

Example 6.1.3. Let $\Gamma_n = (V, E, r)$ be an infinite n-ary tree (where each vertex has exactly n descendants) with additional horizontal edges between every pair of vertices in lexicographic order on the same level.

This family of graphs Γ_n is expansive. To prove this, consider any two vertices $x, y \in V$ such that $d_h(x, y) > 1$, and let $u \in J_1(x)$ and $v \in J_1(y)$. If $|x| = |y|$, then $|u| = |v| = |x| + 1$, and $d_h(u, v) > 1$ because the horizontal distance between parents/children of non-adjacent vertices at the same level is always greater than 1. If $|x| \neq |y|$, then $d_h(u, v) = \infty > 1$ by the definition of horizontal distance. Therefore, Γ_n satisfies the expansive condition for all $n \geq 2$.

Let $\Gamma = (V, E)$ be a locally finite connected graph. Let x, y and z be vertices in V, then the three geodesics joining them are called *sides* and form a *geodesic triangle*. If each of the sides of geodesic triangle is contained in δ-neighbourhood of the union of the other two sides, for some non-negative δ, then such triangle is called δ-*thin*. A locally finite connected graph is called δ-*hyperbolic* if each geodesic triangle in Γ is δ-thin. We call the smallest such $\delta \geq 0$ the *hyperbolicity constant* of Γ. A graph is called *hyperbolic in the Gromov sense* or simply *hyperbolic* if there exists a $\delta \geq 0$ such that each subgraph of Γ is δ-hyperbolic. [22].

Theorem 6.1.4. *[25, Theorem 1.1] Let $\Gamma = (V, E, r)$ be an expansive rooted graph. The following statements are equivalent:*

1. *(V, E, r) is hyperbolic;*
2. *There exists a constant $P < \infty$ such that the lengths of all horizontal geodesics are bounded by P;*
3. *(V, E, r) is (m, k)-departing for $m, k \in \mathcal{N}$.*

6.2 Preliminary on quasi-isometries

We now state some definitions and key properties of quasi-isometries, as referenced from various sources. These concepts are fundamental in geometric group theory and play a crucial role in understanding large-scale geometry and spaces with hyperbolic properties.

Definition 6.2.1. [14, Definition 3.11] Let (X, d_X) and (Y, d_Y) be metric spaces, and let $f\colon X \to Y$ be a map. We say that f is a *quasi-isometry* if there exist constants $A \geq 1$, $B \geq 0$ and $C \geq 0$ such that the following conditions hold:

1. $\forall x, y \in X : \frac{1}{A} d_X(x, y) - B \leq d_Y\big(f(x), f(y)\big) \leq A d_X(x, y) + B$;
2. $\forall z \in Y : \exists x \in X : d_Y\big(z, f(x)\big) \leq C$.

The constants A and B are called the *quasi-isometric constants* of the embedding f. When only the first condition is satisfied f is called a *quasi-isometric embedding*.

Remark 6.2.2. [26, Definition 3.1.1] Let f be a surjective quasi-isometric embedding from X to Y then f is a quasi-isometry.

Theorem 6.2.3. *[26, Theorem 1.2.3] Let X and Y be geodesic metric spaces. Let $f : X \to Y$ be a quasi-isometry. Then X is hyperbolic if and only if Y is hyperbolic.*

6.3 Hyperbolicity of the Cayley Graph $\mathcal{M}$

Hyperbolicity in graphs is a fundamental concept in geometric group theory. It indicates tree-like properties on a large scale and often leads to efficient solutions to word problems [22]. In the context of monoids, a hyperbolic Cayley graph suggests negatively curved geometry. This has profound implications for the structure's behavior, including efficient word problems and connections to finitely generated groups. This section proves the following.

Theorem 6.3.1. *The Cayley graph $\mathcal{M} = (V, E, r)$ of the monoid $M = \langle L, R : LR^2 = RL^2 \rangle$ is hyperbolic.*

We will prove this theorem by a number of lemmas. The first step is to construct a modified graph $\mathcal{M}'$ that includes additional horizontal edges between every pair of adjacent vertices, that are not present in $\mathcal{M}$. The next step will be establish that the two graphs are quasi-isometric. Then using Theorem 6.2.3, we show that $\mathcal{M}$ is hyperbolic by proving that $\mathcal{M}'$ is hyperbolic.

