LOW COMPLEXITY AMONG PRINCIPAL FULLY IRREDUCIBLE ELEMENTS OF $\mathrm{Out}(F_3)$

NAOMI ANDREW, PAIGE HILLEN, RYLEE ALANZA LYMAN, AND CATHERINE PFAFF

Abstract. We find the shortest realized stretch factor for a fully irreducible $\varphi \in \mathrm{Out}(F_3)$ and show that it is realized by a "principal" fully irreducible element. We also show that it is the only principal fully irreducible produced by a single fold in any rank.

1. Introduction

Let F_r denote the free group of rank $r \geq 2$, and consider its outer automorphism group $\mathrm{Out}(F_r)$. Each element of $\mathrm{Out}(F_r)$ has an associated stretch factor (also called growth rate or dilatation)

$$\sup_{w \in F_r} \limsup \sqrt[n]{\|\varphi^n(w)\|}.$$

Here $\|w\|$ denotes the cyclically reduced word length with respect to some fixed basis. The stretch factor records how fast elements grow under iteration of φ, and is independent of the chosen basis and the representative of the outer class.

For a fully irreducible element, where no power φ^k preserves a proper free factor, Bestvina and Handel [BH92] construct an "irreducible train track representative" for φ, a self homotopy equivalence of a graph with good behavior under iteration which induces φ on the fundamental group. They show that the stretch factor of φ is realised as the Perron–Frobenius eigenvalue of the transition matrix, a non-negative integer matrix associated to the train track map.

Here we are concerned with two aspects of this theory. The first aspect is the minimal stretch factors attained by fully irreducible outer automorphisms. The set of possible values is studied, for example, in [AKR15], [DKL15], [DKL17], and [Thu14] (exposited upon in [DDH$^+$22]). The second aspect explores properties of train track representatives for certain fully irreducible outer automorphisms, and the dynamics of their action on the Culler–Vogtmann Outer space CV_r.

In both aspects, we investigate the outer automorphism

$$\psi = \begin{cases} x \mapsto y \\ y \mapsto z \\ z \mapsto zx^{-1}. \end{cases}$$

With a suitable marking, it is represented by the train track map:

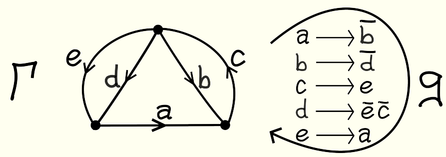

Its stretch factor is the real root of $x^5 - x - 1$, approximately 1.167. In fact, ψ is a fully irreducible outer automorphism, and we show:

Theorem A. *The stretch factor of ψ is minimal among fully irreducible elements in* $\mathrm{Out}(F_3)$.

The ingredients in the Theorem A proof are Lemma 5.1, showing stretch factors of fully irreducible elements are Perron numbers, and verification that $\mathfrak{g}$ is an irreducible train track representative of a fully irreducible outer automorphism realising the minimum achievable Perron number.

We also prove the map ψ is *principal* in the sense of [AKKP19], where it is proved they possess a stability property mimicking that held by pseudo-Anosov mapping classes and used in [KMPT22b] and [KMPT22a] to understand a typical outer automorphism and tree in $\partial\mathrm{CV}_r$. Their action on Culler–Vogtmann Outer space generally more closely resembles a hyperbolic setting: While all fully irreducible elements act loxodromically, they only have an axis *bundle* rather than a unique axis [HM11]. But principal fully irreducible outer automorphisms do have a unique axis, reflecting the uniqueness of the Stallings fold decompositions of their irreducible train track representatives. This axis passes through the highest dimensional simplices of Outer space.

The name is chosen to reflect similarities with the principal pseudo-Anosov mapping classes of a surface, as used for instance by Masur in [Mas82]. Every pseudo-Anosov mapping class has a representative leaving invariant a pair of transverse measured singular minimal foliations, one of which is expanded and the other contracted by the homeomorphism. Pseudo-Anosovs act loxodromically on Teichmüller space, and these two foliations provide the endpoints of its axis in the Thurston boundary. A pseudo-Anosov is said to be *principal* if these foliations have only 3-pronged singularities. Gadre and Maher [GM17] proved principal pseudo-Anosov mapping classes random-walk generic. The "singularity structure" of a fully irreducible outer automorphism, namely its ideal Whitehead graph, is similarly controlled by its attracting endpoint, and in the case of a principal element will be a union of the maximal number of triangles. As such, principal elements are the subset of those proved generic in [KMPT22b] that possess the stability property of [AKKP19].

Given a train track map, we can factor it as a series of *folds* followed by a homeomorphism. Up to a reasonable equivalence relation, the number of folds in this decomposition is an invariant of a principal fully irreducible outer automorphism, and so we can try to minimise it. It turns out that the same automorphism achieves the minimum complexity — over all ranks — in this sense.

Theorem B. *Up to edge relabeling, the map $\mathfrak{g}$ is the only train track map representing a principal fully irreducible outer automorphism in any rank whose Stallings fold decomposition consists of only a single fold and then a graph-relabeling isomorphism.*

The tool we use for characterising rank-3 principal fully irreducible outer automorphisms is a directed graph $\widehat{\mathcal{A}_3}$ we call the *principal stratum automaton* (§3). This is a refinement of the lonely direction automaton of [GP23], which was used to determine the graphs carrying train track representatives of principal fully irreducible elements of $\mathrm{Out}(F_3)$. That theorem was proved by passing to a *rotationless* power of φ, which is unhelpful for minimising other kinds of complexity such as the stretch factor or length of a Stallings fold decomposition. Our refinement allows us to work with the automorphism as it comes, allowing for the periodic behavior absent in the rotationless case.

Theorem C. *Suppose g is a train track representative of a principal fully irreducible $\varphi \in \mathrm{Out}(F_3)$. Then the Stallings fold decomposition of g is partial-fold conjugate to one determining a directed loop in $\widehat{\mathcal{A}_3}$.*

In the process, we also prove in Proposition 3.4 that any principal axis in CV_3 passes through the $\mathrm{Out}(F_3)$ orbit of simplices given by the graph shown above.

Structure of the paper. In §3 we introduce the principal stratum automaton, and prove that it captures the principal fully irreducible elements of $\mathrm{Out}(F_3)$. §4 introduces the map $\mathfrak{g}$ which is the subject of both theorems, and proves that it induces a principal fully irreducible outer automorphism. §5 contains the proof of Theorem A and §6 the proof of Theorem B.

Acknowledgements. We are grateful to the Women in Groups Geometry and Dynamics (WiGGD) program from which this paper arose, and to Anna Parlak for discussions on related questions. We are grateful to Lee Mosher for providing his talents as a rabbit-hole preventing sounding board, and to Ilya Kapovich for helping us track down a reference and his ongoing interest in our work. This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 850930), an NSF postdoctoral fellowship, and an NSERC Discovery Grant.

2. Background

Assume throughout this section Γ is a finite oriented graph where each vertex has valence at least 3 and F_r is a free group of rank $r \geq 3$.

2.1. Edge Maps on Graphs.

Suppose Γ has positively oriented edges $\{e_1, e_2 \ldots, e_n\}$ and vertices $\{v_1, v_2, \ldots, v_m\}$. We use the notation $E\Gamma := \{e_1, e_2 \ldots, e_n\}$, and $V\Gamma := \{v_1, v_2, \ldots, v_m\}$, and $E^{\pm}\Gamma = \{e_1, \overline{e_1}, \ldots, e_n, \overline{e_n}\}$, with an overline indicating a reversal of orientation.

Given $v \in V\Gamma$, a *direction at* v will mean an element of $E^{\pm}\Gamma$ with initial vertex v. We let $\mathcal{D}\Gamma$ denote the set of directions at vertices in Γ. A *turn at* v will mean an unordered pair $\{d_1, d_2\}$ of directions at v. The turn is *degenerate* if $d_1 = d_2$.

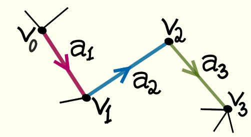

An *edge path* (or just *path*) ρ in Γ is a finite sequence $(a_1, a_2, \ldots, a_\ell) \in (E^{\pm}\Gamma)^\ell$ such that there exists a sequence $(v_1, v_2, \ldots, v_{\ell-1}) \in (V\Gamma)^{\ell-1}$ satisfying that the turn $\{\overline{a_j}, a_{j+1}\}$ is a turn at v_j for each $j \in \{1, 2, \ldots, \ell-1\}$. For such a path $(a_1, a_2, \ldots, a_m)$ we write $\gamma = a_1 a_2 \ldots a_m$ and say γ *contains* the oriented edges $a_1, a_2, \ldots, a_m$ and *takes* the turns $\{\overline{a_1}, a_2\}$, $\{\overline{a_2}, a_3\}$, $\ldots$, $\{\overline{a_{n-1}}, a_n\}$. We call γ *tight* if it takes no degenerate turns, which we colloquially describe as there being no "backtracking."

An *edge (or graph) map* $g : \Gamma \to \Gamma'$ will mean

- a map $\mathcal{V} : V\Gamma \to V\Gamma'$, where we write $g(v)$ for $\mathcal{V}(v)$, together with
- for each $e \in E^{\pm}\Gamma$, an assignment of a path $g(e)$ in Γ' such that
  (1) if the initial vertex of e is v, then the initial vertex of $g(e)$ is $g(v)$, and
  (2) if $g(e)$ is the edge path $g(e) = a_1 a_2 \ldots a_m$, then $g(\overline{e})$ is the edge path concatenation $g(\overline{e}) = \overline{a_m} \ldots \overline{a_2}\ \overline{a_1}$.

Viewing Γ and Γ' as topological spaces, g is a continuous map sending vertices to vertices. We say a turn $\{d_1, d_2\}$ is g-taken if it appears in the image $g(e)$ of some edge of Γ, and call g *tight* if the image of each edge is a tight path. In particular, no degenerate turns are g-taken.

If $\gamma = a_1 a_2 \ldots a_n$ is a path in Γ for some $a_1, a_2, \ldots, a_n \in E^{\pm}\Gamma$, then $g(\gamma)$ will mean the concatenation of edge paths $g(\gamma) = g(a_1)g(a_2)\ldots g(a_n)$. Note that $g(\gamma)$ is tight if and only if γ is tight and g is locally injective on γ.

To g we associate a *direction map* $Dg : \mathcal{D}\Gamma \to \mathcal{D}\Gamma'$ such that if $g(e) = a_1 a_2 \ldots a_m$, for some $m \geq 1$ and $a_1, a_2, \ldots, a_m \in E^{\pm}\Gamma$, then $Dg(e) = a_1$. We call a direction e *periodic* if $Dg^k(e) = e$ for some $k > 0$, and *fixed* if $k = 1$. When $g\colon \Gamma \to \Gamma$ is a self-map, the turn $\{d_1, d_2\}$ is called an *illegal turn* for g if $\{Dg^k(d_1), Dg^k(d_2)\}$ is degenerate for some k. Defining an equivalence relation on $\mathcal{D}\Gamma$ by $d_1 \sim d_2$ when $\{d_1, d_2\}$ is an illegal turn, the equivalence classes are called *gates*. Note that each gate contains a unique periodic direction.

