

Thurston’s theorem: entropy in dimension one

RYAN DICKMANN, GEORGE DOMAT, THOMAS HILL, SANGHOON KWAK,
CARLOS OSPINA, PRIYAM PATEL, AND REBECCA RECHKIN

In his paper [13], Thurston shows that a positive real number h is the topological entropy for an ergodic traintrack representative of an outer automorphism of a free group if and only if its expansion constant $\lambda = e^h$ is a weak Perron number. This is a powerful result, answering a question analogous to one regarding surfaces and stretch factors of pseudo-Anosov homeomorphisms. However, much of the machinery used to prove this seminal theorem on traintrack maps is contained in the part of Thurston’s paper on the entropy of postcritically finite interval maps and the proof is difficult to parse. In this expository paper, we modernize Thurston’s approach, fill in gaps in the original paper, and distill Thurston’s methods to give a cohesive proof of the traintrack theorem. Of particular note is the addition of a proof of ergodicity of the traintrack representatives, which was missing in Thurston’s paper.

1 Introduction	128
2 Background	131
3 Star Maps	140
4 Prototype Graph	152
5 Splitting graphs	157
6 Conclusion: Proof of Thurston’s Theorem	166
Acknowledgements	172
References	172

1. Introduction

Topological entropy describes the complexity of the orbit structure of a dynamical system, and is an invariant of topological conjugacy classes. A classical problem in dynamics is to characterize the numbers that can arise as the topological entropy for a particular family of dynamical systems. In his paper [13], Thurston proved that a positive real number h is the topological entropy of a postcritically finite self-map of the unit interval if and only if $\lambda = e^h$ is a weak Perron number, i.e., λ is an algebraic integer that is at least as large as the absolute value of any conjugate of λ . He uses the tools developed for studying interval maps to then prove the following theorem about outer automorphisms of free groups, answering a prominent question in geometric group theory.

Theorem 1.1 ([13, Theorem 1.9]). *A positive real number h is the topological entropy for an ergodic traintrack representative of an outer automorphism of a free group if and only if $\lambda = e^h$ is an algebraic integer that is at least as large as the absolute value of any conjugate of λ , i.e., λ is a weak Perron number.*

Unfortunately, Thurston fell ill while writing a draft of the manuscript [13], and there are some gaps in the final version of the paper. In this expository article, we modernize Thurston’s approach, fill in gaps in the original paper, and distill Thurston’s methods to give a cohesive proof of the traintrack theorem that is especially readable for geometric group theorists. Though the motivation for the proof of Theorem 1.1 comes from the dynamics of postcritically finite interval maps, we exclude details about such maps below since the purpose of this paper is to give a complete and concise proof of Thurston’s traintrack theorem.

To understand the content of Theorem 1.1, recall that for a finitely generated free group \mathbb{F}_n , the **outer automorphism group** of \mathbb{F}_n is $\text{Out}(\mathbb{F}_n) = \text{Aut}(\mathbb{F}_n)/\text{Inn}(\mathbb{F}_n)$. Bestvina and Handel [3, Theorem 1.7] showed that certain outer automorphism of a free group can be represented by a special map between graphs called a **traintrack map** (see [4, Chapter 6.3]). Roughly, a traintrack map is a continuous graph map which has particularly nice cancellation properties with respect to iterations. A traintrack map $f : \Gamma \rightarrow \Gamma$ is called *irreducible* if f does not admit an invariant proper subgraph which is not a tree. We formally define and give all relevant background on traintrack maps in Section 2.5. The notion of irreducibility is compatible with that of ergodicity of the traintrack map. In particular, the map being ergodic

as a dynamical system essentially means the system cannot be reduced or factored into smaller components, which in this case would be proper subgraphs. Thus, ergodicity of the traintrack map implies irreducibility.

Since their introduction, traintrack maps have become a standard tool for understanding the geometry and dynamics of automorphisms of free groups. It is also easy to calculate the expansion constant of a traintrack map since traintrack maps eliminate backtracking when iteratively applying the map to an edge or edge path. In fact, the expansion constant of the traintrack map f can be calculated by finding the Perron-Frobenius eigenvalue, λ , of the transition matrix for f . The topological entropy of f is then exactly $\log(\lambda)$. See [Section 2.2], [Section 2.3], and [3], Remark 1.8] for more details.

Thus, one direction of [Theorem 1.1] follows almost immediately from [Theorem 2.3] (the Perron-Frobenius Theorem). In particular, an ergodic traintrack representative of an outer automorphism of a free group has a transition matrix with a positive leading eigenvalue λ that is not smaller than the magnitude of the other eigenvalues (see the last paragraph of [Section 2.3]). Therefore, the entropy of the traintrack map is $h(f) = \log(\lambda)$ and $e^h = \lambda$ is a weak Perron number as desired. Proving the other direction is far more difficult, but we provide a sketch of Thurston's (and our) argument here.

Fix a Perron number λ , i.e. an algebraic integer that is *strictly* larger than the absolute value of its conjugates. Thurston uses two main ingredients to construct a traintrack map with growth rate λ . First, he defines a collection of prototype traintrack maps ϕ_n for all odd, positive integers n . Second, he defines a star map on a star graph that is *uniformly λ -expanding*, called f_λ below. (For the definition of star graphs and star maps, see [Section 3.1]; these are called *asterisk* maps on *asterisk* graphs in [13].) This is a delicate process that requires an understanding of the arithmetic of Perron numbers. In fact, there is a somewhat serious number theoretic error in the version of [Lemma 2.15] (The Even Lemma) that appears in Thurston's paper [13], p.359, proof of Theorem 6.2]. We correct this error and correct the construction of the star maps accordingly (see [Remark 2.16] and [Section 3.3] respectively). It was noted by the referee that Richard Webb and his reading group at Cambridge were also aware of the error and how to fix the argument.

Next, given the star map $f_\lambda : \Gamma \rightarrow \Gamma$, Thurston defines the split graph $S(\Gamma)$ and the split map $S(f_\lambda) : S(\Gamma) \rightarrow S(\Gamma)$, which are defined using *both* the prototype traintrack maps ϕ_n and the λ -expanding star map f_λ . The desired traintrack map for [Theorem 1.1] is $S(f_\lambda)$, and the definition of split

maps ensures that the expansion constant of $S(f_\lambda)$ is λ because the expansion constant of f_λ is λ . It then remains to show that the traintrack map actually represents an outer automorphism of a free group, and that the map $S(f_\lambda)$ is in fact ergodic. Unfortunately, the details on these two points are sparse in Thurston's paper and we remedy this as described in the next paragraph. Finally, to extend the results to *weak* Perron numbers, Thurston uses the fact that for a weak Perron number λ , there exists $N \in \mathbb{Z}^+$ such that λ^N is Perron.

The main contributions of this expository paper are as follows. First and foremost, we distill and streamline all of the pieces needed for Thurston's traintrack theorem from [13], which contains a variety of additional theorems regarding interval maps and other topics. Second, we correct the number theoretic error in the Even Lemma mentioned above. Third, we add a proof of the ergodicity of the traintrack maps, which was completely missing in Thurston's paper. Finally, we use Stallings folds to thoroughly prove that the traintrack maps represent elements of $\text{Out}(\mathbb{F}_n)$. For this last piece, Thurston outlines an algebraic argument, but provides only a few details. We found the argument using Stallings folds more straightforward and rigorous.

We conclude with some remarks. First, the notion of traintrack maps for $\text{Out}(\mathbb{F}_n)$ is motivated by Thurston's traintracks on surfaces. In particular, elements of $\text{Out}(\mathbb{F}_n)$ that admit traintrack representatives are the analog of pseudo-Anosov homeomorphisms of surfaces, and expansion constants for these maps are the analog of stretch factors for pseudo-Anosovs. Despite the fact that traintrack theory for free groups is more complicated than for surfaces, Thurston gave a complete answer for which algebraic integers can arise as expansion constants for traintrack representatives of outer automorphisms. Thurston also showed that, for surfaces, every stretch factor of a pseudo-Anosov homeomorphism is an algebraic unit, but it is still unknown exactly which units can appear as stretch factors. There has been partial progress in this direction. For example, Fried [5, Theorem 1] proved every stretch factor λ of a pseudo-Anosov homeomorphism is a bi-Perron unit. Fried further conjectured [5, Problem 2] that every bi-Perron unit has some power such that it is realized as a stretch factor of a pseudo-Anosov homeomorphism. The work of Pankau [11] and Liechti–Pankau [6] made some progress towards the conjecture, but the question is still far from resolved.

In the surface case, the stretch factors for a pseudo-Anosov homeomorphism and its inverse are always equal. This is no longer the case for outer automorphisms of \mathbb{F}_n . See [3, Remark p.9] for an example of this. In [13],

Thurston asks the question: Which pairs of numbers can appear as the expansion constants of an outer automorphism and its inverse? He gives a partial answer when one restricts to the class of *bipositive* outer automorphisms, but he notes that his result cannot generalize to all outer automorphisms. He also conjectures that every pair of weak Perron numbers greater than 1 is a pair of expansion constants for an outer automorphism and its inverse. We are not aware of a proof of this claim in the literature as of this moment, but a proof of this claim would be a good first step towards answering Thurston's question. For the sake of brevity, we do not discuss this portion of Thurston's paper and direct interested readers to [13, Section 11 & 12].

1.1. Outline of paper

In [Section 2](#) we provide all relevant background on entropy, λ -expanding graph maps, traintrack structures, Stallings folds, and Perron numbers, in that order. In [Section 3](#), we define star maps and use the geometry of Perron numbers to construct uniformly λ -expanding star maps for all Perron and weak Perron numbers λ . In [Section 4](#), we define Thurston's prototype traintrack maps ϕ_n and prove that they are indeed homotopy equivalences so that they represent elements of $\text{Out}(\mathbb{F}_n)$. Then, in [Section 5](#), we define split graphs, split maps, and prove that our split star maps are traintrack representatives of elements of $\text{Out}(\mathbb{F}_n)$. Finally, in [Section 6](#), we prove that the split star maps are ergodic and finish the proof of [Theorem 1.1](#).

2. Background

Thurston uses a variety of tools from dynamics, geometric group theory, and algebraic number theory throughout his proof of [Theorem 1.1](#). All relevant background material on these three topics is given in this section.

2.1. Entropy

The value of the topological entropy describes the complexity of the orbit structure of a dynamical system, and is an invariant of topological conjugacy classes. Let (X, d) be a compact metric space, and let $f : X \rightarrow X$ be a continuous map. For $n \geq 1$, set

$$d_n(x, y) := \max_{0 \leq j \leq n-1} d(f^j(x), f^j(y)).$$

Denote by $B_\varepsilon^n(x)$ the open ε -ball $\{y \in X \mid d_n(x, y) < \varepsilon\}$ with respect to d_n .

Since X is a compact space, for each $\varepsilon > 0$, X can be covered by a finite collection of open sets of the form $B_\varepsilon^n(x_i)$ for $x_i \in X$. Then, let $N(\varepsilon, n)$ be the minimum number of such open sets that cover X .

Definition 2.1. Let $f : X \rightarrow X$ be a continuous map on a compact metric space X . The **topological entropy** $h(f)$ is

$$h(f) = \lim_{\varepsilon \rightarrow 0} \limsup_{n \rightarrow \infty} \frac{1}{n} \log(N(\varepsilon, n)).$$

2.2. Graph Maps and Entropy

Now, we turn our attention to the category of *graphs*. A **graph** Γ is a 1-dimensional CW complex, whose 0-simplices are called the **vertices**, and whose 1-simplices are called the **edges**. We will always assume our graphs have finitely many vertices and edges. If a graph is given an orientation on each of the edges (a choice of left and right endpoints), then it is called an **directed graph**.

A **graph map** $f : \Gamma_1 \rightarrow \Gamma_2$ between graphs is a continuous map that sends vertices to vertices and edges to edge paths.

Endowing the edges of a graph with lengths produces a **metric graph**. With a metric, the length of any edge path can be measured as the sum of the lengths of edges in the path. Namely, this induces a **length function** ℓ on the set of edge paths in the graph. A typical choice of such a metric is the **combinatorial metric**: the metric that gives every edge length 1.

For a given metric graph Γ with a length function ℓ , we define the **total variation** of a graph map $f : \Gamma \rightarrow \Gamma$ as:

$$\text{Var}(f) = \sum_{e \in E(\Gamma)} \ell(f(e)),$$

where $E(\Gamma)$ denotes the set of edges of Γ .

Computing the topological entropy for a graph map is simple due to the 1-dimensionality of the graph, and we will use the following theorem to compute the entropy of graph maps throughout the paper.

Theorem 2.2 ([1]). *Let $f : \Gamma \rightarrow \Gamma$ be a graph map on a finite metric graph that has finitely many points at which f is not a local homeomorphism. Then*

$$h(f) = \lim_{n \rightarrow \infty} \frac{1}{n} \log(\text{Var}(f^n)).$$

2.3. Perron-Frobenius Theorem

Thurston's argument uses the following form of the well-known Perron-Frobenius theorem. We say a matrix is nonnegative (or positive) when its entries are nonnegative (or positive, respectively). A nonnegative matrix is said to be **ergodic** if the sum of a finite number of its consecutive positive powers is a positive matrix. A nonnegative matrix is said to be **mixing** if some power is a positive matrix. These two notions describe that the dynamics on a space will not decompose into subspaces that have independent dynamics. Mixing is stronger than ergodicity; a mixing map exhibits more uniform orbit dynamics. We note that Thurston uses the definition of ergodicity above.

However, in the literature, this definition of ergodicity is commonly referred to as *irreducible*. What we defined as mixing is also commonly referred to as *primitive*, see for instance [7], definitions 4.2.2 and 4.5.7.

Theorem 2.3 (Perron-Frobenius). *Let M be an $n \times n$ matrix with integer entries. If M is nonnegative, then M has at least one eigenvector such that*

(i) *The corresponding eigenvalue λ is nonnegative.*

(ii) *$\lambda \geq |\lambda_i|$ for all the other eigenvalues λ_i .*

Furthermore, if M is ergodic, there is a unique eigenvector whose corresponding eigenvalue λ is strictly positive.

To use the Perron-Frobenius theorem in our setting, we relate a (self) graph map f with the *transition matrix* M , defined as follows. Enumerate the edges of Γ by positive integers, say $1, \dots, n$. Then the **transition matrix** of $f : \Gamma \rightarrow \Gamma$ is an $n \times n$ integer matrix, whose (i, j) entry is determined by the number of times the i -th edges appear in the edge path $f(j)$. In light of the definitions of ergodic and mixing matrices, we call a graph map **ergodic** (or **mixing**) if its transition matrix is ergodic (or mixing, respectively).

