Cutting Sequences and the p-adic Littlewood Conjecture

BLACKMAN, JOHN,EDWARD

Cutting Sequences and the p-adic Littlewood Conjecture

John Blackman

A Thesis presented for the degree of
Doctor of Philosophy

Pure Mathematics
Department of Mathematical Sciences
Durham University
United Kingdom

October 2020

Cutting Sequences and the p-adic Littlewood Conjecture

John Blackman

Submitted for the degree of Doctor of Philosophy

October 2020

Abstract: The main aim of this thesis is to use the geometric setting of cutting sequences to better understand the behaviour of continued fractions under integer multiplication. We will use cutting sequences to construct an algorithm that multiplies continued fractions by an integer n. The theoretical aspects of this algorithm allow us to explore the interesting properties of continued fractions under integer multiplication. In particular, we show that an eventually recurrent continued fractions remain eventually recurrent when multiplied by a rational number. Finally, and most importantly, we provide a reformulation the p-adic Littlewood Conjecture in terms of a condition on the semi-convergents of a real number α.

Declaration

The work in this thesis is based on research carried out in the Department of Mathematical Sciences at Durham University. No part of this thesis has been submitted elsewhere for any degree or qualification. This studentship (EP/N509462/1) was funded by the EPSRC Doctoral Training Partnership.

Acknowledgements

First and foremost, I would like to thank my supervisor, Dr. Anna Felikson, for her continued support and encouragement throughout my Ph.D. studies. Anna has given me a large amount of freedom to pursue my mathematical interests, which has helped me to grow as a mathematician. I would also like to thank Dr. Pavel Tumarkin and Prof. John Parker for their advice and support over the years.

Additionally, I would like to thank all those in the department who organise the seminars, help the department run smoothly and, all-in-all, make the department a nice welcoming environment. I would like to give a special thanks to the Maths office, as well as Pam and all of the cleaning staff.

I would also like to my office mates Irene, Will, Oliver, Phil and Samuel, for useful mathematical discussions and creating a warm, friendly atmosphere. In addition, I would like to thank my friends Ollie, Lucy, Clare, Christian, Rob, Scott, Kieran, Connor, Matthew and Katie. I would also like to give a special thank you to Matt Northey. Matt is an amazing friend and mathematician, and his friendship has been invaluable over the last seven years.

Finally I would like to thank all my family. In particular, I would like to thank my parents and my sisters for all that they have done over the years. I would also like to thank my partner Abby, who has become a new addition to my family, and I look forward to what the future holds for us.

Contents

Chapter 1

Introduction

1.1 Motivation

The main aim of this thesis is to further understand and develop the link between cutting sequences and continued fraction expansions. In particular, we wish to use cutting sequences to better understand how continued fractions behave when iteratively multiplied by an integer.

The question of how continued fractions behave when iteratively multiplied by an integer is important to several open problems in Diophantine approximation. The most important motivation for us comes from the p-adic Littlewood Conjecture (pLC). This conjecture was first proposed by de Mathan and Teulié in 2004 in [MT04], as a one-dimensional analogue of the classical Littlewood Conjecture. We state the conjecture as follows:

The p-adic Littlewood Conjecture. *For every real number $\alpha \in \mathbb{R}$, we have:*

$$m_p(\alpha) := \inf_{q \in \mathbb{N}} \{q \cdot |q|_p \cdot \|q\alpha\|\} = 0,$$

where $|\cdot|_p$ is the p-adic norm and $\|\cdot\|$ is the distance to the nearest integer function.

Reformulating the p-adic Littlewood Conjecture

Whilst it is not initially clear from the above formulation how integer multiplication of continued fractions relates to the p-adic Littlewood Conjecture, we can reformulate pLC to highlight the importance of integer multiplication of continued fractions in this setting. Before we do this, we first introduce some definitions. Note that this will be covered in more detail in Chapter 2.

If α is a real number, then we will denote its continued fraction expansion as $\overline{\alpha} = [a_0; a_1, \ldots]$ with $a_0 \in \mathbb{N} \cup \{0\}$ and $a_i \in \mathbb{N}$. We refer to the terms a_k appearing in the continued fraction as the *partial quotients* of α. We define the *height function* $B(\alpha)$ of a real number α to be the largest partial quotient in the continued fraction expansion, excluding a_0. More explicitly, we have $B(\alpha) := \sup_{k \in \mathbb{N}} \{a_k : \overline{\alpha} = [a_0; a_1, a_2, \ldots]\}$. We will say that α is *badly approximable* if the height function $B(\alpha)$ is finite. We denote the set of all *badly approximable numbers* as:

$$\mathbf{Bad} := \{\alpha \in \mathbb{R} : B(\alpha) < \infty\}.$$

It is then a classical result of Diophantine approximation that for any $\alpha \in \mathbb{R}$, we can bound the value $\inf_{k \in \mathbb{N}} \{q \cdot \|q\alpha\|\}$ above and below in terms of the height function $B(\alpha)$. More specifically:

$$\frac{1}{B(\alpha) + 2} \le \inf_{k \in \mathbb{N}} \{q \cdot \|q\alpha\|\} \le \frac{1}{B(\alpha)}.$$

Since the p-adic norm is bounded above by $|q|_p \le 1$ for all $q \in \mathbb{N}$, it is not hard to see that if $\alpha \notin \mathbf{Bad}$, then $m_p(\alpha) = 0$. With a fair bit of work - the details of which we present in Chapter 2 - we can then show that for all $\alpha \in \mathbf{Bad}$, we have:

$$m_p(\alpha) = \inf_{\ell \in \mathbb{N} \cup \{0\}} \left\{ \inf_{q \in \mathbb{N}} q \cdot \|qp^{\ell}\alpha\| \right\}.$$

This is then bounded above and below as:

$$\inf_{\ell \in \mathbb{N} \cup \{0\}} \left\{ \frac{1}{B(p^{\ell}\alpha) + 2} \right\} \le \inf_{\ell \in \mathbb{N} \cup \{0\}} \left\{ \inf_{q \in \mathbb{N}} q \cdot \|qp^{\ell}\alpha\| \right\} \le \inf_{\ell \in \mathbb{N} \cup \{0\}} \left\{ \frac{1}{B(p^{\ell}\alpha)} \right\}.$$

This leads to the following reformulation of pLC:

Corollary 2.2.14. (Folklore) *Let* $\alpha \in \mathbf{Bad}$*. Then* α *satisfies pLC if and only if:*

$$\sup_{\ell \in \mathbb{N} \cup \{0\}} \left\{ B(p^{\ell}\alpha) \right\} = \infty.$$

Since $\alpha \in \mathbf{Bad}$, the value of $B(p^{\ell}\alpha)$ is finite for each $\ell \in \mathbb{N} \cup \{0\}$. Therefore, using this reformulation, we can think of the p-adic Littlewood Conjecture as saying that for all $\alpha \in \mathbf{Bad}$ there is some increasing subsequence $\{\ell_m\}_{m \in \mathbb{N}}$ such that the limit of $B(p^{\ell_m}\alpha)$ is unbounded. Alternatively, if α is a counterexample to pLC, then there exists some $K \in \mathbb{N}$ such that for all $\ell \in \mathbb{N} \cup \{0\}$, we have:

$$B(p^{\ell}\alpha) \leq K.$$

It is for this reason that the (theoretical) set of counterexamples to pLC is often referred to as the *multiplicatively badly approximable numbers*:

$$\mathbf{Mad}_p := \{\alpha \in \mathbb{R} : m_p(\alpha) > 0\}.$$

Multiplying Continued Fractions

The above reformulation of pLC is heavily dependent on the behaviour of continued fraction expansions under iterative integer multiplication - more precisely, the behaviour of the size of the largest partial quotient. Therefore, to better understand pLC, it would be useful to have a way of directly computing the continued fraction expansion $\overline{n\alpha}$ from the continued fraction expansion $\overline{\alpha}$, for an arbitrary $n \in \mathbb{N}$.

For every $\alpha \in \mathbb{R}$, we can construct a unique map between α and $\overline{\alpha}$, (if $\alpha \in \mathbb{Q}$, take $\overline{\alpha}$ with 1 as the final partial quotient - otherwise the map is not unique), and a unique map between α and $n\alpha$. Therefore, we expect to be able to construct the map $\overline{n} : \overline{\alpha} \to \overline{n\alpha}$ to get the following commutative diagram:

$$\begin{array}{ccc} \alpha & \xrightarrow{\;n\;} & n\alpha \\ \downarrow & & \downarrow \\ \overline{\alpha} & \overset{}{\underset{\overline{n}}{\dashrightarrow}} & \overline{n\alpha} \end{array}$$

From a theoretical point of view we could avoid the $\overline{n}$ map - denoted by a dashed arrow - and go the "long way" around the diagram. More specifically, we could first take the continued fraction $\overline{\alpha}$ and map this to the real number α. We could then easily multiply α by n to obtain $n\alpha$, and then finally, compute the continued fraction expansion $\overline{n\alpha}$. However, on a more practical level, this produces issues. Whenever $\overline{\alpha}$ is an infinite continued fraction (and not periodic), we can not compute the precise value of α in real time. Instead, we must truncate $\overline{\alpha}$ to some fixed precision and then compute an approximation of α. Since we are only dealing with an approximation of α, any computation done in this way is likely to produce numerical errors. Whilst these errors can be minimised, they can not be avoided in all cases and iterative multiplication compounds the issue. Instead, we wish to directly compute $\overline{n\alpha}$ from $\overline{\alpha}$ in an algorithmic way.

There are a few benefits to expressing the $\overline{n}$ map in an algorithmic way. Firstly, it makes the most heuristic sense. Continued fractions are a way of encoding a continuous object (the real number α) as a collection of discrete data (the partial quotients). If we want to see how this discrete data behaves under integer multiplication, it is more natural to construct a map which acts directly on this discrete data, instead of the continuous object. Secondly, since algorithms are iteratively defined, we can construct them in a way such that they never introduce computational errors. In particular, if we do not have enough data to complete the final step of the algorithm, the algorithm will not produce an output for this final step - as opposed to producing an erroneous output. Finally, if we compute the multiplication map on, say, the first million partial quotients in the continued fraction expansion, and then decide that this is not enough precision for our uses, we can continue the multiplication algorithm from the millionth partial quotient without introducing any errors, provided that we remember the last state we encountered. In the previously explained continuous setting, we would have to completely redo the multiplication step, and then recompute the corresponding continued fraction $\overline{n\alpha}$.

Previous Work on Multiplying Continued Fractions

The topic of constructing an algorithm for integer multiplication of continued fraction expansions has been previously investigated, most notably by Raney in [Ran73] and Liardet and Stambul in [LS98]. In fact, both of these previous works tackled a slightly more general question; they showed how for any real number α and any matrix $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ with non-zero determinant (i.e. $ad - bc \neq 0$), one could use the continued fraction expansion of α to compute the continued fraction expansion of $\beta = M \cdot \alpha := \frac{a\alpha+b}{c\alpha+d}$. These works were themselves inspired by an earlier paper of Hall [Hal47], which outlined how one could construct such multiplication algorithms. However, Hall's algorithm was not very efficient, even for matrices with small determinant. On the other hand, both Raney's algorithm and Liardet and Stambul's algorithm are very efficient for explicit computation. However, for theoretical purposes, these algorithms are not necessarily. The main reason for this is that the algorithm is solely defined for each matrix $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ with non-zero determinant. In particular, if we have $m, n \in \mathbb{N}$ with $m \neq n$, knowing how the $\overline{m}$ multiplication algorithm works does not really tell us much information about the $\overline{n}$ multiplication algorithm.

This is not to say that Raney's algorithm and Liardet and Stambul's algorithm can not be used for theoretical results. Liardet and Stambul show in [LS98] that if $\overline{\alpha} = [0; \overline{k}]$ has period length 1, then $\overline{n\alpha}$ has period length at most 5, and Vandehey notably used Liardet and Stambul's algortihm in [Van17] to show that CF-normality is preserved by non-trivial matrix multiplication. However, in this thesis we present an algorithm which has a more natural link between the continued fraction expansion $\overline{\alpha}$ and corresponding real number α - they are both encoded by the same object in our setting. The most important theoretical aspect to come out of our construction is the concept of *infinite loop* mod n. An infinite loop mod n is any positive real number $\alpha \in \mathbb{R}_{>0}$ which has no semi-convergent denominators divisible by n. As we will see in Chapter 5, infinite loops mod n have very interesting properties with regards to multiplication by n and are very important in understanding the structure of potential counterexamples to pLC. As far as the author is aware, infinite loops as described in this thesis have not been previously discussed. We use the nomenclature “infinite loop”, since when we perform our multiplication by n algorithm on an infinite loop mod n, the process will go on indefinitely, however, it never returns to the initial state.

1.2 Summary and Main Results

The main purpose of Chapter 2 is to introduce some classical results of Diophantine approximation and use these preliminary results to discuss the p-adic Littlewood Conjecture (pLC) and mixed Littlewood Conjecture (mLC). This chapter summarises the current literature available on these topics, and from that point of view, should be considered as background instead of novel material. We begin by briefly discussing the history of continued fractions, of which, more detail can be found in [Bre80]. From there, we give a brief summary of continued fractions and outline the properties of continued fractions that we will use throughout the text. We then move on to discussing the classical Littlewood conjecture. This serves two main purposes: firstly, it provides motivation as to why one would want to study the mixed and p-adic Littlewood conjectures, and secondly, it allows us to discuss the analogies between the Littlewood conjecture and these conjectures. After discussing some key results pertaining to the Littlewood conjecture, we then formally introduce the mixed and p-adic Littlewood conjecture. We discuss the main developments with regards to these conjectures and finish the chapter by proving the folklore reformulation of pLC. A good source which outlines the major developments for both the Littlewood and mixed Littlewood Conjecture can be found in [Bug14].

In Chapter 3, we introduce the notion of cutting sequences and describe how one can use this geometric setting to understand how integer multiplication affects continued fractions. Section 3.1 is more background, with the main aim being to introduce the notion of a cutting sequence and discuss the main properties of cutting sequences. We start by defining a cutting sequence of a geodesic ray ζ intersecting an ideal triangulation T of $\mathbb{H}$. We then replace the geodesic ray with an arbitrary path, so that we can see that homotopy only affects cutting sequences in a fairly trivial way. We then introduce the Farey tessellation $\mathcal{F}$ as an ideal triangulation of $\mathbb{H}$. This section ends by providing the link between cutting sequences relative to $\mathcal{F}$ and continued fractions. In particular, we state the following theorem:

Theorem 3.1.22. ([Ser85b, Theorem A]) *Let ζ_α be a geodesic ray in $\mathbb{H}$, starting at the y-axis and terminating at a point $\alpha \in \mathbb{R}_{>0}$. Then, if $(\zeta_\alpha, \mathcal{F}) = L^{n_0} R^{n_1} \cdots$ is the cutting sequence of ζ_α relative to $\mathcal{F}$, the continued fraction expansion of α is given by $\overline{\alpha} = [n_0; n_1, \ldots]$, for $n_0 \in \mathbb{N} \cup \{0\}$ and $n_i \in \mathbb{N}$ otherwise.*

This theorem is the foundation of our work on multiplying continued fractions by an integer, since it allows us to encode continued fractions as cutting sequences in a very natural way.

Section 3.2 is predominantly novel material, with the exception of section 3.2.2 which recaps the work in [Kul91]. We begin by defining the $\frac{1}{n}$-scaled Farey Tessellation, $\frac{1}{n}\mathcal{F}$. This is given by $\frac{1}{n}\mathcal{F} := (n^*)^{-1}(\mathcal{F})$, where $n^* := \begin{pmatrix} \sqrt{n} & 0 \\ 0 & \frac{1}{\sqrt{n}} \end{pmatrix} \in PSL_2(\mathbb{R})$. We then note that if ζ_α is any geodesic ray starting at the y-axis and terminating at the point $\alpha \in \mathbb{R}_{>0}$, and $\zeta_{n\alpha}$ is a geodesic ray starting at the y-axis and terminating at the point $n\alpha$, then the cutting sequence $(\zeta_\alpha, \frac{1}{n}\mathcal{F})$ is equivalent to the cutting sequence $(n^*(\zeta_\alpha), n^*\left(\frac{1}{n}\mathcal{F}\right)) = (n^*(\zeta_\alpha), \mathcal{F})$. This is in turn equivalent to the cutting sequence $(\zeta_{n\alpha}, \mathcal{F})$. In particular, $(\zeta_\alpha, \mathcal{F})$ corresponds to the the continued fraction $\overline{\alpha}$ and $(\zeta_\alpha, \frac{1}{n}\mathcal{F})$ corresponds to the continued fraction expansion of $\overline{n\alpha}$. Since $\alpha \in \mathbb{R}_{>0}$ was arbitrarily chosen, replacing the triangulation $\mathcal{F}$ by the triangulation $\frac{1}{n}\mathcal{F}$ represents integer multiplication of continued fractions. However, this process is not algorithmic.

To construct such an algorithm, we will observe that for any fundamental domain P_n of $\Gamma_0(n)$, we can define $T_{\{1,n\}}$ to be $P_n \cap \mathcal{F}$ and $T_{\{n,n\}}$ to be $P_n \cap \frac{1}{n}\mathcal{F}$. The action of $\Gamma_0(n)$ on the fundamental domain P_n tessellates the plane $\mathbb{H}$. Moreover, the action of $\Gamma_0(n)$ on the locally defined $T_{\{1,n\}}$ and $T_{\{n,n\}}$ recovers the triangulations $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$, respectively. More specifically, $\Gamma_0(n) \cdot T_{\{1,n\}} = \mathcal{F}$ and $\Gamma_0(n) \cdot T_{\{n,n\}} = \frac{1}{n}\mathcal{F}$. Therefore, in order to describe the structure of $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ simultaneously, it is sufficient to describe their structure inside some fundamental domain P_n of $\Gamma_0(n)$. It is ultimately this property that allows us to construct the integer multiplication algorithm. In Section 3.2.2, we take a brief hiatus from our own work to describe the work done by Kulkarni in [Kul91]. Here, we discuss how Kulkarni uses the structure of $\mathcal{F}$ to produce fundamental domains of Φ, where Φ is any subgroup of $PSL_2(\mathbb{Z})$. In particular, one can use this work to find fundamental domains of $\Gamma_0(n)$. Using the fundamental domains as described in Kulkarni's work, we formally outline the construction of the integer multiplication algorithm in Section 3.2.3. This Chapter ends with examples of how to construct and implement the algorithm for $n = 2$ and 3.

We begin Chapter 4 by generalising the notion of cutting sequences to include paths intersecting triangulations of a quotient space $\Phi\backslash\mathbb{H}$, where Φ is a finite index subgroup of $PSL_2(\mathbb{Z})$. To do this, we define a triangulation of this quotient space and prove that if an ideal triangulation T of $\mathbb{H}$ is invariant under the action of Φ - i.e. $\Phi \cdot T = T$ - then T induces a triangulation $\widehat{T}$ of $\Phi\backslash\mathbb{H}$. We then discuss the properties of cutting sequences relative to $\Phi\backslash\mathbb{H}$ and compare these properties to cutting sequences relative $\mathbb{H}$. Since the triangulations $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ are both preserved under the action of $\Gamma_0(n)$, these triangulations induce the triangulations $\widehat{\mathcal{F}}$ and $\widehat{\frac{1}{n}\mathcal{F}}$ of $\Gamma_0(n)\backslash\mathbb{H}$. In particular, the following theorem highlights how we can view integer multiplication of a continued fraction as a triangulation replacement on $\Gamma_0(n)\backslash\mathbb{H}$.

Theorem 4.1.8. *For every geodesic ray ζ_α in $\mathbb{H}$ starting at the y-axis with endpoint $\alpha > 0$, there is a canonical projection $\widehat{\zeta_\alpha}$ onto $\Gamma_0(n)\backslash\mathbb{H}$ such that $(\zeta_\alpha, \mathcal{F}) = (\widehat{\zeta_\alpha}, \widehat{\mathcal{F}})$ and $(\zeta_\alpha, \frac{1}{n}\mathcal{F}) = (\widehat{\zeta_\alpha}, \widehat{\frac{1}{n}\mathcal{F}})$.*

In Section 4.2, we discuss how if λ is a path on $\Phi\backslash\mathbb{H}$, then the cutting sequence

(λ, T) satisfies certain properties if and only if the path λ satisfies certain properties, where T is any triangulation of $\Phi\backslash\mathbb{H}$. For example, we will show how the cutting sequence (λ, T) is periodic if and only if λ is homotopic to a path that goes around a closed curve infinitely often. Similarly, the cutting sequence (λ, T) will be eventually recurrent - see Section 2.1.2 for definition - if and only if λ satisfies some geometric property - which we refer to as *geometric recurrence*. Since these geometric properties are independent of triangulation T, we get the following corollary:

Corollary 4.2.4. *Let $\Phi\backslash\mathbb{H}$ be an orbifold and let ζ be a geodesic ray in $\Phi\backslash\mathbb{H}$, starting at some arc γ_ζ. If there is some triangulation T containing the arc γ_ζ such that (ζ, T) is eventually recurrent/periodic, then (ζ, T') is eventually recurrent/periodic, where T' is any other triangulation of $\Phi\backslash\mathbb{H}$ containing the starting arc γ_ζ.*

By Theorem 4.1.8, we can represent integer multiplication of a continued fraction by n as the triangulation replacement of $\widehat{\mathcal{F}}$ by $\widehat{\frac{1}{n}\mathcal{F}}$ on the space $\Gamma_0(n)\backslash\mathbb{H}$. We can then use Corollary 4.2.4 deduce that if the continued fraction expansion of α is eventually recurrent, then the continued fraction expansion of $n\alpha$ will also be eventually recurrent for all $n \in \mathbb{N}$, since the corresponding cutting sequences $(\widehat{\zeta_\alpha}, \widehat{\mathcal{F}})$ and $(\widehat{\zeta_\alpha}, \widehat{\frac{1}{n}\mathcal{F}})$ will be eventually recurrent. The continued fraction expansion α will also be eventually recurrent if we add an arbitrary integer, - i.e. if we take $\beta = \alpha + k$ for some $k \in \mathbb{N}$ - or if we invert α - i.e. if we take $\beta = \frac{1}{\alpha}$. Combining together all of these operations, we end the chapter with the following corollary:

Corollary 4.2.6. *Let $\alpha \in \mathbb{R}$, let $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ be a non-trivial integer matrix (i.e. $a, b, c, d \in \mathbb{Z}$, $ad - bc \neq 0$), and let $\beta = M \cdot \alpha = \frac{a\alpha+b}{c\alpha+d}$. If the continued fraction expansion $\overline{\alpha}$ is eventually recurrent and $c\alpha + d \neq 0$, then the continued fraction $\overline{\beta}$ is eventually recurrent.*

In Chapter 5 we introduce the notion of an infinite loop mod n. These objects have geometric origins, but can be nicely expressed as the positive real numbers, which have no semi-convergent denominators divisible by n. We begin by discussing the properties of infinite loops and describing why these objects are interesting. Using these properties, we see that infinite loops behave very badly with regards to the integer multiplication algorithm. On the other hand, if α is not an infinite loop mod n, then we can deduce some nice properties of $n\alpha$. In particular, if α is not an infinite loop mod n, then we can guarantee that the partial quotients of $B(\alpha)$ and $B(n\alpha)$ can not both be simultaneously small relative to $\sqrt{n}$. More precisely, we have the following lemma:

Lemma 5.1.10. *Assume that α is not an infinite loop mod n. Then we have:*

$$\max\{B(\alpha), B(n\alpha)\} \geq \lfloor 2\sqrt{n} \rfloor - 1.$$

We will use this property to show that if there is some sequence of natural numbers $\{\ell_m\}_{m\in\mathbb{N}}$ such that $p^{\ell_m}\alpha$ is not an infinite loop mod p^m, then α satisfies pLC. This is summarised in the following corollary:

Corollary 5.1.12. *Let $\alpha \in$ **Bad** and assume there is some sequence of natural numbers $\{\ell_m\}_{m\in\mathbb{N}}$ such that $p^{\ell_m}\alpha$ is not an infinite loop mod p^m. Then α satisfies pLC.*

The reverse to this statement is also true. In particular, we have the following lemma:

Lemma 5.1.14. *Let $\alpha \in$ **Bad** and assume there exists an $m \in \mathbb{N}$ such that $p^{\ell}\alpha$ is an infinite loop mod p^m for all $\ell \in \mathbb{N}$. Then α is a counterexample to pLC.*

Combining these statements together, we get the following reformulation of pLC in terms of infinite loops mod n.

Theorem 5.1.15. *Let $\alpha \in$ **Bad**. Then, α satisfies pLC if and only if there is a sequence of natural numbers $\{\ell_m\}_{m\in\mathbb{N}}$ such that $p^{\ell_m}\alpha$ is not an infinite loop mod p^m.*

We end this chapter by discussing two algorithms one could use to construct infinite loops mod n, for any $n \in \mathbb{N}$. One of these algorithms is geometric in nature and the other is arithmetic. Due to the reformulation of the p-adic Littlewood Conjecture in terms of infinite loops mod n, we hope that these algorithms will be useful in further investigating the set of potential counterexamples to pLC.

In the final chapter, we discuss the main results and findings of this thesis, as well as the ways that this work could be developed in the future.

Chapter 2

Motivation and Context

The main focus of this chapter is to introduce the motivation for the main question of this thesis: "*How do continued fraction expansions behave when multiplied by an integer?*" This question has strong ties to several open problems in Diophantine approximation, such as the mixed and p-adic Littlewood Conjectures. In order to introduce these conjectures, we will briefly discuss the history of continued fractions and a range of the classical results pertaining to continued fractions. We will also discuss why the mixed and p-adic Littlewood conjectures are interesting and outline some of the more important work done in this area.

2.1 The History of Continued Fractions and Important results

2.1.1 A Brief History of Continued Fractions

We begin by briefly discussing the history of continued fractions. The main source used for this overview is [Bre80], but we cite other sources when applicable.

The origins of continued fractions are often attributed to Euclid's algorithm (c. 306-283BCE). This algorithm is a process to find the greatest common divisor of two integers p and q, and slight modification of this algorithm can be used to determine the continued fraction expansion of $\frac{p}{q}$. However, it is unlikely that Euclid or his contemporaries used the algorithm in this way. Following Euclid, there were several notable instances of convergents of continued fractions being implicitly used in mathematics. For example, Āryabhata the elder (c. 475–550CE) detailed a process he called "kuttaka" to find integer solutions x and y to indeterminate equations of the form $ax + by = c$, for a, b, c fixed variables in $\mathbb{Z}$. Convergents were also used to approximate square roots of integers and to find solutions to *Pell equations*. These are equations of the form $x^2 - Dy^2 = 1$, where D is a fixed positive square-free integer and x and y are positive integers variables. Theon of Smyrna (c. 70–135CE)[Hat07] found solutions to the equation $d^2 - 2a^2 = \pm 1$ and noted that the solutions gave "good approximations" of $\sqrt{2}$.

However, it was not until the 17^{th} century that the phrase *continued fraction* was first introduced by John Wallis. Wallis explicitly investigated continued fractions and their general properties in the book *Arithmetica Infinitorum*, which was published in 1655. However, it was the 18^{th} century, in which the theory of continued fractions flourished into the topic we know today. In particular, Euler's book *De Fractionibus Continius* laid the groundwork for much of the modern theory of continued fractions. In this book, Euler proved that rational numbers correspond to finite continued fractions, irrational numbers correspond to infinite continued fractions and that e is irrational. He also gave a proof that eventually periodic continued fractions correspond to quadratic irrationals - however, he did not show that the reverse is true. In the latter half of the 18^{th} century, Lambert expressed $\tan(x)$ as a generalised continued fraction and used this to show that $\frac{\pi}{4}$ is irrational, and, by extension, so is π. In 1766, Lagrange proved that Pell's equation has infinitely many integer solutions for D a square-free positive integer. In 1767 and 1768, Lagrange published two papers which completed the identification of quadratic irrationals and periodic continued fractions, by showing that quadratic irrationals all have eventually periodic continued fraction expansions.

2.1.2 Classical Results Surrounding Continued Fractions

In this section, we will discuss several classical results pertaining to (simple) continued fractions. These results can be found in a number of introductory texts, but the main sources we use are [Khi63] and chapter 10 of [HW38]. We start by defining simple continued fractions.

Definition 2.1.1. A (simple) *continued fraction* $\overline{\alpha}$ is an expression of the form

$$\overline{\alpha} := a_0 + \cfrac{1}{a_1 + \cfrac{1}{a_2 + \cfrac{1}{\dots}}},$$

where $a_0 \in \mathbb{Z}$ and $a_i \in \mathbb{N}$ for $i \geq 1$.

Throughout this thesis, we will solely look at simple continued fractions, which we will from now on just refer to as "continued fractions" for brevity.

We will usually identify continued fractions with their sequence of a_i's, $\overline{\alpha} = [a_0; a_1, \dots]$ and refer to the a_i's as *partial quotients*. We note that explicit evaluation of the continued fraction expansion produces a real number α. Similarly, for any real number α we can find an associated continued fraction expansion by using Euclid's algorithm:

Euclid's Algorithm for Computing Continued Fractions.

Let $\alpha \in \mathbb{R}$ and let $\overline{\alpha}$ be an empty list.

1. Set $i = 0$ and $\alpha_0 = \alpha$:
2. While $\alpha_i \neq 0$:
    (a) Find $a_i \in \mathbb{N}$ and $r_i \in [0, 1]$ such that $\alpha_i = a_i + r_i$.
    (b) Append a_i to the list of partial quotients $\overline{\alpha}$.
    (c) If $r_i \neq 0$:

- Set $\alpha_{i+1} = \frac{1}{r_i}$.

Else:

- Set $\alpha_{i+1} = 0$.

3. End of algorithm.

We note that for $\alpha \in \mathbb{Q}$, the above process terminates. In particular, for $\alpha \in \mathbb{Q} \setminus \{0\}$ there is some $k \in \mathbb{N}$ such that $\alpha_k \in \mathbb{N}$. However, since our *remainder* r_k lies in the interval $[0, 1]$, this results in two equivalent continued fraction expansions given by $[a_0; a_1, \ldots, a_k]$ and $[a_0; a_1, \ldots, a_k - 1, 1]$, where $a_k > 1$. For $\alpha \in \mathbb{R} \setminus \mathbb{Q}$, the above process does not terminate and so there is no value $k \in \mathbb{N} \cup \{0\}$ with $\alpha_k \in \mathbb{N}$. In particular, for $\alpha \in \mathbb{R} \setminus \mathbb{Q}$ there is a unique infinite continued fraction expansion.

Definition 2.1.2. Let $\overline{\alpha} = [a_0; a_1, a_2, \ldots]$ be a continued fraction. We define the *k-th convergent* of $\overline{\alpha}$ to be $\frac{p_k}{q_k} := [a_0; a_1, \ldots, a_k]$. We can define this iteratively where:

$$\begin{aligned} p_{-1} &= 1 & p_0 &= a_0 & p_k &= a_k p_{k-1} + p_{k-2} \\ q_{-1} &= 0 & q_0 &= 1 & q_k &= a_k q_{k-1} + q_{k-2} \end{aligned}$$

We refer to the term p_k as the *k-th convergent numerator* of α and q_k as the *k-th convergent denominator*.

Remark 2.1.3. It is worth noting that since $a_i \in \mathbb{N}$ for $i \geq 1$ and $a_0 \in \mathbb{Z}$, only the initial partial quotient a_0 can be negative. From this set-up, the convergent denominators will always be non-negative and the convergent numerators (possibly excluding p_{-1}) will have the same sign as α.

Continued fractions as Good Rational Approximations

One of the main reasons why continued fractions have been so well-studied, is that the convergents of a real number $\alpha \in \mathbb{R}$ are very "good rational approximations" of α. This notion of a "good rational approximation" takes several different forms.

Given a real number α, we define a rational number $\frac{p}{q}$ to be the *best rational approximation of the first type*, if for every $n \in \mathbb{N} \cup \{0\}$ and $d \in \mathbb{N}$ with $d \leq q$, we have:

$$\left|\alpha - \frac{p}{q}\right| \leq \left|\alpha - \frac{n}{d}\right|.$$

This is equivalent to saying that there is no rational number with a denominator smaller than q, which is closer to α. We similarly define $\frac{p}{q}$ to be the *best rational approximation of the second type*, if for every $n \in \mathbb{N} \cup \{0\}$ and $d \in \mathbb{N}$ with $d \leq q$, we have:

$$|q\alpha - p| \leq |d\alpha - n|.$$

Here, it will be useful to introduce the *distance to the nearest integer function* $\|\cdot\| : \mathbb{R} \to \left[0, \frac{1}{2}\right]$, which is the function given by:

$$\|\alpha\| = \min\left\{|\alpha - n| : n \in \mathbb{Z}\right\}.$$

It is easy to see that for a rational $\frac{p}{q}$ to be a best approximation of the second type, we must necessarily have that:

$$\|q\alpha\| = |q\alpha - p|.$$

Otherwise, this would imply that there is some $n \in \mathbb{Z}$ with $n \neq p$, such that:

$$\|q\alpha\| = |q\alpha - n| < |q\alpha - p|.$$

In this case, $\frac{n}{q}$ would be a better rational approximation than $\frac{p}{q}$. We note that if $\frac{p_k}{q_k}$ is a convergent of α, then:

$$\|q_k\alpha\| = |q_k\alpha - p_k|.$$

Therefore, the convergents of α satisfy this necessary condition. In fact, the following theorem shows that the best rational approximations of the second type are precisely the convergents of α:

Theorem 2.1.4. ([Khi63, Theorem 16 and 17]) *Every convergent is a best approximation of the second kind and every best approximation of the second kind is a* *convergent.*

Remark 2.1.5. There is one trivial exception to this theorem, which is when $\alpha = a + \frac{1}{2}$ for $a \in \mathbb{Z}$. Here, both $\frac{a}{1}$ and $\frac{a+1}{1}$ give equally good approximations of α, but only one of them will be a convergent (if $a \in \mathbb{Z}_{\geq 0}$, then $\frac{p_0}{q_0} = a$, and $\frac{p_0}{q_0} = a + 1$, otherwise).

With a bit of work we can see that if $\frac{p}{q}$ is a best approximation of the second kind, then $\frac{p}{q}$ is a best approximation of the first kind. In particular, for all $n \in \mathbb{N} \cup \{0\}$ and $d \in \mathbb{N}$ with $d \leq q$ we have the following:

$$q\left|\alpha - \frac{p}{q}\right| = |q\alpha - p| \leq |d\alpha - n| = d\left|\alpha - \frac{n}{d}\right| \leq q\left|\alpha - \frac{n}{d}\right|,$$

and so:

$$\left|\alpha - \frac{p}{q}\right| \leq \left|\alpha - \frac{n}{d}\right|.$$

Therefore, we can conclude that every convergent is also a best approximation of the first kind. However, not every best approximations of the first kind is a convergent. In order to get all best approximations of the first kind, we first must generalise the notion of a convergent. To do this, we define the *semi-convergents* of α, also known as the *intermediate convergents* or *secondary convergents.*

Definition 2.1.6. Let $\overline{\alpha} = [a_0; a_1, a_2, \dots]$ be a continued fraction expansion of some real number α. We define the $\{k, m\}$*-th semi-convergent* of $\overline{\alpha}$ to be $\frac{p_{\{k,m\}}}{q_{\{k,m\}}} := [a_0; a_1, \dots, a_k, m]$, where $0 \leq m \leq a_{k+1}$. We can define this iteratively using the standard convergents:

$$p_{\{k,m\}} = mp_k + p_{k-1},$$

$$q_{\{k,m\}} = mq_k + q_{k-1}.$$

We refer to the term $p_{\{k,m\}}$ as the $\{k, m\}$*-th semi-convergent numerator* of α and $q_{\{k,m\}}$ as the $\{k, m\}$*-th semi-convergent denominator.*

We can view the k-th convergents as truncations of the continued fraction expansion $\overline{\alpha}$ after the k-th partial quotient. The $\{k, m\}$-th semi-convergents can similarly be viewed as the truncation of $\overline{\alpha}$ which occurs m integers into the $(k+1)$-th partial quotient. By construction, the $\{k,0\}$-th semi-convergent of α will be the $(k-1)$-th convergent of α. Similarly, the $\{k, a_{k+1}\}$-th semi-convergent of α will be the $(k+1)$-th convergent of α. Using the notion of a semi-convergent, one can obtain the following result:

Theorem 2.1.7. ([Khi63, Theorem 15]) *The fraction $\frac{n}{d}$ is a best approximation to a real number α of the first kind if and only if it is a k-th convergent or a $\{k, m\}$-th semi-convergent of α, where either:*

- $m = \frac{a_{k+1}}{2}$ *and* $[a_{k+1} : a_k, \ldots, a_0] > [a_{k+1}; a_{k+2}, \ldots]$, *or*
- $m > \frac{a_{k+1}}{2}$.

Convergents can also be viewed as good rational approximations of real numbers in an alternative way. In particular, the difference between α and one of its convergents is very small relative to the size of the convergent denominator. This concept is emphasised by the following theorem:

Theorem 2.1.8. ([HW38, Theorem 163 and 171], [Khi63, Theorem 9]) *The k-th convergent of a real number α satisfies the following property:*

$$\left|\alpha - \frac{p_k}{q_k}\right| < \frac{1}{q_k q_{k+1}} < \frac{1}{a_{k+1} q_k^2} \leq \frac{1}{q_k^2}.$$

Using this notion of a good approximation, it is natural to ask whether or not any rational number $\frac{n}{d}$ satisfying $\left|\alpha - \frac{n}{d}\right| < \frac{1}{d^2}$ is a convergent of α. However, this is not quite true. If we have a slightly stronger condition - that $\left|\alpha - \frac{n}{d}\right| < \frac{1}{2d^2}$ - then $\frac{n}{d}$ is a convergent of α.

Legendre's Theorem. ([Khi63, Theorem 19], [HW38, Theorem 184]) *Every irreducible rational number $\frac{n}{d}$ satisfying $\left|\alpha - \frac{n}{d}\right| < \frac{1}{2d^2}$ is a convergent of α.*

For α an irrational number, α has infinitely many partial quotients, and therefore, has infinitely many convergents. One can conclude from Theorem 2.1.8 that:

$$\left|\alpha - \frac{p}{q}\right| < \frac{1}{q^2},$$

for infinitely many $p \in \mathbb{Z}$ and $q \in \mathbb{N}$ with $gcd(p,q) = 1$. In fact, one can say something even stronger:

Hurwitz's Theorem. ([Khi63, Theorem 21]) *For every $\alpha \in \mathbb{R} \setminus \mathbb{Q}$, there are infinitely many $p \in \mathbb{Z}$ and $q \in \mathbb{N}$ with $gcd(p,q) = 1$ such that:*

$$\left|\alpha - \frac{p}{q}\right| < \frac{1}{\sqrt{5}q^2}.$$

Hurwitz's Theorem shows that for every $\alpha \in \mathbb{R} \setminus \mathbb{Q}$ and $\nu = \frac{1}{\sqrt{5}}$, we have:

$$\left|\alpha - \frac{p}{q}\right| < \frac{\nu}{q^2},$$

for infinitely many $p \in \mathbb{Z}$ and $q \in \mathbb{N}$ with $gcd(p,q) = 1$. From this, it is natural to ask if $\nu = \frac{1}{\sqrt{5}}$ is the best constant for this upper bound. In fact, for a general $\alpha \in \mathbb{R}$, $\nu = \frac{1}{\sqrt{5}}$ is the best that we can do. For example, the golden ratio $\varphi := \frac{1+\sqrt{5}}{2}$ only has finitely many rationals $\frac{p}{q}$ which satisfy $\left|\alpha - \frac{p}{q}\right| < \frac{\nu}{q^2}$, when $\nu < \frac{1}{\sqrt{5}}$.

In a similar vein of questioning, we can ask "What is the largest value of $c \in \mathbb{R}$ such that:

$$\frac{c}{q^2} < \left|\alpha - \frac{p}{q}\right|,$$

for every rational $\frac{p}{q}$?". We will use the function $c(\alpha)$ to denote the largest value of c in this case. We can rearrange the equation and take the infimum, to see that:

$$\begin{aligned} c(\alpha) &= \inf_{\frac{p}{q} \in \mathbb{Q}} \left\{ q^2 \left|\alpha - \frac{p}{q}\right| \right\} \\ &= \inf_{\frac{p}{q} \in \mathbb{Q}} \left\{ q \left| q\alpha - q\frac{p}{q} \right| \right\} \\ &= \inf_{\frac{p}{q} \in \mathbb{Q}} \left\{ q \left| q\alpha - p \right| \right\} \\ &= \inf_{q \in \mathbb{N}} \left\{ q \| q\alpha \| \right\}. \end{aligned}$$

The function $c(\alpha)$ gives some kind of idea of how well an irrational number is approximated by the rationals. For example, if $c(\alpha) = 0$, then for every $\epsilon > 0$ there is some rational number $\frac{p}{q}$ which satisfies:

$$\left|\alpha - \frac{p}{q}\right| < \frac{\epsilon}{q^2}.$$

However, if $c(\alpha) > \epsilon > 0$, then for every rational number $\frac{p}{q}$ we have:

$$\frac{\epsilon}{q^2} < \left|\alpha - \frac{p}{q}\right|.$$

We refer to the class of real numbers with $c(\alpha) > 0$ as the set of *badly approximable numbers* (also referred to as the *badly approximables*), denoted **Bad**. Explicitly, we have that:

$$\mathbf{Bad} := \left\{\alpha \in \mathbb{R} : \inf_{\frac{p}{q} \in \mathbb{Q}} \left\{q^2 \left|\alpha - \frac{p}{q}\right|\right\} > 0\right\}.$$

By Hurwitz's Theorem, we know that $c(\alpha) < \frac{1}{\sqrt{5}}$ for all real α. Since $\frac{1}{\sqrt{5}} < \frac{1}{2}$, we can use Legendre's Theorem to see that the rationals $\frac{p}{q}$ which minimise $q^2 \left|\alpha - \frac{p}{q}\right|$ must be the convergents of α. As a result, we can rewrite $c(\alpha)$ in terms of the convergents of α:

$$c(\alpha) = \inf_{k \in \mathbb{N}} \{q_k \|q_k \alpha\|\}.$$

Therefore, we can rewrite the set **Bad** in terms of the convergents of α:

$$\mathbf{Bad} := \left\{\alpha \in \mathbb{R} : \inf_{k \in \mathbb{Q}} \{q_k \|q_k \alpha\|\} > 0\right\}.$$

We can similarly define **Bad** by using a condition on the partial quotients of the continued fraction expansions. In order to do this, we introduce the following theorem which gives us a lower bound on the value of $|\alpha - \frac{p_k}{q_k}|$ in terms of the partial quotients:

Theorem 2.1.9. ([Khi63, Theorem 13]) *The k-th convergent of a real number α satisfies the following property:*

$$\frac{1}{(a_{k+1} + 2)q_k^2} < \frac{1}{q_k(q_{k+1} + q_k)} < \left|\alpha - \frac{p_k}{q_k}\right|.$$

Combining together Theorem 2.1.8 and Theorem 2.1.9, we see that:

$$\frac{1}{a_{k+1}+2} < q_k\|q_k\alpha\| < \frac{1}{a_{k+1}}.$$

As a result, we can bound $c(\alpha)$ in terms of the partial quotients of α:

$$\inf_{k\in\mathbb{N}\cup\{0\}}\left\{\frac{1}{a_{k+1}+2}\right\} < c(\alpha) < \inf_{k\in\mathbb{N}\cup\{0\}}\left\{\frac{1}{a_{k+1}}\right\}.$$

Here, we see that for $\alpha\in\mathbb{R}\setminus\mathbb{Q}$, we have $c(\alpha)>0$ if and only if $\inf\limits_{k\in\mathbb{N}}\left\{\frac{1}{a_k}\right\}>0$. We note that $\inf\limits_{k\in\mathbb{N}}\left\{\frac{1}{a_k}\right\}>0$ if and only if $\sup\limits_{k\in\mathbb{N}}\{a_k\}<\infty$, i.e. partial quotients are bounded. We define the *height function* $B(\alpha)$ to be the largest partial quotient of α, i.e.:

$$B(\alpha)=\sup_{k\in\mathbb{N}}\{a_k : \overline{\alpha}=[a_0;a_1,a_2\ldots]\}.$$

This allows us to redefine **Bad** in terms of the height function of α:

$$\mathbf{Bad} := \{\alpha\in\mathbb{R}\setminus\mathbb{Q} : B(\alpha)<\infty\}.$$

Remark 2.1.10. Note that in the definition of the height function, we ignore the term $k=0$ and only look at $k\in\mathbb{N}$. This is because for $\{\alpha\}=\alpha-\lfloor\alpha\rfloor$, we have $\|q\alpha\|=\|q\{\alpha\}\|=\|\{q\alpha\}\|$, and so it is natural drop the integer part of α.

Here we should also note the link between the constant $c(\alpha)$ and the *Markov constant* $\mu(\alpha)$, where:

$$\mu(\alpha) := \liminf_{q\to\infty}\{q\|q\alpha\|\}.$$

We can view the Markov constant as the smallest value of u such that

$$\left|\alpha-\frac{p}{q}\right|<\frac{u}{q^2}$$

is satisfied for infinitely many rational numbers $\frac{p}{q}$. Occasionally, the badly approximable numbers are defined as the set of real numbers $\alpha\in\mathbb{R}$ satisfying $\mu(\alpha)>0$ [BV11]. However, as seen in [Bur00, Theorem 7.4], we see that $c(\alpha)=0$ if and only if $\mu(\alpha)=0$ and so these definitions are equivalent. See [Bur00] for more information about $c(\alpha)$ and $\mu(\alpha)$.

There are a few reasons why we chose to define the badly approximable numbers using the constant $c(\alpha)$, instead of the Markov constant $\mu(\alpha)$. The most notable reason why we do this is that the constant $c(\alpha)$ corresponds to a nicer definition of infinite loops in Chapter 5. This in turn results in nicer statements of Corollary 5.1.12 and Lemma 5.1.14.

Finally, we note that if $\alpha \in \mathbf{Bad}$, then so is $n\alpha$ for all $n \in \mathbb{N}$. Since α is an element of **Bad**, we know that there is some $\varepsilon > 0$ such that:

$$q \cdot \|q\alpha\| > \varepsilon,$$

for all $q \in \mathbb{N}$. As a result, it follows that for a fixed n, we have:

$$nq \cdot \|nq\alpha\| > \varepsilon,$$

for all $q \in \mathbb{N}$. Finally, this implies that:

$$q \cdot \|q(n\alpha)\| > \frac{\varepsilon}{n},$$

for all $q \in \mathbb{N}$, and the result follows.

Continued Fraction Expansions as Words

In order to discuss some of the results pertaining to the Littlewood-type conjectures, it will often be useful to view continued fractions as their associated *words*.

Let $\mathbb{A}$ be some set, then a *finite word* w (in $\mathbb{A}$) of *length* N is a composition of N elements in $\mathbb{A}$. The set $\mathbb{A}$ is referred to as the *alphabet* and the elements of this set are referred to as *letters*. The set of all words of length N is denoted by $\mathbb{A}^N$ and the set of all finite word will be denoted by $\mathbb{A}^*$. An *infinite word* is a composition of countably many elements in $\mathbb{A}$. The set of all infinite words is denoted $\mathbb{A}^{\mathbb{N}}$. The left shift map $T : \mathbb{A}^{\mathbb{N}} \to \mathbb{A}^{\mathbb{N}}$ is the map induced by removing the first letter from the word. For example, if $w = w_0w_1w_2\ldots$, then $Tw = w_1w_2\ldots$.

A *sub-word* of $w = w_1w_2w_3\ldots$ is any finite word u which appears at least once in w, i.e. there is some $i, k \in \mathbb{N}$ with $u = w_{i+1} \dots w_{i+k}$. Given a word $w = w_1 w_2 w_3 \dots$ of length $n \in \mathbb{N} \cup \{\infty\}$, a *prefix* is any sub-word $p = w_1 w_2 \dots w_m$ formed by truncating w after the first m terms, where $m < n$. If w is a finite word, then a *suffix* is any sub-word v of w such that $w = uv$. If w is an infinite word, then a *tail* is an infinite word v such that $w = uv$, where u is a prefix. [AS03].

Given a real number $\alpha \in \mathbb{R}$, we define $w_{CF}(\alpha)$ to be the (potentially infinite) word formed by iteratively appending the non-zero partial quotients of the continued fraction expansion of α [Bad15], i.e.:

$$w_{CF}(\alpha) = \begin{cases} a_0 a_1 a_2 \dots & \text{if } \overline{\alpha} = [a_0; a_1, a_2, \dots] \\ a_1 a_2 a_3 \dots & \text{if } \overline{\alpha} = [0; a_1, a_2, a_3, \dots] \end{cases}$$

For $\alpha \in \mathbb{R} \setminus \mathbb{Q}$, we see that that $w_{CF}(\alpha) \in \mathbb{N}^{\mathbb{N}}$. However, if $\alpha \in \mathbf{Bad}$ there exists some $K \in \mathbb{N}$ such that $B(\alpha) \leq K$, and so, in this case, $w_{CF}(\alpha) \in \mathbb{A}_K^{\mathbb{N}}$ where $\mathbb{A}_K := \{1, \dots, K\}$. This K is not unique, but we will often take K to be the minimum possible value, i.e. $K = B(\alpha)$. We will typically only be interested in words corresponding to infinite continued fractions and the properties that these words have. The two main properties we will discuss are *periodicity* and *recurrency*.

An infinite word $w = w_1 w_2 \dots$ is *strictly periodic* if there exists an $m \in \mathbb{N}$ such that $w_i = w_{m+i}$ for all $i \in \mathbb{N}$. We will use the notation $\overline{w_1 \dots w_m}$ to indicate that $w_i = w_{m+i}$ for all $i \in \mathbb{N}$. The *period* of w is the smallest m such that $w_i = w_{m+i}$ for all $i \in \mathbb{N}$. An infinite word w is *eventually periodic*, if there exists some $n \in \mathbb{N} \cup \{0\}$ such that $T^n w$ is periodic, i.e. one can write $w = u_1 \dots u_n \overline{w_1 \dots w_m}$.

Theorem 2.1.11 (Euler and Lagrange). *Let $\alpha \in \mathbb{R}$. Then $w_{CF}(\alpha)$ is eventually periodic if and only if α is a quadratic irrational, i.e. α is a real number of the form $\frac{a+b\sqrt{c}}{d}$, where $a, b, c, d \in \mathbb{Z} \setminus \{0\}$ and c is square-free.*

An infinite word w is *strictly recurrent* if every sub-word that appear once in w, appears infinitely often in w. An infinite word w is *eventually recurrent*, if there exists some $n \in \mathbb{N} \cup \{0\}$ such that $T^n w$ is strictly recurrent. Examples of strictly recurrent words include: Sturmian words, Thue-Morse words and periodic words; See [AS03] for more details on each of these objects.

Remark 2.1.12. In general, we will drop the phrase "strictly" when referring to "strictly recurrent" and "strictly periodic" words, and will instead refer to them as being recurrent or periodic, respectively. We will also commonly abuse notation and refer to the continued fraction expansion as recurrent/periodic if the corresponding word is recurrent/periodic.

Given two irrational numbers α and β, we will say that $\overline{\alpha}$ has the same tail as $\overline{\beta}$ if there is some word v which is a tail for both $w_{CF}(\alpha)$ and $w_{CF}(\beta)$. In other words, there exist some $k, \ell \in \mathbb{N} \cup \{0\}$ such that $\overline{\alpha} = [a_0; \dots, a_k, c_0, c_1, \dots]$ and $\overline{\beta} = [b_0; \dots, b_\ell, c_0, c_1, \dots]$. The following theorem tells us exactly when two real numbers will have the same tails:

Theorem 2.1.13. *Let $\alpha, \beta \in \mathbb{R}$ be two irrational numbers. Then, α and β have the same tails if and only if there is some matrix $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in GL_2(\mathbb{Z})$ such that $\beta = \frac{a\alpha+b}{c\alpha+d}$.*

Another useful tool to look at will be the *complexity function*, $P(w_{CF}(\alpha), n)$. The *complexity function* $P(w, n)$ is the function which counts the number of unique sub-words of length n, i.e. $P(w, n) := |\{w_{l+1} \dots w_{l+n} : l \geq 0\}|$, where $w = w_1 w_2 \dots$. For $\alpha \in \mathbb{Q}$, there exists some $r \in \mathbb{N}$ such that $P(w_{CF}(\alpha), n) = 0$ for all $n \geq r$. If $\alpha \in \mathbb{R} \setminus (\mathbb{Q} \cup \mathbf{Bad})$, then $P(w_{CF}(\alpha), n) = \infty$ for all $n \geq 1$. However, if $\alpha \in \mathbf{Bad}$ with height function $B(\alpha) \leq K$ for some $K \in \mathbb{N}$, then $P(w_{CF}(\alpha), n) \leq K^n$. In [MH38], Morse and Hedlund showed that if $\alpha \in \mathbf{Bad}$ and $w_{CF}(\alpha)$ not ultimately periodic, then $P(w_{CF}(\alpha), n) \geq n + 1$ for all $n \geq 1$. In the case that α is ultimately periodic, there exists a $C \in \mathbb{N}$ such that $P(w_{CF}(\alpha), n) \leq C$ for $n \geq 1$.

When the alphabet of a word w is finite, then it follows by the definition of the complexity function that $P(w, n+m) \leq P(w, n)P(w, m)$ for all $n, m \in \mathbb{N}$. This also implies that $\log(P(w, n+m)) \leq \log(P(w, n)) + \log(P(w, m))$, and so $\log(P(w, n+m))$ is subadditive. In particular, for $\alpha \in \mathbf{Bad}$, the limit $E(\alpha) := \lim_{n\to\infty} \frac{\log(P(w_{CF}(\alpha), n))}{n}$ exists and is finite. We refer to the value $E(\alpha)$ as the *entropy* of α.

2.2 The Littlewood-type Conjectures

2.2.1 The Littlewood Conjecture

The Littlewood conjecture (LC) is a classical open problem in Diophantine approximation. According to Montgomery in Littlewood's obituary [Mon79], the problem was first (officially) stated by Spencer - one of Littlewood's students - in [Spe42] in 1942.

The Littlewood Conjecture. *For any pair of real numbers numbers $(\alpha, \beta) \in \mathbb{R}^2$, we have:*

$$\liminf_{q \in \mathbb{N}} \{q \cdot \|q\alpha\| \cdot \|q\beta\|\} = 0.$$

It follows from the definition of the badly approximable numbers

$$\mathbf{Bad} := \left\{ x \in \mathbb{R} : \inf_{q \in \mathbb{N}} q \cdot \|qx\| > 0 \right\},$$

that if either α or β do not lie in the set of **Bad**, then the pair (α, β) must satisfy the Littlewood conjecture.

The first major result for the Littlewood Conjecture was due to Cassels and Swinnerton-Dyer in 1955, where they showed that if the pair (α, β) each belong to the same cubic field, then they satisfy the Littlewood Conjecture [CSD55]. Note that it is still an open problem in Diophantine approximation, whether general algebraic numbers, and by extension cubics, are badly approximable or not.

The next major result was due to Pollington and Velani in 2000 [PV00]. In this paper they showed the following result:

Theorem 2.2.1 (Pollington and Velani [PV00])**.** *Given any $\alpha \in$ **Bad** there exists a subset $G(\alpha)$ of **Bad** with full Hausdorff dimension such that for all $\beta \in G(\alpha)$ and infinitely many $q \in \mathbb{N}$, we have:*

$$q \cdot \|q\alpha\| \cdot \|q\beta\| < \frac{1}{\log(q)}.$$

If $\{q_k\}_{k\in\mathbb{N}}$ *is the set of convergent denominators of* α*, then the set* $G(\alpha)$ *is given by:*

$$G(\alpha) := \left\{\beta \in \mathbf{Bad} : \|q_k\beta\| < \frac{1}{\log q_k} \text{ for infinitely many } q_k\right\}.$$

Since this theorem holds for infinitely many $q \in \mathbb{N}$, it follows as a corollary that, given any $\alpha \in \mathbf{Bad}$ and $\beta \in G(\alpha)$, the pair (α, β) satisfies the Littlewood Conjecture.

The next big result was due to Einsiedler, Katok and Lindenstruass [EKL06]. In particular, they showed that the set of potential counterexamples to the Littlewood Conjecture is relatively small.

Theorem 2.2.2 (Einsiedler, Katok and Lindenstruass [EKL06])**.** *Let* Θ *be the set of counterexamples to the Littlewood Conjecture, i.e.:*

$$\Theta := \left\{(\alpha, \beta) \in \mathbb{R} : \liminf_{q\in\mathbb{N}} q \cdot \|q\alpha\| \cdot \|q\beta\| > 0\right\}.$$

Then, Θ *has Hausdorff dimension* 0.

2.2.2 The mixed and p-adic Littlewood Conjectures

The *mixed Littlewood Conjecture* (mLC) was first proposed by de Mathan and Teulié in 2004, as a 1-dimensional analogue of the classical Littlewood Conjecture [MT04]. The main purpose of this conjecture was to gain further insight into the Littlewood Conjecture. However, this problem has proved very interesting in its own right, and whilst significant progress has been made, the conjecture remains open. As we will see later, many of the results pertaining to the Littlewood Conjecture have analogues in the mixed/p-adic setting. It is also worth noting that definitive answers for associated problems have been found. In particular, the *t-adic Littlewood Conjecture* - an analogue of pLC over function fields - was proven to be false for $\mathbb{F}_3$ [ANL18]. This provides some credence to the notion that pLC (or mLC) may also be false.