Viewing g as a continuous map of graphs, we say g *represents* φ when $\pi_1(\Gamma)$ has been identified with F_r (that is, Γ is *marked*) and φ is the induced map of fundamental groups. When a marking is not explicitly given, we mean that "there exists a marking such that."

2.2. Train track maps and fully irreducible outer automorphisms.

Suppose $g : \Gamma \to \Gamma$ is an edge map. We call g a *train track (tt) map* if g^k is tight for each $k \in \mathbb{Z}_{>0}$. We call the train track map g *expanding* if for each edge $e \in E\Gamma$ we have $|g^n(e)| \to \infty$ as $n \to \infty$, where for a path

γ we use $|\gamma|$ to denote the number of edges γ traverses (with multiplicity). Note that, apart from our not requiring a "marking," these definitions coincide with those in [BH92] when g is in fact a homotopy equivalence of graphs (viewed topologically).

The *transition matrix* $M(g)$ of a tt map $g\colon \Gamma \to \Gamma$ is the square $|E\Gamma| \times |E\Gamma|$ matrix $[a_{ij}]$ such that a_{ij}, for each i and j, is the number of times $g(e_i)$ contains either e_j or $\overline{e_j}$. Note that each transition matrix is a nonnegative integer matrix.

A nonnegative integral matrix $A = [a_{ij}]$ is *irreducible* if for each (i, j), there is a $k \in \mathbb{Z}_{>0}$ so that the ij^{th} entry of A^k is positive, and so in particular at least 1. If A^k is strictly positive for some $k \in \mathbb{Z}_{>0}$ then A is *primitive*. Further, A is *Perron–Frobenius (PF)* if there exists an N such that, for each $k \geq N$, we have that A^k is strictly positive.

For nonnegative integral matrices, being primitive is equivalent to being irreducible aperiodic. If $M = M(g)$ is primitive, then by Perron–Frobenius theory, M has a unique eigenvalue $\lambda(g)$ of maximal modulus, $\lambda(g)$ is real, and $\lambda(g) > 1$. This $\lambda(g)$ is the Perron–Frobenius (PF) eigenvalue of M and is called the *stretch factor* (or *dilatation*) of g.

A tt map is *irreducible* if its transition matrix is irreducible. Not every element of $\mathrm{Out}(F_r)$ is represented by a train track map, and even fewer by irreducible train track maps. An outer automorphism $\varphi \in \mathrm{Out}(F_r)$ is *fully irreducible* if no positive power preserves the conjugacy class of a proper free factor of F_r. Bestvina and Handel [BH92] proved that every fully irreducible outer automorphism admits expanding irreducible train track representatives and that the stretch factor of these representatives is an invariant of the outer automorphism.

By [FH11, Corollary 4.43], for each $r \geq 2$, there exists an $R(r) \in \mathbb{N}$ such that for each expanding irreducible train track representative g of a fully irreducible $\varphi \in \mathrm{Out}(F_r)$, among other properties, each periodic direction is fixed by $g^{R(r)}$. This power R is called the *rotationless* power.

2.3. **Whitehead graphs & lamination train track (ltt) structures.** Local Whitehead graphs, stable Whitehead graphs, and ideal Whitehead graphs were introduced in [HM11]. We give definitions here only in the circumstance of no periodic Nielsen paths (PNPs), as this will always be the case for us. (PNPs only impact the ideal Whitehead graph definition; since we do not explicitly use them, we refer the reader to [BH92, BFH00] for definitions.)

Let $g : \Gamma \to \Gamma$ be a tt map. The *local Whitehead graph* $\mathrm{LW}(g; v)$ at a $v \in V\Gamma$ has a vertex for each direction at v and edge connecting the vertices corresponding to a pair of directions $\{d_1, d_2\}$ at v precisely when the turn $\{d_1, d_2\}$ is g^k-taken for some $k \in \mathbb{Z}_{>0}$. Given a fixed vertex v, the *stable Whitehead graph* $\mathrm{SW}(g; v)$ is the restriction of $\mathrm{LW}(g; v)$ to the periodic direction vertices and the edges between them. In terms of gates, $\mathrm{SW}(g; v)$ has a vertex for each gate at v.

In the absence of PNPs, if g represents a fully irreducible outer automorphism φ, then the *ideal Whitehead graph* $\mathrm{IW}(\varphi)$ for φ is defined as

$$\mathrm{IW}(\varphi) \cong \bigsqcup_{v \in V\Gamma} \mathrm{SW}(g; v),$$

but with components containing only 2 vertices removed.

The ideal Whitehead graph is an invariant of the conjugacy class of the outer automorphism represented by g and $\mathrm{IW}(\varphi^k) = \mathrm{IW}(\varphi)$ for each $k \in \mathbb{Z}_{>0}$ [HM11, Pfa12].

The *lamination train track (ltt) structure* $G(g)$ is obtained from its *underlying graph* Γ by replacing each vertex $v \in V\Gamma$ with $\mathrm{LW}(g; v)$: replace v with a vertex for each directed edge at v labeled with that direction. Then identify each of these new vertices with the corresponding vertex of $\mathrm{LW}(g; v)$. Vertices and edges of $\mathrm{SW}(g; v)$ are colored purple and the remaining vertices and (open) edges are colored red. Alternatively, one could start with $\bigsqcup_{v \in V\Gamma} \mathrm{LW}(g; v)$, color the $\mathrm{LW}(g; v)$ as just described, and then include a directed edge $[e, \overline{e}]$ for each directed edge $e \in E\Gamma$. To simplify figures, if the local Whitehead graph at a vertex of Γ is complete we do not always draw it.

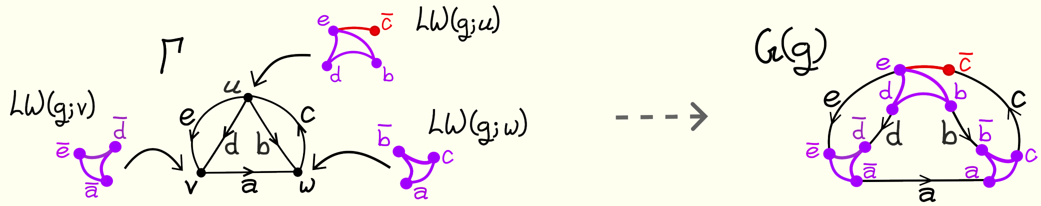

FIGURE 1. The right-hand image is the ltt structure $G(\mathfrak{g})$ for the map $\mathfrak{g}$ of (11) with the taken turns as in (12). The 3 local Whitehead graphs are colored, with the stable Whitehead graphs in purple (missing only $\bar{c}$). The ideal Whitehead graph is the union of the purple graphs, i.e. of the stable Whitehead graphs.

2.4. **Full irreducibility criterion.** We use the following criterion for proving that a tt map represents a fully irreducible outer automorphism.

Proposition 2.1 (Full Irreducibility Criterion, [Pfa13, Proposition 4.1])**.** *Suppose that $g\colon \Gamma \to \Gamma$ is a PNP-free, irreducible train track representative of $\varphi \in \mathrm{Out}(F_r)$ such that $M(g)$ is Perron-Frobenius and all the local Whitehead graphs are connected. Then φ is fully irreducible.*

2.5. **Principal fully irreducible outer automorphisms.** Principal fully irreducible outer automorphisms (sometimes just called principal outer automorphisms) were introduced in [AKKP19] as analogues of pseudo-Anosov homeomorphisms with only 3-pronged singularities: As in [AKKP19], we call a fully irreducible $\varphi \in \mathrm{Out}(F_r)$ *principal* if $\mathrm{IW}(\varphi)$ is the disjoint union of $2r-3$ triangles.

We include now a short description of properties of principal fully irreducible elements omitting definitions not used elsewhere in this paper. The interested reader can find the definitions in [MP16], which also includes a nice explanation of the history at the start of §2.9. For a principal $\varphi \in \mathrm{Out}(F_r)$, the rotationless index is $i(\varphi) = \frac{3}{2} - r$. This implies φ is ageometric and, by [MP16, Theorem 4.5], all its tt representatives are stable. In particular, none of its tt representatives has a PNP. By [MP16, Theorem 4.7] one has that φ will also have only a single axis.

2.6. **Folds & Stallings fold decompositions.** Suppose Γ and Γ' are graphs viewed topologically and $e_0, e_1 \in E^{\pm}\Gamma$ are distinct edges emanating from a common vertex. We say Γ' is obtained from Γ by a *proper full fold of e_1 over e_0* if there exist orientation-preserving homeomorphisms $\sigma_0\colon [0,1] \to e_0$ and $\sigma_1\colon [0,2] \to e_1$ so that $\Gamma' = \Gamma\backslash \sim$ is the topological quotient of Γ with respect to the equivalence relation $\sim$ defined by $\sigma_0(t) = \sigma_1(t)$ for each $t \in [0,1]$. We say that Γ' is obtained from Γ by a *complete fold of e_0 and e_1* if instead $\sigma_1\colon [0,1] \to e_1$ and a *partial fold of e_0 and e_1* if instead $\sigma_0\colon [0,2] \to e_0$.

Since they appear often, we describe here the notational conventions we use for proper full folds. Suppose that $f\colon \Gamma \to \Gamma'$ is a single proper full fold of an edge e_1 over an edge e_0, as depicted to the right. Apart from e_1, each edge $e_k \in E\Gamma$ is mapped to a single edge of Γ', which we call e_k'. The image of e_1 is an edge-path in Γ' consisting of 2 edges, the latter of which we call e_1'. The map f is then defined by

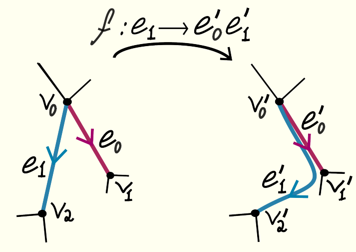

$$f = \begin{cases} e_1 \mapsto e_0' e_1' \\ e_k \mapsto e_k' \text{ for } k \neq 1 \end{cases}$$

and we just write $f\colon e_1 \mapsto e_0 e_1$. We call the edge-labeling of Γ' just described the *induced edge-labeling.* Note that it can be more convenient, especially where an orientation was chosen in advance, to write $f\colon e_1 \mapsto e_1 e_0$, where this was induced by folding $\bar{e}_1$ over $\bar{e}_0$. We may drop primes in complicated diagrams.