We will now introduce the modified graph $\mathcal{M}'$. Since there are no horizontal edges in $\mathcal{M} = (M, E, 1)$, we allow $\mathcal{M} = (M, E_v, 1)$. Let $\mathcal{M}' = (M, E_v \cup E_h, 1)$, where E_h is the set of horizontal edges defined by:

$$E_h = \{(x, y) : \exists m \in V,\ x = mL, y = mR\} \cup \{(x, y) : \exists m \in V,\ x = mLR, y = mRL\}.$$

See Figure 6.1 for an illustration.

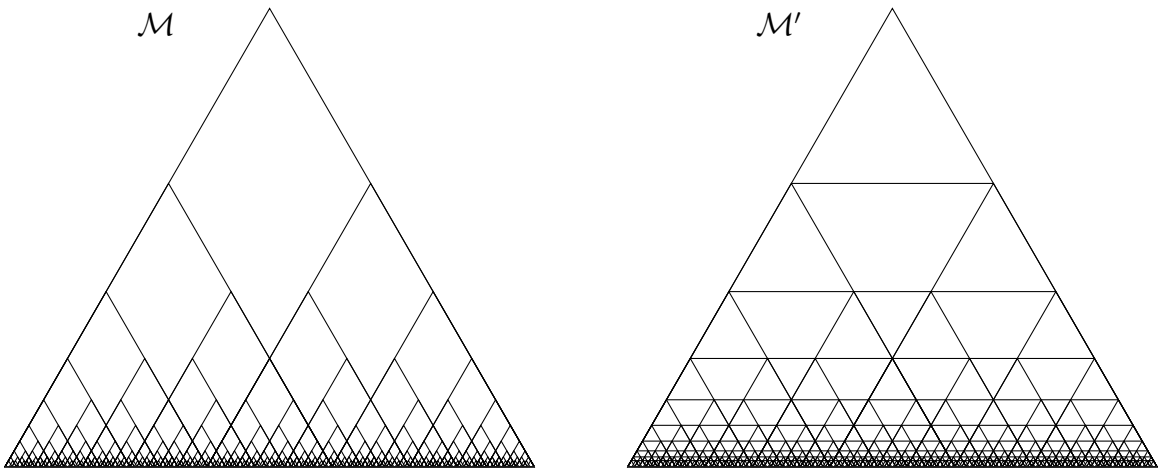

Figure 6.1: The graphs $\mathcal{M}$ and $\mathcal{M}'$

Observe that the graph $\mathcal{M}$ is contained in $\mathcal{M}'$. We now have enough tools to prove that the large-scale geometry of the graphs $\mathcal{M}$ and $\mathcal{M}'$ are essentially the same.

Proposition 6.3.2. *Let $f : \mathcal{M} \to \mathcal{M}'$ be the natural inclusion map, then f is a quasi-isometry.*

Proof. Let $f\colon \mathcal{M} \to \mathcal{M}'$ be the natural inclusion map. Since $\mathcal{M}'$ contains every path of $\mathcal{M}$, we conclude that for any pair of vertices $x, y \in \mathcal{M}$, the following holds:

$$d_{\mathcal{M}'}\big(f(x), f(y)\big) \leq d_{\mathcal{M}}(x, y).$$

On the other hand, every added horizontal edge in $\mathcal{M}'$ decreases the minimal distance between the points it connects by a factor of 2. Hence, for every pair of vertices $x, y \in \mathcal{M}$, we have

$$\frac{1}{2} d_{\mathcal{M}}(x, y) \leq d_{\mathcal{M}'}\big(f(x), f(y)\big).$$

Combining these two inequalities, we obtain

$$\frac{1}{2} d_{\mathcal{M}}(x, y) \leq d_{\mathcal{M}'}\big(f(x), f(y)\big) \leq 2 d_{\mathcal{M}}(x, y),$$

which shows that the natural inclusion map $f\colon \mathcal{M} \to \mathcal{M}'$ is a quasi-isometric embedding with quasi-isometric constants $A = 2$ and $B = 0$. Therefore, $\mathcal{M}$ quasi-isometrically embeds into $\mathcal{M}'$. Since f is surjective we conclude that it is a quasi-isometry. □

Having established the existence of a quasi-isometry between the spaces $\mathcal{M}$ and $\mathcal{M}'$, we can now investigate certain large-scale geometric properties of $\mathcal{M}$ by studying their counterparts in $\mathcal{M}'$. This approach is justified by the fact that quasi-isometries preserve many coarse geometric features, allowing us to transfer certain structural insights from one space to the other.