In order to explicitly state the mixed (and p-adic) Littlewood Conjecture, we must first introduce some definitions. Let $\mathcal{C} = (c_k)_{k\in\mathbb{N}}$ be a sequence of integers with $c_k \geq 2$

for all k. Then we set $d_0 = 1$ and $d_k = c_k d_{k-1}$ for all $k \in \mathbb{N}$, i.e. $d_k = c_1 \cdot c_2 \cdot \ldots \cdot c_k$. We refer to any sequence $\mathcal{D} := (d_k)_{k\in\mathbb{N}}$ which can be defined in this way as a *pseudo-absolute sequence*. If we define $v_{\mathcal{D}}(q) := \sup_{n\in\mathbb{N}} \{d_n : d_n \mid q\}$, then the *$\mathcal{D}$-adic norm* (or *pseudo-absolute norm*) is given by:

$$|q|_{\mathcal{D}} := \frac{1}{v_{\mathcal{D}}(q)}.$$

The mixed Littlewood Conjecture is then stated as follows:

The Mixed Littlewood Conjecture 2.2.3. *For every real number $\alpha \in \mathbb{R}$ and every pseudo-absolute sequence $\mathcal{D}$, we have:*

$$m_{\mathcal{D}}(\alpha) := \inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} = 0.$$

Remark 2.2.4. Note that if $\beta = \alpha + k$ for some integer $k \in \mathbb{Z}$, then β satisfies mLC if and only if α satisfies mLC. This follows since:

$$\begin{aligned} m_p(\alpha) &= \inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} \\ &= \inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha + qk\|\} \\ &= \inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q(\alpha + k)\|\} \\ &= \inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\beta\|\} \\ &= m_p(\beta). \end{aligned}$$

Remark 2.2.5. Frequently, the literature will write the condition given in mLC as:

$$\liminf_{q\to\infty} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} = 0,$$

however, both conditions are equivalent. This is obviously true for $\alpha \in \mathbb{Q}$. For $\mathbb{R} \setminus \mathbb{Q}$, we note that $|q|_{\mathcal{D}} > 0$ for $q \in \mathbb{N}$ and $q_k \cdot \|q_k\alpha\| > \frac{1}{a_{k+1}+2}$, for q_k any convergent denominator of α. Therefore, since the convergent denominators give the best (lower) approximation of $q \cdot \|q\alpha\|$, for each fixed $q \in \mathbb{N}$, we have:

$$q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\| > 0.$$

As a result, $\inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\| = 0\}$ if and only if there is some monotonically increasing subsequence $\{n_j\}_{j\in\mathbb{N}}$ such that $\lim_{j\to\infty} \left\{n_j \cdot |n_j|_{\mathcal{D}} \cdot \|n_j\alpha\|\right\} = 0$, i.e.:

$$\liminf_{q\to\infty} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} = 0.$$

It is worth noting that $|q|_{\mathcal{D}} \leq 1$ for every $q \in \mathbb{N}$ and every pseudo-absolute sequence $\mathcal{D}$. It follows that for every $\alpha \notin \mathbf{Bad}$, we have $\inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} = 0$. For reasons which will become more obvious in the next section, we will refer to the set of counterexamples to mLC as the set of *multiplicatively badly approximable numbers* [BV11]. We will denote this set as:

$$\mathbf{Mad}_{\mathcal{D}} := \{\alpha \in \mathbb{R} : m_{\mathcal{D}}(\alpha) > 0\}.$$

When $\mathcal{C}$ is the constant sequence (i.e. every $c_k = a$ for some $a \geq 2$), then we will write $|\cdot|_a$ to mean $|\cdot|_{\mathcal{D}}$, where $\mathcal{D}$ is the corresponding pseudo-absolute sequence. In this case, we have $\mathcal{D} = \left\{1, a, a^2, a^3, \ldots\right\}$. When $a = p$ is a prime, the $\mathcal{D}$-adic norm $|\cdot|_p$ is just the standard *p-adic norm*. For a fixed prime p, we obtain a specific case of the mixed Littlewood conjecture, known as the *p-adic Littlewood Conjecture* (pLC). We state the conjecture as follows:

The p-adic Littlewood Conjecture. *For every real number $\alpha \in \mathbb{R}$, we have:*

$$m_p(\alpha) := \inf_{q\in\mathbb{N}} \left\{q \cdot |q|_p \cdot \|q\alpha\|\right\} = 0.$$

2.2.3 Known Results for mLC and pLC

As part of the paper which introduced the problem [MT04], de Mathan and Teulié managed to show the following:

Theorem 2.2.6 (de Mathan and Teulié [MT04])**.** *Let $\mathcal{D}$ be a pseudo-absolute sequence. Then, if there exists $C \in \mathbb{N}$ such that $\frac{d_{k+1}}{d_k} < C$, every quadratic irrational α satisfies:*

$$\liminf_{q\to\infty} \{q \cdot \log q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} < \infty.$$

In particular, every quadratic irrational satisfies mLC for such sequences.

In 2007, Einsiedler and Kleinbock showed that the set of counterexamples to pLC is very small [EK07].

Theorem 2.2.7 (Einsiedler and Kleinbock [EK07]). *For every pseudo-absolute sequence $\mathcal{D}$, the set of counterexamples $\mathbf{Mad}_{\mathcal{D}}$ has Hausdorff dimension 0.*

In the same year Bugeaud, Drmota and de Mathan showed that if for a given $\alpha \in \mathbf{Bad}$ the corresponding word $w_{CF}(\alpha)$ "limits to a periodic word", then α satisfies pLC (for every p) [BDM07].

Theorem 2.2.8 (Bugeaud, Drmota and de Mathan [BDM07]). *Let $\alpha \in \mathbf{Bad}$ and let $w_{CF}(\alpha) = a_0 a_1 \ldots$ be the word corresponding to the continued fraction expansion of α. Let $T \in \mathbb{N}$ and let $u = \overline{b_1 \ldots b_T}$ be a periodic word in $\mathbb{N}^{\mathbb{N}}$. If there exists an unbounded sequence $(m_k)_{k\in\mathbb{N}}$ such that $w_{CF}(\alpha)$ and u have common sub-words of length m_k for every $k \in \mathbb{N}$, then α satisfies pLC for every p.*

The next big result was due to Harrap and Haynes in 2013, when they showed that a weak form of the mixed Littlewood Conjecture is true [HH13].

Theorem 2.2.9 (Harrap and Haynes [HH13]). *Let $a \geq 2$ and let $\mathcal{D}$ be a pseudo-absolute sequence such that there exists $C \in \mathbb{N}$ with $\frac{d_{k+1}}{d_k} < C$. Then for every $\alpha \in \mathbb{R}$, we have:*

$$\inf_{q\in\mathbb{N}} \{q \cdot |q|_a \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} = 0.$$

The final results that we will discuss are due to Badziahin, Bugeaud, Einsiedler and Kleinbock in 2015 [BBEK15]. In this paper they managed to show that if the complexity function of a real number α grows too quickly, then α satisfies pLC.

Theorem 2.2.10 (Badziahin *et al.* [BBEK15]). *If for a real number α, the entropy $E(\alpha) = \lim_{n\to\infty} \frac{\log(P(w_{CF}(\alpha),n))}{n} > 0$, then α satisfies pLC for every prime number.*

This paper also showed that if the word corresponding to the continued fraction expansion $w_{CF}(\alpha)$ was eventually recurrent then α satisfies mLC.

Theorem 2.2.11 (Badziahin *et al.* [BBEK15])**.** *If for a real number α, we have $w_{CF}(\alpha)$ is eventually recurrent, then for every pseudo-absolute sequence, we have:*

$$\inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} = 0.$$

A corollary of the above theorem is that if the complexity function of a real number α grows too slowly, then α satisfies mLC.

Corollary 2.2.12 (Badziahin *et al.* [BBEK15])**.** *Let $\alpha \in \mathbb{R}$ such that:*

$$\lim_{n\to\infty} P(w_{CF}(\alpha), n) - n < \infty.$$

Then for every pseudo-absolute sequence, we have:

$$\inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} = 0.$$

2.2.4 Reformulating mLC and pLC

Since we defined **Bad** both as the set of irrational numbers satisfying $\liminf_{q\in\mathbb{N}} \{q \cdot \|q\alpha\|\} > 0$ and the set of irrational numbers with bounded partial quotients in their continued fraction expansions, it is natural to ask whether we can reformulate mLC and pLC from the condition $\inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} = 0$ to a condition on partial quotients. This is indeed possible and is commonly known to experts in the field, however, there seems to be no formal statement or proof in the literature. For completeness, we will include a formal statement and a proof of this statement.

Proposition 2.2.13. *Let $\alpha \in$ Bad and let $\mathcal{D} = (d_j)_{j\in\mathbb{N}}$ be a pseudo-absolute sequence. Then α satisfies mLC if and only if:*

$$\sup_{j\in\mathbb{N}} \left\{B(d_j\alpha)\right\} = \infty.$$

Corollary 2.2.14. *Let $\alpha \in$ Bad. Then α satisfies pLC if and only if:*

$$\sup_{\ell\in\mathbb{N}} \left\{B(p^\ell\alpha)\right\} = \infty.$$

Remark 2.2.15. It is (partly) due to this reformulation that we refer to the set of counterexamples to mLC as the set of multiplicatively badly approximable numbers.

To prove these reformulations, we first introduce the following lemmas:

Lemma 2.2.16. *For every $\alpha \in \boldsymbol{Bad}$ and every pseudo-absolute norm $\mathcal{D}$, we have:*

$$\inf_{j\in\mathbb{N}}\left\{\frac{1}{B(d_j\alpha)+2}\right\} < m_{\mathcal{D}}(\alpha).$$

Before we begin this proof, we first make the following claim:

Claim: Let $\beta \in \mathbf{Bad}$ and let $B(\beta)$ be the height function. Then, we have:

$$\frac{1}{B(\beta)+2} < \inf_{k\in\mathbb{N}}\{q_k \cdot \|q_k\beta\|\} = \inf_{q\in\mathbb{N}}\{q\cdot\|q\beta\|\}.$$

Proof of claim. By a reformulation of Hurwitz's Theorem, if $\beta \in \mathbb{R}\setminus\mathbb{Q}$, then there are infinitely many $q \in \mathbb{N}$ such that $q\|q\beta\| < \frac{1}{\sqrt{5}}$. We can similarly reformulate Legendre's Theorem to see that if $q\|q\beta\| < \frac{1}{2}$ for some $q \in \mathbb{N}$, then q is a convergent denominator of β. As a result, we can conclude that the $q \in \mathbb{N}$ which minimise $q\|q\beta\|$ are the convergent denominators of β. Therefore, we have that:

$$\inf_{q\in\mathbb{N}} q\|q\beta\| = \inf_{k\in\mathbb{N}} q_k\|q_k\beta\|,$$

where q_k are the convergent denominators of β. By Theorem 2.1.9 for each convergent denominator q_k, we have:

$$\frac{1}{a_{k+1}+2} < q_k\cdot\|q_k\beta\|,$$

where a_{k+1} is the $(k+1)$-th partial quotient of β. Combining this information together, we see that:

$$\frac{1}{B(\beta)+2} = \inf_{k\in\mathbb{N}}\left\{\frac{1}{a_{k+1}+2}\right\} < \inf_{k\in\mathbb{N}}\{q_k\cdot\|q_k\beta\|\} = \inf_{q\in\mathbb{N}}\{q\cdot\|q\beta\|\}.$$

□

Proof of Lemma 2.2.16. Assume that $\alpha \in \mathbf{Bad}$ and $\mathcal{D}$ is a pseudo-absolute sequence. Recall that $v_{\mathcal{D}}(q) := \sup_{n\in\mathbb{N}}\left\{d_j : d_j \mid q\right\}$. Then, for any $q \in \mathbb{N}$, we write q in terms of the function $v_{\mathcal{D}}(q)$, i.e. $q = v_{\mathcal{D}}(q)q'$, for some $q' \in \mathbb{N}$. Since the $\mathcal{D}$-adic norm is given by $|q|_{\mathcal{D}} = \frac{1}{v_{\mathcal{D}}(q)}$, we have:

$$\begin{aligned} q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\| &= v_{\mathcal{D}}(q)q' \cdot |v_{\mathcal{D}}(q)q'|_{\mathcal{D}} \cdot \|v_{\mathcal{D}}(q)q'\alpha\| \\ &= v_{\mathcal{D}}(q)q' \cdot \frac{1}{v_{\mathcal{D}}(q)} \cdot \|v_{\mathcal{D}}(q)q'\alpha\| \\ &= q'\|q'(v_{\mathcal{D}}(q)\alpha)\|. \end{aligned}$$

We can then apply the above claim, replacing β with $(v_{\mathcal{D}}(q)\alpha)$, to see that:

$$\begin{aligned} m_{\mathcal{D}}(\alpha) &= \inf_{q\in\mathbb{N}} \{q \cdot |q|_{\mathcal{D}} \cdot \|q\alpha\|\} \\ &= \inf_{q\in\mathbb{N}} \left\{q'\|q'(v_{\mathcal{D}}(q)\alpha)\|\right\} \\ &> \inf_{q\in\mathbb{N}} \left\{\frac{1}{B(v_{\mathcal{D}}(q)\alpha) + 2}\right\}. \end{aligned}$$

Since the function $v_{\mathcal{D}}(q)$ only outputs values $d_j \in \mathcal{D}$, we can conclude that:

$$\inf_{j\in\mathbb{N}} \left\{\frac{1}{B(d_j\alpha) + 2}\right\} = \inf_{q\in\mathbb{N}} \left\{\frac{1}{B(v_{\mathcal{D}}(q)\alpha) + 2}\right\}.$$

In particular, if $\inf_{j\in\mathbb{N}} \left\{\frac{1}{B(d_j\alpha)+2}\right\} > 0$, then α does not satisfy mLC. □

Lemma 2.2.17. *For every $\alpha \in$ **Bad** and every pseudo-absolute norm $\mathcal{D}$, we have:*

$$m_{\mathcal{D}}(\alpha) < \inf_{j\in\mathbb{N}} \left\{\frac{1}{B(d_j\alpha)}\right\}.$$

Proof. Given a pseudo-absolute sequence $\mathcal{D}$, we fix j and consider d_j. We can then take q_k to be the k-th convergent denominator for $d_j\alpha$. We can then construct the sequence of natural numbers $\{d_jq_k\}_{k\in\mathbb{N}}$. For each, d_jq_k, we know that $d_j \mid d_jq_k$ and so $|d_jq_k|_{\mathcal{D}} \le \frac{1}{d_j}$. Using this information, we see that:

$$\begin{aligned} d_jq_k \cdot |d_jq_k|_{\mathcal{D}} \cdot \|d_jq_k\alpha\| &\le d_jq_k \cdot \frac{1}{d_j} \cdot \|d_jq_k\alpha\| \\ &= q_k \cdot \|q_k(d_j\alpha)\| \end{aligned}$$

We know that, if a_k is a partial quotient of $d_j\alpha$ and q_k is a convergent denominator, then we have:

$$\inf_{k\in\mathbb{N}}\{q_k\cdot\|q_k(d_j\alpha)\|\}<\inf_{k\in\mathbb{N}}\left\{\frac{1}{a_k}\right\}<\frac{1}{B(d_j\alpha)}.$$

Therefore, we can conclude that:

$$\inf_{k\in\mathbb{N}}\{d_jq_k\cdot|d_jq_k|_{\mathcal{D}}\cdot\|d_jq_k\alpha\|\}<\left\{\frac{1}{B(d_j\alpha)}\right\}.$$

Since $\{d_jq_k\}_{k\in\mathbb{N}}$ is a sub-sequence of $\{q\}_{q\in\mathbb{N}}$, we can also conclude that:

$$m_{\mathcal{D}}(\alpha)=\inf_{q\in\mathbb{N}}\{q\cdot|q|_{\mathcal{D}}\cdot\|q\alpha\|\}\le\inf_{k\in\mathbb{N}}\{d_jq_k\cdot|d_jq_k|_{\mathcal{D}}\cdot\|d_jq_k\alpha\|\}<\frac{1}{B(d_j\alpha)}.$$

Finally, since d_j was arbitrarily chosen from the pseudo-absolute sequence $\mathcal{D}$, we can conclude that:

$$m_{\mathcal{D}}(\alpha)<\inf_{j\in\mathbb{N}}\left\{\frac{1}{B(d_j\alpha)}\right\}.$$

□

Proof of Proposition 2.2.13. Combining together Lemma 2.2.16 and Lemma 2.2.17, we see that for every $\alpha\in\mathbf{Bad}$, we have:

$$\inf_{j\in\mathbb{N}}\left\{\frac{1}{B(d_j\alpha)+2}\right\}<m_{\mathcal{D}}(\alpha)<\inf_{j\in\mathbb{N}}\left\{\frac{1}{B(d_j\alpha)}\right\}.$$

As a result, α satisfies mLC, i.e. $m_{\mathcal{D}}(\alpha)=0$, if and only if

$$\inf_{j\in\mathbb{N}}\left\{\frac{1}{B(d_j\alpha)}\right\}=0.$$

This is equivalent to saying:

$$\sup_{j\in\mathbb{N}}\{B(d_j\alpha)\}=\infty,$$

as required. □

Chapter 3

Cutting Sequences and Integer Multiplication of Continued Fractions

In this chapter we will introduce the notion of cutting sequences, discuss the link between cutting sequences and continued fractions, and explain how replacing one triangulation of the hyperbolic plane with another can be used to represent integer multiplication of continued fractions. This chapter will be split into two sections.

The first section - Section 3.1 - is predominantly a brief overview of previous work done in this area. In Section 3.1.1, we will recall the classical definition of the cutting sequence (ζ, T) of a geodesic ray ζ relative to some ideal triangulation T. We will use the terminology and general construction introduced by Series in [Ser85a; Ser85b], which itself was influenced by the work of Humbert [Hum16]. We will then extend this definition to include paths - not just geodesics - and discuss how homotopy affects the cutting sequence of paths. In Section 3.1.2, we will introduce the Farey tessellation $\mathcal{F}$, an ideal triangulation of the hyperbolic plane. Having constructed the Farey tessellation and discussed some of its properties, we then state Theorem A of [Ser85b], which is foundational for the rest of the work in this thesis:

Theorem 3.1.22. ([Ser85b, Theorem A]) *Let ζ_α be a geodesic ray in $\mathbb{H}$, starting at the y-axis and terminating at a point $\alpha \in \mathbb{R}_{>0}$. Then, if $(\zeta_\alpha, \mathcal{F}) = L^{n_0}R^{n_1}\cdots$ is the cutting sequence of ζ_α relative to $\mathcal{F}$, the continued fraction expansion of α is given by $\overline{\alpha} = [n_0; n_1, \ldots]$, for $n_0 \in \mathbb{N} \cup \{0\}$ and $n_i \in \mathbb{N}$ otherwise.*

The second section - Section 3.2 - is predominantly new material, building on the work introduced in Section 3.1. The main aim of this section is to construct an algorithm that multiplies continued fractions by some integer n. This sets up the framework for Chapters 4 and 5, which discuss how certain classes of continued fractions behave when multiplied. We begin this section by showing that if ζ_α is a geodesic ray starting at the y-axis I and terminating at a point $\alpha \in \mathbb{R}_{>0}$, and $\frac{1}{n}\mathcal{F}$ is the $\frac{1}{n}$-scaled Farey tessellation, then the cutting sequence $(\zeta_\alpha, \frac{1}{n}\mathcal{F})$ is equivalent to the continued fraction expansion $\overline{n\alpha}$. In particular, replacing the triangulation $\mathcal{F}$ by $\frac{1}{n}\mathcal{F}$ induces integer multiplication by n on the corresponding continued fraction expansion. We discuss how if P_n is a fundamental domain of $\Gamma_0(n) := \{\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in PSL_2(\mathbb{Z}) : c \equiv 0 \mod n\}$, then there are two canonical triangulations of P_n - which we denote $T_{\{1,n\}}$ and $T_{\{n,n\}}$ - such that these triangulations "generate" $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ under the action of $\Gamma_0(n)$. In particular, the triangulation replacement of $T_{\{1,n\}}$ by $T_{\{n,n\}}$ in P_n is equivalent to the triangulation replacement of $\mathcal{F}$ by $\frac{1}{n}\mathcal{F}$ in $\mathbb{H}$. In Section 3.2.2, we give a brief overview of the work done by Kulkarni in [Kul91], which discusses some of the general properties for fundamental domains of $\Gamma_0(n)$. We end this chapter with Section 3.2.3, which outlines how we can use the triangulation replacement of $T_{\{1,n\}}$ by $T_{\{n,n\}}$ in P_n to construct an explicit integer multiplication algorithm for continued fractions.

3.1 Continued Fractions as Cutting Sequences

In this section we will introduce the notion of a cutting sequence of a geodesic ray relative to some arbitrary triangulation and then generalise this notion to include paths. We will discuss how cutting sequences relate to continued fraction expansions, as well as some more general properties of cutting sequences. The most important properties will be that:

1. Cutting sequences behave "nicely" when performing homotopy on the underlying path.

2. Given certain conditions, we can decompose a path λ into an ordered collection of sub-paths $\{\lambda_i\}_{i\in\mathbb{N}}$, such that the cutting sequence (λ, T) is equal to the product of cutting sequences $\prod_{i\in\mathbb{N}}(\lambda_i, T)$.

It is these properties which allow us to ultimately build the multiplication algorithm.

We then introduce the Farey tessellation $\mathcal{F}$ and discus how the cutting sequences of geodesic rays relative to the Farey tessellation correspond to continued fraction expansions in a natural way. In particular, if ζ_α is a geodesic ray starting at the y-axis and terminating at some point $\alpha \in \mathbb{R}_{>0}$ and $L^{n_0}R^{n_1}\cdots$ is the cutting sequence of ζ_α relative to $\mathcal{F}$, then the continued fraction expansion $\overline{\alpha}$ is $[n_0; n_1, \ldots]$.

3.1.1 Cutting Sequences

Cutting Sequences of Geodesic Rays

Throughout this thesis we will work with the *hyperbolic plane* $\mathbb{H}$. We will represent the hyperbolic plane by the upper half plane model $\mathbb{H} := \{z \in \mathbb{C} \cup \{\infty\} : Im(z) \geq 0\}$ with boundary $\partial\mathbb{H} = \mathbb{R} \cup \{\infty\}$. Geodesic lines are given by Euclidean half-lines of the form $\{a + iy : 0 \leq y \leq \infty\}$ and semicircles centred on $\partial\mathbb{H}$. We define a hyperbolic n-gon P to be the region enclosed by (and including) the *edges* $l_1, \ldots, l_n$, where:

1. each l_i is a geodesic segment,

2. consecutive edges l_i and l_{i+1} intersect only at a common endpoint v_i and no other edges pass through v_i - here, we treat l_{n+1} as l_1,

3. and the edges are otherwise pairwise disjoint, i.e.:

$$l_i \cap l_j = \begin{cases} \varnothing & \text{If } j \neq i-1, i+1, \\ v_{i-1} \text{ or } v_i & \text{otherwise.} \end{cases}$$

Given two consecutive edges l_i and l_{i+1} in P, we refer to the common endpoint of these edges v_i as a vertex of P. A hyperbolic n-gon is *ideal* if all of its vertices lie on the boundary of the hyperbolic plane $\partial\mathbb{H}$. A *tessellation* of $\mathbb{H}$ will be a collection of hyperbolic polygons $\mathcal{P} = \{\tau_i\}_{i\in\mathbb{N}}$ such that the collection of these polygons cover $\mathbb{H}$, i.e. $\bigcup_{i\in\mathbb{N}} \tau_i = \mathbb{H}$, and for any two polygons τ_j, τ_k in $\mathcal{P}$ these polygons either do not intersect, i.e. $\tau_j \cap \tau_k = \varnothing$, intersect only at a common vertex, i.e. $\tau_j \cap \tau_j = z_i$, or intersect along a common edge, i.e. $\tau_j \cap \tau_k = l_i$, where l_i is an edge of both τ_1 and τ_2. If E is an edge of a polygon $\tau \in \mathcal{P}$, we will say that E is an edge of the tessellation $\mathcal{P}$. If these polygons in $\mathcal{P}$ are all ideal 3-gons, then we refer to $\mathcal{P}$ as an *ideal triangulation* of $\mathbb{H}$.

Remark 3.1.1. We will only consider 'nice' tessallations: we will assume that for every open neighbourhood ν in $\mathbb{H}$, there are only finitely many polygons in the tessellation which intersect ν.

Let ζ be an oriented geodesic ray which enters an ideal triangle $\triangle ABC$, labelled clockwise, through the edge AB. Then ζ can leave the triangle $\triangle ABC$ in one of three ways:

1. The geodesic ζ passes through the edge BC. This isolates the vertex B (lying to the left of ζ) from the vertices A and C (which lie to the right of ζ). In this case, we say that ζ cuts $\triangle ABC$ to form a *left triangle*. See Fig. 3.1 (a).

2. The geodesic ζ passes through the edge AC. This isolates the vertex A (lying to the right of ζ) from the vertices B and C (which lie to the left of ζ). In this case, we say that ζ cuts $\triangle ABC$ to form a *right triangle*. See Fig. 3.1 (b).

3. The geodesic terminates at the vertex C. Here, we refer to the vertex C as the *opposing vertex*.

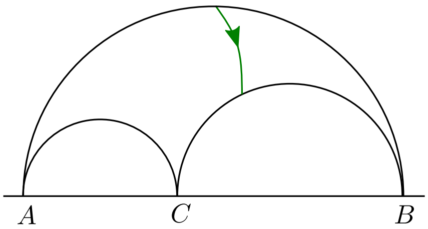

(a) An example of a left triangle.

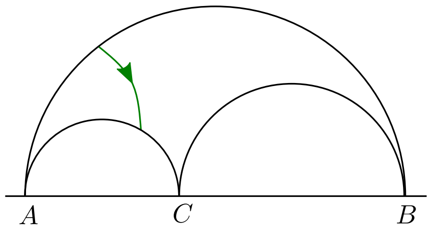

(b) An example of a right triangle.

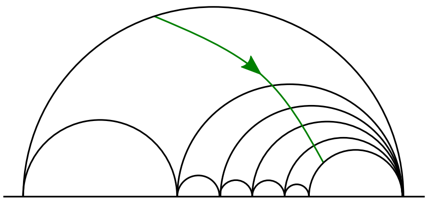

(c) An example of a left fan.

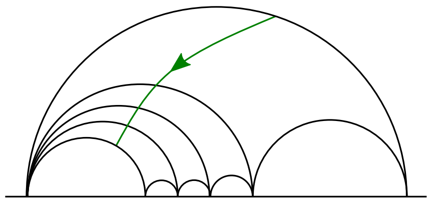

(d) An example of a right fan.

Figure 3.1: Examples of left and right triangles and fans.

Let T be an ideal triangulation of $\mathbb{H}$ and let ζ be an oriented geodesic ray, starting at some edge E of T and terminating at some point $p \in \partial\mathbb{H}$ (where p is not an endpoint of E). We can then form an ordered collection $\{\tau_i\}_{i\in\mathbb{N}\cup\{0\}}$ of the all the triangles in T, which ζ non-trivially intersects, i.e. ζ intersects the interior of each triangle τ_i. For each triangle τ_i, the geodesic ray ζ either cuts τ_i to form a left triangle, a right triangle, or terminates at the opposing vertex. If ζ intersects multiple left triangles in a row, then we refer to the collection of all these triangles as a *left fan*. Similarly, if ζ intersects multiple right triangles in a row, then we refer to the collection of right triangles as a *right fan*. See Fig. 3.1 (c) and (d). If ζ passes through an opposing vertex of a triangle τ, then we could think of this as ζ cutting τ to form either a left triangle or a right triangle - however, for the sake of uniqueness, we will always take this triangle to be a left triangle. If ζ terminates at an opposing vertex, then ζ does not intersect any more triangles in T. In particular, the collection of triangles that

ζ intersects is finite if and only if ζ terminates at some opposing vertex. Using these notions, we can define the *cutting sequence* (ζ, T) of a geodesic ray ζ relative to a triangulation T, as follows:

Definition 3.1.2. Let T be an ideal triangulation of $\mathbb{H}$, let E be any edge of T and let ζ be an oriented geodesic ray starting at E and terminating at some point $p \in \partial\mathbb{H}$. Also, let $\{\tau_i\}_{i\in\mathbb{N}\cup\{0\}}$ be the ordered collection of all triangles in T which ζ non-trivially intersects. Then, the *cutting sequence* of ζ with respect to T, denoted (ζ, T), is the (potentially) infinite word over the alphabet $\{L, R\}$, formed by the following algorithm:

1. Start with $i = 0$ and $(\zeta, T) = \varepsilon$.
2. Repeat the following process until told to stop:
    - If ζ cuts τ_i to form a left triangle:
        - Append the letter L to (ζ, T).
        - Set $i = i + 1$.
    - Else, if ζ cuts τ_i to form a right triangle:
        - Append the letter R to (ζ, T).
        - Set $i = i + 1$.
    - Else, ζ intersects the opposing vertex of τ_i:
        - Append L to (ζ, T).
        - Stop.
3. End of algorithm.

We can write every cutting sequence (ζ, T) in the form $L^{n_0}R^{n_1}L^{n_2}\cdots$, where $n_0 \in \mathbb{N}\cup\{0\}$ and $n_i \in \mathbb{N}$. Each index n_i indicates the size of the i-th fan which ζ forms with T. We will abuse notation and also refer to the term L^{n_i}/R^{n_i} in the cutting sequence as the *i-th fan* of (ζ, T).

Since we can write each cutting sequence in the form $L^{n_0}R^{n_1}L^{n_2}\cdots$ for $n_0 \in \mathbb{N} \cup \{0\}$ and $n_i \in \mathbb{N}$, there is an natural map η between cutting sequences and continued fraction expansions of positive real numbers. This map converts each fan of size of n_i into a partial quotient of size n_i. Explicitly, we have $\eta : L^{n_0}R^{n_1}L^{n_2}\cdots \mapsto [n_0; n_1, n_2, \ldots]$. If the cutting sequence is finite, then it maps to a finite continued fraction. If the cutting sequence is infinite, then it maps to an infinite continued fraction.

If we have the cutting sequence $L^{n_0}R^{n_1}L^{n_2}\cdots L^{n_k}L$, then this would correspond to the continued fraction $[n_0; n_1, n_2, \ldots, n_k+1]$. In our convention, we will always take L to be the final term. This ensures that the cutting sequence is formed in a unique way. However, we could have instead picked R to be our final term, i.e. $L^{n_0}R^{n_1}L^{n_2}\cdots L^{n_k}R$. This would correspond to the continued fraction $[n_0; n_1, n_2, \ldots, n_k, 1]$. In particular, the choice of ending the cutting sequence with either L or R is analogous to the choice of whether the continued fraction expansion is of the form $[n_0; n_1, n_2, \ldots, n_k + 1]$ or $[n_0; n_1, n_2, \ldots, n_k, 1]$.

Cutting Sequences of Paths

It will often be useful to deal with paths starting from some edge E and terminating at some point in $p \in \partial\mathbb{H}$, instead of geodesic rays. This allows us to see how homotopy affects cutting sequences. To do this, we will need extend to the definition of cutting sequences to include paths which may double back on themselves. We will assume that λ is an oriented path, which starts in the interior of some edge γ in T, and terminates at some point $\alpha \in \partial\mathbb{H} = \mathbb{R} \cup \{\infty\}$. Furthermore, we will assume that our path λ is *transverse* to each edge in T, i.e. no point on λ will have a common tangent with any point on any edge that it intersects. Up to homotopy, we can always guarantee that λ is transverse to each edge in T. If the path λ (transversally) intersects an edge AB of a triangle $\triangle ABC$, labelled clockwise, then λ can intersect $\triangle ABC$ in one of six ways:

1. The path λ leaves $\triangle ABC$ through the edge BC, i.e. λ cuts $\triangle ABC$ to form left triangle.

2. The path λ leaves $\triangle ABC$ through the edge AC, i.e. λ cuts $\triangle ABC$ to form right triangle.

3. The path λ doubles back on itself and leaves through the edge AB. In this case we say λ forms a *bigon* with $\triangle ABC$. See Fig. 3.2.

4. The path λ terminates at the vertex A. We refer to A as the *right vertex*.

5. The path λ terminates at the vertex B. We refer to B as the *left vertex*.

6. The path λ terminates at the vertex C. We refer to C as the *opposing vertex*.

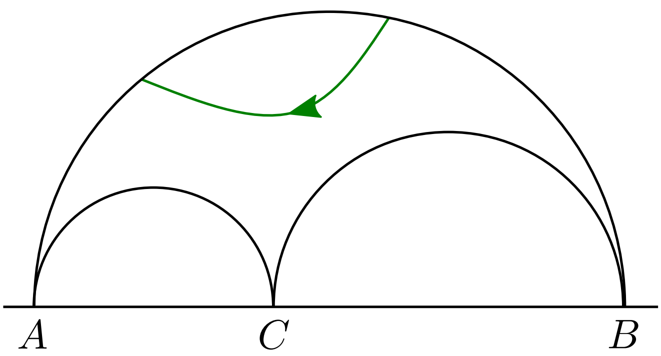

Figure 3.2: An example of a path λ cutting a triangle $\triangle ABC$ to form a bigon.

We will refer to left vertex and right vertex as the *adjacent vertices*. We can then define the *generalised cutting sequence* for λ relative to T, as follows:

Definition 3.1.3. Let T be an ideal triangulation of $\mathbb{H}$, let E be any edge of T and let λ be an oriented path starting at E and terminating at some point in $p \in \partial\mathbb{H}$. Also, let $\{\tau_i\}_{i\in\mathbb{N}\cup\{0\}}$ be the ordered collection of all triangles in T which λ intersects.

Then the *generalised cutting sequence* of λ with respect to T, denoted (λ, T), is the potentially infinite word over the alphabet $\{L, R, X\}$, formed by the following algorithm:

1. Start with $i = 0$ and $(\lambda, T) = \varepsilon$.
2. Repeat the following process until told to stop:
    - If λ cuts τ_i to form a left triangle:
        - Append the letter L to (λ, T).
        - Set $i = i + 1$.
    - Else, if λ cuts τ_i to form a right triangle:
        - Append the letter R to (λ, T).
        - Set $i = i + 1$.
    - Else, if λ cuts τ_i to form a bigon:
        - Append the letter X to (λ, T).
        - Set $i = i + 1$.
    - Else, if λ intersects the left vertex:
        - Append LX.
        - Stop.
    - Else, if λ intersects the right vertex:
        - Append RX.
        - Stop.
    - Else, λ intersects the opposing vertex:
        - Append L.
        - Stop.
3. End of algorithm.

Remark 3.1.4. For a geodesic ray ζ and an ideal triangulation T, the geodesic ray ζ will only ever cut a triangle τ to form a left or right triangle, or intersect the opposing vertex. In particular, the notion of a cutting sequence and a generalised cutting sequence are equivalent for geodesic rays.

Furthermore, for a general path λ and an ideal triangulation T, we may have that the cutting sequence (λ, T) ends with a fan of infinite size, i.e. L^∞ or R^∞. In this case, each triangle in the fan has a common vertex p. Let $\{E_i\}_{i\in\mathbb{N}}$ are the edges of the fan that λ intersects. Then each edge E_i has p as one of its endpoints and some point v_i as the other endpoint. We can then conclude that the sequence of endpoints $\{v_i\}_{i\in\mathbb{N}}$ must limit to p. Otherwise, if v is the limit of the sequence $\{v_i\}_{i\in\mathbb{N}}$, then the edges $\{E_i\}_{i\in\mathbb{N}}$ limit to an edge E with distinct endpoints v and p. In this case, any point on the edge E will always have infinitely many triangles in any open neighbourhood - and therefore, we do not have a 'nice' triangulation of $\mathbb{H}$. As a result, any path λ which passes through a fan of infinite size, must terminate at the common vertex p of this fan.

Homotopy and Cutting Sequences

One of the main reasons why we introduced the notion of a generalised cutting sequence is so that we can discuss the effect that homotopy has on cutting sequences. In actual fact, we will want a slightly restricted form of homotopy.

Definition 3.1.5. Let λ and λ' be two paths which both start at the same edge E. Then, λ and λ' are *homotopic relative to E*, if:

- The paths λ and λ' are homotopic.
- This homotopy preserves the interior of the edge E.
- This homotopy fixes the endpoints of λ and λ' in $\partial\mathbb{H}$.

If λ and λ' start at the same edge E and are homotopic to E, we will denote this as $\lambda \sim_E \lambda'$. Note that in the hyperbolic plane, the paths λ and λ' are homotopic relative to E if and only if they have the same endpoint $p \in \partial\mathbb{H}$. As a result, we will denote the class of all paths which are homotopic to λ relative to E as $[\lambda]_p^E$, where E is the starting edge and $p \in \partial\mathbb{H}$ is the terminal point. However, this is not true when discussing paths on surfaces, as in Chapter 4.

Given an initial starting edge E, there are two possible ways a path λ can leave this edge, corresponding to the two sides of each edge. We will arbitrarily refer to one of these sides as the *positive side* of E, which label we label "$+$", and refer to the other side as the *negative side* of E, which we label "$-$". If λ leaves E via the positive side then we say that λ has *positive direction of departure*, and if λ leaves E via the negative side we say that λ has *negative direction of departure*. We denote the direction of departure by the pair $(E, \pm)$, where E is the starting edge and $\pm \in \{+, -\}$ represents whether λ leaves E via the positive side or negative side.

Given a starting edge and direction of departure, the generalised cutting sequence completely encodes the path λ. In particular, if we had an arbitrary word W over the alphabet $\{L, R, X\}$, an initial starting edge and direction of departure, we could reverse-engineer a path λ_W such that $(\lambda_W, T) = W$. Generally speaking, each letter of W iteratively tells us how λ_W intersects each triangle in T, however, if the word is finite and ends in an X, then we have to consider the last two letters as a pair. This path λ_W is unique up to some minimal homotopy, i.e. for any triangle τ that λ_W intersects, we are free to homotope the path λ_W within this triangle, as long as the path still cuts τ in the same way. From this we can deduce that if two paths λ and λ' have the same starting edge, direction of departure and generalised cutting sequence, then they must be homotopic relative to E.

Lemma 3.1.6. *Let λ and λ' be two paths in $\mathbb{H}$ starting at some edge E in T with the same cutting sequence and direction of departure. Then λ and λ' are homotopic relative to E.*

The reverse of this statement is not always true: if λ and λ' are homotopic relative to E, they do not necessarily have the same generalised cutting sequence.

Reduced Paths and Minimal Position

In general, it would be useful to be able to tell when two paths are homotopic relative to T just by looking at their generalised cutting sequences. In order to do this, we first must introduce a bit of general theory.

Definition 3.1.7. Let λ and μ be paths in some space X and let $[\lambda]$ and $[\mu]$ be the homotopy classes of λ and μ which fix the start and endpoints of λ and μ, respectively. We define the *intersection number* $i(\lambda, \mu)$ of λ and μ to be the minimum number of times any representatives of the homotopy classes $[\lambda]$ and $[\mu]$ (transversally) intersect. In other words:

$$i(\lambda, \mu) := \min \left\{ \#(\lambda' \cap \mu') : \lambda' \in [\lambda], \mu' \in [\mu] \right\}.$$

If two paths λ and μ intersect each other minimally, i.e. $i(\lambda, \mu) = \#(\lambda \cap \mu)$, then we say that these paths are in *minimal position*. Finally, given a path λ and a collection of paths C, which are disjoint except for at ∂X, we say that λ and C are in *minimal position*, if λ is in minimal position with each path in C.

In our case, we will take $\mathbb{H}$ to be our space X and will want to know exactly when a path λ_p is in minimal position relative to some ideal triangulation T. We will also use the homotopy class $[\lambda_p]_p^E$ instead of the homotopy class $[\lambda_p]$, i.e. our starting point is allowed to move along the interior of the initial edge E. The first thing that we notice is that geodesic rays ζ_p are always in minimal position relative to a triangulation T.

Lemma 3.1.8. *Let T be a triangulation of $\mathbb{H}$ and let ζ_p be a geodesic ray starting at some edge E of T and terminating at some point $p \in \partial\mathbb{H}$. Then ζ_p is in minimal position relative to T.*

Proof. Two geodesics in $\mathbb{H}$ can either intersect exactly once, not intersect at all or be concurrent. As a result, if we assume that p is not an endpoint of the starting edge, then ζ_p intersects edges of T either exactly once or not at all. Let $\{E_i\}_{i\in\mathbb{N}}$ be the sequences of edges in T that ζ_p intersects non-trivially, with the index indicating the order that ζ_p intersects these edges. Each edge E_i splits $\mathbb{H}$ into two regions, which we can arbitrarily label as positive and negative. Without loss of generality, we can assume that ζ_p intersects each edge E_i by going from the negative region to the positive region. The positive region of each E_i must contain p. Otherwise, for some edge E_k, either the endpoint p lies in the negative region - this can not happen since ζ_p would have to pass through E_k again and therefore ζ_p could not be a geodesic - or p is also an endpoint of E_k - this can not happen since ζ_p would not need to intersect the interior of E_k and so this edge would not be part of our sequence of edges. Since each path λ_p in the homotopy class $[\zeta_p]_p^E$ starts in the negative region of each E_i and p lies in the positive region of each E_i, the path λ_p must also intersect all edges in $\{E_i\}_{i\in\mathbb{N}}$ (and potentially some others). Since ζ_p only intersects the edges $\{E_i\}_{i\in\mathbb{N}}$ exactly once and no other edges, we can conclude that ζ_p is in minimal position. □

As a consequence of the above proof, if a geodesic ray ζ_p intersects a sequence of edges $\{E_i\}_{i\in\mathbb{N}}$ in a triangulation T, then a path $\lambda_p \in [\zeta_p]_p^E$ will only be in minimal position if it intersects the same sequence of edges in the same order without intersecting any other edges. As a result, all paths in $[\zeta_p]_p^E$ in minimal position must have the same cutting sequence.

Corollary 3.1.9. *Let λ_p and λ_p' be two homotopic paths in $\mathbb{H}$ starting at some edge E in T and terminating at some point $p \in \partial\mathbb{H}$. Assume that λ_p and λ_p' are both in minimal position with T. Then the corresponding cutting sequences are equal, i.e. $(\lambda_p, T) = (\lambda_p', T)$.*

We will say that a path λ_p is *reduced* relative to T, if λ_p does not form a bigon with any triangle in T, terminate at an adjacent vertex of any triangle in T or have a final fan of infinite size. As a result, a path λ_p is reduced if and only if the cutting sequence (λ_p, T) does not contain the letter X or end with the terms L^∞ or R^∞. In this case, we will also say that the cutting sequence (λ_p, T) is *reduced.* Furthermore, if λ_p is not in minimal position with T, we can guarantee that λ_p either forms a bigon with T or terminates at an adjacent vertex. In particular, we can reduce a path λ_p relative to a triangulation T by using homotopy to remove bigons, adjacent vertices and terminal fans of infinite size. As we will see in the next lemma, a path λ_p is reduced if and only if it is in minimal position (relative to triangulation T).

Lemma 3.1.10. *Let λ_p be a path in $\mathbb{H}$ starting at some edge E in T and terminating at some point $p \in \partial\mathbb{H}$. Then, the path λ_p is in minimal position with T if and only if λ_p is reduced relative to T.*

Remark 3.1.11. This lemma is analogous to the *bigon criterion* in [FM11].

Proof. $(\Rightarrow):$ As previously discussed, every geodesic ray ζ_p is reduced relative to every triangulation T. Since every path $\lambda_p \in [\zeta_p]_p^E$ which is in minimal position with T has the same cutting sequence with T, we can conclude that all paths which are in minimal position with T are reduced relative to T.

$(\Leftarrow:)$ We will prove this direction by proving the contrapositive, i.e. if λ_p is not in minimal position with T, then λ_p is not reduced relative to T.

Let λ_p be a path starting at an edge E in T and let ζ_p be a geodesic ray in $[\lambda_p]_p^E$. We will assume that λ_p is not in minimal position. As a result, we can guarantee that there is some edge E' in T that λ_p intersects but ζ_p does not.

We can assume that λ_p passes through E' from the positive region to the negative region. Since ζ_p does not intercept E' and starts within the positive region of E', the endpoint p of λ_p and ζ_p must either lie in positive region of E' or be an endpoint of E'.

Case 1. The point p lies in the positive region of E'.

In this case we can guarantee that λ_p must intersect E' again - otherwise λ_p is trapped in the negative region of E'. Let x be a point on λ_p just before λ_p intersects E' the first time and let y be a point just after λ_p intersects E' for the second time. Then x and y both lie to the positive side of E'. Let μ be the sub-path of λ_p starting at x and terminating at y. The path μ runs from one point in $\mathbb{H}$ to another point in

$\mathbb{H}$ and, therefore, is finite length. As a result, we can find an open neighbourhood of μ. Since we are only looking at 'nice' triangulations of $\mathbb{H}$, we know that this open neighbourhood can only intersect finitely many triangles in T. By extension, the path μ can only intersect finitely many triangles in T and, furthermore, it can only intersect each of these triangles finitely many times (since λ_p is continuous and finite). As a result, we can label each edge that μ intersects, using some a canonical ordering, i.e. $E_1 = E'$ is the first edge that μ intersects, E_2 is the second edge that μ intersects, etc. If for some $i \in \mathbb{N}$ we have $E_i = E_{i+1}$, then this tells us that μ passes through an edge E_i and then immediately passes through E_i again without passing through any other edges in T. In particular, E_i forms a bigon with μ.

Assume that μ passes through n edges - counted with multiplicity. Then we can conclude that E_1 and E_n are both E', since this is the first and last edge that μ intersects. If $n = 2$, then $E_1 = E_2$ and, therefore, λ_p forms a bigon. Otherwise, we take E_2 to be the second edge that μ passes through. Without loss of generality, we can assume that μ passes from the positive region of E_2 to the negative region. Since the edge $E_1 = E_n = E'$ lies in the positive region of E_2, we can conclude that μ must pass through E_2 again - and it must do this before passing through E_n. In particular, we can find some $2 < k < n$ with $E_2 = E_k$. If $E_1 = E_2$, then λ_p forms a bigon.

Since μ can only pass through a finite number of edges - counted with multiplicity - we can conclude that by repeating this procedure, we will find some $1 \le i \le n$ with $E_i = E_{i+1}$. At this point, we can conclude that μ forms bigon with the edge E_i.

Case 2. The endpoint p is also an endpoint of E'.

Let x_1 be the other endpoint of E'. Then, we can assume that when λ_p intersected E' it passed from a triangle τ_0, with endpoints x_0, x_1 and p, to a triangle τ_1, with endpoints x_1, x_2 and p. We know that λ_p can not go to the points x_1 or x_2, since the path would have to terminate at these points. Therefore, there are four possible ways that λ_p can intersect τ_1.

1. The path λ_p passes through the edge $E^{'}$ again. In this case, λ_p forms a bigon with τ_1.

2. The path λ_p passes through the edge F_1 between x_1 and x_2. In this case, we can assume that λ_p goes from the positive region of F_1 to the negative region. Since p lies in the positive region of F_1, we can apply procedure used in case 1 to deduce that λ_p intersects some triangle in T to form a bigon.

3. The path λ_p terminates at p. In this case, λ_p terminates at an adjacent vertex of τ_1.

4. The path λ_p passes through the edge between x_2 and p.

If λ_p intersects τ_1 in any of the first three ways, then we can conclude that λ_p is not reduced, as required. As a result, we will assume that λ_p passes through the edge between x_2 and p. Here, we can assume that λ_p goes from τ_1 to a triangle τ_2, with endpoints x_2, x_3 and p. However, we can now see that λ_p can only intersect τ_2 in one of the same four ways that λ_p can intersect τ_1.

By repeating the above procedure, we see that either λ_p forms a bigon with some triangle τ_k in T in T or λ_p terminates at an adjacent vertex of some triangle τ_k in T or λ_p passes through infinitely many unique triangles $\{\tau_i\}_{i\in\mathbb{N}}$, each of which has p as a common vertex. In this case, since each of the triangles has p as a common vertex, we can conclude that λ_p cuts each of these triangles in the same way and, therefore, the cutting sequence (λ_p, T) ends with a fan of infinite size. In all of these cases, the path λ_p is not reduced and the result follows. □

Since any path λ_p can be reduced by removing bigons, adjacent vertices and fans of infinite size and a path is reduced if any only if it is in minimal position, we can conclude that the homotopy required to put a path λ_p in minimal position can also be realised by removing bigons, adjacent vertices and fans of infinite size. Furthermore, this reduction process introduces a well-defined series of relations on the letters L, R and X of the corresponding cutting sequences. We refer to these relations as the *reduction relations* and denote these relations using the using the symbol $\sim$. These relations are listed below and the corresponding homotopy moves can be seen in Fig. 3.3.

For prefixes:

- Let W be an arbitrary word over the alphabet $\{L, R, X\}$. If $W = XV$, for some V word over the alphabet $\{L, R, X\}$, then $W \sim V$. If λ_W is the path with cutting sequence W and λ_V is the path with cutting sequence V, then λ_W and λ_V have opposing directions of departure.

For arbitrary sub-words:

- $X^2 \sim \varepsilon$.
- $LXL \sim R$ (or $LXL \sim L$, if LXL is a suffix).
- $LXR \sim X$.
- $RXR \sim L$.
- $RXL \sim X$.

For suffixes:

Let W be an arbitrary word in $\{L, R, X\}^*$.

- If $W = VLLX$, for some $V \in \{L, R, X\}^*$, then $W \sim VLX$.
- If $W = VLRX$, for some $V \in \{L, R, X\}^*$, then $W \sim VL$.
- If $W = VRRX$, for some $V \in \{L, R, X\}^*$, then $W \sim VRX$.
- If $W = VRLX$, for some $V \in \{L, R, X\}^*$, then $W \sim VL$.
- If $W = VRL^\infty$, for some $V \in \{L, R, X\}^*$, then $W \sim VL$.
- If $W = VLR^\infty$, for some $V \in \{L, R, X\}^*$, then $W \sim VL$.

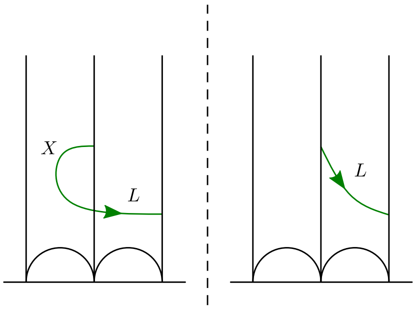

(a) An example of how homotopy can remove the starting letter X from a cutting sequence. Here, the cutting sequence XL reduces to L.

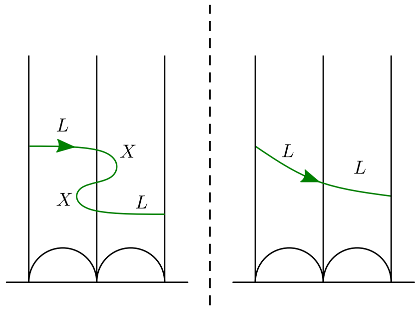

(b) An example of how homotopy induces the reduction relation $X^2 \sim \varepsilon$. Here, the cutting sequence LX^2L reduces to L^2.

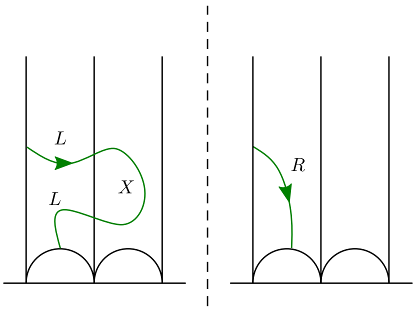

(c) An example of how homotopy induces the reduction relation $LXL \sim R$. By taking the mirror image, we get the reduction relation $RXR \sim L$.

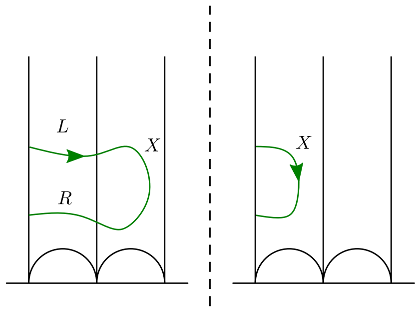

(d) An example of how homotopy induces the reduction relation $LXR \sim X$. By taking the mirror image, we get the reduction relation $RXL \sim X$.

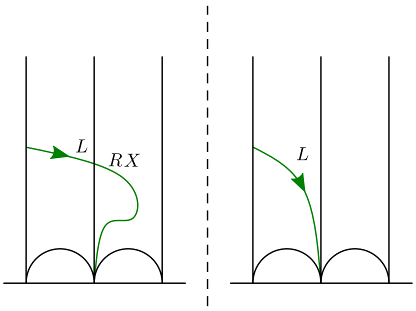

(e) An example of how homotopy induces the reduction relation on the suffix LRX. Here, $LRX \sim L$. By taking the mirror image, we get the reduction relation $RLX \sim L$, where RLX is a suffix.

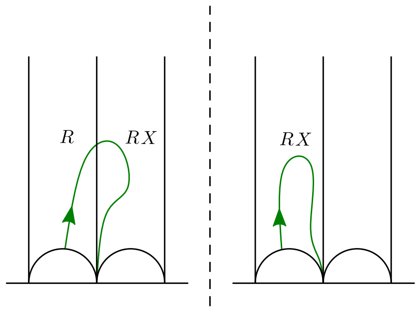

(f) An example of how homotopy induces the reduction relation on the suffix RRX. Here, $RRX \sim RX$. By taking the mirror image, we get the reduction relation $LLX \sim LX$, where LLX is a suffix.

Figure 3.3: A collection of examples of how homotopy induces the reduction relations.

For two words W and W' over the alphabet $\{L, R, X\}$, we will write $W \sim W'$ if these words are equivalent under the above reduction relations. This lead to the following corollary:

Corollary 3.1.12. *Let λ_p and λ_p' be two homotopic paths in $\mathbb{H}$ starting at some edge E in T and terminating at some point $p \in \partial\mathbb{H}$. Then the corresponding cutting sequences are equivalent under reduction relations, i.e. $(\lambda_p, T) \sim (\lambda_p', T)$.*

Remark 3.1.13. Typically, we will only be interested in reduced paths. For ease, we will often take these paths to be geodesics, but any reduced path would work.

As preciously discussed, given a starting edge E, direction of departure and a word W over the alphabet $\{L, R, X\}$, we can form a path λ_W such that this path has cutting sequence W relative to T. However, if we picked the opposite direction of departure, then this could produce a path λ_W' with cutting sequence W, and λ_W' would not be homotopic to λ_W. Furthermore, we could produce the path λ_{XW} which has cutting sequence XW and has the same direction of departure to λ_W, but λ_{XW} would be homotopic to λ_W', and by extension not homotopic to λ_W. This is because the reduction relations do not take into account the fact that homotopy can change the direction of departure. However, we note that the each edge E in the triangulation T will split $\mathbb{H}$ into two regions. We will label these regions E_+ and E_- such that this labelling is consistent with the direction of departure. Instead of describing our paths λ_W by using the direction of departure, we can instead describe our paths based on where the endpoints lie. We can then make the following statement:

Lemma 3.1.14. *Let W and W' be two words over the alphabet $\{L, R, X\}$, which are equivalent up to the reduction relations and assume that these words are not equivalent to LX or RX. Let T be an ideal triangulation and E some edge of the triangulation. If λ and λ' are two paths starting at E with cutting sequences $W = (\lambda, T)$ and $W' = (\lambda', T)$, then λ is homotopic to λ' relative to E if and only if the endpoints of λ and λ' both lie in the region E_+ or both lie in the region E_-.*

Note that if the word W does not contain the letter X, then for any triangulation T, any edge E and any direction of departure, the constructed path λ_W will be reduced relative to T. In this case, the path λ_W does not double back on itself, and so whether the endpoint lies in E_+ or E_- only depends on the initial direction of departure. Therefore, we get the following corollary:

Corollary 3.1.15. *Let W be a word over the alphabet $\{L, R\}$. Let T be an ideal triangulation and E some edge of the triangulation. If λ and λ' are two paths starting at E with cutting sequences $W = (\lambda, T)$ and $W = (\lambda', T)$, then λ is homotopic to λ' relative to E if and only if λ and λ' have the same direction of departure.*

Composing Paths and Cutting Sequences

To describe infinite length paths, it will often be useful to split them into a union of infinitely many finite length paths.

If two finite length paths λ and λ' both start at an edge E_1 and terminate at an edge E_2, then we will say that these paths are *homotopic relative to* E_1 and E_2 if they are homotopic, and this homotopy preserves the interior of both E_1 and E_2. We denote λ and λ' being homotopic relative to E_1 and E_2 as $\lambda \sim_{E_1,E_2} \lambda'$. We will say that a finite-length path λ has a *well-defined cutting sequence relative to* T if the path λ starts at one edge E_1 of a triangulation T and terminates at another edge E_2 of T. If λ is a path with a well-defined cutting sequence, then for every triangle in T that λ passes through, λ cuts this triangle to form a left triangle, a right triangle or a bigon, since λ starts and terminates in the interior of two edges of T. As a result, we can think of the paths with well-defined cutting seqeunces as being the paths which are completely determined by their cutting sequence.

Lemma 3.1.16. *Assume that two finite paths λ and λ' both start at the edge E_1 in T and terminate at the edge E_2 in T. Then, the paths λ and λ' are homotopic relative E_1 and E_2 if and only if the cutting sequences (λ, T) and (λ', T) are equivalent up to reduction relations.*

Proof. This follows from the proof of Lemma 3.1.10. □

Remark 3.1.17. It is worth noting that in this setting, we require both a starting edge E_1 in T and a terminal edge E_2 in T. This terminal edge will either lie in the region $E_{1,+}$ or the region $E_{1,-}$, but since E_2 is assumed to be the terminal edge for both λ and λ', the "if and only if" statement in Lemma 3.1.14 is automatically satisfied.

In analogy to the previous section, given a finite word W, a starting edge E of T and a direction of departure, we can construct a finite length path λ_W which terminates in the interior of some edge E' in T and has cutting sequence $W = (\lambda_W, T)$. This path will be unique up to homotopy relative to the starting edge E and terminal edge E'.

If a finite length path λ terminates at an edge E_2 in T, we can define the *direction of approach* in a similar way to how we define the direction of departure. In particular, if λ approaches E_2 from the positive side, we will say that λ has *positive direction of approach*, and if λ approaches E_2 via the negative side, then we will say that λ has *negative direction of approach.* As we did for the direction of departure, we can express the direction of approach as the pair $(E_2, \pm)$. Here, we have a nice duality between the direction of departure and direction of approach. In particular, given a finite word W, a terminal edge E of T and a direction of approach, we can also construct a a finite length path λ_W which starts at some edge E' in T and has cutting sequence $W = (\lambda_W, T)$. This path will be unique up to homotopy relative to the starting edge E' and terminal edge E.

If we let λ_1 and λ_2 be two finite length paths such that λ_1 terminates at an edge E and λ_2 starts at E, then we can homotope the endpoint of λ_1 to coincide with the starting point of λ_2. This effectively allows us to concatenate these paths to form a new path $\lambda = \lambda_1 \circ \lambda_2$. Given two such paths λ_1 and λ_2, we will say that these paths are *compatible* relative to E if the direction of approach of λ_1 is opposite to the direction departure for λ_2. If these paths are compatible and E is an edge in T, then the cutting sequence (λ, T) is equal to the product of the two cutting sequences $(\lambda_1, T) \cdot (\lambda_2, T)$. As a result, the concatenated path λ will be reduced if and only if both λ_1 and λ_2 are reduced. Alternatively, if E is an edge in T and λ_1 and λ_2 are not compatible, then the concatenated path λ trivially intersects the edge E. In particular, if two paths λ_1 and λ_2 are not compatible, then the concatenated path λ will never be reduced. In this case, the cutting sequence (λ, T) is equivalent to the cutting sequence $(\lambda_1, T) \cdot X \cdot (\lambda_2, T)$ - using the convention that a path that intersects an edge E and then immediately turns back contributes an X to the cutting sequence. See Fig 3.4.

If λ is an infinite reduced path, then we can cut λ into sub-paths by cutting along the an arbitrary set of edges in T which intersect λ, and each of these sub-paths have a well-defined cutting sequence relative to $\mathcal{F}$. Note that, since λ is oriented, there is an inherent ordering of the edges of T that λ intersects. As an example, we could cut along a set of edges $E_1, E_2, \ldots, E_{n-1}$ to obtain a collection of sub-paths $\lambda_1, \lambda_2, \ldots, \lambda_n$. The sub-path λ_i meets the sub-path λ_{i+1} at the edge E_i and these sub-paths will be pairwise compatible (since λ was assumed to be reduced). The paths $\lambda_1, \lambda_2, \ldots, \lambda_{n-1}$ will all be finite, but the path λ_n will be infinite (since λ was assumed to be infinite). Since these paths are all compatible, we can conclude that the cutting sequence of the original path (λ, T) is equal to the (ordered) product of the cutting sequences for each sub-path, i.e. $(\lambda, T) = (\lambda_1, T) \cdot (\lambda_2, T) \cdot \ldots \cdot (\lambda_n, T)$. If λ does not terminate at a vertex of T, then it must intersect infinitely many edges of T. In particular, we can decompose λ into an infinite collection of finite paths, i.e. $\lambda = \lambda_1 \circ \lambda_2 \circ \ldots$. In this case we have the following lemma.