Stallings [Sta83] showed that a surjective homotopy equivalence graph map $g\colon \Gamma \to \Gamma'$ factors as a composition of folds and a final homeomorphism, giving a *Stallings fold decomposition.*

$$\begin{tikzcd}
\Gamma = \Gamma_0 \arrow[r, "g_1"'] \arrow[rrrr, bend left, "g=\mathfrak{g}_0"] & \Gamma_1 \arrow[r, "g_2"'] \arrow[rrr, bend left, "\mathfrak{g}_1"] & \Gamma_2 \arrow[r, "g_3"'] \arrow[rr, bend left, "\mathfrak{g}_2"] & \dots \arrow[r, "g_n"'] & \Gamma_n = \Gamma'
\end{tikzcd}$$

In a Stallings fold decomposition the folds g_{k+1} (sometimes called *Stallings folds*) of e_0 and e_1 in Γ_k additionally satisfy $\mathfrak{g}_k(\sigma_0(t)) = \mathfrak{g}_k(\sigma_1(t))$ for each $t \in [0,1]$ and that the terminal vertices of e_0 and e_1 are distinct points in $\mathfrak{g}_k^{-1}(V\Gamma')$. The latter property ensures g is a homotopy equivalence.

In general, each fully irreducible $\varphi \in \mathrm{Out}(F_r)$ has many train track representatives, each of which can have several distinct Stallings fold decompositions. By the main theorem of [MP16, Theorem 4.7], principal outer automorphisms (and their powers) each have a unique Stallings fold decomposition. Further, none of their train track representatives has a PNP.

2.7. **The Outer space CV_r & its (principal) geodesics.** Culler–Vogtmann Outer space was first defined in [CV86]. We refer the reader to [FM11, Bes14, Vog15] for background on Outer space, giving only an abbreviated discussion here. For $r \geq 2$ denote the (volume-1 normalized) Outer space for F_r by CV_r. Points of CV_r are equivalence classes of volume-1 marked metric graphs $h : R_r \to \Gamma$ where R_r is the r-rose, and Γ is a finite volume-1 metric graph with betti number $b_1(\Gamma) = r$ and with all vertices of degree at least 3, and h is a homotopy equivalence called a *marking*. Outer space CV_r has a simplicial complex structure with some faces missing: there is an open simplex for each marked graph (obtained by varying the lengths on the edges of that graph). The faces of a simplex are obtained by collapsing the edges of a forest (so that their lengths become zero). There is an asymmetric metric d_{CV_r} on CV_r.

There is an action (by isometries) of $\mathrm{Out}(F_r)$ on CV_r, given by changing the marking. We denote the quotient by $\mathcal{M}_r$.

Given a Stallings fold decomposition of a train track map g, one can define a "periodic fold line" in CV_r using Skora's [Sko89] interpretation of the decomposition as a sequence of folds performed continuously. In [AKKP19, Lemma 2.27] it is proved that periodic fold lines associated to train track maps are geodesics in the sense that, given 3 points $\gamma(t_1), \gamma(t_2), \gamma(t_3)$ on the geodesic γ with $t_1 < t_2 < t_3$, we have $d_{\mathrm{CV}_r}(\gamma(t_1), \gamma(t_2)) + d_{\mathrm{CV}_r}(\gamma(t_2), \gamma(t_3)) = d_{\mathrm{CV}_r}(\gamma(t_1), \gamma(t_3))$.

Since a general fully irreducible $\phi \in \mathrm{Out}(F_r)$ can have several distinct Stallings fold decompositions, it can have several distinct periodic fold lines. However, the principal outer automorphisms that we are interested in will have just a single periodic fold line, its "lone axis."

2.8. **Fold-conjugate decompositions.** Since an axis for a fully irreducible $\varphi \in \mathrm{Out}(F_r)$ has a periodic structure, it becomes useful to view its Stallings fold decompositions cyclically. With some work one can see that starting at a different fold in a decomposition now yields a tt map representing an outer automorphism $\mathrm{Out}(F_r)$-conjugate to φ and with the same axis. It is also possible to start the tt map "in the middle of a fold." We formalize here these different notions of cyclically permuting a Stallings fold decomposition or, equivalently, shifting along an axis.

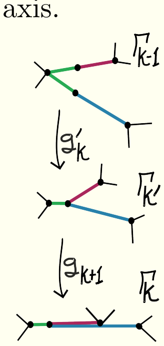

A *subdivided fold* will mean a fold written as a composition of 2 folds, as depicted to the right. Suppose $\Gamma_0 \xrightarrow{g_1} \Gamma_1 \xrightarrow{g_2} \cdots \xrightarrow{g_{k-1}} \Gamma_{k-1} \xrightarrow{g_k} \Gamma_k$ and $\Gamma'_0 \xrightarrow{h_1} \Gamma'_1 \xrightarrow{h_2} \cdots \xrightarrow{h_{n-1}} \Gamma'_{n-1} \xrightarrow{h_n} \Gamma'_n$ are Stallings fold decompositions of homotopy equivalence tt maps $g \colon \Gamma \to \Gamma$ and $h \colon \Gamma' \to \Gamma'$. We say the decompositions are *fold-conjugate* if, possibly after a fold subdivision of g_k into $\Gamma_{k-1} \xrightarrow{g'_k} \Gamma_{k'} \xrightarrow{g_{k+1}} \Gamma_k$ or h_n into $\Gamma'_{n-1} \xrightarrow{h'_n} \Gamma'_{n'} \xrightarrow{h_{n+1}} \Gamma'_n$, we have for some j that one the following diagrams (1)-(3) commutes, where the vertical arrows are label-preserving graph isomorphisms.

(1)
$$\begin{tikzcd}
\Gamma_j \arrow[r, "g_{j+1}"] \arrow[d] & \Gamma_{j+1} \arrow[r, "g_{j+2}"] \arrow[d] & \dots \arrow[r, "g_k"] & \Gamma_k = \Gamma_0 \arrow[r, "g_1"] \arrow[d] & \dots \arrow[r, "g_{j-1}"] & \Gamma_{j-1} \arrow[r, "g_j"] \arrow[d] & \Gamma_j \arrow[d] \\
\Gamma'_0 \arrow[r, "h_1"] & \Gamma_1 \arrow[r, "h_2"] & \dots \arrow[r, "h_{k-j}"] & \Gamma'_{k-j} \arrow[r, "h_{k-j+1}"] & \dots \arrow[r, "h_{n-1}"] & \Gamma'_{n-1} \arrow[r, "h_n"] & \Gamma'_n
\end{tikzcd}$$

or

(2)
$$\begin{tikzcd}
\Gamma_0 \arrow[r, "g_1"] \arrow[d] & \Gamma_1 \arrow[r, "g_2"] \arrow[d] & \dots \arrow[r, "g_{k-j+1}"] & \Gamma_{k-j+1} \arrow[r, "g_{k-j+2}"] \arrow[d] & \dots \arrow[r, "g'_k"] & \Gamma_{k'} \arrow[r, "g_{k+1}"] \arrow[d] & \Gamma_k \arrow[d] \\
\Gamma'_j \arrow[r, "h_{j+1}"] & \Gamma'_{j+1} \arrow[r, "h_{j+2}"] & \dots \arrow[r, "h_n"] & \Gamma'_n = \Gamma'_0 \arrow[r, "h_1"] & \dots \arrow[r, "h_{j-1}"] & \Gamma'_{j-1} \arrow[r, "h_j"] & \Gamma'_j
\end{tikzcd}$$

or

(3)
$$\begin{tikzcd}
\Gamma_j \arrow[r, "g_{j+1}"] \arrow[d] & \Gamma_{j+1} \arrow[r, "g_{j+2}"] \arrow[d] & \dots \arrow[r, "g_k"] & \Gamma_k = \Gamma_0 \arrow[r, "g_1"] \arrow[d] & \dots \arrow[r, "g_{j-1}"] & \Gamma_{j-2} \arrow[r, "g_{j-1}"] \arrow[d] & \Gamma_{j-1} \arrow[r, "g_j"] \arrow[d] & \Gamma_j \arrow[d] \\
\Gamma'_0 \arrow[r, "h_1"] & \Gamma_1 \arrow[r, "h_2"] & \dots \arrow[r, "h_{k-j}"] & \Gamma'_{k-j} \arrow[r, "h_{k-j+1}"] & \dots \arrow[r, "h_{n-1}"] & \Gamma'_{n-1} \arrow[r, "h'_n"] & \Gamma'_{n'} \arrow[r, "h_{n+1}"] & \Gamma'_n
\end{tikzcd}$$

Fold-conjugate decompositions are *partial-fold conjugate* if we are in the case of (2) or (3) with $j = 0$.

Fold-conjugate tt maps represent $\mathrm{Out}(F_r)$-conjugate outer automorphisms, hence share all conjugacy class invariants, such as ideal Whitehead graphs, and whether or not a map is fully irreducible, and then also principal.

We can also consider an equivalence relation generated by (a)-(c) on a more general class of fold sequences. This class of fold sequences would include pathological examples with many "unnecessary" subdivisions of one or more folds, as well as allowing us to start in the middle of any fold in the sequence. The benefit of such a relation is that it would allow one to define and understand an $\mathrm{Out}(F_r)$-conjugacy class invariant that is the minimal number of folds in an equivalence class with respect to the partial-fold conjugacy relation.

3. The Rank-3 Principal Stratum Automaton $\widehat{\mathcal{A}_3}$

The "Lonely Direction Automaton" of [GP23] includes as directed loops only the Stallings fold decompositions of rotationless principal fully irreducible $\varphi \in \mathrm{Out}(F_3)$. To include directed loops for all principal fully irreducible outer automorphisms, instead of just their rotationless powers, we construct a new directed graph we call the "Rank-3 Principal Stratum Automaton" $\widehat{\mathcal{A}_3}$. We construct $\widehat{\mathcal{A}_3}$ by adding in "permutation arrows" where ltt structure permutation relabelings exist in the following sense:

Definition 3.1 (Graph/permutation relabeling)**.** Suppose that Γ is a directed labeled graph and σ is a permutation of $E^{\pm}\Gamma$ so that $\sigma(\bar{e}) = \overline{\sigma(e)}$ for each $e \in E^{\pm}\Gamma$. By $\sigma \cdot \Gamma$ we mean the directed labeled graph obtained from Γ by replacing each edge-label e with $\sigma(e)$. By a *graph relabeling (by σ)* we mean the graph isomorphism $g_\sigma \colon \Gamma \to \sigma \cdot \Gamma$ defined by that for each $e \in E\Gamma$ the image of e is the directed edge $\sigma(e) \in E(\sigma \cdot \Gamma)$. Suppose that G is an ltt structure with underlying graph Γ, we let $\sigma \cdot G$ denote the ltt structure on $\sigma \cdot \Gamma$ obtained from G by relabeling the black edges and vertices via σ. We then consider $\sigma \cdot G$ to be a *permutation relabeling* of G by σ.

The following lemma provides a concrete description of the $\mathcal{A}_r$, and $\widehat{\mathcal{A}_r}$, nodes.

Lemma 3.2 (Lonely Direction Lemma). *Suppose $\varphi \in \mathrm{Out}(F_r)$ is a principal fully irreducible outer automorphism, then each tt representative of φ is partial-fold conjugate to a tt map $g\colon \Gamma \to \Gamma$ satisfying:*

(a) each $SW(v, g)$ is a triangle and there are $2r - 3$ such triangles, and
(b) all vertices of Γ have valence 3 except a single vertex which has valence 4, and
(c) all but one direction in Γ is g-periodic and this direction is at the valence-4 vertex of Γ, and
(d) the nonperiodic direction is contained in precisely 1 turn taken by g.