Perron-Frobenius implies that an ergodic graph map $f : \Gamma \rightarrow \Gamma$ has a transition matrix with a *positive* leading eigenvalue λ that is no smaller than the magnitude of the other eigenvalues. Therefore, $\text{Var}(f^n) \sim \lambda^n$ and it follows that $h(f) = \lim_{n \rightarrow \infty} \frac{1}{n} \log(\text{Var}(f^n)) = \log(\lambda)$ by [Theorem 2.2](#). This proves the forward direction of [Theorem 1.1](#) since, as we will discuss in [Section 2.5](#), a traintrack representative of an outer automorphism of a free group is a special case of a graph map.

2.4. Uniform λ -Expanders

Definition 2.4. Let $\lambda > 0$ be an algebraic integer. Let Γ be a metric graph, and ℓ be a length function on the edge paths in Γ . We say f is **uniformly λ -expanding** if every edge is scaled by the same factor λ ; namely, $\ell(f(e)) = \lambda\ell(e)$ for every edge e of Γ .

Throughout the paper, we construct many maps of this form while working towards the main theorem due to the following useful proposition.

Proposition 2.5. *Let λ be an algebraic integer and f be a uniformly λ -expanding graph map on a finite metric graph. Then $h(f) = \log \lambda$.*

Proof. Let $f : \Gamma \rightarrow \Gamma$ and ℓ denote the length function on Γ . Observe that $\text{Var}(f) = \lambda\ell(\Gamma)$, and more generally $\text{Var}(f^n) = \lambda^n\ell(\Gamma)$. Using [Theorem 2.2](#), we compute

$$h(f) = \lim_{n \rightarrow \infty} \frac{1}{n} \log(\text{Var}(f^n)) = \log(\lambda) + \lim_{n \rightarrow \infty} \frac{1}{n} \log(\ell(\Gamma)) = \log(\lambda).$$

□

2.5. Traintrack Structures

Now, we will describe the necessary background required to understand the statement of [Theorem 1.1](#). The celebrated work of Bestvina and Handel [\[3\]](#) gives a geometric approach to understanding $\text{Out}(\mathbb{F}_n)$, the outer automorphism group of a finitely generated free group \mathbb{F}_n . See [\[2, 4\]](#) and for a general introduction to $\text{Out}(\mathbb{F}_n)$ and traintrack maps. Recall

$$\text{Out}(G) = \text{Aut}(G)/\text{Inn}(G),$$

where $\text{Aut}(G)$ is the group of automorphisms and $\text{Inn}(G)$ is the group of inner automorphisms of G , i.e., the subgroup of $\text{Aut}(G)$ consisting of automorphisms of the type $\varphi : x \mapsto gxg^{-1}$ for $g \in G$.

Bestvina and Handel [\[3, Theorem 1.7\]](#) showed that certain outer automorphism can be represented by a special map between graphs called a traintrack map (see also [\[4, Chapter 6.3\]](#)), which we now define.

A graph map is **taut** if it restricts to a local embedding on the interior of each edge. The term taut was chosen to represent the fact that all backtracking in the graph map has been removed. Indeed, any graph map is homotopic to a taut graph map.

A self-graph map $f : \Gamma \rightarrow \Gamma$ that is also a homotopy equivalence induces an automorphism of the fundamental group, which is well-defined up to postcomposing an inner automorphism depending on the choice of path connecting base points. The fundamental group of a graph is \mathbb{F}_n for some n , so the map f corresponds to an element $\psi \in \text{Out}(\mathbb{F}_n)$. The notion of tautness of a graph map is analogous to cyclically reduced words in free groups.

The standard topology on the intervals descends to a topology on the graph. Consider a small closed neighborhood of a vertex in a directed graph, which we can view as a directed graph itself. The oriented edges in this neighborhood are called **directions**. Formally, a **turn** in the graph is a 2-element subset of directions from a single vertex. More intuitively, we can think of a turn as a segment of a path passing through the vertex. In an oriented graph, we can refer to a turn with a 2-letter long word, where an uppercase letter represents traversing an edge backward. For example, the turn bA will refer to the turn a path makes after first traveling along the edge b in the forward direction followed by a in the reverse direction. Note that reversing the order of the word and swapping the case of the letters gives the same turn, but now the path representing the turn is traveling in the opposite direction; e.g., aB is the same turn as bA . In either case, we say that the path takes the given turn. In general, a locally embedded path takes a turn at every vertex it crosses, and the turn is determined by the incoming and outgoing edges the path takes at a given vertex.

Now we partition the set of turns into two collections, a set of legal turns and a set of illegal turns. We will always assume that every *backtracking*, a turn of the form xX , is illegal. Such a partition is called a **traintrack structure** on Γ . A path is **legal** if it is a local embedding, and it takes a legal turn at each vertex on the path. A path is **illegal** if it is not legal. Note a graph map sends a turn to another turn, so we have the following definition.

Definition 2.6. A graph map $f : \Gamma \rightarrow \Gamma$ is a **traintrack map** when it is taut and there exists a traintrack structure on Γ such that

- (i) legal turns are sent to legal turns, and
- (ii) every edge is sent to a legal path.

The conditions (i) and (ii) together imply that legal paths are sent to legal paths under a traintrack map. Thurston defines a traintrack map to be a map of a graph such that all iterates are local embeddings on each edge. The fact that legal paths are sent to legal paths shows that our definition

implies Thurston's. A traintrack map $f : \Gamma \rightarrow \Gamma$ is called **irreducible** if f does not admit an invariant proper subgraph that is not a tree.




Figure 1: Left: A graph with labeled oriented edges. Right: A traintrack structure on the graph with legal turns da , dA , Dc , DB , bc , and cb .

A traintrack structure on a graph is often graphically represented by blowing up each vertex to a disk. Then, for each legal turn, we draw a smooth path within the corresponding disk, which connects the endpoints of edges from the legal turn. The legal turns then correspond to turns which an actual train could make traveling along a track modeled after the picture, i.e., avoiding sharp turns. See [Figure 1](#) for an example. In the given traintrack structure from [Figure 1](#), no legal path can travel along the a edge in either direction twice in succession (the turn aa is illegal), nor can it travel along d immediately after traveling along c in the forward direction (the turn cd is illegal), etc.

2.6. Stallings Folds

In [Section 4](#) and [Section 5](#), we will make use of Stallings folds, first introduced in [\[12\]](#), to verify that our graph maps are homotopy equivalences. Note that this differs from Thurston's approach in the original proof of [Lemma 4.1](#). We felt that using Stallings folds was more intuitive for a rigorous proof.

Definition 2.7. A **morphism** of graphs is a continuous map of graphs that sends vertices to vertices and edges to edges. An **immersion** of graphs is a morphism that is locally injective.

Note that a morphism of graphs is different from a graph map since we require only that edges be sent to *edge-paths* for graph maps. Beginning with a graph map, we can subdivide the edges of the domain graph at the complete pre-image of the vertices of the codomain graph in order to obtain a graph morphism in the style of Stallings. Note that the property of a morphism being an immersion needs only to be checked at the vertices of the domain.

Definition 2.8. Let Γ be a graph and x_1, x_2 be two edges of Γ sharing a vertex. Let $\Gamma' = \Gamma/(x_1 \sim x_2)$. A **fold** is the natural quotient morphism $\Gamma \rightarrow \Gamma'$.

Folds come in two flavors depending on whether the two edges x_1, x_2 share a single vertex or share both vertices. Folds between edges that share only a single vertex are called **Type I** folds and are homotopy equivalences. Folds between edges that share both vertices are **Type II** folds and fail to be homotopy equivalences (they reduce the rank of the fundamental group by one). See [Figure 2](#) for examples of Type I and Type II folds.




Figure 2: Examples of Type I and Type II folds.

Theorem 2.9 ([12]). *Let $\phi : \Gamma \rightarrow \Delta$ be a graph morphism. Then ϕ factors as*

$$\Gamma = \Gamma_0 \xrightarrow{F_1} \Gamma_1 \xrightarrow{F_2} \dots \xrightarrow{F_n} \Gamma_n \xrightarrow{\psi} \Delta,$$

where each of the F_i are folds and ψ is an immersion.

This theorem gives an algorithm for verifying that a given graph map is a homotopy equivalence: First, one subdivides the domain in order to obtain a graph morphism. Then, one performs all possible folds. If all of the folds performed are Type I folds and the final immersion, ψ , is a homeomorphism (or, equivalently, graph isomorphism), then the original map is a homotopy equivalence. In fact, for our applications, we will always only perform Type I folds, and ψ will be a graph automorphism.

2.7. Perron Numbers

In this section, we compile relevant definitions and facts in algebraic number theory, toward the introduction of Perron numbers.

Definition 2.10. An **algebraic integer** is a complex number that is a root of some monic polynomial with integer coefficients. A polynomial with integer coefficients is also called an integer polynomial.

Given an algebraic integer α , the **minimal polynomial** p_α of α is the integer monic polynomial of the least degree that has α as a root. Then the **degree** of α denoted by $\deg \alpha$, is the degree of its minimal polynomial p_α . The **Galois conjugates**, sometimes simply called the conjugates, of α are the other roots of p_α .

Definition 2.11. A **weak Perron number** is a real algebraic integer $\lambda = \lambda_1$, whose Galois conjugates $\lambda_2, \dots, \lambda_d$ have modulus no larger than λ :

$$\lambda \geq |\lambda_i|, \quad \text{for all } i = 1, \dots, d.$$

A **(strong) Perron number** is a real algebraic integer $\lambda = \lambda_1$ that satisfies the strict inequality for the moduli of Galois conjugates $\lambda_2, \dots, \lambda_d$, namely:

$$\lambda > |\lambda_i|, \quad \text{for all } i = 2, \dots, d.$$

The following fact will be used in [Section 6](#) to expand our result on Perron numbers to weak Perron numbers.

Proposition 2.12 ([8, Theorem 3]). *Let λ be an algebraic integer. Then λ is weak Perron if and only if λ^n is Perron for some positive integer n .*

Now let λ be a Perron number. Denote by $\mathbb{Q}(\lambda)$ the **number field** that is the smallest field extension over \mathbb{Q} containing λ . As λ is an algebraic integer, we have $\mathbb{Q}(\lambda) = \mathbb{Q}[x]/(p_\lambda)$. Note $\deg \mathbb{Q}(\lambda) = \deg \lambda$. Let \mathcal{O}_λ be the **ring of integers** in the field $\mathbb{Q}(\lambda)$, defined as the set of all algebraic integers in $\mathbb{Q}(\lambda)$. Then, the ring of integers will be realized as a submodule of the number field with exactly $\deg \lambda$ basis elements:

Fact 2.13 ([9, Theorem 9, pp.20–21]). *If $\mathbb{Q}(\lambda)$ is a number field of degree d , then its ring of integers \mathcal{O}_λ is a free \mathbb{Z} -submodule of $\mathbb{Q}(\lambda)$ of rank d . In other words, there are d elements $\alpha_1, \dots, \alpha_d \in \mathcal{O}_\lambda$ such that*

$$\mathcal{O}_\lambda = \{m_1\alpha_1 + \dots + m_d\alpha_d \mid m_1, \dots, m_d \in \mathbb{Z}\}.$$

Remark 2.14. For an algebraic integer μ of degree d , $\mathbb{Z}[\mu] \subset \mathcal{O}_\mu$ will also be a rank d submodule of $\mathbb{Q}(\mu)$, but $\mathbb{Z}[\mu]$ may not be *maximal* in the sense that, for some μ , the containment is proper $\mathbb{Z}[\mu] \subsetneq \mathcal{O}_\mu$.

For example, when $\mu = \sqrt{5}$, then $\mathcal{O}_\mu = \mathbb{Z} \left[\frac{1+\sqrt{5}}{2} \right]$ as $5 \equiv 1 \pmod{4}$. (refer to [9, Chapter 2, Corollary 2]). This properly contains $\mathbb{Z}[\sqrt{5}]$ and both of these are rank-2 submodules of $\mathbb{Q}[\sqrt{5}]$. Even when λ is a Perron number, for example take $\lambda = 1 + \sqrt{5}$, it is possible that $\mathbb{Z}[\lambda] \subsetneq \mathcal{O}_\lambda$. Additionally, there are many examples where $\mathcal{O}_\lambda \neq \mathbb{Z}[\mu]$ for every algebraic integer $\mu \in \mathcal{O}_\lambda$. (For an example, refer to [10, pp 64–65].) Such are said to be **non-monogenic**.

For [Section 3.3](#), we need the following lemma.

Lemma 2.15 (Even Lemma). *Let λ be an algebraic integer. Then, there exist $n \neq n_0 \in \mathbb{Z}^+$ such that*

$$\lambda^n \equiv \lambda^{n_0} \pmod{2\mathcal{O}_\lambda}.$$

Proof. Since $\mathcal{O}_\lambda/2\mathcal{O}_\lambda \cong (\mathbb{Z}/2\mathbb{Z})^d$ is finite, by the pigeonhole principle there must be some $n \neq n_0 \in \mathbb{Z}^+$ for which $\lambda^n \equiv \lambda^{n_0} \pmod{2\mathcal{O}_\lambda}$. □

Remark 2.16 (Counterexample for $n_0 = 0$). In [13], Thurston stated the Even Lemma with $n_0 = 0$, but that statement is false in general. Indeed, consider $\lambda = \frac{3+\sqrt{17}}{2}$. This is a degree 2 Perron number with minimal polynomial $p_\lambda(x) = x^2 - 3x - 2$. Because $17 \equiv 1 \pmod{4}$, $\mathcal{O}_\lambda = \mathbb{Z}[\lambda]$. We can write

the basis elements $1, \lambda \in \mathbb{Z}[\lambda] = \mathcal{O}_\lambda$ as ordered pairs with respect to the basis for $\mathbb{Q}(\lambda) = \mathbb{Q}^2 = \langle 1, \lambda \rangle$; that is $\mathbb{Z}[\lambda] = \langle 1, \lambda \rangle = \langle (1, 0), (0, 1) \rangle$.

