Faithful semigroup diagram representations of homeomorphism groups

Liam Kevin Stott

A thesis submitted for the degree of PhD
at the
University of St Andrews

2024

Identifier to use to cite or link to this thesis:

DOI: https://doi.org/10.17630/sta/1166

Candidate's declaration

I, Liam Kevin Stott, do hereby certify that this thesis, submitted for the degree of PhD, which is approximately 28,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for any degree. I confirm that any appendices included in my thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

I was admitted as a research student at the University of St Andrews in September 2019.

I received funding from an organisation or institution and have acknowledged the funder(s) in the full text of my thesis.

Date 11/11/2024 Signature of candidate

Supervisor's declaration

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of PhD in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree. I confirm that any appendices included in the thesis contain only material permitted by the 'Assessment of Postgraduate Research Students' policy.

Date 11/11/2024 Signature of supervisor

Permission for publication

In submitting this thesis to the University of St Andrews we understand that we are giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. We also understand, unless exempt by an award of an embargo as requested below, that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or research worker, that this thesis will be electronically accessible for personal or research use and that the library has the right to migrate this thesis into new electronic forms as required to ensure continued access to the thesis.

I, Liam Kevin Stott, confirm that my thesis does not contain any third-party material that requires copyright clearance.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis:

Printed copy

No embargo on print copy.

Electronic copy

No embargo on electronic copy.

Date 11/11/2024 Signature of candidate

Date 11/11/2024 Signature of supervisor

Underpinning Research Data or Digital Outputs

Candidate's declaration

I, Liam Kevin Stott, hereby certify that no requirements to deposit original research data or digital outputs apply to this thesis and that, where appropriate, secondary data used have been referenced in the full text of my thesis.

Date 11/11/2024 Signature of candidate

Abstract

In this thesis we investigate two different classes of objects. The first is the class of Dénes cycles of a tree; we show a 'strong' correspondence between such cycles and certain partial orders on the edges of the tree. This allows us to find a slick new proof of a classical result, as well as give an algorithm for computing the multiplicity of a Dénes cycle. The second class of objects we investigate, and the main focus of this thesis, is groups generated by geometrically fast sets of bumps. We show that the class of such groups on the interval coincides with a certain subclass of diagram groups, and then exploit this connection to answer a question of Matthew Brin. We then go on to consider groups generated by fast bumps in a general setting and, in so doing, find a necessary and sufficient condition for when groups generated by fast sets of bumps on the circle are isomorphic to certain annular diagram groups. To finish, we deduce a couple of properties of fast groups of the interval using their diagram group representations. In aid of this, we define a group structure on infinite tree diagrams.

Dedication

I dedicate this thesis to my late mother, Rosie, without whom none of this would have been possible.

Acknowledgements

The author has been partially funded by the EPSRC through the University of St Andrews doctoral training program during the creation of this thesis.

The author would like to thank James Belk, Matthew Brin, Justin Moore, James Hyde and Peter Cameron for helpful discussions during the creation of this thesis. In particular, the author would like to thank his supervisor, Collin Bleak, for his invaluable guidance and support throughout his time on the program.

Contents

Chapter 1

Introduction

Around sixty years ago, Richard Thompson introduced three infinite groups emerging from his work on logic. These groups, now commonly known as Thompson's groups $F \subseteq T \subseteq V$, have since been the subject of extensive study within pure mathematics. This widespread interest can, perhaps, be explained by two factors. Firstly, they exhibit a range of interesting (and, at times, seemingly competing) properties; for example, the groups T and V were the first examples of finitely presented infinite simple groups, while the group F was one of the first potential counterexamples to the Von Neumann conjecture (and, although the Von Neumann conjecture was disproven over forty years ago, it remains an open question whether F indeed forms a counterexample). Secondly, they appear as fundamental objects in a plethora of areas of mathematics, ranging from group theory to homotopy theory, to the study of dynamical systems and, of course, logic. As such, it can be said that the study of these groups, and their many other relatives that have since been introduced, is a driving force for advancements in geometric group theory, which lies at the intersection of many of the aforementioned areas of mathematics. It is within that final area which this thesis finds itself, and to which it hopes to contribute. More explicitly, we will concern ourselves with Thompson's groups constructed as certain groups of homeomorphisms of Cantor sets, which are among the most popular representations of these groups. In particular, we construct F as a group of order preserving homeomorphisms of the unit interval.

More recently, around thirty years ago, the concept of diagram groups was first explicitly introduced in the literature by Victor Guba and Mark Sapir in their seminal paper [29] (although some particulars of the resulting theory could be found implicitly in work of Adian, Pride, and Squier). Diagram groups and their elements, known as diagrams, described in detail in Section 2.3, are defined over semigroup presentations. Among the first fundamental results detailed in their paper, Guba and Sapir show that the diagram group over the semigroup presentation $\langle x \mid x = x^2 \rangle$ is isomorphic to Thompson's group F. Minor variations on the definition of a (planar) diagram group, reflecting the variation between F, T and V themselves, find that T and V can be obtained as (non-planar) diagram groups over this same semigroup presentation. This suggests a deep connection between diagram groups and Thompson's groups. Later, in their paper [31], Guba and Sapir show that any (planar) diagram group has a homomorphic image in the group of all order preserving homeomorphisms of the unit interval and in many cases this image is isomorphic to the diagram group. This, in turn, suggests a deep connection between diagram groups and certain groups of homeomorphisms of Cantor sets.

Given one of these 'connection suggestions' the other is not surprising; indeed the connection between Thompson's groups and the aforementioned homeomorphism groups is already well-studied. One particularly interesting result is the now classic Ubiquity Theorem due to Matthew Brin [9] which shows that a group of piecewise linear order preserving homeomorphisms of the unit interval is guaranteed to contain a subgroup isomorphic to F under a remarkably weak condition on the action of its elements. In recent years work has been developed which produces results in much the same spirit as this. According to this work, given two order preserving homeomorphisms g, h each with connected support such that their supports overlap, there is a k such that $\langle g^i, h^j \rangle$ is isomorphic to F for all $i, j \geq k$; in particular, a subgroup of $\langle g, h \rangle$ is isomorphic to F. Indeed, this is simply a minor corollary of said work, published by Bleak, Brin, Kassabov, Moore and Zaremsky in their paper [7], which we give details of in Section 2.2. For now, suffice it to say that they provide a sufficient condition for two sets of order preserving homeomorphisms to generate isomorphic groups. Sets satisfying this condition are known as geometrically fast.

In this thesis, we begin, in Chapter 3, by investigating Dénes cycles of a tree T [15], which arise as the products of the edges of T when considered as transpositions of edges. This chapter comes out of joint work with Peter Cameron. We prove a correspondence between Dénes cycles and certain partial orders on the edges of T (Theorem 3.3.2) which allows one to recover a classical theorem (Theorem 3.4.1) as well as giving an algorithm for determining the number of products of edges which produce a given Dénes cycle.

In Chapter 4, we establish a connection between geometrically fast sets and diagram groups. More specifically, we define an isomorphism between a group generated by a geometrically fast set of homeomorphisms, each with connected support, and a particular diagram group; indeed, this diagram group is rich with structure reflecting that of the action of the fast group G. In particular, we may construct an infinite rooted tree which both represents a family of partitions of the interval arising from said action and completely describes how this action extends to these partitions; the diagram group arises from this extension of the action of G and we may use the infinite rooted tree to construct it explicitly. This construction efficiently highlights the features of the action of G identified as central to the proofs from [7] in a visual yet rigourous manner. In joint work with Jim Belk, we use this, along with some established theory on diagram groups, to show that two particular fast groups are isomorphic, thereby answering an open question of Brin [11].

The action on the infinite rooted tree described above is, in fact, a particular restriction of a group action determined by a local similarity system, a concept first introduced by Hughes [34]. Farley and Hughes later show that if a local similarity system is 'small' then the group action it determines is isomorphic to a symmetric diagram group [21] and this can be used to recover the fact that fast groups are isomorphic to diagram groups. Indeed, this makes it natural to increase the scope of investigation to include groups generated by certain geometrically fast sets acting on the circle, which in turn expands the collection of fast sets acting on the interval that this theory can be applied to. Specifically, in Chapter 5 we show that a group generated by a certain type of fast set is isomorphic to a particular annular diagram group if the set does not contain any closed stretched transition chains. Many such fast sets acting on the circle can be realised as fast sets acting on the interval; for example, it follows that the Brin-Navas group is an annular diagram group.