Lemma 3.1.18. *Let λ be an infinite length path, let T be an ideal triangulation and let $\{\lambda_i\}_{i \in \mathbb{N}}$ be a collection of finite length paths such that $\lambda = \lambda_1 \circ \lambda_2 \circ \ldots$ and each pair of consecutive paths λ_i and λ_{i+1} are compatible (relative to T). Then, the cutting sequence (λ, T) is equal to the infinite product $\prod_{i=1}^{\infty} (\lambda_i, T)$.*

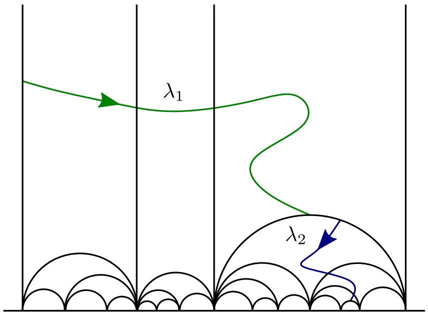

(a) An example of two compatible paths λ_1 and λ_2. These paths have cutting sequences $(\lambda_1, T) = L^2R$ and $(\lambda_2, T) = LRL$.

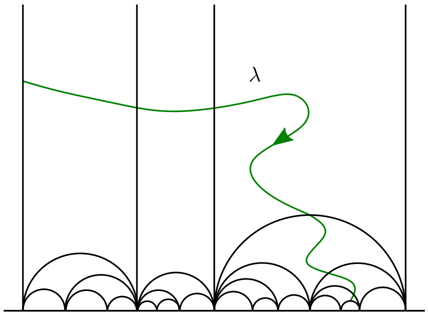

(b) An example of the concatenation λ of two compatible paths λ_1 and λ_2. This path has cutting sequence $(\lambda, T) = L^2RLRL$.

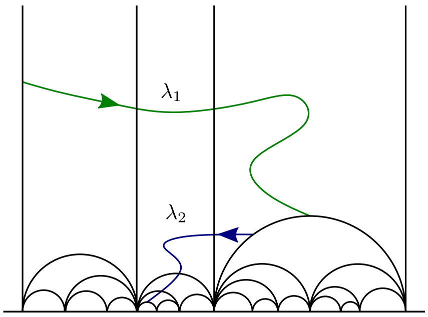

(c) An example of two non-compatible paths λ_1 and λ_2. These paths have cutting sequences $(\lambda_1, T) = L^2R$ and $(\lambda_2, T) = LRL^2$.

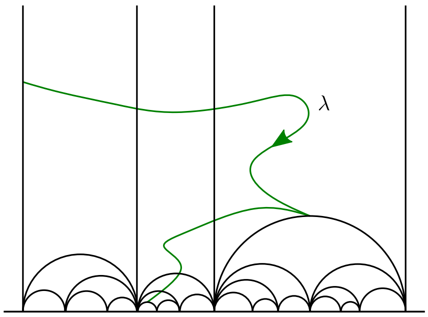

(d) An example of the concatenation λ of two non-compatible paths λ_1 and λ_2. This path has cutting sequence $(\lambda, T) = L^2RXLRL^2$.

Figure 3.4: Examples of how concatenating both compatible paths, (a) and (b), and non-compatible paths, (c) and (d), affects the cutting sequence of their concatenation.

3.1.2 The Farey Tessellation $\mathcal{F}$

The Farey Tessellation $\mathcal{F}$

The Farey tessellation $\mathcal{F}$ is an ideal triangulation of the upper-half plane $\mathbb{H}$. The vertices are the set $\mathbb{Q} \cup \{\infty\}$. Two vertices A and B have a geodesic edge between them if once written in reduced form, $A = \frac{p}{q}$ and $B = \frac{r}{s}$, we have $\mid ps - qr \mid= 1$. We will say that two vertices are *neighbours*, if they have an edge between them. In this definition, we treat ∞ as $\frac{1}{0}$.

Given two vertices $A = \frac{p}{r}$ and $B = \frac{q}{s}$ in $\mathbb{Q} \cup \{\infty\}$ in reduced form, we can define *Farey addition* $\oplus$ and *Farey subtraction* $\ominus$, as follows:

$$A \oplus B := \frac{p+r}{q+s} = \frac{r+p}{s+q} =: B \oplus A$$

$$A \ominus B := \frac{p-r}{q-s} = \frac{r-p}{s-q} =: B \ominus A$$

The first thing to note is that if $A = \frac{p}{q}$ and $B = \frac{r}{s}$ are neighbours in the Farey tessellation, i.e. $\mid ps - qr \mid = 1$, then the point $A \oplus B = \frac{p+r}{q+s}$ is a neighbour of both A and B. The points A and $A \oplus B$ are neighbours since:

$$\mid p \cdot (q+s) - q \cdot (p+r) \mid \; =\mid pq + ps - qp - qr \mid \; =\mid ps - qr \mid = 1,$$

and the points B and $A \oplus B$ are neighbours since:

$$\mid r \cdot (q+s) - s \cdot (p+r) \mid \; =\mid rq + rs - ps - sr \mid \; =\mid -ps + qr \mid = 1.$$

As a result, the points A, B and $A \oplus B$ each have a geodesic edge between them, and, therefore, form a triangle in $\mathcal{F}$. Similarly, if A and B are neighbours in the Farey tessellation, then the point $A \ominus B$ is also a neighbour of both A and B (and is not a neighbour of $A \oplus B$).

Given any point $z \in \mathbb{H}$ and any matrix $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in PSL_2(\mathbb{R})$, we can define the action of M on the point z as follows:

$$M \cdot z := \frac{az+b}{cz+d}.$$

The group $PSL_2(\mathbb{R})$ with action as defined above is isomorphic to the group of *orientation preserving isometries* of $\mathbb{H}$, denoted $Isom^+(\mathbb{H})$. From this perspective, if we take $M = \left(\begin{smallmatrix} p & r \\ q & s \end{smallmatrix}\right) \in PSL_2(\mathbb{Z}) < PSL_2(\mathbb{R})$, and we take the line I between 0 and ∞, then the action of M on I maps I to an edge between the points $M \cdot 0 = \frac{r}{s}$ and $M \cdot \infty = \frac{p}{q}$. Since $M \in PSL_2(\mathbb{Z})$, we have that $\det(M) = ps - qr = 1$. As a result, M maps I to an edge of $\mathcal{F}$. Alternatively, if $A = \frac{p}{q}$ and $B = rs$ are neighbours in $\mathcal{F}$, then since $\mid ps - rq \mid = 1$, we can deduce that either $\left(\begin{smallmatrix} p & r \\ q & s \end{smallmatrix}\right)$ or $\left(\begin{smallmatrix} p & -r \\ q & -s \end{smallmatrix}\right)$

is an element of $PSL_2(\mathbb{Z})$. In particular, the set of edges of the Farey tessellation is equivalent to the set of images of I under the action of $PSL_2(\mathbb{Z})$, i.e. the set of edges $PSL_2(\mathbb{Z}) \cdot I$. This allows us to deduce that $\mathcal{F}$ is *preserved* under the action of $PSL_2(\mathbb{Z})$, i.e. $M \cdot \mathcal{F} = \mathcal{F}$ for all $M \in PSL_2(\mathbb{Z})$. Furthermore, $PSL_2(\mathbb{Z})$ is the *maximal orientation-preserving group* which preserves $\mathcal{F}$, i.e. $M \cdot \mathcal{F} \neq \mathcal{F}$ for any $M \in PSL_2(\mathbb{R}) \setminus PSL_2(\mathbb{Z})$. We write $Isom^{+}(\mathcal{F}) = PSL_2(\mathbb{Z})$ to indicate that $PSL_2(\mathbb{Z})$ is the maximal orientation-preserving group which preserves $\mathcal{F}$. See Fig. 3.5 for a truncated picture of the Farey tessellation.

In fact, not only is the set of edges $PSL_2(\mathbb{Z}) \cdot I$ equivalent to the 1-skeleton of $\mathcal{F}$, but given any two Farey neighbours A and B there is a unique map $M \in PSL_2(\mathbb{Z})$ such that $M \cdot \infty = A$ and $M \cdot 0 = B$.

Proposition 3.1.19. *Given any two Farey neighbours A and B there is a unique matrix $M \in PSL_2(\mathbb{Z})$ such that $M \cdot \infty = A$ and $M \cdot 0 = B$.*

Proof. We start by writing the points A and B in reduced form, $A = \frac{p}{q}$ and $B = \frac{r}{s}$. Since A and B are Farey neighbours we can know that these points satisfy the relations:

$$|ps - rq| = 1,$$

and:

$$\gcd(p, r) = \gcd(s, q) = \gcd(p, q) = \gcd(r, s) = 1.$$

Let $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in PSL_2(\mathbb{Z})$ be a matrix satisfying $M \cdot \infty = A$ and $M \cdot 0 = B$. We know by the above arguments that such a matrix exists and will now show that it is unique. Since $ad - bc = 1$, we can conclude that:

$$\gcd(a, b) = \gcd(a, c) = \gcd(b, d) = \gcd(c, d) = 1.$$

Claim 1: Let $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in PSL_2(\mathbb{Z})$ such that $M \cdot \infty = A$. Then:

$$a = (-1)^n p \text{ and } c = (-1)^n q,$$

for some $n \in \{0, 1\}$.

Since $M \cdot \infty = A$, we can conclude that:

$$M \cdot \infty = \frac{a \cdot \infty + b}{c \cdot \infty + d} = \frac{a}{c} = \frac{p}{q} = A.$$

There are two cases we need to consider here: either $A = \infty$, i.e. $p = 1$ and $q = 0$, or $A = \frac{p}{q}$ with $q \neq 0$.

In the case that $A = \infty = \frac{1}{0}$, then since $\frac{a}{c} = \frac{1}{0}$ and $b, d \in \mathbb{Z}$, we must have that $c = 0 = q$. Therefore, we can conclude that M is of the form:

$$M = \begin{pmatrix} a & b \\ 0 & d \end{pmatrix}.$$

Since $M \in PSL_2(\mathbb{Z})$, we have:

$$ad = 1.$$

In particular, $a \mid 1$ and so $a = (-1)^n = (-1)^n \cdot 1 = (-1)^n \cdot p$ and $c = 0 = (-1)^n \cdot 0 = (-1)^n \cdot q$, for some $n \in \{0, 1\}$.

Alternatively, if $A \neq \infty$ and $q \neq 0$, then we can also conclude that $c \neq 0$, since otherwise this would imply $M \cdot \infty = \infty \neq A$. As a result, we can multiply both sides of the equation $\frac{a}{c} = \frac{p}{q}$ by cq to determine that:

$$qa = pc.$$

Since $a, c, p, q \in \mathbb{Z}$, we must have that $q \mid pc$. However, we know that $\gcd(p, q) = 1$, and so $q \mid c$. On the other hand, we can also conclude that $c \mid qa$. Furthermore, since we have $\gcd(a, c) = 1$, we can conclude that $c \mid q$. Since $q \mid c$ and $c \mid q$, it follows that

$$c = (-1)^n q,$$

where $n \in \{0, 1\}$. This allows us to deduce that:

$$\begin{aligned} qa &= pc \\ &= p(-1)^n q \\ &= (-1)^n pq. \end{aligned}$$

Dividing both sides by $q \neq 0$, we see that:

$$a = (-1)^n p.$$

The claim follows as required. *QED*

Claim 2: Let $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in PSL_2(\mathbb{Z})$ such that $M \cdot 0 = B$. Then:

$$b = (-1)^m r \text{ and } d = (-1)^m s,$$

for some $m \in \{0, 1\}$.

Similar to claim 1, we start by noting that since $M \cdot 0 = B$, we can conclude that:

$$M \cdot 0 = \frac{a \cdot 0 + b}{c \cdot 0 + d} = \frac{b}{d} = \frac{r}{s} = B.$$

Again, there are two cases we need to consider here: either $B = \infty$, i.e. $r = 1$ and $s = 0$, or $B = \frac{r}{s}$ with $s \neq 0$.

In the case that $B = \infty = \frac{1}{0}$, then since $\frac{b}{d} = \frac{1}{0}$ and $b, d \in \mathbb{Z}$, we must have that $d = 0 = s$. Therefore, we can conclude that M is of the form:

$$M = \begin{pmatrix} a & b \\ c & 0 \end{pmatrix}.$$

Since $M \in PSL_2(\mathbb{Z})$, we have:

$$-bc = 1.$$

In particular, $b \mid 1$. We can therefore write that $b = (-1)^m = (-1)^m \cdot 1 = (-1)^m \cdot r$ and $d = 0 = (-1)^m \cdot 0 = (-1)^m \cdot s$, for some $m \in \{0, 1\}$.

The rest of claim 2 follows, by the same argument of claim 1. *QED.*

Combining together claim 1 and claim 2, we can rewrite M in the form:

$$M = \left(\begin{smallmatrix} (-1)^n p & (-1)^m r \\ (-1)^n q & (-1)^m s \end{smallmatrix}\right),$$

for some $n, m \in \{0, 1\}$. Since $M \in PSL_2(\mathbb{Z})$ and I and $-I$ are equivalent in $PSL_2(\mathbb{Z})$, we can assume that $n = 0$ and $m = \{0, 1\}$, i.e. $M = \left(\begin{smallmatrix} p & (-1)^m r \\ q & (-1)^m s \end{smallmatrix}\right)$. As a result, M has two possible forms (however we will see that given a choice of A and B only one of these choices is viable). We can take the determinant to see that:

$$p \cdot (-1)^m s - q \cdot (-1)^m r = (-1)^m (ps - rq) = 1.$$

Since A and B are Farey neighbours, we know that:

$$|ps - rq| = 1,$$

and so, either:

$$ps - rq = 1 \text{ or } ps - rq = -1.$$

If $ps - rq = 1$, then $m = 0$ and $M = \left(\begin{smallmatrix} p & r \\ q & s \end{smallmatrix}\right)$. Otherwise, we have that $m = 1$ and $M = \left(\begin{smallmatrix} p & -r \\ q & -s \end{smallmatrix}\right)$. In either case, M is uniquely defined. □

Corollary 3.1.20. *Let A and B be a pair of Farey neighbours and let C and D be another pair of Farey neighbours. Then there is a unique map $M \in PSL_2(\mathbb{Z})$ such that $M \cdot A = C$ and $M \cdot B = D$.*

Proof. Let N_1 be the map such that $N_1 \cdot \infty = A$ and $N_1 \cdot 0 = B$ and let N_2 be the map such that $N_2 \cdot \infty = C$ and $N_2 \cdot 0 = B$. Then M is given by:

$$M = N_2 N_1^{-1}.$$

□

Remark 3.1.21. Here, we make a brief comment that often we will abuse notation and use $\mathcal{F}$ to refer the 1-skeleton of Farey tessellation (i.e. all of the edges in the Farey tessellation). However, all of the triangles in the Farey tessellation can be defined either be their edges or their interiors, and so, on a practical level, there is no real distinction. In particular, we will think of $PSL_2(\mathbb{Z}) \cdot I$ as being equivalent to $\mathcal{F}$.

Cutting Sequences and the Farey Tessellation

The following theorem highlights the importance of the Farey tessellation with regards to continued fractions. Recall that η is the map the converts (reduced) cutting sequences into continued fractions expansions, i.e. $\eta : L^{n_0}R^{n_1}\cdots \mapsto [n_0; n_1, \ldots]$.

Theorem 3.1.22. ([Ser85b, Theorem A]) *Let ζ be a geodesic in $\mathbb{H}$ with endpoints $\alpha_1 > 0$ and $\alpha_2 < 0$, and let I be the geodesic line between 0 and ∞. Let I_+ be the region $\{z : Re(z) > 0\}$ and I_- be the region $\{z : Re(z) < 0\}$. Then, for $\zeta^+ = \zeta \cap I_+$ and $\zeta^- = \zeta \cap I_-$ (with implicit orientation), $\eta((\zeta^+, \mathcal{F}))$ is the continued fraction expansion of α_1 and $\eta((\zeta^-, \mathcal{F}))$ is the continued fraction expansion of $\frac{-1}{\alpha_2}$.*

The main point to take away from the above theorem is the following: if λ_α is a path (which is reduced relative to $\mathcal{F}$) starting at the the y-axis I and terminating at the the point $\alpha \in \mathbb{R}_{>0}$, then $\eta((\lambda_\alpha, \mathcal{F})) = \overline{\alpha}$. As a result, we can identify the real number $\alpha \in \mathbb{R}_{>0}$ with any path λ_α starting at I and terminating at the point α, and the cutting sequence $(\lambda_\alpha, \mathcal{F})$ is equivalent to the continued fraction expansion $\overline{\alpha}$. However, this is not the only connection between the cutting sequence of a (reduced) path λ_α with the Farey tessellation and the continued fraction expansion $\overline{\alpha}$.

For every fan that λ_α forms with $\mathcal{F}$, there is a vertex which is in all of the triangles of this fan. In particular, every edge in this fan will have a unique common endpoint. We refer to this vertex as the *fixed vertex* of the fan. Let v_k be the fixed vertex of the $(k+1)$-th fan. Then we can label each edge in the fan $E_{k,i}$, where $0 \leq i \leq n_{k+1}$, using the order that λ_α intersects these edges. As previously mentioned, each of these edges $E_{k,i}$ has a common vertex v_k. For each edge $E_{k,i}$, we label the other vertex $v_{k,i}$. If $v_{k,i}$ is the final "other vertex" in this fan, (i.e. $i = n_{k+1}$), then this vertex is either the endpoint of λ_α or it is the fixed vertex of the next fan, i.e. the $(k+2)$-th fan. Likewise, if $v_{k,0}$ is the first "other vertex" of the $(k+1)$-th fan, then $v_{k,0}$ is the fixed vertex of the previous fan, i.e. the k-th fan. We can now note, that if $(\lambda_\alpha, \mathcal{F}) = L^{n_0}R^{n_1}\cdots$, then the reduced path $\lambda_\alpha^{k,i}$ which starts at I and terminates at the vertex $v_{k,i}$, has cutting sequence $L^{n_0}R^{n_1}\cdots L^{n_k}R^{i-1}L$ or

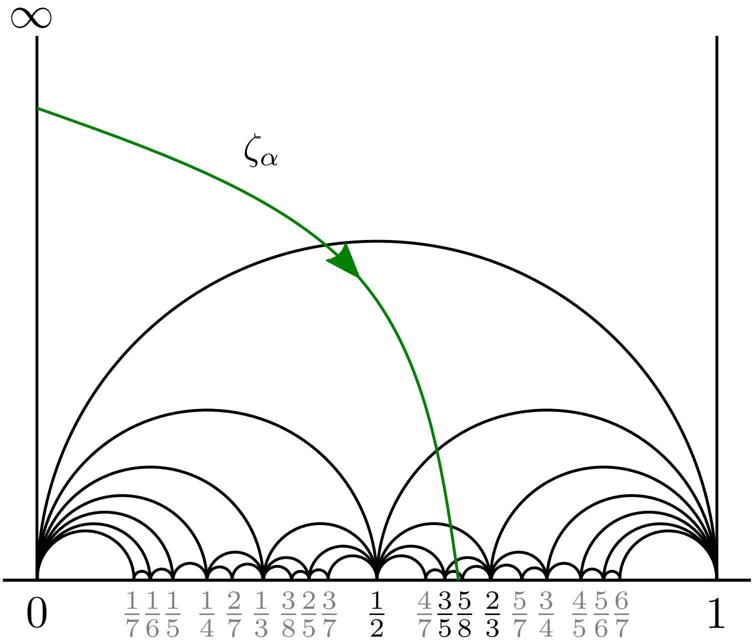

Figure 3.5: An image of a geodesic ray ζ_α intersecting the Farey tessellation $\mathcal{F}$ with (some of the) convergents shown in bold. The endpoint of ζ_α is $\alpha = \frac{\sqrt{5}-1}{2}$. The convergents are $\infty, 0, 1, \frac{1}{2}, \frac{2}{3}, \frac{3}{5}, \frac{5}{8}, \ldots$.

$L^{n_0}R^{n_1}\cdots R^{n_k}L^i$, depending on whether k is even or odd respectively. As a result, we find that the point $v_{k,i}$ has continued fraction expansion $[n_0; n_1, \ldots, n_k, i]$ (up to taking equivalent continued fraction expansions). However, this is simply the $\{k,i\}$-th semi-convergent of α. See Definition 2.1.6. Note that by construction, the point $v_k = v_{k-1,n_k} = v_{k+1,0}$ is the k-th convergent $\frac{p_k}{q_k}$, which can also be written as semi-convergent as $\frac{p_{k-1,n_k}}{q_{k-1,n_k}} = \frac{p_{k+1,0}}{q_{k+1,0}}$. Except for possibly the point $\frac{p_{-1}}{q_{-1}} = \frac{1}{0} = \infty$, every convergent is a fixed point of a fan. This means that each convergent is the endpoint of at least two edges that λ_α intersects. Alternatively, if λ_α intersects two distinct edges, which have the same endpoint, then this endpoint is a fixed point of a fan and, therefore, this point is a convergent. Putting together this information, we get the following corollary:

Corollary 3.1.23. *Let ζ_α be a geodesic ray (or a reduced path) starting at the y-axis I and terminating at the point $\alpha \in \mathbb{R}_{>0}$. Then, the point $v \in \mathbb{Q} \cup \{\infty\}$ is a semi-convergent of α if and only if it is the endpoint of some edge E in $\mathcal{F}$ which* *intersects* ζ_α. *The point* $v \in \mathbb{Q} \cup \{\infty\}$ *is a convergent of* α *if and only if it is the point at* ∞ *or it is the endpoint of at least two edges in* $\mathcal{F}$ *which intersect* ζ_α.

3.2 Integer Multiplication of Continued Fractions and Triangulation Replacement

Let $n^* := \begin{pmatrix} \sqrt{n} & 0 \\ 0 & \frac{1}{\sqrt{n}} \end{pmatrix} \in PSL_2(\mathbb{R})$ and define $\frac{1}{n^*} := (n^*)^{-1}$ for $n \in \mathbb{N}$. These two maps scale both $\mathbb{H}$ and $\mathcal{F}$ by a factor of n and $\frac{1}{n}$, respectively. In particular, they multiply the real axis by n and $\frac{1}{n}$. Since $n^* \notin PSL_2(\mathbb{Z})$ for $n \geq 1$, these maps do not preserve $\mathcal{F}$ and we will refer to the images of $\mathcal{F}$ under these maps as $n\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$, respectively. Both $n\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ will be ideal triangulations of $\mathbb{H}$, since the n^* map will take geodesics to geodesics and triangles to triangles. It is worth noting that both of these maps preserve the line I between 0 and ∞, which is our conventional starting edge for our geodesic rays in $\mathcal{F}$. The initial direction of departure is also preserved, since n^* and $\frac{1}{n^*}$ both preserve the orientation of $\mathbb{H}$. It follows that for any geodesic ray ζ_α starting at I and terminating at $\alpha \in \mathbb{R}_{>0}$, the scaled geodesic ray $n^*(\zeta_\alpha)$ will start at I and terminate at the point $n\alpha \in \mathbb{R}_{>0}$. Note that $n^*(\zeta_\alpha)$ will also be a geodesic ray, since $n^* \in PSL_2(\mathbb{R}) \cong Isom^+(\mathbb{H})$. As a result, the cutting sequence $(n^*(\zeta_\alpha), \mathcal{F})$ will be reduced and equivalent to the continued fraction expansion of $n\alpha$.

Alternatively, we could scale the Farey tessellation by $(n^*)^{-1}$ to get the tessellation $\frac{1}{n}\mathcal{F}$. Relatively speaking, the geodesic ray $n^*(\zeta_\alpha)$ will intersect $\mathcal{F}$ in the same way that ζ_α intersects $\frac{1}{n}\mathcal{F}$. Thus, the cutting sequences will be equivalent, i.e. $(\zeta, \frac{1}{n}\mathcal{F}) = (n^*(\zeta), \mathcal{F})$. Therefore, $\eta(\zeta_\alpha, \frac{1}{n}\mathcal{F}) = \eta(n^*(\zeta_\alpha), \mathcal{F}) = \overline{n\alpha}$. As a result, we can view the map $\overline{n} : \overline{\alpha} \to \overline{n\alpha}$ as being equivalent to replacing the triangulation $\mathcal{F}$ with $\frac{1}{n}\mathcal{F}$ in the corresponding cutting sequence. Explicitly, we can express $\overline{n}$ as the map between the cutting sequences $\overline{n} : (\zeta_\alpha, \mathcal{F}) \to (\zeta_\alpha, \frac{1}{n}\mathcal{F})$.

Since the $\overline{n}$ map in this context is dependent upon the $\frac{1}{n^*}$ map (which is a continuous map), we have only described the $\overline{n}$ map via continuous action on $\mathbb{H}$. Instead we want to describe $\overline{n}$ as a discrete action on local structures. To find such a discrete action, we will claim that for any natural number n, there exists a finite polygon P_n with side pairings and two decorated copies of P_n, $T_{\{1,n\}}$ and $T_{\{n,n\}}$, such that $T_{\{1,n\}}$ tessellates $\mathcal{F}$ and $T_{\{n,n\}}$ tessellates $\frac{1}{n}\mathcal{F}$, under the group action induced by the side pairings of P_n. We will take P_n containing the y-axis I and will take this edge to be our starting edge, unless otherwise stated. Then we express our geodesic ray ζ_α as a collection of ordered sub-paths $\bigcup_{i=1}^{\infty} \zeta_\alpha^{(i)}$ intersecting the tessellation induced by P_n, such that each sub-path $\zeta_\alpha^{(i)}$ is entirely contained in some image of P_n in this tessellation. By slightly expanding our definition of cutting sequences, we find that $(\zeta_\alpha, \mathcal{F}) = \prod_{i=1}^{\infty}(\zeta_\alpha^{(i)}, T_{\{1,n\}})$ and $(\zeta_\alpha, \frac{1}{n}\mathcal{F}) = \prod_{i=1}^{\infty}(\zeta_\alpha^{(i)}, T_{\{n,n\}})$. In particular, replacing $T_{\{1,n\}}$ with $T_{\{n,n\}}$ encodes the multiplication map $\overline{n} : \overline{\alpha} \to \overline{n\alpha}$.

3.2.1 Common Structure of $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$

Given a tessellation T in $\mathbb{H}$ and a matrix $M \in PSL_2(\mathbb{R})$, we say that T is *invariant* under the action of M, if $M \cdot T = T$. We also say that M *preserves* T in this case. Given such a tessellation T, we define the *group of orientation-preserving isometries* of T, $Isom^+(T)$, to be the largest subgroup of $PSL_2(\mathbb{R}) \cong Isom^+(\mathbb{H})$ such that T is invariant under the action of every element of $Isom^+(T)$. Of course, if G is a subgroup of $Isom^+(T)$, then T is also invariant under the action of G. In the case $Isom^+(T)$ is a discrete group, we can find a fundamental domain P such that the action of $Isom^+(T)$ on P tessellates $\mathbb{H}$.

As seen in Section 3.1.2, $Isom^+(\mathcal{F}) = PSL_2(\mathbb{Z})$ is the maximal orientation-preserving group which preserves $\mathcal{F}$. One can then show that $Isom^+\left(\frac{1}{n}\mathcal{F}\right) = \{n^{-1} \circ A \circ n : A \in PSL_2(\mathbb{Z})\}$ is the maximal orientation preserving group which preserves $\frac{1}{n}\mathcal{F}$. We can view each element of $Isom^+(\frac{1}{n}\mathcal{F})$ as a composition of maps: first the map scaling $\frac{1}{n}\mathcal{F}$ to $\mathcal{F}$, followed by an isomorphism of $\mathcal{F}$ and finally, the map scaling $\mathcal{F}$ back to $\frac{1}{n}\mathcal{F}$. By explicit computation, we can see that these elements are of the following form:

$$Isom^+\left(\frac{1}{n}\mathcal{F}\right) = \left\{\left(\begin{smallmatrix} a & \frac{b}{n} \\ nc & d \end{smallmatrix}\right) \in PSL_2(\mathbb{R}) : \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in PSL_2(\mathbb{Z})\right\},$$

and $Isom^+(\frac{1}{n}\mathcal{F})$ takes on a natural group structure induced by $Isom^+(\mathcal{F})$.

We can recover a common subgroup of the maximal invariant subgroups $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ by taking the intersection of $Isom^+(\mathcal{F})$ and $Isom^+(\frac{1}{n}\mathcal{F})$. By explicit computation, we see that $Isom^+(\mathcal{F}) \cap Isom^+(\frac{1}{n}\mathcal{F}) = \Gamma_0(n) := \{\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in PSL_2(\mathbb{Z}) : c \equiv 0 \,(\text{mod n})\}$. The group $\Gamma_0(n)$ is a subgroup of both $Isom^+(\mathcal{F})$ and $Isom^+(\frac{1}{n}\mathcal{F})$ by construction, and, therefore, preserves the structure of both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$. If we take P_n to be a fundamental domain of $\Gamma_0(n)$, then P_n under the action of $\Gamma_0(n)$ tessellates the plane. We note that $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ split P_n into different regions and these regions tessellate P_n. Here, we will explicitly identify the Farey tessellation $\mathcal{F}$ with its set of edges and do the same with the scaled Farey tessellation $\frac{1}{n}\mathcal{F}$. If we take $T_{\{1,n\}} := P_n \cap \mathcal{F}$ and $T_{\{n,n\}} := P_n \cap \frac{1}{n}\mathcal{F}$, we can say more: $T_{\{1,n\}}$ under the action of $\Gamma_0(n)$ will be equivalent to $\mathcal{F}$ and $T_{\{n,n\}}$ under the action of $\Gamma_0(n)$ will be equivalent to $\frac{1}{n}\mathcal{F}$. This follows since:

- The group $\Gamma_0(n)$ preserves $\mathcal{F}$, i.e. it maps edges of $\mathcal{F}$ to edges of $\mathcal{F}$. Therefore, the set of edges in $T_{\{1,n\}}$ under the action $\Gamma_0(n)$ must be contained in the set of edges for $\mathcal{F}$, i.e. $\Gamma_0(n) \cdot T_{\{1,n\}} \subseteq \mathcal{F}$.
- The fundamental domain P_n covers $\mathbb{H}$ under the action of $\Gamma_0(n)$. Therefore the set of edges in $T_{\{1,n\}}$ must cover the set of edges in $\mathcal{F}$, i.e. $\Gamma_0(n) \cdot T_{\{1,n\}} = \mathcal{F}$.
- The group $\Gamma_0(n)$ preserves $\frac{1}{n}\mathcal{F}$, i.e. it maps edges of $\frac{1}{n}\mathcal{F}$ to edges of $\frac{1}{n}\mathcal{F}$. Therefore, the set of edges in $T_{\{n,n\}}$ under the action $\Gamma_0(n)$ must be contained in the set of edges for $\frac{1}{n}\mathcal{F}$, i.e. $\Gamma_0(n) \cdot T_{\{n,n\}} \subseteq \frac{1}{n}\mathcal{F}$.
- The fundamental domain P_n covers $\mathbb{H}$ under the action of $\Gamma_0(n)$. Therefore, the set of edges in $T_{\{n,n\}}$ must cover the set of edges in $\frac{1}{n}\mathcal{F}$, i.e. $\Gamma_0(n) \cdot T_{\{n,n\}} = \frac{1}{n}\mathcal{F}$.

In particular, if we wish to describe how replacing the triangulation $\mathcal{F}$ with the triangulation $\frac{1}{n}\mathcal{F}$ affects cutting sequences of paths in $\mathbb{H}$, it is equivalent to discuss how replacing $T_{\{1,n\}}$ with $T_{\{n,n\}}$ affects cutting sequences. See Fig. 3.6 for images of $T_{\{1,4\}}$ and $T_{\{2,4\}}$ and how these structures tessellate the plane.

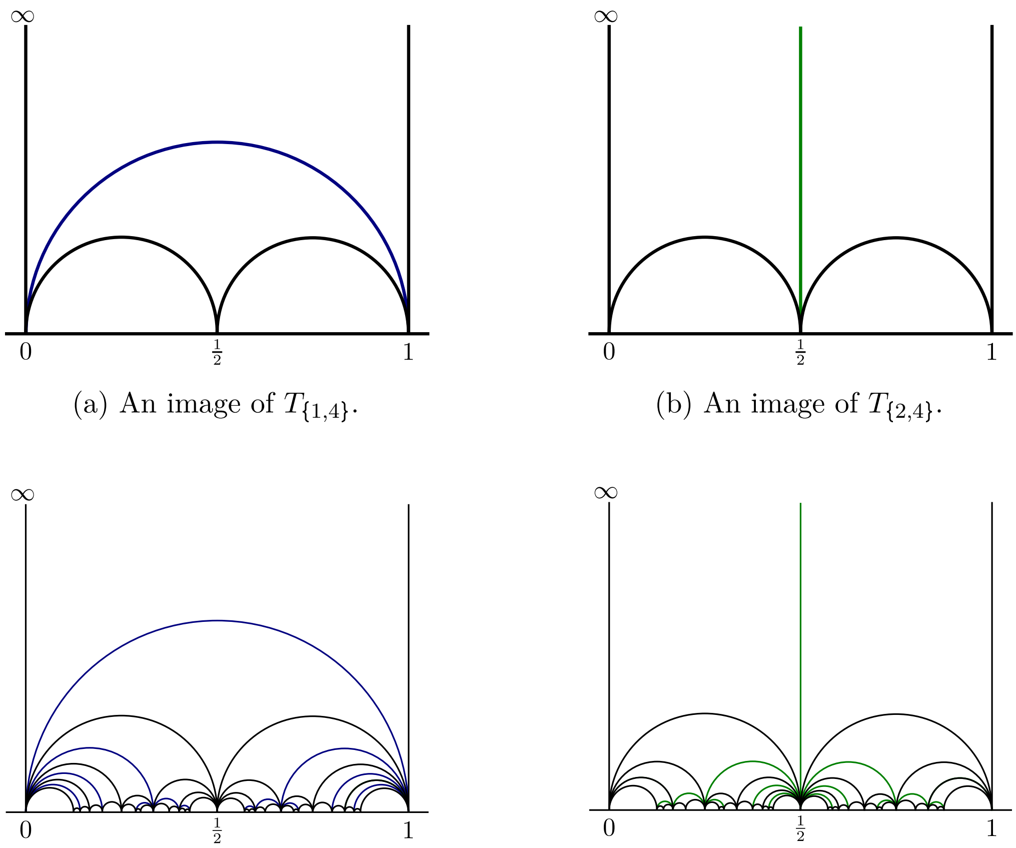

(a) An image of $T_{\{1,4\}}$.

(b) An image of $T_{\{2,4\}}$.

(c) An image of how $T_{\{1,4\}}$ tessellates $\mathbb{H}$ to form $\mathcal{F}$.

(d) An image of how $T_{\{2,4\}}$ tessellates $\mathbb{H}$ to form $\frac{1}{2}\mathcal{F}$.

Figure 3.6: An example of how $T_{\{1,4\}}$ and $T_{\{2,4\}}$ tessellate $\mathbb{H}$ to form $\mathcal{F}$ and $\frac{1}{2}\mathcal{F}$, respectively.

Note that $\Gamma_0(n)$ is a subgroup of $\Gamma_0(d)$ if and only if $d \mid n$. Since $\Gamma_0(n)$ is a subgroup of $\Gamma_0(d)$, $\Gamma_0(n)$ preserves $\frac{1}{d}\mathcal{F}$. As a result, we can also define $T_{\{d,n\}} := P_n \cap \frac{1}{d}\mathcal{F}$ for every $d \mid n$ and the decorated tile $T_{\{d,n\}}$ together with the side pairings induced by $\Gamma_0(n)$ encodes sufficient data to recover $\frac{1}{d}\mathcal{F}$, for all $d \mid n$.

For any geodesic ray ζ_α, we can decompose ζ_α into an ordered collection of subpaths $\bigcup_{i=1}^{\infty} \zeta_\alpha^{(i)}$, such that each $\zeta_\alpha^{(i)}$ is entirely contained in an image of P_n under its tessellation by $\Gamma_0(n)$. We will abuse notation and think of each $\zeta_\alpha^{(i)}$ as a subpath in P_n. Then, the cutting sequence of ζ_α with $\frac{1}{d}\mathcal{F}$ for $d \mid n$, is equivalent to ordered product of the cutting sequences for each $\zeta_\alpha^{(i)}$ with $T_{\{d,n\}}$. In the case that every sub-path $\zeta_\alpha^{(i)}$ has a well defined cutting sequence relative to $T_{\{d,n\}}$, each pair of consecutive sub-paths will be mutually compatible relative to $T_{\{d,n\}}$. As a result, we can apply Lemma 3.1.18 to see that:

$$\prod_{i=1}^{\infty}(\zeta_\alpha^{(i)}, T_{\{d,n\}}) = (\zeta_\alpha^{(1)}, T_{\{d,n\}}) \cdot (\zeta_\alpha^{(2)}, T_{\{d,n\}}) \cdot \ldots = \left(\zeta_\alpha, \frac{1}{d}\mathcal{F}\right).$$

In particular, if each sub-path $\zeta_\alpha^{(i)}$ has a well-defined cutting sequence relative to both $T_{\{1,n\}}$ and $T_{\{n,n\}}$, we can very easily describe the map $\overline{n} : \overline{\alpha} \to \overline{n\alpha}$ by looking at these cutting sequences.

Unfortunately, the paths $\zeta_\alpha^{(i)}$ will rarely have well-defined cutting sequences relative to both $T_{\{1,n\}}$ and $T_{\{n,n\}}$, and so, we need to do a bit more work before we can describe the multiplication algorithm for an arbitrary n. See Section 3.2.3. Before we do this, we will take a brief break to discuss how to construct a fundamental domain P_n of $\Gamma_0(n)$ using the structure of the Farey tessellation. This construction will be very useful, since - for the most part - the edges of P_n will lie in $\mathcal{F}$. As a result, when we split ζ_α into sub-paths $\zeta_\alpha^{(i)}$, each contained in P_n, the endpoints of these sub-paths $\zeta_\alpha^{(i)}$ will lie on the edges of $T_{\{1,n\}}$. Therefore, the paths $\zeta_\alpha^{(i)}$ will have a well-defined cutting sequence relative to $T_{\{1,n\}}$. We then use this nice structure to compensate for the fact that the sub-paths $\zeta_\alpha^{(i)}$ may not have well-defined cutting sequences relative to $T_{\{n,n\}}$.

3.2.2 Fundamental domains of $\Gamma_0(n)$

Fundamental domains of $\Gamma_0(n)$ have been well studied with relation to modular forms. Notably, R.S. Kulkarni outlined how one could explicitly construct a fundamental domain for $\Gamma_0(n)$ (with side pairings) using *Farey symbols* in [Kul91]. This process uses the structure of the Farey tessellation to describe fundamental domains for arbitrary finite index subgroups of $PSL_2(\mathbb{Z})$. These fundamental domains will have very nice structure and, for this reason, we use this section to describe properties of these fundamental domains and main results of [Kul91]. For the sake of brevity, we will not discuss explicitly how these fundamental domains are constructed, but details can be found in [Kul91] and [KL07].

Farey Symbols and Special Polygons

A *Farey Sequence* is a sequence of points in $\mathbb{Q} \cup \{\infty\}$ of the form $\{\infty, x_0, \ldots, x_r, \infty\}$ such that:

- Each consecutive pair of points x_i and x_{i+1} are neighbours in the Farey tessellation $\mathcal{F}$.
- There is some $i \in \{0, \ldots, r\}$ with $x_i = 0$.
- For each $j \in \{0, \ldots, r-1\}$ we have $x_j < x_{j+1}$, i.e. the vertices (excluding ∞) are ordered via the natural ordering of $\mathbb{Q}$.

Given a Farey sequence $\{\infty, x_0, \ldots, x_r, \infty\}$ we can construct a *special polygon.* These special polygons act as the fundamental domains of Φ, where Φ is some finite index subgroup of $PSL_2(\mathbb{Z})$. The Farey sequence will act as our underlying vertex set - from which we will add additional vertices and construct edges.

Definition 3.2.1. A *special polygon* P is a polygon in $\mathbb{H}$ with a set of edge identifications Φ_P in $PSL_2(\mathbb{Z})$ satisfying the following properties:

1. **Underlying vertex set:** The vertices of P lying on $\partial\mathbb{H}$ form a Farey sequence $\{\infty, x_0, \ldots, x_r, \infty\}$ with an induced natural ordering.
2. **Edges:** Each pair of consecutive neighbours x_i and x_{i+1} in the underlying Farey sequence satisfy exactly one of the following conditions:
    - The points x_i and x_{i+1} are connected by a geodesic edge e_i.
    - There is an edge f_i running between x_i and y_i and another edge g_i between y_i and x_{i+1}, where y_i is the centre of the triangle with vertices x_i, x_{i+1} and $x_i \oplus x_{i+1}$.

3. **Edge identifications:** Edges between two Farey neighbours x_i and x_{i+1} have edge identifications of one of the following types:

   - The edge e_i between x_i and x_{i+1} is identified to another edge e_j between two vertices x_j and x_{j+1}, for some $j \neq i$. In this case, we refer to e_i (and e_j) as a *free edge* of P.
   - The edge e_i between x_i and x_{i+1} is identified to itself by a map which takes x_i to x_{i+1} and vice versa. In this case, we refer to e_i as an *even edge* of P.

   Given two vertices x_i and x_{i+1} which are connected by an edge f_i running between x_i and y_i and another edge g_i between y_i and x_{i+1}, the edges f_i and g_i are identified with each other by a map which takes x_i to x_{i+1}, x_{i+1} to $x_i \oplus x_{i+1}$, and $x_i \oplus x_{i+1}$ to x_i. This edge identifications correspond to an elliptic involution of order 3 centred at y_i. We refer to the edges f_i and g_i as *odd edges* of P.

Remark 3.2.2. We should note that the ordering of the vertices in the Farey sequence induces a natural (anti-clockwise) orientation of the edges (and vertices) of a special polygon. We can use this ordering to recover the underlying Farey sequence of a special polygon P. In particular, if we take the sequence of all vertices in P - starting at ∞ with an anti-clockwise ordering - and we remove each of the odd vertices from this sequence, then this process produces the underlying Farey sequence.

Since each of these edge identifications is an element of $PSL_2(\mathbb{Z})$ and each of the edge identifications maps a pair of Farey neighbours to another pair of Farey neighbours, Corollary 3.1.20 tells us that there is in fact a unique edge identification for each edge e_i in P. Furthermore, we should note that we can give the vertices some arbitrary cyclic ordering (this can be induced from the cyclic ordering of the underlying Farey sequence). Given this ordering each of the edge identifications maps an edge e_i of

P to another edge e_j of P with opposing orientation. For even and free sides, the change in orientation can be realised by noting that edge identification φ_i maps the point x_i to a point x_{j+1} and the point x_{i+1} to the point x_j. For an odd edge f_i running between x_i and y_i, the edge identification maps y_i to itself and maps x_i to x_{i+1} and so the image of the edge $\varphi_i(f_i)$ has a different orientation to g_i.

To construct the maps which identify edges in the special polygon P_σ, we note that if we have a pair of vertices in $\mathcal{F}$, $x_i = \frac{a_i}{b_i}$ and $x_{i+1} = \frac{a_{i+1}}{b_{i+1}}$, and we want a map in $PSL_2(\mathbb{Z})$ taking x_i to $x_{j+1} = \frac{a_{j+1}}{b_{j+1}}$ and x_{i+1} to $x_j = \frac{a_j}{b_j}$, then this map will be of the form:

$$\varphi := \begin{pmatrix} a_j b_i + a_{j+1} b_{i+1} & -a_i a_j - a_{i+1} a_{j+1} \\ b_i b_j + b_{i+1} b_{j+1} & -a_i b_j - a_{i+1} b_{j+1} \end{pmatrix} \tag{3.2.1}$$

To make the process of describing special polygons easier, we introduce the notion of a *Farey symbol*. Given a Farey sequence V, we construct a *Farey symbol* σ by identifying each pair of consecutive vertices x_i, x_{i+1} in the Farey symbol with one of the following intervals:

1. A *free interval* with label $a \in \mathbb{N}$ such that there is another pair of consecutive vertices x_j, x_{j+1}, which also form a free interval and have the same label,

$$x_i \underset{a}{\smile} x_{i+1} \quad \text{and} \quad x_j \underset{a}{\smile} x_{j+1}$$

2. An *even interval*,

$$x_i \underset{\circ}{\smile} x_{i+1}$$

3. An *odd interval*.

$$x_i \underset{\bullet}{\smile} x_{i+1}$$

Of course, there is a natural duality between special polygons and Farey symbols. In particular, the free edges in special polygons correspond to free intervals of Farey symbols, the even edges correspond to the even intervals, and the odd edges correspond to odd intervals. In particular, given a special polygon P, one can construct a corresponding Farey symbol σ_P, and given a Farey symbol σ one can construct a corresponding special polygon P_σ.

For any Farey symbol σ (or the corresponding special polygon P_σ), there is a collection of maps (explicitly given by equation 3.2.1) corresponding to the edge identifications. We can therefore define Φ_σ to be group generated by the edge identifications of σ. By the Poincaré Polyhedron Theorem [Ser13, Theorem 6.14], P_σ is a fundamental domain for Φ_σ. Theorem (6.1) in [Kul91] states that the edge identifications for any Farey symbol σ, form an independent set of generators for Φ_σ. Moreover, the following Theorem explains the importance of special polygons.

Theorem 3.2.3. ([Kul91, Theorem 3.2 and 3.3]) *Every special polygon P_σ is a fundamental domain for a finite index subgroup Φ_σ of $PSL_2(\mathbb{Z})$, and this subgroup Φ_σ is generated by the side pairings of P_σ. Every finite index subgroup Φ of $PSL_2(\mathbb{Z})$ admits a special polygon P_Φ as a fundamental domain.*

See Fig. 3.7 for examples of two special polygons, which are fundamental domains of $\Gamma_0(7)$ and $\Gamma_0(11)$. The corresponding Farey symbols are:

$$\infty \underbrace{\quad}_{0} 0 \underbrace{\quad}_{\bullet} \tfrac{1}{2} \underbrace{\quad}_{\bullet} 1 \underbrace{\quad}_{0} \infty\ ,$$

and:

$$\infty \underbrace{\quad}_{0} 0 \underbrace{\quad}_{1} \tfrac{1}{3} \underbrace{\quad}_{2} \tfrac{1}{2} \underbrace{\quad}_{1} \tfrac{2}{3} \underbrace{\quad}_{2} 1 \underbrace{\quad}_{0} \infty\ .$$

Special Polygons as Fundamental Domains of $\Gamma_0(n)$

For any $\Gamma_0(n)$, we can take a special polygon P_σ to be our fundamental domain such that the y-axis I and $I+1$ are paired sides of P_σ. In particular, we can take

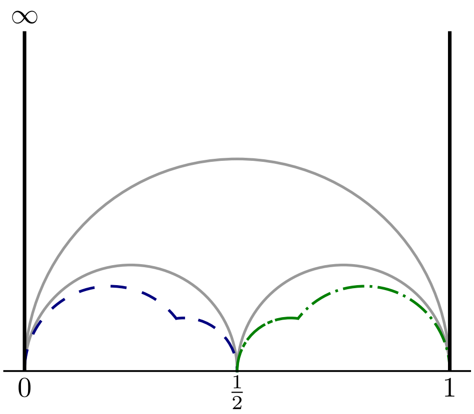

(a) An image of a special polygon which is a fundamental domain for $\Gamma_0(7)$.

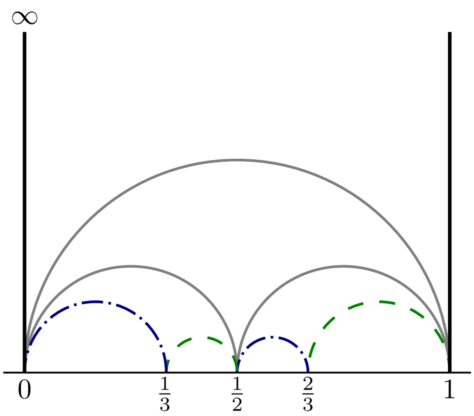

(b) An image of a special polygon which is a fundamental domain for $\Gamma_0(11)$.

Figure 3.7: Special polygons which are fundamental domains $\Gamma_0(7)$ and $\Gamma_0(11)$. Internal edges show the structure of $\mathcal{F}$ in this region. External edges are identified by colour and line type.

$x_0 = 0$ and $x_r = 1$ in the corresponding Farey sequence. This is due to the fact that $\left(\begin{smallmatrix} 1 & 1 \\ 0 & 1 \end{smallmatrix}\right) \in \Gamma_0(n)$, for all $n \in \mathbb{N}_{>1}$.

For p prime, we can find a Farey symbol σ (for which P_σ is a fundamental domain of $\Gamma_0(p)$), in which the vertices are symmetric in the line $x = \frac{1}{2}$ to ∞. In other words, the underlying Farey sequence will be of the form $\{\infty, 0, x_1, x_2, \ldots x_2', x_1', 1, \infty\}$, where $x_i' = 1 - x_i$. The term $\frac{1}{2}$ will be in every Farey symbol of $\Gamma_0(p)$ for $p \geq 5$. This is due to the fact that the line 0 to $\frac{1}{2}$ separates $\mathbb{H}$ into two regions: one containing the vertex 1 and the other containing all other neighbours of 0. Therefore, to get a Farey symbol containing the terms 0 and 1, the underlying Farey sequence must either only contain the vertices ∞, 0 and 1 or the sequence must contain the vertex $\frac{1}{2}$. If we have either an odd or even interval, then the interval identifications are symmetric in the line $\frac{1}{2}$ to ∞. However, for free intervals we have antisymmetry, i.e. the free interval labelled a will be replaced with the label a' in this symmetry. Due to the symmetry of the vertices in the line $x = \frac{1}{2}$ to ∞ and the pseudo-symmetry of the interval identifications, we will shorten the sequence up to the term $\frac{1}{2}$ (for $p \geq 5$, since for $p = 2, 3$ we only use the vertices ∞, 0 and 1). Similarly, due to identification of the y-axis I with $I + 1$, we will not include the terms $\infty \underbrace{}_{0} 0$ or $1 \underbrace{}_{0} \infty$ with the identification between these edges being implicit. For example, we would write:

$$\left\{ 0 \underbrace{}_{1} x_1 \underbrace{}_{1'} x_2 \underbrace{}_{\bullet} x_3 \underbrace{}_{\circ} \tfrac{1}{2} \mid \text{refl.} \right\}$$

instead of

$$\left\{ \infty \underbrace{}_{0} 0 \underbrace{}_{1} x_1 \underbrace{}_{1'} x_2 \underbrace{}_{\bullet} x_3 \underbrace{}_{\circ} \tfrac{1}{2} \underbrace{}_{\circ} x_3' \underbrace{}_{\bullet} x_2' \underbrace{}_{1} x_1' \underbrace{}_{1'} 1 \underbrace{}_{0} \infty \right\}.$$

We can explicitly state how many odd, even and free intervals there will be in each fundamental domain of $\Gamma_0(n)$. This can be derived from the properties of the quotient space $\Gamma_0(n) \backslash \mathbb{H}$. For a prime $p \geq 5$ with $p \equiv 1 \bmod 3$ there are exactly two odd intervals (either side of $\frac{1}{2}$), otherwise there are no odd intervals. Similarly, if $p \equiv 1 \bmod 4$ there are exactly two even intervals (either side of $\frac{1}{2}$), otherwise there are no even intervals. If $p \equiv 1 \bmod 3$, the Farey symbol has $\frac{p+2}{3}$ terms, otherwise the Farey symbol has $\frac{p+4}{3}$ terms.

Given Φ a subgroup of $PSL_2(\mathbb{Z})$, we can use the Riemann-Hurwitz formula to relate geometric invariants the quotient space $\Phi \backslash \mathbb{H}$, as follows:

$$d = 3e_2 + 4e_3 + 12g + 6t - 12$$

where

- d is the index of $[PSL_2(\mathbb{Z}) : \Phi]$

- e_2 is the number orbifold points in $\Phi\backslash\mathbb{H}$ with cone angle π (or equivalently the number of even intervals in a corresponding special polygon)
- e_3 is the number orbifold points in $\Phi\backslash\mathbb{H}$ with cone angle $\frac{2\pi}{3}$ (or equivalently the number of odd intervals in a corresponding special polygon)
- g is the genus of $\Phi\backslash\mathbb{H}$
- t is the number of cusps for $\Phi\backslash\mathbb{H}$

For $\Phi = \Gamma_0(n)$:

$$d = n\prod_{q|n}\left(1 + \frac{1}{q}\right),$$
$$t = \sum_{a|n}\varphi\left(gcd\left(a, \frac{n}{a}\right)\right),$$

where q is a prime number, $a \in \mathbb{N}$ and φ is the Euler totient function.

Calculating this information for $\Gamma_0(p)$, we observe that the quotient space $\Gamma_0(p)\backslash\mathbb{H}$ will have 2 punctures, e_2 even intervals, e_3 odd intervals and genus g. The above relation then reduces to:

$$p + 1 = 3e_2 + 4e_3 + 12g.$$

3.2.3 Constructing the Integer Multiplication Algorithm for Cutting Sequences.

In order to discuss how we construct the integer multiplication algorithm, we should first explain exactly how we want this algorithm to work. The algorithm will take the form of a *deterministic finite automaton with output* (DFAO) [AS03]. This is a sextuple $\mathcal{M}_n = (\mathcal{S}_n, \mathcal{I}_n, \mathcal{O}_n, \delta_n, \tau_n, q_0)$, where:

- $\mathcal{S}_n$ is a finite set of *states*.
- $\mathcal{I}_n$ is the *input alphabet* (a finite set of all possible inputs).
- $\mathcal{O}_n$ is the *output alphabet* (a finite set of all possible outputs).

- $\delta_n : \mathcal{S}_n \times \mathcal{I}_n \to \mathcal{S}_n$ is the *transition function.*
- $\tau_n : \mathcal{S}_n \times \mathcal{I}_n \to \mathcal{O}_n$ is the *output function.*
- q_0 is the initial state.

Keeping the above notions in mind, we will produce the following motivating construction. There are a few issues in this construction, which we will discuss and resolve later. However, this construction gives a good overview and idea of how the multiplication algorithm will work.

A Motivating Construction of the Multiplication Algorithm

We take P_n to be a special polygon which is a fundamental domain of $\Gamma_0(n)$ containing I and $I+1$ as paired sides. In this example, the edges on the boundary of P_n represent the *states* $\mathcal{S}_n$ of our algorithm. Given two edges E_i and E_j of ∂P_n, we can construct a path λ_{E_i,E_j}, which starts at E_i and terminates at E_j and this path is unique up to homotopy relative to E_i and E_j. Given a starting edge E_i, we can then construct the set of paths $\Lambda_n(E_i)$, which start at E_i and terminate at some arbitrary edge E_j (considered up to relative homotopy). Since P_n has only finitely many sides, the set $\Lambda_n(E_i)$ is finite. Since the paths we wish to construct start at the y-axis I and I is on the boundary of P_n, the edge I will be our *initial state* q_0.

For each starting edge E_i, we define the *input alphabet* to be:

$$\mathcal{I}_n(E_i) := \{(\lambda_{E_i,E_j}, T_{\{1,n\}}) : \lambda_{E_i,E_j} \in \Lambda_n(E_i)\}.$$

This is the set of cutting sequences of paths which start at E_i, taken relative to $T_{\{1,n\}}$.

The *output alphabet* is similarly given by:

$$\mathcal{O}_n(E_i) := \{(\lambda_{E_i,E_j}, T_{\{n,n\}}) : \lambda_{E_i,E_j} \in \Lambda_n(E_i)\},$$

This is the set of cutting sequences of paths which start at E_i, taken relative to $T_{\{n,n\}}$, i.e. we have performed triangulation replacement on the cutting sequences of the input alphabet.

One thing to note is that the paths $\Lambda_n(E_i)$ encode both the input alphabet $\mathcal{I}_n(E_i)$ and the output alphabet $\mathcal{O}_n(E_i)$. As a result, it will be more natural for us to consider the output function τ_n and transition function δ_n as a function on $\Lambda_n(E_i)$ instead of as a function on the input alphabet, i.e. $\tau_n : \Lambda_n(E_i) \to \mathcal{O}_n(E_i)$ and $\delta_n : \Lambda_n(E_i) \to \mathcal{S}_n$. With the above observations, we can define the *output function* $\tau_n : \Lambda_n(E_i) \to \mathcal{O}_n(E_i)$ to be the function given by $\tau_n(\lambda_{E_i,E_j}) = (\lambda_{E_i,E_j}, T_{\{n,n\}})$.

In order to construct the transition function δ_n, we need to observe how the fundamental domains in the tessellation of P_n piece together. This in turn allows us to observe how the paths in Λ_n can be concatenated. If we take λ_{I,E_1} to be our initial path in P_n, which starts at I and terminates at some edge $E_1 \in \mathcal{S}_n$, then our next path $\lambda_{\overline{E_1},E_2}$, should start at $\overline{E_1}$, where $\overline{E_1}$ is the edge paired to E_1 under the edge identifications of P_n. If φ_1 is the map taking $\overline{E_1}$ to E_1, then we see that this allows us to concatenate the paths λ_{I,E_1} and $\varphi_1(\lambda_{\overline{E_1},E_2})$ in a natural way:

- The map φ_1 maps $\overline{E_1}$ to E_1 and, therefore, λ_{I,E_1} terminates at the edge E_1 and $\varphi_1(\lambda_{\overline{E_1},E_2})$ starts at the edge E_1.
- The map φ_1 does not map P_n to P_n and, therefore, the paths λ_{I,E_1} and $\varphi_1(\lambda_{\overline{E_1},E_2})$ are contained in adjacent copies of P_n.
- Therefore, the paths λ_{I,E_1} and $\varphi_1(\lambda_{\overline{E_1},E_2})$ must have compatible directions of approach/departure.

Here, we will slightly abuse notation and denote the concatenated path as $\lambda_{I,E_1} \circ \lambda_{\overline{E_1},E_2}$, instead of $\lambda_{I,E_1} \circ \varphi_1(\lambda_{\overline{E_1},E_2})$. By a similar argument, if the path $\lambda_{\overline{E_1},E_2}$ terminates at E_2, then the next path $\lambda_{\overline{E_2},E_3}$ should start at the edge $\overline{E_2}$, where $\overline{E_2}$ is the edge of P_n identified to E_2. In this case, the paths $\lambda_{\overline{E_1},E_2}$ and $\varphi_2(\lambda_{\overline{E_2},E_3})$ will be compatible. As a result, we can concatenate $\varphi_1\varphi_2(\lambda_{\overline{E_2},E_3})$ to our path $\lambda_{I,E_1} \circ \lambda_{\overline{E_1},E_2}$.

Again, we will abuse notation and write this path as $\lambda_{I,E_1} \circ \lambda_{\overline{E_1},E_2} \circ \lambda_{\overline{E_2},E_3}$, instead of $\lambda_{I,E_1} \circ \varphi_1(\lambda_{\overline{E_1},E_2}) \circ \varphi_1\varphi_2(\lambda_{\overline{E_2},E_3})$. This process generalises and we can construct an arbitrarily long path λ in $\mathbb{H}$ by repeatedly concatenating paths $\{\lambda_k\}$ in P_n, as long as we have that λ_k terminates at the edge E_k and λ_{k+1} starts at the edge $\overline{E_k}$. Here, we denote this path as $\lambda = \lambda_{I,E_1} \circ \lambda_{\overline{E_1},E_2} \circ \lambda_{\overline{E_2},E_3} \circ \dots$. Since the edges are the states of this algorithm, we see that the above process induces the *transition function* $\delta_n : \Lambda_n(E_i) \to \mathcal{S}_n$. In particular, if $\lambda_{\overline{E_{k-1}},E_k}$ terminates at the edge E_k, then $\delta_n(\lambda_{\overline{E_{k-1}},E_k}) := \overline{E_k}$.

Assuming we have constructed all the above structures and these structures are well defined, the algorithm would work as follows:

The Motivating Cutting Sequence Multiplication Automaton

Let W word over the alphabet $\{L, R\}$ corresponding to the cutting sequence which we wish to multiply by n, let V be the empty word $\varepsilon \in \{L, R\}^*$, and let I be our initial state.