We claim that there is no $n \in \mathbb{Z}^+$ for which $\lambda^n \equiv 1 = (1, 0) \pmod{2\mathcal{O}_\lambda}$, where here $2\mathcal{O}_\lambda = \langle (2, 0), (0, 2) \rangle$. In fact, we assert that $\lambda^n \equiv (0, 1) \pmod{2\mathcal{O}_\lambda}$ for all $n \in \mathbb{Z}^+$. Using the relationship given from the minimal polynomial $\lambda^2 = 3\lambda + 2$, we see that this is true for the first few powers of λ :

$$\begin{aligned}\lambda^1 &= (0, 1) \not\equiv (1, 0) \pmod{2\mathcal{O}_\lambda}, \\ \lambda^2 &= (2, 3) \equiv (0, 1) \not\equiv (1, 0) \pmod{2\mathcal{O}_\lambda}, \\ \lambda^3 &= (6, 11) \equiv (0, 1) \not\equiv (1, 0) \pmod{2\mathcal{O}_\lambda}.\end{aligned}$$

More generally, as we will see in [Lemma 3.5](#), the multiplication by λ on an element in $\mathbb{Q}(\lambda) \supset \mathcal{O}_\lambda$ can be realized as the multiplication by the companion matrix $C_\lambda = \begin{pmatrix} 0 & 2 \\ 1 & 3 \end{pmatrix}$ on the corresponding ordered pair written as a column vector. It follows inductively that for $n > 1$:


$$\lambda^n = C_\lambda \cdot \lambda^{n-1} \equiv \begin{pmatrix} 0 & 2 \\ 1 & 3 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 2 \\ 3 \end{pmatrix} \equiv \begin{pmatrix} 0 \\ 1 \end{pmatrix} \pmod{2\mathcal{O}_\lambda},$$

thus $\lambda^n \equiv (0, 1) \not\equiv (1, 0) \pmod{2\mathcal{O}_\lambda}$ for all $n \in \mathbb{Z}^+$.

3. Star Maps

3.1. Definition of a Star Map

Definition 3.1. A **star** (referred to as an asterisk graph in [\[13\]](#)) with n -tips is the complete bipartite graph $*_n = K_{1,n}$.




A **star map** is a self graph map of a star, $f: *_n \rightarrow *_n$ such that:

- (a) f fixes the center vertex (in this way f preserves the bipartite structure of $*_n$), and
- (b) the map that sends each edge to the first edge of its image is a permutation. We will refer to this map as the **first edge map**.

We note that a star map can always be homotoped to be taut. Namely, the image of each edge can be reduced by canceling back-tracking, only leaving the last letter. However, we will let our star maps have back-trackings, and will not pass to a simpler representative in its homotopy class. Such redundancy will be crucial when we are defining the *split* of a star map in [Section 5](#).

Example 3.2. Let $f: *_3 \rightarrow *_3$ be the star map defined by:




Since f maps each edge in $*_3$ to an edge path of odd (unsigned) length, it must fix the center vertex. Also, letting f_1 be the first edge map, we have that $\text{Im}(f_1) = \{a, b, c\}$. Thus, f is a star map.

This section aims to prove the existence of a uniformly λ -expanding star map for every Perron number λ .

Theorem 3.3 ([\[13, Theorem 6.2\]](#), **Uniformly Expanding Star Maps for Perron Numbers**). *Let λ be a Perron number. Then there exists $n > 0$ and a star map $f: *_n \rightarrow *_n$ such that f is a uniform λ -expander with mixing transition matrix.*

Example 3.4. The following is an example of a star map that satisfies [Theorem 3.3](#). Suppose $\lambda = 5$. Consider the star graph $*_4$ and endow each edge with length 1. Let $f: *_4 \rightarrow *_4$ be defined by



We see that f is a star map since the first edge map f_1 permutes the edges ($a \rightarrow b \rightarrow a$ and $c \rightarrow d \rightarrow c$),

and f respects the bipartite structure since it maps each edge to an edge path of odd length. Furthermore, f is a 5-uniform expander since $\ell(f(e)) = 5 = 5\ell(e)$ for all edges e .

It is straightforward to check that the cube of the transition matrix is positive, so f is mixing. This example can easily be generalized to a construction for λ equal to any *odd* integer.

3.2. Geometry of Perron Numbers

Let λ be a Perron number with $\deg \lambda = d$. In this section, we will study the dynamics of λ -multiplication on $\mathbb{Q}(\lambda)$.

Viewing $\mathbb{Q}(\lambda)$ as a d -dimensional \mathbb{Q} -vector space, we *extend scalars* by tensoring $\mathbb{Q}(\lambda)$ with \mathbb{R} :

$$V_\lambda := \mathbb{Q}(\lambda) \otimes_{\mathbb{Q}} \mathbb{R} \cong \mathbb{R}^d,$$

where the latter isomorphism comes from realizing $\mathbb{Q}(\lambda)$ as a d -dimensional \mathbb{Q} -vector space with basis $1, \lambda, \dots, \lambda^{d-1}$. Then, we have an isomorphism:

$$V_\lambda = \mathbb{Q}(\lambda) \otimes_{\mathbb{Q}} \mathbb{R} \longrightarrow \mathbb{R}^d, \quad \sum_{i=0}^{d-1} (\lambda^i \otimes_{\mathbb{Q}} r_i) \longmapsto (r_0, \dots, r_{d-1}).$$

Using this identification, we can view a number $q \in \mathbb{Q}(\lambda)$ as a d -dimensional column vector $\mathbf{q} \in V_\lambda$. Namely, with the basis $\{1, \lambda, \dots, \lambda^{d-1}\}$ of $\mathbb{Q}(\lambda)$, write $q = \sum_{i=0}^{d-1} (\lambda^i \cdot q_i) \in \mathbb{Q}(\lambda)$ with $q_i \in \mathbb{Q}$ for all $i = 0, \dots, d - 1$. Then q corresponds to a d -dimensional column vector $\mathbf{q} = [q_0, \dots, q_{d-1}]^t$ in V_λ .

Now we convert the λ -multiplication action on $\mathbb{Q}(\lambda)$ into matrix multiplication by the *companion matrix*, C_λ , of λ on the corresponding vector space $V_\lambda \cong \mathbb{R}^d$. Recall the **companion matrix** C_μ of an algebraic integer

μ of degree d is the $d \times d$ matrix:

$$C_\mu = \begin{pmatrix} 0 & 0 & \cdots & 0 & -c_0 \\ 1 & 0 & \cdots & 0 & -c_1 \\ 0 & 1 & \cdots & 0 & -c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & -c_{d-1} \end{pmatrix}$$

where the c_i 's are the integer coefficients of the minimal polynomial of μ ; $p_\mu(x) = c_0 + c_1x + \dots + c_{d-1}x^{d-1} + x^d$.

Lemma 3.5. *Let λ be an algebraic integer, and $q \in \mathbb{Q}(\lambda)$. Then $\lambda \cdot q \in \mathbb{Q}(\lambda)$ corresponds to $C_\lambda \cdot \mathbf{q} \in V_\lambda$, where C_λ is the companion matrix of λ .*

Proof. Let $\deg \lambda = d$ with the minimal polynomial $p_\lambda(x) = c_0 + c_1x + \dots + c_{d-1}x^{d-1} + x^d$. Since $p_\lambda(\lambda) = 0$, we have

$$\lambda^d = \sum_{i=0}^{d-1} -c_i \lambda^i.$$

Writing $q \in \mathbb{Q}(\lambda)$ as $\sum_{i=0}^{d-1} (\lambda^i \cdot q_i)$ for $q_i \in \mathbb{Q}$, we identify q with $\mathbf{q} = [q_0, \dots, q_{d-1}]^t$ in V_λ . Now

$$\begin{aligned} \lambda \cdot q &= \lambda \sum_{i=0}^{d-1} (\lambda^i \cdot q_i) = \left\{ \sum_{i=0}^{d-2} (\lambda^{i+1} \cdot q_i) \right\} + \lambda^d \cdot q_{d-1} \\ &= \left\{ \sum_{i=1}^{d-1} \lambda^i \cdot q_{i-1} \right\} + \sum_{i=0}^{d-1} (-c_i \lambda^i) \cdot q_{d-1} \\ &= -c_0 q_{d-1} \cdot \lambda^0 + \sum_{i=1}^{d-1} (q_{i-1} - c_i q_{d-1}) \cdot \lambda^i, \end{aligned}$$

where the last line can be identified with the matrix product:

$$\begin{pmatrix} 0 & 0 & \cdots & 0 & -c_0 \\ 1 & 0 & \cdots & 0 & -c_1 \\ 0 & 1 & \cdots & 0 & -c_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & -c_{d-1} \end{pmatrix} \begin{pmatrix} q_0 \\ q_1 \\ q_2 \\ \vdots \\ q_{d-1} \end{pmatrix} = C_\lambda \mathbf{q},$$

concluding the proof. □

Given [Lemma 3.5](#), we will interchangeably use C_λ -action and λ -action by identifying $\lambda \cdot q \in \mathbb{Q}(\lambda)$ and $C_\lambda \cdot \mathbf{q} \in V_\lambda$.

Next, we find a subset of V_λ that is invariant under C_λ -multiplication. As this is a linear action, the natural choice of such a subset is an eigenspace of C_λ . However, as λ is a Perron number, we can find a more useful invariant space K_λ as follows. First, by definition of the companion matrix, the minimal polynomial of C_λ is exactly the minimal polynomial p_λ of λ . Since any finite extension over \mathbb{Q} is separable, p_λ has no repeating roots. Hence, we can enumerate the eigenvectors of C_λ as v_1, \dots, v_d , where v_1 is associated with the leading eigenvalue λ . Normalize v_1, \dots, v_d to have norm 1. Each $\mathbf{v} \in V_\lambda$ can be expressed as $\mathbf{v} = a_1 v_1 + \dots + a_d v_d$ for some $a_i \in \mathbb{R}$. Now define the invariant open cone K_λ as follows:

$$K_\lambda := \{a_1 v_1 + \dots + a_d v_d \in V_\lambda \mid a_1 > 0, \quad a_1 > |a_i|, \text{ for all } i = 2, \dots, d\}.$$

Namely, K_λ is the set of all points in V_λ whose projection to the λ -eigenspace of C_λ is positive and larger than the size of projection to any of the other eigenspaces.

Note K_λ is **polyhedral**, namely, the cone is generated by finitely many vectors in V_λ . Indeed, the $2(d-1)$ bisectors, $\{v_1 \pm v_i\}_{i=2}^d$, will generate K_λ . We label these by $\{w_1, \dots, w_{2d-2}\}$ in the remainder of this section. Note that on each of these bisectors, the projections onto either v_1 or v_i have the same magnitude. See [Figure 3](#) for K_λ with $d = 2$.




Figure 3: $v_1 \pm v_2$ generate K_λ when $d = 2$.

The fact that λ is Perron makes K_λ invariant under C_λ multiplication. Moreover, we have the proper containment $C_\lambda \cdot K_\lambda \subsetneq K_\lambda$. To see why, let $\mathbf{v} = a_1 v_1 + \dots + a_d v_d$, with λ_i being the associated eigenvalue corresponding

to the eigenvector v_i for $i = 1, \dots, d$. Note $\lambda_1 = \lambda$. Then,

$$C_\lambda \cdot \mathbf{v} = a_1 \lambda_1 v_1 + \dots + a_d \lambda_d v_d.$$

Because λ is Perron, we have that $|\lambda_1| > |\lambda_i|$ for all $i = 2, \dots, d$. Therefore,

$$|a_1 \lambda_1| > |a_i \lambda_i|, \quad \text{for all } i = 2, \dots, d,$$

showing that $C_\lambda \cdot \mathbf{v} \subset K_\lambda$. To see that the containment is proper, it suffices to show that the faces of K_λ get mapped to the interior of K_λ . Indeed, let $\mathbf{v} = av_1 + av_i$ with $a \neq 0$ be a point on a face of K_λ . Then $C_\lambda \cdot \mathbf{v} = a\lambda_1 v_1 + a\lambda_i v_i$. Since $|a\lambda_1| > |a\lambda_i|$, it follows that $C_\lambda \cdot \mathbf{v}$ has *strictly* larger projection on $\langle v_1 \rangle$ than on $\langle v_i \rangle$, showing $C_\lambda \cdot \mathbf{v} \in K_\lambda$.

Recall by [Fact 2.13](#), we can embed \mathcal{O}_λ into $\mathbb{Q}(\lambda) \subset V_\lambda$ as a rank- d lattice. In fact, we would like to trim the invariant cone K_λ to have faces passing through such nice lattice points. Any such cone generated by vectors in the lattice will be called **rational**.

Proposition 3.6 (Rational Cone; cf. [\[13\]](#), [Proposition 3.4](#)). *There is a rational polyhedral convex cone KR_λ contained in K_λ and containing $\lambda \cdot K_\lambda$.*

Proof. Consider the projective space $\mathbb{P}(V_\lambda) \cong \mathbb{P}(\mathbb{R}^d)$. We first claim that $\mathbb{P}(\mathcal{O}_\lambda)$ is dense in $\mathbb{P}(V_\lambda)$. This follows from the fact that $\mathbb{Q}(\lambda)$ is dense in V_λ and that $\mathbb{P}(\mathcal{O}_\lambda) = \mathbb{P}(\mathbb{Q}(\lambda))$. Indeed, by definition of the ring of integers, for any $x \in \mathbb{Q}(\lambda)$ there exists $m \in \mathbb{Z}^+$ such that $mx \in \mathcal{O}_\lambda$ where we can pick m to be the least common multiple of the denominators of coefficients of the (monic and rational) minimal polynomial of x .

Recall we have the proper containment $\lambda \cdot K_\lambda \subsetneq K_\lambda$. Let w_1, \dots, w_{2d-2} be the generators of the cone K_λ described above. Then $\lambda \cdot w_1, \dots, \lambda \cdot w_{2d-2}$ are the generators of the cone $\lambda \cdot K_\lambda$ where $w_j \neq \lambda \cdot w_j$. Now, project the generators onto $\mathbb{P}(V_\lambda)$. Note that w_j and $\lambda \cdot w_j$ are *not* identified in $\mathbb{P}(V_\lambda)$ because here $\lambda \cdot w_j$ corresponds to the vector $C_\lambda \cdot w_j$ and $w_j = v_1 \pm v_i$, where v_1 and v_i are eigenvectors for different eigenvalues of \mathbb{C}_λ .

Thus, each geodesic segment between $\overline{w_j}$ and $\overline{\lambda w_j}$ in $\mathbb{P}(V_\lambda)$ is nondegenerate. As $\mathbb{P}(\mathcal{O}_\lambda)$ is dense in $\mathbb{P}(V_\lambda)$, for each $j = 1, \dots, 2d-2$ we can pick $\overline{u_j} \in \mathbb{P}(\mathcal{O}_\lambda)$ arbitrarily close to the midpoint of the geodesic joining $\overline{w_j}$ and $\overline{\lambda w_j}$ in $\mathbb{P}(V_\lambda)$.




Figure 4: The figure on the left shows the choice of \overline{u}_1 and \overline{u}_2 from $\mathbb{P}(\mathcal{O}_\lambda)$ so that each \overline{u}_i lies arbitrarily close to the midpoint of the two corresponding projections \overline{w}_i and $\overline{\lambda w}_i$. This is possible as $\mathbb{P}(\mathcal{O}_\lambda)$ is dense in $\mathbb{P}(V_\lambda)$. Such a choice of \overline{u}_1 and \overline{u}_2 gives rise to the desired proper containment: $\lambda \cdot K_\lambda \subset KR_\lambda \subset K_\lambda$, shown on the right.