In Chapter 6, looking toward future analysis, we introduce a precise definition of infinite diagrams (extending that of Genevois [23]) such that we may define a group structure on them (in a restricted set of cases) and thereby construct a group in which the diagram group (in the usual sense) is contained as a normal subgroup. We hope to use this in future work to determine the isomorphism types of certain subgroups of fast groups.

Having established these basic facts about certain fast groups, in Chapter 7, we proceed to investigate some further properties. We produce a method for deducing information about the group presentation of a fast bump group using its diagram group representation. This method begins with the observation that by using a given reduced diagram we can algorithmically reverse engineer precisely which freely reduced words on the generators are equal to this diagram in the group. From there, we argue that a particular set of reduced diagrams is sufficient to generate all the relators of the group - the relations obtained from this set are all commutators.

Chapter 2

Preliminaries

2.1 Order theory

While this thesis is not, with the exception of Chapter 3, directly concerned with order theory, it is always lurking in the background and, as such, the definitions and results discussed in this subsection will be relevant, at least indirectly, in every subsequent part. Thus, we recall some of the basics of order theory here.

Let X be a set. Formally, we call a set $R \subseteq X \times X$ a (binary) **relation** on X. If S is a subset of R then we say S is a **subrelation** of R. If $Y \subseteq X$ then we call the subrelation defined as $R|_Y = R \cap (Y \times Y)$ the **induced subrelation** of R on Y. Given a relation R on X we define its **converse** to be the relation $R^C = \{(y, x) \mid (x, y) \in R\}$. If R_X, R_Y are relations on sets X, Y respectively then we say a function $f : X \to Y$ is **relation preserving** if $(x_1 f, x_2 f) \in R_Y$ for all $(x_1, x_2) \in R_X$ and we say it is **relation reflecting** if $(x_1, x_2) \in R_X$ for all $(x_1 f, x_2 f) \in R_Y$. If f is an injection satisfying both these conditions then f is a **relation embedding**. If f is a bijection satisfying both of these properties then we say f is a **relation isomorphism**, in which case we say the relations R_X, R_Y have the same **relation type**. A relation isomorphism $f : X \to X$ is called a **(relation) automorphism**. If there exists a relation automorphism f on X such that $Rf = R^C$ then we say X is **self-converse** and we call f a **converse automorphism** of X; if an automorphism is not converse then we say it is **non-converse**. If $f : X \to Y$ is a function then we call $R_X f = \{(x_1 f, x_2 f) \mid (x_1, x_2) \in R_X\}$ the **relation induced by** f on Y.

There is a suite of properties, each member of which a given relation may satisfy, and certain subsets of which define particular types of relations that we will consider over the next few definitions. Some such relations are typically referred to as orders. We recall the relevant details.

Let R be a relation on a set X.

- if $(x, x) \in R$ for all $x \in X$ then R is **reflexive**;
- if $(x, x) \notin R$ for all $x \in X$ then R is **irreflexive**;
- if for all $(x, y) \in R$ we have $(y, x) \in R$ then R is **symmetric**;
- if for all $(x, y) \in R$ where $x \neq y$ we have $(y, x) \notin R$ then R is **anti-symmetric**;
- if for all $(x, y), (y, z) \in R$ we have $(x, z) \in R$ then R is **transitive**;
- if for all $x, y \in X$ where $x \neq y$ we have either $(x, y) \in R$ or $(y, x) \in R$ then R is **total**.

For certain properties, we will find it useful to consider, given a relation Q, what is the smallest relation C which satisfies one such property. We refer to C as the closure of Q with respect to said property. More specifically, we may consider

- the **reflexive closure** Q^r, the smallest reflexive relation containing Q;
- the **symmetric closure** Q^s, the smallest symmetric relation containing Q; and
- the **transitive closure** Q^t, the smallest transitive relation containing Q.

This will be particularly useful when defining, say, naturally occurring partial orders in the chapters that follow. We now recall the definition of a partial order as well as various other types of order.

Let R be a relation on a set X,

- if R is reflexive and transitive then R is a **preorder**;
- if R is reflexive, anti-symmetric and transitive then R is a (non-strict) **partial order**;
- if R is irreflexive, anti-symmetric and transitive then R is a **strict partial order**
- if R is reflexive, symmetric and transitive then R is an **equivalence relation**;
- if R is reflexive, anti-symmetric, transitive and total then R is a **total order**;
- if R is irreflexive, anti-symmetric, transitive and total then R is a **strict total order**.

When discussing relations of one of these types we will often describe them as structuring its underlying set in some way. For example, if R is a preorder on X then we may say that X is **preordered** by R; similarly, if R is a partial order we may say X is **partially ordered** by R. Relatedly, in such cases we may refer to the relation R as an **order** on X and similarly for subsidiary terms (e.g. subrelation becomes suborder, relation isomorphism becomes order isomorphism, etc.).

One may notice that some of the definitions of different types of order contain superfluous properties (e.g. R being symmetric and transitive implies it is reflexive). This is done to make the connections between different types of relation more explicit. In particular, it is important to notice that R is a partial order if and only if it is a symmetric preorder; and, similarly, R is an equivalence relation if and only if it is an anti-symmetric preorder. In fact, any given preorder contains an equivalence relation as a subrelation whose complement is a strict partial order which induces a strict partial order on the corresponding quotient set.

Proposition 2.1.1. *Let X be a set preordered by a relation R. Define $E = \{(x, y) \in R \mid (x, y), (y, x) \in R\}$ and $P = R \setminus E$. Then E is an equivalence relation and P is a strict partial order such that the relation Pf induced by the canonical surjection $f : X \to X/E$ is a strict partial order.*

Proof. Since R is reflexive we have $(x, x) \in R$ for all $x \in X$ and so $(x, x) \in E$. If $(x, y), (y, z) \in E$ then $(x, y), (y, x), (y, z), (z, y) \in R$ which implies $(x, z), (z, x) \in R$ by transitivity. Thus E is an equivalence relation since it is symmetric by construction. On the other hand, it is clear that P is irreflexive and if $(x, y) \in P$ we have $(y, x) \notin R$ by definition and in particular $(y, x) \notin P$. If $(x, y), (y, z) \in P$ then $(x, z) \in R$ but if $(z, x) \in R$ then $(y, x) \in R$ which contradicts $(x, y) \in P$ and thus P is a strict partial order. By construction, f is order preserving and order reflecting when X is equipped with P and the quotient X/E is equipped with Pf. If $([x], [y]) \in Pf$ then we have $(t, u) \in P$ which implies $(u, t) \notin P$ for all $t \in [x], u \in [y]$ and so we conclude $([y], [x]) \in Pf$ (where $[x]$ is the equivalence class in X/E containing $x \in X$). If, in addition, we have $([y], [z]) \in Pf$ then $(u, v) \in P$ for all $u \in [y], v \in [z]$ and thus $(t, v) \in P$ for all $t \in [x], v \in [z]$ by transitivity of P. $\square$

Remark. We may refer to E as the canonical equivalence relation of R, and Pf as the canonical partial order of R.

Notice that if we have a partial order on X we may define a strict partial order on X and vice versa, simply by respectively removing or adding the diagonal set $\{(x, x) \mid x \in X\}$. We say these two orders **correspond**. With an order or an equivalence relation R on X we may sometimes use an arbitrary symbol, such as $\sim_R$, to refer to it, where by convention $x \sim_R y$ is understood to mean $(x, y) \in R$. In particular, it may be convenient to refer to a partial order P by the symbol $\leq$ and to refer to its corresponding strict partial order by the symbol $<$.

Topologies arise naturally from relations. Let R be a relation on X. For $x \in X$ define subsets

$$I_x = \{i \in X \mid (i, x) \in R,\, i \neq x\},\; D_x = \{d \in X \mid (x, d) \in R,\, d \neq x\}$$

and let T be the topology with subbasis $\{I_x, D_x \mid x \in X\}$. We may call T the **relation topology** of R on X. In particular, if R is a partial order with symbol $\leq$ then T is generated by the sets

$$I_x = \{i \in X \mid i < x\},\; D_x = \{d \in X \mid x < d\}$$

and we call T the **order topology**.