1. Set $k = 0$ and take $W_0 = W$ and $E_0 = I$.
2. Repeat the following steps:
   (a) Find a word $U_k \in \mathcal{I}_n(E_i)$ which is a prefix of W_k and let λ_k be the corresponding path.
   (b) Find $V_k = \tau_n(\lambda_k) = (\lambda_k, T_{\{n,n\}}) \in \mathcal{O}_n$ and append V_k to V.
   (c) Set $E_{k+1} = \delta_n(\lambda_k)$
   (d) Write $W_k = U_k W_{k+1}$.
   (e) Set $k = k + 1$.
3. End of algorithm.

The initial word W can then be decomposed into a product of sub-words, i.e. $W = U_0U_1U_2\cdots$, and the output word can be written as $V = V_0V_1V_2\cdots$.

Issues with the Motivating Construction

Whilst the above construction is very good for presenting the general ideas behind the algorithm, there are two main issues:

Issue 1: For every edge E_i and every word $W \in \{L,R\}^{\mathbb{N}}$ there may not always be a prefix U_k of W which is also in the input alphabet $\mathcal{I}_n(E_i)$. Furthermore, this prefix may not correspond to a unique path in P_n.

To resolve Issue 1, we will consider special polygons P_n without odd edges and show a slightly stronger condition: that each of the input alphabets $\mathcal{I}_n(E_i)$ is a *base* over $\{L,R\}$. Here, we define a *base* over the alphabet $\{L,R\}$ to be a finite set of words $\mathcal{B} = \{U_k \in \{L,R\}^*\}$ such that for any word $W \in \{L,R\}^{\mathbb{N}}$, there is a unique element in $\mathcal{B}$ which is a prefix of W. An example of a base over $\{L,R\}$ would be $\mathcal{B} = \{L, RL, RRL, RRR\}$. We then discuss special polygons with odd edges and see that the input alphabets $\mathcal{I}_n(E_i)$ are also base - relative to some small concession when E_i is an odd edge. In doing this, we will also show that the cutting sequences $(\lambda_k, T_{\{1,n\}})$ are well-defined for all paths λ_k in P_n - again up to some small concession when P_n contains odd edges.

Issue 2: The cutting sequences of the paths λ_k are not necessarily well-defined relative to $T_{\{n,n\}}$. As a result, the path $\lambda_k \circ \lambda_{k+1}$, formed by concatenating two paths λ_k and λ_{k+1} in P_n, is not necessarily reduced relative to $\frac{1}{n}\mathcal{F}$. This may be the case, even if both λ_k and λ_{k+1} are reduced relative $T_{\{n,n\}}$. See Fig. 3.8.

To resolve Issue 2, we will show that if we subdivide the edges to produce some additional states, then we only have to consider paths λ_k which intersect $T_{\{n,n\}}$ in a minimal way. In this case, we have a natural way to interpret the cutting sequence and the concatenated paths will be reduced to $\frac{1}{n}\mathcal{F}$. We will also show that in this construction, the paths λ_k will still satisfy nice properties relative to $T_{\{1,n\}}$.

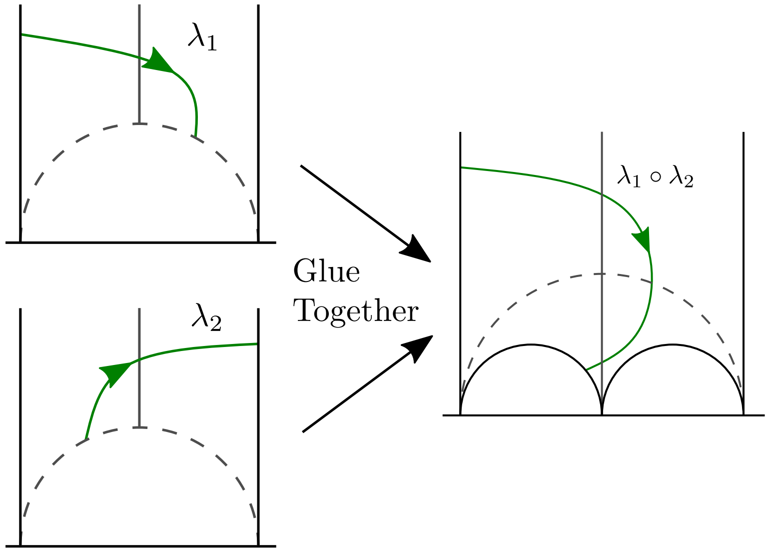

Figure 3.8: An image showing how two paths λ_1 and λ_2, which are both reduced relative $T_{\{2,2\}}$, can be concatenated together to form a path $\lambda_1 \circ \lambda_2$ which is not reduced relative to $\frac{1}{2}\mathcal{F}$.

Resolving Issue 1 – Special Polygons without Odd Edges

Issue 1: For every edge E_i and every word $W \in \{L, R\}^{\mathbb{N}}$ there may not always be a prefix U_k of W which is also in the input alphabet $\mathcal{I}_n(E_i)$.

In the case that P_n is a special polygon of $\Gamma_0(n)$, which does not contain any odd edges, we can explicitly show that all of the paths in $\Lambda_n(E_i)$ have a well-defined cutting sequence relative to $T_{\{1,n\}}$, for every edge E_i in ∂P_n. This follows since every edge on the boundary of P_n is also an edge of $\mathcal{F}$ by construction. Thus, every edge of ∂P_n is also in $T_{\{1,n\}}$. Therefore, any path from one edge E_i of ∂P_n to another edge E_j of ∂P_n will have a well-defined cutting sequence relative to $T_{\{1,n\}}$, since E_i and E_j are both edges of $T_{\{1,n\}}$.

Using these ideas, we can show that the input alphabet $\mathcal{I}_n(E_i)$ is a base, and thus, Issue 1 is resolved quite simply in this case.

Lemma 3.2.4. *Assume that P_n is a special polygon which is a fundamental domain of $\Gamma_0(n)$ that does not contain any odd edges. Then, for any edge E_i the input alphabet $\mathcal{I}_n(E_i)$ forms a base.*

Proof. Assume that P_n is a special polygon of $\Gamma_0(n)$ that does not contain any odd edges. Since the polygon P_n is formed by using edges of the Farey Tessellation $\mathcal{F}$, every edge in the boundary of P_n is also an edge of $\mathcal{F}$. As seen in Section 3.1.1, given an infinite word W, an edge E in $\mathcal{F}$ and a direction of departure, we can construct an infinite path λ_W such that $W = (\lambda_W, \mathcal{F})$, and this path is unique up to homotopy relative to E.

We will take E_i to be any edge on the boundary of P_n and assume that the direction of departure heading into P_n is positive. We will then construct the reduced path λ_W, which starts at the edge E_i and has positive direction of departure, such that $(\lambda_W, \mathcal{F}) = W$, for some arbitrary word $W \in \{L, R\}^{\mathbb{N}}$. The special polygon P_n is made up of finitely many triangles in $\mathcal{F}$ and so the path λ_W must intersect the boundary of P_n at a unique edge $E_j \neq E_i$ or at a vertex V. We can then cut λ_W along the edge E_j to produce two paths $\lambda_{E_i,E_j} \in \Lambda_n(E_i)$ and λ'_W such that $\lambda_W = \lambda_{E_i,E_j} \circ \lambda'_W$. Note that λ_{E_i,E_j} is completely determined (up to relative homotopy) by its initial edge and terminal edge. Therefore, if we have an alternative decomposition $\lambda_W = \lambda_{E_i,E_k} \circ \lambda'_V$, then we must have $E_k = E_j$ and $\lambda'_V = \lambda_W$.

Since E_j is an edge of $\mathcal{F}$ and λ_W is assumed to be reduced relative to $\mathcal{F}$, the path λ_{E_i,E_j} has a well-defined cutting sequence. Furthermore, the paths λ_{E_i,E_j} and λ'_W will be compatible. Therefore, it follows that $W = (\lambda, \mathcal{F}) = (\lambda_{E_i,E_j}, \mathcal{F}) \cdot (\lambda'_W, \mathcal{F})$. In particular, $(\lambda_{E_i,E_j}, \mathcal{F})$ is a prefix of W. Since $T_{\{1,n\}} := P_n \cap \mathcal{F}$ and λ_{E_i,E_j} is a path in P_n, we can conclude that $(\lambda_{E_i,E_j}, \mathcal{F}) = (\lambda_{E_i,E_j}, T_{\{1,n\}}) \in \mathcal{I}_n(E_i)$ and $(\lambda_{E_i,E_j}, \mathcal{F})$ is a unique prefix of W. Since W was arbitrarily chosen, this must be true for all $W \in \{L, R\}^{\mathbb{N}}$, and so $\mathcal{I}_n(E_i)$ is a base. □

Resolving Issue 1 – Special Polygons Containing Odd Edges

Recall from Section 3.2.2 that to form the odd edges, we first take two Farey neighbours x_i and x_{i+1} and find their Farey sum $x_i \oplus x_{i+1}$ (see Section 3.1.2). The triple of points x_i, x_{i+1} and $x_i \oplus x_{i+1}$ form a triangle in $\mathcal{F}$. We can then take the point y_i

to be the centre of the triangle. The *odd triangle* τ is then the triangle in $\mathbb{H}$ with vertices x_i, x_{i+1} and y_i, and the edges between x_i and y_i and y_i and x_{i+1} are referred to as *odd edges*. Note that these odd edges are identified under the side pairings. It will be useful to also consider the edge in $\mathcal{F}$ between the points x_i and x_{i+1}, which we refer to as a *supporting edge*. An odd triangle is then made up of two odd edges and one supporting edge.

If P_n is a special polygon which contains some odd edges, then it is worth emphasising that these odd edges are not part of the Farey tessellation and, therefore, not an edge in $T_{\{1,n\}}$. As a result, if E_j and $\overline{E_j}$ are the two odd edges of an odd triangle and E_i is any other edge of ∂P_n, then the paths λ_{E_i,E_j} and $\lambda_{E_i,\overline{E_j}}$ do not have a well-defined cutting sequence relative to $T_{\{1,n\}}$. The same is true if our paths start at an odd edge. To remedy this, we will extend our definition of $T_{\{1,n\}}$ to include these odd edges.

We note that if E_s is the supporting edge of the odd triangle containing the odd edges E_j and $\overline{E_j}$, then since E_s is in the Farey tessellation, the paths λ_{E_i,E_s} will have a well defined cutting sequence (assuming that E_i is not an odd edge). One can then note that if the odd triangle has edges E_j, E_s and $\overline{E_j}$ labelled clockwise, any reduced path λ_{E_s,E_j} from E_s to E_j will always form a right triangle with the odd triangle, and any reduced path $\lambda_{E_s,\overline{E_j}}$ from E_s to $\overline{E_j}$ will always form a left triangle. We will interpret the paths λ_{E_s,E_j} and $\lambda_{E_s,\overline{E_j}}$ as forming "half triangles" with $T_{\{1,n\}}$ and give these paths the cutting sequences $(\lambda_{E_s,E_j}, T_{\{1,n\}}) = R^{\frac{1}{2}}$ and $(\lambda_{E_s,\overline{E_j}}, T_{\{1,n\}}) = L^{\frac{1}{2}}$. As a result, we can decompose the path λ_{E_i,E_j} from E_i to E_j into two compatible sub-paths: λ_{E_i,E_s} and λ_{E_s,E_j}. This would have the cutting sequence $(\lambda_{E_i,E_j}, T_{\{1,n\}}) = (\lambda_{E_i,E_s}, T_{\{1,n\}}) \cdot (\lambda_{E_s,E_j}, T_{\{1,n\}}) = (\lambda_{E_i,E_s}, T_{\{1,n\}}) \cdot R^{\frac{1}{2}}$.

When considering paths λ_{E_j,E_k} which start at an odd edge E_j, we will similarly break these into a path λ_{E_j,E_s} and followed by the path λ_{E_s,E_k}, where E_s is the supporting edge. The path λ_{E_j,E_s} forms a left triangle with the odd triangle and the path $\lambda_{\overline{E_j},E_s}$ forms a right triangle with the odd triangle. Again, we consider these as half triangles in the cutting sequence relative to $T_{\{1,n\}}$. We will not consider the

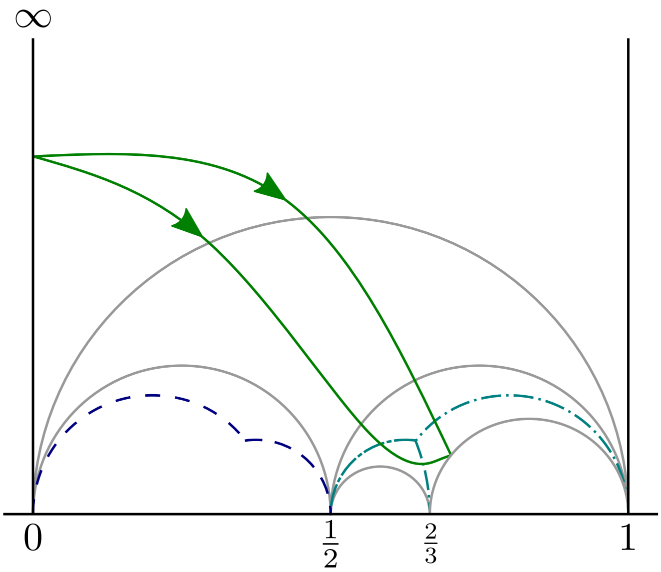

Figure 3.9: An image showing how we can homotope a path so that it only ever intersects one odd edge in a row. The dashed lines represent the odd edges.

path $\lambda_{E_j,\overline{E_j}}$ or the path $\lambda_{\overline{E_j},E_j}$ in this setting, because, up to homotopy, the path formed by concatenating λ_{E_i,E_j} with the path $\lambda_{\overline{E_j},E_j}$ is equivalent to the path $\lambda_{E_i,\overline{E_j}}$ - the path that went to the edge $\overline{E_j}$ in the first place. Given any path in $\mathbb{H}$, we can homotope the path so that it only ever intersects one odd edge in a row. See Fig. 3.9. Since the edges E_j and $\overline{E_j}$ are identified via the edge identifications, the path λ_{E_s,E_j} with cutting sequence $(\lambda_{E_s,E_j}, T_{\{1,n\}}) = R^{\frac{1}{2}}$ is always followed by the path $\lambda_{\overline{E_j},E_s}$ with cutting sequence $(\lambda_{\overline{E_j},E_s}, T_{\{1,n\}}) = R^{\frac{1}{2}}$. As a result, the path $\lambda_{\overline{E_j},E_s}$ "completes the triangle" relative to the path λ_{E_s,E_j}. Similarly, the path $\lambda_{E_s\overline{E_j}}$ with cutting sequence $(\lambda_{E_s,\overline{E_j}}, T_{\{1,n\}}) = L^{\frac{1}{2}}$ is always followed by the path λ_{E_j,E_s} with cutting sequence $(\lambda_{E_j,E_s}, T_{\{1,n\}}) = L^{\frac{1}{2}}$. We can view these odd edges as being part of an extended tessellation of $T_{\{1,n\}}$, and this guarantees that the paths in P_n have well-defined cutting sequences relative to $T_{\{1,n\}}$.

This construction allows us to distinguish the paths λ_{E_i,E_j} and $\lambda_{E_i,\overline{E_j}}$ by their cutting sequences relative to $T_{\{1,n\}}$. More precisely, we have $(\lambda_{E_i,E_j}, T_{\{1,n\}}) = (\lambda_{E_i,E_s}, T_{\{1,n\}}) \cdot R^{\frac{1}{2}}$ and $(\lambda_{E_i,\overline{E_j}}, T_{\{1,n\}}) = (\lambda_{E_i,E_s}, T_{\{1,n\}}) \cdot L^{\frac{1}{2}}$. Using this information, we can extend the proof of Lemma 3.2.4 to fundamental domains with odd edges.

Corollary 3.2.5. *Let P_n be a special polygon, which is a fundamental domain of $\Gamma_0(n)$ containing odd edges. Then, for any edge E_i which is not an odd edge, the input alphabet $\mathcal{I}_n(E_i)$ forms a base. If E_i is an odd edge, then the input alphabet $\mathcal{I}_n(E_i)$ will be of the form $L^{\frac{1}{2}}\mathcal{B}$ or $R^{\frac{1}{2}}\mathcal{B}$, where $\mathcal{B}$ is a base.*

Proof. Assume that E_i is not an odd edge and λ_W is an arbitrary path starting at E_i with positive direction of departure (relative to P_n) with cutting sequence $W = (\lambda_W, \mathcal{F})$, where W is some arbitrary word in $\{L, R\}^{\mathbb{N}}$. Then by the arguments in the proof of Lemma 3.2.4, we can uniquely decompose λ_W into two sub-paths: $\lambda_{E_i,E_j} \in \Lambda_n(E_i)$ and λ'_W. Since λ_{E_i,E_j}, and λ'_W are compatible, we have $W = (\lambda_W, \mathcal{F}) = (\lambda_{E_i,E_j}, \mathcal{F}) \cdot (\lambda'_W, \mathcal{F}) = (\lambda_{E_i,E_j}, T_{\{1,n\}}) \cdot (\lambda'_W, \mathcal{F})$. Therefore, the cutting sequence $(\lambda_{E_i,E_j}, T_{\{1,n\}})$ is a prefix of W. Since paths in $\Lambda_n(E_i)$ each have distinct cutting sequences relative to $T_{\{1,n\}}$, the input alphabet $\mathcal{I}_n(E_i)$ is a base.

If E_i is an odd edge, then we can take $\overline{E_i}$ to be the other odd edge of the odd triangle, and take E_s to be supporting edge. If $\Lambda_n(E_s)$ is the set paths in P_n starting at E_s and terminating at E_k with $E_k \neq E_i$ and $E_k \neq \overline{E_i}$, then we can use the above arguments to see that the set of cutting sequences $\{(\lambda_{E_s,E_k}, T_{\{1,n\}}) : \lambda_{E_s,E_k} \in \Lambda_n(E_s)\}$ is a base. In the notation of the corollary, $\mathcal{B} := \{(\lambda_{E_s,E_k}, T_{\{1,n\}}) : \lambda_{E_s,E_k} \in \Lambda_n(E_s)\}$. The paths λ_{E_i,E_s} and $\lambda_{\overline{E_i},E_s}$ contribute either an initial $L^{\frac{1}{2}}$ term or $R^{\frac{1}{2}}$ term, depending on whether the paths λ_{E_i,E_s} and $\lambda_{\overline{E_i},E_s}$ intersect the odd triangle to form a left triangle or a right triangle. □

Resolving Issue 2

Issue 2: The cutting sequences of the paths λ_k are not necessarily well-defined relative to $T_{\{n,n\}}$. As a result, the path $\lambda_k \circ \lambda_{k+1}$, formed by concatenating two paths λ_k and λ_{k+1} in P_n, is not necessarily reduced relative to $\frac{1}{n}\mathcal{F}$. This may be the case, even if both λ_k and λ_{k+1} are reduced relative $T_{\{n,n\}}$.

The main problem here is that the fundamental domain P_n is unlikely to intersect the edges in $\frac{1}{n}\mathcal{F}$ in a nice way and so the polygons in the tessellation $T_{\{n,n\}}$ of P_n, will not all be triangles. As a result, it is not immediately clear how one could compute the cutting sequence $(\lambda_k, T_{\{n,n\}})$ for λ_k some path in P_n (or even what a cutting sequence would mean in this case).

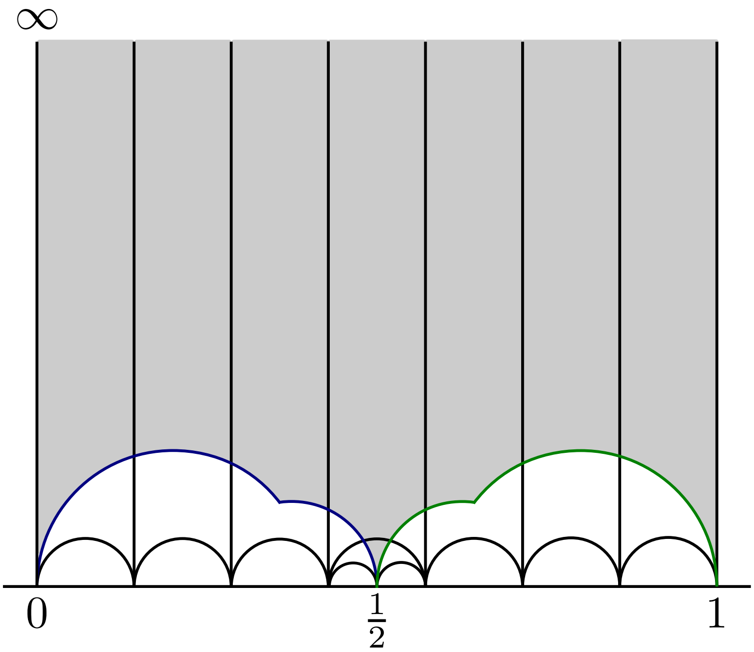

Figure 3.10: An example of the finite set of triangles $\mathcal{C}_7$ which are contained in $\frac{1}{7}\mathcal{F}$ and cover P_7. The special polygon P_7 is shaded grey.

In order to define a cutting sequence for λ_k, we note that for each edge E_i on the boundary of P_n only a finite number of edges in $\frac{1}{n}\mathcal{F}$ will intersect E_i. As a result, we can find a finite number of triangles in $\frac{1}{n}\mathcal{F}$, which cover P_n. We denote this covering as $\mathcal{C}_n$. See Fig. 3.10. This covering is given by taking all of the triangles which intersect P_n non-trivially, i.e. the triangles that intersect P_n, and this intersection is not just an edge or a vertex of P_n. For any path λ_k in P_n, we can take E_1 to be the first edge of $\mathcal{C}_n$ that λ_k intersects, and take E_ℓ to be the last edge of $\mathcal{C}_n$ that λ_k intersects. Therefore, we can use the structure this covering $\mathcal{C}_n$ to determine the cutting sequence of the sub-path $\widetilde{\lambda_k}$ of λ_k, which runs from E_1 to E_ℓ. Whilst this gives us some structure to form a cutting sequence, we discard part of the path λ_k and so concatenated terms $\lambda_k \circ \lambda_{k+1}$ will not necessarily have the cutting sequence $(\widetilde{\lambda_k}, T_{\{n,n\}}) \cdot (\widetilde{\lambda_{k+1}}, T_{\{n,n\}})$. In particular, λ_k will terminate inside some triangle τ of

$\frac{1}{n}\mathcal{F}$ and λ_{k+1} will start inside this triangle, and so by only considering the cutting sequences of the paths $\widetilde{\lambda_k}$ and $\widetilde{\lambda_{k+1}}$, we miss out on how the concatenated path $\lambda_k \circ \lambda_{k+1}$ intersects τ. As a result, the cutting sequence of $(\lambda_k \circ \lambda_{k+1}, \frac{1}{n}\mathcal{F})$ is given by $(\widetilde{\lambda_k}, T_{\{n,n\}}) \cdot Y \cdot (\widetilde{\lambda_{k+1}}, T_{\{n,n\}})$, where Y encodes how $\lambda_k \circ \lambda_{k+1}$ intersects τ. In general, we want to avoid paths which introduce bigons when concatenated. Fortunately, we will see that if we put some restrictions on the choices of λ_k, we can do exactly that. Not only that, but under these restrictions, the way λ_{k+1} intersects τ completely determines how $\lambda_k \circ \lambda_{k+1}$ intersects τ.

Restricting the paths in P_n

Assume that E is an edge of P_n and E' is and edge of $\frac{1}{n}\mathcal{F}$, and these edges non-trivially intersect. We also assume that λ_k and λ_{k+1} are two paths which are compatible relative to E and want to see how the concatenated path $\lambda_k \circ \lambda_{k+1}$ intersects both E and E'. One thing to note is that if $\lambda_k \circ \lambda_{k+1}$ intersects E' and does not form a bigon, then we can always homotope this intersection to occur relative to the path λ_{k+1}. See Fig. 3.11. In particular, we will never need the path λ_k to intersect the edge E', since if we do not need the concatenated path $\lambda_k \circ \lambda_{k+1}$ to intersect E', then having λ_k intersecting E', will force us to introduce a bigon - see Fig. 3.11 and if we need the concatenated path $\lambda_k \circ \lambda_{k+1}$ to intersect E', we can make sure that this happens along the path λ_{k+1} - see Fig. 3.12.

In particular, if the path λ_k terminates at an edge E in ∂P_n, then λ_k should not intersect any edge of $\frac{1}{n}\mathcal{F}$ that intersects E. If λ_k satisfies this condition, then we will say that λ_k is a *minimal path* relative to E.

Recall that every geodesic splits the plane $\mathbb{H}$ into two regions. If γ is a geodesic intersecting P_n (and γ is not a boundary edge), then γ also splits P_n into two regions. In this setting, if λ_k is a path in P_n terminating at an edge E_i, then the starting point and terminal point of λ_k should lie in the same region relative to all edges in $\frac{1}{n}\mathcal{F}$ which intersect the edge E_i. In fact, if the starting point of λ_k and the terminal

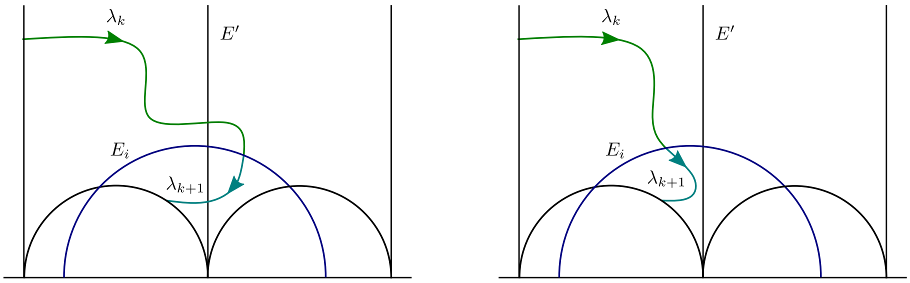

Figure 3.11: An example of how if λ_k intersects an edge $E^{'}$ and the concatenated path $\lambda_k \circ \lambda_{k+1}$ forms a bigon with $E^{'}$ (left), then we could have taken representatives of λ_k and λ_{k+1}, such that λ_k never intersected $E^{'}$ in the first place (right).

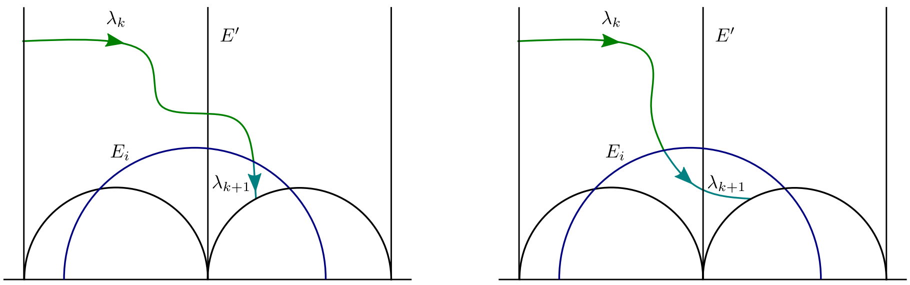

Figure 3.12: An example of how if the concatenated path $\lambda_k \circ \lambda_{k+1}$ intersects $E^{'}$ (left) and does not form a bigon, then, we can homotope $\lambda_k \circ \lambda_{k+1}$ such that λ_{k+1} intersects $E^{'}$ and λ_k does not (right).

point λ_k both lie in the same region relative to edges in $\frac{1}{n}\mathcal{F}$ which intersect the edge E_i, and λ_k is reduced relative to $\frac{1}{n}\mathcal{F}$, i.e. it does not form bigons, then λ_k will be a minimal path relative to E_i.

The edges in $T_{\{n,n\}}$ split each edge E_i of ∂P_n into a finite number of distinct regions, which we label $E_{i,j}$. It is worth noting that if E_i and $\overline{E_i}$ are edges paired under the edge identifications, then up to relabelling each segment $E_{i,j}$ will be identified to a corresponding region $\overline{E_{i,j}}$ under this edge identification. Note that these segments nicely match up because the edge identifications are elements of $\Gamma_0(n)$ and $\Gamma_0(n)$ acting on $T_{\{n,n\}}$ produces the triangulation $\frac{1}{n}\mathcal{F}$. For each segment $E_{i,j}$, we can assign a base point $B_{i,j}$ and refer to the set of points for a given edge as the *base points for* E_i. We can do this in such a way that the base point $B_{i,j}$ is paired to the base point

$\overline{B_{i,j}}$, where $\overline{B_{i,j}}$ is the base point corresponding to the region $\overline{E_{i,j}}$. It is these points $B_{i,j}$ which will represent our states in the modified algorithm, and so, we denote the set of all base points as $\mathcal{S}_n$.

If $\lambda_{B_{i,j},E_k}$ is a path starting at the base point $B_{i,j}$ and terminating at the edge E_k, we can homotope $\lambda_{B_{i,j},E_k}$ fixing the initial point $B_{i,j}$ such that $\lambda_{B_{i,j},E_k}$ is a minimal path relative to E_k, i.e. $\lambda_{B_{i,j},E_k}$ does not intersect any of the edges in $\frac{1}{n}\mathcal{F}$ which intersect E_k. For each edge E_k, there will be a unique segment $E_{k,\ell}$ that the path $\lambda_{B_{i,j},E_k}$ can terminate at such that $\lambda_{B_{i,j},E_k}$ is a minimal path relative to E_k. Existence of such a region occurs because the edges in $\frac{1}{n}\mathcal{F}$ intersect E_k disjointly. As a result, the edges of $\frac{1}{n}\mathcal{F}$ which intersect E_k split P_n into finitely many disjoint regions and each of these regions contains a unique segment of E_k. Since these regions contain unique segments of E_k, we can deduce that this segment $E_{k,\ell}$ is unique. As a result, for each base point $B_{i,j}$, we can compute $\Lambda_n(B_{i,j})$ to be the set of minimal paths in P_n, which start at $B_{i,j}$.

As previously mentioned, if we have two paths λ_k and λ_{k+1} in P_n which are compatible relative to some edge E_i on the boundary of P_n (and E_i is not in $\frac{1}{n}\mathcal{F}$), then these paths will meet in the interior of some triangle $\triangle ABC$ in $\frac{1}{n}\mathcal{F}$. If Y encodes how $\lambda_k \circ \lambda_{k+1}$ intersects τ, then the cutting sequence of $(\lambda_k \circ \lambda_{k+1}, \frac{1}{n}\mathcal{F})$ is given by $(\widetilde{\lambda_k}, T_{\{n,n\}}) \cdot Y \cdot (\widetilde{\lambda_{k+1}}, T_{\{n,n\}})$. The triangle $\triangle ABC$ can intersect E_i in one of the following ways:

1. Two edges of $\triangle ABC$ intersect the interior of E_i. In this case, a segment $E_{i,j}$, runs between these edges of $\triangle ABC$ which intersect E_i.

2. One edge of $\triangle ABC$ intersects the interior of E_i and one of the endpoints of E_i is a vertex of τ. In this case, there is a segment $E_{i,j}$ of E_i between this vertex of $\triangle ABC$ and the edge of τ intersecting E_i.

3. One edge of $\triangle ABC$ intersects an odd edge E_i, another edge of $\triangle ABC$ intersects the paired odd edge $\overline{E_i}$ in P_n. Let D be the odd vertex of the

corresponding odd triangle. Then, there is a segment $E_{i,j}$ which runs from D to the intersection of τ with E_i.

If λ_k is some minimal path that terminates at $E_{i,j}$, and λ_{k+1} is the next path that we wish to construct, then we will show that the way the concatenated path $\lambda_k \circ \lambda_{k+1}$ intersects $\triangle ABC$ is completely determined by how λ_{k+1} intersects $\triangle ABC$ relative to $E_{i,j}$.

Case 1: Two edges of $\triangle ABC$ intersect the interior of E_i. In this case, a segment $E_{i,j}$, runs between these edges of $\triangle ABC$ which intersect E_i.

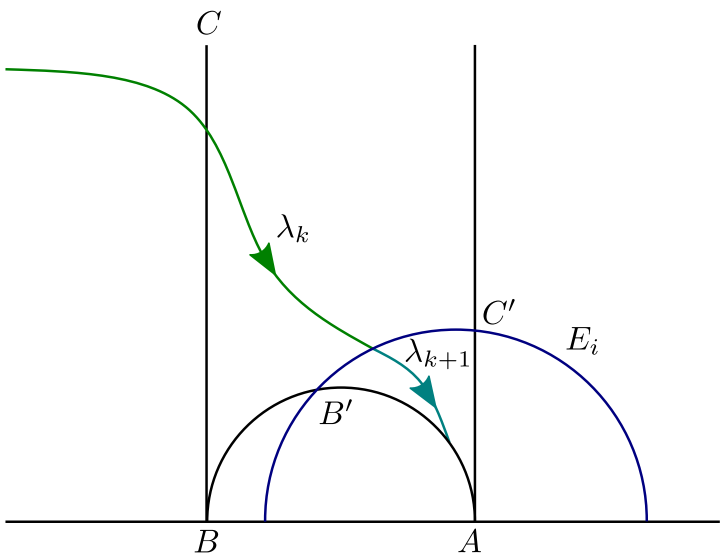

Figure 3.13: An image showing how the way the concatenated path $\lambda_k \circ \lambda_{k+1}$ intersects the triangle $\triangle ABC$ in Case I is dictated by the way that λ_{k+1} intersects the triangle $\triangle AB'C'$.

Without loss of generality, we assume that the edges of $\triangle ABC$ which intersect E_i, meet at the vertex A. We will denote the intersection of the line between A and B as B' and label the intersection of the line between A and C as C'. The points A, B' and C' form a triangle. If λ_k is a minimal path in P_n terminating at the region $E_{i,j}$, then λ_k can not have intersected the edge between A and B or the edge between A and C, as this would contradict λ_k being a minimal path. As a result, λ_k must have entered the triangle through the edge between B and C. Therefore, we can deduce that λ_k intersects the quadrilateral $BB'C'C$, before terminating at the edge between B' and C', and likewise we can deduce that λ_{k+1} intersects the triangle $\triangle AB'C'$. In this case, we see that if the path λ_{k+1} forms a right triangle relative to the triangle $\triangle AB'C'$, i.e. if λ_k passes through the edge between A and C', then the concatenated path $\lambda_k \circ \lambda_{k+1}$ forms a right triangle relative to $\triangle ABC$. Similarly, if the path λ_{k+1} forms a left triangle relative to the triangle $\triangle AB'C'$, then the concatenated path $\lambda_k \circ \lambda_{k+1}$ forms a left triangle relative to $\triangle ABC$. As a result, the triangle that λ_{k+1} forms with $\triangle AB'C'$ tells us how the concatenated path should behave. See Fig. 3.13.

Case 2: One edge of $\triangle ABC$ intersects the interior of E_i and one of the endpoints of E_i is a vertex of τ. In this case, there is a segment $E_{i,j}$ of E_i between this vertex of $\triangle ABC$ and the edge of τ intersecting E_i.

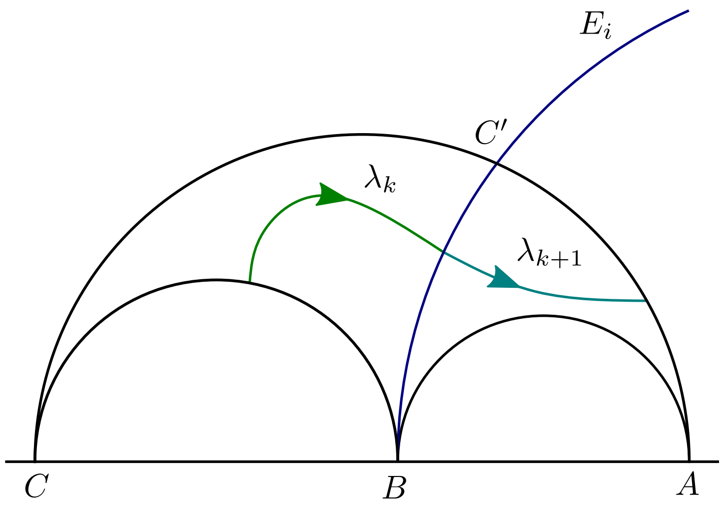

Figure 3.14: An image showing how the way the concatenated path $\lambda_k \circ \lambda_{k+1}$ intersects the triangle $\triangle ABC$ in Case II is dictated by the way that λ_{k+1} intersects the triangle $\triangle ABC'$.

In this case, we will assume that B is the vertex of the triangle which is the endpoint of E_i. Then the edge E intersects the edge between A and C at a point C'. If λ_k is a minimal path, then λ_k can not have intersected the edge between A and C, since this edge intersects E_i. As a result, λ_k must have entered the triangle $\triangle ABC$

through the edge between B and C or the edge between B and A. Without loss of generality, we assume that λ_k entered the triangle through the edge between B and C. Then λ_k intersects the triangle $\triangle BCC'$ and λ_{k+1} intersects the triangle $\triangle ABC'$. We then see that if the path λ_{k+1} forms a left or right triangle relative to the triangle $\triangle ABC'$, then the concatenated path $\lambda_k \circ \lambda_{k+1}$ forms a triangle of the same type relative to $\triangle ABC$. Again, we see that the triangle that λ_{k+1} forms with $\triangle AB'C'$ tells us how the concatenated path should behave. See Fig. 3.14.

Case 3: One edge of $\triangle ABC$ intersects an odd edge E_i, another edge of $\triangle ABC$ intersects the paired odd edge $\overline{E_i}$ in P_n. Let D be the odd vertex of the corresponding odd triangle. Then, there is a segment $E_{i,j}$ which runs from D to the intersection of τ with E_i.

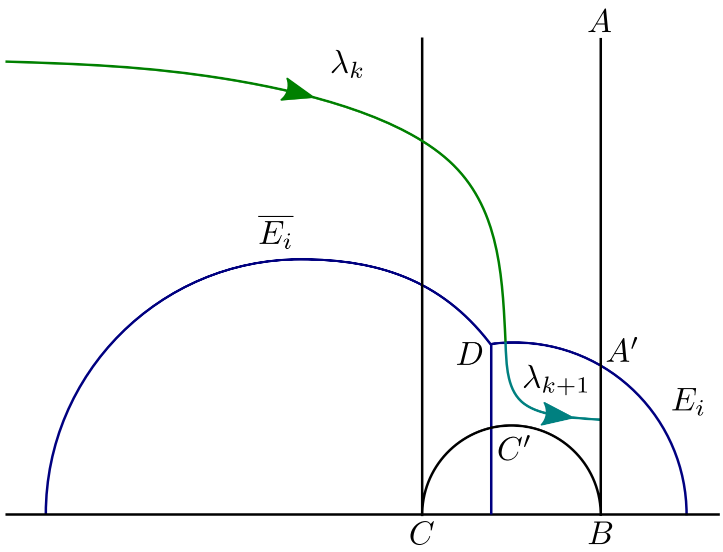

Figure 3.15: An image showing how the way the concatenated path $\lambda_k \circ \lambda_{k+1}$ intersects the triangle $\triangle ABC$ in Case III is dictated by the way that λ_{k+1} intersects the quadrilateral $A'BC'D$.

We will take $\widetilde{E_i}$ to be the other odd edge (which is not part of P_n) which is paired with E_i and $\overline{E_i}$ under the elliptic involution of order 3 around D. Then, without loss of generality, the edge between A and B will intersect E_i at a point A' and the edge between B and C will intersect $\widetilde{E_i}$ at a point C'. The path λ_k can not have intersected the edge between A and B, and can not have started at the edge $\overline{E}$ (since we saw that these paths were redundant in the previous section). As a result, λ_k must have passed through the edge between A and C. The next path λ_{k+1} lies in a quadrilateral $A'BC'D$. The path λ_{k+1} can not pass through the edge between D and C', as this would also lead to a redundant path. In particular, λ_{k+1} must either pass through the edge between B and A', in which case the concatenated path $\lambda_k \circ \lambda_{k+1}$ forms a right triangle with $\triangle ABC$, or λ_{k+1} passes through the edge between B and C', in which case the concatenated path $\lambda_k \circ \lambda_{k+1}$ forms a left triangle with $\triangle ABC$. If we treat the edges between A' and D and D and B' as if they are a single edge, then induces a natural notion λ_{k+1} forming a left or right "triangle" with $A'BC'D$, and this notion is consistent with the concatenated paths, i.e. if λ_{k+1} cuts $A'BC'D$ to form a left "triangle" then the concatenated path $\lambda_k \circ \lambda_{k+1}$ forms a left triangle with $\triangle ABC$. See Fig. 3.15.

As a result of looking at each of the above three cases, we see that if λ_k and λ_{k+1} are both minimal paths in P_n which are compatible relative to some edge E_i on the boundary of P_n (and E_i is not in $\frac{1}{n}\mathcal{F}$), then the way that λ_{k+1} intersects $\triangle ABC$ relative to $E_{i,j}$ completely determines how $\lambda_k \circ \lambda_{k+1}$ intersects $\triangle ABC$. As a result, in this case we will take $(\lambda_{k+1}, T_{\{n,n\}}) = Y \cdot (\widetilde{\lambda_{k+1}}, T_{\{n,n\}})$, where Y corresponds to the way that λ_{k+1} intersects the initial "triangle" that λ_{k+1} is contained in. Therefore, we can deduce that the cutting sequence $\left(\bigcup_{k \in \mathbb{N} \cup \{0\}} \lambda_k, \frac{1}{n}\mathcal{F}\right)$ is equivalent to:

$$
\begin{aligned}
\prod_{k=0}^{\infty} (\lambda_k, T_{\{n,n\}}) &= (\lambda_0, T_{\{n,n\}}) \cdot (\lambda_1, T_{\{n,n\}}) \cdot (\lambda_2, T_{\{n,n\}}) \dots \\
&= (\widetilde{\lambda_0}, T_{\{n,n\}}) \cdot Y_1 \cdot (\lambda_1, T_{\{n,n\}}) \cdot Y_2 \cdot (\lambda_2, T_{\{n,n\}}) \dots,
\end{aligned}
$$

where Y_k is ε if λ_k starts at an edge in $\frac{1}{n}\mathcal{F}$ and otherwise Y_k is determined by how λ_k intersects the connecting triangle.

To summarise, for every base point $B_{i,j}$ and every edge E_j, we can find a unique base point $B_{k,\ell}$ on the edge E_j such that the path $\lambda_{B_{i,j},B_{k,\ell}} \in \Lambda_n(B_{i,j})$. We can take $\lambda_{B_{i,j},B_{k,\ell}}$ to be a geodesic representative which starts at $B_{i,j}$ and terminates at

$B_{k,\ell}$, in which case $\lambda_{B_{i,j},B_{k,\ell}}$ will be reduced relative to both $T_{\{1,n\}}$ and $T_{\{n,n\}}$. For each path $\lambda_{B_{i,j},B_{k,\ell}}$ in $\Lambda_n(B_{i,j})$, there is a unique homotopic path in $\Lambda_n(E_i)$. This is given by λ_{E_i,E_j}. The edges E_i and E_j are either in $\mathcal{F}$ or an odd edge of P_n, and so the cutting sequence relative to $T_{\{1,n\}}$ is preserved under relative homotopy of these edges. As a result, the extended input alphabet $\mathcal{I}_n(B_{i,j})$ is equivalent to the initial input alphabet $\mathcal{I}_n(E_i)$ and so $\mathcal{I}_n(B_{i,j})$ is a base. By the above arguments, we can guarantee that the cutting sequences $(\lambda_{B_{i,j},B_{k,\ell}}, T_{\{n,n\}})$ are all well defined and concatenate nicely. Therefore, the output alphabet $\mathcal{O}_n(B_{i,j})$ is also well-defined.

The Cutting Sequence Multiplication Algorithm

Let $n \in \mathbb{N}$ be the integer by which we wish to multiply.

1. Construct the special polygon P_n for $\Gamma_0(n)$.

2. Find $T_{\{1,n\}}$ and $T_{\{n,n\}}$ by taking the local structure of $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ in P_n.

3. By looking at how $T_{\{n,n\}}$ intersects the boundary of P_n, construct the set of base points $\mathcal{S}_n$.

4. For each base point $B_{i,j} \in \mathcal{S}_n$ construct the set of all non-trivial minimal paths starting at this points $\Lambda_n(B_{i,j})$.

5. For each base point $B_{i,j}$ construct the input alphabet:

$$\mathcal{I}_n(B_{i,j}) := \left\{(\lambda_{B_{i,j},B_{k,\ell}}, T_{\{1,n\}}) : \lambda_{B_{i,j},B_{k,\ell}} \in \Lambda_n(B_{i,j})\right\}.$$

6. For each base point $B_{i,j}$ construct the output alphabet:

$$\mathcal{O}_n(B_{i,j}) := \left\{(\lambda_{B_{i,j},B_{k,\ell}}, T_{\{n,n\}}) : \lambda_{B_{i,j},B_{k,\ell}} \in \Lambda_n(B_{i,j})\right\}.$$

7. If $\lambda_{B_{i,j},B_{k,\ell}} \in \Lambda_n(B_{i,j})$, then the transition function *sigma* is given by:

$$\sigma(\lambda_{B_{i,j},B_{k,\ell}}) = \overline{B_{k,\ell}},$$

where $\overline{B_{k,\ell}}$ is the base point which is paired with $B_{k,\ell}$ under the side pairings.

8. If $\lambda_{B_{i,j},B_{k,\ell}} \in \Lambda_n(B_{i,j})$, then the output function τ is given by:

$$\tau(\lambda_{B_{i,j},B_{k,\ell}}) = (\lambda_{B_{i,j},B_{k,\ell}}, T_{\{n,n\}}).$$

9. End of algorithm.

3.2.4 Explicit Constructions of the Multiplication Algorithm

For relatively low integers, the multiplication algorithm can be nicely expressed as automata. See Fig 3.19 and 3.24 . However, for $n \geq 5$ this gets a bit too complicated and so expressing the automaton as a table of the initial state, the input, the output and the next state is much more appropriate.

An Explicit Algorithm for $p = 2$

For $p = 2$, we use the special polygon shown in Fig. 3.16. By overlapping the triangulations $T_{\{1,2\}}$ and $T_{\{2,2\}}$, we see that there are four base points $\mathcal{S}_2 := \{B_0, \overline{B_0}, B_1, \overline{B_1}\}$.

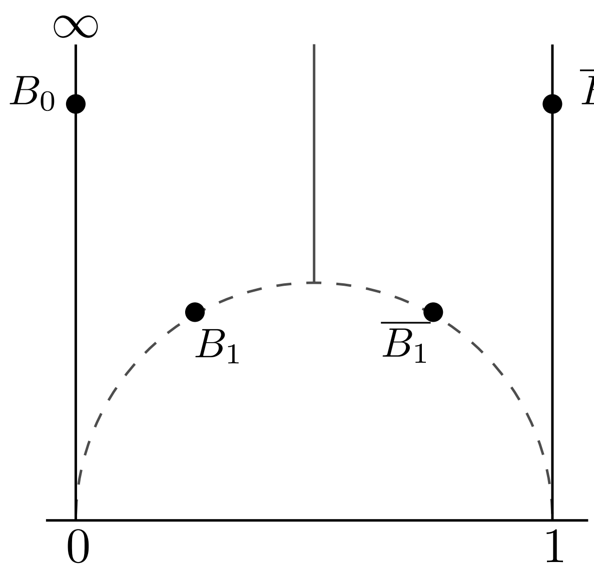

Figure 3.16: An image of the special polygon P_2 with the embedded structure of $T_{\{2,2\}}$. The base points are indicated on the diagram.

We first look at the paths starting at B_0. From Fig. 3.17, we can see that there are two minimal paths: One from B_0 to $\overline{B_0}$ with cutting sequences $(\lambda_{B_0,\overline{B_0}}, T_{\{1,2\}}) = L$ and $(\lambda_{B_0,\overline{B_0}}, T_{\{2,2\}}) = L^2$, and the other from B_0 to B_1 with cutting sequences $(\lambda_{B_0,B_1}, T_{\{1,2\}}) = R$ and $(\lambda_{B_0,B_1}, T_{\{2,2\}}) = \varepsilon$.

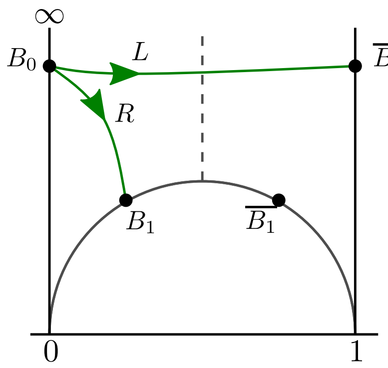

(a) An image of the minimal paths starting at the base point B_0 in P_2. These paths are taken relative to $T_{\{1,2\}}$ and labelled with their cutting sequence.

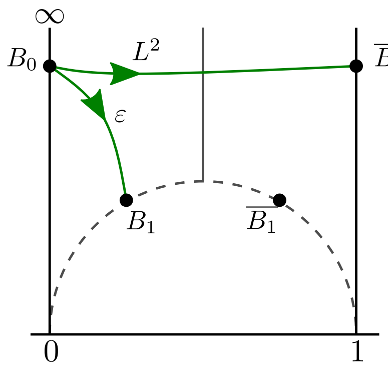

(b) An image of the minimal paths starting at the base point B_0 in P_2. These paths are taken relative to $T_{\{2,2\}}$ and labelled with their cutting sequence.

Figure 3.17

We can then look at the paths starting at B_1. From Fig. 3.18, we can see that there are two minimal paths: One from B_1 to B_0 with cutting sequences $(\lambda_{B_1,B_0}, T_{\{1,2\}}) = L$ and $(\lambda_{B_1,B_0}, T_{\{2,2\}}) = L$, and the other from B_1 to $\overline{B_0}$ with cutting sequences $(\lambda_{B_1,\overline{B_0}}, T_{\{1,2\}}) = R$ and $(\lambda_{B_1,\overline{B_0}}, T_{\{2,2\}}) = RL$.

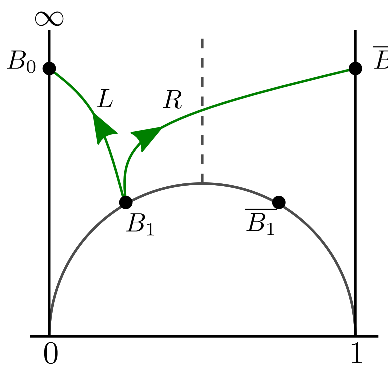

(a) An image of the minimal paths starting at the base point B_1 in P_2. These paths are taken relative to $T_{\{1,2\}}$ and labelled with their cutting sequence.

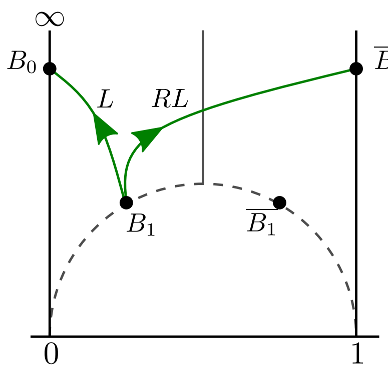

(b) An image of the minimal paths starting at the base point B_1 in P_2. These paths are taken relative to $T_{\{2,2\}}$ and labelled with their cutting sequence.

Figure 3.18

We note that P_2 is symmetric in the line $x = \frac{1}{2}$. Under this symmetry, $\overline{B_0}$ maps to B_0 and $\overline{B_1}$. This symmetry takes left triangles to right triangles and vice versa. Therefore, having constructed the minimal paths and cutting sequences for the paths starting at B_0 and B_1, we can deduce the minimal paths and cutting sequences for the paths starting at $\overline{B_0}$ and $\overline{B_1}$ by taking the mirror image, and the cutting sequences are given by replacing L with R and vice versa. This produces Table 3.1. We can produce an equivalent diagram by drawing all the states as nodes and all of the transition maps as labelled arrows. The labelling indicates which input is swapped with which output. See Fig. 3.19.

Initial State	Transition Map	Next State
S_0	$L : L^2$ $R : \varepsilon$	S_0 $\overline{S_1}$
S_1	$R : RL$ $L : L$	S_0 $\overline{S_0}$
$\overline{S_0}$	$R : R^2$ $L : \varepsilon$	$\overline{S_0}$ S_1
$\overline{S_1}$	$L : LR$ $RL : \varepsilon$	$\overline{S_0}$ S_0

Table 3.1: The transition table of the cutting sequences corresponding to multiplication by 2.

We finish the $p = 2$ case with a short example of how the algorithm would work for $\alpha = \frac{116}{165}$. This has continued fraction expansion $\overline{\alpha} = [0; 1, 2, 2, 1, 2, 1, 1, 2]$ corresponding to the cutting sequence $RL^2R^2LR^2LRL^2$.

$S_0 = B_0$	$W_0 = (R)L^2R^2LR^2LRL^2$	$U_0 = R$	$\tau_2(U_0) = \varepsilon$	$V = \varepsilon$
$S_1 = \overline{B_1}$	$W_1 = (L)LR^2LR^2LRL^2$	$U_1 = L$	$\tau_2(U_1) = LR$	$V = LR$
$S_2 = \overline{B_0}$	$W_2 = (L)R^2LR^2LRL^2$	$U_2 = L$	$\tau_2(U_2) = \varepsilon$	$V = LR$
$S_3 = B_1$	$W_3 = (R)RLR^2LRL^2$	$U_3 = R$	$\tau_2(U_3) = RL$	$V = LR^2L$
$S_4 = B_0$	$W_4 = (R)LR^2LRL^2$	$U_4 = R$	$\tau_2(U_4) = \varepsilon$	$V = LR^2L$
$S_5 = \overline{B_1}$	$W_5 = (L)R^2LRL^2$	$U_5 = L$	$\tau_2(U_5) = LR$	$V = LR^2L^2R$
$S_6 = \overline{B_0}$	$W_6 = (R)RLRL^2$	$U_6 = R$	$\tau_2(U_6) = R^2$	$V = LR^2L^2R^3$
$S_7 = \overline{B_0}$	$W_7 = (R)LRL^2$	$U_7 = R$	$\tau_2(U_7) = R^2$	$V = LR^2L^2R^5$
$S_8 = \overline{B_0}$	$W_8 = (L)RL^2$	$U_8 = L$	$\tau_2(U_8) = \varepsilon$	$V = LR^2L^2R^5$
$S_9 = \overline{B_0}$	$W_9 = (R)L^2$	$U_9 = R$	$\tau_2(U_9) = RL$	$V = LR^2L^2R^6L$
$S_{10} = \overline{B_0}$	$W_{10} = (L)L$	$U_{10} = L$	$\tau_2(U_{10}) = L^2$	$V = LR^2L^2R^6L^3$
$S_{11} = \overline{B_0}$	$W_{11} = (L)$	$U_{11} = L$	$\tau_2(U_{11}) = L^2$	$V = LR^2L^2R^6L^5$

The final output is $V = LR^2L^2R^6L^5$, which corresponds to the continued fraction expansion $[1; 2, 2, 6, 5]$. The corresponding real number is $\frac{232}{165}$, which is indeed $2 \times \frac{116}{165}$.

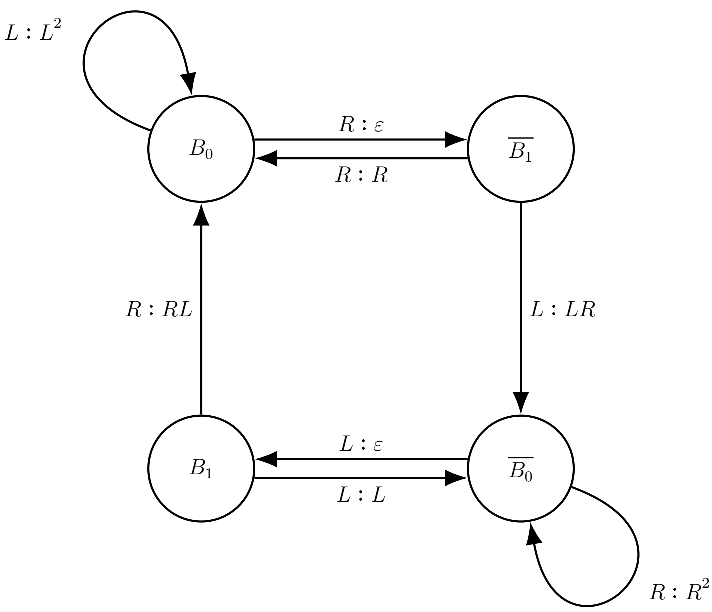

Figure 3.19: An automata which multiplies continued fractions by 2. The initial state is B_0.

An Explicit Algorithm for $n = 3$

For $n = 3$, we use the special polygon shown in Fig. 3.20. By overlapping the triangulations $T_{\{1,3\}}$ and $T_{\{3,3\}}$, we see that there are six base points:

$$\mathcal{S}_3 := \{B_0, B_{1,1}, B_{1,2}, \overline{B_0}, \overline{B_{1,1}}, \overline{B_{1,2}}\}.$$

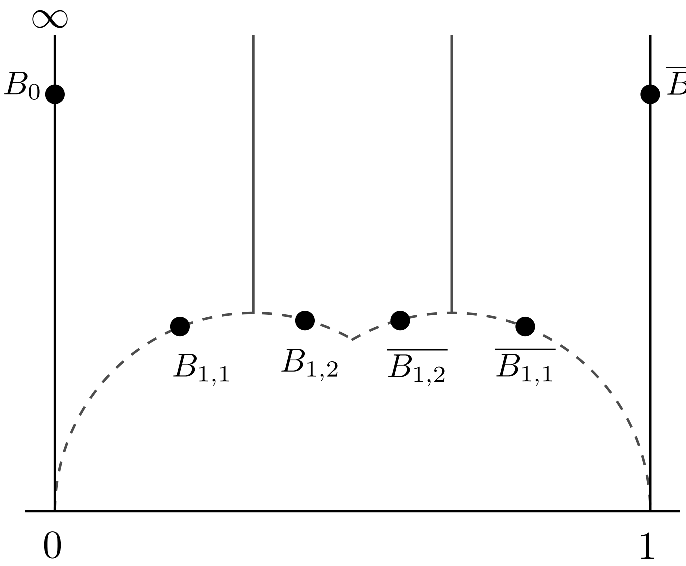

Figure 3.20: An image of the special polygon P_3 with the embedded structure of $T_{\{3,3\}}$. The base points are indicated on the diagram.

We first look at the paths starting at B_0. From Fig. 3.21, we can see that there are three minimal paths. One goes from B_0 to $\overline{B_0}$ with cutting sequences $(\lambda_{B_0,\overline{B_0}}, T_{\{1,3\}}) = L$ and $(\lambda_{B_0,\overline{B_0}}, T_{\{3,3\}}) = L^3$. Another goes from B_0 to $\overline{B_{1,2}}$ with cutting sequences $(\lambda_{B_0,\overline{B_{1,2}}}, T_{\{1,3\}}) = RR^{\frac{1}{2}}$ and $(\lambda_{B_0,\overline{B_{1,2}}}, T_{\{3,3\}}) = L$. The final path goes from B_0 to $B_{1,1}$ with cutting sequences $(\lambda_{B_0,B_{1,1}}, T_{\{1,3\}}) = RL^{\frac{1}{2}}$ and $(\lambda_{B_0,B_{1,1}}, T_{\{3,3\}}) = \varepsilon$.

We can then look at the paths starting at $B_{1,1}$. From Fig. 3.22, we can see that there are two minimal paths: One from $B_{1,1}$ to B_0 with cutting sequences $(\lambda_{B_{1,1},B_0}, T_{\{1,3\}}) = L^{\frac{1}{2}}L$ and $(\lambda_{B_{1,1},B_0}, T_{\{3,3\}}) = L$, and the other from $B_{1,1}$ to $\overline{B_0}$ with cutting sequences $(\lambda_{B_{1,1},\overline{B_0}}, T_{\{1,3\}}) = L^{\frac{1}{2}}R$ and $(\lambda_{B_{1,1},\overline{B_0}}, T_{\{3,3\}}) = RL^2$.