Finally, let KR_λ denote the closed rational polyhedral cone generated by u_1, \dots, u_{2d-2} . Illustrated in Figure 4, by construction we have

$$\lambda \cdot K_\lambda \subset KR_\lambda \subset K_\lambda,$$

which concludes the proof. □

Define $S_\lambda := (\mathcal{O}_\lambda \cap KR_\lambda) \setminus \{0\}$, the set of lattice points in the rational cone KR_λ in V_λ , minus the origin. Note S_λ is equipped with a semigroup structure, as both \mathcal{O}_λ and KR_λ are closed under the addition. Now we show it is also *finitely generated*.

Proposition 3.7 (Gordan’s Lemma; cf. [13, Proposition 3.5]). *S_λ is a finitely generated semigroup.*

Proof. Say KR_λ is generated by $u_1, \dots, u_k \in \mathcal{O}_\lambda$ constructed in the proof of Proposition 3.6. Let $H := \{\sum_{i=1}^k a_i u_i \mid a_i \in [0, 1]\} \subset KR_\lambda$. Since H is compact and \mathcal{O}_λ is discrete, it follows that $H \cap \mathcal{O}_\lambda$ is finite. We claim that this finite set, which we label $\{s_1, \dots, s_m\}$, is the desired generating set for S_λ , which we now show. Note that $u_i \in H \cap \mathcal{O}_\lambda$ by taking $a_i = 1$ and $a_j = 0$ for $j \neq i$. Therefore, the set $\{s_1, \dots, s_m\}$ contains the set $\{u_1, \dots, u_k\}$.

Now, pick any $u = \sum_{i=1}^k b_i u_i \in S_\lambda$, where $b_i \in \mathbb{Q}_{\geq 0}$ and $u \neq 0$. Then each b_i can be decomposed into an integer and non-integer part $b_i = n_i + r_i$, where

$n_i \in \mathbb{Z}_{\geq 0}$ and $r_i \in [0, 1)$, and

$$u = \left(\sum_{i=1}^k n_i u_i \right) + \left(\sum_{i=1}^k r_i u_i \right).$$

We aim to show that u can be written as a non-negative integer sum of $\{s_1, \dots, s_m\}$. Since $u_i \in \{s_1, \dots, s_m\}$ for all i , the left part of the summation, $\sum_{i=1}^k n_i u_i$, is a non-negative integer sum of $\{s_1, \dots, s_m\}$. In addition, $\sum_{i=1}^k r_i u_i \in H$ by definition. However, $u - \sum_{i=1}^k n_i u_i \in \mathcal{O}_\lambda$, so that $\sum_{i=1}^k r_i u_i \in H \cap \mathcal{O}_\lambda$. This means that $\sum_{i=1}^k r_i u_i$ is one of the s_j itself, which concludes the proof. \square

Proposition 3.8 (Slim Cone Lemma; cf. [13], Section 4). *Let $s \in S_\lambda$. Then there exists $N_s > 0$ such that $n \geq N_s$ implies that $\lambda^n \cdot S_\lambda \subset s + S_\lambda$.*

Proof. Since S_λ is finitely generated (Proposition 3.7), we can let $S_\lambda = \langle s_1, \dots, s_m \rangle$. To show the conclusion, it suffices to prove the following claim:

Claim 1. *Fix $s \in S_\lambda$. For each $i = 1, \dots, m$, there exists $N_i > 0$ such that $n \geq N_i$ implies that $\lambda^n \cdot s_i \in s + S_\lambda$.*

Claim 1 then implies the conclusion by taking $N_s := \max\{N_1, \dots, N_m\}$. Then for any $t = a_1 s_1 + \dots + a_m s_m \in S_\lambda$ with $a_i \in \mathbb{Z}_{\geq 0}$ for all i , it follows that for $n > N_s$,

$$\lambda^n \cdot t = \lambda^n \cdot \left(\sum_{i=1}^m a_i s_i \right) = \sum_{i=1}^m a_i (\lambda^n \cdot s_i) \in s + S_\lambda$$

because $\lambda^n \cdot s_i \in s + S_\lambda$ for each i and $s + S_\lambda$ is a semigroup.

To prove Claim 1, we need another claim:

Claim 2. *For any $s \in S_\lambda$, the projection of $s + KR_\lambda$ subjects onto the interior of $\mathbb{P}(KR_\lambda)$.*

Proof of Claim 2. Write $s = b_1 s_1 + \dots + b_m s_m$ for some nonnegative integers b_1, \dots, b_m . Any interior point of $\mathbb{P}(KR_\lambda)$ can be represented by an interior point of KR_λ . Let t be an interior point of KR_λ . We will find $t' \in s + KR_\lambda$ such that $\bar{t} = \bar{t}'$. Write $t = a_1 s_1 + \dots + a_m s_m$. Then $a_1, \dots, a_m > 0$, because t does not lie on any face of KR_λ . Therefore, there exists a large R

such that $Ra_i > b_i$ for every $i = 1, \dots, m$. Then we claim that

$$Rt = Ra_1s_1 + \dots + Ra_ms_m \in s + KR_\lambda.$$

Indeed, $Rt - s = \sum_{i=1}^m (Ra_i - b_i)s_i$ has nonnegative coefficients, so $Rt - s \in KR_\lambda$. Therefore, $\overline{Rt} \in \mathbb{P}(s + KR_\lambda)$ and $\overline{Rt} = \bar{t} \in \text{int}(\mathbb{P}(KR_\lambda))$, as desired. \triangle

Now we prove [Claim 1](#). Note the leading eigenvector \bar{v}_1 is the unique attracting fixed point in $\mathbb{P}(KR_\lambda)$ under the λ -multiplication action. [Claim 2](#) implies that $\bar{v}_1 \in \mathbb{P}(s + KR_\lambda)$. Then $\{\overline{\lambda^n \cdot s_i}\}_{n \in \mathbb{Z}^+}$ converges to \bar{v}_1 . Hence, eventually for some large $N_s \in \mathbb{Z}^+$, we can say whenever $n \geq N_s$, it follows that $\overline{\lambda^n \cdot s_i}$ is arbitrarily close to \bar{v}_1 in $\mathbb{P}(s + KR_\lambda)$. Together with the fact that $\lambda^n \cdot s_i \in S_\lambda$ for all $n > 0$, we obtain

$$n \geq N_s \implies \lambda^n \cdot s_i \in s + S_\lambda,$$

proving [Claim 1](#) and concluding the proof. \square

3.3. Uniformly λ -expanding Star Map – The proof of [Theorem 3.3](#)

In this section, we construct a uniformly λ -expanding star map for any Perron number λ to prove [Theorem 3.3](#), which we restate here:

Theorem 3.4 (Uniformly Expanding Star Maps for Perrons). *Let λ be a Perron number. Then there exists $n > 0$ and a star map $f : *_n \rightarrow *_n$ such that f is a uniform λ -expander with mixing incidence matrix.*

Proof. Use [Lemma 2.15](#) to find $N, n_0 \in \mathbb{Z}^+$ with $N > n_0$ such that $\lambda^N \equiv \lambda^{n_0} \pmod{2\mathcal{O}_\lambda}$. Recall by [Proposition 3.7](#), we can write $S_\lambda = \langle s_1, \dots, s_m \rangle$. Set $T = s_1 + \dots + s_m$, and for each $k = 1, \dots, m$ define:

$$g_k = \lambda^{n_0} \cdot s_k + 2(T + \lambda \cdot s_k) \in S_\lambda.$$

For each $k = 1, \dots, m$, [Proposition 3.8](#) with $s = g_k$ yields $N_k \in \mathbb{Z}^+$ such that $n \geq N_k$ implies that $\lambda^n \cdot s_k \in g_k + S_\lambda$.

Now, letting $M := p(N - n_0) + N$, where p is sufficiently large integer so that $M \geq \max\{N_1, \dots, N_m\}$, we have that $\lambda^M \cdot s_k \in g_k + S_\lambda$, for all

$k = 1, \dots, m$. We claim that we also have that $\lambda^M \equiv \lambda^{n_0} \pmod{2\mathcal{O}_\lambda}$ since

$$\begin{aligned} \lambda^{p(N-n_0)+N} &\equiv \lambda^{p(N-n_0)+n_0} \\ &\equiv \lambda^{pN-(p-1)n_0} \\ &\equiv \lambda^{(p-1)(N-n_0)+N} \\ &\quad \vdots \\ &\equiv \lambda^{(p-2)(N-n_0)+N} \\ &\quad \vdots \\ &\equiv \lambda^N \equiv \lambda^{n_0} \pmod{2\mathcal{O}_\lambda}. \end{aligned}$$

Hence, for each $k = 1, \dots, m$,

$$\lambda^M \cdot s_k - \lambda^{n_0} \cdot s_k \in 2(T + \lambda \cdot s_k) + S_\lambda,$$

where the left hand side is $(\lambda^M - \lambda^{n_0}) \cdot s_k \equiv 0 \pmod{2\mathcal{O}_\lambda}$. This implies that $(\lambda^M - \lambda^{n_0})s_k$ is written as a linear combination of the $\{s_i\}_{i=1}^m$ with *even* coefficients. However, as $2(T + \lambda \cdot s_k)$ already has even coefficients so that for each $k = 1, \dots, m$, there must exist $e_1^{(k)}, \dots, e_m^{(k)} \in \mathbb{Z}_{\geq 0}$ such that

$$\begin{aligned} (\lambda^M - \lambda^{n_0}) \cdot s_k &= 2(T + \lambda \cdot s_k) + \sum_{i=1}^m 2e_i^{(k)} s_i \\ &= 2\lambda \cdot s_k + \sum_{i=1}^m (2e_i^{(k)} + 2)s_i. \end{aligned}$$

Thus, we obtain the following key identity that will be used to construct the desired uniformly λ -expanding star map:

(†)

$$\lambda^M \cdot s_k = \left[\sum_{i=1}^m (2e_i^{(k)} + 2)s_i \right] + (2\lambda \cdot s_k) + (\lambda^{n_0} \cdot s_k), \quad \text{for } k = 1, \dots, m.$$

Now, we begin the construction of the star map using [Equation \(†\)](#). Consider a star graph $*_{mM}$ with mM tips. For $k = 1, \dots, m$ and $i = 1, \dots, M$, label each edge by (s_k, i) and set its length to be:

$$\|(s_k, i)\| = \lambda^{i-1} \cdot s_k \in S_\lambda \subset \mathbb{R}^+,$$

where the last containment comes from the usual embedding $S_\lambda = (KR_\lambda \cap \mathcal{O}_\lambda) \setminus \{\mathbf{0}\} \hookrightarrow (KR_\lambda \cap \mathbb{Q}(\lambda)) \setminus \{\mathbf{0}\} \hookrightarrow \mathbb{R}^+$.

Now define the star map $f_\lambda : *_mM \rightarrow *_mM$ by mapping each edge (s_k, i) to $(s_k, i + 1)$ when $i < M$ and mapping (s_k, M) to an edge path as follows:

$$(s_k, M) \xrightarrow{f_\lambda} (2e_k^{(k)} + 2)(s_k, 1) \sqcup \bigsqcup_{i \neq k} (2e_i^{(k)} + 2)(s_i, 1) \sqcup 2(s_k, 2) \sqcup (s_k, n_0 + 1)$$

More precisely, f_λ sends each edge $\{(s_k, i)\}_{i=1}^{M-1}$ to the next edge corresponding to s_k , and sends the M -th edge (s_k, M) to an edge path traversed in the following order:

- (i) First, it maps $2e_k^{(k)} + 2$ times over the edge $(s_k, 1)$,
- (ii) next, for each $j \neq k$ it goes $2e_j^{(k)} + 2$ times over the edge $(s_j, 1)$,
- (iii) then, it maps 2 times over the edge $(s_k, 2)$,
- (iv) and finally, it maps over the edge $(s_k, n_0 + 1)$ once.

The choices of lengths of edges in $*_mM$ guarantee that f_λ is uniformly λ -expanding. We immediately have $\|f_\lambda((s_k, i))\| = \|(s_k, i + 1)\| = \lambda^i s_k = \lambda \|(s_k, i)\|$ for $1 \leq i < M$. For $i = M$, by [Equation \(†\)](#) we have

$$\begin{aligned} \|f_\lambda((s_k, M))\| &= \left[\sum_{i=1}^m (2e_i^{(k)} + 2)s_i \right] + (2\lambda \cdot s_k) + (\lambda^{n_0} \cdot s_k) \\ &= \lambda^M s_k = \lambda \|(s_k, M)\|. \end{aligned}$$

Now we verify that f_λ is a star map with mixing incidence matrix. First, the fact that f_λ traverses all edges but one an even number of times implies that f_λ fixes the center vertex (or equivalently, preserves the bipartite structure). By construction, it follows that every edge is the first element of the image edge path of some edge (i.e., the first edge map f_1 is a permutation).

Finally, to see that f_λ has mixing incidence matrix, we observe that for each edge (s_k, i) (regarding i as an integer modulo M):

$$\begin{aligned} f_\lambda^M((s_k, i)) &\supset \bigsqcup_{k=1}^m (s_k, i) \\ f_\lambda^{2M}((s_k, i)) &\supset \bigsqcup_{k=1}^m (s_k, i) \cup \bigsqcup_{k=1}^m (s_k, i + 1) \end{aligned}$$

$$\begin{aligned}
f_\lambda^{3M}((s_k, i)) &\supset \bigsqcup_{k=1}^m (s_k, i) \cup \bigsqcup_{k=1}^m (s_k, i+1) \cup \bigsqcup_{k=1}^m (s_k, i+2) \\
&\vdots \\
f_\lambda^{M^2}((s_k, i)) &\supset \bigsqcup_{i=1}^M \bigsqcup_{k=1}^m (s_k, i) = *_m M,
\end{aligned}$$

which implies that f_λ has mixing incidence matrix. Indeed, denoting by A_λ the incidence matrix of f_λ , its power $A_\lambda^{M^2}$ is a positive matrix by the above observation. This concludes the proof. \square

The fact that the incidence matrix for f_λ is mixing will be used to prove [Theorem 6.1](#), which states that the incidence matrix of its *split* $S(f_\lambda)$ is as well.

Next, we expand [Theorem 3.3](#) to *weak Perron numbers* via the following lemma.

Lemma 3.9 (Uniformly Expanding Star Maps for Weak Perrons).

*Let λ be a weak Perron number. Then there exists $k > 0$ and a star map $f : *_k \rightarrow *_k$ such that f is a uniform λ -expander.*

Proof. If λ is a weak Perron number, then by [Proposition 2.12](#), there exists some N such that λ^N is a Perron number. Then by [Theorem 3.3](#), we can construct a star map $f : *_n \rightarrow *_n$ for some $n > 0$, which is λ^N -uniformly expanding.