The class of relation topologies is very broad due to the lack of restriction on a relation in general. On the other hand, the class of order topologies is much more restricted - for many sets X, the order topology of any partial order P on X is the discrete topology. In general, however, some important topologies appear; in particular, the order topology of $\mathbb{R}$ with the usual total order coincides with the standard topology of $\mathbb{R}$.

2.1.1 Well orders, ordinal numbers and beyond

An important class of total orders is that of well orders. The definitions in this subsection will be most relevant for the objects defined in Chapter 6.

Let X be a set with strict total order R. We say that R is a **well order** if there does not exist an infinite sequence $x_0, x_1, \ldots \in X$ such that $(x_n, x_{n-1}) \in R$ for all $n \geq 1$. If $Y \subseteq X$ is such that there does not exist $x \in X \setminus Y$ satisfying $(x, y) \in R$ for any $y \in Y$ then Y is well ordered by $R|_Y$ and we say Y is a **prefix** of X, and we denote this relationship by $Y \leq X$ or $Y < X$ if, in addition, $Y \neq X$.

Ordinals are, in spirit, the order types of well ordered sets. In order to define them we opt for a canonical construction originally due to Von Neumann.

Let X be a set and define $R_\in = \{(x, y) \in X \times X \mid x \in y\}$. We say X is an **ordinal** or an **ordinal number** if $x \subseteq X$ for all $x \in X$ and $R_\in$ is a well order.

Remark. A given well ordered set S is order isomorphic to a unique ordinal α as defined above. As such, ordinal numbers give us a canonical representative for the order type of a well ordered set.

The idea here is that an ordinal α is the well ordered set of all ordinals strictly smaller than α; notice that an ordinal β is a proper prefix of α if and only if β is an element of α, which is to say $\beta \in \alpha$ if and only if $\beta < \alpha$. More explicitly, we start by defining $0 = \emptyset$ and inductively define

$$n + 1 = n \cup \{n\}$$

to obtain every finite ordinal. Then, taking the union of every finite ordinal we obtain

$$\omega = \bigcup_n n$$

and thereby define the first (i.e. smallest) countably infinite ordinal. We may then continue as before, defining $\omega+1 = \omega \cup \{\omega\}$ and then $\omega+n+1 = (\omega+n) \cup \{\omega+n\}$ to define a countably infinite collection of distinct countably infinite ordinals. Of course, we can then define

$$\omega + \omega = \bigcup_n (\omega + n)$$

and we may continue this process indefinitely. In general, an ordinal of the form $\alpha \cup \{\alpha\}$ for some ordinal α is called a **successor ordinal** and any non-empty ordinal not of this form is called a **limit ordinal**.

Arithmetic

It is possible to define several different forms of arithmetic on ordinal numbers. The 'standard' operations on ordinals, known simply as ordinal arithmetic, give us notions of addition, multiplication and exponentiation which, in general, behave rather differently to operations by the same name on finite cardinal numbers. These operations are, perhaps, the most natural for ordinals, since they are rooted in the treatment of ordinals as well ordered sets. We will have use of said addition on totally ordered sets more generally, and so we give a more general definition here.

Let S, T be totally ordered sets. Define the totally ordered set $S + T$ to be the disjoint union $S \sqcup T$ with relation equal to the transitive closure of $R_S \sqcup R_T \sqcup \{(s,t) \mid s \in S, t \in T\}$ where R_S, R_T are the orders on S, T respectively. If S, T are well ordered with order types α, β respectively then $S + T$ is well ordered and we define the sum $\alpha + \beta := \gamma$ where γ is the order type of $S + T$.

Ordinal addition is associative; however, in general, it is not commutative. For instance, the sum $\omega+n$ for a finite ordinal n defines the nth successor ordinal of ω as discussed above, but the sum $n+\omega$ is equal to ω since there is an order isomorphism between their corresponding well ordered sets.

Let S, T be well ordered sets with order types α, β respectively. Define a well order $<$ on $S \times T$ by

$$(s_1, t_1) < (s_2, t_2) \text{ if } s_1 <_S s_2 \text{ or } s_1 = s_2,\ t_1 <_T t_2$$

where $<_S, <_T$ are the well orders on S, T respectively. We define $\beta \cdot \alpha := \gamma$ where γ is the order type of the well ordered set $S \times T$.

Ordinal multiplication is associative but not commutative, since $2 \cdot \omega = \omega \neq \omega + \omega = \omega \cdot 2$.

Exponentiation is not so straightforward to define in this manner. We opt to define it by transfinite induction.

Let α, β be ordinal numbers. We define α^β by inducting on β:

- $\alpha^0 = 1$;
- $\alpha^{\beta+1} = (\alpha^\beta) \cdot \alpha$;
- $\alpha = \bigcup_{\delta \in \beta} \alpha^\delta$ if β is a limit ordinal.

Nordinal numbers and cordinal numbers

One may notice that ordinal numbers are, in a sense, defined to always be 'positive'. However, it will be fruitful to consider their 'negative' counterparts, as well as combinations of these, which we will now define formally.

Let X be a set totally ordered by a relation R, and let α be an ordinal. If the converse relation R^{-1} is a well order with order type α then we denote the order type of R by α^{-1} and we call α^{-1} a **nordinal**. If β is a prefix of α then we say β^{-1} is a **suffix** of α^{-1} and denote this by $\beta^{-1} \geq \alpha^{-1}$ or $\beta^{-1} > \alpha^{-1}$ if, in addition, $\alpha^{-1} \neq \beta^{-1}$.

Remark. Notice that $n^{-1} = n$ for a finite ordinal n. The relation $\beta^{-1} \geq \alpha^{-1}$ is equivalent to $\beta^{-1} \in \alpha^{-1}$.

Naturally, ordinal addition translates to the converse nordinal addition. For instance, $n + \omega^{-1} \neq \omega^{-1} = \omega^{-1} + n$. We can also use this operation to define a more general class of orders which extends the addition operation of both ordinals and nordinals.

Let α, β be ordinals; then α^{-1}, β^{-1} are nordinals and the sums $\alpha^{-1} + \beta, \alpha + \beta^{-1}, \beta^{-1} + \alpha, \beta + \alpha^{-1}$ are totally ordered sets. We define a **cordinal number** to be an order type expressible as a finite sum of order types of any of these forms. If γ is a cordinal number we denote its total order by R_γ.

The addition operation is associative on cordinal numbers. For instance, $(\omega^{-1} + n) + \omega = \omega^{-1} + \omega = \omega^{-1} + (n + \omega)$ for a finite (n)ordinal n. Notice also that cordinals are closed under addition, since the sum of two finite sums is a finite sum.

If $\gamma = \alpha_1^{\epsilon_1} + \ldots + \alpha_n^{\epsilon_n}$ for ordinals α_i and $\epsilon_i = \pm 1$ then we may collect adjacent terms together where $\epsilon_i = \epsilon_{i+1}$ to express γ as an sum alternating between ordinals and nordinals $\gamma = \beta_1^\epsilon + \beta_2^{-\epsilon} + \ldots + \beta_n^{\eta\epsilon}$ for $\epsilon, \eta = \pm 1$. Notice this alternating sum is not unique for a given cordinal since.

A slightly different way we may decompose a given cordinal γ is to identify intervals of the form $\omega^{-1}+\omega$ in γ and call them $\zeta_1, \ldots, \zeta_n$ in increasing order (notice this only occurs at a nordinal followed by an ordinal, and so there are finitely many such intervals). We can thus express $\gamma = \alpha_1+\eta_1+\zeta_1+\ldots+\alpha_n+\eta_n+\zeta_n+\alpha_{n+1}+\eta_{n+1}$ where α_i, η_i are ordinals and nordinals respectively, all maximally chosen, and this expression is unique.

Definition 2.1.2. Let γ be a cordinal number and let $\zeta_1, \ldots, \zeta_n$ be the ordered sequence of distinct intervals of γ with order type $\omega^{-1} + \omega$. We may write

$$\gamma = \alpha_1 + \eta_1 + \zeta_1 + \ldots + \alpha_n + \eta_n + \zeta_n + \alpha_{n+1} + \eta_{n+1}$$

where α_i, η_i are maximally chosen ordinals and nordinals respectively. We call this expression the **ζ-decomposition** of γ and refer to ζ_i as the ζ **summands** of γ.