Proof. Suppose $h\colon \Gamma' \to \Gamma'$ is a tt representative of φ and let R denote the rotationless power of φ. Then $h^R\colon \Gamma' \to \Gamma'$ is a tt representative of φ^R, and each h-periodic vertex and direction is fixed by h^R. By the proof of [MP16, Corollary 3.8], h has at most 1 nonperiodic vertex and that vertex has precisely 2 gates. As in the proof of [MP16, Corollary 3.8], by performing a fold at that vertex, one obtains a tt representative g' of φ^R for which all vertices are fixed and have at least 3 fixed directions. Since φ and φ^R are lone axis fully irreducibles, this fold is within their shared axis $\mathcal{A}$.

Because $\mathcal{A}$ is both h and h^R-periodic, with both periods starting and ending at (re-marked copies of) Γ', the [MP16, Corollary 3.8] fold, let us call it f, is a fold of Γ' occurring at the start of the Stallings fold decompositions of both h and h^R. If f is a partial fold, since it occurs within the axis, there is some full fold F in $\mathcal{A}$ and partial fold f' in $\mathcal{A}$ so that $F = f' \circ f$. For similar reasons, f' would appear at the end of the Stallings fold decompositions. Hence, in fact, letting g denote the conjugation of h by the fold, $g' = g^R$.

We now show g is a tt representative of φ. Since $g' = g^R$ is a tt map, no power of g can cause backtracking on an element of $E\Gamma$. Since quotienting by a partial fold cannot change the induced map of fundamental groups, we are left to show g maps vertices to vertices. But each vertex has 3 g'-gates, so cannot be mapped within its Stallings fold decomposition to a nonvertex point.

Since $\varphi \in \mathrm{Out}(F_r)$ is principal $\mathrm{IW}(\varphi)$, hence also $\mathrm{IW}(\varphi^k)$ for each $k \in \mathbb{Z}_{>0}$, is the disjoint union of $2r-3$ triangles. By [MP16, Lemma 4.5] and since $\mathrm{IW}(\varphi) = \mathrm{IW}(\varphi^k)$ has no cut vertices for each $k \in \mathbb{Z}_{>0}$, no tt representative τ of any φ^k has a PNP. Thus, since $IW(\varphi^k)$ is the disjoint union of the $SW(\tau, v)$ having at least 3 vertices, for any such τ and vertex v with at least 3 periodic directions, $SW(f, v)$ is a triangle. Since each vertex of g^R has at least 3 fixed directions, each vertex of g has at least 3 periodic directions. Hence, each $SW(g, v)$ is a triangle and there are $2r - 3$ such triangles, proving (a).

We now prove (b)-(c). Since each $SW(g, v)$ is a triangle, there are precisely 3 periodic directions at each vertex $v \in V\Gamma$. By [MP16, Lemma 3.6] g^R, hence g, has precisely one illegal turn. Thus, Γ has precisely one g-nonperiodic direction. And we have that Γ has a unique vertex of valence 4, and all other vertices have valence 3.

(d) follows from [GP23, Lemma 2], after observing that each periodic direction must be contained in 2 turns (by the structure of the stable Whitehead graphs) so that the turn described in [GP23, Lemma 2] must contain the unique nonperiodic direction. □

We now define the rank-3 principal stratum automaton $\widehat{\mathcal{A}_3}$:

Definition 3.3 (Rank-3 Principal Stratum Automaton $\widehat{\mathcal{A}_3}$). The *Rank-3 Principal Stratum Automaton* $\widehat{\mathcal{A}_3}$ is obtained from the "Lonely Direction Automaton" of [GP23] by adding a bi-directed edge labeled with σ for each ltt structure permutation re-labeling $G \to \sigma \cdot G$, which implicitly records the directed edge σ labeling the edge $[G, \sigma \cdot G]$ and σ^{-1} labeling the edge $[\sigma \cdot G, G]$. For each $r \geq 3$, the *Rank-r Principal Stratum Automaton* $\widehat{\mathcal{A}_r}$ is defined analogously.

Figure 3 depicts $\widehat{\mathcal{A}_3}$ with the permutation edges depicted in green. While we refer the reader to [GP23] for the details of its construction, note that it is the maximal strongly connected components (in rank 3 there is only one such component containing a loop, up to permutation relabeling) of a graph where

- vertices are ltt structures of the appropriate rank satisfying the "Lonely Direction Property" dictated by Lemma 3.2 and
- directed edges are either folds compatible with the ltt structures at the endpoints, or permutation relabelings as discussed in our definition.

Following a directed loop thus gives a sequence of folds and permutations which could be a decomposition of a principal fully irreducible outer automorphism, though one can easily produce directed loops that do not give such elements.

Consider a directed loop in the automaton with the edge labels reading as follows (the f_k are folds and the σ_k are graph-relabeling permutations):

$$\Gamma_0' \xrightarrow{g_{\sigma_0}} \Gamma_0 \xrightarrow{f_1} \Gamma_1' \xrightarrow{g_{\sigma_1}} \Gamma_1 \xrightarrow{f_2} \Gamma_2' \xrightarrow{g_{\sigma_2}} \Gamma_2 \xrightarrow{f_3} \cdots \xrightarrow{g_{\sigma_{n-1}}} \Gamma_{n-1} \xrightarrow{f_n} \Gamma_n' \xrightarrow{g_{\sigma_n}} \Gamma_n = \Gamma. \tag{4}$$

Using the notation

$$\Gamma_k^{(j)} := (\sigma_k \circ \cdots \circ \sigma_{k-j})^{-1} \cdot \Gamma_k \quad \text{and} \tag{5}$$

$$f_k^{(j)} := g^{-1}_{\sigma_{k-j-1}\circ\cdots\circ\sigma_0} \circ f_k \circ g_{\sigma_{k-j-1}\circ\cdots\circ\sigma_0} \colon \Gamma_{k-1}^{(j)} \to \Gamma_k^{(j+1)}, \tag{6}$$

we have the following commutative diagram indicating how the map of (4) can be rewritten as folds and then possibly a single ltt structure permutation relabeling at the end.

$$\begin{tikzcd}
 & & & & & \Gamma_n = \Gamma_0 \\
 & & & & \Gamma_{n-1} \arrow[r, "f_n"] & \Gamma_n' \arrow[u, "g_{\sigma_n}"'] \\
 & & \Gamma_2 \arrow[rr, dashed] & & \Gamma_{n-1}^{(n-3)} \arrow[r, "f_n^{(n-3)}"] \arrow[u, dashed] & \Gamma_n^{(n-2)} \arrow[u, dashed] \\
 & \Gamma_1 \arrow[r, "f_2"] & \Gamma_2' \arrow[u, "g_{\sigma_2}"'] \arrow[rr, dashed] & & \Gamma_{n-1}^{(n-2)} \arrow[r, "f_n^{(n-2)}"] \arrow[u, "g_{\sigma_2}"'] & \Gamma_n^{(n-1)} \arrow[u, "g_{\sigma_2}"'] \\
\Gamma_0 \arrow[r, "f_1"] & \Gamma_1' \arrow[u, "g_{\sigma_1}"'] \arrow[r, "f_2'"] & \Gamma_2'' \arrow[u, "g_{\sigma_1}"'] \arrow[rr, dashed] & & \Gamma_{n-1}^{(n-1)} \arrow[r, "f_n^{(n-1)}"] \arrow[u, "g_{\sigma_1}"'] & \Gamma_n^{(n)} \arrow[u, "g_{\sigma_1}"'] \\
\Gamma_0' \arrow[u, "g_{\sigma_0}"] \arrow[r, "f_1'"] & \Gamma_1'' \arrow[u, "g_{\sigma_0}"'] \arrow[r, "f_2''"] & \Gamma_2''' \arrow[u, "g_{\sigma_0}"'] \arrow[rr, dashed] & & \Gamma_{n-1}^{(n)} \arrow[r, "f_n^{(n)}"] \arrow[u, "g_{\sigma_0}"'] & \Gamma_n^{(n+1)} \arrow[u, "g_{\sigma_0}"']
\end{tikzcd} \tag{7}$$

We strengthen here [GP23, Proposition 6]:

Theorem C. *Suppose g is a train track representative of a principal fully irreducible $\varphi \in \mathrm{Out}(F_3)$. Then the Stallings fold decomposition of g is partial-fold conjugate to one determining a directed loop in $\widehat{\mathcal{A}_3}$.*

Proof. Suppose g is a tt representative of a principal fully irreducible $\varphi \in \mathrm{Out}(F_3)$. By [GP23, Proposition 6], some power g^p of g is fold-conjugate (really partial-fold conjugate) to a map given by a directed loop in $\mathcal{A}_3$, and we explain how to fine-tune the proof of [GP23, Proposition 6] for this level of precision that we need. Since tt representatives of the same principal (hence lone axis) outer automorphism have Stallings fold decompositions yielding the same axis, they are in fact fold-conjugate through folds of the axis. Hence, if the folds for one of these tt representatives is represented by a loop in $\mathcal{A}_3$, then so is the other, provided they both start and end on the kinds of ltt structures represented in the automata, i.e those whose underlying graph satisfies the "Lonely Direction property" of [GP23]. This is addressed in Lemma 3.2.

Since φ is principal, hence lone axis, g has a unique Stallings fold decomposition

$$\Gamma = \Gamma_0 \xrightarrow{g_1} \Gamma_1 \xrightarrow{g_2} \Gamma_2 \xrightarrow{g_3} \cdots \xrightarrow{g_{t-2}} \Gamma_{t-1} \xrightarrow{g_{t-1}} \Gamma_t \xrightarrow{g_\sigma} \Gamma_0 = \Gamma. \tag{8}$$

If necessary, we can increase p so that σ^p is the identity. Since g^p can be decomposed as

$$(\Gamma_0 \xrightarrow{g_1} \Gamma_1 \xrightarrow{g_2} \Gamma_2 \xrightarrow{g_3} \cdots \xrightarrow{g_{t-2}} \Gamma_{t-1} \xrightarrow{g_{t-1}} \Gamma_t \xrightarrow{g_\sigma} \Gamma_0)^p, \tag{9}$$

in light of (7), we obtain a Stallings fold decomposition of g^p as

$$\begin{aligned} \Gamma = \Gamma_0 \xrightarrow{g_1} \Gamma_1 \xrightarrow{g_2} \cdots \xrightarrow{g_{t-2}} \Gamma_{t-1} \xrightarrow{g_{t-1}} \Gamma_t = \Gamma'_0 \xrightarrow{g'_1} \Gamma'_1 \xrightarrow{g'_2} \cdots \xrightarrow{g'_{t-2}} \Gamma'_{t-1} \xrightarrow{g'_{t-1}} \Gamma'_t = \Gamma''_0 \cdots \\ \cdots \xrightarrow{g^{(p-2)}_{t-1}} \Gamma^{(p-2)}_t = \Gamma^{(p-1)}_0 \xrightarrow{g^{(p-1)}_1} \Gamma^{(p-1)}_1 \xrightarrow{g^{(p-1)}_2} \cdots \xrightarrow{g^{(p-1)}_{t-2}} \Gamma^{(p-1)}_{t-1} \xrightarrow{g^{(p-1)}_{t-1}} \Gamma^{(p-1)}_t \xrightarrow{g^p_\sigma = Id} \Gamma_0 = \Gamma. \end{aligned} \tag{10}$$

Since g^p must also be principal, hence lone axis, this is the unique Stallings fold decomposition for g^p, so must be partial-fold conjugate to a sequence represented by a directed loop in $\mathcal{A}_3$.