Now take N -copies C^0, \dots, C^{N-1} of $*_n$, and glue them along the center vertices to form a star $*_k = *_N n$ with Nn tips.

Decide the lengths on each edge in C^0 according to the construction of $*_n$ from [Theorem 3.3](#). Then for $i = 1, \dots, N-1$, set the edge length of edges in C^i to be exactly λ^i times the length of the corresponding edge of C^0 . Now define $f_N : *_k \rightarrow *_k$ as follows. For $0 \leq i \leq N-2$, send each edge of C^i to the corresponding edge in the next copy, C^{i+1} , i.e. f_N simply shifts C^i to C^{i+1} . Then by construction, f_N is uniformly λ -expanding on C^0, \dots, C^{N-2} . Next, to make f_N uniformly λ -expanding on C^{N-1} as well, define f_N by mapping each edge of C^{N-1} to an edge path of C^0 following the recipe given by the map $f : *_N \rightarrow *_N$ in [Theorem 3.3](#), which is λ^N -uniformly expanding. In particular, we apply f to the edge of C^{N-1} , uniformly expanding the length of the edge by λ^N , and then shift the image of the edge path to the corresponding edge path in C_0 , which shrinks the length of the edge path




Figure 5: Gluing four copies of $*_4$ to form $*_{16}$. The length of each edge in C^{i+1} is given by multiplying by λ to the corresponding edge in C^i .

by λ^{N-1} . Therefore, each edge e in C^{N-1} will be mapped to an edge path in C^0 whose length is $(\lambda^N)/\lambda^{N-1} = \lambda$ times the length of the edge e in C^{N-1} . Given a weak Perron number λ , this concludes the construction of $f_N : *_k \rightarrow *_k$ which is uniformly λ -expanding. \square

4. Prototype Graph

4.1. Constructing the prototype maps

Let P_7 denote the bipartite graph with two vertices v_0, v_1 and seven edges as in [Figure 6](#). We will refer to this graph as the **prototype graph**. Orient the edges according to the bipartite structure, *e.g.*, with initial vertex v_0 and terminal vertex v_1 ; let a, b, \dots, g label the edges with respect to this orientation, and A, B, \dots, G denote the opposite orientation. We endow P_7 with a traintrack structure as shown in [Figure 6](#). Note that all of the turns between the a, b , and c edges are legal at both of the vertices.

Let ϕ_1 be the identity map on P_7 . Now we define a set of traintrack maps $\{\phi_{3+2m}\}_{m=0}^\infty$ on the prototype graph P_7 . For $m \geq 0$ let $\phi_{3+2m} : P_7 \rightarrow P_7$ be the taut graph map defined as follows:




Figure 6: The traintrack structure on P_7 from Thurston's original paper [13]. The red paths between edges indicate the legal turns.

$$\phi_{3+2m} : \begin{cases} a \mapsto aG(aB)^m a, \\ b \mapsto bD(bC)^m b, \\ c \mapsto cF(cA)^m c, \\ d \mapsto aB(aB)^m a, \\ e \mapsto cB(aB)^m a, \\ f \mapsto aC(aB)^m a, \\ g \mapsto bE(bA)^m b. \end{cases}$$

Since edges are mapped to paths of odd length, each ϕ_{3+2m} preserves the bipartite structure of P_7 , and in fact fixes the vertices v_0, v_1 pointwise. We will call these maps $\{\phi_n\}$, for n odd and positive, the **prototype maps**.

We claim the prototype maps are all traintrack maps. Since ϕ_1 is the identity map it is a traintrack map. One can then check that for ϕ_{3+2m} every legal turn is sent to a legal turn. Note that in determining the action of a map on turns, we only need to consider the first and last edges of the image edge-paths. This allows us to check that ϕ_{3+2m} is a traintrack map for all $m \geq 0$. For example, we check ϕ_{3+2m} maps the legal turns at v_0 that involve

the a edge to legal turns at v_0 .

$$Ba \mapsto Ba$$

$$Ca \mapsto Ca$$

$$Ga \mapsto Ba.$$

Now, using [Figure 6](#), one can check that each edge is mapped to a legal path. This check is possible for all $m \geq 0$ since at most only *four* distinct turns occur in the image of each edge. For example, $\phi_3(a)$ consists of two kinds of turns aG, Ga and $\phi_{3+2m}(a)$ with $m \geq 1$ consists of four kinds of turns: aG, Ga, aB and Ba , which are all legal as shown in [Figure 6](#) and this is independent of the value of $m \geq 0$. Lastly, we note that the prototype maps are indeed taut since they are local embeddings on the interiors of edges. Again, this is independent of m and is checked directly by noting that the image of each edge is *reduced* as a word in the free group $\pi_1(P_7)$.

4.2. Homotopy equivalence via folding.

In this section, we prove that the prototype maps are, in fact, homotopy equivalences and, therefore, induce automorphisms of $\pi_1(P_7) \cong \mathbb{F}_6$, the free group of rank 6.

Lemma 4.1. *All of the prototype maps $\{\phi_n\}$, for n positive and odd, are homotopy equivalences.*

Proof. The map ϕ_1 is the identity map and hence a homotopy equivalence. We will check directly using Stallings folds that ϕ_{3+2m} is a homotopy equivalence for $m \geq 0$. Recall that if we perform only Type I folds while decomposing a graph map á la Stallings and the resulting immersion, ψ , is a homotopy equivalence, then the original map is a homotopy equivalence. This argument is largely a “proof by picture,” and so we direct the reader to [Figure 7](#) throughout the proof. We first show that ϕ_3 is a homotopy equivalence. Begin by subdividing the edges of the domain graph via the full pre-image of the vertices. Then ϕ_3 is described by [Figure 7](#). Note the edges d, e, f, g are in black as they will never be folded.

We begin with all of the possible folds on edges adjacent to the left- and rightmost vertices. That is to say, we fold all edges with the same label and orientation at a vertex. The first four collection of folds, labeled F_1, F_2, F_3 and F_4 , are performed in [Figure 7](#). For readability, we will omit the a, b, c labels on the yellow, green, and pink edges in the figures from now on. We




Figure 7: Folding ϕ_3 . All folds F_1, \dots, F_7 are of Type I.

point out that for these first four folding maps, we are only ever folding edges that are distance at most two from the leftmost vertex or one from the rightmost vertex. This is unimportant right now, but will be important

in the proof of [Proposition 5.6](#), when we verify that our *split maps* are homotopy equivalences.

[Figure 7](#) shows the next three folds, F_5 , F_6 , and F_7 that are performed. The map ψ from [Theorem 2.9](#) after these seven folding maps is the identity up to permuting edges. Thus, ϕ_3 is a homotopy equivalence.

Next, we turn our attention to an arbitrary ϕ_{3+2m} with $m \geq 1$. [Figure 8](#) shows the map ϕ_{3+2m} graphically.

We begin with the first four folding maps that are effectively the same folds as in the ϕ_3 case. The main difference is that for F_1 , we fold not just the first edge adjacent to the rightmost vertex, but instead fold the entire edge path of length $2m + 1$ adjacent to this vertex. For example, $\phi_{3+2m}(a)$ and $\phi_{3+2m}(d)$ both end in $(aB)^m a$, so the portions of the edge paths corresponding to these $2m + 1$ characters are folded. The next three maps, F_2, F_3, F_4 , are exactly the same as in the ϕ_3 case. (Compare F_2, F_3, F_4 in [Figure 7](#) with those in [Figure 8](#).)

Next, we perform a new type of folding map called a *wrapping* map. See [Figure 8](#). After performing $F_4 \circ F_3 \circ F_2$, there is a loop labeled by aB beginning at the left-most vertex. The map F_5 will consist of $2m + 1$ Type I folds that take an edge path labeled by $(aB)^m a$ and first “wrap” it around this loop labeled by aB a total of m -times before folding the last a -edge. These folds are possible because the a -edge at the start of the edge path labeled by $(aB)^m$ and the a -edge of the loop aB we are wrapping the edge path around begin at the same vertex. Next, F_6 and F_7 are simple single folds along a b -edge and a c -edge respectively. Then F_8, F_9 , and F_{10} are again “wrapping” maps. Namely, the map F_8 folds m times around cB , the map F_9 folds m times around aB , and the map F_{10} folds m times around aC . The single fold performed by the F_6 map eliminates the need to do a single fold in F_8 and F_9 . Note that it is essential to perform the single fold in the F_6 map before F_8 in order to ensure that the edge path labeled by $(cB)^m$ and the c -edge of the loop cB begin at the same vertex so that folding is possible. The same applies to the edge path labeled by $(aB)^m$ that is wrapped around aB in F_9 . Similarly, F_{10} , wrapping $(aC)^m$ over aC requires the single-edge fold F_7 . In this way, we see that we obtain the original graph and that ψ from [Theorem 2.9](#) is the identity map up to permuting edges after performing all of these Type I folds. We conclude that ϕ_{3+2m} is also a homotopy equivalence for all $m \geq 1$. \square

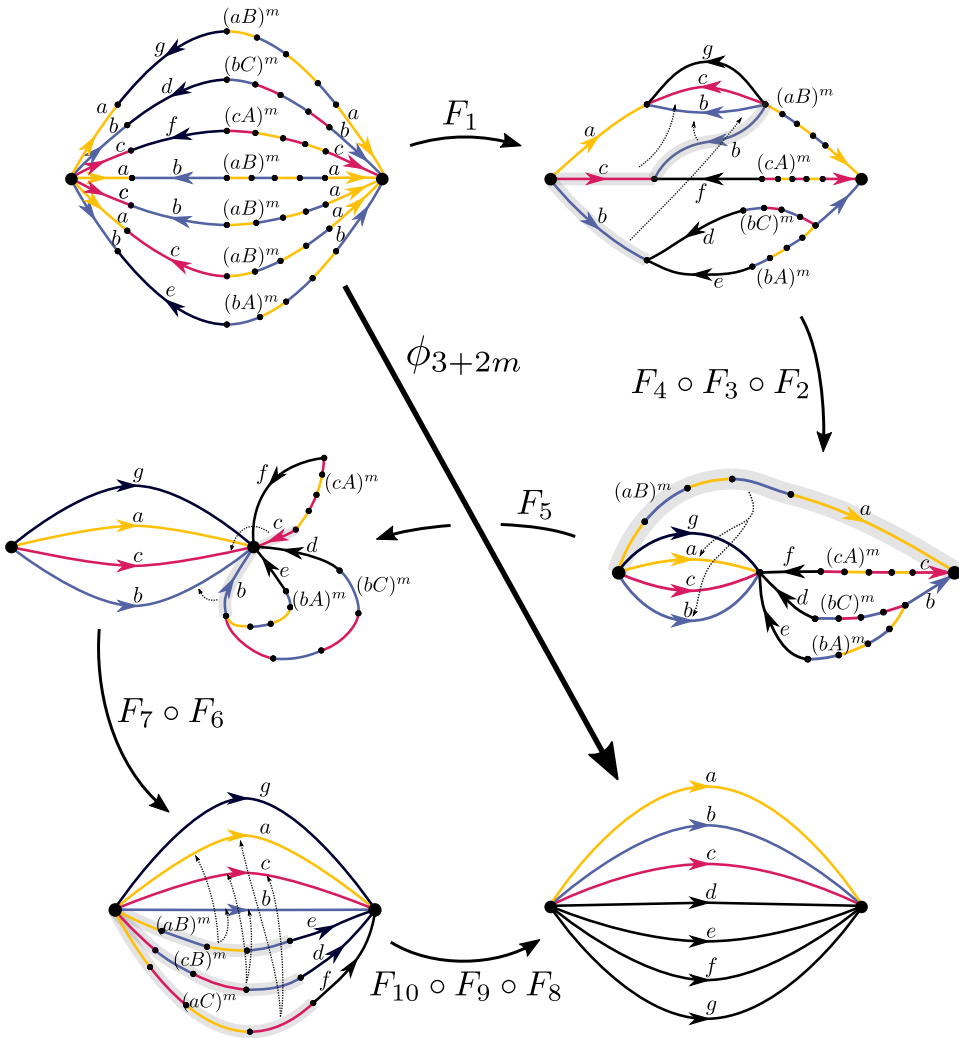


Figure 8: Folding ϕ_{3+2m} for $m \geq 1$. All folds F_1, \dots, F_{10} are of Type I. The highlighted grey edge paths at each stage represent edges involved in the folding maps being described.

5. Splitting graphs

Given a weak Perron number λ , we constructed a graph map f_λ on a star graph that is uniformly λ -expanding and a homotopy equivalence in [Section 3](#). Though it is close to being the desired map for [Theorem 1.1](#), it is far

from being a traintrack map because of the large amounts of backtracking in the star maps. In this section, we edit f_λ by “blowing up” each edge of the star graph to the prototype graph P_7 and use the prototype maps of [Section 4](#) to resolve the backtracking and obtain a traintrack map ([Proposition 5.5](#)). During this process, we carefully preserve certain properties of f_λ , specifically uniformly λ -expanding ([Proposition 5.4](#)) and being a homotopy equivalence ([Proposition 5.6](#)). This process of blowing up a graph is named *splitting* by Thurston in [[13](#), Section 9], and we will discuss the process in [Section 5.1](#).

5.1. Split Graphs and Split Maps

Definition 5.1. Given a bipartite metric graph Γ , we define the **split graph**, $S(\Gamma)$, by replacing each edge of Γ with a copy of the prototype graph P_7 from [Section 4](#) whose edges have the same length as the replaced edge. That is, if the edges of Γ are enumerated x_1, x_2, \dots , replace the edge x_i between vertices v_0 and v_1 with 7 edges labelled a_i, b_i, \dots, g_i each of which has the same length as x_i . See [Figure 9](#) for an example of a split graph.




Figure 9: Example of a split star graph. The initial graph, Γ , is on the left, and the split graph, $S(\Gamma)$, is on the right.

$S(\Gamma)$ inherits a traintrack structure from the traintrack structure on the prototype graph in the following way: turns are legal in $S(\Gamma)$ if and only if the turns without subscripts define a legal turn in the prototype. For example, a

turn of the form $a_i B_j$ is legal as the turn aB is legal in the prototype graph P_7 , and a turn of the form $a_i F_j$ is not, as the turn aF is illegal in P_7 .

Definition 5.2. Given a self-map $f : \Gamma \rightarrow \Gamma$ which preserves the bipartite structure and does not collapse any edges, we define the **split map**

$$S(f) : S(\Gamma) \rightarrow S(\Gamma)$$

as follows. Let y_i be an edge of $S(\Gamma)$ corresponding to the edge x_i of Γ , where $y \in \{a, b, \dots, g\}$. For an edge path w in Γ , denote by $\|w\|$ the length of w when each edge of Γ is endowed with length 1. Now, if $\|f(x_i)\| = \ell$, then $S(f)$ sends y_i to the edge path given by $\phi_\ell(y)$ with subscripts identical to the subscripts of the edge path $f(x_i)$. Note that the word length of the edge path $f(x_i)$ determines which prototype map is used to define the image of y_i for all $y \in \{a, b, \dots, g\}$. Symbolically:

$$S(f)(y_i) := [\phi_{\|f(x_i)\|}(y)]_{f(i)},$$

where for a path \mathcal{P} in the prototype graph P_7 , the notation $[\mathcal{P}]_{f(i)} \subset S(\Gamma)$ means that the path \mathcal{P} is given the subscripts of the edge path $f(x_i)$.