This expression is particularly useful when considering automorphisms of γ since any non-converse automorphism of γ must restrict to a non-converse automorphism of each summand in its ζ-decomposition. If α is an ordinal or a nordinal then it has precisely one non-converse automorphism, thus a non-converse automorphism of γ is fully specified by choosing non-converse automorphisms for each ζ_i summand, each of which has order type $\omega^{-1} + \omega$. Further, a non-converse automorphism of $\omega^{-1} + \omega$ is fully specified by choosing the image of a single point. Thus, a non-converse automorphism of γ is fully specified by choosing the image of a suitable finite subset.

We will, in particular, be interested in ordered sets labelled by some disjoint set.

Let X and Σ be disjoint sets. We call a surjection $l : X \to \Sigma$ a **labelling of X by $\boldsymbol{\Sigma}$** and refer to the triple (X, Σ, l) as a set **labelled by $\boldsymbol{\Sigma}$**, though we will usually denote it simply by X once a labelling has been specified. If $(X, l_X), (Y, l_Y)$ are sets labelled by Σ then we say a function $f : X \to Y$ is **label-preserving** if $l_X = f \circ l_Y$.

If γ is a labelled by some set then its label-preserving non-converse automorphisms can similarly be specified.

Suppose T is a totally ordered set with order type $\omega^{-1} + \omega$ labelled by some set Σ. We say the labelling of T is **periodic** if there exists $w \in \Sigma^+$ such that there exists a label-preserving order isomorphism from T to the bi-infinite word $w^{\mathbb{Z}} := \Pi_{n\in\mathbb{Z}} w = \ldots w_{-1}w_0w_1 \ldots$ and we call a minimal such word a **period** of the labelling. Given a periodic labelling of T with period w we refer to a particular label-preserving order isomorphism $t : T \to w^{\mathbb{Z}}$ a **tiling** and refer to the preimages $t^{-1}(w_i)$ as **tiles** where w_i is the ith factor in the product $\Pi_{n\in\mathbb{Z}} w$ for $i \in \mathbb{Z}$. Otherwise, we say the labelling of T is **aperiodic**.

Remark. In general, a period w for a periodic labelling is not unique; each cyclic rotation of w is a period of the same labelling. Notice, then, that the length of a period $|w|$ is uniquely determined.

From this definition we can see that there is only one label-preserving non-converse automorphism of an aperiodically labelled set T with order type $\omega^{-1}+\omega$, while if the labelling of T is periodic with period $w = u_1 \dots u_n$ then any label-preserving non-converse automorphism is fully specified by choosing which tile any particular tile is mapped to. Indeed, the specification made by this choice is independent of the chosen period and tiling; if, given a period and tiling, we specify a map by saying "map the 0th tile to the nth tile" then this statement specifies the same map with respect to any other period-tiling pair. This is because, no matter the period and tiling in question, the result of this statement is to shift each element of T up by $n|w|$.

A central use of cordinals for our purposes is to index sequences. Let X be a set, let γ be a non-finite cordinal number. Given a bijection $f : \gamma \to X$ we define the **transfinite sequence** defined by f to be X ordered by the relation $R_\gamma f$ induced by f. We usually denote αf by x_α for $\alpha \in \gamma$ and denote the transfinite sequence by $(x_\alpha)_{\alpha\in\gamma}$, leaving the bijection f implicit. We typically say $(x_\alpha)_{\alpha\in\gamma}$ is a **sequence indexed by** γ. We may say a sequence indexed by γ has **length** γ.

Remark. Transfinite sequences indexed by ω are precisely those which are sequences under the usual definition; being indexed by the natural numbers $\mathbb{N}$. Similarly, transfinite sequences indexed by $\omega^{-1}+\omega$ are precisely those which are bi-infinite sequences; being indexed by $\mathbb{Z}$. Indeed, the class of transfinite sequences of a set X is contained in the class of nets of X.

Transfinite rooted trees

Finally, we will require a generalisation of a rooted tree to include transfinitely many levels. We borrow and adapt a definition from Ehrlich [17].

Definition 2.1.3. A partially ordered class $(A, <_T)$ is called a **tree** if for every $x \in A$ the induced suborder on its class of predecessors $\mathrm{pr}_A(x) = \{y \in A \mid y <_T x\}$ is a well order. If the members of A are labelled by some disjoint set then we say it is a labelled tree; for a given $x \in A$ we call the well ordered sequence of labels of $\mathrm{pr}_A(x)$ the **legacy** of x. We say a subclass B of A is an **initial subtree** if $\mathrm{pr}_B(x) = \mathrm{pr}_A(x)$ for all $x \in B$. For $x \in A$, the order type of the induced suborder on $\mathrm{pr}_A(x)$ is known as the **tree rank** of x, denoted $\rho_A(x)$, and we define **level** α of $(A, <)$ to be the set $L_A(\alpha) := \{x \in A \mid \rho_A(x) = \alpha\}$. The members of $L_A(0)$ are known as the **roots** and we say A has **height** α if $\alpha = \min\{\beta \mid L_A(\beta) = \emptyset\}$. If $x \in A$ such that there does not exist $y \in A$ where $x < y$ then we call x an **end** of A and we denote the set of ends of A by $\mathrm{end}(A)$. We say y is an **immediate successor** of x if $x <_T y$ and

$\rho_A(y) = \rho_A(x) + 1$ and we say it is the **immediate successor of a chain** $(x_\alpha)_{\alpha<\beta}$ of $<_T$ if $x_\alpha <_T y$ for all $\alpha < \beta$ and $\rho_A(y) = \min\{\sigma \mid \rho_A(x_\alpha) < \sigma \text{ for } \alpha < \beta\}$. Given $x \in A$ or a chain $c := (x_\alpha)_{\alpha<\beta}$ we denote its set of immediate successors in A by $\text{is}_A(x)$ and $\text{is}_A(c)$ respectively. Notice that if a chain c has a maximal element $y \in A$ then $\text{is}(c) = \text{is}(y)$.

We say a tree $(A, <_T)$ is **labelled** by a set Σ if every element $x \in A$ is assigned some label $s \in \Sigma$ and we denote $l(x) := s$. If $<$ is a (strict) total order on A then we say $(A, <_T, <)$ is an **ordered tree**. If for each chain $c := (x_\alpha)_{\alpha<\beta}$ of $<_T$ the class $\text{is}_A(c)$ of its immediate successors is totally ordered by $<_c$, as well as the roots $L_A(0)$ being totally ordered by $<_r$, then we say $(A, <_T, <_r, (<_c)_{c \text{ is a chain}})$ is a **locally ordered tree**.

Given a locally ordered tree $(A, <_T, <_r, (<_c)_{c \text{ is a chain}})$ we may define the **lexicographic order** $<_L$ on A, thus making $(A, <_T, <_L)$ an ordered tree, as follows. Given $x, y \in A$, let p be the maximal common prefix of their predecessor sets $\text{pr}(x), \text{pr}(y)$. We say $x <_L y$ if either:

- $p = \text{pr}(x)$ (in other words, $x <_T y$), or;
- $p = \emptyset$ and $x_r <_r y_r$ where $x_r \in \text{pr}(x) \cap L_A(0)$ and $y_r \in \text{pr}(y) \cap L_A(0)$, or;
- $x_p <_p y_p$ where $<_p$ is the linear order on $\text{is}(p)$ and x_p, y_p are the immediate successors of the chain p inside the chains $\text{pr}(x), \text{pr}(y)$ respectively.

A nice property of $<_L$ is that the identity function $\text{id} : A \to A$ on A defines an order-preserving bijection $(A, <_T, <_r, (<_c)_{c \text{ is a chain}}) \to (A, <_L)$.

Notice that this definition of tree would include what some may wish to call a forest - that is, a tree in the above sense may be disconnected.

2.2 Geometrically fast groups

In [7], Bleak et al. define a new class of groups of homeomorphisms of the real line and other spaces. Here we recall the definition of fast groups given in [7] in detail, as well as some of its surrounding theory.

Consider the unit interval $I = ([0, 1], \leq, T)$ with the usual total order $\leq$ and equipped with the order topology T. Define

$$\text{Homeo}_+(I) = \{f : I \to I \mid f \text{ is a non-converse order isomorphism}\}$$

Remark. Notice that, since I is equipped with the order topology, the order isomorphisms of I are precisely the homeomorphisms of I.