If Γ_0 has a valence-4 vertex then, by the proof of [GP23, Proposition 6] and the first paragraph of this proof, g^p itself is represented by a directed loop in $\mathcal{A}_3$. Now (8) is, but for the addition of g_σ, a subpath of the directed loop (10) in $\mathcal{A}_3$. And the Rank-3 Principal Stratum Automaton $\widehat{\mathcal{A}_3}$ is obtained from $\mathcal{A}_3$ by adding in edges for all possible ltt structure permutation relabelings, closing up the path via g_σ to a loop. So g will be represented by a directed loop in $\widehat{\mathcal{A}_r}$.

Now suppose that Γ_0 is fully trivalent. As in (7),

$$\Gamma_0 \xrightarrow{g_1} \Gamma_1 \xrightarrow{g_2} \Gamma_2 \xrightarrow{g_3} \cdots \xrightarrow{g_{t-2}} \Gamma_{t-1} \xrightarrow{g_\sigma} \sigma \cdot \Gamma_{t-1} \xrightarrow{g_\sigma \circ g_{t-1} \circ g_\sigma^{-1}} \sigma \cdot \Gamma_t = \Gamma_0$$

composes to be the same map as (8). We can see Γ_1 has a valence-4 vertex as follows. Complete the fold of the illegal turn at Γ_0. If it fully identifies two edges (upper right-hand image), the terminal vertex would have valence-5 and neither g nor its fold-conjugates could have represented a principal $\varphi \in \text{Out}(F_3)$, a contradiction. So g_1 must be a proper full fold (lower right-hand image) and Γ_1 now has a valence-4 vertex at the image of w. So the previous paragraph indicates that

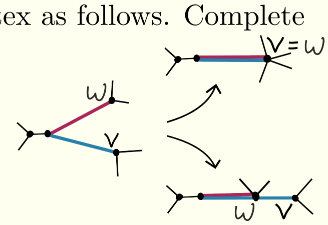

$$\Gamma_1 \xrightarrow{g_2} \Gamma_2 \xrightarrow{g_3} \cdots \xrightarrow{g_{t-2}} \Gamma_{t-1} \xrightarrow{g_\sigma} \sigma \cdot \Gamma_{t-1} \xrightarrow{g_\sigma \circ g_{t-1} \circ g_\sigma^{-1}} \sigma \cdot \Gamma_t = \Gamma_0 \xrightarrow{g_1} \Gamma_1$$

is represented by a directed loop in $\widehat{\mathcal{A}_3}$. So in this case also the Stallings fold decomposition of g is partial-fold conjugate to one determining a directed loop in $\widehat{\mathcal{A}_3}$. □

Proposition 3.4. *Every principal fully irreducible $\varphi \in \text{Out}(F_3)$ contains in its Stallings fold decomposition the left-hand graph. Thus, every principal axis must pass through the simplex with underlying graph both the second and third graphs in the image.*

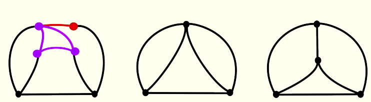

Proof. By Theorem C, the Stallings fold decomposition of any representative g of a principal fully irreducible $\varphi \in \mathrm{Out}(F_3)$ is partial-fold conjugate to one determining a directed loop in $\widehat{\mathcal{A}_3}$.

The image to the right schematically depicts $\widehat{\mathcal{A}_3}$, with the folds changing the labeling on the ltt structure included as dashed. The numbers on the arrows indicate how many distinct maps connect the isomorphism classes of ltt structures. It is an interesting observation that, as drawn to the right and in Figure 3, $\widehat{\mathcal{A}_3}$ has a clockwise orientation.

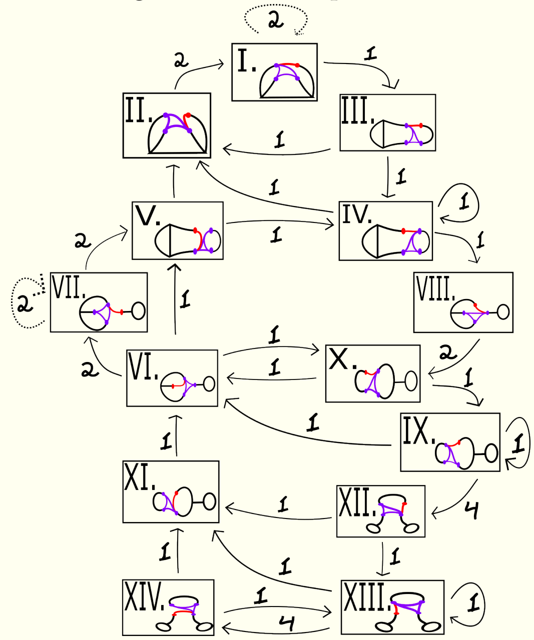

One can then observe that, once Node I is removed the maximal strongly connected component, which we call here $\mathcal{M}_I$, includes none of Node I, Node II, or Node III. Further, in the edge-labeling scheme of Figure 3, no permutation within $\mathcal{M}_I$ contains c and no fold within $\mathcal{M}_I$ maps c over another edge.

Hence, no directed loop within $\mathcal{M}_I$ can define an irreducible map. And we thus have that any irreducible map represented by a directed loop in the automata would need to contain Node I, which is the left-hand image in the statement of the proposition. The underlying graph of this ltt structure is the middle image. Inspection of all 4 folds entering Node I (including the self-maps), indicates that the axis must also pass through the third graph in the image. $\square$

4. The single-fold map

Consider the graph map $\mathfrak{g}$ defined by:

$$\mathfrak{g} = \begin{cases} a \mapsto \bar{b} \\ b \mapsto \bar{d} \\ c \mapsto e \\ d \mapsto \bar{e}\,\bar{c} \\ e \mapsto a \end{cases} \tag{11}$$

In $\widehat{\mathcal{A}_3}$, one can find $\mathfrak{g}$ represented by the directed loop:

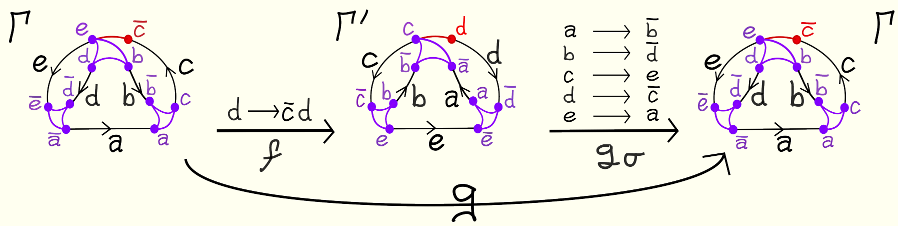

Figure 2. The Stallings fold decomposition of $\mathfrak{g}$ is a fold f, then homeomorphism g_σ.

Lemma 4.1. *The map $\mathfrak{g}$ of Figure 2 represents a principal fully irreducible $\psi \in \mathrm{Out}(F_3)$.*

Proof. We first use the criterion of Proposition 2.1 to prove that $\mathfrak{g}$ represents an ageometric fully irreducible outer automorphism ψ.

The direction map $D\mathfrak{g}$ is given by:

$$\begin{tikzcd} \overline{c} \arrow[r] & \overline{e} \arrow[r] & \overline{a} \arrow[r] & b \arrow[r] & \overline{d} \arrow[r] & c \arrow[r] & e \arrow[r] & a \arrow[r] & \overline{b} \arrow[r] & d \arrow[llllllll, bend right] \end{tikzcd}$$

So one can check as in [GP23, Lemma 5] that the turns taken by $\{\mathfrak{g}^p(e) \mid p \in \mathbb{Z}_{>0}, e \in E\Gamma\}$ are:

$$\{e,\overline{c}\},\ \{a,\overline{e}\},\ \{\overline{b},\overline{a}\},\ \{d,b\},\ \{\overline{e},\overline{d}\},\ \{\overline{a},c\},\ \{b,e\},\ \{\overline{d},a\},\ \{c,\overline{b}\},\ \{e,d\}. \tag{12}$$

Since the only illegal turn $\{d,\overline{c}\}$ is not taken, $\mathfrak{g}$ is a tt map. Since an adequately high power of the transition map is positive, $\mathfrak{g}$ is Perron–Frobenius. Alternatively, one can note that $\mathfrak{g}^{13}(d)$ contains all edges of Γ and d is in $\mathfrak{g}^2(a)$, $\mathfrak{g}(b)$, $\mathfrak{g}^4(c)$, and $\mathfrak{g}^3(e)$. So that all edges of Γ are in $\mathfrak{g}^{15}(a)$, $\mathfrak{g}^{16}(b)$, $\mathfrak{g}^{17}(c)$, and $\mathfrak{g}^{16}(e)$. So that $\mathfrak{g}^{17}$ maps every edge over every edge.

In light of (12), the local Whitehead graphs are triangles at each valence-3 vertex, and then the union of the Figure 2 colored edges at the valence-4 vertex. Hence, each local Whitehead graph is connected. Using the sage package of Coulbois [Cou14], for example, one can check that there are indeed no PNPs.

By Proposition 2.1, $\mathfrak{g}$ thus represents an ageometric fully irreducible outer automorphism ψ. Further, since all directions but $\overline{c}$ are periodic, the ideal Whitehead graph would be a union of 3 triangles. Hence, φ is in fact principal. $\square$

5. Perron numbers & Minimal stretch factors

A *weak Perron number* is an algebraic integer λ that is at least $|\alpha|$ for each conjugate α of λ. A weak Perron number where the inequality is strict is called a *Perron number*. Perron numbers are precisely the spectral radii of nonnegative aperiodic integral matrices.

We will need:

Lemma 5.1. *For each $r \geq 2$, the stretch factor of each fully irreducible $\varphi \in \mathrm{Out}(F_r)$ is a Perron number.*

Proof. Suppose that $\varphi \in \mathrm{Out}(F_r)$ is a fully irreducible outer automorphism. Its stretch factor is the PF eigenvalue $\lambda(g)$ of the transition matrix $M(g)$ for any expanding irreducible train track map g. By definition, $M(g)$ is a nonnegative integer matrix. By [Kap14, Lemma 2.4(2)], $M(g)$ is primitive. Hence, $M(g)$ is aperiodic. Since Perron numbers are precisely the spectral radii of nonnegative aperiodic integral matrices, $\lambda(g)$ is a Perron number, as desired. $\square$

Theorem A. *The stretch factor of ψ is minimal among fully irreducible elements in* $\mathrm{Out}(F_3)$.