Note, since f preserves the bipartite structure, the word length of edge paths of the form $f(x_i)$ is always *odd*, which makes our definition using prototype maps $\{\phi_n\}$ well-defined. Moreover, $S(f)$ preserves the induced bipartite structure on $S(\Gamma)$.

Example 5.3. Let Γ be the star graph $*_3$. Define a star map $f : \Gamma \rightarrow \Gamma$ by:

$$\begin{aligned} f(x_1) &= x_2 X_2 x_3, \\ f(x_2) &= x_3, \\ f(x_3) &= x_1 X_1 x_3 X_3 x_2. \end{aligned}$$

Then, for the split map $S(f) : S(\Gamma) \rightarrow S(\Gamma)$, we need to use three different prototype maps for edges with three different subscripts. Namely, for $y \in \{a, b, \dots, g\}$, we need to use ϕ_3 for $S(f)(y_1)$ as $\|f(x_1)\| = 3$, use $\phi_1 = \text{id}$ for $S(f)(y_2)$ as $\|f(x_2)\| = 1$, and use ϕ_5 for $S(f)(y_3)$ as $\|f(x_3)\| = 5$.

More precisely,

$$\begin{aligned} S(f)(a_1) &= a_2 G_2 a_3, & S(f)(b_1) &= b_2 D_2 b_3, & S(f)(c_1) &= c_2 F_2 c_3, \\ S(f)(a_2) &= a_3, & S(f)(b_2) &= b_3, & S(f)(c_2) &= c_3, \\ S(f)(a_3) &= a_1 G_1 a_3 B_3 a_2, & S(f)(b_3) &= b_1 D_1 b_3 C_3 b_2, & S(f)(c_3) &= c_1 F_1 c_3 A_3 c_2, \end{aligned}$$

and so on for $y_i = d_i, e_i, f_i$, and g_i , with $i = 1, 2, 3$.

Proposition 5.4. *If $f : \Gamma \rightarrow \Gamma$ is uniformly λ -expanding, then so is $S(f) : S(\Gamma) \rightarrow S(\Gamma)$.*

Proof. Essentially, this follows from the definition of the split map. Let ℓ be the length function on Γ . If $f(x_i) = x_{i_1} \cdots x_{i_k}$, then $S(f)(y_i) = z_{i_1}^{(1)} \cdots z_{i_k}^{(k)}$ for some $z^{(1)}, \dots, z^{(k)} \in \{a^{\pm 1}, b^{\pm 1}, \dots, g^{\pm 1}\}$. By definition of $S(\Gamma)$, we have $\ell(x_{i_j}) = \ell(z_{i_j}^{(j)})$ for all j , so

$$\ell(S(f)(y_i)) = \sum_{j=1}^k \ell(z_{i_j}^{(j)}) = \sum_{j=1}^k \ell(x_{i_j}) = \ell(f(x_i)) = \lambda \ell(x_i) = \lambda \ell(y_i),$$

which proves that $S(f)$ is uniformly λ -expanding. \square

5.2. Splitting Maps are Traintrack Maps

It turns out that, with some mild conditions, $S(f)$ is always a traintrack map.

Proposition 5.5. *Let $f : \Gamma \rightarrow \Gamma$ be a map that preserves the bipartite structure and maps each edge to an edge-path of length at least 1. Then $S(f)$ is a traintrack map.*

Proof. As before, we will denote by $\{x_i\}$ the edge set of Γ , and by y_i an edge of $S(f)$ with $y \in \{a, \dots, g\}$, corresponding to an edge x_i of Γ .

Recall that the traintrack structure on $S(\Gamma)$ is defined so that a turn is legal if and only if the corresponding turn in the prototype graph P_7 obtained by forgetting subscripts is legal. Hence, it automatically follows that every edge y_i of $S(\Gamma)$ is sent to a legal path via $S(f)$, because forgetting the subscripts of the image $S(f)(y_i)$ is exactly $\phi_{\|f(x_i)\|}(y)$, which is a legal path since the prototype maps $\{\phi_n\}$ are traintrack maps. (See [Section 4.1](#).)

Therefore, to show $S(f)$ is a traintrack map, it suffices to show that $S(f)$ sends a legal turn to a legal turn. First note that only turns of the form $y_i Z_j$, where $y \in \{a, \dots, g\}$ and $Z \in \{A, \dots, G\}$ are legal in $S(\Gamma)$. Take such a legal turn. Then, again by definition, yZ is a legal turn in P_7 . Since all the legal turns in P_7 consist of at least one of the a -, b -, or c -edges, and zY is legal if and only if yZ is legal, we may assume without loss of generality that $y \in \{a, b, c\}$ (otherwise swap the roles of y and z in this proof). On the

other hand, the turn $y_i Z_j$ will be mapped to the following turn, which is a concatenation of two edges:

$$[\text{Last edge of } S(f)(y_i)] \cdot [\text{First edge of } S(f)(Z_j)].$$

This turn is legal if and only if the corresponding turn in P_7

$$[\text{Last edge of } \phi_{\|f(x_i)\|}(y)] \cdot [\text{First edge of } \phi_{\|f(x_j)\|}(Z)]$$

(i.e., deleting the subscripts) is legal in P_7 . The key observation is that by the definition of prototype maps $\{\phi_{3+2m}\}$ on the a , b , and c edges of P_7 , the last edge of $\phi_{3+2m}(y)$ is the same as y regardless of the value $3 + 2m$. Therefore, we can replace the prototype map $\phi_{\|f(x_i)\|}$ for y to $\phi_{\|f(x_j)\|}$, so the turn is actually identical to the following turn:

$$[\text{Last edge of } \phi_{\|f(x_j)\|}(y)] \cdot [\text{First edge of } \phi_{\|f(x_j)\|}(Z)],$$

which is legal because yZ is a legal turn and $\phi_{\|f(x_j)\|}$ is a traintrack map. This concludes that $S(f)$ is a traintrack map. □

5.3. Splitting Star Maps

Next we verify that $S(f)$ induces an automorphism of $\pi_1(S(\Gamma))$. That is, we check that $S(f)$ is a homotopy equivalence.

Proposition 5.6. *Let $*_n$ be a star graph and let $f : *_n \rightarrow *_n$ be a star map. Then $S(f)$ is a homotopy equivalence.*

Proof. We use Stallings folds as in the proof for prototype maps ([Lemma 4.1](#)). First, note the following two properties of star maps. A star map is always a permutation on the first edge map and each edge is mapped to an edge-path of odd length. In particular, the image of an edge under a star map is of the form $f(x_r) = x_{j_1} X_{j_1} \cdots x_{j_{m+1}} X_{j_{m+1}} x_{j_{m+2}}$, where r and j_1, \dots, j_{m+2} are in $\{1, \dots, n\}$ so that the edges in the image come in pairs (of an edge and its inverse) until the final edge.

The split graph, $S(*_n)$, is a “flower” with n “petals,” each of which is a copy of the prototype graph, P_7 . We will first focus on a single petal corresponding to an edge x_r of $*_n$ and begin folding there. The map on a petal is similar to those given by [Figure 8](#), with the differences arising from the fact that the labels on the edges come equipped with subscripts. The

subscripts are also determined by the map f on the original edges of the star graph. For instance, if $f(x_r) = x_{j_1}X_{j_1} \cdots x_{j_{m+1}}X_{j_{m+1}}x_{j_{m+2}}$, then

$$S(f)(a_r) = a_{j_1}G_{j_1}a_{j_2}B_{j_2} \cdots a_{j_{m+1}}B_{j_{m+1}}a_{j_{m+2}},$$

since $\phi_{3+2m}(a) = aG(aB)^ma$. This means that if $\|f(x_r)\| \geq 3$, then $S(f)(a_r)$ will traverse a -, B -, and G -edges in different petals. To initiate the folding process, we subdivide the edges according to their image under $S(f)$. On the r -th petal, we see the same sequence of subscripts on every single edge since the subscripts for the images $S(f)(a_r), S(f)(b_r), \dots, S(f)(g_r)$ are all determined by $f(x_r)$. In case $\|f(x_r)\| = 1$, say $f(x_r) = x_{r'}$, then $S(f)$ will just map the r -th petal to the r' -th petal. In this case, there is no need to subdivide the edges in the r -th petal, but we need to relabel the edges y_r as y'_r for $y \in \{a, \dots, g\}$ before we fold.

Recall the first four folding maps we performed in the proof of [Lemma 4.1](#) only fold edges that have the same subscript. To be precise, every fold that is performed through the maps F_1, \dots, F_4 is between edges that lie in the two leftmost edges or in the rightmost edges in each petal. We refer again to [Figure 8](#) for a picture of these folds. Thus, after performing F_1, \dots, F_4 on each petal, we obtain a picture as in the middle of [Figure 10](#). When $\|f(x_r)\| = 1$, the maps F_1, \dots, F_4 are just identity maps on the r -th petal.

The last few folding maps in the proof of [Lemma 4.1](#) were “wrapping” maps that performed a sequence of folds wrapping edge paths of the form $(aB)^m, (cB)^m$, and $(Ca)^m$ around loops aB, cB , and aC , respectively. We could perform the first fold of F_5 because the first a -edge of $(aB)^m$ and the a -edge in the aB loop that the path is wrapped around share an initial vertex. Then, each subsequent fold of b - or a -edges has the same property. The subtlety in performing similar wrapping maps in [Figure 10](#) is that the edge paths now have subscripts, and we must verify that the appropriate edges with the same subscript share an initial vertex so that we can perform each fold in the wrapping maps.

The wrapping map F_5 in [Lemma 4.1](#) folds the edge path corresponding to $(aB)^ma$ in $\phi_{3+2m}(a)$ (all but the first two characters). Recall that

$$S(f)(a_r) = a_{j_1}G_{j_1}a_{j_2}B_{j_2} \cdots a_{j_{m+1}}B_{j_{m+1}}a_{j_{m+2}},$$

and that f is a permutation on the first edges, which implies that after performing $F_4 \circ \cdots \circ F_1$ all of the a_i, b_i , and c_i edges are adjacent to the central vertex in [Figure 10](#) for $i = 1, \dots, n$. This means that we can perform the first two folds on the $a_{j_2}B_{j_2}$ edges of the folding map F_5 , regardless of the value of j_2 . After folding these two edges, the next pair $a_{j_3}B_{j_3}$ in the




Figure 10: ($j_1 = 1, j_{m+2} = 2$). The graphs obtained after doing the four folds F_1, \dots, F_4 on each petal, followed by performing F_5 on a single petal. Note that since the star map f acts as a permutation on the first edges, all of the $a_i, b_i,$ and c_i appear and share the middle vertex as their initial vertex.

edge path now starts at the center vertex in [Figure 10](#) so that we can fold regardless of the value of j_3 . Continuing in this way, we can perform $2m$ of the $2m + 1$ single folds that constitute F_5 . The final fold in F_5 in the proof of [Lemma 4.1](#) is along an a -edge. This final a -edge is labeled by $a_{j_{m+2}}$ and has the center vertex of the flower as its initial vertex after the first $2m$ folds of F_5 (see [Figure 10](#)). Thus, the final fold of F_5 can indeed be performed. The result of the folding map F_5 on the petal corresponding to x_r is shown in [Figure 10](#) with the assumption that $j_1 = 1$ and $j_{m+2} = 2$ for simplicity of the picture. The figure is far more complicated than that in the proof of [Lemma 4.1](#) since edges in the r -th petal are folded with edges in petals corresponding to other x_k 's.

Moving forward, the F_6 map in the proof of [Lemma 4.1](#) is a fold involving the last b -edge in the edge path corresponding to $\phi_{3+2m}(b) = bD(bC)^m b$. The last b -edge in $S(f)(b_r)$ is labelled by $b_{j_{m+2}}$. In fact, due to the folds that were already performed, this edge has been identified with the last b -edge of $S(f)(g_r)$. Due to the final fold in F_5 described in the previous paragraph, this $b_{j_{m+2}}$ edge shares an initial vertex with another $b_{j_{m+2}}$ edge, and so we

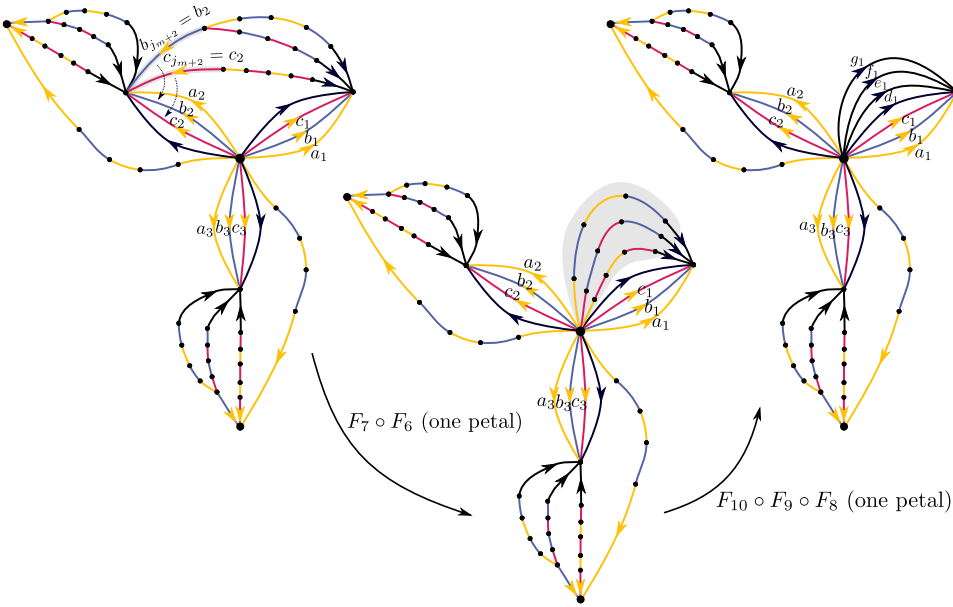


Figure 11: ($j_1 = 1, j_{m+2} = 2$). Continuing from [Figure 10](#), we perform F_6 and F_7 each to fold single edges $b_{j_{m+2}}$ and $c_{j_{m+2}}$ to obtain the middle graph. Then we do three “wrapping-up” folds F_8, F_9 and F_{10} to obtain the rightmost graph, where the petal we chose to fold now looks like a copy of P_7 .

perform the single fold and call it F_6 (see [Figure 11](#)). Similarly, the final edge of the edge path corresponding to $\phi_{3+2m}(c) = cF(cA)^m c$ is labeled by $c_{j_{m+2}}$ and shares an initial vertex with another $c_{j_{m+2}}$ due to the folds of F_5 . We perform this single fold, which we call F_7 . Again, we refer the reader to [Figure 11](#) where we assume $j_1 = 1$ and $j_{m+2} = 2$ for simplicity.