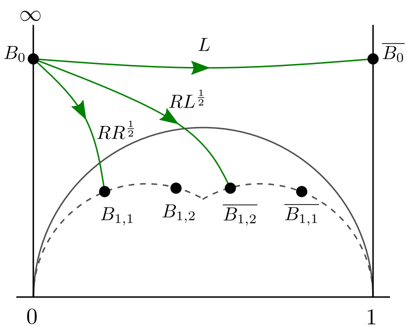

(a) An image of the minimal paths starting at the base point B_0 in P_3. These paths are taken relative to $T_{\{1,3\}}$.

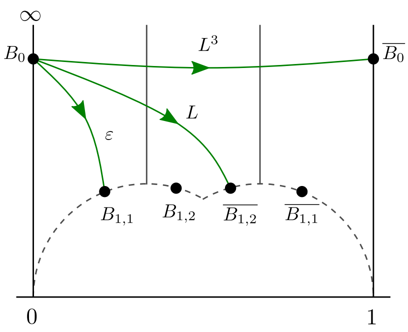

(b) An image of the minimal paths starting at the base point B_0 in P_3. These paths are taken relative to $T_{\{3,3\}}$.

Figure 3.21

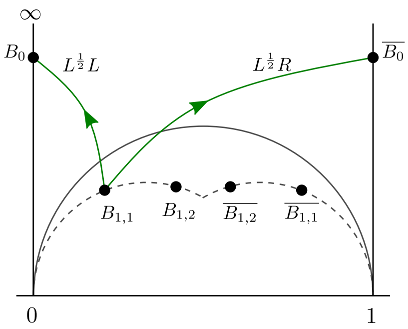

(a) An image of the minimal paths starting at the base point $B_{1,1}$ in P_3. These paths are taken relative to $T_{\{1,3\}}$.

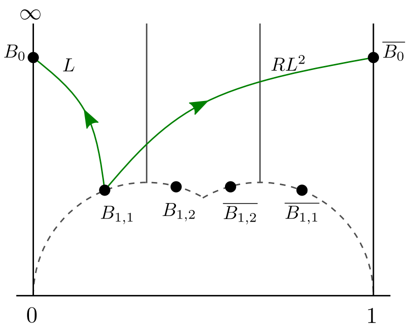

(b) An image of the minimal paths starting at the base point $B_{1,1}$ in P_3. These paths are taken relative to $T_{\{3,3\}}$.

Figure 3.22

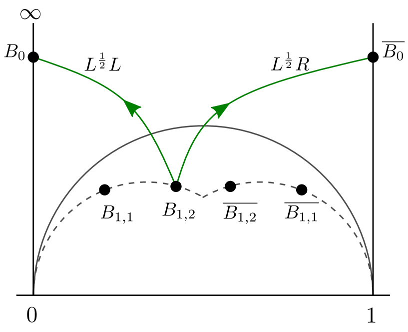

(a) An image of the minimal paths starting at the base point $B_{1,2}$ in P_3. These paths are taken relative to $T_{\{1,3\}}$.

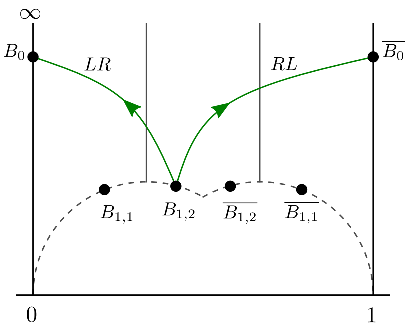

(b) An image of the minimal paths starting at the base point $B_{1,2}$ in P_3. These paths are taken relative to $T_{\{3,3\}}$.

Figure 3.23

Finally, we look at the paths starting at $B_{1,2}$. From Fig. 3.23, we can see that there are two minimal paths: One from $B_{1,2}$ to B_0 with cutting sequences $(\lambda_{B_{1,2},B_0}, T_{\{1,3\}}) = L^{\frac{1}{2}}L$ and $(\lambda_{B_{1,2},B_0}, T_{\{3,3\}}) = LR$, and the other from $B_{1,2}$ to $\overline{B_0}$ with cutting sequences $(\lambda_{B_{1,2},\overline{B_0}}, T_{\{1,3\}}) = L^{\frac{1}{2}}R$ and $(\lambda_{B_{1,2},\overline{B_0}}, T_{\{3,3\}}) = RL$.

We note that P_3 is symmetric in the line $x = \frac{1}{2}$. Under this symmetry, $\overline{B_0}$ maps to B_0, $\overline{B_{1,1}}$ maps to $B_{1,1}$, and $\overline{B_{1,2}}$ maps to $B_{1,2}$. This symmetry takes left triangles to right triangles and vice versa. Therefore, having constructed the minimal paths and cutting sequences for the paths starting at B_0 and B_1, we can deduce the minimal paths and cutting sequences for the paths starting at $\overline{B_0}$ and $\overline{B_1}$ by taking the mirror image, and the cutting sequences are given by replacing L with R and vice versa. This produces Table 3.2.

Initial State	Transition	Next State	Initial State	Transition	Next State
B_0	$L : L^3$ $RL^{\frac{1}{2}} : L$ $RR^{\frac{1}{2}} : \varepsilon$	B_0 $B_{1,2}$ $B_{1,1}$	$\overline{B_0}$	$R : R^3$ $LR^{\frac{1}{2}} : R$ $LL^{\frac{1}{2}} : \varepsilon$	$\overline{B_0}$ $\overline{B_{1,2}}$ $B_{1,1}$
$B_{1,1}$	$L^{\frac{1}{2}}L : L$ $L^{\frac{1}{2}}R : RL^2$	$\overline{B_0}$ B_0	$\overline{B_{1,1}}$	$R^{\frac{1}{2}}R : R$ $R^{\frac{1}{2}}L : LR^2$	B_0 $\overline{B_0}$
$B_{1,2}$	$L^{\frac{1}{2}}L : LR$ $L^{\frac{1}{2}}R : RL$	$\overline{B_0}$ B_0	$\overline{B_{1,2}}$	$R^{\frac{1}{2}}R : RL$ $R^{\frac{1}{2}}L : LR$	B_0 $\overline{B_0}$

Table 3.2: The transition table of the cutting sequences corresponding to multiplication by 3.

We then produce the corresponding automaton - see Fig. 3.24 - as we did for $p = 2$.

We will finish this the $p = 3$ case with a short example of how the algorithm would work for $\alpha = \frac{116}{165}$. This has continued fraction expansion $\overline{\alpha} = [0; 1, 2, 2, 1, 2, 1, 1, 2]$ corresponding to the cutting sequence $RL^2R^2LR^2LRL^2$.

$S_0 = B_0$	$W_0 = (RL^{\frac{1}{2}})L^{\frac{3}{2}}R^2LR^2LRL^2$	$U_0 = RL^{\frac{1}{2}}$	$\tau_2(U_0) = L$	$V = L$
$S_1 = B_{1,2}$	$W_1 = (L^{\frac{1}{2}}L)R^2LR^2LRL^2$	$U_1 = L^{\frac{1}{2}}L$	$\tau_2(U_1) = LR$	$V = L^2R$
$S_2 = \overline{B_0}$	$W_2 = (R)RLR^2LRL^2$	$U_2 = R$	$\tau_2(U_2) = R^3$	$V = L^2R^4$
$S_3 = \overline{B_0}$	$W_3 = (R)LR^2LRL^2$	$U_3 = R$	$\tau_2(U_3) = R^3$	$V = L^2R^7$
$S_4 = \overline{B_0}$	$W_4 = (LR^{\frac{1}{2}})R^{\frac{3}{2}}LRL^2$	$U_4 = LR^{\frac{1}{2}}$	$\tau_2(U_4) = R$	$V = L^2R^8$
$S_5 = \overline{B_{1,2}}$	$W_5 = (R^{\frac{1}{2}}R)LRL^2$	$U_5 = R^{\frac{1}{2}}R$	$\tau_2(U_5) = R$	$V = L^2R^9$
$S_6 = B_0$	$W_6 = (L)RL^2$	$U_6 = L$	$\tau_2(U_6) = L^3$	$V = L^2R^9L^4$
$S_7 = B_0$	$W_7 = (RL^{\frac{1}{2}})L^{\frac{3}{2}}$	$U_7 = RL^{\frac{1}{2}}$	$\tau_2(U_7) = L$	$V = L^2R^9L^5$
$S_8 = B_{1,2}$	$W_8 = (L^{\frac{1}{2}}L)$	$U_8 = L^{\frac{1}{2}}L$	$\tau_2(U_8) = L$	$V = L^2R^9L^6$

The final output is $V = L^2R^9L^6$, which corresponds to the continued fraction expansion $[2; 9, 6]$. The corresponding real number is $\frac{348}{165}$, which is indeed $3 \times \frac{116}{165}$.

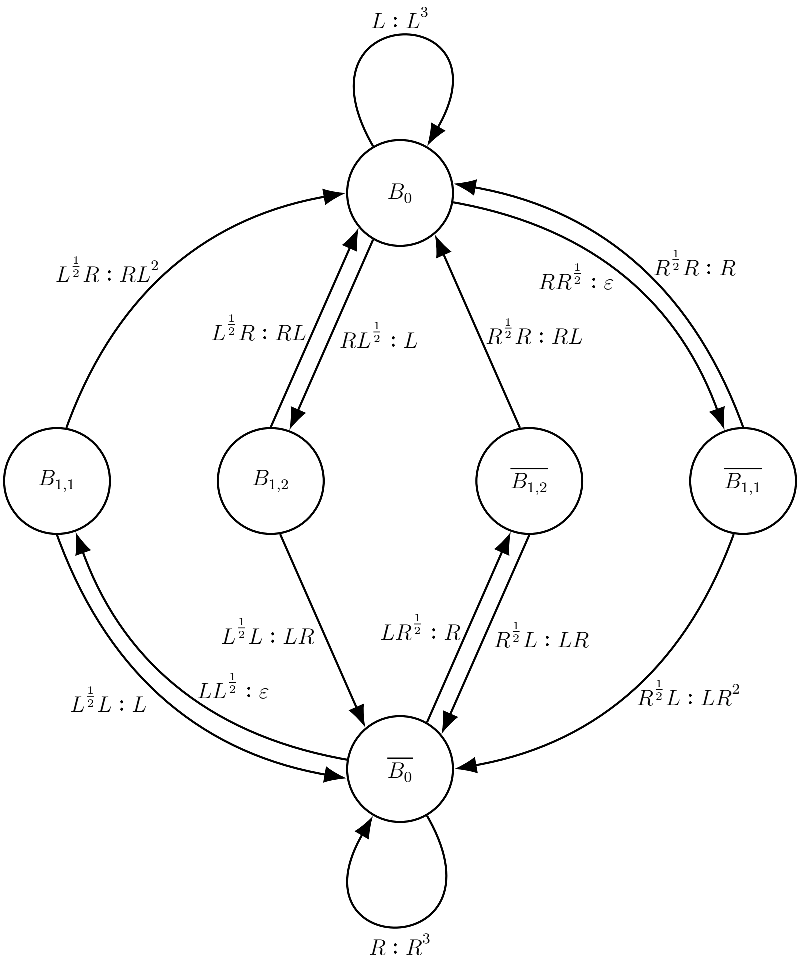

Figure 3.24: An automata which multiplies continued fractions by 3. The initial state is B_0.

Chapter 4

Cutting sequences on $\Phi\backslash\mathbb{H}$

4.1 Cutting Sequences on $\Phi\backslash\mathbb{H}$

In Section 3.1, we defined cutting sequences of geodesic rays with respect to ideal triangulations of $\mathbb{H}$. The concept of a geodesic ray intersecting a triangle to form a left or right triangle is still well-defined for any triangle on any surface, and so, for any triangulated surface we can define the cutting sequence of a geodesic ray relative to this triangulation. Similarly, we can define the generalised cutting sequence of a path relative to a triangulated surface, since the notions of a bigon, a left triangle and a right triangle all still hold.

However, if we take Φ to be a finite index subgroup of $PSL_2(\mathbb{Z})$, the quotient space $\Phi\backslash\mathbb{H}$ is not strictly a surface. Instead, we obtain a two-dimensional *orbifold*. Here, we define a *two-dimensional orbifold* to be a surface S (possibly with boundary), with a set of marked points M and a potentially empty set of orbifold points Q. In the case that Q is empty, the orbifold will simply be a surface. When we supply the orbifold with a metric, each element of M will correspond to *cusps* with angle 0 and each element of Q will correspond to a *cone point/orbifold point* with angle $\frac{2\pi}{k}$ for some $k \in \mathbb{N}_{\geq 2}$. As discussed in Section 3.2.2, these orbifolds will have empty boundary, at least one cusp (element of M) and a potentially empty set of orbifold points Q, which can be decomposed into two disjoint subsets: Q_2, the set of orbifold points with cone angle π and Q_3, the set of orbifold points with cone angle $\frac{2\pi}{3}$. Following the definitions in Section 3.2.2, the number of cusps is exactly given by $|M| = t$, the number of orbifold points with cone angles π is $|Q_2| = e_2$, and the number of orbifold points with cone angle $\frac{2\pi}{3}$ is $|Q_3| = e_3$.

Given some subgroup Φ of $PSL_2(\mathbb{Z})$, we can find the special polygon P_Φ which is the fundamental domain for Φ. Taking the quotient $\Phi \backslash \mathbb{H}$ is equivalent to taking the corresponding special polygon P_Φ and identifying sides via the side pairings. We denote the special polygon with edges identified as $P_\Phi / \sim$. When we take the corresponding special polygon P_Φ and identify sides via its side pairings, we see that the central points of the even edges directly correspond to elements of Q_2, the interior vertices of P_Φ connecting two odd edges correspond to elements of Q_3 and the vertices on the boundary of $\mathbb{H}$ correspond to elements of M (considered up to the side pairings).

4.1.1 Triangulations of $\Phi \backslash \mathbb{H}$

We define an *arc* γ on an orbifold $\Phi \backslash \mathbb{H}$ to be a geodesic path, which is disjoint from $M \cup Q$ except from its endpoints with the following properties:

- The endpoints of γ are contained in $M \cup Q$ and at least one endpoint is in M,
- The only self-intersections of γ occur at the endpoints of γ, if at all,
- If γ bounds a monogon (i.e. both endpoints of γ are at the same point in M), then this monogon either contains one element of M, one element of Q_3 or two elements of Q_2.

If γ has one endpoint in Q_3 (and the other in M), then we will say that γ is an *odd arc*. We will say that a pair of arcs γ, γ' are *admissible*, if $\gamma \cap \gamma' \subset M$ (i.e. γ and γ' only intersect at endpoints which are also marked points). If the number of marked points M and the number of orbifold points Q are finite, then there is a maximum number of admissible arcs that we can have on orbifold. We define a *quotient triangulation* T of an orbifold ${}_{\Phi}\backslash^{\mathbb{H}}$ to be a maximal collection of pairwise admissible arcs on ${}_{\Phi}\backslash^{\mathbb{H}}$. This maximal collection arcs separates our space into a collection of triangles. There are six possible of types triangle that can arise from a quotient triangulation, which we list in Table 4.1.

Type	Name	Diagram	Lift in $\mathbb{H}$
(I)	Standard triangle	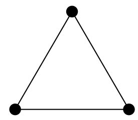	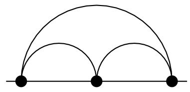
(II)	Self-folded triangle	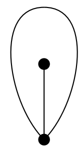	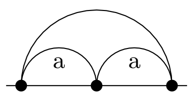
(IIIa)	Quotient-2 triangle (a)	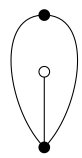	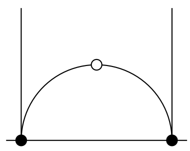
(IIIb)	Quotient-2 triangle (b)	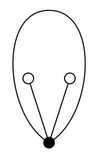	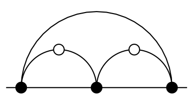
(IIIc)*	Quotient-2 triangle (c)	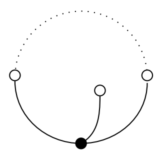	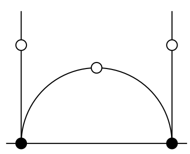
(IV)	Quotient-3 triangle	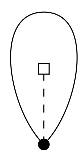	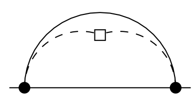

Table 4.1: A table of the six possible types of triangles that can appear in a quotient triangulation and their lifts in $\mathbb{H}$. Elements of P are indicated by $\bullet$, elements of Q_2 are indicated by $\circ$, and elements of Q_3 are indicated by $\square$. Dashed lines indicate odd arcs and their lifts.

Remark 4.1.1. * The quotient-2 triangle (IIIc) occurs as a triangulation for exactly one orbifold. This orbifold has three elements in Q_2 and a single cusp, and the triangle is formed by taking an arc between each point in Q_2 and the cusp. There is only one subgroup Φ of $PSL_2(\mathbb{Z})$, given by $\Gamma_3 = \langle \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}, \begin{pmatrix} 1 & -2 \\ 1 & -1 \end{pmatrix}, \begin{pmatrix} 1 & -1 \\ 2 & -1 \end{pmatrix} \rangle$, which induces a quotient space $_{\Phi}\backslash^{\mathbb{H}}$ that allows such a triangle. Furthermore, only triangles of this type appear on this orbifold. There are two special polygons of Γ_3 with Farey symbols $\{\, \infty \underset{\circ}{\smile} 0 \underset{\circ}{\smile} 1 \underset{\circ}{\smile} \infty \,\}$ and $\{\, \infty \underset{\circ}{\smile} -1 \underset{\circ}{\smile} 0 \underset{\circ}{\smile} \infty \,\}$.

When we take the Farey tessellation $\mathcal{F}$ relative to some special polygon P_Φ, since the special polygon is built using of Farey neighbours, it is not hard to see that $\mathcal{F}$ decomposes P_Φ into triangles of type $(\tilde{\text{I}}) - (\widetilde{\text{IV}})$ minus the lift of any odd arcs (where $(\tilde{\text{I}})$ is the lift of (I) in $\mathbb{H}$ etc.). As a result, the projection of $\mathcal{F}$ decomposes $_{\Phi}\backslash^{\mathbb{H}}$ into triangles of type (I)-(IIIc) or into monogons containing a single element of Q_3. For the monogons containing a single element of Q_3, we can construct a unique odd arc between this element of Q_3 and the element of M on the boundary of this monogon. In particular, the projection of $\mathcal{F}$ induces a unique quotient triangulation for all the quotient spaces $_{\Phi}\backslash^{\mathbb{H}}$, for Φ any finite index subgroup of $PSL_2(\mathbb{Z})$. However, in general it is not immediately clear when an arbitrary ideal triangulation T will induce a quotient triangulation of $_{\Phi}\backslash^{\mathbb{H}}$. In particular, it is not obvious whether $\frac{1}{d}\mathcal{F}$ induces a quotient triangulation of $_{\Gamma_0(n)}\backslash^{\mathbb{H}}$ for $d \mid n$. The following lemma gives a sufficient condition for an ideal triangulation T of $\mathbb{H}$ to induce a unique quotient triangulation on $_{\Phi}\backslash^{\mathbb{H}}$.

Lemma 4.1.2. *Let Φ be a non-trivial finite subgroup of $PSL_2(\mathbb{Z})$ (excluding Γ_3), let P_Φ be a special polygon which is a fundamental domain for Φ, and let T be an ideal triangulation of $\mathbb{H}$, which is invariant under Φ. Then, the projection of T decomposes $_{\Phi}\backslash^{\mathbb{H}}$ into triangles of type (I)-(IIIb) or into monogons containing a single element of Q_3. In particular, the projection of T induces a unique quotient triangulation T_Φ of $_{\Phi}\backslash^{\mathbb{H}}$.*

Proof. We will split this proof into three cases: Firstly, we show that if a triangle τ in T does not contain the lift of an orbifold point, then τ projects to a triangle of type (I) or (II) in $_{\Phi}\backslash^{\mathbb{H}}$. Secondly, we will consider the case that the triangle τ contains lifts of a points in Q_2. In this case τ projects to a triangle of type (IIIa)-(IIIc) in

$\Phi\backslash\mathbb{H}$. Finally, we will show that if τ contains the lift of a point in Q_3, then τ projects to a triangle of type (IV) in $\Phi\backslash\mathbb{H}$.

1. We first show that if a triangle τ in the ideal triangulation T of $\mathbb{H}$ does not contain a lift of an orbifold point in $\Phi\backslash\mathbb{H}$, then τ projects to a triangle of type (I) or (II) in $\Phi\backslash\mathbb{H}$.

Let τ be a triangle in the ideal triangulation T of $\mathbb{H}$ which does not contain a lift of an orbifold point in $\Phi\backslash\mathbb{H}$. Then, the projection of τ in $\Phi\backslash\mathbb{H}$ will not contain any elements of $Q = Q_2 \cup Q_3$. Since T is invariant under Φ, geodesics in $\mathbb{H}$ will project to geodesic arcs in $\Phi\backslash\mathbb{H}$ and these geodesic arcs will be pairwise disjoint except for at P. As a result, the projection of τ will be triangles of type (I) or (II).

2. We will now show that if τ contains the lift of an orbifold point in Q_2, then τ projects to a triangle of type (IIIa)-(IIIc) in $\Phi\backslash\mathbb{H}$.

Claim: If P_Φ contains an even edge e_2, then any triangulation T preserved by Φ must contain an edge that runs through m_2, where m_2 is the fixed point of φ_2, the side pairing induced by the even edge e_2.

Proof of claim: First, we assume the opposite, that m_2 is not intersected by any edge of T. Then, since m_2 lies in the interior of $\mathbb{H}$, the point m_2 must lie in the interior of τ, some triangle in T. Two vertices of τ will lie on one side e_+ of the even edge e_2 and one vertex of τ will lie on the other side e_-. Since φ_2 is an elliptic involution of order 2 with fixed point m_2, the image of $\varphi_2(\tau)$ will contain m_2 and have two vertices in e_- and one vertex in e_+. Since φ_2 is an element of Φ, it follows that $\varphi_2(\tau)$ must be a triangle in T (since T is invariant under Φ). Both triangles τ and $\varphi_2(\tau)$ contain the point m_2, however these triangles τ and $\varphi_2(\tau)$ are not equivalent, since the number of endpoints in e_+ and e_- are different for τ and $\varphi_2(\tau)$. See Fig. 4.1. This implies τ and $\varphi_2(\tau)$ have non-trivial intersection and do not intersect along a common edge (since then m_2 would lie on this edge). Therefore, T can not be an ideal triangulation and this is a contradiction to our initial assumptions. *QED.*

It follows from the above claim, that if a triangle τ in T contains the point m_2, then

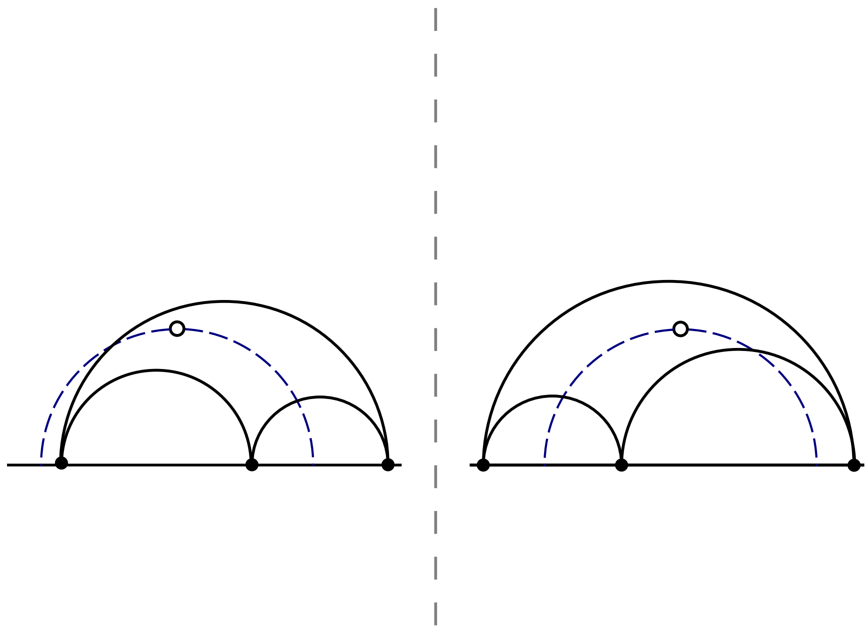

Figure 4.1: An example of a triangle which contains the point m_2 (left) in its interior, and its image under the action of φ_2 (right), where φ_2 is an elliptic involution of order 2 with fixed point m_2.

this point lies on one of the edges of τ. Such a triangle can either have one, two or three edges, which each contain the lift of a point in Q_2. These triangles will project to a quotient triangle in $_{\Phi}\backslash^{\mathbb{H}}$ of type (IIIa), (IIIb) or (IIIc) (which occurs only for $\Phi = \Gamma_3$), respectively.

3. Finally, we show that if τ contains the lift of an orbifold point in Q_3, then τ projects to a monogon containing a single element of Q_3. As seen above, we can then construct a unique odd arc between this element of Q_3 and the element of M which lies on the boundary of this monogon.

Claim: No edge in T projects to an odd arc in $_{\Phi}\backslash^{\mathbb{H}}$.

Proof of Claim. Assume that E is an edge of an ideal triangulation T in $\mathbb{H}$, which projects to an odd arc in $_{\Phi}\backslash^{\mathbb{H}}$. Then E must intersect the interior vertex c_3 of an odd triangle τ_3 in P_Φ. We define φ_3 to be the side pairing induced by this odd triangle. Then φ_3 is an elliptic involution of order 3 with fixed point c_3. In particular, the images of E under Id, φ_3 and φ_3^{-1} are three geodesic which all intersect at c_3. Since

T is invariant under Φ, all of the images of E under Φ (and therefore under Id, φ_3 and φ_3^{-1}) must be edges in T. See Fig. 4.2a. However, E, $\varphi_3(E)$ and $\varphi_3^{-1}(E)$ intersecting inside $\mathbb{H}$ and therefore, T can not be an ideal triangulation. This is a contradiction to our initial assumptions. *QED.*

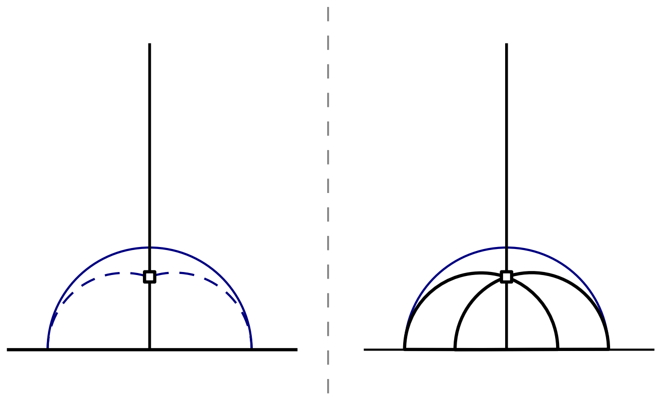

(a) A geodesic line (bold) passing through the point c_3 (left), and its images under Id, φ_3 and φ_3^{-1} (right).

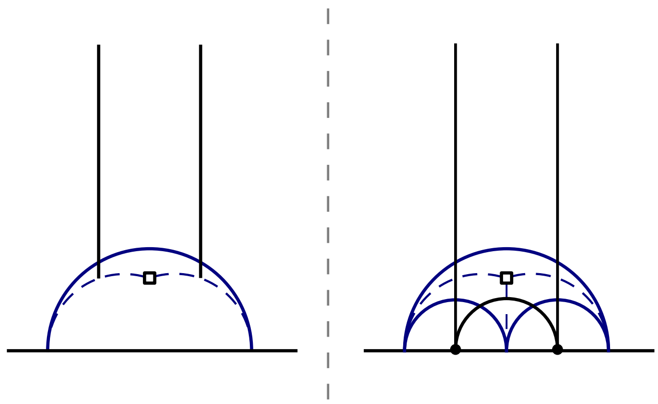

(b) A pair of geodesic rays (left), which form a triangle under the action of Id, φ_3 and φ_3^{-1} (right).

Figure 4.2: Examples of edges and their images under the actions of Id, φ_3 and φ_3^{-1}, where φ_3 is an elliptic involution of order 3 with fixed point c_3. The odd triangle τ_3 is drawn in for structure.

Following this claim, the point c_3 must lie in the interior of some triangle τ in T. The elliptic involution about c_3 will split $\mathbb{H}$ into three different regions, each containing a vertex of τ. Therefore, the projection of τ on to $\Phi\backslash\mathbb{H}$ will be a monogon containing a single orbifold point with cone angle $\frac{2\pi}{3}$. See Fig. 4.2(b). □

4.1.2 Paths and Cutting Sequence on $\Phi\backslash\mathbb{H}$

For the most part, the concepts we introduced in Section 3.1.1 all still apply. However, we will recap the main concepts for convenience and clarity.

Let $\Phi\backslash\mathbb{H}$ be an orbifold (as described in the previous section), and let T be a quotient triangulation of $\Phi\backslash\mathbb{H}$. Infinite paths will start at some arc γ of the triangulation T (which is not an odd arc) and will limit to a marked point on $\Phi\backslash\mathbb{H}$ (here, we treat marked points a like points at ∞). We will say that two infinite paths λ and λ' are *homotopic relative to* γ if they both start at the same arc γ of T, limit to the same marked point P and there is a homotopy between these paths which preserves the starting arc γ and the final vertex P. In this case, we write $\lambda \sim_\gamma \lambda'$. Alternatively, if λ and λ' two finite length paths, then we will say that these paths are *homotopic relative to* γ_1 and γ_2 if they both start at the same arc γ_1, terminate at the same arc γ_2, and there is a homotopy between them which preserves both γ_1 and γ_2. In this case, we write $\lambda \sim_{\gamma_1,\gamma_2} \lambda'$.

Cutting Sequences relative to Odd Arcs

When defining the cutting sequence, we note that each path λ can cut each triangle τ in T to either form a left triangle, a right triangle or a bigon (the path leaves the triangle from the same edge it entered), or it can terminate at a vertex of a triangle. In particular, the cutting sequence over the alphabet $\{L, R, X\}$ can generally be formed as described in Section 3.1.1. However, for triangles of type (IV), there are a few alterations we need to make. Firstly, if λ cuts an odd arc to form a left/right triangle then we will append either $L^{\frac{1}{2}}$ or $R^{\frac{1}{2}}$, respectively. Whenever a path λ intersects an odd arc, we will require that λ does not “loop around” the cone point of order 3. See Fig 4.3 (a) and (b). As a result, the path λ must either form a bigon

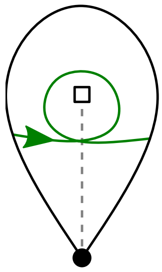

(a) An example of a path "looping around" the cone point of order 3.

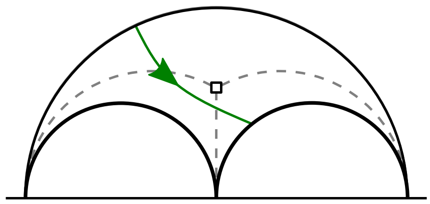

(b) The lift of a path which "loops around" a cone point of order 3 in $\mathbb{H}$.

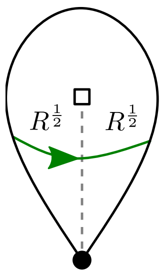

(c) An example of a path intersecting an odd arc to "complete" the triangle. It has cutting sequence $R^{\frac{1}{2}}R^{\frac{1}{2}} = R$.

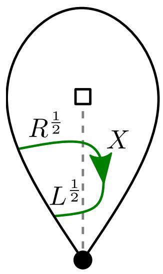

(d) An example of a path intersecting an odd arc, which then forms a bigon with this odd arc. It has cutting sequence $R^{\frac{1}{2}}XL^{\frac{1}{2}}$.

Figure 4.3: Examples of the different ways a path can intersect an odd arc.

with the odd arc or λ must "complete" the half triangle. If λ forms a bigon with the odd arc, then this induces the reduction relations $L^{\frac{1}{2}}XR^{\frac{1}{2}} \sim X$ and $R^{\frac{1}{2}}XL^{\frac{1}{2}} \sim X$, and if λ "completes" the triangle, then we have the reduction relations $L^{\frac{1}{2}}L^{\frac{1}{2}} \sim L$ or $R^{\frac{1}{2}}R^{\frac{1}{2}} \sim R$. See Fig. 4.3 (c) and (d).

Direction of Departure and Direction of Approach

Each arc γ in the triangulation T has two sides, which we can arbitrarily label "+" and "−". Every path λ (finite or infinite) which starts at the arc γ must leave either via the positive side or negative side. If λ leaves via the positive side, we say that λ

has *positive direction of departure*, and if λ leaves via the negative side, we say that λ has *negative direction of departure*. We will express the direction of departure as a pair $(\gamma, \pm)$, where γ is the initial edge and $\pm$ tells us how the path leaves this edge, i.e. from the positive side or from the negative side. If λ is a finite path, then it will approach its terminal edge γ' from either the positive side, in which case we will say that λ has *positive direction of approach*, or λ will approach γ' from the negative side, in which case we say that λ has *negative direction of approach*. Again, we will express the direction of approach as the pair $(\gamma', \pm)$, where γ' is the terminal edge and $\pm$ encodes whether λ approaches from the negative side or positive side. If λ_1 is a finite length path terminating at an edge γ in T and λ_2 is any path starting at γ, we can concatenate these paths to form the path $\lambda = \lambda_1 \circ \lambda_2$. We will say that λ_1 and λ_2 are *compatible* if λ_1 approaches γ from one direction and λ_2 leaves γ from the opposite directions. If two paths λ_1 and λ_2 are compatible and reduced, then their concatenated product $\lambda = \lambda_1 \circ \lambda_2$ will be reduced, and the cutting sequence (λ, T) is equal to the product of the cutting sequences $(\lambda_1, T) \cdot (\lambda_2, T)$.

Reduction Relations and Other Path Equivalences

Let Q be an orbifold point of order 2 and let E be an edge terminating at Q. If a path λ intersects E and then loops around Q to intersect E again, then we will consider the loop that this path forms to be a bigon - since it intersects the edge E twice in a row and lifts to a bigon in $\mathbb{H}$. See Fig. 4.4. In particular, we will consider the path λ to be equivalent to the path λ' which has this loop removed.

As in Section 3.1.1, we can remove bigons from our paths and this induces reduction relations on the cutting sequence. All the previous reduction relations still hold. We will say that a path λ is *reduced* if it does not form bigons with T - and if it is not homotopic to a sub-path of its starting arc, i.e. it does not have the cutting sequence LX or RX considered up to reduction relations. The corresponding cutting sequence of a reduced path will also be *reduced*, in the sense that it is only made up of the letters L and R. As in Section 3.1.1, if two reduced paths are homotopic, then they

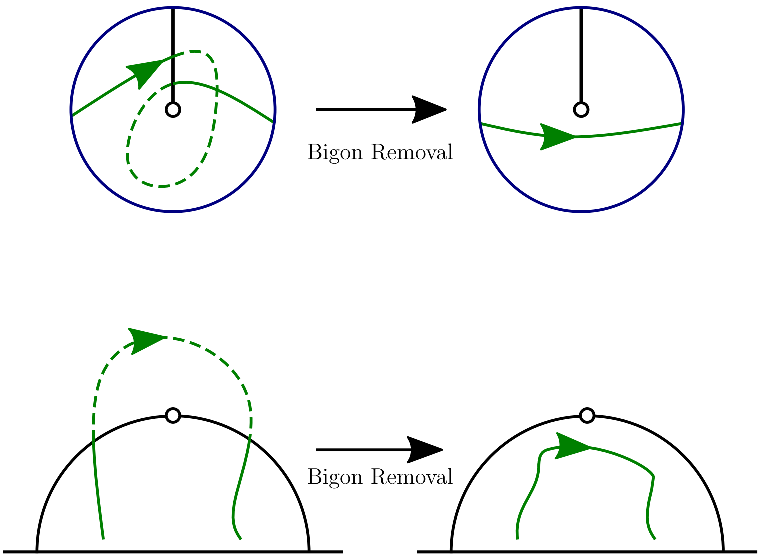

Figure 4.4: An example of how to remove a bigon which loops around an orbifold point of order 2, along with the corresponding lifts in $\mathbb{H}$. The arcs which form the bigon are indicated by a dashed line.

will have the same cutting sequence.

Lemma 4.1.3. *Let T be a triangulation of some orbifold $_{\Phi}\backslash^{\mathbb{H}}$ and let λ and λ' be two homotopic reduced paths, starting at the same edge γ in T. Then $(\lambda_1, T) = (\lambda_2, T)$.*

However, given two paths λ and λ' in $_{\Phi}\backslash^{\mathbb{H}}$, starting at the same edge S in T with the same cutting sequence, it is not necessarily true that λ_1 and λ_2 are homotopic. Given an initial starting edge γ, a direction of departure, and a reduced word $W \in \{L, R\}^{\mathbb{N}}$, we can produce a path λ_W such that $(\lambda_W, T) = W$. Inside triangles of type $(I), (II)$ and (IV), this path is uniquely defined up to homotopy (relative to T). However, if the path λ_W intersects an even edge, then it can do this in two different ways, and these two different ways induce the same cutting sequence. This is because a sub-path which terminates an even edge will add the same letter to the cutting sequence regardless of the direction of approach. See Fig. 4.5.

As a result, given a reduced word W and a triangulation T, we can construct many different paths $\lambda_{W,i}$, each with the cutting sequence $(\lambda_{W,i}, T) = W$. These paths will

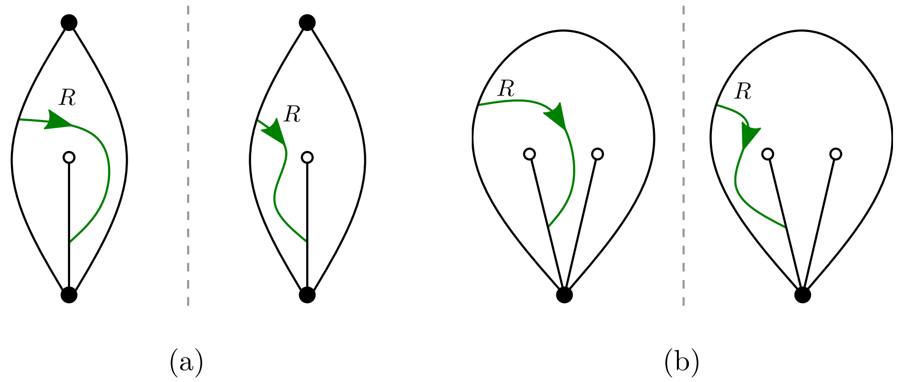

Figure 4.5: An example of how paths can terminate at an even edge with different directions of approach but still have the same cutting sequence. The left figure (a) shows this for paths in the triangle (IIIa) and the right figure (b) shows this for paths in the triangle (IIIb)

intersect the same edges in the same order, but they can intersect each even edge in one of two different ways corresponding to the two different directions of approach.

The fact that we can have multiple non-homotopic paths $\lambda_{W,i}$ with the same starting edge direction of departure and cutting sequence is not ideal when considering how triangulation replacement affects cutting sequences. In particular, if λ and λ' have the same starting edge direction of departure and cutting sequences (λ, T) and (λ', T), then it is not obvious that the cutting sequences (λ, T') and (λ', T') will be equivalent, where T' is some other triangulation. Fortunately, the next lemma proves exactly this.

Lemma 4.1.4. *Let $_{\Phi}\backslash^{\mathbb{H}}$ be some orbifold, with triangulation T. Assume that λ and λ' are two reduced paths which start at the same edge S with the same direction of departure and that these paths have the same (reduced) cutting sequence. If T' is some other triangulation of $_{\Phi}\backslash^{\mathbb{H}}$, then the cutting sequences (λ, T') and (λ', T') are equivalent (up to reduction relations).*

Remark 4.1.5. The above lemma is true for both finite paths and infinite paths. However, if λ and λ' are finite paths which terminate at an edge F, then we will assume that the edge F is in both T and T'. The main reason why do this is because otherwise the cutting sequences may not be well-defined.

Proof. Since λ and $\lambda^{'}$ are reduced paths, and have the same cutting sequence, we know that λ and $\lambda^{'}$ intersect the same edges of T in the same order, however they may intersect even edges in different ways.

We describe the proof assuming that λ and $\lambda^{'}$ only differ in how they intersect a single even edge and they only differ in how they intersect this edge once. However, since the edges of triangulations of $_\Phi\backslash^{\mathbb{H}}$ do not intersect each other (except at M the set of marked points) and our paths λ and $\lambda^{'}$ are disjoint from M except at the endpoints, the paths λ and $\lambda^{'}$ only ever intersect one edge of a triangulation at a time. In particular, this result follows even in the case that λ and $\lambda^{'}$ intersect arbitrarily many even edges in different ways.

Let Q be an orbifold point of order 2 and let E be the edge in T between Q and some marked point M of $_\Phi\backslash^{\mathbb{H}}$. We will assume that λ and $\lambda^{'}$ differ in how they intersect E exactly once. We can construct an open neighbourhood N around Q and homotope λ and $\lambda^{'}$ such that the intersections that λ and $\lambda^{'}$ have with E lie within this neighbourhood.

Remark 4.1.6. Note that the marked points and orbifold points on $_\Phi\backslash^{\mathbb{H}}$ appear discretely - i.e. there are not points which are arbitrarily close to each other. Furthermore, there is exactly one edge E in T which has Q as an endpoint. Similarly, there is exactly one edge $E^{'}$ in $T^{'}$ with Q as an endpoint. Since no other edges in T or $T^{'}$ can get arbitrarily close to Q, we can freely chose the neighbourhood N such that it contains no edge in T or $T^{'}$ except for E and $E^{'}$. Similarly, we can also assume that N contains no boundary points except for Q.

Since λ and $\lambda^{'}$ only differ in how they intersect E - and we have homotoped this intersection to appear inside N - outside this neighbourhood N we can homotope these paths to be concurrent. Let μ be the sub-path of λ inside this neighbourhood for the section of λ where λ and $\lambda^{'}$ intersect E in different ways and let $\mu^{'}$ be the corresponding sub-path of $\lambda^{'}$. We can then take λ_1 to be the sub-path of λ (and $\lambda^{'}$) going from the start of λ to the start of μ (or equivalently $\mu^{'}$). Similarly, we can take λ_2 to be the sub-path of λ (and λ') which starts at the end of μ. Then we have $\lambda = \lambda_1 \circ \mu \circ \lambda_2$ and $\lambda' = \lambda_1 \circ \mu' \circ \lambda_2$.

Let s be the starting point of μ and μ', let t be the end point of μ and μ' and let e be the be the point on E where E intersects ∂N. By going clockwise round ∂N, we induce a cyclic ordering points s, t and e. The two possible orderings are $\{s,t,e\}$ and $\{s,e,t\}$. We will assume that the ordering is $\{s,t,e\}$, however the argument we make holds for the ordering $\{s,e,t\}$ by symmetry. Since λ and λ' are reduced, μ and μ' intersect E exactly once. Up to homotopy, there are two possible paths in N with the ordering $\{s,t,e\}$. See Fig 4.6. We will call the section of ∂N lying in the region going clockwise from s to t the *top half of* ∂N and will label the region going anticlockwise from s to t the *bottom half of* ∂N. In our ordering, e lies in the bottom half of ∂N. Without loss of generality, we will take μ to be the path goes in an anti-clockwise direction from s to t without forming a loop around Q and intersecting E exactly once. Whereas, we will take μ' to be the path that starts at the s, loops around P once in a clockwise direction and then terminates at t. See Fig.4.6.

Let T' be a new triangulation of $\mathcal{O}$. Then, there is exactly one edge E' in T' going from some marked point to the orbifold point Q.

If E' comes from the bottom half of ∂N, then E' will intersect μ and μ' once each in the same way that the edge E intersected μ and μ'. Since $\lambda = \lambda_1 \circ \mu \circ \lambda_2$ and $\lambda' = \lambda_1 \circ \mu' \circ \lambda_2$, outside of the neighbourhood N, the paths λ and λ' intersect the same edges in the same order and in the same way. Inside this N, μ and μ' intersect both intersect E' exactly once. As a result, λ and λ' intersect the same edges in the same order and therefore have the same cutting sequence.

Alternatively, if E' comes from the top half of ∂N, then μ intersects this edge twice in a row but μ' does not intersect E' at all. See Fig. 4.7. The path μ forms a bigon with E', since it μ' intersects the even edge E' twice in a row. As such, the path μ' is equivalent to the path with the bigon removed. However, this path is simply μ.

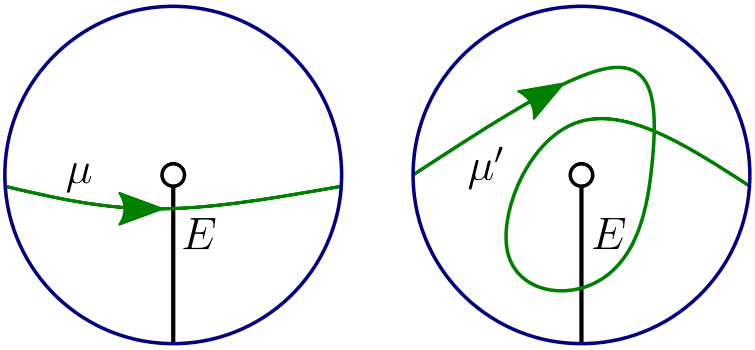

Figure 4.6: An image showing the two different possible ways (up to symmetry) that a reduced path can intersect an even edge.

As a result, the paths λ and λ' are equivalent up to bigon removal. Therefore, they have equivalent reduced cutting sequence.

Finally, if E' intersects N at one of the end points of μ of μ' then we can slightly shift either of these end points, s or t, via homotopy to end up in one of the above cases. In homotoping these endpoints, we will also need to homotope that paths λ_1 and λ_2 such that we still have $\lambda = \lambda_1 \circ \mu \circ \lambda_2$ and $\lambda' = \lambda_1 \circ \mu' \circ \lambda_2$. $\square$

Remark 4.1.7. Note that we will still think of even edges as having two directions of approach and two directions of departure, since this will allow us to talk about compatible arcs in a natural way. For simplicity, we will always assume that paths which start at even edges have positive direction of departure and paths which terminate at even edges have negative direction of approach. In particular, these paths are compatible.

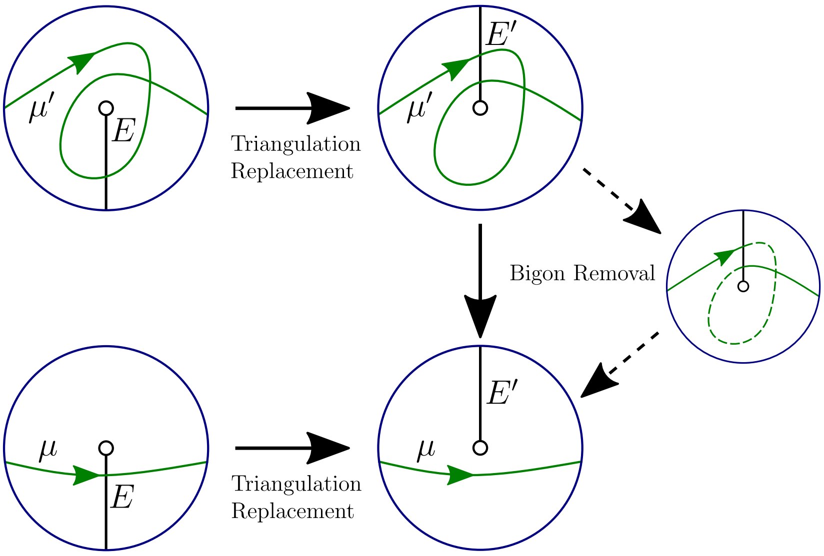

Figure 4.7: An example of how two paths, which only differ in how they intersect an even edge, are still equivalent (up to reduction) regardless of triangulation replacement. In this figure, the even edge E - lying in the bottom half of ∂N) - is replaced by another even edge $E^{'}$ - which lies in the top half of ∂N.

4.1.3 Triangulation Replacement on $\Gamma_0(n)\backslash\mathbb{H}$

We will end this section by discussing how we can use triangulation replacement on the orbifold $\Gamma_0(n)\backslash\mathbb{H}$ to represent integer multiplication of a continued fraction. This statement is essentially equivalent to the statement that triangulation replacement on a special polygon P_n represents integer multiplication of a continued fraction. However, we will still explicitly state this, since it will be useful in the next section for showing that if a continued fraction is recurrent, then any rational multiple of this continued fraction will also be recurrent.

When discussing the effect that triangulation replacement has on cutting sequences, we will want to make sure that the induced action is well-defined. If T and $T^{'}$ are our chosen triangulations, we will require that each path λ starts at an edge E and this edge is in both T and $T^{'}$. If the initial path is infinite, then we can freely homotope the path λ relative to E. In particular, we can homotope the path λ to make sure the cutting sequences (λ, T) and (λ, T') are both reduced (though not necessarily at the same time).

To discuss how triangulation replacement on $\Gamma_0(n)\backslash\mathbb{H}$ can represent integer multiplication by d for $d \mid n$, we start by noting that the scaled Farey complexes $\frac{1}{d}\mathcal{F}$ are preserved under the action of $\Gamma_0(n)$. Therefore, by Lemma 4.1.2 these triangulations induce quotient triangulations of $\Gamma_0(n)\backslash\mathbb{H}$. We denote these triangulations as $\widehat{\frac{1}{d}\mathcal{F}}$, since these triangulations are the projections of $\frac{1}{d}\mathcal{F}$ on to $\Gamma_0(n)\backslash\mathbb{H}$ with the addition of any required odd arcs. If $\widehat{I}$ is the projection of the y-axis I in $\Gamma_0(n)\backslash\mathbb{H}$, then since every triangulation $\frac{1}{d}\mathcal{F}$ contains the y-axis, each of their projections $\widehat{\frac{1}{d}\mathcal{F}}$ contains the arc $\widehat{I}$. If ζ_α is a geodesic ray in $\mathbb{H}$ starting at the y-axis I with endpoint $\alpha > 0$, then we can take $\widehat{\zeta_\alpha}$ to be its projection in $\Gamma_0(n)\backslash\mathbb{H}$. This path $\widehat{\zeta_\alpha}$ has a well-defined starting edge $\widehat{I}$, which is in all of the triangulations $\widehat{\frac{1}{d}\mathcal{F}}$, and so the cutting sequence $(\widehat{\zeta_\alpha}, \widehat{\frac{1}{d}\mathcal{F}})$ is well-defined, for all $d \mid n$. Using this information, we get the following theorem:

Theorem 4.1.8. *For every geodesic ray ζ_α in $\mathbb{H}$ starting at the y-axis I with endpoint $\alpha > 0$, there is a canonical projection $\widehat{\zeta_\alpha}$ onto $\Gamma_0(n)\backslash\mathbb{H}$ such that $(\zeta_\alpha, \frac{1}{d}\mathcal{F}) = (\widehat{\zeta_\alpha}, \widehat{\frac{1}{d}\mathcal{F}})$, for all $d \mid n$.*

Proof. Since the y-axis I is an edge in $\frac{1}{d}\mathcal{F}$ for all $d \in \mathbb{N}$, the projection $\widehat{\zeta_\alpha}$ of ζ_α in $\Gamma_0(n)\backslash\mathbb{H}$ is unique and has a well-defined starting edge $\widehat{I}$ and direction of departure, for all $\widehat{\frac{1}{d}\mathcal{F}}$. Taking the projection of ζ_α and $\frac{1}{d}\mathcal{F}$ does not change how these objects intersect. In particular, $(\zeta_\alpha, \frac{1}{d}\mathcal{F}) = (\widehat{\zeta_\alpha}, \widehat{\frac{1}{d}\mathcal{F}})$, for all $d \mid n$. □

4.2 Parallels between Geometric and Combinatorial Properties

Given some path λ on an orbifold $\Phi\backslash\mathbb{H}$ and some triangulation T of $\Phi\backslash\mathbb{H}$, we can then define the cutting sequence (λ, T) as described above. Throughout this section, we shall assume that λ is reduced relative to to T and therefore, the cutting sequence (λ, T) is also reduced. Since, we will primarily be dealing with reduced paths, we will often take the geodesic representative, i.e. we will often work with a geodesic ray ζ. If this cutting sequence (ζ, T) has some combinatorial property (i.e. it is periodic or recurrent), it is natural to ask whether this induces a geometric property on ζ. In particular, if (ζ, T) has some property P, can we show that (ζ, T') has the same property P, where T' is any other triangulation of $\Phi\backslash\mathbb{H}$? This is possible for certain properties, and recurrency is an example of such a property. We use this to show that eventually recurrent continued fractions remain eventually recurrent under rational multiplication and addition.

4.2.1 Properties of Cutting Sequences

Recurrent Cutting Sequences

Let $\Phi\backslash\mathbb{H}$ be an orbifold with triangulation T, let γ be an edge of T, and let ζ be a geodesic ray starting at the arc γ. If ζ transversely intersects γ infinitely often, then we can homotope each of these intersection points to a common point lying on γ, which we denote P. Note that this process does not introduce new intersections of γ with T, and so we do not have to worry about forming bigons with T. This process decomposes ζ into a potentially infinite collection of closed curves $\mathcal{C}(\zeta, \gamma)$, each based at the point P. If two elements in $\mathcal{C}(\zeta, \gamma)$ are homotopic (or equivalent under reduction), we will consider them as the same element. Due to construction, each of these closed curves will be reduced relative to γ, i.e. each closed curve will only intersect γ at the start and end point P and none of these closed curves will be homotopic to a point. Since ζ intersects γ infinitely often, the geodesic ray ζ will be homotopic to the infinite product of closed curves $c_1 \circ c_2 \circ \ldots$, where $c_i \in \mathcal{C}(\zeta, \gamma)$ and $\circ$ is the standard loop concatenation, i.e. $c_1 \circ c_2$ is the loop which goes round the curve c_1 and then the curve c_2. For each geodesic ray ζ, we can take $\mathcal{C}(\zeta, \gamma)$ to be an alphabet and construct the infinite word $w(\zeta, \gamma) := c_1 c_2 \ldots$, which is formed by appending (the labels of) the closed curves appearing in ζ, in order. Given any sub-word of $w(\zeta,\gamma)$, we refer to the path formed by composing the loops that appear in this sub-word as a *product of loops in* ζ. We denote the set of all products of loops by $\mathcal{P}(\zeta,\gamma)$. We say that ζ is *strictly geometrically recurrent* relative to γ, if we can form the word $w(\zeta,\gamma)$, and the word $w(\zeta,\gamma)$ is strictly recurrent. We say that ζ is *eventually geometrically recurrent* relative to an arc γ, if ζ (with a fixed starting point) is homotopic to a finite length path ρ followed by a geodesic ray ξ, where ξ is recurrent relative to γ, i.e. $\zeta \simeq \rho \circ \xi$. We refer to ξ as the *recurrent component.*

We want to show that the above definition is well defined, i.e. if ζ is eventually geometrically recurrent relative to some arc γ, then ζ is eventually geometrically recurrent relative to all other arcs which the recurrent component intersects. The next lemma proves exactly that.

Lemma 4.2.1. *Let ζ be a geodesic ray on an orbifold $\Phi\backslash\mathbb{H}$ and assume that ζ is strictly geometrically recurrent relative to some arc γ. Then ζ is eventually geometrically recurrent relative to γ', where γ' is any other arc that ζ intersects.*

Proof. Let ζ be strictly geometrically recurrent relative to some arc γ. Let c be some closed curve in $\mathcal{C}(\zeta,\gamma)$, which intersects γ' at least once. Since ζ is strictly geometrically recurrent, the closed curve c is a sub-path of ζ infinitely often. Therefore, the arc γ' is intersected infinitely often by ζ. Let ρ be the shortest (initial) path along ζ connecting γ to γ' and let ξ be the geodesic ray such that $\zeta \simeq \rho \circ \xi$. We define $\mathcal{C}(\xi,\gamma')$ to be the collection of closed curves that ξ decomposes into relative to γ'.

Let π be any product of curves in $\mathcal{P}(\zeta,\gamma)$ which intersects γ' at least once. We can express π as $\pi \simeq p_\pi \circ \nu \circ s_\pi$, where p_π is a path from γ to γ', ν is a product of closed curves in $\mathcal{P}(\xi,\gamma')$ and s_π is a path from γ' to γ. Here p_π and s_π are both disjoint from γ' except at the end/start point respectively. Since π is a product of curves in $\mathcal{P}(\zeta,\gamma)$ and ζ is geometrically recurrent relative to γ, the sub-word $w(\pi,\gamma)$ must occur in $w(\zeta,\gamma)$ infinitely often. It follows from this, that the sub-word $w(\nu,\gamma')$ occurs infinitely often in $w(\xi,\gamma')$. Given any such product of closed curves π in

$\mathcal{P}(\zeta, \gamma)$ we can construct a product of curves $\pi^{'}$, which contains all loops along ζ until the end of π, i.e. if $\zeta = c_1 \circ c_2 \circ \ldots \circ c_n \circ \pi \circ c_{n+1}$, we take $\pi^{'} = c_1 \circ c_2 \circ \ldots \circ c_n \circ \pi$. Thus, either $\pi \simeq \pi^{'}$, in which case p_π is equivalent to the initial path ρ, or p_π is a sub-path of some loop based at $P^{'}$, as required.

It follows from the above statements that every product of curves ν in $\mathcal{P}(\xi, \gamma^{'})$ is induced by a product of curves π in $\mathcal{P}(\zeta, \gamma)$. Since every product of curves in $\mathcal{P}(\zeta, \gamma)$ that occurs once, occurs infinitely often, every product of curves in $\mathcal{P}(\xi, \gamma^{'})$ must also occur infinitely often. Therefore, it follows that ζ is eventually geometrically recurrent relative to $\gamma^{'}$. □

The above lemma shows if a geodesic ray ζ is (strictly) geometrically recurrent relative to one arc, then it is eventually geometrically recurrent relative to all arcs it passes through. In this case, we refer to ζ as *eventually geometrically recurrent* and drop the phrase "relative to".

Theorem 4.2.2. *Let $_{\Phi}\backslash^{\mathbb{H}}$ be an orbifold and let ζ be a geodesic ray in $_{\Phi}\backslash^{\mathbb{H}}$, starting at some arc γ_ζ. Then ζ is eventually geometrically recurrent if and only if the cutting sequence (ζ, T) is eventually recurrent, where T is any triangulation of $_{\Phi}\backslash^{\mathbb{H}}$ containing the arc γ_ζ.*

Proof. $(\Rightarrow)$: Let T be any triangulation of $_{\Phi}\backslash^{\mathbb{H}}$ and assume that ζ is geometrically recurrent relative to γ, some edge in T. As seen in the proof of the previous lemma, we can decompose ζ into an initial non-recurrent sub-path ρ, which terminates at γ, followed by a geodesic ray ξ, which start at γ and is (strictly) geometrically recurrent relative to γ. We can then construct the collection of curves $\mathcal{C}(\xi, \gamma)$ and the word $w(\xi, \gamma)$. We can write $w(\xi, \gamma) = \pi_1 \pi_2 \ldots$, where each π_i represents a curve in $\mathcal{C}(\xi, \gamma)$. The cutting sequence of ξ relative to T, can similarly broken up into the product of the cutting sequences of the closed curve decomposition, i.e. $(\xi, T) = (\pi_1, T) \cdot (\pi_2, T) \cdot \ldots$. Since each closed curve π_i starts and ends at γ, which is an edge in T, the cutting sequences $(\pi_{,i}, T)$ are all well-defined. We now wish to show that if any sub-word $W \in \{L, R\}^*$ appears in the cutting sequence (ξ, T) once, then it appears in the cutting sequence appears infinitely often.