Let $f \in \text{Homeo}_+(I)$. We define the **support** of f to be the set $\text{supt}(f) = \{x \in I \mid xf \neq x\}$ and refer to the maximal connected subspaces of $\text{supt}(f)$ as the **orbitals** of f. For a subset $A \subseteq \text{Homeo}_+(I)$ we define $\text{supt}(A) = \bigcup_{f\in A} \text{supt}(f)$ and we call the connected components of this set the orbitals of A. If f has precisely one orbital (that is, $\text{supt}(f)$ is connected) then we say f is a **bump**. If f is a bump such that $x < xf$ for all $x \in \text{supt}(f)$ then we say f is **positive**; symmetrically, we say f is **negative** if $x > xf$ for all $x \in \text{supt}(f)$. If f is a bump then $\text{supt}(f)$ is an open interval and thus has precisely two boundary points $a < b$ which we call the **left** and **right endpoints** of f respectively and we sometimes write $\text{supt}(f) = (a, b)$. If B is a set of bumps then we say $b \in B$ is **isolated** in B if $\text{supt}(b)$ does not contain any endpoints of elements of B.

Remark. If f is a bump then it is either positive or negative.

We are interested in subgroups of $\text{Homeo}_+(I)$ generated by sets of a particular form, which we now define. Given a positive bump b, we refer to a point $m \in \text{supt}(b) = (a, c)$ as a **marker** of b. Given a marker m of b, we define the intervals $\text{src}(b) = (a, m)$ and $\text{dest}(b) = [mb, c)$, respectively called the **source** and **destination** of b and collectively called the **feet** of b. If a bump b is negative then we define $\text{src}(b) := \text{dest}(b^{-1})$ and $\text{dest}(b) := \text{src}(b^{-1})$. Given a collection of bumps $(b_i)_i$ we call a corresponding collection of markers $(m_i)_i$ a **marking** of $(b_i)_i$. Finally, if there exists a marking for a collection $B = (b_i)_i$ of bumps such that its set of feet $F_B(m_i) = \bigcup_i \{\text{src}(b_i), \text{dest}(b_i)\}$ satisfy $J \cap K = \emptyset$ for all distinct $J, K \in F_B(m_i)$ then we say $(b_i)_i$ is **geometrically fast**, or simply **fast**.

Remark. These definitions can be extended to arbitrary $A \subseteq \text{Homeo}_+(I)$ by instead considering the set of bumps $\{f|_O \mid O \text{ is an orbital of } f\}$ for each $f \in A$. Notice that a geometrically fast set of bumps is naturally totally ordered by correspondence with the induced total order on its set of left endpoints.

As one might guess, we are interested in the groups which are generated by such sets. More specifically, in this thesis, we are interested in those which are finitely generated. So, let B be a finite fast set of positive bumps. We call $G = \langle B \rangle$ a **fast bump group**.

Remark. We assume that the bumps of B are positive since this does not affect the isomorphism type of G. We will often leave this assumption implicit.

In studying these groups we will find the following concepts useful.

Let B be a set of bumps and consider a totally ordered subset $\mathcal{T} = (b_i)_i \subseteq B$ where $\text{supt}(b_i) = (a_i, c_i)$. If $a_i < a_{i+1} < c_i < c_{i+1}$ for all i then we say $\mathcal{T}$ is a **transition chain**. If, in addition, each interval (a_{i+1}, c_i) does not contain any endpoints of B then we say $\mathcal{T}$ is **stretched**. If B is fast (with a fixed marking) we denote by $S(\mathcal{T})$ and $D(\mathcal{T})$ the sets of sources and destinations of the transition chain

$\mathcal{T}$ respectively. These sets may inherit the total order from $\mathcal{T}$, which we denote by $\overrightarrow{S(\mathcal{T})}$ and $\overrightarrow{D(\mathcal{T})}$, or its converse, denoted $\overleftarrow{S(\mathcal{T})}$ and $\overleftarrow{D(\mathcal{T})}$. We denote the minimal element of $\overrightarrow{S(\mathcal{T})}$ by $s(\mathcal{T})$ and the maximal element of $\overrightarrow{D(\mathcal{T})}$ by $d(\mathcal{T})$, and we may say the transition chain $\mathcal{T}$ **starts at** $s(\mathcal{T})$ and **ends at** $d(\mathcal{T})$.

Remark. As a convenient convention, in particular for the purposes of Lemma 4.2.5, we consider the empty transition chain $\mathcal{T} = \emptyset$ to start at every source of B and end at every destination of B

Remark. The maximal stretched transition chains of a fast set of bumps B partition B.

For the purposes of studying fast groups more broadly, the authors in [7] introduce a graph constructed with respect to such a generating set. We define them as is sufficient for our purposes. An example is shown in Figure 2.1.

Definition 2.2.1. Let B be a finite fast set of positive bumps with marking M. Construct a graph on the totally ordered set $F_B(M)$ by defining an oriented edge from $\mathrm{src}(b)$ to $\mathrm{dest}(b)$ for each $b \in B$. We denote this graph by D_B (or D) and refer to it as the **dynamical diagram** of B. We consider two dynamical diagrams D_1, D_2 for sets B_1, B_2 with markings M_1, M_2 to be **isomorphic** if there is a non-converse order isomorphism $F_{B_1}(M_1) \to F_{B_2}(M_2)$ which induces a graph isomorphism $D_1 \to D_2$.

Remark. Notice, then, that the isomorphism type of D_B is independent of the choice of marking M for B.

We are now in a position to state (a restriction of) the main theorem from [7], which lies at the heart of this thesis as a whole.

Theorem 2.2.2 (Fast group theorem [7])**.** *Let B_1, B_2 be fast sets of bumps with dynamical diagrams D_1, D_2. Then $\langle B_1 \rangle \cong \langle B_2 \rangle$ if D_1, D_2 are isomorphic.*

Remark. It will be convenient - and, in light of this theorem, coherent - to sometimes refer to a fast bump group $\langle B \rangle$ as being generated by its dynamical diagram D_B.

The converse of this theorem is not true - as we will see in Chapter 4, it is rather easy to find non-isomorphic dynamical diagrams which generate isomorphic groups.

It will be convenient to give the following definition, making use of a general marking for a set B as defined in [7, Proposition 4.3] known as the **canonical marking** for B. This marking is efficient in the sense that the markers of a stretched transition chain are determined by the marker of the first bump in said chain, where the marker of the first bump f in B is the first left end point in $\mathrm{supt}(f)$ and the remaining markers are determined inductively. The canonical marking M_c has the property that if a set B is fast then it will be witnessed by M_c.

Definition 2.2.3. Let B be a fast set of bumps with the canonical marking M_c. If $\mathrm{supt}(B) \setminus F_B(M_c)$ consists only of isolated points then we say B is **canonical**.

Remark. By Theorem 2.2.2 we may always assume that B is canonical when considering an isomorphism type of fast bump group. In general, we will assume that a fast set of bumps B has the canonical marking M_c and, from here, denote $F_B := F_B(M_c)$.

Central to the proof of Theorem 2.2.2 in [7], as well as our own proofs, is the idea of *local reductions* of a word on a fast set B.

Definition 2.2.4. Given a geometrically fast set of bumps B we say that a word $w \in (B^{\pm})^*$ is **simply locally reduced at a point** $x \in I$ if for any prefix of the form ub for $b \in B^{\pm}$ we have $xu \neq xua$. Given a word $w \in (B^{\pm})^*$ we define the **simply local reduction** w_x **of** w **at** $x \in I$ to be the word obtained by deleting b from w whenever $xu = xub$ for some prefix ub of w. We denote the set of simply local reductions of a word by $L(w) = \{w_x \mid x \in I\}$. For a word w we denote its free reduction by $w^{\vee}$; additionally, if a word is simply locally reduced at x and freely reduced we say it is **locally reduced at** x and we say $(w_x)^{\vee}$ is a **local reduction** of w while denoting its **set of local reductions** by $L^{\vee}(w)$.

Remark. Notice that the definition of a local reduction given here is not arbitrary, since $(w_x)^{\vee} = ((w_x)^{\vee})_x$ for any word w on B and point $x \in [0, 1]$.