After F_6 and F_7 have been applied, the edge path $c_{j_{m+1}}B_{j_{m+1}} \cdots c_{j_2}B_{j_2}$ starts at the center vertex of the petal, and therefore, we can begin the folds that wrap around the appropriate loops of the form c_iB_i , just as we did in the proof of [Lemma 4.1](#). Again, we are using the fact here that all of the $a_i, b_i,$ and c_i edges begin at the center vertex (see the argument for F_5 above). In the same way, we can perform the F_9 map once F_6 has been applied and wrap the edge path labeled by $a_{j_{m+1}}B_{j_{m+1}} \cdots a_{j_2}B_{j_2}$ around the appropriate a_iB_i loops. Lastly, F_{10} wraps the edge path labeled by $a_{j_{m+1}}C_{j_{m+1}} \cdots a_{j_2}C_{j_2}$ around the appropriate a_iC_i loops, which is possible due to the single fold of F_7 .

Finally, we perform the analogous folds F_5, \dots, F_{10} on each of the other $n - 1$ petals in the rightmost graph of [Figure 11](#). The resulting map ψ

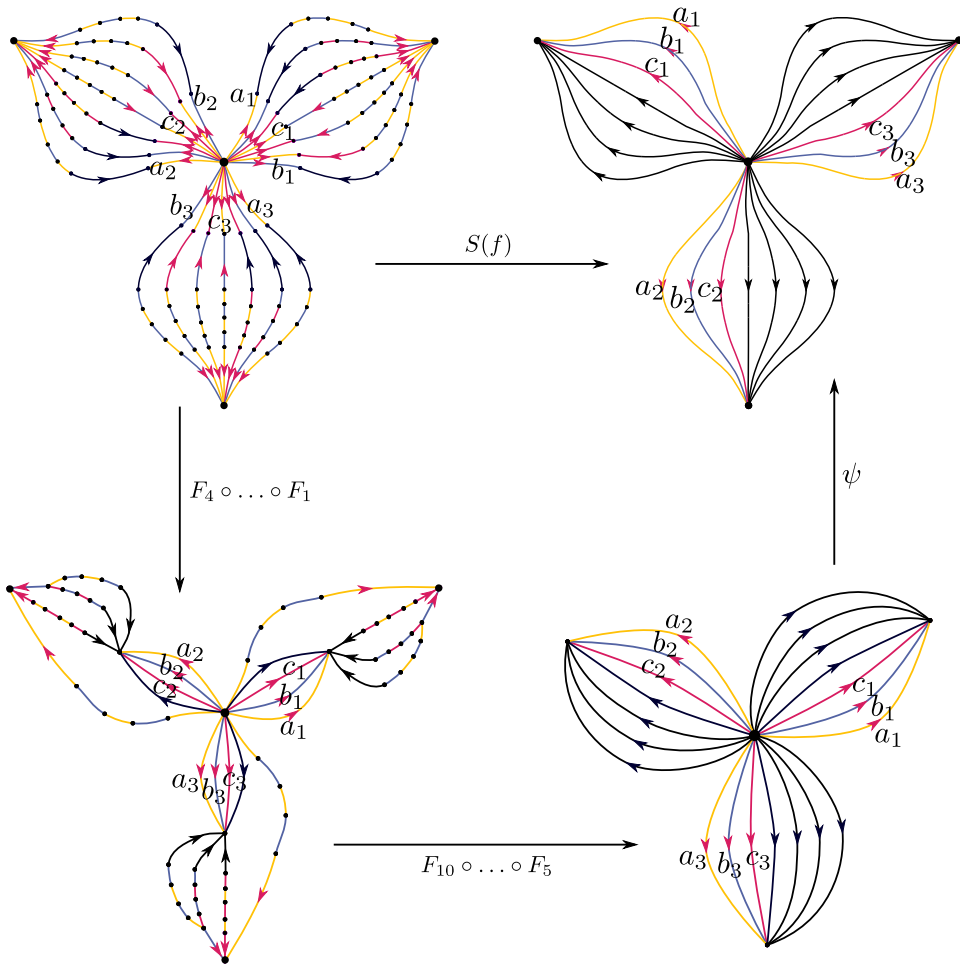


Figure 12: The decomposition of $S(f)$. Note the immersion ψ is the identity map up to permuting petals or edges within petals.

of [Theorem 2.9](#) is the identity map on $S(*_n)$ up to permuting the petals and permuting edges within a petal (a graph automorphism), which is indeed a homotopy equivalence. Since only folds of Type I were performed in F_1, \dots, F_{10} , we conclude that $S(f)$ is a homotopy equivalence. See [Figure 12](#) for the summary. \square

6. Conclusion: Proof of Thurston's Theorem

It remains to show that for any weak Perron or Perron number λ , $S(f_\lambda)$ is ergodic. Note that the proof of this fact was missing in Thurston's paper. We first prove the case where λ is Perron and then extend to weak Perron numbers.

We retain the notation from [Section 3.3](#) so that $*_n = *_{mM}$ consists of m families of M tips. We call each family of M tips a *fan*. We label the tips by $I = (s_k, i)$ where $k \in \{1, \dots, m\}$ and $i \in \{1, \dots, M\}$. Also, when $i < M$ for $I = (s_k, i)$ we write $I + 1 = (s_k, i + 1)$.

Theorem 6.1. *Let λ be a Perron number and f_λ be the star map constructed in [Theorem 3.3](#). Then its split map $S(f_\lambda) : S(*_n) \rightarrow S(*_n)$ is mixing, i.e., for an edge y_I in $S(*_n)$, there exists an L such that $S(f_\lambda)^L(y_I) \supseteq S(*_n)$.*

There are two key aspects of our constructions thus far that will be pivotal in the proof of this lemma: the way the star map f_λ constructed in the proof of [Theorem 3.3](#) acts on the edges (s_k, M) of $*_n$ (the last tip in each fan), and the way prototype maps act on edges in the prototype graph, with particular focus on the fact that for large $m > 0$, $\phi_{3+2m}(a) = aG(aB)^m a$ maps over aB a large number of times. Using these two facts together, we break down the proof of [Theorem 6.1](#) into the following five steps:

- **Step 1:** We first show that the image of y_I under a large enough power of $S(f_\lambda)$ contains *some* a -edge in $S(*_n)$.
- **Step 2:** Then we show that a bigger power of $S(f_\lambda)$ applied to y_I contains *all* a -edges in $S(*_n)$.
- **Step 3:** Next, we show that any further power of y_I under $S(f_\lambda)$ will *still* contain every a -edge in $S(*_n)$.
- **Step 4:** We use this to conclude that there is a larger power of $S(f_\lambda)$ so that the image of y_I contains all a - and g -edges in $S(*_n)$.
- **Step 5:** We then obtain all of the b -, e -, c -, d -, and f -edges in $S(*_n)$ using further powers of $S(f_\lambda)$, in that order, to obtain the mixing conclusion.

Proof. In what follows, we use Γ to denote $*_n = *_{mM}$. We drop the subscript λ from f_λ and simply call its split $S(f)$. Take a random edge in $S(\Gamma)$ labelled by y_I , where $y \in \{a, b, \dots, g\}$ and $I = (s_k, i)$ for some $k \in \{1, \dots, m\}$ and

$i \in \{1, \dots, M\}$. Then, there exists $p \geq 0$ such that $S(f)^p(y_I) = y_K$ where $K = (s_k, M)$. To be precise, we can take $p = M - i$. That is to say, we shift y_I along its fan until it lives on the prototype corresponding to the last edge in this fan. Here we are using the fact that when $i < M$ and $I = (s_k, i)$, then $S(f)$ sends y_I to y_{I+1} . We do this so that the next iterate of $S(f)$ applied to y_I involves the application of a non-identity prototype map.

Recall the prototype maps ϕ_{3+2m} for $m \geq 0$ are given by:

$$\phi_{3+2m} : \begin{cases} a & \mapsto aG(aB)^m a, \\ b & \mapsto bD(bC)^m b, \\ c & \mapsto cF(cA)^m c, \\ d & \mapsto aB(aB)^m a, \\ e & \mapsto cB(aB)^m a, \\ f & \mapsto aC(aB)^m a, \\ g & \mapsto bE(bA)^m b. \end{cases}$$

Step 1: Since $f((s_k, M))$ always maps over at least 5 edges, we use the prototype maps ϕ_{3+2m} with $m \geq 1$ to construct the split map $S(f)$ from f . In addition, the image of an edge under such prototypes will always contain a portion of the form $(wZ)^m$ where $w, z \in \{a, b, c\}$ are distinct. Recall that $f((s_k, M))$ in Γ maps over $(s_k, 1)$ an even number of times, then maps over $(s_\ell, 1)$ for all $\ell \neq k$ an even number of times in some order, then maps over $(s_k, 2)$ exactly twice, and finishes by mapping over $(s_k, n_0 + 1)$ exactly once. Therefore, the image of an edge in $S(\Gamma)$ under $S(f)$ will map over the w - and z -edges of the prototypes corresponding to $(s_\ell, 1)$ for all $\ell \neq k$ and $(s_k, 2)$.

From this, we conclude that, so long as $y \neq b$, $S(f)^{p+1}(y_I)$ contains a_J for some J ; in fact, this image contains several a -edges. This is due to the fact that the image of every other edge besides b in the prototype contains a word of the form $(wZ)^m$ where either w or z is a . If $y = b$, then $\phi_{3+2m}(b) = bD(bC)^m b$, so that $S(f)^{p+1}(y_I)$ contains $d_{(s_k, 1)}$. Therefore, $S(f)^{p+2}(y_I)$ contains $d_{(s_k, 2)}$, $S(f)^{p+M}(y_I)$ contains $d_{(s_k, M)}$, and $S(f)^{p+M+1}(y_I)$ contains a_J for $J = (s_k, 1)$ since $\phi_{3+2m}(d) = aB(aB)^m a$.

Step 2a: However, we claim that since there is a power of $S(f)$ so that the image of y_I contains a_{J_0} for *some* J_0 , then there is a further power, call it $S(f)^q$, so that the image of y_I contains a_J for all $J = (s_j, 1)$, where $j = 1, \dots, m$. Note that this is still a subset of the set of *all* a -edges in $S(*_n)$. Without loss of generality, assume $J_0 = (s_\ell, M)$ for some ℓ (making the second component M can be achieved by applying a few more iterates of $S(f)$). Then, the claim follows from the fact that $\phi_{3+2m}(a) = aG(aB)^m a$ so

that $S(f)(a_{J_0})$ first maps over $a_J G_J$ for $J = (s_\ell, 1)$ and then maps over $a_J B_J$ for $J = (s_j, 1)$ for all $j \neq \ell$. In fact, after mapping over $a_J B_J$ for $J = (s_j, 1)$ and $j \neq \ell$, $S(f)(a_{J_0})$ then maps over $a_J B_J$ for $J = (s_\ell, 2)$. The fact that the image contains $a_{(s_\ell, 2)}$ is essential for showing that the map is mixing and not just ergodic, as we will now see.

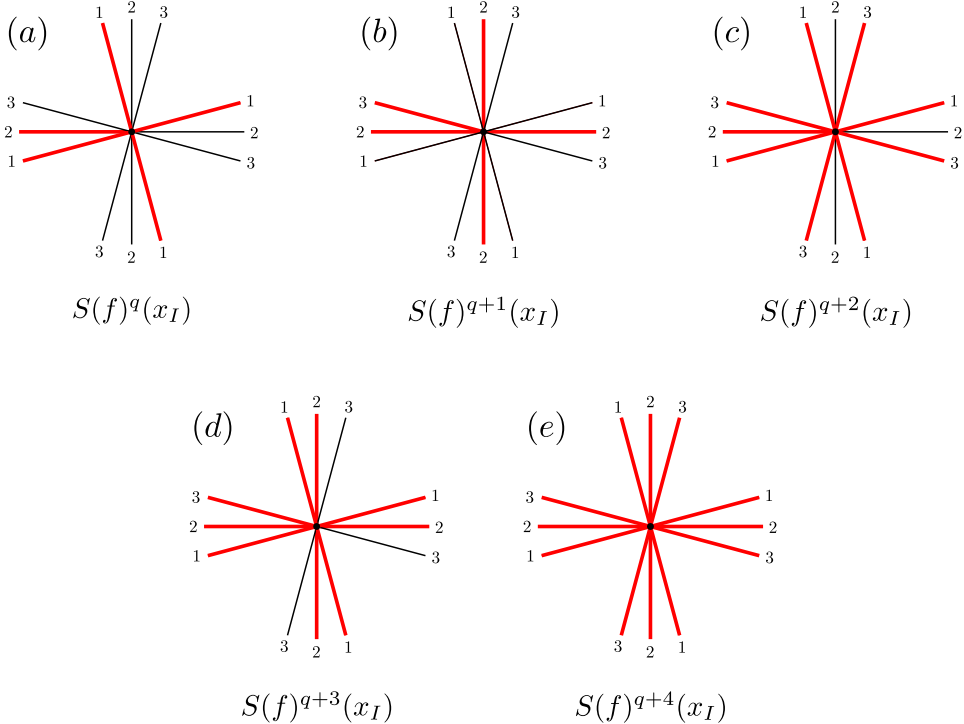


Figure 13: $M = 3, m = 4$. The K -th edge in Γ is thickened and colored in red when $S(f)^n(y_I)$ contains a_K . The figure illustrates how the subsequent powers of $S(f)^q$ will map y_I over *all* edges of the form a_J .

Step 2b: Given the fact that $S(f)^q(y_I)$ contains a_J for all $J = (s_j, 1)$ and $J = (s_\ell, 2)$ for some ℓ , we now show that another further power of $S(f)$ applied to y_I contains *all* a -edges of $S(*_n)$. We outline the argument in the bullet points below for ease of readability.

- First, $S(f)^{q+1}(y_I)$ contains a_J where $J = (s_j, 2)$ for all j and for $J = (s_\ell, 3)$.
- Then, $S(f)^{q+M-2}(y_I)$ contains a_J where $J = (s_j, M - 1)$ for all j and for $J = (s_\ell, M)$.