Let W be any sub-word of (ξ, T). Then, we can write the cutting sequence (ξ, T) as a product of words UWV, for some words U and V, since W is a sub-word of (ξ, T). We can equivalently write the cutting sequence (ξ, T) as the infinite product of cutting sequences $(\pi_1, T) \cdot \ldots \cdot (\pi_j, T) \cdot (\pi_{j+1}, T) \cdot \ldots \cdot (\pi_{j+k}, T) \cdot \ldots$, where $\xi = \pi_1 \circ \pi_2 \circ \ldots$. The word W is finite, and so there is some finite product of cutting sequences $(\pi_j, T) \cdot (\pi_{j+1}, T) \cdot \ldots \cdot (\pi_{j+k}, T)$, which contains W. Since ξ is (strictly) geometrically recurrent relative to γ, the corresponding sub-word $\pi_j \pi_{j+1} \ldots \pi_{j+k}$ must occur in $w(\xi, \gamma)$ infinitely often. By extension, the word W must occur in the cutting sequence (ξ, T) infinitely often. Since W was arbitrarily chosen, the cutting sequence (ξ, T) must be recurrent. Because ζ can be decomposed into the paths $\rho \circ \xi$, the cutting sequence can similarly be decomposed into $(\zeta, T) = (\rho, T) \cdot (\xi, T)$. The path ρ is a finite length path and so the cutting sequence (ρ, T) is a finite word. It follows that the cutting sequence (ζ, T) is eventually recurrent, as required.

$(\Longleftarrow)$: Assume that that cutting sequence (ζ, T) is eventually recurrent. We split (ζ, T) into a finite non-recurrent word P and an infinite recurrent word W. We can similarly split ζ along some γ in T such that $\zeta \simeq \rho \circ \xi$, where $(\rho, T) = P$ and $(\xi, T) = W$.

Without loss of generality, we can assume that ξ has positive direction of departure from γ, i.e. the direction of departure is $(\gamma, +)$. Note that T is a finite triangulation, i.e. it is made up of finitely many edges. As a result, if we take $\mathbf{D}_T$ to be the set of all possible directions of departure for paths starting at some edge of T, then we see that this set is also finite - there are K arcs of the triangulation and each arc can be approached in twos ways, so this set is of size $2K$. Similarly, the set of directions of approach $\mathbf{A}_T$ will also be finite.

Note that for simplicity, we will freely assume that any path which starts at an even edge has positive direction of departure and any path which terminates at an even edge has negative direction of approach. In practice, this means that we will never use some directions of departure, but this does not affect the overall constructions.

We note that if V_1 is an arbitrary sub-word of W with some implicit positioning, then we can associate a sub-path λ_{V_1} of ξ to it. This sub-path is precisely the section of ξ which contributes the word V_1 to the cutting sequence (in the required position). In order to prove this direction of the proof, we will show that the sub-path λ_{V_1} lies within a loop/product of loops, and that this loop/product of loops occurs infinitely often in $w(\xi, \gamma)$. Since the initial sub-word V_1 is arbitrarily chosen, the corresponding sub-path represents all possible sub-paths of ξ. In particular, every sub-path of ξ lies within a product of loops, and this product of loops occurs in $w(\xi, \gamma)$ infinitely often.

To make this process somewhat easier, instead of dealing with arbitrary sub-words V_1 of W, we will instead deal with prefixes W_1 of W. Note that if V_1 is not initially a prefix of W, then we can always find a prefix U of W, such that UV_1 is a prefix. In this case, we can simply take $W_1 = UV_1$. Also note, that since W_1 is a prefix of W, the corresponding sub-path λ_{W_1} starts at the edge γ and has direction of departure $(\gamma, +)$. Copies of the prefix W_1 occur infinitely often in W, since W is recurrent. To distinguish these copies, we will add a subscript based off of the natural ordering of these copies, starting with $W_1 = W_{1,0}$. In particular, we can decompose the word W as follows, $W = W_{1,0}V_1W_{1,1}V_2W_{1,2}\dots$. In general, the induced paths $\lambda_{W_{1,k}}$ will rarely start at $(\gamma, +)$ and will therefore lead to a collection of different paths (when considered up to homotopy relative to their starting edges and terminal edges). However, since there finitely many directions of departure and there are an infinite number of copies of W_1 in W, we can conclude that infinitely many of the paths $\lambda_{W_{1,k}}$ must start at the same edge with the same direction of departure. We define $D(W_1)$ to be the set of all pairs $(\gamma', \pm)$ which are directions of departure for infinitely many paths $\lambda_{W_{1,k}}$. By construction $D(W_1)$ is a non-empty subset of $\mathbf{D}_T$. Note that we can similarly define the set $D(V)$, where V is an arbitrary sub-word of W.

We can also construct the set $A(W_1)$ to be the set of all pairs $(\gamma', \pm)$ which are directions of approach for infinitely many paths $\lambda_{W_{1,k}}$. Note that there is a duality between the direction of approach and direction of departure; if $(\gamma_1, \pm_1) \in D(W_1)$ is the direction of departure for infinitely many copies of $\lambda_{W_{1,k}}$, then the direction of approach $(\gamma_2, \pm_2)$ will be the same for each of these copies $\lambda_{W_{1,k}}$ and so $(\gamma_2, \pm_2) \in A(W_1)$. In fact we can say something stronger, if $\lambda_{W_{1,j}}$ has direction of departure $(\gamma_1, \pm_1)$ and direction of approach $(\gamma_2, \pm_2)$, then $(\gamma_1, \pm_1) \in D(W_1)$ if and only if $(\gamma_2, \pm_2) \in A(W_1)$. Note that here we have used the assumption that all paths that start at the an even edge have positive direction of departure and all paths that terminate at an even edge have negative direction of approach.

Claim: Let W_1 be any prefix of W. Then, $(\gamma, +)$ is the direction of departure for infinitely many path $\lambda_{W_{1,k}}$. In other words, $(\gamma, +) \in D(W_1)$.

Proof of Claim. We will prove this claim by contradiction.

Assume W_1 is a prefix of W with only finitely many associated paths $\lambda_{W_{1,k}}$ starting at $(\gamma, +)$. Let $(\gamma_1, \pm_1)$ be the direction of approach for $\lambda_{W_{1,0}}$. Since we have assumed that $(\gamma, +) \notin D(W_1)$, it follows that $(\gamma_1, \pm_1) \notin A(W_1)$. However, W contains infinitely many copies of W_1, and so $A(W_1)$ and $D(W_1)$ must both be non-empty. In particular, there must be some pair $(\gamma_2, \pm_2) \in A(W_1)$ which is the direction of approach for infinitely many copies of $\lambda_{W_{1,k}}$. Let $W_{1,j}$ be the first copy of W_1 such that the path $\lambda_{W_{1,j}}$ has direction of approach $(\gamma_2, \pm_2)$. We can now construct a new sub-word of W of the form $W_2 = W_{1,0}V_1 \dots V_j W_{1,j}$. Note that W_2 is a prefix of W, since it starts with $W_{1,0}$, which is a prefix of W. Copies of W_2 appear infinitely often in the word W, since W is recurrent. Again, we will add a subscript based on the natural ordering to distinguish these copies, starting with $W_2 = W_{2,0}$.

Since W_2 has W_1 as a prefix, any direction of departure $(\gamma', \pm')$, which is the direction of departure for infinitely many paths $\lambda_{W_{2,k}}$, must also be a direction of departure for infinitely many paths of the form $\lambda_{W_{1,k}}$. It follows that $D(W_2) \subset D(W_1)$. In particular, since only finitely many paths of the form $\lambda_{W_{1,k}}$ have $(\gamma, +)$ as their direction of departure, only finitely many paths $\lambda_{W_{2,k}}$ have $(\gamma, +)$ as their direction of departure. We can then conclude - by the duality of the direction of approach and direction of departure - that only finitely many copies of $\lambda_{W_{2,k}}$ have $(\gamma_2, \pm_2)$ as their direction of approach. Similarly, since W_2 has W_1 as a suffix, we can conclude that any direction of approach $(\gamma', \pm')$, which is the direction of approach for infinitely many paths $\lambda_{W_{2,k}}$, must also be a direction of approach for infinitely many paths of the form $\lambda_{W_{1,k}}$. As a result, we see that $A(W_2) \subset A(W_1)$. In particular, only finitely many $\lambda_{W_{2,k}}$ have $(\gamma_1, \pm_1)$ as their direction of approach. Gathering together this information, we can now conclude that $(\gamma_1, \pm_1)$ and $(\gamma_2, \pm_2)$ are not elements of $A(W_2)$. Note that $(\gamma_2, \pm_2)$ is an element of $A(W_1)$ and not $A(W_2)$, and therefore $A(W_2)$ is a proper subset of $A(W_1)$.

We now define the word W_3 in a similar way. Let $(\gamma_3, \pm_3)$ be the direction of approach for infinitely many copies of $\lambda_{W_{2,k}}$ and let $W_{2,j}$ be first instance of W_2, such that $\lambda_{W_{2,j}}$ has direction of approach $(\gamma_3, \pm_3)$. We can then take W_3 to be the sub-word of W of the form $W_3 = W_{2,0}V_0W_{2,1}\dots V_jW_{2,j}$. Note that W_3 is a prefix of W, since it starts with $W_{2,0}$, and contains a copy of W_2 as both a prefix and a suffix. Since W_2 is a prefix of W_3 and only finitely many paths $\lambda_{W_{2,k}}$ have $(\gamma, +)$ as their direction of departure, we can conclude that only finitely many paths $\lambda_{W_{3,k}}$ have $(\gamma, +)$ as their direction of departure. Since $(\gamma_3, \pm_3)$ is the direction of approach for $\lambda_{W_{3,0}}$, we can conclude by duality of the direction of approach and direction of departure, that only finitely many paths $\lambda_{W_{3,k}}$ have $(\gamma_3, \pm_3)$ as their direction of approach, i.e. $(\gamma_3, \pm_3) \notin A(W_3)$. Since W_2 is a suffix of W_3, we can also conclude that $A(W_3) \subset A(W_2)$ and so, $(\gamma_1, \pm_1)$ and $(\gamma_2, \pm_2)$ are not elements of $A(W_3)$. Note that $(\gamma_3, \pm_3)$ is an element of $A(W_2)$ and not $A(W_3)$, and therefore $A(W_3)$ is a proper subset of $A(W_2)$.

In the above procedure, we have that $A(W_3) \subsetneq A(W_2) \subsetneq A(W_1) \subsetneq \mathbf{A}_T$. Since the triangulation T is made up of finitely many edges, the set of directions of approach $\mathbf{A}_T$ is finite. If we iterate the above procedure (i.e. construct a word W_{i+1} to be a word with prefix $W_{i,0}$ and suffix $W_{i,j}$, where $\lambda_{W_{i,j}}$ has direction of approach $(\gamma_i, \pm_i) \in A(W_i)$), we see that each set $A(W_{i+1})$ is a proper subset of

$A(W_i)$. Therefore, there must be some $\ell \in \mathbb{N}$ such that $A(W_\ell)$ contains no elements. However, this implies that the word W_ℓ can only appear as a sub-word of W finitely often. This is a contradiction to the fact that W is recurrent. Therefore our initial assumption must be false, and so infinitely many paths of the form $\lambda_{W_{1,k}}$ must have $(\gamma, +)$ as their direction of departure. This prove the claim. *QED.*

By the above claim, we can now assume that $(\gamma, +) \in D(W_1)$. As a result, we can form the word $U_1 = W_{1,0}V_1W_{1,1} \dots V_j$ where $W = W_{1,0}V_1 \dots V_jW_{1,j} \dots$ and $\lambda_{W_{1,j}}$ is the next occurrence W_1 which has direction of departure $(\gamma, +)$. Since $\lambda_{W_{1,j}}$ has direction of departure $(\gamma, +)$ and $W = U_1W_{1,j} \dots$, the corresponding path λ_{U_1} has direction of departure $(\gamma, +)$ and direction of approach $(\gamma, -)$. We can therefore conclude that λ_{U_1} is a product of loops. Note that by construction U_1 is also a prefix of W. As a result, we can conclude that $(\gamma, +) \in D(U_1)$. In particular, the product of loops λ_{U_1} occurs in the word of loops $w(\xi, \gamma)$ infinitely often, as required. □

Periodic Cutting Sequences

A similar process can be used to look at paths with periodic cutting sequences. If a path ζ goes around a closed curve infinitely many times without deviating, then the cutting sequence (ζ, T) will be strictly periodic for any triangulation T. The reverse is also true: a path with strictly periodic cutting sequences is homotopic to a path which goes around a closed curve infinitely often (up to the equivalence relations on even edges). Note that this statement is not obvious, but follows from a slight adjustment in the proof of Theorem 4.2.2:

Proposition 4.2.3. *Let $\Phi\backslash\mathbb{H}$ be an orbifold and let ζ be a geodesic ray in $\Phi\backslash\mathbb{H}$, starting at some arc γ. Then ζ is homotopic to a path (relative to γ) which goes around a closed curve infinitely many times (up to equivalence relations) if and only if the cutting sequence (ζ, T) is strictly periodic, where T is any triangulation of $\Phi\backslash\mathbb{H}$ containing the arc γ.*

Proof. $(\Rightarrow)$: Since (up to homotopy and equivalence relations) ζ goes around a closed curve π infinitely many times, the word $w(\zeta,\gamma)$ can be written as the infinite product of π, i.e. $w(\zeta,\gamma) = \pi^{\infty}$. Since γ is in T, the cutting sequence (π,T) is well-defined relative to T. In particular, $(\zeta,T) = \prod_{i=1}^{\infty}(\pi,T)$. Let $W = (\pi,T)$, then $(\zeta,T) = \overline{W}$ and, therefore, is periodic.

$(\Leftarrow)$: If $(\zeta,T) = \overline{L^{n_0}R^{n_1}\cdots L^{n_k}}$, take W_1 to be the prefix $L^{n_0}R^{n_1}\cdots L^{n_\ell}$. If ζ has direction of departure $(\gamma,+)$, then so does our initial path λ_{W_1}. If λ_{W_1} has direction of approach $(\gamma,-)$, then the path λ_{W_1} forms a loop. Since the cutting sequence (ζ,T) can be represented as a product of infinitely many copies of W_1, ζ can be viewed as the path which goes around the loops λ_{W_1} infinitely many times, i.e. ζ is homotopic to a path which goes around a closed curve infinitely many times.

Assume λ_{W_1} has direction of approach $(\gamma_1,\pm_1) \neq (\gamma,-)$, i.e. assume it does not loop back up. Then we can take W_2 to be the prefix containing two copies of the period W_1, i.e. $W_2 = W_{1,0}W_{1,1}$. If λ_{W_2} has direction of approach $(\gamma_2,\pm_2) \neq (\gamma,-)$, then λ_{W_2} does produce a loop either. In this case, λ_{W_2} can not have direction of approach $(\gamma_1,\pm_1)$, since this would imply that $\lambda_{W_{1,1}}$ has direction of approach $(\gamma_1,\pm_1)$. This could only occur if $\lambda_{W_{1,0}}$ and $\lambda_{W_{1,1}}$ had the same direction of departure, i.e. $\lambda_{W_{1,0}}$ is a loop. If we then take W_3 to be the prefix of (ζ,T) containing three copies of W_1, i.e. $W_{1,0}W_{1,1}W_{1,2}$, then we start see a pattern. Either λ_{W_3} is a loop or it can not have the same directions of approach as either λ_{W_2} or λ_{W_1}. By the pigeonhole principle, there must be some W_k and $W_{k'}$ with $k < k'$, such that the path λ_{W_k} and $\lambda_{W_{k'}}$ have the same direction of approach. If $k = 0$, we are done. Otherwise, we note that we can express the path λ_{W_k} as the concatenation of paths $\lambda_{W_{1,0}} \circ \lambda_{W_{1,1}} \circ \ldots \circ \lambda_{W_{1,k}}$, and $\lambda_{W_{k'}}$ as the concatenation of paths $\lambda_{W_{1,0}} \circ \lambda_{W_{1,1}} \circ \ldots \circ \lambda_{W_{1,k'}}$. We then see that, since λ_{W_k} and $\lambda_{W_{k'}}$ have the same direction of approach, the paths $\lambda_{W_{1,k}}$ and $\lambda_{W_{1,k'}}$ are equivalent up to homotopy and reduction, and these paths have the same direction of departure. Therefore, we can deduce that $\lambda_{W_{k-1}}$ and $\lambda_{W_{k'-1}}$ have the same direction of approach. Iterating this process we get that $\lambda_{W_{1,0}}$ has the same direction of approach as $\lambda_{W_{1,k'-k}}$. In particular, $\lambda_{W_{1,k'-k-1}}$, must have direction of approach $(\gamma,-)$ and so $\lambda_{W_{1,k'-k-1}}$ forms a loop. Since (ζ,T) can also be expressed as product of infinitely many copies of $W_{k'-k-1}$, we see that ζ goes round the loop $\lambda_{W_{1,k'-k-1}}$ infinitely many times, as required. □

4.2.2 Properties of Cutting Sequences under Triangulation Replacement

As seen in Theorem 4.2.2, if ζ starts at an arc γ, then ζ is eventually geometrically recurrent if and only if the cutting sequence (ζ, T) is eventually recurrent, where T is any triangulation of $_{\Phi}\backslash^{\mathbb{H}}$ containing the arc γ. We can therefore deduce that if (ζ, T) is eventually recurrent for some triangulation T on an orbifold $_{\Phi}\backslash^{\mathbb{H}}$, then ζ is eventually geometrically recurrent. If we then replace the triangulation T with a triangulation T' containing the starting arc γ, then since ζ is eventually geometrically recurrent, we see that the cutting sequence (ζ, T') is eventually recurrent. This gives us the following corollary:

Corollary 4.2.4. *Let $_{\Phi}\backslash^{\mathbb{H}}$ be an orbifold and let ζ be a geodesic ray in $_{\Phi}\backslash^{\mathbb{H}}$, starting at some arc γ_ζ. If there is some triangulation T containing the arc γ such that (ζ, T) is eventually recurrent, then (ζ, T') is eventually recurrent, where T' is any other triangulation of $_{\Phi}\backslash^{\mathbb{H}}$ containing the starting arc γ_ζ.*

We can similarly use Proposition 4.2.3 to see that if a geodesic ray ζ has an eventually periodic cutting sequence (ζ, T), for some triangulation T on an orbifold $_{\Phi}\backslash^{\mathbb{H}}$, then (ζ, T') will be eventually periodic, where T' is any other triangulation of $_{\Phi}\backslash^{\mathbb{H}}$.

Corollary 4.2.5. *Let $_{\Phi}\backslash^{\mathbb{H}}$ be an orbifold and let ζ be a geodesic ray in $_{\Phi}\backslash^{\mathbb{H}}$, starting at some arc γ_ζ. If there is some triangulation T containing the arc γ_ζ such that (ζ, T) is eventually periodic, then (ζ, T') is eventually periodic, where T' is any other triangulation of $_{\Phi}\backslash^{\mathbb{H}}$ containing the starting arc γ_ζ.*

The same observations also hold for geodesic rays with strictly periodic and strictly recurrent cutting sequences. We can therefore deduce that strictly periodic, eventually periodic, strictly recurrent and eventually recurrent cutting sequences each form a closed class of with regards to triangulation replacement on a triangulated orbifold. Note that given a strictly periodic/strictly recurrent cutting sequence, the corresponding continued fractions need not necessarily be strictly periodic/strictly recurrent. For strictly periodic cutting sequences, this is because the periodic block in the cutting sequence may start and end with the same letter. As an example, if we take the cutting sequence $(\zeta, T) = \overline{L^3R^2LRL^2}$, then the corresponding continued fraction expansion would be $[3; \overline{2, 1, 1, 5}]$. We can give such continued fraction expansions a general form:

$$\overline{\alpha} := \begin{cases} [a_0; \overline{a_1, a_2, \ldots, a_{2n}}] & a_0 \leq a_{2n} \\ [0; a_1, \overline{a_2, \ldots, a_{2n+1}}] & a_1 \leq a_{2n+1} \end{cases},$$

where $a_i \in \mathbb{N}$ for $i \geq 0$. If a continued fraction expansion $\overline{\alpha}$ gives rise to a strictly periodic cutting sequence, we refer to it as *essentially periodic*. Similarly, we refer to continued fractions leading to recurrent cutting sequences as *essentially recurrent*.

Triangulation Replacement and Continued Fractions

In general, it is not known how triangulation replacement explicitly affects cutting sequences or the corresponding continued fraction expansions. However, we do know that specific triangulation replacement on specific orbifolds represents integer multiplication of the associated continued fraction, as seen in Theorem 4.1.8. We can also represent integer division by taking this triangulation replacement in reverse. From Proposition 4.2.2, we can deduce that if the cutting sequence of a geodesic ray ζ relative to some triangulation T is recurrent, then the cutting sequence (ζ, T') is recurrent for T' any other triangulation of $\Phi\backslash\mathbb{H}$. In particular, a cutting sequence being (eventually) recurrent is independent of the triangulation. If a continued fraction expansion $\overline{\alpha}$ is eventually recurrent, then this implies that the cutting sequence of $(\widehat{\zeta_\alpha}, \widehat{\mathcal{F}})$ will be eventually recurrent. Since this is independent of triangulation, the cutting sequence $(\widehat{\zeta_\alpha}, \widehat{\frac{1}{n}\mathcal{F}})$ is also eventually recurrent. In particular, if $\overline{\alpha}$ is an eventually recurrent continued fraction expansion then so is $\overline{q\alpha}$, where $q \in \mathbb{Q}$. The same is true if we replace the cutting sequence being eventually recurrent with being strictly recurrent/strictly periodic/eventually periodic.

For eventually recurrent continued fractions, we see that rational addition also preserves the property of being eventually recurrent. As does taking $\overline{\frac{1}{\alpha}}$, since this is equivalent to replacing all R's with L's in the corresponding cutting sequence and vice versa. From this, we obtain the following corollary of Theorem 4.2.2. Note that in the proof of this corollary, we repeat (and add to) some of the above arguments for completeness.

Corollary 4.2.6. *Let $\alpha \in \mathbb{R}$, let $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ be a non-trivial integer matrix (i.e. $a, b, c, d \in \mathbb{Z}$, $ad - bc \neq 0$), and let $\beta = M \cdot \alpha = \frac{a\alpha+b}{c\alpha+d}$. If the continued fraction expansion $\overline{\alpha}$ is eventually recurrent and $c\alpha + d \neq 0$, then the continued fraction $\overline{\beta}$ is eventually recurrent.*

Proof. For such a $\alpha \in \mathbb{R}$, let ζ_α be a corresponding geodesic ray in $\mathbb{H}$. Since the continued fraction expansion is equivalent to the cutting sequence $(\zeta_\alpha, \mathcal{F})$, this cutting sequence will also be eventually recurrent. We take $\Gamma_0(n)\backslash\mathbb{H}$ to be our orbifold, for some arbitrary $n \in \mathbb{N}$. Then the two triangulations of $\Gamma_0(n)\backslash\mathbb{H}$ given by $\widehat{\mathcal{F}}$ and $\widehat{\frac{1}{n}\mathcal{F}}$, encode the structure of $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$. In particular, if we take $\widehat{\zeta_\alpha}$ to be the projection of ζ_α on $\Phi\backslash\mathbb{H}$, then $(\widehat{\zeta_\alpha}, \widehat{\mathcal{F}}) = (\zeta_\alpha, \mathcal{F})$ and $(\widehat{\zeta_\alpha}, \widehat{\frac{1}{n}\mathcal{F}}) = (\zeta_\alpha, \frac{1}{n}\mathcal{F})$ by Theorem 4.1.8. Since the cutting sequence $(\zeta_\alpha, \mathcal{F})$ is eventually recurrent, so is the associated cutting sequence $(\widehat{\zeta_\alpha}, \widehat{\mathcal{F}})$ on $\Phi\backslash\mathbb{H}$. It follows from Theorem 4.2.2 that the cutting sequence $(\widehat{\zeta_\alpha}, \widehat{\frac{1}{n}\mathcal{F}})$ is also eventually recurrent. In particular, the continued fraction expansion of $\overline{n\alpha} = \eta(\zeta_\alpha, \frac{1}{n}\mathcal{F})$ will also be eventually recurrent. Therefore, we can conclude that recurrency of the continued fraction expansion is preserved under integer multiplication. Using a similar argument, we can extend this result to show that recurrency is preserved under non-zero integer division, and by extension rational multiplication. Recurrency of the continued fraction expansion is also preserved under integer addition, and so by composing these operations we can conclude that if $\alpha \in \mathbb{R}$ has an eventually recurrent continued fraction expansion, then $q\alpha + r$ will also has an eventually recurrent continued fraction expansion, where $q, r \in \mathbb{Q}$.

Let $a, b, c, d \in \mathbb{Z}$ with $ad - bc \neq 0$. If $c = 0$, then the result follows trivially. Instead, assume $c \neq 0$. It follows that, if $\alpha \in \mathbb{R}$ has an eventually recurrent continued fraction expansion, then $c\alpha + d$ will have an eventually recurrent continued fraction expansion, since $c, d \in \mathbb{Z}$. We know that if $\beta \in \mathbb{R} \setminus \{0\}$ admits an eventually recurrent continued fraction expansion, then so will $\frac{1}{\beta}$. This follows, since $\overline{\frac{1}{\beta}} = [0; b_0, b_1, \ldots]$, where $\overline{\beta} = [b_0; b_1, \ldots]$. Therefore, for $c\alpha + d \neq 0$, $\frac{1}{c\alpha+d}$ admits a recurrent continued fraction expansion. Let $r = \frac{a}{c}$ and $q = b - \frac{ad}{c}$. It follows that $q, r \in \mathbb{Q}$, and $q \neq 0$, since $ad - bc \neq 0$. Since recurrency of the continued fraction expansion is preserved by rational multiplication and addition, the continued fraction expansion of $q \cdot \frac{1}{c\alpha+d} + r$ is also eventually recurrent. Then:

$$
\begin{aligned}
q \cdot \frac{1}{c\alpha + d} + r &= (b - \frac{ad}{c}) \cdot \frac{1}{c\alpha + d} + \frac{a}{c} \\
&= \frac{b - \frac{ad}{c} + \frac{a}{c} \cdot (c\alpha + d)}{c\alpha + d} \\
&= \frac{b - \frac{ad}{c} + a\alpha + \frac{ad}{c}}{c\alpha + d} \\
&= \frac{a\alpha + b}{c\alpha + d}
\end{aligned}
$$

And the result follows. □

One consequence is that if a property of the cutting sequence is preserved by triangulation replacement, by adding a L^m term to the start of the sequence for $m \in \mathbb{Z}$ and by replacing all L's with R's and all R's with L's, then this property is preserved by the transformation $\frac{a\alpha+b}{c\alpha+d}$, for all $a, b, c, d \in \mathbb{Z}$ with $ad - bc \neq 0$.

This leaves two interesting avenues open for future research: Firstly, what transformations of continued fraction expansions are induced by triangulation replacement? In particular, can we express any other of these transformations explicitly (as we did for integer multiplication and division)? Secondly, what other properties of cutting sequences are preserved by triangulation replacement?

Chapter 5

Infinite Loops as Potential Counterexamples to the p-adic Littlewood Conjecture

In this chapter we will discuss *infinite loops* mod n. Infinite loops can be viewed in two equivalent ways. Firstly, we can define an infinite loop mod n to be a geodesic ray ζ_α which starts at the y-axis and terminates at some point $\alpha \in \mathbb{R}_{>0}$ and is disjoint from the set $\Gamma_0(n) \cdot I$ except for the edges of the form $I + k$, where $k \in \mathbb{Z}_{\geq 0}$. Secondly, we can define an infinite loop mod n to be a real number $\alpha \in \mathbb{R}_{>}0$, which has no semi-convergent denominators divisible by n (except for $q_{-1} = 0$). Since semi-convergents will be frequently used throughout this chapter, we recall the definition of a semi-convergent for the benefit of the reader.

Definition 2.1.6. Let $\overline{\alpha} = [a_0; a_1, a_2, \ldots]$ be a continued fraction expansion of some real number α. We define the $\{k,m\}$*-th semi-convergent* of $\overline{\alpha}$ to be $\frac{p_{\{k,m\}}}{q_{\{k,m\}}} := [a_0; a_1, \ldots, a_k, m]$, where $0 \leq m \leq a_{k+1}$. We can define this iteratively using the standard convergents:

$$p_{\{k,m\}} = mp_k + p_{k-1}$$

$$q_{\{k,m\}} = mq_k + q_{k-1}$$

We refer to the term $p_{\{k,m\}}$ as the $\{k,m\}$*-th semi-convergent numerator* of α and $q_{\{k,m\}}$ as the $\{k,m\}$*-th semi-convergent denominator*.

Before discussing infinite loops in full generality, we will first discuss the structure of $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$ and the geodesic rays which start at I and are otherwise disjoint from $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$. We will see that the cutting sequences of these geodesic rays behave "badly" when multiplied by n. We will then see that the set of edges $\Gamma_0(n) \cdot I$ is a subset of the edges in $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$, and that these sets of edges are equal if and only if n is a prime power. As a result, if n not a prime power, a geodesic ray being an infinite loop mod n is actually a slightly weaker condition than this geodesic ray being disjoint from $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$ (except at $I+k$ for $k \in \mathbb{Z}_{\geq 0}$). However, by duality if ζ_α is not an infinite loop mod n, then this is a stronger condition than saying ζ_α intersects $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$. Moreover, if ζ_α is not an infinite loop mod n, then we will see that it behaves very nicely when multiplied by n. Therefore, we can view the property of being an infinite loop mod n as some kind of minimal condition for the multiplication algorithm to not behave nicely.

After discussing the structure of $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$, we will then discuss the structure of $\Gamma_0(n) \cdot I$ and the preliminary properties of infinite loops mod n. We will then use these properties to look at the role infinite loops play in the p-adic Littlewood Conjecture. Firstly, we have the following lemma:

Lemma 5.1.10. *Assume that α is not an infinite loop mod n. Then:*

$$\max\{B(\alpha), B(n\alpha)\} \geq \lfloor 2\sqrt{n} \rfloor - 1.$$

Remark 5.0.1. Recall that $B(\alpha)$ is the *height function*, which outputs the largest partial quotient in the continued fraction expansion, i.e:

$$B(\alpha) := \sup_{k\in\mathbb{N}} \{a_i : \overline{\alpha} = [a_0; a_1, \ldots]\}.$$

This lemma tells us that if α is not an infinite loop mod n, then the largest partial quotients for α and $n\alpha$ can not both be "small" relative to $\sqrt{n}$. As we will see in more detail later, this lemma allows us show that if there is a sequence of natural numbers $\{\ell_m\}_{m\in\mathbb{N}}$ such that $p^{\ell_m}\alpha$ is not an infinite loop mod p^m, then α satisfies pLC. Here, we note that the sequence $\{\ell_m\}$ need not be increasing. This is a fact which we will discuss later in Section 5.1.3. However, if no such sequence exists, α is a counterexample to pLC. In particular, we have the following lemma:

Lemma 5.1.14. *Let $\alpha \in$ **Bad** and assume there exists an $m \in \mathbb{N}$ such that $p^{\ell}\alpha$ is an infinite loop mod p^m for all $\ell \in \mathbb{N}$. Then α is a counterexample to pLC.*

Combining these statements together, we get the following reformulation of pLC, written as a condition of infinite loops mod p^m:

Theorem 5.1.15. *Let $\alpha \in$ **Bad**. Then, α satisfies pLC if and only if there is a sequence of natural numbers $\{\ell_m\}_{m\in\mathbb{N}}$ such that $p^{\ell_m}\alpha$ is not an infinite loop mod p^m.*

In Section 5.2.1 we will provide two ways of constructing infinite loops mod n. One of these ways is geometric in nature and the other is arithmetic. We will describe the theoretical process of constructing both of these algorithms for some arbitrary n. We finish this section by constructing the algorithm of $p = 5$ and 7. We do this using both the geometric algorithm and the arithmetic algorithm.

5.1 Infinite Loops and the p-adic Littlewood Conjecture

We begin this section by defining an *infinite loop mod n* (as a geodesic ray).

Definition 5.1.1. Let ζ_α be a geodesic ray starting at the y-axis I and terminating at the point $\alpha \in \mathbb{R}_{>0}$. Then ζ_α is an *infinite loop mod n*, if ζ_α is disjoint from $\Gamma_0(n)\cdot I$ except for the edges of the form $I + k$, for $k \in \mathbb{Z}_{\geq 0}$.

Remark 5.1.2. Here we allow ζ_α to intersect the the lines between $a \in \mathbb{N} \cup \{0\}$ and ∞. This is because the integer part of α does not affect how α multiplies. In particular, if $\{\alpha\}$ is the fractional part of α, then $\{n\{\alpha\}\}$ will be the fractional part of $n\alpha$, i.e. for $a = \lfloor\alpha\rfloor$ and $\{x\} := x - \lfloor x\rfloor$, we have $\{n\alpha\} = \{n\alpha - na\} = \{n\{\alpha\}\}$. As remarked in Chapter 2, α satisfies mLC if and only if $\{\alpha\}$ satisfies mLC.

As we will soon see, the set of edges $\Gamma_0(n)\cdot I$ is a subset of the set of edges in $\mathcal{F}\cap\frac{1}{n}\mathcal{F}$. In order to motivate the definition of an infinite loop mod n, we will initially discuss geodesic rays which are disjoint from $\mathcal{F}\cap\frac{1}{n}\mathcal{F}$. We will then explain why we loosen this condition when defining infinite loops.

5.1.1 Initial Motivation for Looking at Infinite Loops

If λ is a finite length path, running from one edge E_i of $\mathcal{F}$ to another edge E_j in $\mathcal{F}$, the cutting sequence of $(\lambda,\mathcal{F})$ and direction of departure completely defines the path λ up to relative homotopy. For example, if λ has direction of departure $(E_i,+)$ and cutting sequence $W = (\lambda,\mathcal{F})$, we can then construct the path λ_W with direction of departure $(E_i,+)$. The path λ_W will terminate at the edge E_j and λ_W is homotopic to λ relative to E_i and Ej. However, when we take the cutting sequence of λ relative to $\frac{1}{n}\mathcal{F}$, the cutting sequence $(\lambda,\frac{1}{n}\mathcal{F})$ is only well-defined if E_i and E_j are both edges of $\frac{1}{n}\mathcal{F}$. In particular, if E_i and E_j are not edges $\frac{1}{n}\mathcal{F}$, then we can not reconstruct the path λ just by knowing the direction of departure and cutting sequence $(\lambda,\frac{1}{n}\mathcal{F})$.

As we saw in Section 3.2.3, if we build a path λ_W starting at I out of smaller sub-paths $\{\lambda_i\}_{i\in\mathbb{N}}$ (each with well defined cutting sequence relative to $\mathcal{F}$), then even if none of the sub-paths λ_i have a well-defined cutting sequence relative to $\frac{1}{n}\mathcal{F}$, the next sub-path λ_{i+1} determines how $\lambda_i\circ\lambda_{i+1}$ intersects $\frac{1}{n}\mathcal{F}$, and so the concatenated path λ_W will still form a well-defined cutting sequence $(\lambda_W,\frac{1}{n}\mathcal{F})$ relative to $\frac{1}{n}\mathcal{F}$. In the case that none of the sub-paths λ_i terminate at an edge E of $\mathcal{F}\cap\frac{1}{n}\mathcal{F}$, we see that no sub-path of the form $\lambda_1\circ\lambda_2\circ\ldots\circ\lambda_k$ has a well defined cutting sequence relative to both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ simultaneously. In particular, no sub-path of λ_W is simultaneously a good approximation of $(\lambda_W,\mathcal{F})$ and $(\lambda_W,\frac{1}{n}\mathcal{F})$. On the other hand, if there is some λ_j which has an endpoint in both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$, then this tells us that the path $\lambda_1 \circ \lambda_2 \circ \ldots \circ \lambda_j$ has a well-defined cutting sequence relative to both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$. As we will see later, this implies that some pair of semi-convergents of $\overline{\alpha}$ induce a pair of semi-convergents of $\overline{n\alpha}$. We can think of this pair of semi-convergents in $\overline{\alpha}$ as good approximations of α relative to multiplication by n. Since pLC and mLC are deeply tied to the behaviour of continued fractions under iterative multiplication, the real numbers which correspond to geodesic rays that are disjoint from $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$ seem like a natural place to find potential counterexamples.

The Structure of $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$

Recall from Section 3.1.2 that two points $A = \frac{p}{q}$ and $B = \frac{r}{s}$ in $\mathbb{Q} \cup \{\infty\}$ are neighbours in $\mathcal{F}$ if and only if $|ps - rq| = 1$. This in turn implies that there is some element $M \in PSL_2(\mathbb{Z})$ such that $M \cdot 0 = A$ and $M \cdot \infty = B$. This matrix M is either of the form $\left(\begin{smallmatrix} p & r \\ q & s \end{smallmatrix}\right)$ or $\left(\begin{smallmatrix} p & -r \\ q & -s \end{smallmatrix}\right)$, depending on whether $ps - rq = 1$ or $ps - rq = -1$, respectively. It is important to note that A and B can only be neighbours in $\mathcal{F}$ if $\gcd(ps, rq) = 1$. By extension we must have that $\gcd(p, r) = \gcd(q, s) = 1$. We can also show that every edge in $\mathcal{F}$ is of the form $M \cdot I$, where $M \in PSL_2(\mathbb{Z})$. Note that if $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ and $M' = \left(\begin{smallmatrix} b & -a \\ d & -c \end{smallmatrix}\right)$, the lines $M \cdot I$ and $M' \cdot I$ will be equivalent with opposite orientation.

Using this information about $\mathcal{F}$, we can deduce similar information about $\frac{1}{n}\mathcal{F}$ by simply scaling $\mathcal{F}$ by the $(n^*)^{-1}$ map. Using this structure, we obtain the following lemma:

Lemma 5.1.3. *Two points A and B are neighbours in both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ if and only if they have reduced form $\frac{a}{n_1 c_1}$ and $\frac{b}{n_2 d_1}$, with $n = n_1 n_2$ and $|a n_2 d - b n_1 c| = 1$.*

Proof. $(\Rightarrow)$: Assume that $A = \frac{a}{c}$ and $B = \frac{b}{d}$ are neighbours in $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$. Since A and B are neighbours in $\mathcal{F}$, we can conclude that $|ad - bc| = 1$, and more importantly for us:

$$\gcd(c, d) = 1.$$

Since $\frac{1}{n}\mathcal{F}$ is a scaled version of the Farey tessellation, A and B are neighbours in $\frac{1}{n}\mathcal{F}$ if and only if $n^*(A) = n \cdot A = \frac{na}{c}$ and $n^*(B) = n \cdot B = \frac{nb}{d}$ are neighbours in $\mathcal{F}$. Of course, $n \cdot A = \frac{na}{c}$ and $n \cdot B = \frac{nb}{d}$ will not necessarily be in reduced form. We will take $g := \gcd(c, n)$ and $h := \gcd(d, n)$. In this case, we can rewrite c, d and n in the following ways:

$$c = n_1 c_1, \quad n = n_1 g,$$

$$d = n_2 d_1, \quad n = n_2 h.$$

We can then rewrite $n \cdot A$ and $n \cdot B$ in reduced form as:

$$n \cdot A = \frac{n_1 g a}{n_1 c_1} = \frac{ga}{c_1},$$

$$n \cdot B = \frac{n_2 h b}{n_2 d_1} = \frac{hb}{d_1}.$$

Since $n \cdot A$ and $n \cdot B$ are neighbours in $\mathcal{F}$, we see that $|gad_1 - hbc_1| = 1$. Necessarily, we can not have $\gcd(g, h) = r \neq 1$, since this would imply that $|gad_1 - hbc_1| \equiv 0 \mod r$ and so $|gad_1 - hbc_1| \neq 1$. Therefore, we can conclude that:

$$\gcd(g, h) = 1.$$

Since we know that $gcd(c, d) = 1$, $c = n_1 c_1$, and $d = n_2 d_1$, we can conclude that:

$$\gcd(c, d) = 1 = \gcd(n_1 c_1, n_2 d_1) = \gcd(n_1, n_2).$$

Using this equality, we see that:

$$\begin{aligned} n_1 &= \gcd(n_1, n) \\ &= \gcd(n_1, n_2 h) \\ &= \gcd(n_1, n_2) \cdot \gcd(n_1, h) \\ &= 1 \cdot \gcd(n_1, h) \\ &= \gcd(n_1, h). \end{aligned}$$

However, since $\gcd(g,h) = 1$, we can also deduce that:

$$\begin{aligned} h &= \gcd(h,n) \\ &= \gcd(h,n_1 g) \\ &= \gcd(h,n_1) \cdot \gcd(h,g) \\ &= \gcd(h,n_1) \cdot 1 \\ &= \gcd(h,n_1), \end{aligned}$$

and so:

$$n_1 = \gcd(n_1,h) = \gcd(h,n_1) = h.$$

Since $n = n_1 g = n_2 h$, we can now conclude that $g = n_2$, and so:

$$n = n_1 n_2.$$

Combining this information all together, we can now write $A = \frac{a}{n_1 c_1}$ and $B = \frac{b}{n_2 d_1}$ with $|an_2 d_1 - bn_1 c_1| = 1$ and $n = n_1 n_2$, as required.

$(\Leftarrow)$: Let $A = \frac{a}{n_1 c}$ and $B = \frac{b}{n_2 d}$, with $n = n_1 n_2$ and $|an_2 d - bn_1 c| = 1$. Since $|an_2 d - bn_1 c| = 1$, we see that A and B are neighbours in $\mathcal{F}$. Writing $n \cdot A$ and $n \cdot B$ in reduced form, we have that:

$$n \cdot A = \frac{n_2 a}{c},$$

and

$$n \cdot B = \frac{n_1 b}{d}.$$

We can now check to see if $n \cdot A$ and $n \cdot B$ are neighbours in $\mathcal{F}$ by computing the value of $|an_2 d - bn_1 c|$. Here, we have $|an_2 d - bn_1 c| = |adn_2 - bcn_1| = 1$, and so $n \cdot A$ and $n \cdot B$ are indeed neighbours in $\mathcal{F}$. By rescaling by a factor of $(n^*)^{-1}$, we now see that A and B are neighbours in $\frac{1}{n}\mathcal{F}$, as required. □

In the above lemma, requiring the condition that A and B have reduced form $\frac{a}{cn_1}$ and $\frac{b}{dn_2}$ with $n = n_1 n_2$ and $|adn_2 - bcn_1| = 1$, is equivalent to saying that if A and B are neighbours of this form in either $\mathcal{F}$ or $\frac{1}{n}\mathcal{F}$, then necessarily they are neighbours in both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$.

Geodesics Intersecting $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$

If ζ_α is a geodesic ray starting at the y-axis I and terminating at some point $\alpha \in \mathbb{R}_{>0}$, and ζ_α intersects $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$, then we can deduce some nice properties corresponding to the semi-convergents of $\overline{\alpha}$ and $\overline{n\alpha}$. In order to make these statements, we recall Corollary 3.1.23.

Corollary 3.1.23. *Let ζ_α be a geodesic ray starting at the y-axis I and terminating at the point $\alpha \in \mathbb{R}_{>0}$. Then the point $v \in \mathbb{Q} \cup \{\infty\}$ is a semi-convergent of α if and only if it is the endpoint of some edge E in $\mathcal{F}$ which intersects ζ_α. The point $v \in \mathbb{Q} \cup \{\infty\}$ is a convergent of α if and only if it is the point at ∞ or it is the endpoint of at least two edges in $\mathcal{F}$ which intersects ζ_α.*

Using the above proposition, we see that if ζ_α intersects an edge E in $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$, then both of the endpoints of E must be semi-convergents for $\overline{\alpha}$. As a result, $\overline{\alpha}$ has two semi-convergents of the form $\frac{a}{n_1 c_1}$ and $\frac{b}{n_2 d_1}$ with $n = n_1 n_2$ and $|an_2 d - bn_1 c| = 1$. In fact, one of these points will be the *fixed point* of a fan and so will be a convergent, not just a semi-convergent. Without loss of generality, we can take $\frac{p_k}{q_k} = \frac{a}{n_1 c_1}$ to be the k-th convergent of $\overline{\alpha}$ and let $\frac{p_{k,m}}{q_{k,m}} = \frac{b}{n_2 d_1}$ be the $\{k, m\}$-th semi-convergent of $\overline{\alpha}$. Since E is also an edge of $\frac{1}{n}\mathcal{F}$, we can rescale our space using the n^* map. This allows us to see that $n^*(\zeta_\alpha)$ intersects $n^*(E)$, which is an edge in $\mathcal{F}$ with end points $\frac{an_2}{c_1}$ and $\frac{bn_1}{d_1}$. Since $n^*(\zeta_\alpha)$ is homotopic relative to I to $\zeta_{n\alpha}$, the geodesic ray starting at I and terminating at $n\alpha$, we can conclude that both $\frac{an_2}{c_1}$ and $\frac{bn_1}{d_1}$ will be semi-convergents of $\overline{n\alpha}$. In fact, we can guarantee that one of $\frac{an_2}{c_1}$ or $\frac{bn_1}{d_1}$ will be a convergent of $\overline{n\alpha}$. It is worth noting that if $A = \frac{a}{n_1 c_1}$ was a convergent of $\overline{\alpha}$, this does not necessarily mean that $n \cdot A = \frac{n_2 a}{c_1}$ is a convergent of $\overline{n\alpha}$ and we may find that $n \cdot B$ is a convergent of $\overline{n\alpha}$. We reformulate this information in the following proposition:

Proposition 5.1.4. *Let $\alpha \in \mathbb{R}_{>0}$ and assume that $\overline{\alpha}$ has a convergent of the form $\frac{p_k}{q_k} = \frac{a}{n_1 c}$ and semi-convergent of the form $\frac{p_{k,m}}{q_{k,m}} = \frac{b}{n_2 d}$ such that $n = n_1 n_2$ and $|an_2 d - bn_2 c| = 1$. Then $\frac{n_2 a}{c}$ and $\frac{n_1 b}{d}$ are both semi-convergents of $\overline{n\alpha}$. In fact, at least one of $\frac{n_2 a}{c}$ or $\frac{n_1 b}{d}$ will be a convergent for $\overline{n\alpha}$.*

5.1.2 Infinite Loops and their Properties

As we saw in the last section, if a geodesic ray ζ_α intersects $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$, then α has nice properties when multiplied by n. However, as we will soon see, if ζ_α intersects $\Gamma_0(n) \cdot I \subset \mathcal{F} \cap \frac{1}{n}\mathcal{F}$, we can recover even more properties. In particular, the geodesic rays which intersect $\Gamma_0(n) \cdot I$ behave extremely nicely with regards to integer multiplication by n and the geodesic rays which do not intersect $\Gamma_0(n) \cdot I$ behave badly. When $n = p^\ell$ is some prime power, we have $\Gamma_0(p^\ell) \cdot I = \mathcal{F} \cap \frac{1}{p^\ell}\mathcal{F}$. In this case, not only do infinite loops behave badly with respect to multiplication by p^ℓ, but they behave as badly as possible.

The Structure of $\Gamma_0(n) \cdot I$

In Lemma 5.1.3 we saw that two points A and B are neighbours in both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ if and only if they have reduced form $\frac{a}{n_1 c_1}$ and $\frac{b}{n_2 d_1}$ with $n = n_1 n_2$ and $|an_2 d - bn_1 c| = 1$. Given two points $A = \frac{a}{n_1 c_1}$ and $B = \frac{b}{n_2 d_1}$ which are neighbours in both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$, then since we know that $|an_2 d_1 - bn_1 c_1| = 1$, we necessarily must have $\gcd(n_1 c_1, n_2 d_1) = 1$ and, by extension, $\gcd(n_1, n_2) = 1$. As a result, if n is a prime power - i.e. $n = p^\ell$ for p some prime and $\ell \in \mathbb{N}$ - then all edges which are in both $\mathcal{F}$ and $\frac{1}{p^\ell}\mathcal{F}$ have endpoints of the reduced form $A = \frac{a}{p^\ell c}$ and $B = \frac{b}{d}$ with $\gcd(p^\ell c, d) = 1$. If we more generally look at neighbours in both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ of the form $A = \frac{a}{nc}$ and $B = \frac{b}{d}$ with $\gcd(nc, d) = 1$, then, by definition of neighbours in $\mathcal{F}$, we have that $|ad - bnc| = 1$. As a result, we can conclude that either $\left(\begin{smallmatrix} a & b \\ nc & d \end{smallmatrix}\right)$ is an element of $\Gamma_0(n)$ or $\left(\begin{smallmatrix} a & -b \\ nc & -d \end{smallmatrix}\right)$ is an element of $\Gamma_0(n)$. Alternatively, if $\varphi = \left(\begin{smallmatrix} a & b \\ nc & d \end{smallmatrix}\right)$ is some element of $\Gamma_0(n)$, then the endpoints of $\varphi \cdot I$ will be $\varphi \cdot 0 = \frac{b}{d}$ and $\varphi \cdot \infty = \frac{a}{nc}$. As a result, we get the following corollaries about the structure of $\Gamma_0(n) \cdot I$:

Corollary 5.1.5. *Let $A = \frac{a}{nc}$ and $B = \frac{b}{d}$ be two (reduced) elements of $\mathbb{Q} \cup \{\infty\}$. Then A and B are neighbours in $\mathcal{F}$ if and only if the matrix $\varphi = \left(\begin{smallmatrix} a & \pm b \\ nc & \pm d \end{smallmatrix}\right)$ is an element of $\Gamma_0(n) \subset PSL_2(\mathbb{Z})$.*

Note that in the above corollary, if $A = \frac{a}{nc}$ and $B = \frac{b}{d}$ are neighbours in $\mathcal{F}$, then we can use Lemma 5.1.3 to see that A and B are also neighbours in $\frac{1}{n}\mathcal{F}$. This essentially leads to the following corollary:

Corollary 5.1.6. *The set of edges $\Gamma_0(n) \cdot I$ is a subset of $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$. These sets are equivalent if and only if n is a prime power.*

Redefining an Infinite Loop mod n

Recall our previous definition of an infinite loop mod n.

Definition 5.1.1. Let ζ_α be a geodesic ray starting at the y-axis I and terminating at the point $\alpha \in \mathbb{R}_{>0}$. Then ζ_α is an *infinite loop mod n*, if ζ_α is disjoint from $\Gamma_0(n) \cdot I$ except for the edges of the form $I + k$, for $k \in \mathbb{Z}_{\geq 0}$.

As stated in Corollary 5.1.5, two points $A = \frac{a}{cn}$ and $B = \frac{b}{d}$ in $\mathbb{Q} \cup \{\infty\}$ are neighbours in $\mathcal{F}$ if and only if $\varphi = \left(\begin{smallmatrix} a & \pm b \\ nc & \pm d \end{smallmatrix}\right)$ is an element of $\Gamma_0(n)$. In this case, the element φ maps ∞ to $\frac{a}{cn}$ and 0 to $\frac{b}{d}$. By extension, φ maps the y-axis I to the edge between $\frac{a}{cn}$ and $\frac{b}{d}$. Since $\frac{a}{cn}$ and $\frac{b}{d}$ can only be neighbours in $\mathcal{F}$ if $\gcd(cn, d) = \gcd(n, d) = 1$, every edge in $\mathcal{F}$ with endpoint of the form $\frac{a}{nc}$ can be expressed in the form $\varphi \cdot I$, for some element $\varphi \in \Gamma_0(n)$. In particular, $\Gamma_0(n) \cdot I$ is the collection of all edges in $\mathcal{F}$ with one endpoint of the form $\frac{a}{nc}$.

Viewing this information through the lens of infinite loops, we see that if ζ_α is an infinite loop mod n, then ζ_α can not intersect any edge in $\mathcal{F}$ which has an endpoint with denominator divisible by n (except for the point at ∞). However, as seen in Proposition 3.1.23, the semi-convergents of α are exactly the endpoints of the edges in $\mathcal{F}$ which ζ_α intersects. This leads to an equivalent definition of an infinite loop mod n (as a real number).

Definition 5.1.1 (b). An *infinite loop* mod n is any real number $\alpha \in \mathbb{R}_{>0}$ with no semi-convergent denominators which are by divisible n (other than $q_{-1} = 0$).

Remark 5.1.7. Here, we should note that if $\alpha \in \mathbb{Q}$, we will assume that the continued fraction expansion $\overline{\alpha}$ ends in a partial quotient of size ∞. The real number α still produces two separate continued fraction expansions of the form $[a_0; a_1, \ldots, a_m + 1, \infty]$ and $[a_0; a_1, \ldots, a_m, 1, \infty]$. The reason why we do this is because we may have rational numbers which are the endpoint of some edge in $\Gamma_0(n) \cdot I$, but do not have a semi-convergent denominator divisible by n, unless we include the final partial quotient of size ∞. Note that since:

$$\lim_{k\to\infty} a_0 + \cfrac{1}{a_1 + \cfrac{1}{\ldots + \cfrac{1}{a_m + 1 + \cfrac{1}{k}}}} = a_0 + \cfrac{1}{a_1 + \cfrac{1}{\ldots + \cfrac{1}{a_m + 1}}},$$

the continued fraction expansion $[a_0; a_1, \ldots, a_m + 1, \infty]$ and $[a_0; a_1, \ldots, a_m + 1]$ are equivalent.

Properties of Infinite Loops

We begin by listing the properties of infinite loops mod n for future reference and then discuss each property individually providing a proof of each statement. For some of these properties we will include further discussion and results.

Properties of Infinite Loops mod n

1. *If α is an infinite loop mod n, then α is an infinite loop mod kn, where $k \in \mathbb{N}$.*

2. *If ζ_α is not an infinite loop mod n, then there is some map $\varphi \in \Gamma_0(n)$ such that ζ_α intersects $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ in the same way that $\varphi^{-1} \cdot \zeta_\alpha$ intersects both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$.*

3. *If α is not an infinite loop mod n, then there is a tail $\overline{\beta}$ of $\overline{\alpha}$ such that $\overline{n\beta}$ is a tail of $\overline{n\alpha}$.*

4. *If ζ_α is not an infinite loop mod n, then either:*

   (a) *It can be decomposed into infinitely many sub-paths such that each sub-path behaves nicely when multiplied by n, or*

   (b) *It can be decomposed into a finite number of sub-paths (each of which behaves nicely when multiplied by n), followed by a sub-path which is equivalent to an infinite loop mod n.*

5. *Let $\frac{a}{cn}$ and $\frac{b}{d}$ be two points lying in the interval $[0,1]$ which satisfy $|ad - bcn| = 1$. Then, for all $\alpha \in \mathbb{R}_{>0}$ satisfying $\min\left\{\frac{a}{cn}, \frac{b}{d}\right\} \leq \alpha \leq \max\left\{\frac{a}{cn}, \frac{b}{d}\right\}$, the corresponding geodesic ray ζ_α is not an infinite loop mod n.*

6. *If $n \in \mathbb{N}$ and $n \geq 4$, then there exist infinite loops mod n.*

Property 1. *If α is an infinite loop mod n, then α is an infinite loop mod kn, where $k \in \mathbb{N}$.*

Proof. This property comes from the definition of infinite loops mod n. Since α is an infinite loop mod n, it has no semi-convergent denominators which are divisible by n. By extension, α has no semi-convergent denominators divisible by kn, where $k \in \mathbb{N}$. □

Property 2. *If ζ_α is not an infinite loop mod n, then there is some map $\varphi \in \Gamma_0(n)$ such that ζ_α intersects $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ in the same way that $\varphi^{-1} \cdot \zeta_\alpha$ intersects both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$.*

Proof. If we assume that ζ_α is not an infinite loop mod n, then we can conclude that there is some non-trivial element $\varphi \in \Gamma_0(n)$ such that ζ_α intersects $\varphi \cdot I$. We can then perform the inverse map $\varphi^{-1} \in \Gamma_0(n)$ on our space $\mathbb{H}$, and since $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ are both invariant under the action of $\Gamma_0(n)$, the map φ^{-1} preserves the structure

Points of the form	Number of neighbours in $\mathcal{F}$ between consecutive neighbours in $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$	Number of neighbours in $\frac{1}{n}\mathcal{F}$ between consecutive neighbours in $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$
$\frac{a}{nc}$	0	$n-1$
$\frac{b}{d}, \quad gcd(n,d) = 1$	$n-1$	0

Table 5.1: A table of the number of neighbours in $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ that points of the form $\frac{a}{nc}$ and $\frac{b}{d}$ (with $gcd(n,d) = 1$) have between consecutive neighbours in $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$

of both $\mathcal{F}$ or $\frac{1}{n}\mathcal{F}$. In particular, the way that ζ_α intersects $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$ is equivalent to the way that $\varphi^{-1} \cdot \zeta_\alpha$ intersects $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$. □

This turns out to be a very powerful tool when looking at how integer multiplication affects cutting sequences. Using a slightly stronger version of this condition, we can gain even more information. For any point of the form $A = \frac{a}{nc}$ and any two consecutive neighbours $B_1 = \frac{b_1}{d_1}$ and $B_2 = \frac{b_2}{d_2}$ of A in $\Gamma_0(n) \cdot I$, we can always find a map $\varphi \in \Gamma_0(n)$ such that $\varphi \cdot 0 = A$, $\varphi \cdot \infty = B_1$ and $\varphi \cdot \frac{1}{n} = B_2$ (up to relabelling). Since $\Gamma_0(n)$ preserves both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$, the number of neighbours that A has between B_1 and B_2 in $\mathcal{F}$ (or equivalently in $\frac{1}{n}\mathcal{F}$) will be equivalent to the number of neighbours that 0 has between ∞ and $\frac{1}{n}$ in $\mathcal{F}$ (or in $\frac{1}{n}\mathcal{F}$). The point 0 has $n-1$ neighbours in $\mathcal{F}$ between ∞ and $\frac{1}{n}$, and each neighbour is of the form $\frac{1}{i}$ with $i \in \{1, 2, \ldots, n-1\}$. In $\frac{1}{n}\mathcal{F}$, 0 has no neighbours between ∞ and $\frac{1}{n}$. Similarly, for any point of the form $B = \frac{b}{d}$ ($gcd(n,d) = 1$) and any two consecutive neighbours $A_1 = \frac{a_1}{nc_1}$ and $A_2 = \frac{a_2}{nc_2}$ of B in $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$, the number of neighbours that B has between A_1 and A_2 in $\mathcal{F}$ (or in $\frac{1}{n}\mathcal{F}$) will be equivalent to the number of neighbours that ∞ has between 0 and 1 in $\mathcal{F}$ (or in $\frac{1}{n}\mathcal{F}$). The point ∞ has no neighbours in $\mathcal{F}$ between 0 and 1 and $n-1$ neighbours in $\frac{1}{n}\mathcal{F}$ between 0 and 1, and each neighbour is of the form $\frac{i}{n}$ with $i \in \{1, 2, \ldots, n-1\}$. We summarise these results in Table 5.1.

This information is used to prove the following result.