Proof. The transition matrix for $\mathfrak{g}$ is:

$$M(\mathfrak{g}) = \begin{bmatrix} 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \end{bmatrix}$$

Thus the stretch factor of $\mathfrak{g}$ is the largest real root of the characteristic polynomial of $M(\mathfrak{g})$, which is $q(x) = x^5 - x - 1$.

Graphs in CV_3 have between 3 and 6 edges. Suppose $\varphi \in \mathrm{Out}(F_3)$ has a tt representative on a graph with 6 edges. By an Euler characteristic argument, each graph with 6 edges is trivalent. If φ is fully irreducible, then its Stallings fold decomposition must consist of at least one fold. As in the proof of Theorem C (and its corresponding image), the graph at the completion of the fold will either have a valence-4 vertex (if it is a proper full fold) or a valence-5 vertex (if it completely identifies 2 edges). In either case an Euler characteristic argument indicates the number of edges would have decreased from 6. Now, fold-conjugate outer automorphisms are in the same conjugacy class, so have the same stretch factor. Thus, if, a fully irreducible stretch factor is achieved by a tt map on a 6-edge graph, it is also achieved by a train track map on a graph with fewer edges. In fact, if a given stretch factor is achieved by a principal fully irreducible outer automorphism with a tt map on a 6-edge graph, then the same can be said on some graph of fewer edges.

In light of the previous paragraph, since the stretch factor of any element of $\mathrm{Out}(F_3)$ is the largest eigenvalue of the transition matrix of a tt representative, the algebraic degree of the stretch factor is between 1 and 5. Moreover, the stretch factor must be a Perron number by Lemma 5.1. The smallest few Perron numbers of degrees 2, 3, 4, and 5 are known. Approximately, these are 1.618, 1.325, 1.221, and 1.124 respectively. ([Boy85, Table 3]) Since $\mathfrak{g}$ has stretch factor less than 1.32 and 1.22, the smallest stretch factor must be a degree-5 Perron number.

The two smallest degree-5 Perron numbers are approximately $\gamma \approx 1.124$, and $\lambda \approx 1.167$, the largest real root of $x^5 + x^4 - x^2 - x - 1$ and $x^5 - x - 1$ respectively ([Boy85, Table 3]). Since $\mathfrak{g}$ has stretch factor exactly λ, we need only rule out γ as a possible stretch factor. Notice that the trace of $x^5 + x^4 - x^2 - x - 1$ is -1. Thus any 5 dimensional matrix with this characteristic polynomial must have negative trace. Since transition matrices have non-negative integer entries, this is impossible. Thus no element of $\mathrm{Out}(F_3)$ has stretch factor γ. So $\mathfrak{g}$ has the minimal stretch factor among fully irreducible elements of $\mathrm{Out}(F_3)$, as desired. ☐

6. Uniqueness of the single-fold principal train track map

The goal of this section is to prove Theorem B. Assume throughout that Γ is a graph of rank $r \geq 3$. Suppose that $h = g_\sigma \circ f$ is a tt representative of a principal fully irreducible $\varphi \in \mathrm{Out}(F_r)$, such that $f\colon \Gamma \to \Gamma'$ is a single proper full fold defined by $f\colon e_1 \mapsto e_0'e_1'$ and $g_\sigma\colon \Gamma' \to \Gamma$ is a graph relabeling isomorphism. Notice that one could not have $h = g_\sigma \circ f$ if f were instead a partial fold, because partial folds change the graph-isomorphism type of a graph.

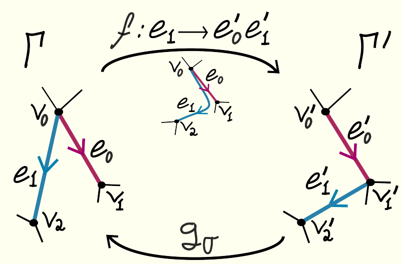

By Lemma 3.2, Γ will have a single valence-4 vertex and all other vertices will have valence 3. Further, the fold must occur at the unique valence-4 vertex in Γ. Let v_1 denote the terminal vertex of e_0 and v_2 the terminal vertex of e_1, so that $e_0 = [v_0, v_1]$ and $e_1 = [v_0, v_2]$. We label the vertices of Γ' so that $f(v) = v'$ and $e' = [v', w']$ for each $e = [v, w] \in E\Gamma - \{e_1\}$ and $e_1' = [v_1', v_2']$.

The following lemma will help us further set notation for the graphs Γ and Γ', as well as aid in the proof of Theorem B.

Lemma 6.1. *Suppose Γ is a graph with rank r and $h = g_\sigma \circ f$ is an irreducible train track representative of a principal fully irreducible $\varphi \in \mathrm{Out}(F_r)$ such that $f\colon \Gamma \to \Gamma'$ is a single proper full fold defined by $f\colon e_1 \mapsto e_0'e_1'$ and $g_\sigma\colon \Gamma' \to \Gamma$ is a graph relabeling isomorphism, as described above. Then*

(a) $|V\Gamma| = 2r - 3$ *and* $|E\Gamma| = 3r - 4$.
(b) the vertices v_0, v_1, *and* v_2 *are distinct*
(c) $g_\sigma(v_1') = v_0$, $g_\sigma(e_1') = e_0$, *and* $g_\sigma(v_2') = v_1$
(d) h *is transitive on the vertex set of* Γ.

Proof. We prove (a) - (d) in order, one at a time.

(a) Since the rank of Γ is r, the Euler characteristic satisfies $\chi(\Gamma) = r - 1$. So

$$r - 1 = |V\Gamma| - |E\Gamma|.$$

Since Γ has a single valence-4 vertex and all other vertices are valence-3, we have:

$$V\Gamma = 2r - 3 \qquad \text{and} \qquad E\Gamma = 3r - 4.$$

(b) If $v_0 = v_1 = v_2$, both e_0 and e_1 would be single-edge loops. Since valence$(v_0) = 4$ and part (a) implies Γ must have more than one vertex, Γ would be disconnected. Thus the three vertices are not all equal.

If $v_0 = v_1$, then e_0 is a single-edge loop, as is its image in Γ'. Since $f(v_0) = v_0'$ is the unique valence-4 vertex in Γ', we must have $h(v_0) = g_\sigma(v_0') = v_0$. However, if v_0 is fixed by h, the set of edges incident to v_0 is h-invariant. Since by part (a) there are necessarily more edges in Γ not incident to v_0, this contradicts the irreducibility of h.

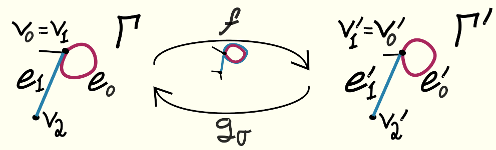

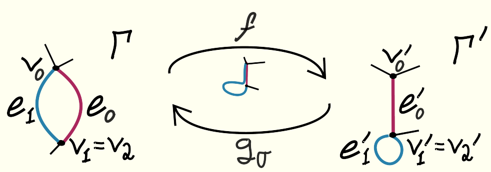

If $v_1 = v_2$, then e_1' is a single-edge loop in Γ'. Now Γ' has one more edge which is a single-edge loop than Γ does, contradicting that g_σ is a graph isomorphism.

If $v_0 = v_2$, then e_1 is a single-edge loop. However e_1' is not a single-edge loop in Γ', so Γ has one more edge which is a single-edge loop than Γ', again contradicting that g_σ is a graph isomorphism.

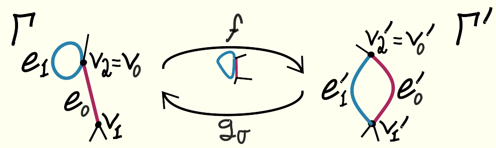

(c) In light of (b) we have the following picture (and will show the vertex colors reflect their orbit):

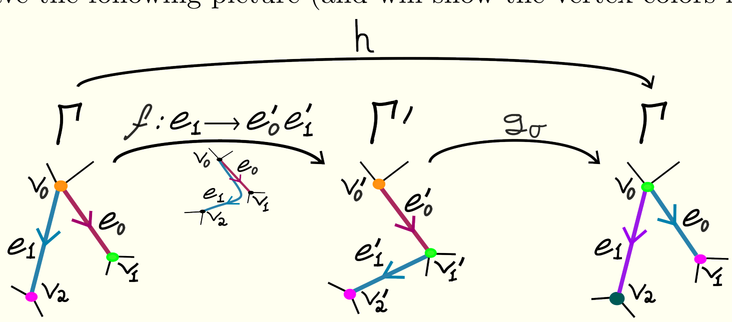

Since v_0 is the unique valence-4 vertex of Γ, v_1' is the unique valence-4 vertex of Γ', and g_σ is a graph isomorphism, we must have $g_\sigma(v_1') = v_0$.

By Lemma 3.2, all but one direction of Γ must be periodic under h. Hence

$$|Dh^k(\mathcal{D}\Gamma)| = |\mathcal{D}\Gamma| - 1 \tag{13}$$

for all powers $k \geq 1$. Observe that e_1' is the single direction in Γ' not in the image of Df.

We claim $Dg_\sigma(e_1') \in \{e_0, e_1\}$. Suppose not. Since g_σ is a graph isomorphism, Dg_σ is bijective on the set of directions of Γ'. Hence there exists a pair of distinct directions $a', b' \in \mathcal{D}\Gamma'$ such that $Dg_\sigma(a') = e_1$ and $Dg_\sigma(b') = e_0$. Since $g_\sigma(v_1') = v_0$ and g_σ is a graph-isomorphism, a' and b'

emanate from v_1'. Since we assumed neither a' nor b' are equal to the Df-unachieved direction e_1', the definition of Γ' ensures there exist $a, b \in \mathcal{D}\Gamma$ so that $Df(a) = a'$ and $Df(b) = b'$. Further, since a' and b' are incident to v_1' and $f(v_1) = v_1'$, we have that a and b are incident to v_1. Thus,

$$Dh^2(a) = Dg_\sigma(Df(Dg_\sigma(Df(a)))) = Dg_\sigma(Df(Dg_\sigma(a'))) = Dg_\sigma(Df(e_1)) = Dg_\sigma(e_0')$$

and

$$Dh^2(b) = Dg_\sigma(Df(Dg_\sigma(Df(b)))) = Dg_\sigma(Df(Dg_\sigma(b'))) = Dg_\sigma(Df(e_0)) = Dg_\sigma(e_0').$$

Hence $Dh^2(a) = Dh^2(b)$, which implies $|Dh^2(\mathcal{D}\Gamma)| \leq |\mathcal{D}\Gamma| - 2$, contradicting 13. Thus, $Dg_\sigma(e_1') \in \{e_0, e_1\}$, as desired.