Proposition 6.3.3. *The graph $\mathcal{M}' = (V, E_v \cup E_h, r)$ is expansive.*

Proof. We will prove this using the contrapositive statement

$$\forall u, v \in V, \text{ and } \forall x \in J_{-1}(u),\ \forall y \in J_{-1}(v) \text{ s.t. } d_h(u,v) \leq 1 \Rightarrow d_h(x,y) \leq 1.$$

We will consider two possible cases corresponding to the possible horizontal distances between u and v.

Let $d_h(u,v) = 0$, i.e. u and v are the same vertex. In this case u and v have the same predecessors due to the monoid property $LR^2 = RL^2$. If for some predecessor pair x and y, $d_h(x,y) > 1$, the restriction $LR^2 = RL^2$ would force $d_h(u,v) > 1$, contradicting our assumption that $d_h(u,v) = 0$.

Now let $d_h(u,v) = 1$. Due to the monoid property $LR^2 = RL^2$, within the next two levels above u and v, there always exists a common ancestor $w \in J_{-2}(u)$, $w \in J_{-2}(v)$. This common ancestor w ensures that all predecessors x and y are at most 1 apart, as any greater distance would contradict the existence of the common ancestor within two levels.

Both cases show that $d_h(u,v) \leq 1 \implies d_h(x,y) \leq 1$. Hence, this proves the contrapositive statement, which is equivalent to the definition of expansiveness. Hence, the graph $\mathcal{M}'$ is expansive. □

Our next goal is to prove that the graph $\mathcal{M}'$ is (m,k)-departing. But first we need some definitions. In the graph $\mathcal{M}'$, we define two types of horizontal edges of length 1:

- let U be a horizontal edge where the right predecessor of the left vertex and the left predecessor of the right vertex coincide;
- let D be a horizontal edge where the right descendant of the left vertex and the left descendant of the right vertex coincide.

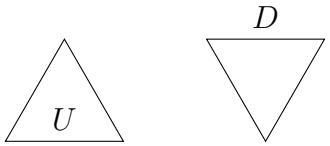

In the graph M', the following properties hold:

1. an edge cannot be both U and D simultaneously;
2. the pattern of U and D edges follows a specific structure:
   - Level 1 (from the root): U
   - Level 2: UDU
   - Level 3: $UDUUDU$
   - Level 4: $UDUUDUDUUDU$

The patterns above are obtained from substitution structure. By analyzing the structure we obtain the following substitution rules:

1. a single edge U, that is not adjacent to any other U's, at level n is transformed into a sequence of UDU edges at level $n+1$;

2. a sequence of UU edges at level n is transformed into a sequence of $UDUDU$ edges at level $n+1$;

3. an edge D is ignored in the transformation to the next level. The purpose for D's is to separate U's into singles and doubles.

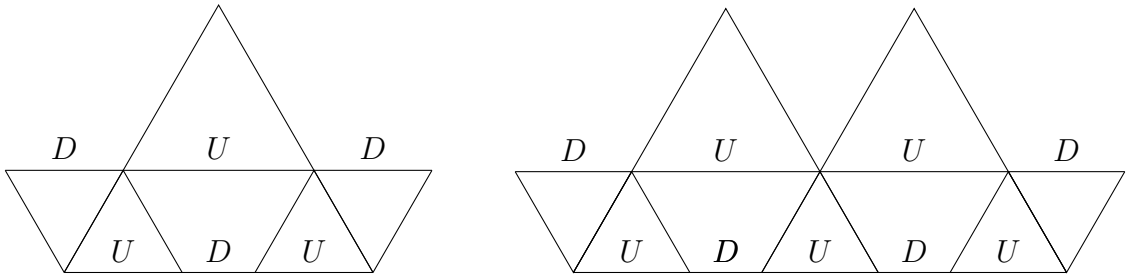

We invite the reader to expand these pictures with various combinations of types of horizontal edges.