Considering local reductions of a given $g \in \langle B \rangle$ allows us to capture the action of g on any particular point. Conversely, it will also be fruitful to consider, given a particular point x, which actions of $\langle B \rangle$ act on x. More specifically, we are interested in which actions contain x in their image.

Definition 2.2.5. Given a fast set of bumps B and a marking with set of feet F we say $x \in \mathrm{supt}(B) \setminus \bigcup_{J \in F} J$ has **trivial history** and define $\bar{x} = \{b \mid x \in \mathrm{supt}(b)\}$. We denote $M = \{\bar{x} \mid x \text{ has trivial history}\}$. If $x \in \bigcup_{J \in F} J$ we define its **history** to be the set $\eta(x)$ containing the following words on $B \sqcup M$

- $u \in (B^{\pm})^*$ where there is some $t \in I$ such that $tu = x$, u is locally reduced at t and t is not in the source of the first letter of u
- $\bar{t}u$ for $u \in (B^{\pm})^*$ and some $t \in I$ with trivial history, $tu = x$ and u is locally reduced at t.

Let us now consider an example. In Figure 2.1 we see the dynamical diagram of a fast pair of bumps; denoting the left bump by a and the right by b, consider the word aba^{-1}.

Notice that if $t > ma^{-1}$ is close enough to ma^{-1} then we have $t \neq ta \neq tab \neq taba^{-1}$ and so aba^{-1} is simply locally reduced at t. Since it is also freely reduced,

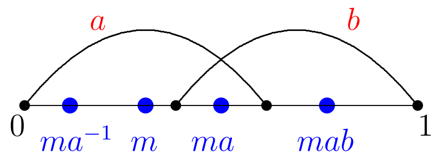

Figure 2.1: The dynamical diagram of a fast pair of bumps with support and marker m for a included.

it is locally reduced at t. If $x \in (0, ma^{-1}]$ then $xa \in (0, m]$ and it follows that $xab = xa$. We now see that aa^{-1} is the simply local reduction of aba^{-1} at x, and aa^{-1} is, of course, simply locally reduced. Thus ϵ is the local reduction of aba^{-1} at x. Similarly, if $y \in [m, 1) \cap \text{supt}(a)$ then it follows that $y < ya < yab$ and since $mab < yab$ we then have $yab = yaba^{-1}$. So the simply local reduction and local reduction of aba^{-1} at y coincide and are equal to ab. If $z \notin \text{supt}(a)$ then $za = z$ but $zb > z$, so b is the (simply) local reduction of aba^{-1} at z. We conclude that $L(aba^{-1}) = \{b, ab, aa^{-1}, aba^{-1}\}$ while $L^{\vee}(w) = \{\epsilon, b, ab, aba^{-1}\}$.

Let us now consider the histories of points with respect to this pair of bumps. By definition, m has trivial history since it is a marker. Take ma^{-1} and consider what length one words w on $B^{\pm}$ might form a history; by definition, ma^{-1} must be in the support of w but the point t such that $tw = ma^{-1}$ cannot be in its source. This only leaves a^{-1} and so $a^{-1} \in \eta(ma^{-1})$. It follows that any longer words on $B^{\pm}$ forming a history must have the form $a^{-1}w'$ where w' is a history of m, but m has trivial history. It follows that $\eta(ma^{-1}) = \{\epsilon, \bar{m}a^{-1}, a^{-1}\}$. More generally, if w is locally reduced at m then $\eta(mw) = \{\bar{m}w\} \cup \{s \mid s \text{ is a suffix of } w\}$. Indeed, points in the orbits of markers are precisely those points with finite history.

Analysing histories of points turns out to be particularly elucidating. Indeed, in [7] the authors go on to use the idea of histories to obtain an axiomatic, partial generalisation of fastness to arbitrary sets of permutations.

2.2.1 Abstract ping-pong systems and the space of histories

The following definitions are introduced in [7, Section 11].

Definition 2.2.6. Let A be a collection of permutations (without inverses) of a set X. A **ping-pong system** on A is an assignment $a \mapsto \text{dest}(a)$ of subsets of X to permutations in $A^{\pm}$ satisfying for each $a, b \in A^{\pm}$ and $x \in X$

(a) $\text{dest}(a) \subseteq \text{supt}(a)$ and if $\text{dest}(a) \cap \text{supt}(b) \neq \emptyset$ then $\text{dest}(a) \subseteq \text{supt}(b)$.

(b) if $a \neq b$ then $\text{dest}(a) \cap \text{dest}(b) = \emptyset$;

(c) if $x \in \text{supt}(a)$ then there exists a non-negative integer k such that $xa^k \in \text{dest}(a)$;

(d) if $x \in \mathrm{supt}(a)$ then $xa \in \mathrm{dest}(a)$ if and only if $x \notin \mathrm{src}(a)$;

where we define $\mathrm{src}(a) := \mathrm{dest}(a^{-1})$ for $a \in A^{\pm}$. We refer to the sets $\mathrm{dest}(a)$ for $a \in A^{\pm}$ as the **feet** of A.

As mentioned, we may generalise the idea of histories of a point to this setting.

Definition 2.2.7. Let A be a set of permutations of X equipped with a ping-pong system. If $x \in X \setminus \bigcup_{a\in A^{\pm}} \mathrm{dest}(a)$ we define $\bar{x} = \{a \in A \mid x \in \mathrm{supt}(a)\}$. Further define $M_A = \{\bar{x} \mid x \in X \setminus \bigcup_{a\in A^{\pm}} \mathrm{dest}(a)\}$. Consider a family η of words on $A^{\pm} \cup M_A$ which has the following properties:

(a) The family η is suffix closed;

(b) there is at most one element of length n for each n;

(c) each $s \in \eta$ does not contain subwords of the form $a\bar{x}$ for $a \in A^{\pm}$ and $\bar{x} \in M_A$;

(d) if $\bar{x}a$ is a subword of some $s \in \eta$ for $\bar{x} \in M_A$ and $a \in A^{\pm}$ then we have $a \in \bar{x}$;

(e) If $s \in \eta$ contains a subword ab where $a, b \in A^{\pm}$ then $\mathrm{dest}(a) \subseteq \mathrm{supt}(b) \setminus \mathrm{src}(b)$;

(f) if $s \in \eta$ doesn't contain any elements from M_A then s is a proper prefix of some $s' \in \eta$.

We denote the set of all such η by K_A. Further, we topologise K_A by stipulating the set of all sets of the form $[w] = \{\eta \in K_A \mid w \in \eta\}$ as a clopen basis and call it the **space of histories** of A.

In addition, we call an element of $L_A = \bigcup_{\eta\in K_A} \eta$ a **local reduction** of A, and refer to L_A as the set of local reductions. Notice that if $X \subseteq L_A$ then X satisfies conditions (c)-(e) and, then, $X \in K_A$ if it also satisfies (a),(b),(f).

Remark. The space of histories K_A is compact if A if finite.

Just as fast groups are fruitfully studied via their histories we may study an abstract ping-pong system via its space of histories. In particular, we obtain a representation of A via its action on K_A.

Definition 2.2.8. Let A be a set of permutations equipped with a ping-pong system and let K_A be its space of histories. For each $a \in A^{\pm}$ we define the homeomorphism $\hat{a} : K_A \to K_A$ by

$$\eta\hat{a} = \begin{cases} \{ua \mid u \in \eta\} \sqcup \{\varepsilon\} & \text{if } \mathrm{dest}(\eta) \subseteq \mathrm{supt}(a) \setminus \mathrm{src}(a) \\ \{u \mid ua^{-1} \in \eta\} & \text{if } \mathrm{dest}(\eta) = \mathrm{src}(a) \\ \eta & \text{if } \mathrm{dest}(\eta) \cap \mathrm{supt}(a) = \emptyset \\ \{\bar{x}, \varepsilon\} & \text{if } \eta = \{\bar{x}, \varepsilon\} \ \&\ a \notin \bar{x} \\ \{\bar{x}a, a, \varepsilon\} & \text{if } \eta = \{\bar{x}, \varepsilon\} \ \&\ a \in \bar{x} \end{cases}$$

where $\operatorname{dest}(\eta) := \operatorname{dest}(a)$ where $a \in A$ is the common suffix of every non-trivial element of η. We denote $\hat{A} = \{\hat{a} \mid a \in A\}$ and refer to it as the **ping-pong representation** of A, and say $\hat{A}$ is **faithful** if $\{\eta \in K_A \mid \eta \text{ is finite}\}$ is dense in K_A. We refer to $\langle \hat{A} \rangle$ as a **ping-pong group**, or a **faithful ping-pong group** if $\hat{A}$ is faithful.