Proposition 5.1.8. *If a continued fraction $\overline{\alpha}$ has a convergent denominator q_k, such that $q_k = nq_k^{'}$, for some $n \in \mathbb{N}$ and some $q_k^{'} \in \mathbb{N}_{>1}$, then $B(n\alpha) \geq n$. Furthermore, if $\frac{p_k}{q_k} = \frac{p_k}{nq_k^{'}}$ is a convergent of $\overline{\alpha}$, $\frac{p_k}{q_k^{'}}$ is a convergent of $\overline{n\alpha}$.*

Proof. Let $A = \frac{p_k}{q_k}$ be a convergent of $\overline{\alpha}$ with geodesic representative ζ_α in $\mathbb{H}$ such that $q_k = nq_k^{'}$, for some $n \in \mathbb{N}$ and some $q_k^{'} \in \mathbb{N}_{>1}$. Then A is a common vertex of a fan in the cutting sequence of ζ_α with $\mathcal{F}$. By Proposition 3.1.23, at least two edges of the cutting sequence have A as an endpoint. Let $B = \frac{r}{s}$ and $C = \frac{t}{u}$ be the other two endpoints of two such edges with $\triangle ABC \in \mathcal{F}$. Since A is a neighbour of both B and C in $\mathcal{F}$, $gcd(q_k, s) = gcd(n, s) = gcd(q_k, u) = gcd(n, u) = 1$. From Lemma 5.1.3, the edges AB and AC are in $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$. By the above paragraph, there exists a map in $\Gamma_0(n)$ which takes ∞ to A, 0 to B and 1 to C (up to relabelling B and C). Since AB and AC are both in $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$, there are $n-1$ edges between AB and AC in $\frac{1}{n}\mathcal{F}$ - as mentioned in Table 5.1 - all of which the geodesic ray ζ_α passes through. All these edges share A as an endpoint and so, they are all edges in the same fan. It follows from this, that the fan ζ_α forms with $\frac{1}{n}\mathcal{F}$ containing both AB and AC, contains at least n triangles. Therefore, the corresponding continued fraction expansion $(\zeta_\alpha, \frac{1}{n}\mathcal{F})$ contains a partial quotient with value at least n.

By the above argument $A = \frac{p_k}{q_k}$ is a common vertex of a fan in the cutting sequence $(\zeta_\alpha, \frac{1}{n}\mathcal{F})$. When we rescale using the n^* map, the vertex $\frac{p_k}{q_k}$ in $\frac{1}{n}\mathcal{F}$ maps to $\frac{p_k}{q_k^{'}}$ in $\mathcal{F}$, which is a fixed vertex of some fan in the cutting sequence $(n^*(\zeta_\alpha), \mathcal{F})$. Therefore, $\frac{p_k}{q_k^{'}}$ is a convergent of $\overline{n\alpha} = [a_0^{(n)}; a_1^{(n)}, \ldots]$. Since $q_k^{'} > 1$, it follows that $\frac{p_k}{q_k^{'}}$ is not the common vertex of the first fan, but necessarily of some fan after. As a result, the partial quotient of $\overline{n\alpha}$ with value at least n is not the first partial quotient. By definition $B(n\alpha) := \max\left\{a_i^{(n)} : \overline{n\alpha} = [a_0^{(n)}; a_1^{(n)}, \ldots]\right\}$ and so $B(n\alpha) \geq a_i^{(n)}$ for all $i \in \mathbb{N}$. Since there exists an $a_i^{(n)} \geq n$ for $i \in \mathbb{N}$, it follows that $B(n\alpha) \geq n$, as required. □

We can improve on this result by taking all such triangles in the fan with the common vertex $\frac{p_k}{q_k}$ with $n \mid q_k$. Since this point can be mapped to ∞ by a map which preserves $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$, each of the triangles in the fan is effectively subdivided each by n when replacing $\mathcal{F}$ with $\frac{1}{n}\mathcal{F}$. There are a_{k+1} such triangles, where a_{k+1} is the $k+1$-th partial quotient. Note that there may be extra terms in this fan (added either side) but, nevertheless, we can still guarantee that this fan contains at least na_{k+1} triangles. In particular, we get the following corollary:

Corollary 5.1.9. *If a continued fraction $\overline{\alpha}$ has a convergent denominator q_k, such that $n \mid q_k$ for $n \in \mathbb{N}$ and $n < q_k$, then $B(n\alpha) \geq na_{k+1}$.*

Property 3. *If α is not an infinite loop mod n, then there is a tail $\overline{\beta}$ of $\overline{\alpha}$ such that $\overline{n\beta}$ is a tail of $\overline{n\alpha}$. If α is an infinite loop mod n, then there is no tail $\overline{\beta}$ of $\overline{\alpha}$ such that $\overline{n\beta}$ is a tail of $\overline{n\alpha}$.*

Proof. If ζ_α is not an infinite loop mod n, then we can decompose ζ_α into a finite path ρ_α, which runs from I to $\varphi \cdot I$ for some $\varphi \in \Gamma_0(n)$, followed by a infinite path ξ_α starting at $\varphi \cdot I$. Since ξ_α starts at $\varphi \cdot I$, we can perform φ^{-1} on $\mathbb{H}$ to find a geodesic ray $\zeta_\beta = \varphi^{-1} \cdot \xi_\alpha$. Since ξ_α started at $\varphi \cdot I$, the geodesic ray ζ_β starts at the y-axis I and terminates at the point $\beta = \varphi^{-1} \cdot \alpha$. The map φ^{-1} preserves both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$, since it is an element of $\Gamma_0(n)$. We can therefore see that $(\xi_\alpha, \mathcal{F}) = (\varphi^{-1} \cdot \xi_\alpha, \varphi^{-1} \cdot \mathcal{F}) = (\zeta_\beta, \mathcal{F})$ and $(\xi_\alpha, \frac{1}{n}\mathcal{F}) = (\varphi^{-1} \cdot \xi_\alpha, \varphi^{-1} \cdot \frac{1}{n}\mathcal{F}) = (\zeta_\beta, \frac{1}{n}\mathcal{F})$. Since $(\xi_\alpha, \mathcal{F})$ is a tail of $(\zeta_\alpha, \mathcal{F})$ and $(\xi_\alpha, \frac{1}{n}\mathcal{F})$ is a tail of $(\zeta_\alpha, \frac{1}{n}\mathcal{F})$, the result follows. □

Property 4. *If ζ_α is not an infinite loop mod n, then either:*

1. *It can be decomposed into infinitely many sub-paths such that each sub-path behaves nicely when multiplied by n, or*

2. *It can be decomposed into a finite number of sub-paths (each of which behaves nicely when multiplied by n), followed by a sub-path which is equivalent to an infinite loop mod n.*

Proof. If ζ_α intersects infinitely many lines in $\Gamma_0(n)\cdot I$, then we decompose ζ_α into an infinite collection of paths $\{\lambda_i\}_{i\in\mathbb{N}}$ such that each λ_i runs from one edge in $\Gamma_0(n)\cdot I$ to another. Since $\Gamma_0(n)\cdot I \subset \mathcal{F}\cap\frac{1}{n}\mathcal{F}$, we can conclude that both cutting sequences $(\lambda_i,\mathcal{F})$ and $(\lambda_i,\frac{1}{n}\mathcal{F})$ are well-defined for every $i\in\mathbb{N}$.

If ζ_α only intersects $\Gamma_0(n)\cdot I$ finitely many times, then we can take $\psi\cdot I$ to be the last time that ζ_α intersects $\Gamma_0(n)\cdot I$. If we take $\beta=\psi^{-1}\cdot\alpha\in\mathbb{R}_{>0}$, then the corresponding geodesic ray ζ_β starting at I and terminating at β is an infinite loop mod n. Since ψ preserves both $\mathcal{F}$ and $\frac{1}{n}\mathcal{F}$, not only is ζ_α equivalent to a unique infinite loop ζ_β mod n, but the multiplication of ζ_α is determined by the multiplication of ζ_β by Property 3. □

Property 5. *Let $\frac{a}{cn}$ and $\frac{b}{d}$ by two points lying in the interval $[0,1]$ which satisfy $|ad-bcn|=1$. Then, for all $\alpha\in\mathbb{R}_{>0}$ satisfying $\min\left\{\frac{a}{cn},\frac{b}{d}\right\}\le\alpha\le\max\left\{\frac{a}{cn},\frac{b}{d}\right\}$, the geodesic ray ζ_α is not an infinite loop mod n.*

Proof. Here, the edge between $\frac{a}{cn}$ and $\frac{b}{d}$ lies in $\Gamma_0(n)\cdot I$ by Corollary 5.1.5. Furthermore, this edge separates $\mathbb{H}$ into two regions: one containing I, and the other containing α. Since the geodesic ray ζ_α runs from I to α, ζ_α must necessarily intersect the edge between $\frac{a}{cn}$ and $\frac{b}{d}$. Therefore, ζ_α can not be an infinite loop mod n. □

Property 6. *If $n\in\mathbb{N}$ and $n\ge 4$, then there exist infinite loops mod n.*

Proof. In order to prove this statement, it is equivalent to show that there is no finite set of edges in $\Gamma_0(n)\cdot I$ connecting 0 to 1. As seen in Property 5, if we have two points $\frac{a}{cn}$ and $\frac{b}{d}$ in the interval $[0,1]$ which satisfy $|ad-bcn|=1$, then for all $\alpha\in\mathbb{R}_{>0}$ satisfying $\min\left\{\frac{a}{cn},\frac{b}{d}\right\}\le\alpha\le\max\left\{\frac{a}{cn},\frac{b}{d}\right\}$, the geodesic ray ζ_α is not an infinite loop mod n. If we assume that $\frac{a}{cn}<\frac{b}{d}$ and assume that there is another point of the form $\frac{e}{nf}>\frac{b}{d}$ with $|ed-bnf|=1$, then we can further conclude that for all $\alpha\in\mathbb{R}_{>0}$ satisfying $\frac{a}{cn}\le\alpha\le\frac{e}{nf}$, the geodesic ray ζ_α is not an infinite loop mod n. If there is a finite set of edges connecting 0 to 1, then we can use Property 5 on each of these edges to see that there is no infinite loop mod n, for all $0 \le \alpha \le 1$. However, if no such finite path exists, then there must be a non-empty set of points in $[0,1]$ which do not lie between any neighbours in $\Gamma_0(n) \cdot I$. If α is one of these points, then the corresponding geodesic ray ζ_α does not intersect $\Gamma_0(n) \cdot I$. Therefore, α is an infinite loop mod n.

To find this set of edges, it is equivalent to find a finite sequence of rational points between 0 and 1 such that each consecutive pair of rational points are neighbours in $\Gamma_0(n) \cdot I$. This sequence of rational numbers will be of the form $\left\{\frac{0}{1} = \frac{b_0}{d_0}, \frac{a_1}{c_1 n}, \frac{b_1}{d_1}, \ldots, \frac{a_k}{c_k n}, \frac{b_k}{d_k} = \frac{1}{1}\right\}$, where $\frac{b_{i-1}}{d_{i-1}} < \frac{a_i}{c_i n} < \frac{b_i}{d_i}$, $a_i, b_i, c_i, d_i \in \mathbb{N}$ and $\gcd(n, d_i) = 1$. Given two points A and B and a sequence of rationals $\{A = A_0, A_1, A_2, \ldots, A_k = B\}$, we will say that this sequence is *a sequence of neighbours in* $\Gamma_0(n) \cdot I$ *connecting* A *and* B if $A_i < A_{i+1}$ and A_i and A_{i+1} are all neighbours in $\Gamma_0(n) \cdot I$ for all $i \in \{0, 1, \ldots, k-1\}$. Similarly, if we have two points A and B and a sequence of rationals $\{A = A_0, A_1, A_2, \ldots, A_k = B\}$, we will say that this sequence is *a sequence of neighbours in* $\mathcal{F}$ *connecting* A *and* B if $A_i < A_{i+1}$ and A_i and A_{i+1} are all neighbours in $\mathcal{F}$ for all $i \in \{0, 1, \ldots, k-1\}$.

Since $\Gamma_0(n) \cdot I$ is a sub-graph of $\mathcal{F} \cap \frac{1}{n}\mathcal{F}$, which is in turn a sub-graph of $\mathcal{F}$, each edge E in the finite set of edges in $\Gamma_0(n) \cdot I$ connecting 0 to 1, must also be an edge of $\mathcal{F}$. As a result, we will start with a sequence of neighbours in $\mathcal{F}$, and insert additional Farey neighbours to this sequence, until this sequence is also a sequence of neighbours in $\Gamma_0(n) \cdot I$. To show that this constructs a minimal sequence of neighbours in $\Gamma_0(n) \cdot I$ (should a minimal sequence exist), we will use the following claim:

Claim: Assume that $\frac{a}{c}, \frac{b}{d} \in \mathbb{Q} \cap [0,1]$ are neighbours in $\mathcal{F}$ with $\frac{a}{c} < \frac{b}{d}$. Then any sequence of neighbours in $\mathcal{F}$ of the form $\left\{\frac{a}{c} = \frac{a_0}{c_0}, \frac{a_1}{c_1}, \frac{a_2}{c_2}, \ldots, \frac{a_k}{c_k} = \frac{b}{d}\right\}$ satisfying $\frac{a_{i-1}}{c_{i-1}} < \frac{a_i}{c_i} < \frac{a_{i+1}}{c_{i+1}}$ must either:

1. Only contain the points $\left\{\frac{a}{c}, \frac{b}{d}\right\}$, or

2. Contain the point $\frac{a}{c} \oplus \frac{b}{d} = \frac{a+b}{c+d}$.

Proof of claim. Since $\frac{a}{c}$ and $\frac{b}{d}$ are neighbours in $\mathcal{F}$ we know that there is an edge E in $\mathcal{F}$ connecting these points. This edge separates the plane $\mathbb{H}$ into two regions: E_+, containing the interval $(\frac{a}{c}, \frac{b}{d})$, and E_-, containing the intervals $[-\infty, \frac{a}{c})$ and $(\frac{b}{d}, \infty]$. The sequence of neighbours $\left\{\frac{a}{c} = \frac{a_0}{c_0}, \frac{a_1}{c_1}, \frac{a_2}{c_2}, \dots, \frac{a_k}{c_k} = \frac{b}{d}\right\}$, must all lie in the interval $[\frac{a}{c}, \frac{b}{d}]$, since we assumed that $\frac{a_{i-1}}{c_{i-1}} < \frac{a_i}{c_i} < \frac{a_{i+1}}{c_{i+1}}$. In particular, the edges between each of these vertices must either be contained in E_+ or be the edge E itself, i.e. the sequence of neighbours in $\mathcal{F}$ is just $\left\{\frac{a}{c}, \frac{b}{d}\right\}$.

If this is not the case, then we can assume that the sequence of neighbours in $\mathcal{F}$, given by $\left\{\frac{a}{c} = \frac{a_0}{c_0}, \frac{a_1}{c_1}, \frac{a_2}{c_2}, \dots, \frac{a_k}{c_k} = \frac{b}{d}\right\}$, contains a vertex $\frac{a_j}{b_j}$ which is not $\frac{a}{c}, \frac{b}{d}$ or $\frac{a+c}{b+d}$. Since $\frac{a}{c} \oplus \frac{b}{d} = \frac{a+b}{c+d}$ is a neighbour of both $\frac{a}{c}$ and $\frac{b}{d}$ in $\mathcal{F}$, the vertices $\frac{a+b}{c+d}, \frac{a}{c}$ and $\frac{b}{d}$ form a triangle in $\mathcal{F}$. Furthermore, since $\frac{a_j}{c_j} \neq \frac{a+c}{b+d}$, the vertex $\frac{a_j}{c_j}$ can either lie in the interval $(\frac{a}{c}, \frac{a+b}{c+d})$ or $(\frac{a+b}{c+d}, \frac{b}{d})$. We assume that $\frac{a_j}{c_j}$ lies in the interval $(\frac{a}{c}, \frac{a+b}{c+d})$ - a similar argument can be made if $\frac{a_j}{c_j}$ lies in the interval $(\frac{a+b}{c+d}, \frac{b}{d})$. Then, we take E' to be the edge between $\frac{a}{c}$ and $\frac{b}{d}$, and assume E'_+ is the region containing the interval $(\frac{a}{c}, \frac{a+b}{c+d})$. By assumption, the vertex $\frac{a_j}{c_j}$ is contained in the region E'_+. In the sequence $\left\{\frac{a}{c} = \frac{a_0}{c_0}, \frac{a_1}{c_1}, \frac{a_2}{c_2}, \dots, \frac{a_k}{c_k} = \frac{b}{d}\right\}$, there must be a subsequence of neighbours in $\mathcal{F}$ given by $\left\{\frac{a_j}{c_j}, \frac{a_{j+1}}{c_{j+1}}, \dots, \frac{a_k}{c_k} = \frac{b}{d}\right\}$ which connects $\frac{a_j}{c_j}$ to $\frac{b}{d}$. However, $\frac{a_j}{c_j}$ lies in E'_+ and $\frac{b}{d}$ lies in E'_-. As a result, the corresponding sequence of edges in $\mathcal{F}$ must either contain the point $\frac{a+c}{b+d}$ or non-trivially intersect the edge E'. However, since E' and the sequence of edges connecting $\frac{a_j}{c_j}$ to $\frac{b}{d}$ are all edges in the Farey tessellation, none of these edges can non-trivially intersect. Therefore, the subsequence of edges must pass through the point $\frac{a+c}{b+d}$ and so, the sequence of neighbours $\left\{\frac{a_j}{c_j}, \frac{a_{j+1}}{c_{j+1}}, \dots, \frac{a_k}{c_k} = \frac{b}{d}\right\}$ must contain the point $\frac{a+c}{b+d}$. Finally, since this subsequence contains the point $\frac{a+b}{c+d}$, so must our original sequence of neighbours $\left\{\frac{a}{c} = \frac{a_0}{c_0}, \frac{a_1}{c_1}, \frac{a_2}{c_2}, \dots, \frac{a_k}{c_k} = \frac{b}{d}\right\}$. *QED.*

Given two Farey neighbours $\frac{a}{c}$ and $\frac{b}{d}$ with $\frac{a}{c} < \frac{b}{d}$, we can use this claim to construct a minimal sequence of neighbours in $\Gamma_0(n) \cdot I$ between these points. We denote this minimal sequence $\widetilde{V}$. Firstly, we take the sequence of neighbours in $\mathcal{F}$ given by $V_0 := \left\{\frac{a}{c}, \frac{b}{d}\right\}$ to be our initial sequence. If $\frac{a}{c}$ and $\frac{b}{d}$ are neighbours in $\Gamma_0(n) \cdot I$, then we will take $\widetilde{V} = V_0$, and we are done. Otherwise, by the above claim, the set $\widetilde{V}$

must include the point $\frac{a+b}{c+d}$. We know that $\frac{a+b}{c+d}$ is a Farey neighbour of both $\frac{a}{c}$ and $\frac{b}{d}$ and $\frac{a}{c} < \frac{a+b}{c+d} < \frac{b}{d}$. As a result, we can replace our initial sequence of neighbours $V_0 = \left\{\frac{a}{c}, \frac{b}{d}\right\}$ with the sequence of neighbours $V_1 := \left\{\frac{a}{c}, \frac{a+b}{c+d}, \frac{b}{d}\right\}$. Since each consecutive pair of vertices in V_1 are neighbours in $\mathcal{F}$, we can consider each pair of vertices in the set V_1 individually and apply the same process on each of these pairs. For example, if $\frac{a}{c}$ and $\frac{a+b}{c+d}$ are neighbours in $\Gamma_0(n) \cdot I$, then we do not need to construct any more vertices between them. However, if they are not neighbours in $\Gamma_0(n) \cdot I$, then our sequence of neighbours in $\Gamma_0(n) \cdot I$ must include their Farey neighbour $\frac{2a+b}{2c+d}$. As a result, we can replace the subsequence $\left\{\frac{a}{c}, \frac{a+b}{c+d}\right\}$ by the subsequence $\left\{\frac{a}{c}, \frac{2a+b}{2c+d}, \frac{a+b}{c+d}\right\}$. We can then apply the same procedure on the subsequence $\left\{\frac{a+b}{c+d}, \frac{b}{d}\right\}$ to form our next iterated set of neighbours in $\mathcal{F}$, which we denote V_2. We can then perform this procedure on each pair of vertices in V_2 to form a new set V_3, and then perform this procedure on the set V_3, and so on. Since we only add in additional neighbours between two points A and B when A and B are not neighbours in $\Gamma_0(n) \cdot I$, this process will form a minimal sequence of neighbours in $\Gamma_0(n) \cdot I$ between the points A and B - provided such a sequence of vertices exist. Starting with the initial set of vertices $V_0 = \{0, 1\}$, the process can be described algorithmically as follows:

1. Start with the set of vertices $V_0 = \left\{\frac{0}{1}, \frac{1}{1}\right\}$.
2. While V_i is not of the required form, repeat the following process:
   (a) Take $V_{i+1} = \left\{\frac{0}{1}\right\}$.
   (b) For each pair of vertices v_i and v_{i+1} in V_i:
   If v_i and v_{i+1} are neighbours in $\Gamma_0(n) \cdot I$:
      - Append v_{i+1} into V_{i+1}

   Otherwise:
      - Append $w = v_i \oplus v_{i+1}$ into V_{i+1}.
      - Append v_{i+1} into V_{i+1}.
3. End of algorithm.

If we take $n = 2$ then we have:

$$V_0 = \left\{\frac{0}{1}, \frac{1}{1}\right\},$$
$$V_1 = \left\{\frac{0}{1}, \frac{1}{2}, \frac{1}{1}\right\}.$$

At which point the process stops.

If we instead take $n = 3$ then we have:

$$V_0 = \left\{\frac{0}{1}, \frac{1}{1}\right\},$$
$$V_1 = \left\{\frac{0}{1}, \frac{1}{2}, \frac{1}{1}\right\},$$
$$V_2 = \left\{\frac{0}{1}, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, \frac{1}{1}\right\}.$$

Again, the process stops at this point.

However, for $n = 5$, we have:

$$V_0 = \left\{\frac{0}{1}, \frac{1}{1}\right\},$$
$$V_1 = \left\{\frac{0}{1}, \frac{1}{2}, \frac{1}{1}\right\},$$
$$V_2 = \left\{\frac{0}{1}, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, \frac{1}{1}\right\},$$
$$V_3 = \left\{\frac{0}{1}, \frac{1}{4}, \frac{1}{3}, \frac{2}{5}, \frac{1}{2}, \frac{3}{5}, \frac{2}{3}, \frac{3}{4}, \frac{1}{1}\right\},$$
$$V_4 = \left\{\frac{0}{1}, \frac{1}{5}, \frac{1}{4}, \frac{2}{7}, \frac{1}{3}, \frac{2}{5}, \frac{1}{2}, \frac{3}{5}, \frac{2}{3}, \frac{5}{7}, \frac{3}{4}, \frac{4}{5}, \frac{1}{1}\right\},$$
$$\dots$$

Given two points $\frac{a}{c}$ and $\frac{b}{d}$, which are neighbours in $\mathcal{F}$, we can see from Corollary 5.1.5 that $\frac{a}{c}$ and $\frac{b}{d}$ are neighbours in $\Gamma_0(n) \cdot I$ if and only if exactly one of $c \equiv 0 \mod n$ or $d \equiv 0 \mod n$. Here, we should note that we can not have that both $c \equiv 0 \mod n$ and $d \equiv 0 \mod n$, since we know that $\gcd(c, d) = 1$. In particular, assuming the points $\frac{a}{c}$ and $\frac{b}{d}$ are Farey neighbours, we only need to know the value of c and d mod n to be able to tell if they are neighbours in $\Gamma_0(n) \cdot I$. As a result, for us to construct a finite sequence of neighbours in $\Gamma_0(n) \cdot I$ of the form $\left\{\frac{0}{1} = \frac{b_0}{d_0}, \frac{a_1}{c_1 n}, \frac{b_1}{d_1}, \dots, \frac{a_k}{c_k n}, \frac{b_k}{d_k} = \frac{1}{1}\right\}$

it is a necessary condition that the sequence of denominators (taken $\mod n$) is of the form $\left\{\overline{d_0} = 1, 0, \overline{d_1}, 0, \ldots, 0, \overline{d_k} = 1\right\}$ where each $\overline{d_i} \in \{1, \ldots, n-1\}$. As a result, if we wish to show that the sequence of neighbours in $\Gamma_0(n) \cdot I$ of the form $\left\{\frac{0}{1} = \frac{b_0}{d_0}, \frac{a_1}{c_1 n}, \frac{b_1}{d_1}, \ldots, \frac{a_k}{c_k n}, \frac{b_k}{d_k} = \frac{1}{1}\right\}$ does not exist, then it is sufficient to show that the corresponding sequence $\left\{\overline{d_0} = 1, 0, \overline{d_1}, 0, \ldots, 0, \overline{d_k} = 1\right\}$ does not exist.

If we start with two points $\frac{a}{c}$ and $\frac{b}{d}$ which are Farey neighbours (with $\frac{a}{c}, \frac{b}{d}$), we can replace the sequence $V_0 := \left\{\frac{a}{c}, \frac{b}{d}\right\}$ with the sequence $D_0 := \left\{\overline{c}, \overline{d}\right\}$, where $\overline{c} \equiv c \mod n$, $\overline{d} \equiv d \mod n$ and $\overline{c}, \overline{d} \in \{0, 1, \ldots, n-1\}$. If one of $\overline{c} = 0$ or $\overline{d} = 0$, then we are done. Otherwise, $\frac{a}{c}$ and $\frac{b}{d}$ are not neighbours in $\Gamma_0(n) \cdot I$. In this case, we would replace the sequence $V_0 := \left\{\frac{a}{c}, \frac{b}{d}\right\}$ with the sequence $V_1 := \left\{\frac{a}{c}, \frac{a+b}{c+d}, \frac{b}{d}\right\}$, and so we analogously replace the sequence $D_0 := \left\{\overline{c}, \overline{d}\right\}$ with the sequence $D_1 := \left\{\overline{c}, \overline{c+d}, \overline{d}\right\}$, where $\overline{c+d} \equiv c+d \mod n$ and $\overline{c+d} \in \{0, 1, \ldots, n-1\}$. If $\overline{c+d} = 0$, then we are done. Otherwise, we can consider each consecutive pair in D_1 and perform the same procedure on each pair, i.e. we perform the same procedure on $\left\{\overline{c}, \overline{c+d}\right\}$ and $\left\{\overline{c+d}, \overline{d}\right\}$. Iterating this procedure, we can form a new algorithm to find a sequence of denominators of the required form $\left\{d_0 = \overline{c}, 0, d_1, 0, \ldots, 0, d_k = \overline{d}\right\}$, where each $d_i \in \{1, \ldots, n-1\}$. For our initial set being $D_0 := \{1, 1\}$ (corresponding to the set $V_0 := \left\{\frac{0}{1}, \frac{1}{1}\right\}$), the above procedure is described by the following algorithm.

1. Start with the set of denominators $D_0 = \{1, 1\}$.
2. While D_i is not of the required form, repeat the following process:
   (a) Take $D_{i+1} = \{1\}$.
   (b) For each pair of denominators d_i and d_{i+1} in D_i:
   If d_i and d_{i+1} are neighbours in $\Gamma_0(n) \cdot I$:
      - Append d_{i+1} into D_{i+1}

   Otherwise:
      - Append $e = d_i + d_{i+1} \mod n$ into D_{i+1}.
      - Append d_{i+1} into D_{i+1}.

3. End of algorithm.

For example, for $n = 5$ we would have:

$$D_0 = \{1, 1\},$$

$$D_1 = \{1, 2, 1\},$$

$$D_2 = \{1, 3, 2, 3, 1\},$$

$$D_3 = \{1, 4, 3, 0, 2, 0, 3, 4, 1\},$$

$$D_4 = \{1, 0, 4, 2, 3, 0, 2, 0, 3, 2, 4, 0, 1\},$$

$$D_5 = \{1, 0, 4, 1, 2, 0, 3, 0, 2, 0, 3, 0, 2, 1, 4, 0, 1\},$$

$$\ldots$$

For $n > 2$, we can always guarantee that the above process does not terminate after the first iteration, and so, the above process creates the set $D_1 = \{1, 2, 1\}$. Furthermore, for an arbitrary $n > 3$, we can perform iterative Farey sums between the sub-sequence $\{1, 2\}$ to obtain the sequence $\{1, 0, n-1, n-2, \ldots, 2\}$, and this sequence does not simply reduce to $\{1, 0, 2\}$, since $n - 1 \neq 2 \mod n$ for $n > 3$. If we perform the same process on the sub-sequence $\{n-1, n-2\} \mod n$, we obtain the sequence $\{n-1, 0, 1, 2, \ldots, n-3, n-2\}$. Combining together these sequences, we see that iteratively performing the procedure on sub-sequence $\{1, 2\}$ to produces the sequence $\{1, 0, n-1, 0, 1, 2, \ldots, n-3, n-2, \ldots, 2\}$. However, the sequence $\{1, 0, n-1, 0, 1, 2, \ldots, n-3, n-2, \ldots, 2\}$ contains the sub-sequence $\{1, 2\}$. This in turn implies that for $n > 3$ we can not resolve any sub-sequence of the form $\{1, 2\}$, since any attempt to do so produces another sub-sequence of the form $\{1, 2\}$. As a result, for $n > 3$ we can not find a finite sequence of denominators $\left\{\overline{d_0} = 1, 0, \overline{d_1}, 0, \ldots, 0, \overline{d_k} = 1\right\}$ corresponding to the finite sequence of neighbours in $\Gamma_0(n) \cdot I$ of the form $\left\{\frac{0}{1} = \frac{b_0}{d_0}, \frac{a_1}{c_1 n}, \frac{b_1}{d_1}, \ldots, \frac{a_k}{c_k n}, \frac{b_k}{d_k} = \frac{1}{1}\right\}$. In particular, no such sequence of neighbours in $\Gamma_0(n) \cdot I$ can exist, for $n > 3$. Finally, this implies that there are infinite loops mod n for all $n > 3$. $\square$

5.1.3 Infinite Loops and the p-adic Littlewood Conjecture

We start this section by restating the p-adic Littlewood Conjecture, as well as its reformulation:

The p-adic Littlewood Conjecture. *For every real number $\alpha \in \mathbb{R}$, we have:*

$$m_p(\alpha) := \inf_{q\in\mathbb{N}} \left\{q \cdot |q|_p \cdot \|q\alpha\|\right\} = 0.$$

Corollary 2.2.14. *Let $\alpha \in Bad$. Then α satisfies pLC if and only if:*

$$\sup_{l\in\mathbb{N}\cup\{0\}} B(p^l\alpha) = \infty$$

As seen in the previous sections, infinite loops mod n behave badly when multiplied by n. In fact, infinite loops mod n behave even worse when $n = p^\ell$. Since, the p-adic Littlewood Conjecture is very closely related to the behaviour of the continued fractions expansions $\left\{\overline{p^m\alpha} : m \in \mathbb{N} \cup \{0\}\right\}$, it seems very natural that infinite loops mod p^ℓ may tell us something non-trivial about the p-adic Littlewood Conjecture. Our first confirmation of this fact, comes from the next lemma and its corollary.

This lemma can be viewed as a slightly weaker version of Proposition 5.1.8. Instead of assuming α has a convergent denominator divisible by n, we assume that α is an infinite loop mod n, i.e. it has a semi-convergent denominator divisible by n. This lemma roughly states that if α is not an infinite loop mod n, then $B(\alpha)$ and $B(n\alpha)$ can not both be simultaneously small relative to $\sqrt{n}$.

Lemma 5.1.10. *Assume that $\alpha \in [0,1]$ is not an infinite loop mod n. Then:*

$$\max\{B(\alpha), B(n\alpha)\} \geq \lfloor 2\sqrt{n}\rfloor - 1,$$

where $\lfloor\cdot\rfloor$ is the standard floor function.

Proof. Assume α is not an infinite loop mod n and let ζ_α be the associated geodesic ray in $\mathbb{H}$. Since α is not an infinite loop mod n, there is an element $\varphi \in \Gamma_0(n)$, where φ is not the identity, such that ζ_α intersects the edge $\varphi \cdot I$. We can apply the map φ^{-1} to the whole of $\mathbb{H}$ such that $\varphi^{-1} \cdot \zeta_\alpha$ intersecting $\mathcal{F}$ resembles Fig. 5.1 (a) - up to taking a mirror image in the y-axis. Taking a mirror image has no affect on this argument other than to swap the roles of left and right fans, and so we shall assume that we are oriented as in the figure.

We assume that the geodesic ray $\varphi^{-1} \cdot \zeta_\alpha$ approaches the y-axis I by a right fan of size $r_\alpha \in \mathbb{N} \cup \{0\}$ and leaves by a right fan of size $s_\alpha \in \mathbb{N} \cup \{0\}$. Here, we allow these fans to be of size 0, however, in this case we interpret this fan to be a left fan. In this case, $\varphi \cdot \zeta_\alpha$ either intersects I and $I - 1$ (when $s_\alpha = 0$), or it intersects I and $I + 1$ (when r_α. In either case, the point at infinity is a fixed point of this fan. This tells us that $\varphi \cdot \infty$ is a convergent of α. This point $\varphi \cdot \infty$ will be of the form $\frac{a}{nc}$ and so by Proposition 5.1.8, this case induces $B(n\alpha) \geq n$, in which case the result follows.

We therefore assume $r_\alpha, s_\alpha \geq 1$ and note that $\varphi^{-1} \cdot \zeta_\alpha$ approaches the y-axis from a value less than $-[0; r_\alpha, 1] = \frac{-1}{r_\alpha+1}$. Similarly, we can assume that $\varphi^{-1} \cdot \zeta_\alpha$ departs the y-axis and approaches a point greater than $[0; s_\alpha, 1] = \frac{1}{s_\alpha+1}$. Since $\frac{1}{n}\mathcal{F}$ has vertices between $\frac{i}{n}$ and ∞ for all $i \in \mathbb{N}$, we can ask how many of these lines the geodesic ray $\varphi^{-1} \cdot \zeta_\alpha$ intersects in this neighbourhood. We see that there is some number $R_\alpha \in \mathbb{N} \cup \{0\}$ such that $\frac{-R_\alpha-1}{n} \leq \frac{-1}{r_\alpha+1} \leq \frac{R_\alpha}{n}$. One can then guarantee that $\varphi^{-1} \cdot \zeta_\alpha$ intersects a left fan in $\frac{1}{n}\mathcal{F}$ of size at least R_α directly before approaching the y-axis. Note that here the value R_α can be defined as $R_\alpha := \left\lfloor \frac{n}{r_\alpha+1} \right\rfloor$, where $\lfloor \cdot \rfloor$ is the standard floor function. By a similar process we can see that $\varphi^{-1} \cdot \zeta_\alpha$ intersects a left fan in $\frac{1}{n}\mathcal{F}$ of size at least $S_\alpha := \left\lfloor \frac{n}{s_\alpha+1} \right\rfloor$ in $\frac{1}{n}\mathcal{F}$ directly after leaving the y-axis, see Fig. 5.1 (b). These fans concatenate to form a fan of size $R_\alpha + S_\alpha$ in $\frac{1}{n}\mathcal{F}$. Therefore, we know that $\overline{\alpha}$ has a term of size at least $r_\alpha + s_\alpha$ and $\overline{n\alpha}$ has a term of size at least $R_\alpha + S_\alpha$. We conclude that $B(\alpha) \geq r_\alpha + s_\alpha$ and $B(n\alpha) \geq R_\alpha + S_\alpha$, and by extension $\max\{B(\alpha), B(n\alpha)\} \geq \max\{r_\alpha + s_\alpha,\ R_\alpha + S_\alpha\}$.

We assume that $r_\alpha + s_\alpha \leq \lfloor 2\sqrt{n} \rfloor - 2$, since otherwise we would have $B(\alpha) \geq \lfloor 2\sqrt{n} \rfloor - 1$. If we fix $0 \leq r_\alpha \leq \lfloor 2\sqrt{n} \rfloor - 2$, then $0 \leq s_\alpha \leq \lfloor 2\sqrt{n} \rfloor - 2 - r_\alpha$. For all s_α in this range,

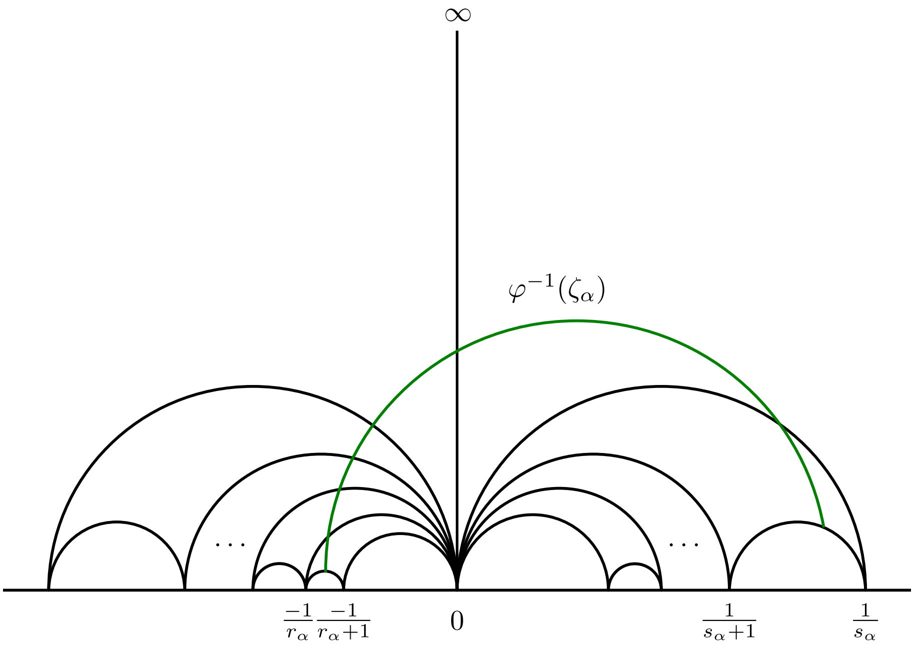

(a) An example of a geodesic $\varphi^{-1} \cdot \zeta_\alpha$ approaching I by a right fan of size r_α in $\mathcal{F}$ and leaving I via a fan of size s_α. This results in a fan of size $(r_\alpha + s_\alpha)$.

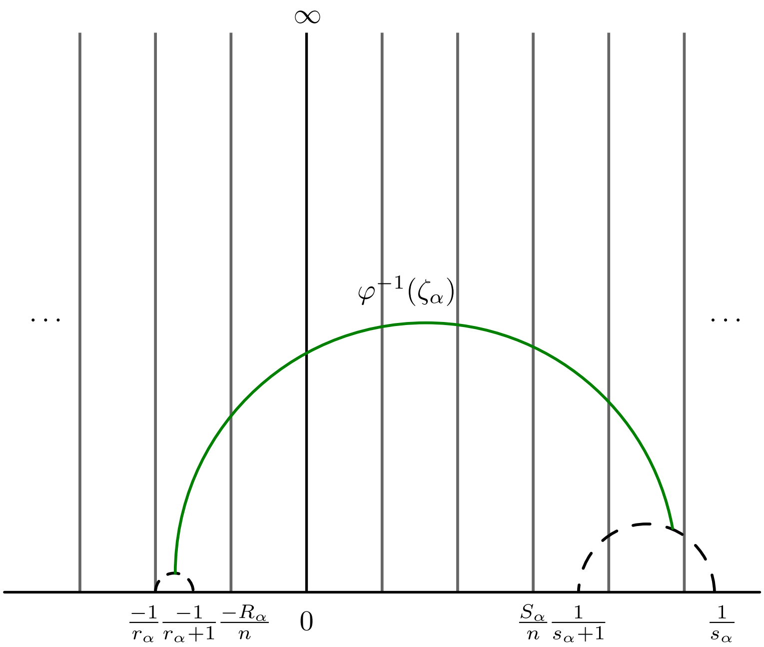

(b) An example of how the geodesic $\varphi^{-1} \cdot \zeta_\alpha$ intersects $\frac{1}{n}\mathcal{F}$. The lines between $\frac{-R_\alpha}{n}$ and $\frac{S_\alpha}{n}$ are necessarily intersected by $\varphi^{-1} \cdot \zeta_\alpha$ and this results in a fan of size $\geq (R_\alpha + S_\alpha)$.

Figure 5.1: An example of a how geodesic ray ζ_α, which is not an infinite loop mod n, intersects both $\mathcal{F}$, (a), and $\frac{1}{n}\mathcal{F}$, (b). This is considered up to re-framing by $\Gamma_0(n)$.

we note:

$$S_\alpha = \left\lfloor \frac{n}{s_\alpha + 1} \right\rfloor \geq \left\lfloor \frac{n}{\lfloor 2\sqrt{n} \rfloor - 2 - r_\alpha + 1} \right\rfloor = \left\lfloor \frac{n}{\lfloor 2\sqrt{n} \rfloor - r_\alpha - 1} \right\rfloor,$$

and:

$$\begin{aligned} R_\alpha + S_\alpha = \left\lfloor \frac{n}{r_\alpha + 1} \right\rfloor + \left\lfloor \frac{n}{s_\alpha + 1} \right\rfloor &\geq \left\lfloor \frac{n}{r_\alpha + 1} \right\rfloor + \left\lfloor \frac{n}{\lfloor 2\sqrt{n} \rfloor - r_\alpha - 1} \right\rfloor \\ &\geq \left\lfloor \frac{n}{r_\alpha + 1} + \frac{n}{\lfloor 2\sqrt{n} \rfloor - r_\alpha - 1} \right\rfloor - 1. \end{aligned}$$

We can find a lower bound estimation for this by considering the following equation:

$$\mathrm{f(x)} := \frac{n}{x+1} + \frac{n}{\lfloor 2\sqrt{n} \rfloor - x - 1}, \text{ for x} \in \left[0, \lfloor 2\sqrt{n} \rfloor - 2\right],$$

and noting that $\lfloor \mathrm{f(x)} \rfloor$ is minimised when f(x) is minimised. The derivative of f(x) is given by:

$$\mathrm{f}'(\mathrm{x}) = n\left(\frac{-1}{(x+1)^2} + \frac{1}{(\lfloor 2\sqrt{n} \rfloor - x - 1)^2} \right).$$

Note that we can write $\lfloor 2\sqrt{n} \rfloor = 2\lfloor \sqrt{n} \rfloor + \delta$ where $\delta = 0, 1$.

1. Assume $\delta = 0$:

In this case, $\mathrm{f}'(\mathrm{x}) = 0$ if and only if $x = \lfloor \sqrt{n} \rfloor - 1$, and so, $x = \lfloor \sqrt{n} \rfloor - 1$ must be either a minima or a maxima (since f(x) is symmetric in x). At $x = \lfloor \sqrt{n} \rfloor - 1$, we have:

$$\begin{aligned} \lfloor \mathrm{f(x)} \rfloor &= \left\lfloor \frac{n}{\lfloor \sqrt{n} \rfloor} + \frac{n}{\lfloor \sqrt{n} \rfloor} \right\rfloor - 1 \\ &= \left\lfloor \frac{2n}{\lfloor \sqrt{n} \rfloor} \right\rfloor - 1 \\ &\geq \left\lfloor \frac{2n}{\sqrt{n}} \right\rfloor - 1 \\ &= \lfloor 2\sqrt{n} \rfloor - 1 \end{aligned}$$

We note that $\lfloor \mathrm{f}(0) \rfloor \geq n$, which is greater than or equal to $2\lfloor \sqrt{n} \rfloor - 1$ for all $n \in \mathbb{N}$. Therefore, $\lfloor \mathrm{f}(\lfloor \sqrt{n} \rfloor - 1) \rfloor$ is a minima.

2. Assume $\delta = 1$:

In this case, $f'(x) = 0$ if and only if $x = \lfloor\sqrt{n}\rfloor - \frac{1}{2}$. At $x = \lfloor\sqrt{n}\rfloor - \frac{1}{2}$, we have:

$$\begin{aligned}\lfloor f(x)\rfloor &= \left\lfloor \frac{n}{\lfloor\sqrt{n}\rfloor + \frac{1}{2}} + \frac{n}{\lfloor\sqrt{n}\rfloor + \frac{1}{2}} \right\rfloor - 1\\ &= \left\lfloor \frac{4n}{2\lfloor\sqrt{n}\rfloor + 1} \right\rfloor - 1\\ &= \left\lfloor \frac{4n}{\lfloor 2\sqrt{n}\rfloor} \right\rfloor - 1\\ &\geq \left\lfloor \frac{4n}{2\sqrt{n}} \right\rfloor - 1\\ &= \lfloor 2\sqrt{n}\rfloor - 1\end{aligned}$$

We note that $\lfloor f(0)\rfloor \geq n$, which is greater than or equal to $2\lfloor\sqrt{n}\rfloor - 1$ for all $n \in \mathbb{N}$. Therefore, $\lfloor f(\lfloor\sqrt{n}\rfloor - \frac{1}{2})\rfloor$ is a minima.

Therefore, for all $r_\alpha + s_\alpha \leq \lfloor 2\sqrt{n}\rfloor - 2$, we have that $R_\alpha + S_\alpha \geq \lfloor 2\sqrt{n}\rfloor - 1$, and so $\max\{r_\alpha + s_\alpha, R_\alpha + S_\alpha\} \geq \lfloor 2\sqrt{n}\rfloor - 1$ for all possible r_α and s_α. □

The above lemma gives us a lower bound for $\max\{B(\alpha), B(n\alpha)\}$ if α is not an infinite loop mod n. We set $m_p(\alpha) := \inf\{q \cdot |q|_p \cdot \|q\alpha\|\}$ and note that pLC is true is equivalent to saying that $m_p(\alpha) = 0$ for every $\alpha \in \mathbb{R}_{>0}$. Combining together Lemma 2.2.16 and Lemma 2.2.17 from Chapter 2, we see that :

$$\inf_{\ell\in\mathbb{N}\cup\{0\}} \left\{\frac{1}{B(p^\ell\alpha) + 2}\right\} \leq m_p(\alpha) \leq \inf_{\ell\in\mathbb{N}\cup\{0\}} \left\{\frac{1}{B(p^\ell\alpha)}\right\}.$$

This leads to the following corollary:

Corollary 5.1.11. *If $\alpha \in \mathbb{R}_{>0}$ is not an infinite loop mod p^m, then:*

$$m_p(\alpha) \leq \frac{1}{\lfloor 2\sqrt{p^m}\rfloor - 1}.$$

Proof. Since we know that:

$$m_p(\alpha) \leq \inf_{\ell\in\mathbb{N}\cup\{0\}} \left\{\frac{1}{B(p^\ell\alpha)}\right\},$$

we can conclude that

$$m_p(\alpha) \leq \frac{1}{B(p^j\alpha)},$$

for any $j \in \mathbb{N} \cup \{0\}$. Since α is not an infinite loop mod p^m, we know by the previous lemma, that:

$$\max\left\{B(\alpha), B(p^m\alpha)\right\} \geq \left\lfloor 2\sqrt{p^m}\right\rfloor - 1.$$

Combining this information together, we see that:

$$m_p(\alpha) \leq \min\left\{\frac{1}{B(\alpha)}, \frac{1}{B(p^m\alpha)}\right\} \leq \frac{1}{\left\lfloor 2\sqrt{p^m}\right\rfloor - 1},$$

as required. □

Corollary 5.1.12. *Let $\alpha \in$ **Bad** and assume there is some sequence of natural numbers $\{\ell_m\}_{m\in\mathbb{N}}$ such that $p^{\ell_m}\alpha$ is not an infinite loop mod p^m. Then α satisfies pLC.*

Proof. From Corollary 5.1.11, we can conclude that for any $\alpha \in \mathbb{R}_{>0}$, if there is a sequence of natural numbers $(\ell_m)_{m\in\mathbb{N}}$ such that $p^{\ell_m}\alpha$ is not an infinite loop mod p^m, then we have:

$$m_p(\alpha) \leq \lim_{m\to\infty} \frac{1}{\left\lfloor 2\sqrt{p^m}\right\rfloor - 1} = 0.$$

Therefore, α satisfies pLC. □

Here, we should note that the sequence $\{\ell_m\}$ need not be monotonically increasing. For example, we may have that α is an infinite loop mod p, but there exist some $K \in \mathbb{N}$, such that $p^K\alpha$ is not an infinite loop mod p^m for all $m \in \mathbb{N}$. In this case, the sequence $l_m = K$ for all $m \in \mathbb{N}$ would allow us to show that α satisfies pLC. The following proposition shows that for all α with eventually recurrent continued fraction expansions, there is some $K \in \mathbb{N}$ such that $p^K\alpha$ is not an infinite loop mod p^m for all $m \in \mathbb{N}$.

Proposition 5.1.13. *Let $\alpha \in \mathbb{R}_{>0}$ be a real number with an essentially recurrent continued fraction expansion and let p be a prime number. Then α is not an infinite loop mod n, for any $n \in \mathbb{N}$. Furthermore, if $\alpha \in \mathbb{R}_{>0}$ is a real number with an eventually recurrent continued fraction expansion, then for any prime number p,* *there exists some $K \in \mathbb{N}$ such that $p^K\alpha$ is not an infinite loop mod p^m, for all $m \in \mathbb{N}$. In particular, eventually recurrent continued fractions satisfy pLC.*

Proof. First we will assume that α is essentially recurrent. Let ζ_α be the geodesic ray starting at I and terminating at α. Then, the projection $\widehat{\zeta_\alpha}$ of ζ_α in $\Gamma_0(n)\backslash\mathbb{H}$, will have a strictly recurrent cutting sequence $(\widehat{\zeta_\alpha}, \widehat{\mathcal{F}})$. Therefore, the geodesic $\widehat{\zeta_\alpha}$ will be strictly recurrent relative to the arc $\widehat{I}$, which is the projection of I in $\Gamma_0(n)\backslash\mathbb{H}$. By Proposition 4.2.1, we can conclude that $\widehat{\zeta_\alpha}$ intersects $\widehat{I}$ infinite often. By lifting $\widehat{\zeta_\alpha}$ back to ζ_α in $\mathbb{H}$, we can see that ζ_α intersects infinitely many edges of $\Gamma_0(n) \cdot I$ - since $\widehat{I}$ lifts to $\Gamma_0(n) \cdot I$. We can therefore conclude that ζ_α, and by extension α, is not an infinite loop mod n.

We will now assume that α has an eventually recurrent continued fraction expansion. Let ζ_α be the geodesic ray starting at I and terminating at α. Then, we can cut ζ_α along some edge in $\mathcal{F}$ to form two paths, ρ_α and ξ_α, such that $\zeta_\alpha \simeq \rho_\alpha \circ \xi_\alpha$ and ξ_α has a strictly recurrent cutting sequence $(\xi_\alpha, \mathcal{F})$. Let z starting point of ξ_α and let $x = \mathrm{Re}(z)$. Since ζ_α is a geodesic ray that approaches α from I, we can that conclude that ξ_α approaches α from the "left", i.e. $x < \alpha$. Therefore, the interval $[x, \alpha)$ is non-empty.

Since this interval is non-empty, we can guarantee that for each prime p, there exists natural numbers $a \in \mathbb{N} \cup \{0\}$ and $K \in \mathbb{N}$ such that $\frac{a}{p^K} \in [x, \alpha)$. Since there is an edge between a and ∞ in $\mathcal{F}$, we can deduce that there is an edge between $\frac{a}{p^K}$ and ∞ in $\frac{1}{p^K}\mathcal{F}$. We will label this edge E. We can then take $\widehat{\zeta_\alpha}$ to be the projection of ζ_α in $\Gamma_0(n)\backslash\mathbb{H}$ and take $\widehat{\xi_\alpha}$ to be the projection of ξ_α in $\Gamma_0(n)\backslash\mathbb{H}$. Similarly, we will take $\widehat{E}$ to be the projection of E in $\Gamma_0(n)\backslash\mathbb{H}$. Since the cutting sequence $(\xi_\alpha, \mathcal{F}) = (\widehat{\xi_\alpha}, \widehat{\mathcal{F}})$ is strictly recurrent, we can use Lemma 4.2.1 to deduce that $\widehat{\xi_\alpha}$ intersects $\widehat{E}$ infinitely often. Furthermore, we can decompose $\widehat{\xi_\alpha}$ into a finite path $\widehat{\pi_\alpha}$, which runs along $\widehat{\xi_\alpha}$ until the first time $\widehat{\xi_\alpha}$ intersects $\widehat{E}$, followed by an infinite path $\widehat{\eta_\alpha}$. This path $\widehat{\eta_\alpha}$ is strictly geometrically recurrent relative to $\widehat{E}$. In particular, the cutting sequence $(\widehat{\eta_\alpha}, \widehat{\frac{1}{p^K}\mathcal{F}})$ is strictly recurrent.

Going back to $\mathbb{H}$, we can see that the lift η_α of $\widehat{\eta_\alpha}$ has a strictly recurrent cutting sequence $(\eta_\alpha, \frac{1}{p^K}\mathcal{F}) = (\widehat{\eta_\alpha}, \widehat{\frac{1}{p^K}\mathcal{F}})$. The path η_α has a well-defined cutting sequence relative to $\frac{1}{p^K}\mathcal{F}$, since the starting edge E is in $\frac{1}{n}\mathcal{F}$. Rescaling our geodesic ray by $(p^K)^*$, we see that $(p^K)^* \cdot \zeta_\alpha$ intersects the line between a and ∞ in $\mathbb{H}$. By rescaling η_α, we see that $(p^K)^* \cdot \eta_\alpha$ starts at this line between a and ∞. The cutting sequence $((p^K)^* \cdot \eta_\alpha, \mathcal{F})$ is equivalent to $(\eta_\alpha, \frac{1}{p^K}\mathcal{F})$ and is therefore strictly recurrent. If we take $\beta = p^K\alpha - a$, then the map $\varphi = \left(\begin{smallmatrix} 1 & -a \\ 0 & 1 \end{smallmatrix}\right)$ takes the starting edge of $(p^K)^* \cdot \eta_\alpha$, which goes from a and ∞, to the edge I. It also maps the endpoint $p^K\alpha$ of $(p^K)^* \cdot \eta_\alpha$ to β. In particular, the cutting sequence $((p^K)^* \cdot \eta_\alpha, \mathcal{F})$ directly corresponds to the continued fraction expansion of β. Since $((p^K)^* \cdot \eta_\alpha, \mathcal{F})$ is strictly recurrent, we can conclude that the continued fraction expansion of β is essentially recurrent. Finally, since β and $p^K\alpha$ only differ by an integer and β has an essentially recurrent continued fraction expansion, we can deduce that $p^K\alpha$ is not an infinite loop mod p^m, for any $m \in \mathbb{N}$. □

In contrast to Corollary 5.1.12, if there exists an $m \in \mathbb{N}$ such that $p^\ell\alpha$ is an infinite loop mod p^m, for all $\ell \in \mathbb{N} \cup \{0\}$, then α is a counterexample to pLC.

Lemma 5.1.14. *Let $\alpha \in$ **Bad** and assume there exists an $m \in \mathbb{N}$ such that $p^\ell\alpha$ is an infinite loop mod p^m, for all $\ell \in \mathbb{N} \cup \{0\}$. Then α is a counterexample to pLC and $m_p(\alpha) \geq \frac{1}{p^m-2}$.*

In order to prove this statement, we will first show that if $\beta \in \mathbb{R}_{>0}$ is a real number such that $p^j\beta$ is an infinite loop mod p^m for all $j \in \mathbb{N} \cup \{0\}$, then $b_{k+1} \leq p^m - 4$, where $\overline{\beta} = [b_0; b_1, \ldots]$. As a result, we can then conclude that $B(\beta) \leq p^m - 4$.

Claim: If $\beta \in \mathbb{R}_{>0}$ is a real number such that $p^j\beta$ is an infinite loop mod p^m for all $j \in \mathbb{N} \cup \{0\}$, then $b_{k+1} \leq p^m - 4$, where $\overline{\beta} = [b_0; b_1, \ldots]$. In particular, we can then conclude that $B(\beta) \leq p^m - 4$.

Proof of claim. Let b_{k+1} be an arbitrary partial quotient of β for some $k \in \mathbb{N}$ and consider the following two cases for the corresponding convergent denominator q_k:

(Case I): The prime p and q_k are coprime.

(Case II): The prime p and q_k are not coprime.

(Case I): Since q_k is coprime with p, we know that there are infinitely many neighbours of $\frac{p_k}{q_k}$ which have a denominator divisible by p^m. This is analogous to the fact the 0 has infinitely many neighbours of the form $\frac{1}{p^m j}$, where $j \in \mathbb{N}$. The corresponding geodesic ray ζ_β must not intersect any of the geodesic arcs from $\frac{p_k}{q_k}$ to these neighbours. As a result, there is a unique pair of neighbours, $\frac{a_1}{c_1 p^m}$ and $\frac{a_2}{c_2 p^m}$, such that the arcs between these points and $\frac{p_k^\ell}{q_k^\ell}$ separate ζ_β from all other neighbours of $\frac{p_k}{q_k}$ whose denominator divisible by p^m. See Fig 5.2.

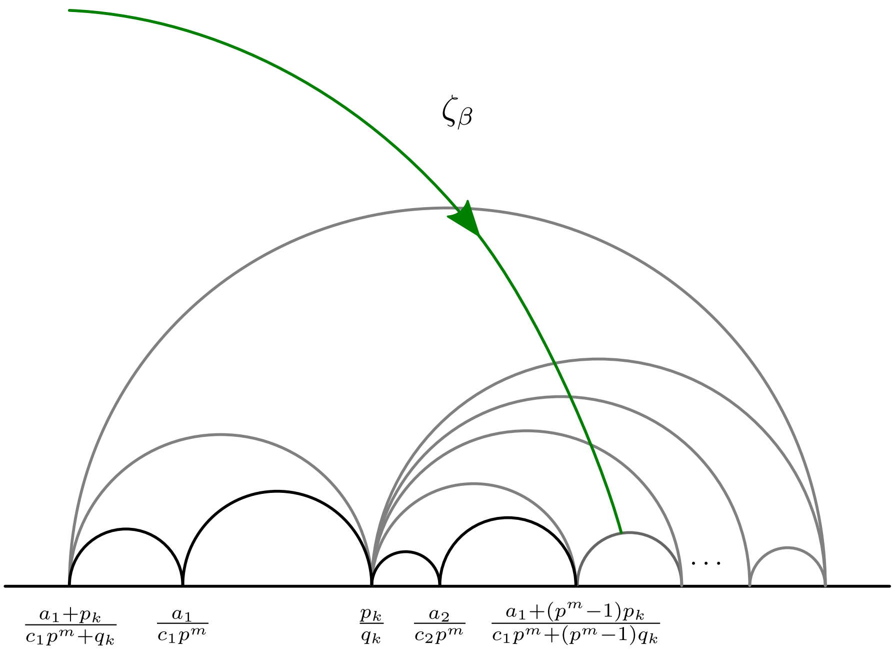

Figure 5.2: An image of ζ_β cutting a fan (relative to $\mathcal{F}$) with fixed point $\frac{p_k}{q_k}$ and $\gcd(p, q_k) = 1$. In this scenario, ζ_β is an infinite loop mod p^m and forms the largest fan possible for this b_{k+1}.

Similarly, we can express all other neighbours of $\frac{p_k}{q_k}$ in this region by using the Farey sum on $\frac{p_k}{q_k}$ and $\frac{a_1}{c_1 p^m}$ (up to relabelling). More explicitly, the neighbours in this region are of the form:

$$n_i = \frac{a_1 + i \cdot p_k^\ell}{c_1 p^m + i \cdot q_k^\ell},$$

where $i \in \{0, 1, \ldots, p^m\}$ and $n_0 = \frac{a_1}{c_1 p^m}$ and $n_{p^m} = \frac{a_2}{c_2 p^m}$.