Note that we can define U and D for the graph $\mathcal{M}$ by considering the same pairs of vertices.

Proposition 6.3.4. *The graph $\mathcal{M}'$ is (m,k)-departing.*

Proof. We will show that M' is at most $(3,2)$-departing, i.e.

$$\forall x, y \in X \text{ and } \forall u \in J_3(x), v \in J_3(y) \text{ s.t. } d_h(x,y) > 2 \Rightarrow d_h(u,v) > 4,$$

We will prove by proving the contrapositive statement:

$$\forall u, v \in X \text{ and } \forall x \in J_{-3}(u), y \in J_{-3}(v) \text{ s.t. } d_h(u,v) \leq 4 \Rightarrow d_h(x,y) \leq 2,$$

Without loss of generality, let u be located to the left of v on the graph, such that $d_h(u,v) \leq 4$. Due to lexicographic order we only need to consider the extreme case when $d_h(u,v) = 4$. There are 3 possible horizontal edge combinations of length 4: $UDUD$, $UUDU$, and $DUUD$. We ignore cases like $UDUU$ due to symmetry.

We start from analyzing the $UDUD$ case. The key here is to observe that in recovering the predecessors an horizontal edge of type U contributes is the appearance of 2 vertical edges, whereas type D lacks contribution. Note the adjacent edges of the sequence $UDUD$ may also contribute in the formation of the horizontal edges at the level above, hence we have to look at all possible combinations in which a $UDUD$ can appear. By analyzing the substitution structure we observe there is a total of 3 such subcases: $UUDUDU$, $UDUDU$, and $DUDUDU$.

The following picture shows two levels of edges above the sequence of edges $UDUD$ inside $UUDUDU$:

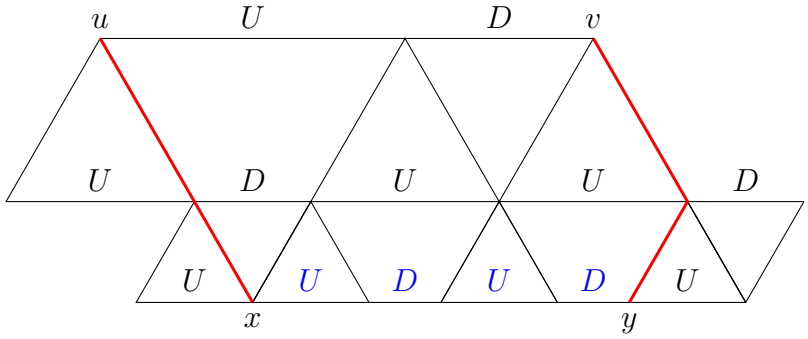

Let x and y be the extreme vertices in the sequence of edges $UDUD$. Then the distance between the extreme predecessors $J_{-2}(x) = u$ and $J_{-2}(y) = v$ is 2. By the monoid property the the distance of $J_{-1}(u)$ and $J_{-1}(v)$ cannot be greater than 2. Note the sequence $UDUD$ is always followed by a U, however it may or may not have a prefix of a U or a D. The absence or the presence of a D in front of the sequence will not contribute in edges in the previous level. The $UDUDU$, and $DUDUDU$ subcases are proven in a similar way.

The next case is $UUDU$. Once again we only look at $UUDU$ is inside: $DUUDUU$, $UUDUU$, $UUUDUD$, $UUDUD$, $DUUDU$ and $UUDU$. The following picture shows two levels of edges above the sequence of edges $UUDU$ inside $DUUDUU$:

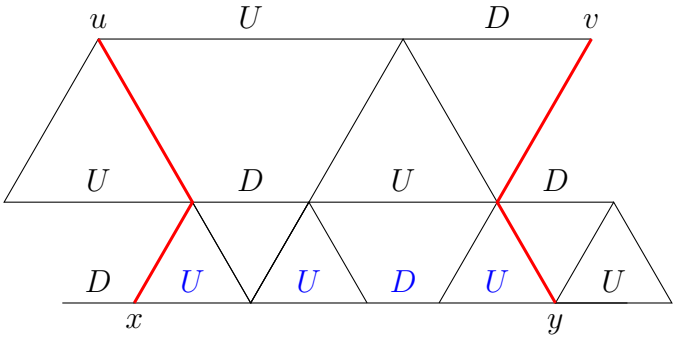

Observe that we obtain the same situation for the extreme predecessors $J_{-2}(x) = u$ and $J_{-2}(y) = v$ as before. The subcases $UUDUU$, $UUUDUD$, $UUDUD$, $DUUDU$ and $UUDU$ are proven in a similar way.