- In the next iterate of $S(f)$, $a_{(s_j, M-1)}$ maps to $a_{(s_j, M)}$. Additionally, $a_{(s_\ell, M)}$ maps over a_J where $J = (s_j, 1)$ for all j and for $J = (s_\ell, 2)$. In summary, $S(f)^{q+M-1}(y_I)$ contains a_J where $J = (s_j, 1), (s_j, M)$ for all j and for $J = (s_\ell, 2)$. See [Figure 13\(c\)](#).
- Then, $S(f)^{q+M}(y_I)$ contains a_J where $J = (s_j, 1)$ and $(s_j, 2)$ for all j and for $J = (s_\ell, 3)$. See [Figure 13\(d\)](#).
- Applying $S(f)^M$ once more, we see that $S(f)^{q+2M}(y_I)$ contains all a_J where $J = (s_j, 1), (s_j, 2), (s_j, 3)$, for all j and for $J = (s_\ell, 4)$.
- Continuing in this fashion, $S(f)^{q+(M-2)M}(y_I)$ contains all a_J where $J = (s_j, 1), \dots, (s_j, M-1)$ for all j and for $J = (s_\ell, M)$.
- Since $S(f)(a_{(s_\ell, M)})$ contains all a_J of the form $J = (s_j, 1)$, it follows that $S(f)^r(y_I)$, for $r = q + (M-2)M + 1$, contains a_J for all J .

Thus, we have that there exists r such that $S(f)^r(y_I)$ contains a_J for all J concluding Step 2.

Step 3: We claim that any further power of y_I under $S(f)$ will still contain every a -edge in $S(*_n)$. In particular, $S(f)(a_J) = a_{J+1}$ for all $J = (s_j, i)$ when $i < M$, and $S(f)(a_{(s_\ell, M)})$ contains all $a_{(s_j, 1)}$. In this way, we do not lose any a -edges in $S(f)^r(y_I) \subset S(*_n)$ by applying further powers of $S(f)$. See [Figure 14](#).

Step 4: Additionally, we claim that for each j , $S(f)(a_{(s_j, M)})$ also contains $g_{(s_j, 1)}$. This follows from the fact that the first two letters in the image of a under any prototype map (except ϕ_1 ; the identity) are aG and the fact that $f((s_j, M))$ first maps over $(s_j, 1)$ at least twice. Therefore, $S(f)^{r+1}(y_I)$ contains g_J for all $J = (s_j, 1)$ (See [Figure 14\(b\)](#)) and $S(f)^{r+2}(y_I)$ contains g_J for all $J = (s_j, 1), (s_j, 2)$ since $g_{(s_j, 1)}$ maps to $g_{(s_j, 2)}$ and $a_{(s_j, M)}$ maps over $g_{(s_j, 1)}$. See [Figure 14\(c\)](#). Finally, $S(f)^{r+M}(y_I)$ contains a_J and g_J for all J .

Step 5: We continue this procedure to obtain, in order, every b_J and e_J , then c_J and d_J , and finally f_J edge in $S(*_n)$. This order is determined by the first two letters in the images of $\{a, b, \dots, g\}$ under the prototype maps. Therefore, there exists an L such that $f^L(y_I)$ contains all of $S(*_n)$. \square

Proceeding to the case of weak Perron numbers λ , we have:

Theorem 6.2. *Let λ be a weak Perron number and f_λ be the star map constructed in [Lemma 3.9](#). Then its split map $S(f_\lambda) : S(*_n) \rightarrow S(*_n)$ is ergodic, i.e., for each pair of edges y_I, z_J in $S(*_n)$, there exists an L such that $S(f_\lambda)^L(y_I) \supset z_J$.*

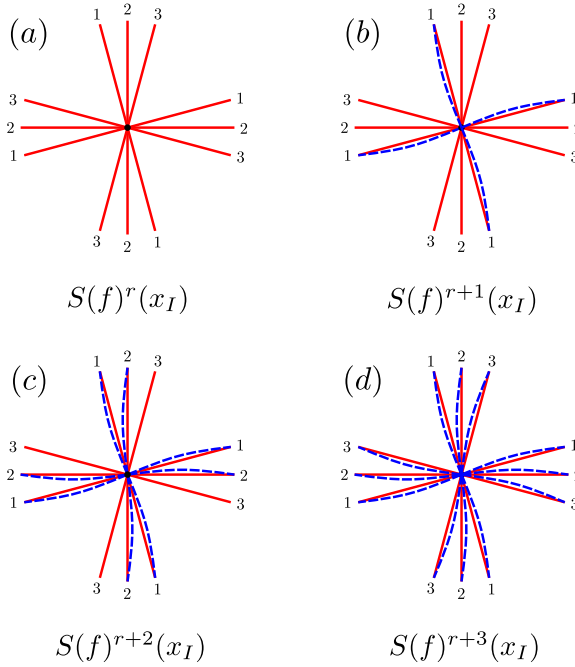


Figure 14: Red(with non-dashed) and blue(with dashed) edges in $S(\Gamma)$ are a - and g -edges in the image of y_I via corresponding power of $S(f)$ written below each graph. The figure illustrates how a subsequent power of $S(f)^r$ will map y_I over *all* edges of the form a_J and g_J . Note the red edges persist throughout the process.

Proof. By Proposition 2.12 there is an integer N such that λ^N is Perron. We use μ to denote λ^N in the remainder of this proof for notational simplicity. Let $f_\mu : *_{n} \rightarrow *_{n}$ be the uniformly μ -expanding star map from Theorem 3.3 and $f_\lambda : *_{Nn} \rightarrow *_{Nn}$ be the uniformly λ -expanding star map constructed using f_μ as in the proof of Lemma 3.9.

Label the edges of $*_{n}$ as $j \in \{1, \dots, n\}$, and the edges of $*_{Nn}$ as (i, j) , where $i \in \{0, \dots, N - 1\}$ and $j \in \{1, \dots, n\}$. Then label the edges of $S(*_{n})$ as y_j and label the edges of $S(*_{Nn})$ as y_j^i 's, where $y \in \{a, b, \dots, g\}$, $i \in \{0, \dots, N - 1\}$ and $j \in \{1, \dots, n\}$.

For each $i = 0, \dots, N - 1$, define C^i as the subgraph of $S(*_{Nn})$ consisting of the edges of the form y_j^i , which is homotopy equivalent to $S(*_{n})$. Call C^0, \dots, C^{N-1} the *fans* of $S(*_{Nn})$. Additionally, for a path γ in $S(*_{n})$, we use $[\gamma]^i$ to denote the copy of this path in the i -th fan C^i of $S(*_{Nn})$.

Recall that $S(f_\mu)$ is mixing (by [Theorem 6.1](#)) and that f_λ is constructed to simply shift C^i to C^{i+1} for all $i < N - 1$ and constructed to apply f_μ to C^{N-1} followed by shifting it to C^0 . We claim that $S(f_\lambda)$ is ergodic.

Now pick two edges y_j^i and $z_{j'}^{i'}$ in $S(*_{Nn})$. Consider the two corresponding edges y_j and $z_{j'}$ of $S(*_n)$. By the fact that $S(f_\mu)$ is mixing (and therefore ergodic), there exists p such that $S(f_\mu)^p(y_j) \supset z_{j'}$. We will show that

$$S(f_\lambda)^{Np-i+i'}(y_j^i) \supset z_{j'}^{i'}.$$

The following diagram breaks down how the power $Np - i + i'$ of $S(f_\lambda)$ is obtained, where the numbers over the arrows denote the power of $S(f_\lambda)$ being applied. We will carefully explain each arrow in what follows:

$$y_j^i \xrightarrow{N-i} [S(f_\mu)(y_j)]^0 \xrightarrow{N(p-1)} [S(f_\mu)^p(y_j)]^0 \xrightarrow{i'} [S(f_\mu)^p(y_j)]^{i'} \supset z_{j'}^{i'}$$

By the construction of f_λ from f_μ , $S(f_\lambda)$ simply translates C^i to C^{i+1} when $i < N - 1$. Thus, we first move y_j^i to y_j^{N-1} using $S(f_\lambda)^{N-1-i}$. Applying $S(f_\lambda)$ one more time brings y_j^{N-1} to $[S(f_\mu)(y_j)]^0$. This is due to the fact that $S(f_\lambda)$ amounts to applying $S(f_\mu)$ to C^{N-1} and then translating the image to C^0 . In summary, $S(f_\lambda)^{N-1-i+1}(y_j^i) = S(f_\lambda)^{N-i}(y_j^i) = [S(f_\mu)(y_j)]^0 \subset C^0$.

Next, notice that applying $S(f_\lambda)^N$ to this path amounts to applying $S(f_\mu)$ once to C^0 . Thus,

$$S(f_\lambda)^{N-i+N(p-1)}(y_j^i) = S(f_\lambda)^{Np-i}(y_j^i) = [S(f_\mu)^p(y_j)]^0.$$

Moreover, $[S(f_\mu)^p(y_j)]^0 \supset z_{j'}^0$, given the fact that $S(f_\mu)^p(y_j) \supset z_{j'}$. Lastly, we apply $S(f_\lambda)^{i'}$ to bring $z_{j'}^0$ to $z_{j'}^{i'}$. Thus, $S(f_\lambda)^{Np-i+i'}(y_j^i) \supset z_{j'}^{i'}$, proving that $S(f_\lambda)$ is ergodic. \square

Remark 6.3. The map $S(f_\lambda)$ from the construction in the proof of [Theorem 6.2](#) is *not* mixing. Indeed, $S(f_\lambda)^k(y_j^i)$ is completely contained in C^m where $m \equiv i + k$ modulo N .

We conclude the paper with the proof of Thurston's main theorem [Theorem 1.1](#).

Proof of [Theorem 1.1](#). The forward direction is covered in the last paragraph of [Section 2.3](#). That is, if h is the topological entropy of an ergodic traintrack representative of an outer automorphism of a free group, then e^h is a weak Perron number.

Conversely, let λ be a weak Perron number. Then using [Lemma 3.9](#) we find a uniformly λ -expanding star map $f_\lambda : *_k \rightarrow *_k$ for some k . Splitting f_λ we obtain our desired map $S(f_\lambda) : S(*_k) \rightarrow S(*_k)$. This is a traintrack map by [Proposition 5.5](#) with respect to the traintrack structure given in [Section 5.1](#), and is indeed a topological representative of an outer automorphism by [Proposition 5.6](#). Also, $S(f_\lambda)$ is still uniformly λ -expanding by [Proposition 5.4](#), and is ergodic by [Theorem 6.2](#). Finally, by [Proposition 2.5](#), the topological entropy of $S(f_\lambda)$ is indeed $\log \lambda$, concluding the proof. \square

Acknowledgements

The authors would like to thank Jack Cook, Sean Howe, Michael Kopreski, and Kurt Vinhage for many helpful discussions and thank Rylee Lyman for helpful comments on the earlier draft of this paper. We would also like to thank Mladen Bestvina for his suggestion that we undertake this project during a yearlong seminar supported by the NSF RTG Grant DMS–1840190. We also thank the referee for numerous insightful comments and suggestions.

The authors acknowledge support from NSF grants RTG DMS–1745583 (Dickmann), DMS–1905720 (Domat, Kwak), RTG DMS–1840190 (Domat, Rechkin), CAREER DMS–2046889 (Kwak, Patel), CAREER DMS–1452762 (Ospina), the University of Utah Faculty Fellows Award (Patel), and a KIAS Individual Grant (HP098501) via the June E Huh Center for Mathematical Challenges at Korea Institute for Advanced Study (Kwak).

References

- [1] L. Alsedà, J. Llibre, and M. Misiurewicz, *Combinatorial dynamics and entropy in dimension one*, Vol. 5 of *Advanced Series in Nonlinear Dynamics*, World Scientific Publishing Co., Inc., River Edge, NJ, second edition (2000), ISBN 981-02-4053-8.
- [2] M. Bestvina, *The Topology of $Out(F_n)$* , Proceedings of the International Congress of Mathematicians (2002, Beijing) **2** (2002) 373–384.
- [3] M. Bestvina and M. Handel, *Train Tracks and Automorphisms of Free Groups*, *Annals of Mathematics* **135** (1992), no. 1, 1–51.
- [4] M. Clay, *Automorphisms of Free Groups*, in M. Clay and D. Margalit, editors, *Office Hours with a Geometric Group Theorist*, chapter 6, Princeton University Press (2017).

- [5] D. Fried, *Growth rate of surface homeomorphisms and flow equivalence*, Ergodic Theory Dynam. Systems **5** (1985), no. 4, 539–563.
- [6] L. Liechti and J. Pankau, *The Geometry of Bi-Perron Numbers with Real or Unimodular Galois Conjugates*, International Mathematics Research Notices (2021).
- [7] D. Lind and B. Marcus, *An Introduction to Symbolic Dynamics and Coding*, Cambridge Mathematical Library, Cambridge University Press, 2 edition (2021).
- [8] D. A. Lind, *The entropies of topological Markov shifts and a related class of algebraic integers*, Ergodic Theory and Dynamical Systems **4** (1984), no. 2, 283–300.
- [9] D. A. Marcus, *Number fields*, Universitext, Springer, Cham (2018), ISBN 978-3-319-90232-6; 978-3-319-90233-3. Second edition.
- [10] W. Narkiewicz, *Elementary and analytic theory of algebraic numbers*, Springer Monographs in Mathematics, Springer-Verlag, Berlin, third edition (2004), ISBN 3-540-21902-1.
- [11] J. Pankau, *Salem number stretch factors and totally real fields arising from Thurston's construction*, Geom. Topol. **24** (2020), no. 4, 1695–1716.
- [12] J. R. Stallings, *Topology of finite graphs*, Inventiones mathematicae **71** (1983), no. 3, 551–565.
- [13] W. P. Thurston, *Entropy in dimension one*, in *Frontiers in complex dynamics*, Vol. 51 of *Princeton Math. Ser.*, 339–384, Princeton Univ. Press, Princeton, NJ (2014).

DEPARTMENT OF MATHEMATICS, GEORGIA INSTITUTE OF TECHNOLOGY
ATLANTA, GA 30332, USA

E-mail address: rdickmann3@gatech.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MICHIGAN
ANN ARBOR, MI 48109, USA

E-mail address: domatg@umich.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF UTAH
SALT LAKE CITY, UT 84112, USA

E-mail address: thill@math.utah.edu

JUNE E HUH CENTER FOR MATHEMATICAL CHALLENGES, KIAS
SEOUL, 02455, SOUTH KOREA

E-mail address: samuelkwak@kias.re.kr

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF UTAH
SALT LAKE CITY, UT 84112, USA

E-mail address: ospina@math.utah.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF UTAH
SALT LAKE CITY, UT 84112, USA

E-mail address: patelp@math.utah.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF UTAH
SALT LAKE CITY, UT 84112, USA

E-mail address: rebecca.rechkin@utah.edu

RECEIVED OCTOBER 3, 2022

ACCEPTED MARCH 29, 2023