Two of these neighbours will be fixed vertices for the previous and subsequent fans, and we label these neighbours as n_s and n_t with $t > s$. The size of the fan b_{k+1} is given by $t - s$. The points $\frac{p_k}{q_k}, n_s$ and n_{s-1} form a triangle in $\mathcal{F}$, and so, since n_s is a convergent denominator of ζ_β, the point n_{s-1} must be a semi-convergent of $p^\ell\alpha$. Similarly, since n_t is the convergent of the next partial quotient, the point n_{t+1} is a semi-convergent of β. If either n_0 or n_{p^m} are semi-convergents of β, then, since they are of the form $\frac{A}{Cp^m}$ with $C \neq 0$, we can conclude that β is not an infinite loop mod p^m. It follows that for ζ_β to be an infinite loop mod p^m, we have $s \in \{2, \ldots, p^m - 3\}$ and $t \in \{3, \ldots, p^m - 2\}$. Therefore, the maximum size of the fan is b_{k+1} is given by $(\max t - \min s) = p^m - 2 - 2 = p^m - 4$, as required. *QED.*

(Case II): In this case, there is some $j \in \mathbb{N}$ such that $p^j \mid q_k$ and $p^{j+1} \nmid q_k$. We will write $q_k = p^j d_k$, where $d_k \in \mathbb{N}$ and $\gcd(q_k^{'}, p) = 1$. Therefore, by Corollary 5.1.9, we can deduce that $B(p^j\beta) \geq p^j b_{k+1}$ and $p_k d_k$ is a convergent of $p^j\beta$. We wish to show that if $b^j_{k'+1}$ is the partial quotient of $p^j\beta$ corresponding to the convergent $p_k d_k$, then we have $b_{k+1} \cdot p^j \leq b^j_{k'+1}$. Since $\gcd(d_k, p) = 1$, we can use Case I to conclude that $b_k < b^j_{k'+1} \leq p^m - 4$.

The geodesic ray ζ_β forms a fan B_{k+1} with $\mathcal{F}$ of size b_{k+1} and this fan has a fixed vertex $\frac{p_k}{q_k}$. Since $p^j \mid q_k$, any neighbour $\frac{a}{c}$ of $\frac{p_k}{q_k}$ in $\mathcal{F}$ must satisfy $\gcd(c, p^j) = 1$. By Corollary 5.1.5, the edge between a neighbour of this form and $\frac{p_k}{q_k}$ must be an edge of $\mathcal{F} \cap \frac{1}{p^j}\mathcal{F}$. Therefore, since every edge that ζ_β intersects in B_{k+1} has $\frac{p_k}{q_k}$ as one of its endpoints, we can conclude that each of these edges lie in $\mathcal{F} \cap \frac{1}{p^j}\mathcal{F}$. Since these edges lie in $\mathcal{F} \cap \frac{1}{p^j}\mathcal{F}$, we can guarantee that $\frac{p_k}{q_k}$ is a fixed point of some fan $B^j_{k'+1}$ in the cutting sequence $(\zeta_\beta, \frac{1}{p^j}\mathcal{F})$. After having corrected for scaling, we can observe that $\frac{p_k}{d_k}$ is a convergent of $p^j\beta$, where $q_k = d_k \cdot p^j$, as above. Each triangle in B_k is sub-divided into p^j triangles when we replace $\mathcal{F}$ with $\frac{1}{p^j}\mathcal{F}$, as described in Proposition 5.1.8 and Table 5.1. Therefore, in analogy to Corollary 5.1.9, if $b^j_{k'+1}$ is the partial quotient of $p^j\beta$ corresponding to the fan $B^j_{k'+1}$ - with corresponding convergent $\frac{p_k}{d_k}$ - then $b^j_{k'+1}$ satisfies:

$$p^j \cdot b_{k+1} \leq b^j_{k'+1}.$$

Since $gcd(d_k, p) = 1$, we can use Case I to see that $b^j_{k'+1} < p^m - 4$. However, since $p^j \cdot b_{k+1} \leq b^j_{k'+1}$, we can conclude $b_{k+1} \leq p^m - 4$, as required. *QED.*

Finally, since the partial quotients b_{k+1} of β are all bounded above by $p^m - 4$, we can conclude that $B(\beta) \leq p^m - 4$ and this completes the proof of the claim. □

Proof of Lemma 5.1.14. For each $\ell \in \mathbb{N} \cup \{0\}$, $p^\ell \alpha$ is not an infinite loop mod p^m. Therefore $p^{\ell+j}\alpha$ is also an infinite loop mod p^m for all $j \in \mathbb{N} \cup \{0\}$. As a result, we can replace β in the above claim by $p^\ell \alpha$, and we see that $B(p^\ell \alpha) \leq p^m - 4$, for all $\ell \in \mathbb{N} \cup \{0\}$. As seen in Lemma 2.2.16, we know that:

$$m_p(\alpha) \geq \inf_{\ell \in \mathbb{N} \cup \{0\}} \frac{1}{B(p^\ell \alpha) + 2}.$$

Finally, we can conclude that:

$$m_p(\alpha) \geq \frac{1}{p^m - 2}.$$

□

Combining Corollary 5.1.11 and Lemma 5.1.14, we get the following theorem.

Theorem 5.1.15. *Let $\alpha \in$ **Bad**. Then α satisfies pLC if and only if there is a sequence of natural numbers $\{\ell_m\}_{m \in \mathbb{N}}$ such that $p^{\ell_m}\alpha$ is not an infinite loop mod p^m.*

5.2 Constructing Infinite Loops

Now that we have discussed the importance of infinite loops mod n and their link to the p-adic Littlewood Conjecture, it is natural to ask "For a given n, how can one construct an infinite loop mod n?". Here we present two ways: one geometric, and the other arithmetic. Both are fairly simple processes, but due to the natures of how these objects are formed, the second is much more practical to be run on computer. There is hope that with further understanding of these objects, we may be able to find some real number $\alpha \in \mathbb{R}_{>0}$, which is an infinite loop mod p^m, such that $p^k\alpha$ is also an infinite loop mod p^m for all $1 \le k \le K$, where K is some large number. If K were suitably large, i.e. $K = 1000$, then such an object would be very strong potential counterexample to pLC.

5.2.1 Theoretical Construction of Infinite Loops

In this section, we provide the groundwork for how one would theoretically construct an infinite loop mod n. As previously discussed, we present two different ways of doing this. In both cases, the algorithm that we present constructs infinite loops mod n in the interval $[0, 1)$. These processes construct all possible infinite loops mod n in this interval. All other infinite loops mod n can be constructed by adding some arbitrary positive integer.

Constructing Infinite Loops via Geometry

As previously mentioned, we can view infinite loops mod n as real numbers α, which have corresponding geodesic rays ζ_α that are disjoint from $\Gamma_0(n) \cdot I$ except at $I + k$, for $k \in \mathbb{Z}_{\ge 0}$. To construct an infinite loop ζ mod n in this setting, we will note that every sub-path path ζ must be disjoint from $\Gamma_0(n) \cdot I$, except at $I + k$ for $k \in \mathbb{Z}_{\ge 0}$. Therefore, we can reverse-engineer an infinite loop mod n by concatenating paths λ_i which are totally disjoint from $\Gamma_0(n) \cdot I$ (except at $I + k$). As we did in Section 3.2.3, we will take these paths λ_i to be completely contained in some fundamental domain P_n. Note that, the concatenated path $\lambda = \lambda_0 \circ \lambda_1 \circ \cdots$ will not be a geodesic ray, however it will be reduced relative to $\mathcal{F}$. By Corollary 3.1.12, λ will be homotopic relative to I to a geodesic ray ζ, and the cutting sequences $(\lambda, \mathcal{F})$ and $(\zeta, \mathcal{F})$ will be equivalent. For ease, we will construct paths which are disjoint from $\Gamma_0(n) \cdot I$ everywhere, except at the starting edge I. In particular, the endpoint of this path will lie in the interval $(0,1)$.

Take P_n to be a special polygon of $\Gamma_0(n)$, which has I and $I+1$ as edges. The edges $\mathcal{E}_n$ on the boundary of P_n will be our set of *states* and I will be our *initial state*. Any (oriented) path λ in P_n which starts at an edge E_i and terminates at an edge E_j will be completely determined - up to relative homotopy - by these initial and final edges. We label such paths λ_{E_i,E_j}, to indicate the initial and final edge. Since the cutting sequence is invariant under homotopy and the boundary edges of P_n are in $\mathcal{F}$ or an odd edge, the cutting sequence $(\lambda_{E_i,E_j}, T_{\{1,n\}})$ is well-defined and unique given some choice of starting and terminal edges. Here, as in Section 3.2.3, we view λ intersecting a odd triangle as being equivalent to cutting a half triangle. In this case, we append $L^{\frac{1}{2}}$ or $R^{\frac{1}{2}}$ to the cutting sequence depending on whether this half-triangle is a left triangle or a right triangle, respectively.

Given a path λ_{E_i,E_j} and a path $\lambda_{\overline{E_j},E_k}$, where $\overline{E_j}$ is the edge paired with E_j in P_n, we can concatenate these paths by gluing together two copies of P_n (which contain these paths) along the identified edge. If W_1 is the cutting sequence $(\lambda_{I,E_i}, T_{\{1,n\}})$ and W_2 is the cutting sequence $(\lambda_{E_i',E_j}, T_{\{1,n\}})$, then the cutting sequence of the concatenated path, relative the induced triangulated polygon, will simply be W_1W_2.

To construct a path which avoids $\Gamma_0(n)\cdot I$, except at I, we can pick some initial sub-path λ_{I,E_i} and then iteratively append sub-paths in P_n which are disjoint from I and $I+1$. We can take $I\Lambda_n(I)$ to be the set of all paths in P_n - considered up to relative homotopy - which start at I, but do not otherwise intersect I or $I+1$. This will be our *input alphabet* for our initial state I. Similarly, we define $I\Lambda_n(E_i)$ to be the set of all paths starting at the edge E_i which are completely disjoint from I and $I+1$. The set $I\Lambda_n(E_i)$ will be our *input alphabet* for the state E_i. For all edges E_i of P_n, we define the *output alphabet* to be $IL_n(E_i) := \{(\lambda_j, T_{\{1,n\}}) : \lambda_j \in I\Lambda_n(E_i)\}$. Using this input alphabet and the notions presented in Section 3.2.3, we can define a definitive finite automaton with output, which constructs the cutting sequence infinite loops mod n, as follows:

The automaton which constructs infinite loops mod n

Let P_n be a special polygon for $\Gamma_0(n)$ containing the edges I and $I+1$:

1. Let $\mathcal{E}_n$ be the set of the edges in P_n. These will be our *states*. The edge I will be our *initial state*.

2. For every $E_i \in \mathcal{E}_n$ construct $I\Lambda_n(E_i)$ and $IL_n(E_i)$. The set $I\Lambda_n(E_i)$ will be our *input alphabet* and the set $IL_n(E_i)$ will be our *output alphabet*.

3. For each path in $I\Lambda_n(E_i)$ from E_i to E_j:

   - The *transition function* $\delta : \mathcal{E}_n \times \mathcal{E}_n \to \mathcal{E}_n$ is given by $\delta(E_i, E_j) \mapsto \overline{E_j}$, where $\overline{E_j}$ is the edge identified to E_j via side pairings.
   - The *output function* $\tau : I\Lambda_n(E_i) \to IL_n(E_i)$ is given by $\lambda_{E_i,E_j} \mapsto (\lambda_{E_i,E_j}, T_{\{1,n\}})$, where λ_{E_i,E_j} is the unique path from E_i to E_j taken up to homotopy.

Remark 5.2.1. The process of constructing the automaton for forming infinite loops mod n is much easier than the construction of the multiplication algorithm mod n. This is because we do not need to know anything about cutting sequence of the paths λ_{E_i,E_j} with $T_{\{n,n\}}$. In particular, we only need states corresponding to the edges - and not to based points of edges.

Constructing Infinite Loops via Arithmetic Methods

We can also construct infinite loops mod n by iteratively constructing partial quotients which do not admit semi-convergent denominators divisible by n. The main tool that we use is the recurrence relation for semi-convergent denominators:

$$q_{\{k,m\}} = mq_k + q_{k-1},$$

where $0 \le m \le a_{k+1}$.

If we assume that $[0; a_1, a_2, \dots, a_k]$ does not admit any semi-convergent denominators divisible by n, then we can ask what values of a_{k+1} can we take, such that

$[0; a_1, a_2, \ldots, a_k, a_{k+1}]$ does not admit any semi-convergent denominators are divisible by n. In particular, for a given value of a_{k+1}, is there some $0 \le m \le a_{k+1}$ such that $q_{\{k,m\}} \equiv 0 \mod n$? Since $q_{\{k,m\}} = mq_k + q_{k-1}$, this is equivalent to asking whether $mq_k \equiv -q_{k-1} \mod n$, for any $0 \le m \le a_{k+1}$. Treating m as an integer variable, and q_k and q_{k-1} as fixed integers, we obtain a simple linear congruence. This linear congruence has a unique solution mod n if and only if $\gcd(q_k, n)$ divides $-q_{k-1}$. We note that if there is a value $M \in \{1, \ldots, n-1\}$ with $Mq_k \equiv -q_{k-1} \mod n$, then as long as $a_{k+1} \ge M$ the continued fraction expansion $[0; a_1, \ldots, a_k, a_{k+1} \ldots]$ will admit the semi-convergent denominator $q_{\{k,M\}}$ which is divisible by n. Since $\gcd(q_k, q_{k-1}) = 1$, if $\gcd(q_k, n) = g \ne 1$, then $g \nmid q_{k-1}$. Therefore, if $\gcd(q_k, n) = g \ne 1$, then we have $Mq_k \not\equiv -q_{k-1} \mod n$ for any value of $M \in \mathbb{Z}$. In particular, if a convergent denominator shares a common factor with n but is not divisible by n, then it can contain a partial quotient of arbitrary size and still be an infinite loop mod n. A good example of this is that $\frac{1}{2}$ is a infinite loop mod 4. Here, we can write the continued fraction expansion of $\frac{1}{2}$ as $[0; 1, 1, \infty]$ or $[0; 2, \infty]$. If we take $[0; 1, 1, \infty]$, the (semi-)convergent denominators are:

$$q_{-1} = 0, q_0 = 1, q_1 = 1, q_2 = 2, q_{2,k} = 2k + 1 \equiv (-1)^k \mod 4.$$

Equivalently, if we take $[0; 2, \infty]$ the semi-convergent denominators are:

$$q_{-1} = 0, q_0 = 1, q_{0,1} = 1, q_{0,2} = 2, q_{1,k} = 2k + 1 \equiv (-1)^k \mod 4.$$

In both cases, none of the semi-convergent denominators are divisible by 4 (except for when $q_{-1} = 0$).

The above observations give us a good basis for constructing an arithmetic automaton for constructing infinite loops mod n. Our *states* are given by the "possible" pairs of consecutive convergent denominators (q_k, q_{k-1}) taken mod n. We could naively express these states as elements in $\mathbb{Z}_n \times \mathbb{Z}_n$, however, this will include some states which are not accessible from other states without introducing semi-convergent denominators that are divisible by n. For $k \ge 2$, q_k and q_{k-1} will both be non-zero elements in $\mathbb{Z}_n$, corresponding to the fact that we do not want our convergent denominators divisible by n, except for q_0. Since q_0 is always equal to 0 and q_1 is always equal to 1, the only pairs of convergent denominators which allow for infinite loops will be contained in the set $\{(a,b) : a,b \in \mathbb{Z}_n^\times\} \cup \{(1,0)\}$. In fact, this set is not minimal and we can further remove redundant states.

If we have a pair of convergent denominators (q_k, q_{k-1}) with $\gcd(q_k, n) = 1$ and $q_k + q_{k-1} \equiv 0$, then the pair (q_{k+1}, q_k) always admits a semi-convergent denominator divisible by n. This is the semi-convergent corresponding to $[0; a_1, \ldots, a_k, 1]$, which has the denominator $q_k + q_{k+1} \equiv 0 \mod n$. We also note that we can not have pairs of the form (q_k, q_{k-1}), where $q_k \equiv q_{k-1} \mod n$, unless $(q_2, q_1) = (1, 1)$. This is because $q_k \equiv a q_k + q_{k-2} \equiv a q_{k-1} + q_{k-2} \mod n$ if and only if $(a-1)q_k + q_{k-2} \equiv (a-1)q_{k-1} + q_{k-2} \equiv 0 \mod n$. The equation $(a-1)q_{k-1} + q_{k-2} \equiv 0 \mod n$ has a unique solution for $a \in \{1, \ldots, n\}$. However, if we take $a-1, a \in \{1, \ldots, n\}$ such that $a - 1 < a$, then the sequence of partial quotients $[0; a_1, \ldots, a_{k-1}, a]$ will admit the semi-convergent of the form $[0; a_1, \ldots, a_{k-1}, a-1]$, which has denominator of the form $(a-1)q_{k-1} + q_{k-2} \equiv 0 \mod n$. The only way that we can resolve this is by having $a \equiv 1 \mod n$ and $a - 1 \equiv n \mod n$. However, this can only occur when $q_{k-2} \equiv 0 \mod n$, which by assumption only occurs when $k = 2$. As a result, $(q_k, q_{k-1}) = (1, 1) \mod n$. We also do not allow elements of the form $(a, -2a)$, since when $a_{k+1} = 1$ this leads to the pair (q_{k+1}, q_k) of the form $(-a, a) \mod n$ and when $a_k \geq 2$, we have that $[0; a_1, \ldots, a_k, 2]$ is a semi-convergent with denominator divisible by n (since $2a + (-2a) \equiv 0 \mod n$). We denote the set $\mathbb{Z}_n \times \mathbb{Z}_n$ excluding elements of the form $(a, -a)$, $(a, -2a)$, (b, b) , $(0, c)$ and $(d, 0)$ for $b, d \neq 1$, as $\mathbf{Pairs}_n$. This set will be the *states* for our algorithm.

For each pair $(q_k, q_{k-1}) \in \mathbf{Pairs}_n$ with $\gcd(q_k, n) = 1$, we can compute:

$$M(q_k, q_{k-1}) := \sup\{m \in \{1, \ldots, n-1\} : m q_k + q_{k-1} \not\equiv 0 \bmod n\}.$$

For pairs (q_k, q_{k-1}) with $\gcd(q_k, n) \neq 1$, we take $M(q_k, q_{k-1}) = \infty$. Given a pair (q_k, q_{k-1}), we can ask whether the induced pair $(\ell q_k + q_{k-1}, q_k)$ with $\ell \in$

$\{1, \ldots, M(q_k, q_{k-1}) - 1\}$ will lie in $\mathbf{Pairs}_n$ or not. This induces a natural *transition function* $\delta : \mathbf{Pairs}_n \to \mathbf{Pairs}_n$ and *output function* $\tau : \mathbb{N} \to \mathbb{N}$, whenever these maps are well defined, i.e. whenever $(\ell q_k + q_{k-1}, q_k)$ is in $\mathbf{Pairs}_n$. We will use x as a stand-in for when this map is not defined. We explicitly express δ and τ, as follows:

$$\delta(q_k, q_{k-1}, \ell) := \begin{cases} (\ell q_k + q_{k-1}, q_k) & \text{if } (\ell q_k + q_{k-1}, q_k) \in \mathbf{Pairs}_n \\ \text{x} & \text{otherwise} \end{cases}$$

$$\tau(q_k, q_{k-1}, \ell) := \begin{cases} \ell & \text{if } (\ell q_k + q_{k-1}, q_k) \in \mathbf{Pairs}_n \\ \text{x} & \text{otherwise} \end{cases}$$

Note that the value ℓ in the range $\{1, 2, \ldots, M(q_k, q_{k-1}) - 1\}$ contains the set of possible partial quotients the pair (q_k, q_{k-1}) can take to avoid admitting a semi-convergent denominator divisible by n. This range is taken since, for $M(q_k, q_{k-1}) < \infty$, we have that $q_k{+}1 \equiv M(q_k, q_{k-1}) \cdot q_k + q_{k-1} \equiv -q_k \mod n$ and this leads to a redundant pair of the form $(-a, a)$. For a fixed pair (q_k, q_{k-1}), the collection of all possible transition functions represents all of the possible values the next convergent denominator q_{k+1} can take, such that the corresponding sequence of partial quotients does not admit a semi-convergent denominator divisible by n. The collection of output functions represents all of the possible partial quotients in this case. Here, the *output alphabet* and *input alphabet* are equivalent, i.e. they are both $\{1, \ldots, \max M(a, b) - 1\}$. Note that if n is prime, then these sets are both finite. However, if n is composite this is not necessarily the case. Although the input and output algorithm are not finite sets, the value of ℓ is only important - with regards to the automaton - modulo n and so the automaton is still well-defined.

We can also note, that if $\gcd(q_k, n) = 1$ and $(q_k, q_{k-1}) \neq (1, 0)$, then $M(q_k, q_{k-1}) < n - 3$. This result essentially follows from the observation that for an arbitrary pair (a, b), we have that $M(a, b) > M(a, b + a)$ and $(a, 2a)$ gives the highest value of $M(a, 2a) = n - 2$. This observation is equivalent to the observation that if $gcd(q_k, p) = 1$ and $a_{k+1} > p^\ell - 4$, then this sequence is not an infinite loop mod p^ℓ.

5.2.2 Constructing Infinite Loops mod 5 and 7

As with the multiplication, the best way to illustrate this process is by example. We do this for reasonably low primes, but the process analogous for all positive integers.

Constructing Infinite Loops mod 5: The Geometric Method

For $p = 5$ the set of edges $\mathcal{E}_5$ is given by the edges $\left\{E_0 = I, \overline{E_0}, E_1, E_1'\right\}$. There are two edges in P_5, E_1 and E_1', which are not $E_0 = I$ or $\overline{E_0} = I + 1$. As a result, the infinite loop alphabet of paths starting at E_0 is given by $I\Lambda_5(E_0) := \{\lambda_{E_0,E1}\lambda_{E_0,E_1'}\}$. These paths have cutting sequence R^2 and RL, respectively, and so our output alphabet is $IL_5(I) = \left\{R^2, RL\right\}$. See Fig. 5.3 (a). For the edge E_1, there is one path $\lambda_{E_1,E_1'}$ in the infinite loop alphabet $I\Lambda_5(E_1)$. This path has cutting sequence R and so our output alphabet is $IL_5(E_1) = \{R\}$. See Fig. 5.3 (b). Similarly, for the edge

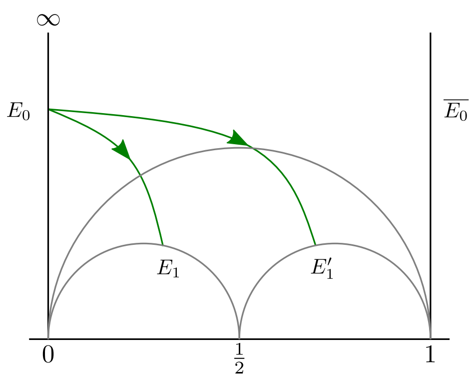

(a) An image of the paths in $I\Lambda_5(E_0)$.

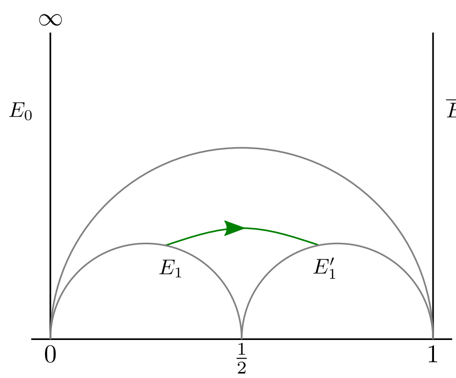

(b) An image of the path in $I\Lambda_5(E_1)$.

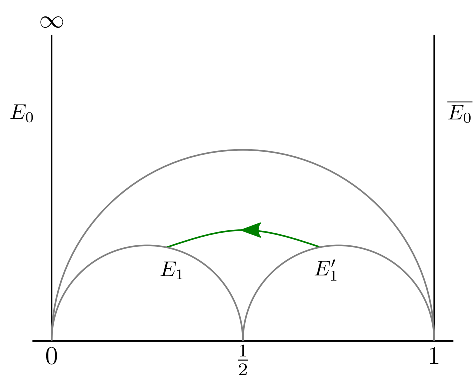

(c) An image of the path in $I\Lambda_5(E_1')$.

Figure 5.3: This figure shows the possible different infinite loop alphabets for $p = 5$.

E_1', there is one path λ_{E_1',E_1} in the infinite loop alphabet $I\Lambda_5(E_1')$. This path has cutting sequence L and so our output alphabet is $IL_5(E_1) = \{L\}$. See Fig. 5.3 (c).

Note that the edges E_1 and E_1' are paired with themselves and so $E_1 = \overline{E_1}$ and $E_1' = \overline{E_1'}$. Using this information, we can recover Table 5.2 and we can construct the corresponding automaton, as in Fig. 5.4.

Initial State	Transition Function	Output Function
E_0	$\lambda_{E_0,E_1'} \to E_1$ $\lambda_{E_0,E_1} \to E_1'$	R^2 RL
E_1	$\lambda_{E_1,E_1'} \to E_1'$	R
E_1'	$\lambda_{E_1',E_1} \to E_1$	L

Table 5.2: A table of the transition functions and output functions used to construct infinite loops mod 5.

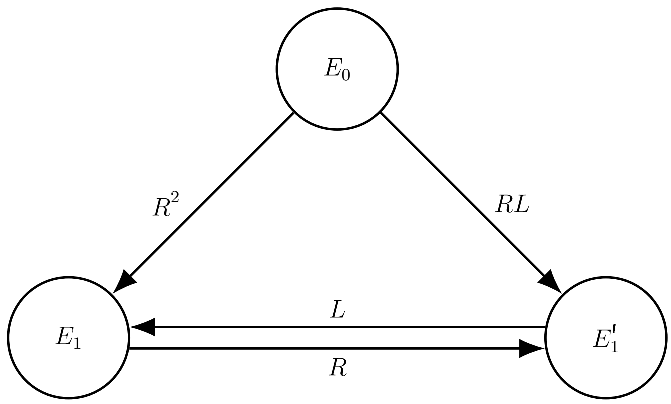

Figure 5.4: An automaton that constructs infinite loops mod 5using the geometric method. The initial state is E_0.

Since there is exactly one path that one can take from E_1 and E_1', there are only two possible infinite loops mod 5. These correspond to the real numbers $\frac{5+\sqrt{5}}{10}$ and $\frac{5-\sqrt{5}}{10}$ with continued fraction expansions $[0;1,2,\overline{1}]$ and $[0;3,\overline{1}]$, respectively. Interestingly, when multiplied by 5, we find that the continued fraction expansions are $5 \cdot \frac{5+\sqrt{5}}{10} = \frac{5+\sqrt{5}}{2} = [3;\overline{1}]$ and $5 \cdot \frac{5-\sqrt{5}}{10} = \frac{5-\sqrt{5}}{2} = [1;2,\overline{1}]$.

Constructing Infinite Loops mod 5: The Arithmetic Method

We start by first constructing the set $\textbf{Pairs}_5 := \{(1,0),(1,1),(1,2),(2,1),(3,1),(4,3)\}$. We note that excluding $(1,0)$ and $(1,1)$, the maximum value of $M(a,b) = 2$, i.e. each possible partial quotient must equal 1, for these cases. For $(1,0)$, we can have 1 or 3 as possible partial quotients and for $(1,1)$ we can have 1 or 2. Using this information, we construct Table 5.3 to represent the infinite loop construction algorithm and construct the automaton as in Fig. 5.5

Next partial quotient a_{k+1}	States					
	$(1,0)$	$(1,1)$	$(1,2)$	$(2,4)$	$(3,1)$	$(4,3)$
1	$(1,1)$	x	$(3,1)$	$(1,2)$	$(4,3)$	$(2,4)$
2	x	$(3,1)$	x	x	x	x
3	$(3,1)$	x	x	x	x	x

Table 5.3: A table showing the possible values the next partial quotient can take to form an infinite loop mod 5.

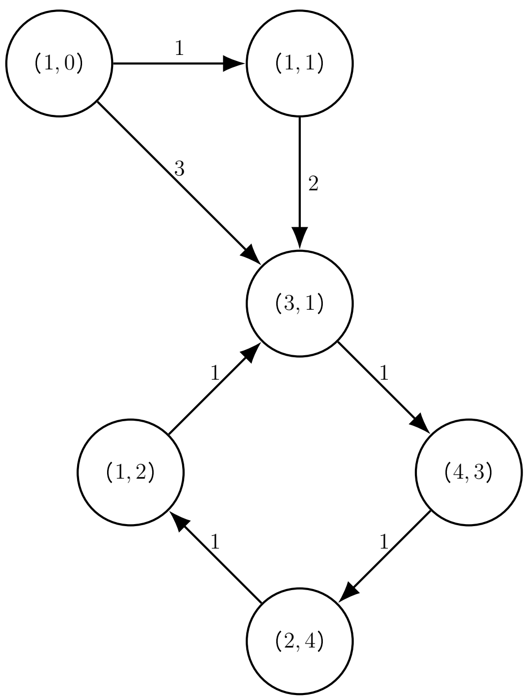

Figure 5.5: An automaton that constructs infinite loops mod 5 using the arithmetic method. The initial state is $(1,0)$.

Here, we see that there are only two paths we can take around this automaton: One which goes from $(1,0)$ to $(1,1)$ to $(3,1)$ and then repeatedly goes from $(3,1)$ to

$(4, 3)$ to $(2, 4)$ to $(1, 2)$ and then back to $(3, 1)$, and another which goes from $(1, 0)$ to $(3, 1)$ and then repeatedly goes from $(3, 1)$ to $(4, 3)$ to $(2, 4)$ to $(1, 2)$ and then back to $(3, 1)$. These paths produce the continued fraction expansions $[0; 1, 2, \overline{1}]$ and $[0; 3, \overline{1}]$, respectively.

Constructing Infinite Loops mod 7: The Geometric Method

For $p = 7$ the set of edges $\mathcal{E}_5$ is given by the edges $\left\{E_0 = I, \overline{E_0}, E_1, \overline{E_1}, E_2, \overline{E_2}\right\}$. There are four edges in P_7, $E_1, \overline{E_1}, E_2$ and $\overline{E_2}$, which are not $E_0 = I$ or $\overline{E_0} = I + 1$. As a result the infinite loop alphabet $I\Lambda_7(E_0)$ is given by $\{\lambda_{E_0,E_1}, \lambda_{E_0,\overline{E_1}}, \lambda_{E_0,E_2}, \lambda_{E_0,\overline{E_2}}\}$. These paths induce the output alphabet $IL_7(I) = \left\{R^2L^{\frac{1}{2}}, R^2R^{\frac{1}{2}}, RLL^{\frac{1}{2}}, RLR^{\frac{1}{2}}\right\}$. See Fig. 5.6 (a). For the edge E_1, there are two paths we can take in the infinite loop alphabet $I\Lambda_7(E_1) = \{\lambda_{E_1,E_2}, \lambda_{E_1,\overline{E_2}}\}$. This induces the output alphabet $IL_7(E_1) =$

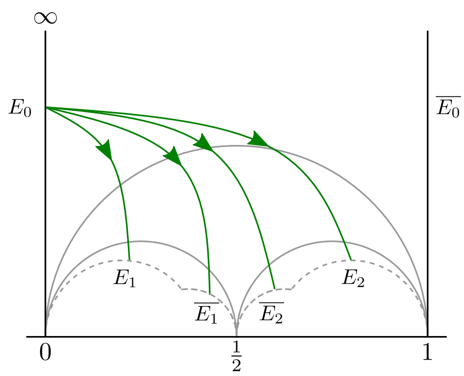

(a) An image of the paths in $I\Lambda_7(E_0)$.

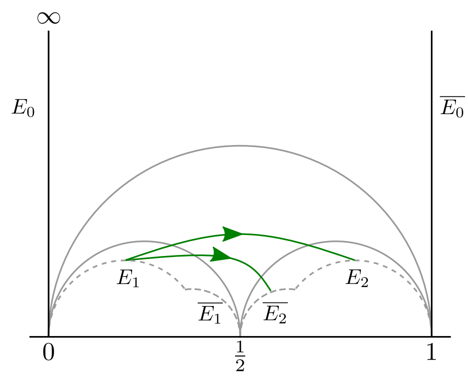

(b) An image of the path in $I\Lambda_7(E_1)$.

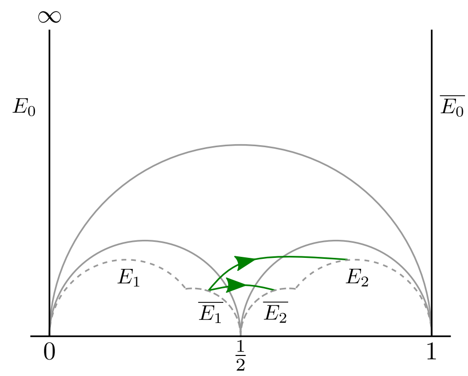

(c) An image of the path in $I\Lambda_7(\overline{E_1})$.

Figure 5.6: This figure shows the possible different infinite loop alphabets for $p = 7$, taken up to symmetry.

$\left\{L^{\frac{1}{2}}RL^{\frac{1}{2}}, L^{\frac{1}{2}}RR^{\frac{1}{2}}\right\}$. See Fig. 5.6 (b). Similarly, for the edge $\overline{E_1}$, there are two paths we can take in the infinite loop alphabet $I\Lambda_7(\overline{E_1}) = \{\lambda_{\overline{E_1},E_2}, \lambda_{\overline{E_1},\overline{E_2}}\}$. This induces the output alphabet $IL_7(\overline{E_1}) = \left\{R^{\frac{1}{2}}RL^{\frac{1}{2}}, R^{\frac{1}{2}}RR^{\frac{1}{2}}\right\}$. See Fig. 5.6 (c). Since E_2 is symmetric to E_1 under the reflection in the line $x = \frac{1}{2}$, we can use this symmetry to deduce the alphabet for $IL_7(E_2)$, by taking the alphabet $IL_7(E_1)$ and swapping the roles of L and R. Similarly, we can use the symmetry of $\overline{E_1}$ and $\overline{E_2}$ to deduce the alphabet $IL_7(\overline{E_2})$ from $IL_7(\overline{E_1})$.

Note that the edges E_1 and $\overline{E_1}$ are paired with each other and E_2 is paired with $\overline{E_2}$. Using this information, we can recover Table 5.4 and the automaton in Fig.5.7.

Initial State	Transition Function	Output Function
E_0	$\lambda_{E_0,\overline{E_1}} \to E_1$ $\lambda_{E_0,E_1} \to \overline{E_1}$ $\lambda_{E_0,\overline{E_2}} \to E_2$ $\lambda_{E_0,E_2} \to \overline{E_2}$	$R^2L^{\frac{1}{2}}$ $R^2R^{\frac{1}{2}}$ $RLR^{\frac{1}{2}}$ $RLL^{\frac{1}{2}}$
E_1	$\lambda_{E_1,\overline{E_2}} \to E_2$ $\lambda_{E_1,E_2} \to \overline{E_2}$	$L^{\frac{1}{2}}RR^{\frac{1}{2}}$ $L^{\frac{1}{2}}RL^{\frac{1}{2}}$
$\overline{E_1}$	$\lambda_{\overline{E_1},\overline{E_2}} \to E_2$ $\lambda_{\overline{E_1},E_2} \to \overline{E_2}$	$R^{\frac{1}{2}}RR^{\frac{1}{2}}$ $R^{\frac{1}{2}}RL^{\frac{1}{2}}$
E_2	$\lambda_{E_2,\overline{E_1}} \to E_1$ $\lambda_{E_2,E_1} \to \overline{E_1}$	$R^{\frac{1}{2}}LL^{\frac{1}{2}}$ $R^{\frac{1}{2}}LR^{\frac{1}{2}}$
$\overline{E_2}$	$\lambda_{\overline{E_2},\overline{E_1}} \to E_1$ $\lambda_{\overline{E_2},E_1} \to \overline{E_1}$	$L^{\frac{1}{2}}LL^{\frac{1}{2}}$ $L^{\frac{1}{2}}LR^{\frac{1}{2}}$

Table 5.4: A table of the transition functions and output functions used to construct infinite loops mod 7.

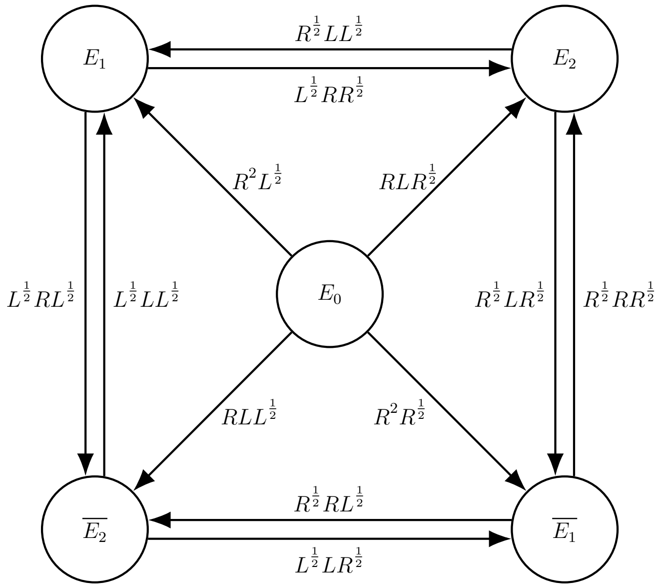

Figure 5.7: An automaton that constructs infinite loops mod 7. The initial state is E_0.

Constructing Infinite Loops mod 7: The Arithmetic Method

We start by first constructing the set $\textbf{Pairs}_7 := \{(1,0),(1,1),(1,2),(1,3),(1,4),$ $(2,1),(2,4),(2,6),(3,2),(3,5),(3,6),(4,1),(4,2),(4,5),(5,1),(5,4),(5,6),(6,3),$ $(6,4),(6,5)\}$. We note that excluding $(1,0)$ and $(1,1)$, the maximum value of $M(a,b) = 4$ and so 3 is the largest value for any partial quotient corresponding to these pairs. For $(1,0)$, we can have 1,2,4 and 5 as possible partial quotients and for $(1,1)$ we can have 1, 3 and 4. We then construct Table 5.5 to represent the algorithm.

Due to the large number of states, it is not practical to produce a diagram of the automaton, but the table can be used instead.

Next partial quotient a_{k+1}	States						
	$(1,0)$	$(1,1)$	$(1,2)$	$(1,3)$	$(1,4)$	$(2,1)$	$(2,4)$
1	$(1,1)$	$(2,1)$	x	$(4,1)$	$(5,1)$	$(3,2)$	$(6,2)$
2	$(2,1)$	x	$(4,1)$	$(5,1)$	x	x	$(1,2)$
3	x	$(4,1)$	$(5,1)$	x	x	x	$(3,2)$
4	$(4,1)$	$(5,1)$	x	x	x	x	x
5	$(5,1)$	x	x	x	x	x	x

Next partial quotient a_{k+1}	States						
	$(2,6)$	$(3,2)$	$(3,5)$	$(3,6)$	$(4,1)$	$(4,2)$	$(4,5)$
1	$(1,2)$	$(5,3)$	$(1,3)$	$(2,3)$	x	$(6,4)$	$(2,4)$
2	$(3,2)$	$(1,3)$	x	$(5,3)$	$(2,4)$	x	$(6,4)$
3	x	x	x	$(1,3)$	$(6,4)$	x	x

Next partial quotient a_{k+1}	States					
	$(5,1)$	$(5,4)$	$(5,6)$	$(6,3)$	$(6,4)$	$(6,5)$
1	$(6,5)$	x	$(4,5)$	$(2,6)$	$(3,6)$	x
2	$(4,5)$	$(6,5)$	x	x	$(2,6)$	$(3,6)$
3	x	$(4,5)$	x	x	x	$(2,6)$

Table 5.5: A table showing the possible values the next partial quotient can take to form an infinite loop mod 7.

Chapter 6

Conclusions and Future Research

The main aim of this thesis was to use the geometric setting of cutting sequences to better understand the behaviour of continued fractions under integer multiplication. Furthermore, we wished to use this setting to further investigate the p-adic Littlewood conjecture. From this perspective, the project was very successful. In particular, in Chapter 3 we outlined a method to construct an algorithm to multiply continued fractions by prime numbers and in Chapter 5 we reformulated pLC using infinite loops: a concept stemming from this geometric setting. In this chapter, we list the most important results and discuss their connection to the literature. We also discuss how this work could be built upon.

6.1 Main Results

The result that most encapsulates the work in this thesis is Theorem 4.1.8 - that we can view integer multiplication of continued fractions as being equivalent to triangulation replacement on some orbifold.

Theorem 4.1.8. *For every geodesic ray ζ_α in $\mathbb{H}$ starting at the y-axis with endpoint $\alpha > 0$, there is a canonical projection $\widehat{\zeta_\alpha}$ onto $\Gamma_0(n)\backslash\mathbb{H}$ such that $(\zeta_\alpha, \mathcal{F}) = (\widehat{\zeta_\alpha}, \widehat{\mathcal{F}})$, which is equivalent to the continued fraction expansion of α, and $(\zeta_\alpha, \frac{1}{n}\mathcal{F}) = (\widehat{\zeta_\alpha}, \widehat{\frac{1}{n}\mathcal{F}})$, which is equivalent to the continued fraction expansion of $n\alpha$.*

By viewing integer multiplication as triangulation replacement, we can learn a lot about the behaviour of continued fractions when they are multiplied. For example, Lemma 4.2.2 tells us that if (ζ, T) is an eventually recurrent cutting sequence on an orbifold $\mathcal{O}$, then (ζ, T') is also eventually recurrent, where T and T' are triangulations of $\mathcal{O}$. Theorem 4.1.8 then allows us to deduce that if $(\widehat{\zeta_\alpha}, \widehat{\mathcal{F}})$ is eventually recurrent, then so is $(\widehat{\zeta_\alpha}, \widehat{\frac{1}{n}\mathcal{F}})$. Since the cutting sequences $(\widehat{\zeta_\alpha}, \widehat{\mathcal{F}})$ and $(\widehat{\zeta_\alpha}, \widehat{\frac{1}{n}\mathcal{F}})$ naturally correspond to the continued fraction expansions of α and $n\alpha$ respectively, we can deduce that for every real number α with an eventually recurrent continued fraction expansion and every integer n, the continued fraction expansion of $n\alpha$ will also be eventually recurrent. This is a main step in proving the following corollary:

Corollary 4.2.6. *Let $\alpha \in \mathbb{R}$, let $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ be a non-trivial integer matrix (i.e. $a, b, c, d \in \mathbb{Z}$, $ad - bc \neq 0$), and let $\beta = M \cdot \alpha = \frac{a\alpha+b}{c\alpha+d}$. If the continued fraction expansion $\overline{\alpha}$ is eventually recurrent and $c\alpha + d \neq 0$, then the continued fraction $\overline{\beta}$ is eventually recurrent.*

Viewing integer multiplication of continued fractions in this way also draws attention to the concept of infinite loops. A geodesic ray ζ_α in $\mathbb{H}$ starting at the y-axis I and terminating at a point $\alpha \in \mathbb{R}$ is an infinite loop if it does not intersect the set of edges $\Gamma_0(n) \cdot I \setminus \{I + a : a \in \mathbb{Z}\}$. Equivalently, an *infinite loop mod n* is any real number with no semi-convergent denominators divisible by n - except for $q_{-1} = 0$. These are very important objects in relation to the p-adic Littlewood Conjecture, since whether α is an infinite loop mod n or not gives a lot of information about the bounds that we can infer from α and $n\alpha$. In particular, if α is not an infinite loop mod n, then at least one of the largest partial quotients of α and $n\alpha$ is large (relative to $\sqrt{n}$):

Lemma 5.1.10. *Assume that α is not an infinite loop mod n. Then we have:*

$$\max\{B(\alpha), B(n\alpha)\} \geq \lfloor 2\sqrt{n} \rfloor - 1.$$

Alternatively, if $p^\ell \alpha$ is an infinite loop mod p^m for all $\ell \in \mathbb{N}$, then the largest partial quotient of $p^\ell\alpha$ is bounded above:

Lemma 5.1.14. *Let $\alpha \in$ **Bad** and assume there exists an $m \in \mathbb{N}$ such that $p^\ell\alpha$ is an infinite loop mod p^m for all $\ell \in \mathbb{N}$. Then $B(p^\ell\alpha) < p^{m-4}$.*

The above lemmas show how viewing multiplication of continued fractions in this geometric way has a lot of power. In particular, this method allows us to nicely recreate the result of Badziahin, Bugeaud, Einsiedler and Kleinbock [BBEK15], namely:

Proposition 5.1.13. *Eventually recurrent continued fractions satisfy pLC.*

Furthermore, we can combine the above lemmas to reformulate pLC in terms of infinite loops.

Theorem 5.1.15. *Let $\alpha \in$ **Bad**. Then, α satisfies pLC if and only if there is a sequence of natural numbers $\{\ell_m\}_{m\in\mathbb{N}}$ such that $p^{\ell_m}\alpha$ is not an infinite loop mod p^m.*

The work presented in this thesis is quite a different way of looking at the p-adic Littlewood conjecture, compared to the rest of the literature. Despite this, not only were we able to recreate some of the previously known results - Proposition 5.1.13 - but we also found some new results: Corollary 4.2.6, Lemmas 5.1.10 and 5.1.14, and Theorem 5.1.15. This shows the power of the methods presented in this thesis, as well as the potential to provide even more information about pLC and the multiplication of continued fractions. In particular, there are several ways that we could build on the results of this thesis. In the next section, we discuss a few of the ways that this work could develop.

6.2 Future Research Aims

There are two main ways I see this work developing:

- To further investigate what this work can tell us about the mixed and p-adic Littlewood conjectures.
- To investigate how triangulation replacement affects cutting sequences in more depth.

6.2.1 Further Investigation of the p-adic Littlewood Conjecture

One of the main aims of the future of this research would be to find a larger class of real numbers which satisfy pLC. As discussed above, the techniques presented in this thesis were shown to be reasonably powerful. As such, the possibility of finding new solutions to pLC looks hopeful. Furthermore, it also seems likely that we can extend the reformulation of pLC to a reformulation of mLC.

Furthermore, I also plan to use the work in this thesis to investigate potential counter examples to pLC. The t-adic Littlewood conjecture - a problem which is analogous to pLC over function fields - was recently proven to be false by Adiceam, Nesharim and Lunnon in [ANL18]. This provides credence to the idea that pLC may also be false.

Finally, I would also like to improve the currently known upper bounds of pLC. Badziahin showed in [Bad16], that for all real numbers $\alpha \in \mathbb{R}$:

$$m_2(\alpha) < \frac{1}{9}.$$

I have already done some investigation of the upper bounds of pLC and have found that for all real numbers $\alpha \in \mathbb{R}$, we have:

$$m_2(\alpha) < \frac{1}{15}.$$

Finding Good Potential Counterexamples to pLC

As mentioned above, I plan to look at finding potential counterexamples to pLC. In particular, Theorem 5.1.15, Section 5.2.1 and previously known results gives us a good idea of where to start looking for potential counterexamples. Firstly, we would use Section 5.2.1 to produce a large number of partial quotients of a continued fraction expansion that is also an infinite loop mod p^k for some prime power p^k. Since we know certain properties of potential counterexamples to pLC - i.e. counterexamples do not limit to a periodic continued fraction [BDM07], counterexamples necessarily have non-recurrent continued fraction expansions [BBEK15], infinite loops mod p^k do not have partial quotients bigger than $p^k - 4$ - the infinite loops that we produce would also be built to satisfy these properties. For a potential counterexample α, we could then use the multiplication algorithm to calculate $p\alpha$ and check whether $p\alpha$ is an infinite loop mod p^k. If there is some large $L \in \mathbb{N}$ such that $p^l\alpha$ is an infinite loop mod p^k for all $l \leq L$, then this would indicate that α is a "good" potential counterexample to pLC.

It is worth noting that such a process would not provide proof that pLC is false. However, being able to find such a result would give us a lot of heuristic information about counterexamples to pLC, and this may be useful for finding other methods and techniques that could prove pLC to be false.

Improving Upper Bounds on the Infimum of the p-adic Littlewood Conjecture

The p-adic Littlewood Conjecture has shown to be very difficult to solve directly. As a result, we can instead ask if we can find upper bounds for $m_p(\alpha) := \inf_{q\geq 1}\{q \cdot |q|_p \cdot \|q\alpha\|\}$ for all $\alpha \in \mathbb{R}$ and a fixed p? To formalise this question we introduce the function $m_{PLC}(p)$:

$$m_{PLC}(p) := \sup_{\alpha\in\mathbb{R}}\{\inf_{q\geq 1}\{q \cdot |q|_p \cdot \|q\alpha\|\}\}.$$

By Hurwitz's Theorem - see Section 2.1.2 - we know that:

$$\inf_{q\geq 1}\{q\cdot\|q\alpha\|\} < \frac{1}{\sqrt{5}}.$$

Furthermore, $|q|_p \leq 1$ for all $p, q \in \mathbb{N}$. Combining these results together, we can guarantee that for all primes p we have:

$$\begin{aligned} m_{PLC}(p) &= \sup_{\alpha\in\mathbb{R}}\{\inf_{q\geq 1}\{q\cdot|q|_p\cdot\|q\alpha\|\}\} \\ &< \sup_{\alpha\in\mathbb{R}}\{\inf_{q\geq 1}\{q\cdot\|q\alpha\|\}\} \\ &< \sup_{\alpha\in\mathbb{R}}\left\{\frac{1}{\sqrt{5}}\right\} \\ &< \frac{1}{\sqrt{5}}. \end{aligned}$$

However, since this initial bound is achieved by "forgetting" the p-adic norm, it does not seem to be very optimal. Indeed, the work of Badziahin in [Bad16] shows that for $p = 2$, we have:

$$m_{PLC}(2) < \frac{1}{9}.$$

This result was achieved by using an algorithm that was computer implemented. We should note that the algorithm required only 3 seconds to give a bound of $m_{PLC}(p) < \frac{1}{9}$. However, when the algorithm was implemented to try to obtain a bound of $m_{PLC}(p) < \frac{1}{10}$, the process ran for over 60 hours without a conclusive result. Badziahin states in this paper [Bad16], that such a large increase in time complexity may be indicative of the existence of a counterexample to pLC.

The work in this thesis provides another way to try and compute upper bounds for $m_{PLC}(p)$. In particular, we can use Corollary 5.1.11 to prove the following proposition:

Proposition 6.2.1. *If for a fixed prime power p^m and every $\alpha \in \mathbb{R}$ there is some $k \in \mathbb{N}\cup\{0\}$ such that $p^k\alpha$ is not an infinite loop mod p^m, then we have:*

$$m_{PLC}(p) \leq \frac{1}{\lfloor 2\sqrt{p^m}\rfloor - 1}.$$

In particular, if we can show that the set of infinite loops mod p^m is empty under finite multiplication by p, then the result holds.

This general idea has been the basis for my own computer based algorithm. Using this algorithm and complementary techniques, I was able to improve on Badziahin's upper bound, and find other values of $m_{PLC}(p)$ as shown in the following table:

Prime p	2	3	5	7	11	13	17	19
$m_{PLC}(p)$	$< \frac{1}{15}$	$< \frac{1}{9}$	$< \frac{1}{9}$	$< \frac{1}{4}$	$< \frac{1}{5}$	$< \frac{1}{6}$	$< \frac{1}{7}$	$< \frac{1}{7}$

Table 6.1: A table of upper bounds on $m_{PLC}(p)$.

These results are still somewhat preliminary and further computation may improve the bounds that we receive.

6.2.2 Further Investigation of Triangulation Replacement

I would also like to investigate how triangulation replacement affects cutting sequences in a more general setting. As seen in Chapter 4, given a specific orbifold $\mathcal{O}$, a specific geodesic ray and two specific triangulations of $\mathcal{O}$, triangulation replacement induces integer multiplication on the level of cutting sequences. However, one can also naturally ask "How do other triangulation replacements on other orbifolds affect the cutting sequence (and it's corresponding continued fraction)?" and "Is there some natural way to describe the transformations on the continued fraction expansions which are induced by these triangulation replacements?"

If one were able to find the set of transformations $\mathcal{T}$ which are realised by triangulation replacement, then, since recurrency and periodicity of the cutting sequence are both tied to geometric properties that are not altered by triangulation replacement, this would imply that recurrent and periodic continued fractions are also preserved by this set of transformations $\mathcal{T}$.

Alternatively, it would also be interesting to investigate whether other properties of cutting sequences are preserved by triangulation replacement. One interesting property to look at would be normality. Does normality correspond naturally to some geometric property? Is normality preserved under triangulation replacement? These questions are extremely pertinent following Vandehey's proof that normality of the continued fraction expansion of a real number x is preserved under the transformation $\frac{ax+b}{cx+d}$, for $\left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right) \in M$ [Van17]. It would be interesting to investigate whether this was due to some geometric property, and if so, whether it is possible to strengthen the result to the larger class of transformations $\mathcal{T}$, as described above.

Bibliography

[ANL18] F. Adiceam, E. Nesharim and F. Lunnon. 'On the t-adic Littlewood Conjecture'. Version 1. In: (12th June 2018). arXiv:1806.04478.

[Art24] E. Artin. 'Ein Mechanisches System mit quasi-ergodischen Bahnen'. In: *Abh. Math. Semin. Univ. Hambg.* vol.3, no. 1. (1924), pp. 170–175.

[AS03] A.P. Allouche and J. Shallit. *Automatic Sequences: Theory, Applications, Generalisations.* Cambridge University Press, 2003.

[Bad15] Dzmitry Badziahin. 'On continued fraction expansion of potential counterexamples to p-adic Littlewood conjecture'. Version 2. In: (2015). arXiv:1406.3594.

[Bad16] Dzmitry Badziahin. 'Computation of the Infimum in the Littlewood Conjecture'. In: *Experimental Mathematics* 25.1 (2016), pp. 100–105.

[BBEK15] D. Badziahin, Y. Bugeaud, M. Einsiedler and D. Kleinbock. 'On the complexity of a putative counterexample to the p-adic Littlewood conjecture'. In: *Compos. Math.* vol. 151, no. 9. (2015), pp. 1647–1662.

[BDM07] Y. Bugeaud, M. Drmota and B. de Mathan. 'On a mixed Littlewood conjecture in Diophantine approximation'. In: *Acta Arith.* vol. 128. (2007), pp. 107–124.

[BHV11] Y. Bugeaud, A. Hayes and S. Velani. 'Metric Considerations Concerning The Mixed Littlewood Conjecture'. In: *Int. J. Number Theory* vol.7, no. 3. (2011), pp. 593–609.

[Bre80] C. Brezinski. *History of Continued Fractions and Pade Approximants.* Springer-Verlag: New York, 1980.

[Bug14] Y. Bugeaud. 'Around the Littlewood conjecture in Diophantine approximation'. In: *Publ. Math. Besançon* vol. 1 (2014), pp. 5–18.

[Bur00] E. B. Burger. *Exploring The Number Jungle: A Journey Into Diophantine Analysis.* American Mathematical Society, 2000.

[BV11] D. Badziahin and S. Velani. 'Multiplicatively badly approximable numbers and generalised Cantor sets'. In: *Adv. Math.* vol. 228, no. 5. (2011), pp. 2766–2796.

[CSD55] J.W.S. Cassels and H.P.F. Swinnerton-Dyer. 'On the product of three homogenous linear forms and indefinite ternary quadratic forms'. In: *Philos. Trans. Roy. Soc. London, Ser. A* vol. 248, no. 940. (1955), pp. 73–96.

[EK07] M. Einsiedler and D. Kleinbock. 'Measure rigidity and p-adic Littlewood-type problems'. In: *Compos. Math.* vol. 143, no. 3. (2007), pp. 689–702.

[EKL06] M. Einsiedler, A. Katok and E. Lindenstrauss. 'Invariant measures and the set of exceptions to Littlewood's conjecture'. In: *Ann. of Math.* vol. 164, no. 2. (2006), pp. 513–560.

[FM11] B. Farb and D. Margalit. *A Primer on Mapping Class Groups.* Princeton University Press, 2011.

[Hal47] M. Hall. 'On the Sum and Product of Continued Fractions'. In: *Ann. of Math.* vol. 48, no. 4. (1947), pp. 966–993.

[Hat07] R.A. Hatch. 'Theon of Smyrna'. In: *The Biographical Encyclopedia of Astronomers.* Ed. by T. Hockey, V. Trimble, T.R. Williams, K. Bracher, R.A. Jarrell, J.D. Marché, F.J. Ragep, J. Palmeri and M. Bolt. Springer, New York, 2007, pp. 73–96.

[Hat91] A. Hatcher. ‘On Triangulations of Surfaces’. In: *Topology Appl.* vol. 40, no. 2. (1991), pp. 189–194.

[HH13] S. Harrap and A. Haynes. ‘The mixed Littlewood conjecture for pseudo-absolute values’. In: *Math. Ann.* vol. 357, no. 3. (2013), pp. 941–960.

[Hum16] M. G. Humbert. ‘Sur les fractions continues et les formes quadratiques binaires indéfinies’. In: *J. Math. Pures Appl.* vol. 7, no. 2. (1916), pp. 104–154.

[HW38] G.H. Hardy and E.M. Wright. *An Introduction to the Theory of Numbers*. 6th Ed. Oxford University Press, 1938.

[Khi63] A. Ya Khinchin. *Continued Fractions*. Dover Publications Inc., 1963.

[KL07] C. A. Kurth and L. Long. *Computations with finite index subgroups of $PSL_2(\mathbb{Z})$ using Farey Symbols*. 2007. `arXiv:0710.1835`.

[Kul91] R. S. Kulkarni. ‘An Arithmetic-Geometric Method in the Study of the Subgroups of the Modular Group’. In: *Amer. J. Math.* vol. 113, no. 6. (1991), pp. 1053–1133.

[LS98] P. Liardet and P. Stambul. ‘Algebraic Computations with Continued Fractions’. In: *J. Number Theory* vol. 73, no. 1. (1998), pp. 92 –121.

[MH38] M. Morse and G. A. Hedlund. ‘Symbolic Dynamics’. In: *Amer. J. Math.* vol. 60, no. 4 (1938), pp. 815–866.

[Moe82] R. Moeckel. ‘Geodesics on modular surfaces and continued fractions’. In: *Ergodic Theory Dynam. Systems* vol. 2, no. 1. (1982), pp. 69–83.

[Mon79] H. L. Montgomery. ‘John Edensor Littlewood: Littlewood’s Work in Number Theory’. In: *Bull. Lond. Math. Soc.* vol. 11, no. 1 (1979), pp. 78–86.

[MT04] B. de Mathan and O. Teulié. 'Problémes Diophantiens simultanés'. In: *Monatsh. Math.* vol. 143, no. 3. (2004), pp. 229–245.

[Old63] C.D. Olds. *Continued Fractions*. Random House:New York, 1963.

[PV00] A.D. Pollington and S.J.L. Velani. 'On a problem in simultaneous Diophantine approximation: Littlewood's conjecture'. In: *Acta Math.* vol. 185, no. 2. (2000), pp. 287–306.

[Ran73] G. Raney. 'On Continued Fractions and Finite Automata'. In: *Math. Ann.* vol. 206, no. 4. (1973), pp. 265–283.

[Ser13] C. Series. *Hyperbolic geometry: MA 448*. 2013. `https://homepages.warwick.ac.uk/~masbb/Papers/MA448.pdf`.

[Ser85a] C. Series. 'The Geometry of Markoff Numbers'. In: *Math. Intelligencer* vol. 7, no. 3. (1985), pp. 20–29.

[Ser85b] C. Series. 'The Modular Surface and Continued Fractions'. In: *J. Lond. Math. Soc. (2)* vol. 31, no. 1. (1985), pp. 69–80.

[Spe42] D.C. Spencer. 'The Lattice Points of Tetrahedra'. In: *J. Math. Phys. Mas. Inst. Tech.* vol. 21. (1942), pp. 189–197.

[Van17] J. Vandehey. 'Non-Trivial Matrix Actions Preserve Normality for Continued Fractions'. In: *Compos. Math.* vol. 153, no. 2. (2017), pp. 274–293.