We are left with the $DUUD$ case. Since D's are not allowed to be adjacent to each other or at the

boundary, we look at $DUUD$ inside $UDUUDU$.

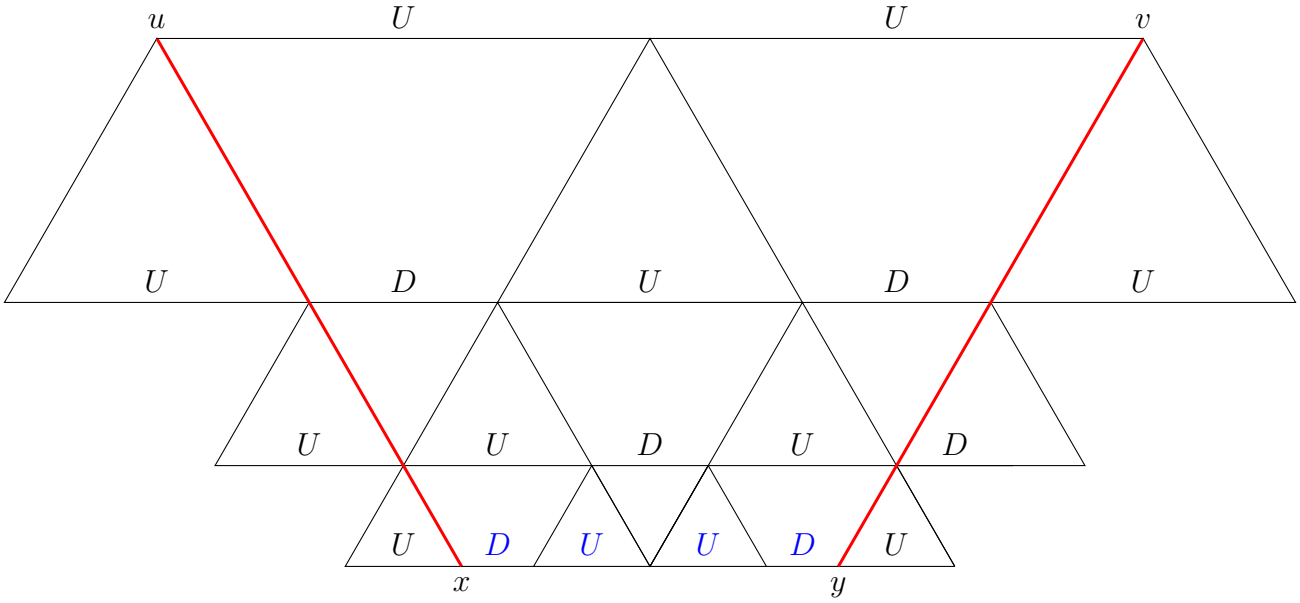

For the extreme case $J_{-3}(x) = u$ and $J_{-3}(y) = v$ we obtain that the distance is at most 2.

Note that we proved only the extreme cases. All other cases are proven in a similar way, moreover most of them are contained in the pictures provided above. Hence, this proves the contrapositive statement, which implies that $\mathcal{M}'$ is at most $(3, 2)$-departing. □

Corollary 6.3.5. *The graph $\mathcal{M}'$ is hyperbolic.*

Proof. Proposition 6.3.3 shows that $\mathcal{M}'$ is expansive. Hence, we can apply of Theorem 6.1.4 to Proposition 6.3.4, which concludes that $\mathcal{M}'$ being (m, k)-departing is equivalent to $\mathcal{M}'$ being hyperbolic. □

We can now prove Theorem 6.3.1.