Remark. Notice that $(\widehat{a^{-1}}) = (\hat{a})^{-1}$ for all $a \in A^{\pm}$.

Notice that the set $\{\eta \in K_A \mid \eta \text{ is finite}\}$ is, in essence, the set of points with finite history. A representation $\hat{A}$ being faithful, then, reflects the fact that the set of points with finite history for a fast set A is dense in $\operatorname{supt}(A)$, this playing an important role in the proof of the main theorem of [7].

As an analogue of the dynamical diagram of a fast set, and, in particular, a generalisation of Definition 2.2.1, we have the following.

Definition 2.2.9. Let A be a set of permutations equipped with a ping-pong system $a \mapsto \operatorname{dest}(a)$. Let $\mathcal{B}$ be a relation on $A^{\pm} \sqcup M_A$ containing:

- (a, b) if $a, b \in A^{\pm}$ and $\operatorname{dest}(a) \subseteq \operatorname{supt}(b)$;
- $(\bar{x}, a)$ if $\bar{x} \in M_A, a \in A$ and $a \in \bar{x}$.

Then we say $\mathcal{B}$ is the **blueprint** of A. We define $\operatorname{supt}_{\mathcal{B}}(a) = \{b \mid (b, a) \in \mathcal{B}\}$ for $a \in A^{\pm} \sqcup M_A$.

We say $\mathcal{B}$ is (totally) **orderable** if there exists a total order T on $A^{\pm} \sqcup M_A$ such that $\operatorname{supt}_{\mathcal{B}}(a)$ forms an interval of T with endpoints a^{-1}, a for every $a \in A$. We say it is **cyclically orderable** if there exists a cyclic order C on $A^{\pm} \sqcup M_A$ such that $\operatorname{supt}_{\mathcal{B}}(a)$ is an interval of C for every $a \in A$, and the endpoints of this interval are a, a^{-1} unless $\operatorname{supt}_{\mathcal{B}}(a) = A^{\pm} \sqcup M_A$.

Naturally, this leads to a generalisation of Theorem 2.2.2.

Theorem 2.2.10 ([7])**.** *Let A_1, A_2 be ping-pong representations with blueprints $\mathcal{B}_1, \mathcal{B}_2$ respectively. Then $\langle A_1 \rangle \cong \langle A_2 \rangle$ if $\mathcal{B}_1$ is relation isomorphic to $\mathcal{B}_2$.*

Ping-pong systems generalise fast bump sets; indeed, choosing a marking for a fast bump set B defines sources and destinations for each bump $b \in B$ giving us the assignment $b \mapsto \operatorname{dest}(b), b^{-1} \mapsto \operatorname{src}(b)$ and thereby defining a ping-pong system. Similarly, we may obtain a blueprint equivalent to the dynamical diagram of such a set B.

Again, let us consider the dynamical diagram in Figure 2.1. We have $A = \{a, b\}$

and $M_A = \{\bar{m}, \bar{m}a, \bar{m}\bar{a}b\}$ and, thus, find that the blueprint is the relation

$$\mathcal{B} = \{(a^{-1}, a^{-1}), (\bar{m}, a^{-1}), (b^{-1}, a^{-1}), (\bar{m}a, a^{-1}), (a, a^{-1}), (a^{-1}, a), (\bar{m}, a), (b^{-1}, a), (\bar{m}a, a), (a, a), (b^{-1}, b^{-1}), (\bar{m}a, b^{-1}), (a, b^{-1}), (\bar{m}\bar{a}b, b^{-1}), (b, b^{-1}), (b^{-1}, b), (\bar{m}a, b), (a, b), (\bar{m}\bar{a}b, b), (b, b)\}$$

and this blueprint is totally orderable as witnessed by the total ordering on $A^{\pm} \sqcup M_A$ corresponding to the total ordering on I in Figure 2.1; explicitly,

$$a^{-1} < \bar{m} < b^{-1} < \bar{m}a < a < \bar{m}\bar{a}b < b$$

witnesses that $\mathcal{B}$ is totally orderable. Indeed, as we will see in Proposition 5.1.8, totally orderable blueprints are equivalent to dynamical diagrams of fast bump sets in general.

2.3 Diagram groups

Diagrams and diagram groups were first studied in the PhD thesis of Kilibarda [36] and, afterwards, the theory around these groups had been developed chiefly by the work of Guba and Sapir (e.g. [29],[30],[31],[32],[33]). Since then, interest in this class of groups, and their various generalisations, has grown steadily, both as an object of study in its own right ([18],[19],[23]) and as avenue for producing methods to study groups, particularly Thompson-like groups ([24],[25],[26]). The class of diagram groups has been found to include a number of separate classes of groups of independent interest ([6],[20],[21]). See also Genevois's recent survey on diagram groups [22].

2.3.1 Planar diagrams

Planar diagrams - the 'classical' diagrams, as one might put it - are perhaps best understood as two-dimensional analogues of words, and diagram groups, similarly, as analogous to free groups. To define a diagram, we may start with a semigroup presentation

$$\langle x_1, \dots, x_m \mid u_1 = v_1, \ \dots, \ u_n = v_n \rangle$$

and consider a word $w = x_{i_1} \dots x_{i_k}$ over this presentation. This word can be represented as an edge-labelled oriented plane graph by a positive path Δ_0 with k edges with the jth edge from left to right labelled x_{i_j}. If w contains a subword u_i then we can replace it with v_i to obtain a word w' equivalent to w. This can be represented in the graph by connecting a positive path labelled v_i beneath the subpath labelled u_i

to obtain a new graph Δ_1. This can be repeated for Δ_1, and so on, indefinitely—each graph $\Delta_0, \Delta_1, \ldots$ obtained is a diagram over the semigroup presentation.

Defining diagrams formally will require a series of definitions, which we now introduce and discuss. First of all, a diagram is a graph in the sense of Serre.

Let V, E be disjoint sets and let $r : E \to E$ and $i, t : E \to V$ be functions. If the collection $\Gamma = (V, E, r, i, t)$ satisfies

- $r(e) \neq e$ and $r(r(e)) = e$ for all $e \in E$;
- $t(r(e)) = i(e)$ and $i(r(e)) = t(e)$ for all $e \in E$;

then we say it is a **graph**. If Γ is a graph then we refer to V as its vertex set and E as its edge set; their elements are called vertices and edges respectively. Given an edge e we call $i(e)$ its initial vertex and $t(e)$ its terminal vertex. We also refer to $r(e)$ as the inverse edge of e and we usually denote $e^{-1} := r(e)$. We call a sequence $p = (e_i)_{i=0,\ldots,n}$ of edges a path if $t(e_i) = i(e_{i+1})$ for every $i = 0, \ldots, n-1$. We say p begins at $i(e_0)$ and ends at $t(e_n)$, and we denote $i(p) := i(e_0), t(p) := t(e_n)$. If $t(e_n) = i(e_0)$ then we say p is **closed** and we call $i(e_0)$ its **base point**; $q = (e_{k+i \mod n})$ for some k then we say q is a **rotation** of p. We also define an **empty path** to be a single vertex $p_\emptyset = v \in V$.

Let $E_O \subseteq E$. If E_O satisfies $E_O \cup E_O^{-1} = E$ and $E_O \cap E_O^{-1} = \emptyset$ then we say E_O is an **orientation** for Γ. If E_O is an orientation for a graph Γ then we call the pair (Γ, E_O) an **oriented graph**, in which case we refer to the edges in E_O as positive and the edges in E_O^{-1} as negative. If a path p consists only of positive edges then we say p is positive. If $p = (e_i)_{i=0,\ldots,n}$ is a path then we define its **inverse path** p^{-1} to be the sequence $(e_{n-i}^{-1})_{i=0,\ldots,n}$.