Since g_σ is a graph isomorphism, the previous paragraph implies that $g_\sigma(e_1') \in \{e_1, e_0\}$. Suppose for the sake of contradiction that $g_\sigma(e_1') = e_1$. Then $g_\sigma(v_2') = v_2$, and so h fixes the vertex v_2. Since $Dh(\overline{e_1}) = Dg_\sigma(Df(e_0')) = Dg_\sigma(\overline{e_1}') = \overline{e_1}'$, the set consisting of the remaining two edges incident to v_2 is invariant under h. This contradicts the irreducibility of h. Therefore $g_\sigma(e_1') = e_0$.

Since $e_0 = [v_0, v_1]$ and $e_1' = [v_1', v_2']$ and g_σ is a graph isomorphism, and we already have $g_\sigma(v_1') = v_0$, we know that $g_\sigma(v_2')$ must be the vertex of e_0 other than v_0, which is exactly v_1. This completes the proof of part (c).

(d) Suppose h is not transitive on the vertex set of Γ. Then $V\Gamma = X \sqcup Y$ for nonempty proper subsets X and Y, both h-invariant. Without loss of generality, suppose $v_2 \in X$. Let E_Y be the edges in Γ with at least one incident vertex in Y. We aim to show E_Y is a nonempty proper subset of $E\Gamma$ which is invariant under $h = g_\sigma \circ f$, contradicting that h must be an *irreducible* tt map.

By part (c), $h(v_2) = g_\sigma(v_2') = v_1$ and $h(v_1) = g_\sigma(v_1') = v_0$. Hence both $v_1, v_0 \in X$ by invariance of X. As Y is disjoint from X, all three vertices $v_0, v_1, v_2 \notin Y$, implying $e_0, e_1 \notin E_Y$. This guarantees that E_Y is indeed a proper subset of the edge set. Clearly E_Y is nonempty, since Y is nonempty and every vertex in Γ has nonzero valence. All that remains to show is that E_Y is h-invariant.

Let $e \in E_Y$. By definition of E_Y, there is a vertex $v \in Y$ incident to e. Since $e_1 \notin E_Y$, we know $e \neq e_1$, so by definition of Γ', we have e' is incident to v' in Γ'. By definition of f, we have

$$h(e) = g_\sigma(f(e)) = g_\sigma(e')$$

and

$$h(v) = g_\sigma(f(v)) = g_\sigma(v').$$

Since g_σ is a graph isomorphism, $g_\sigma(e')$ remains incident to $g_\sigma(v')$. Moreover, $g_\sigma(v') = h(v) \in Y$ by invariance of Y. Hence $h(e) = g_\sigma(e') \in E_Y$, so E_Y is invariant under h, completing the proof of part (d). $\square$

Theorem B. *Up to edge relabeling, the map $\mathfrak{g}$ of Figure 2 is the only train track map representing a principal fully irreducible outer automorphism whose Stallings fold decomposition consists of only a single fold composed with a graph-relabeling isomorphism.*

Proof. Suppose h is a tt representative of a principal fully irreducible $\varphi \in Out(F_3)$. Then, by Proposition 3.4, h is represented by a loop in $\widehat{\mathcal{A}_3}$ including the ltt structure (let us call it G) at Node I. If h contains only a single fold, then this fold must be one of the folds taking G to itself. Up to permutation-relabeling of ltt structures, there are only 2 such folds and the other combination of a fold then graph-relabeling isomorphism leaves invariant the subgraph labeled with b and c in the graph-labeling class of $\widehat{\mathcal{A}_3}$ included in Figure 3. Thus, up to edge relabeling, h is unique.

We now show that for each $r \geq 4$ no such map can exist. Let $n = 2r - 3$ be the number of vertices in Γ. By Lemma 6.1(d), the vertices of Γ are all in the same orbit under h. Hence we can recursively label the vertices of Γ as follows, taking subscripts modulo n:

(1) Let v_0 remain as is and
(2) let $v_{k-1} := h(v_k)$ for $1 \le k \le n$.

On the vertex set, h is now given by:

$$v_0 \mapsto v_{n-1} \mapsto v_{n-2} \mapsto \cdots \mapsto v_2 \mapsto v_1 \mapsto v_0.$$

We maintain the convention that a vertex v in Γ is given the label v' in Γ', so we have a labeling of vertices in Γ', as well. Since $f(v) = v'$ for any $v \in V\Gamma$, the above labeling implies $g_\sigma(v'_k) = v_{k-1}$ and $g_\sigma^{-1}(v_k) = v'_{k+1}$. Observe that by Lemma 6.1(c), our labeling is consistent with our original labeling of v_1 and v_2.

We will reach a contradiction to the existence of h by showing v_2 would have the wrong valence. In what follows, we rely on the fact that the edge-labelings induced by f (in the sense of §2.6) provides a bijection between $E\Gamma \backslash \{e_1\}$ and $E\Gamma' \backslash \{[v'_1, v'_2]\}$ given by mapping $e = [v, w]$ to $e' = [v, w]$. Thus, if e'_j and e'_k have distinct vertices and are not $e'_1 = [v'_1, v'_2]$ or $\overline{e'_1} = [v'_2, v'_1]$, then removing primes on vertices gives distinct edges of Γ.

Let $e'_2 := g_\sigma^{-1}(e_1)$. Since $e_1 = [v_0, v_2]$ and g_σ^{-1} is a graph isomorphism,

$$e'_2 = g_\sigma^{-1}(e_1) = [g_\sigma^{-1}(v_0), g_\sigma^{-1}(v_2)] = [v'_1, v'_3].$$

Since $e'_1 = [v'_1, v'_2]$ and $v'_2 \neq v'_3$, this tells us $e'_2 \notin \{e'_1, \overline{e'_1}\}$. Therefore, by our convention for labeling the vertices of Γ', there is an edge in $E\Gamma$ joining v_1 and v_3. We call this edge e_2. Similarly, let $e'_3 := g_\sigma^{-1}(e_2)$. Since $e_2 = [v_1, v_3]$ and g_σ^{-1} is a graph isomorphism,

$$e'_3 = g_\sigma^{-1}(e_2) = [g_\sigma^{-1}(v_1), g_\sigma^{-1}(v_3)] = [v'_2, v'_4].$$

By Lemma 6.1(a) and the assumption that $r \ge 4$, we have $|V\Gamma'| = |V\Gamma| \ge 5$. Hence $v'_1 \neq v'_4$, so $e'_3 \notin \{e'_1, \overline{e'_1}\}$. So there is an edge in $E\Gamma$ joining v_2 and v_4. We call this edge e_3.

Let $\alpha_1 := g_\sigma(e'_0)$. Since $e'_0 = [v'_0, v'_1]$ and g_σ is a graph isomorphism,

$$\alpha_1 = g_\sigma(e'_0) = [g_\sigma(v'_0), g_\sigma(v'_1)] = [v_{n-1}, v_0]$$

Since $n \ge 5$, we know $v_{n-1} \neq v_2$, and hence also $\alpha_1 \notin \{e_1, \overline{e_1}\}$. By the definition of Γ' and its labeling, there is an edge $\alpha'_1 = [v'_{n-1}, v'_0] \in E\Gamma'$. Similarly, let

$$\alpha_2 := g_\sigma(\alpha'_1) = [g_\sigma(v'_{n-1}), g_\sigma(v'_0)] = [v_{n-2}, v_{n-1}].$$

Again, by the definition of Γ', there is an edge $\alpha'_2 = [v'_{n-2}, v'_{n-1}] \in E\Gamma'$. Recursively define

$$\alpha_k := g_\sigma(\alpha'_{k-1}) = [g_\sigma(v'_{n-k+1}), g_\sigma(v'_{n-k+2})] = [v'_{n-k}, v'_{n-k+1}].$$

for $k \in \{2, \dots n-1\}$.

The final two recursively defined edges $\alpha_{n-2} = [v_2, v_3]$ and $\alpha_{n-1} = [v_1, v_2]$ both contain v_2. Moreover, $e_1 = [v_0, v_2]$ and $e_3 = [v_2, v_4]$ contain v_2. Since v_0, v_1, v_3, and v_4 are distinct vertices, all four of these edges are distinct. Therefore v_2 is contained in at least 4 distinct edges in Γ. However, this contradicts that v_2 should have valence 3. Thus for $r \ge 4$ can there exist a single-fold irreducible tt map h representing a principal fully irreducible element of $Out(F_r)$.