Proof of Theorem 6.3.1. Proposition 6.3.2 states that $f : \mathcal{M} \to \mathcal{M}'$ is a quasi-isometry. Corollary 6.3.5 shows that $\mathcal{M}'$ is hyperbolic. Then by Theorem 6.2.3 $\mathcal{M}$ must also be hyperbolic. □

The proof that the Cayley graph $\mathcal{M}$ of the monoid $M = \langle L, R : LR^2 = RL^2 \rangle$ is hyperbolic establishes its underlying negative curvature and tree-like structure. This result opens the door to applying hyperbolic geometric techniques and provides a strong foundation for further exploration of the algebraic and geometric properties of the monoid.

Chapter 7

Conclusion

In this thesis, we have investigated several key properties of the golden ratio Thompson group V_τ and the monoid $M = \langle L, R : LR^2 = RL^2 \rangle$, focusing on their geometric and algebraic structures through automata theory, Cayley graphs, and horofunction boundaries.

We proved that the group G of homeomorphisms on the Cantor set $\{0,1\}^\omega$, generated by the rational homeomorphisms $X_0, X_1, Y_0, Y_1, C_1, C_2, \Pi_0, \Pi_1$, is isomorphic to the golden ratio Thompson group V_τ. Further, we explored the monoid M by examining the Cayley graph $\mathcal{M}$ and establishing an important distance function between vertices. This result lays the groundwork for understanding the geometry of $\mathcal{M}$, showing that it has a highly nontrivial structure that involves self-similar cones and geodesically convex regions. We also proved that the horofunction boundary of the Cayley graph is homeomorphic to a Cantor-like set with additional isolated points between every pair of breakpoints $\mathcal{D}_\tau$, expanding our understanding of the boundary of such graphs. In our last section we proved the hyperbolicity of the Cayley graph $\mathcal{M}$. This finding is significant because it indicates that $\mathcal{M}$ shares the properties of hyperbolic spaces, which have well-understood geometric and algebraic characteristics.

These results not only bring important insights but also raise a number of questions to be answered:

1. Are all irrational slope Thompson's groups rational similarity groups? If not, then what are the criteria for a Thompson's group to be one?

2. Are all horofunction boundaries of the Cayley graphs of the family of monoids $M_n = \langle L, R : LR^n = RL^n \rangle$, for $n > 1$, Cantor-like sets with additional isolated points in between every pair of breakpoints?

3. Are all Cayley graphs of the family of monoids $M_n = \langle L, R : LR^n = RL^n \rangle$, for $n > 1$, hyperbolic?

4. What is the Gromov boundary of the Cayley graphs of the family of monoids $M_n = \langle L, R : LR^n = RL^n \rangle$, for $n > 1$?

Bibliography

[1] Belk, J. (2004). Thompson's group F. PhD thesis, Cornell University.

[2] Belk, J., and Bleak, C. (2023). Embedding hyperbolic groups into finitely presented infinite simple groups. arXiv preprint arXiv:2306.14863.

[3] Belk, J., and Bleak, C. (2017). Some undecidability results for asynchronous transducers and the Brin-Thompson group 2V. Transactions of the American Mathematical Society, 369(5), 3157-3172.

[4] Belk, J., Bleak, C., Matucci, F., and Zaremsky, M. C. (2023). Hyperbolic groups satisfy the Boone-Higman conjecture. arXiv preprint arXiv:2309.06224.

[5] Belk, J., Bleak, C., and Matucci, F. (2021). Rational embeddings of hyperbolic groups. Journal of Combinatorial Algebra, 5(2), 123-183

[6] Belk, J., Hyde, J., and Matucci, F. (2019). On the asynchronous rational group. Groups, Geometry, and Dynamics, 13(4), 1271-1284.

[7] Bergman, G. (1957). A number system with an irrational base. Mathematics magazine, 31(2), 98-110.

[8] Burillo, J., Nucinkis, B., and Reeves, L. (2021). An irrational-slope Thompson's group. Publicacions matemàtiques, 65(2), 809-839.

[9] Burillo, J. Introduction to Thompson's group F.

[10] Burillo, J., Nucinkis, B., and Reeves, L. (2022). Irrational-slope versions of thompson's groups T and V. Proceedings of the Edinburgh Mathematical Society, 65(1), 244-262.