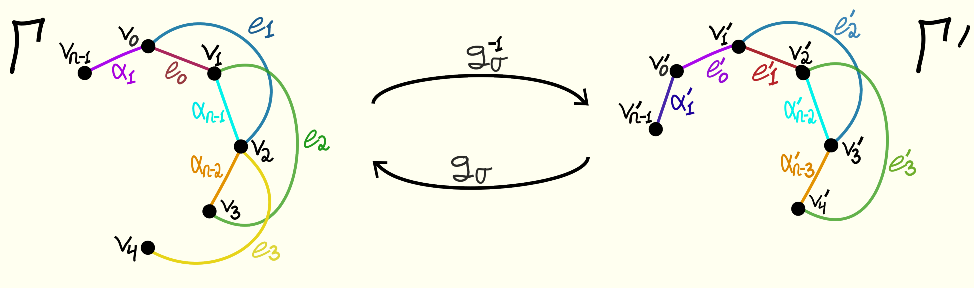

□

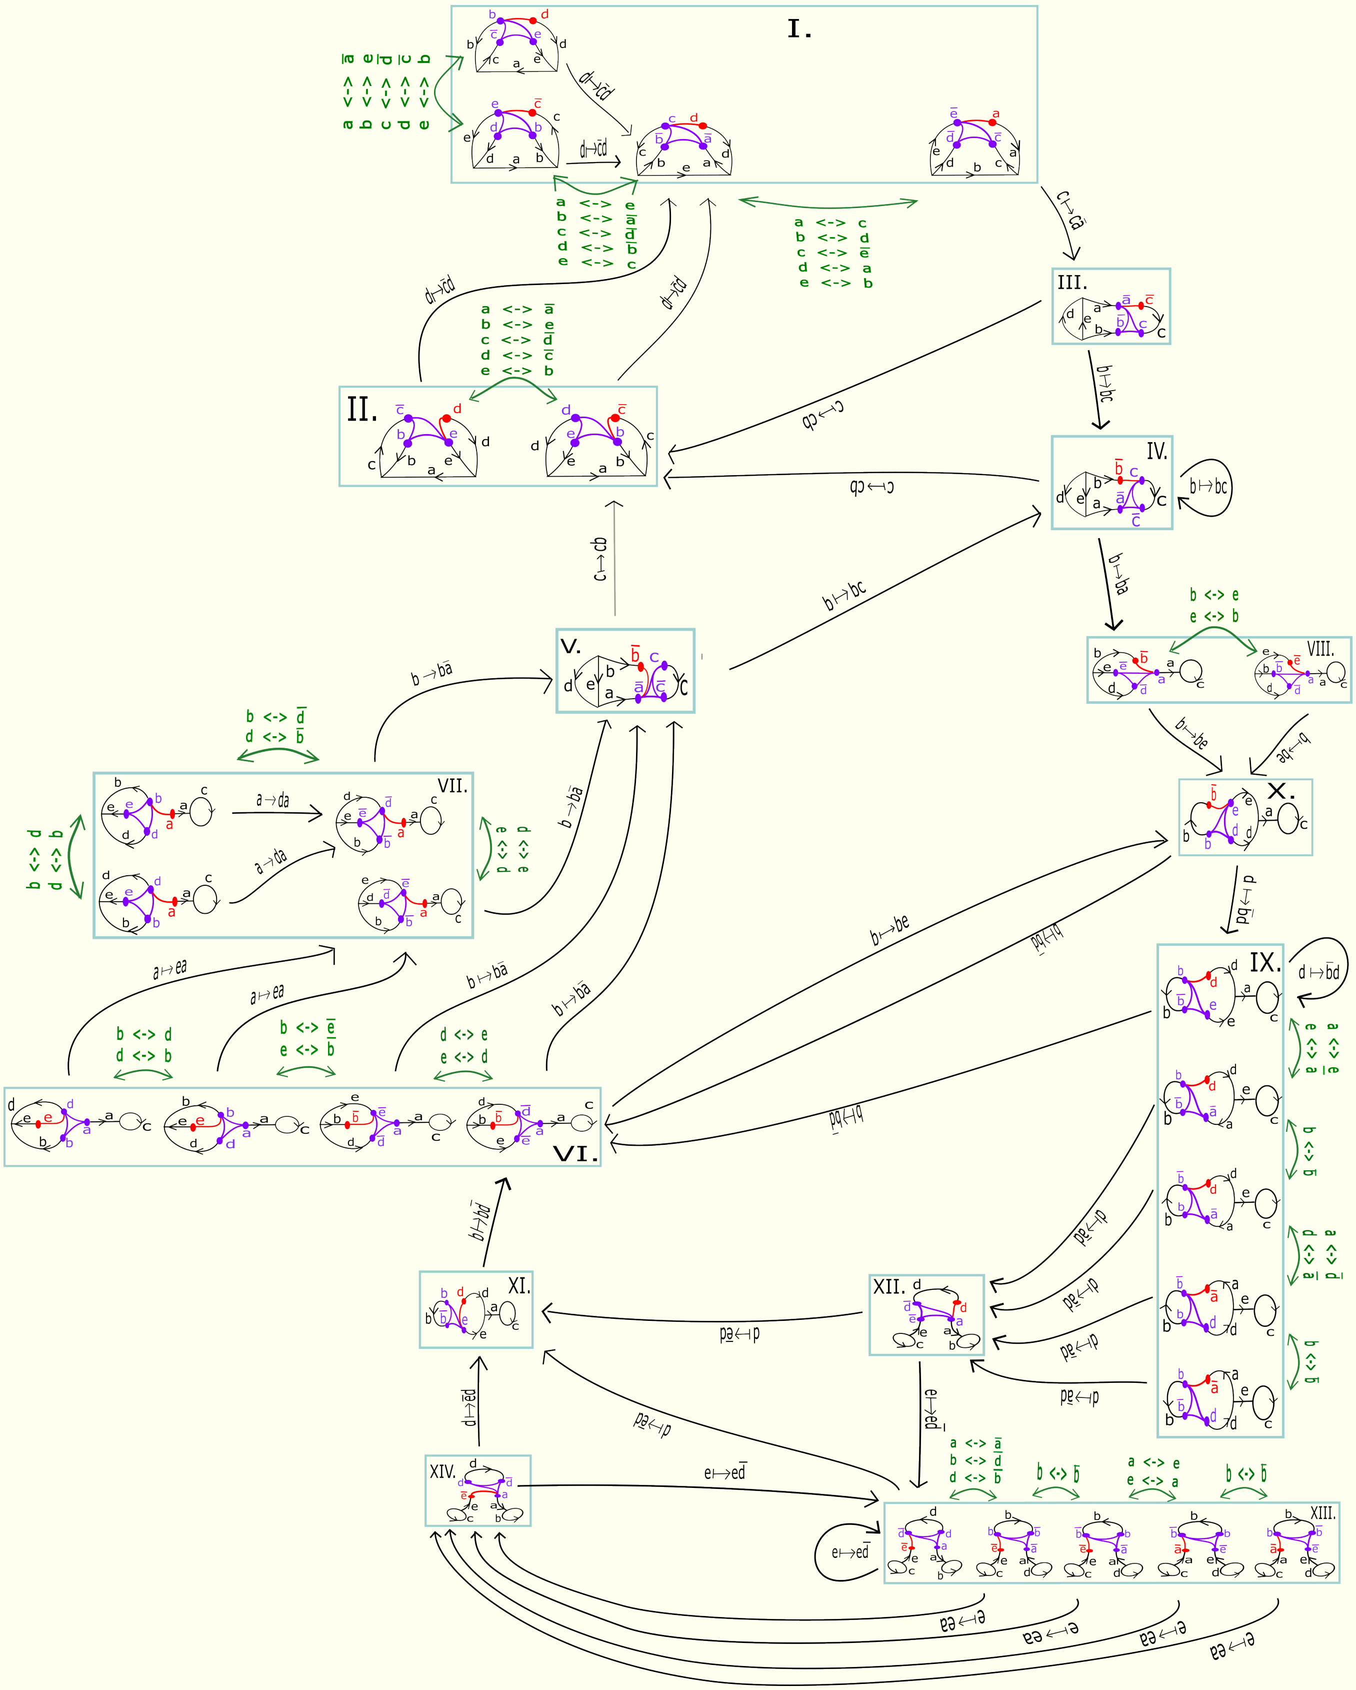

FIGURE 3. This figure depicts the Rank-3 Principal Stratum Automaton. Permutations of the automata are included in green. Compositions of included permutations are implicitly included.

References

[AKKP19] Y. Algom-Kfir, I. Kapovich, and C. Pfaff. Stable strata of geodesics in outer space. *International Mathematics Research Notices*, 2019(14):4549–4578, 2019.

[AKR15] Y. Algom-Kfir and K. Rafi. Mapping tori of small dilatation expanding train-track maps. *Topology and its Applications*, 180:44–63, 2015.

[Bes14] M. Bestvina. Geometry of outer space. In *Geometric group theory*, volume 21 of *IAS/Park City Math. Ser.*, pages 173–206. Amer. Math. Soc., Providence, RI, 2014.

[BFH00] M. Bestvina, M. Feighn, and M. Handel. The Tits Alternative for Out (F_n) I: Dynamics of exponentially-growing automorphisms. *Annals of Mathematics-Second Series*, 151(2):517–624, 2000.

[BH92] M. Bestvina and M. Handel. Train tracks and automorphisms of free groups. *The Annals of Mathematics*, 135(1):1–51, 1992.

[Boy85] D. W. Boyd. Supplement to "The maximal modulus of an algebraic integer". *Math. Comput.*, 45:s17–s20, 1985.

[Cou14] T. Coulbois. Free group automorphisms and train-track representative in python/sage. `https://github.com/coulbois/sage-train-track`, 2012–2014.

[CV86] M. Culler and K. Vogtmann. Moduli of graphs and automorphisms of free groups. *Inventiones mathematicae*, 84(1):91–119, 1986.

[DDH$^+$22] R. Dickmann, G. Domat, T. Hill, S. Kwak, C. Ospina, P. Patel, and R. Rechkin. Thurston's theorem: Entropy in dimension one. *arXiv preprint arXiv:2209.15102*, 2022.

[DKL15] S. Dowdall, I. Kapovich, and C. J. Leininger. Dynamics on free-by-cyclic groups. *Geom. Topol.*, 19(5):2801–2899, 2015.

[DKL17] S. Dowdall, I. Kapovich, and C. J. Leininger. McMullen polynomials and Lipschitz flows for free-by-cyclic groups. *Journal of the European Mathematical Society*, 19(11):3253–3353, 2017.

[FH11] M. Feighn and M. Handel. The recognition theorem for $\mathrm{Out}(F_n)$. *Groups Geom. Dyn.*, 5(1):39–106, 2011.

[FM11] S. Francaviglia and A. Martino. Metric properties of outer space. *Publicacions Matemàtiques*, 55(2):433–473, 2011.

[GM17] V. Gadre and J. Maher. The stratum of random mapping classes. *Ergodic Theory and Dynamical Systems*, pages 1–17, 2017.

[GP23] D. Gagnier and C. Pfaff. Taking the high-edge route of rank-3 outer space. *International Journal of Algebra and Computation*, 33(08):1659–1685, 2023.

[HM11] M. Handel and L. Mosher. *Axes in outer space*. Number 1004. Amer Mathematical Society, 2011.

[Kap14] I. Kapovich. Algorithmic detectability of iwip automorphisms. *Bulletin of the London Mathematical Society*, 46(2):279–290, 2014.

[KMPT22a] I. Kapovich, J. Maher, C. Pfaff, and S. J. Taylor. Random trees in the boundary of outer space. *Geometry & Topology*, 26(1):127–162, 2022.

[KMPT22b] I. Kapovich, J. Maher, C. Pfaff, and S.J. Taylor. Random outer automorphisms of free groups: Attracting trees and their singularity structures. *Transactions of the American Mathematical Society*, 375(01):525–557, 2022.

[Mas82] H. Masur. Interval exchange transformations and measured foliations. *Ann. of Math*, 115(1):169–200, 1982.

[MP16] L. Mosher and C. Pfaff. Lone axes in outer space. *Algebr. Geom. Topol.*, 16(6):3385–3418, 2016.

[Pfa12] C. Pfaff. *Constructing and Classifying Fully Irreducible Outer Automorphisms of Free Groups*. PhD thesis, Rutgers University, 2012. PhD Thesis,`http://www.math.rutgers.edu/~cpfaff/Thesis.pdf`.

[Pfa13] C. Pfaff. Ideal Whitehead graphs in $\mathrm{Out}(F_r)$ II: the complete graph in each rank. *Journal of Homotopy and Related Structures*, 10(2):275–301, 2013.

[Sko89] R. Skora. Deformations of length functions in groups, preprint. *Columbia University*, 1989.

[Sta83] J.R. Stallings. Topology of finite graphs. *Inventiones Mathematicae*, 71(3):551–565, 1983.

[Thu14] W. Thurston. Entropy in dimension one. *arXiv preprint arXiv:1402.2008*, 2014.

[Vog15] K. Vogtmann. On the geometry of outer space. *Bull. Amer. Math. Soc. (N.S.)*, 52(1):27–46, 2015.

Mathematical Institute, University of Oxford
https://naomigandrew.wordpress.com/,
Email address: Naomi.Andrew@maths.ox.ac.uk

University of California - Santa Barbara Department of Mathematics
https://sites.google.com/view/paigehillen/home,
Email address: paigehillen@ucsb.edu

Department of Mathematics, Rutgers University - Newark
https://ryleealanza.org/,
Email address: rylee.lyman@rutgers.edu

Department of Mathematics & Statistics, Queen's University
https://mast.queensu.ca/~cpfaff/,
Email address: c.pfaff@queensu.ca