[11] Cannon, J. W., Floyd, W. J., Kenyon, R., and Parry, W. R. (1997). Hyperbolic geometry. Flavors of geometry, 31(2),59-115.

[12] Cannon, J. W. (1996). Introductory notes on Richard Thompson's groups. Enseignement Mathematique. 42(2), 215-256.

[13] Cleary, S. (2000). Regular subdivision in $\mathbb{Z}[\frac{1+\sqrt{5}}{2}]$. Illinois Journal of Mathematics, 44(3), 453-464.

[14] Druţu, C., and Sapir, M. (2005). Tree-graded spaces and asymptotic cones of groups. Topology, 44(5), 959-1058.

[15] Duboc, C. (1986). Mixed product and asynchronous automata. Theoretical Computer Science, 48, 183-199.

[16] Dydak, J. (1977). 1-movable continua need not be pointed 1-movable. Bulletin de l'Academie polonaise des sciences-serie des sciences mathematiques, astronomiques, et physiques, 25(6), 559-562.

[17] Elder, M. (2012). A short introduction to self-similar groups. The Australian Mathematical Society Gazette, 39(3), 125-133.

[18] Eilenberg, S. (1974). Automata, languages, and machines. Academic press.

[19] Freyd, P., and Heller, A. (1993). Splitting homotopy idempotents II. Journal of pure and applied algebra, 89(1-2), 93-106.

[20] Geoghegan, R., and Brown, K. S. (1984). An infinite-dimensional torsion-free FP_∞ group. Inventiones Mathematicae, 77, 367-381.

[21] Grigorchuk, R. I., Nekrashevych, V. V., and Sushchansky, V. I. (2000). Automata, dynamical systems, and groups. Trudy Matematicheskogo Instituta Imeni VA Steklova, 231, 134-214.

[22] Gromov, M. (1987). Hyperbolic groups. In Essays in group theory. Mathematical Sciences Research Institute Publications 8, 75-263.

[23] Gromov, M. (1981). Hyperbolic manifolds, groups and actions, Riemannian Surfaces and Related Topics. Princeton University Press, 97, 213-307.

[24] Jech, T. (2003). Set theory: The third millennium edition, revised and expanded. Springer Berlin Heidelberg.

[25] Kong, S. L., Lau, K. S., and Wang, X. Y. (2021). Gromov hyperbolic graphs arising from iterations. arXiv:2006.12916.

[26] Lanfranco, J. (2019). An Introduction to Quasi-isometry and Hyperbolic Groups. Master's thesis. University of Pennsylvania.

[27] McKenzie, R., and Thompson, R. J. (1973). An elementary construction of unsolvable word problems in group theory. In Studies in Logic and the Foundations of Mathematics, 71, 457-478.

[28] Meier, J. (2008). Groups, graphs and trees, an Introduction to the Geometry of Infinite Groups. Cambridge University Press.

[29] Mihalik, M. L. (1985). Ends of groups with the integers as quotient. Journal of Pure and Applied Algebra, 35, 305-320.

[30] Molyneux, L., Nucinkis, B., and Rego, Y. S. (2024). The sigma invariants for the golden mean Thompson group. New York Journal of Mathematics, 30, 532–549.

[31] Nekrashevych, V. (2024). Self-similar groups. American Mathematical Society, 117.

[32] Perego, D. (2023). Rationality of Boundaries. PhD thesis, Universita degli Studi di Milano-Bicocca.

[33] Rover, C. E. (1999). Constructing finitely presented simple groups that contain Grigorchuk groups. Journal of Algebra, 220(1), 284-313.

[34] Thompson, R. J. (1980). Embeddings into finitely generated simple groups which preserve the word problem. In Studies in Logic and the Foundations of Mathematics, 95, 401-441.

[35] Webster, S. (2014). The path space of a directed graph. Proceedings of the American Mathematical Society, 142(1), 213-225.

[36] Webster, C., and Winchester, A. (2005). Boundaries of hyperbolic metric spaces. Pacific journal of mathematics, 221(1), 147-158.

[37] Willard, S. (1970). General Topology Addison-Wesley. Reading, MA.