Finally, if Γ is a graph let L be a set disjoint from V and E and let $l : E \to L \sqcup L^{-1}$ be a function where $L^{-1} = \{l^{-1} \mid l \in L\}$ is a set of formal symbols disjoint from L. If l satisfies $l(e^{-1}) = l(e)^{-1}$ for all $e \in E$ then we say the pair (L, l) is a **labelling** for Γ and the collection (Γ, L, l) is a **labelled graph** or a **graph labelled by** L. If $p = (e_i)_i$ is a path of a labelled graph Γ then we say the label $l(p)$ of p is equal to the concatenation $l(e_0)l(e_1)\ldots l(e_n)$. If the graph Γ is oriented then we assume $l(e) \in L$ for all positive edges $e \in E_O$.

More specifically, diagrams are isotopy classes of plane graphs; that is, particular embeddings of graphs into $\mathbb{R}^2$ considered identical up to continuous transformations of $\mathbb{R}^2$ which preserve the graphs' structure.

Consider $\mathbb{R}^2$. Let $V \subseteq \mathbb{R}^2$, let

$$E \subseteq \{\gamma : [0,1] \to \mathbb{R}^2 \text{ continuous embedding} \mid \gamma(0), \gamma(1) \in V\}$$

and define functions i, t, r by $i(e) = e(0)$, $t(e) = e(1)$ and $r(e(a)) = e(1-a)$ for

$e \in E, a \in [0,1]$. If $\Gamma = (V, E, r, i, t)$ is a graph satisfying

- if $e(a) \in V$ then $a = 0$ or $a = 1$;
- if $e(a) = f(b)$ for some a, b with at least one not in $\{0,1\}$ then $f = e$ or $f = e^{-1}$;

then we say Γ is a **plane graph**. A plane graph can be oriented or labelled just as a graph can.

Further, let $\delta : \mathbb{R}^2 \times [0,1] \to \mathbb{R}^2$ be a continuous function. We say δ is a **homotopy** of $\mathbb{R}^2$ if for every $a \in [0,1]$ the function $\delta_a : \mathbb{R}^2 \to \mathbb{R}^2$ defined by $\delta_a(x) = \delta(x, a)$ is continuous; we say δ is an **isotopy** if δ_a is a homeomorphism for all $a \in [0,1]$. If $\gamma_1, \gamma_2 : [0,1] \to \mathbb{R}^2$ are continuous embeddings such that $\gamma_1(0) = \gamma_2(0)$ and $\gamma_1(1) = \gamma_2(1)$ then we say γ_1 is **above** γ_2 (equivalently, γ_2 is **below** γ_1) if there exists an isotopy δ of $\mathbb{R}^2$ such that $\delta_0 \circ \gamma_1 = \gamma_1, \delta_0 \circ \gamma_2 = \gamma_2$ and $\delta_1 \circ \gamma_1(x) = (x, f(x))$, $\delta_1 \circ \gamma_2(x) = (x, g(x))$ for continuous functions f, g where $f(x) \geq g(x)$ for all $x \in [0,1]$. Notice that this formalism can be applied to oriented paths in general. If the only solutions to $\gamma_1(a) = \gamma_2(a)$ are $a = 0, 1$ for continuous embeddings γ_1, γ_2 then either γ_1 is above γ_2 or γ_2 is above γ_1.

Finally, let $\Gamma_1 = (V_1, E_1, r, i, t)$ and Γ_2 be plane graphs with vertex sets V_1, V_2 and edge sets E_1, E_2 respectively. We say Γ_1 and Γ_2 are **isotopic** if there exists an isotopy δ of $\mathbb{R}^2$ such that $\delta_0(v) = v, \delta_0 \circ e = e$ for all $v \in V_1, e \in E_1$ while $\delta_1(V_1) = V_2, \delta_1 \circ E_1 = E_2$ and $\Gamma_a = (\delta_a(V_1), \delta_a \circ E_1, r^{\delta_a}, i^{\delta_a}, t^{\delta_a})$ is a plane diagram for all $a \in [0,1]$. If Γ_1 is labelled or oriented we, in addition, require that its orientation and labelling is preserved by the isotopy to consider it isotopic to Γ_2.

Remark. Isotopy defines an equivalence relation on the set of plane graphs, since an isotopy $\delta(x,t)$ can be have its orientation reversed $\delta(x, 1-t)$ and still be an isotopy, while two isotopies can be 'concatenated' to form a new isotopy. Further, we can see from this definition that an isotopy between two (labelled, oriented) plane graphs induces a (labelled, oriented) graph isomorphism between them, since we have a map between their vertices which respects their adjacency relations. This is important since, as we will see, we will want to distinguish between certain diagrams which are isomorphic as graphs.

This is enough background for us to define diagrams in specific. We begin by defining the most basic 'pieces' out of which a general diagram will be 'built'.

Definition 2.3.1. Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and let $r := (u, w) \in \mathcal{R}$ be a relation. We define the **elementary diagram over $\mathcal{P}$ labelled by** r to be the isotopy class of oriented plane graphs labelled by Σ containing Ψ_r such that

Figure 2.2: The two elementary diagrams over the semigroup presentation $\langle a, b \mid ab = ba \rangle$. They are isomorphic, but not isotopic.

(a) there are two distinguished vertices denoted $i(\Psi_r)$ and $t(\Psi_r)$, called the **initial** and **terminal** vertices respectively;

(b) Ψ_r has precisely two maximal positive paths, both beginning at $i(\Psi_r)$, ending at $t(\Psi_r)$ and otherwise disjoint from each other - one positive path is above and the other is below, so we call them the **top path** and the **bottom path** and denote them by $\lceil \Psi_r \rceil$, $\lfloor \Psi_r \rfloor$ respectively;

(c) the positive paths $\lceil \Psi_r \rceil$, $\lfloor \Psi_r \rfloor$ are labelled by u, w respectively.

We also define the **elementary diagram over $\mathcal{P}$ labelled by** r^{-1} in exactly the same way, except the top path is labelled u and the bottom path is labelled w.

Given a word $w \in \Sigma^*$ we define the **trivial diagram over $\mathcal{P}$ labelled by** w to be the isotopy class of oriented plane graphs labelled by Σ which contains the graph ε_w satisfying (1) and has precisely one maximal positive path which begins at $i(\varepsilon_w)$, ends at $t(\varepsilon_w)$ and is labelled by w.

Remark. We may still refer to the top and bottom path of a trivial diagram, but in this case they both refer to the unique maximal positive path. Indeed, a trivial diagram ε_w may be thought of as corresponding to the trivial relation (w, w). Notice that there is a trivial diagram labelled by the empty word ε_ϵ which consists of a single vertex, and its unique maximal positive path is the empty path at this vertex.

Notice that the difference between isotopy and isomorphism has already become important, since isotopy allows us to distinguish between the elementary diagrams labelled r and r^{-1} while isomorphism does not. An example is shown in Figure 2.2. By convention, when sketching a diagram we will always place the initial vertex to the left of the terminal vertex (with both usually lying on a common horizontal line) and leave the orientation on the edges implicit.

Elementary diagrams are analogous to letters where diagrams are analogous to words. The second dimension, so to speak, comes from the fact that, given a diagram, there may be multiple different 'locations' where an elementary diagram can be appended, each forming a distinct diagram, whereas given a word there is only one word which can be formed by appending a given letter.

We make this precise in the following definitions. First, we require an operation.

Definition 2.3.2. Let Ψ_1, Ψ_2 be either elementary or trivial diagrams over $\mathcal{P}$. We define **the sum** $\Psi_1 + \Psi_2$ to be the isotopy class of oriented graphs labelled by Σ containing the graphs formed by identifying (via a homotopy of $\mathbb{R}^2$, assuming they are initially disjoint) the terminal vertex $t(\Psi_1)$ with the initial vertex $i(\Psi_2)$ of any representatives for Ψ_1, Ψ_2. We consider this graph to have two distinguished vertices and two distinguished paths as before. By definition

- $i(\Psi_1 + \Psi_2) = i(\Psi_1)$;
- $t(\Psi_1 + \Psi_2) = t(\Psi_2)$;
- $\lceil \Psi_1 + \Psi_2 \rceil = \lceil \Psi_1 \rceil \circ \lceil \Psi_2 \rceil$;
- $\lfloor \Psi_1 + \Psi_2 \rfloor = \lfloor \Psi_1 \rfloor \circ \lfloor \Psi_2 \rfloor$.

We can similarly extend this to any sum of finitely many elementary or trivial diagrams $\Psi_1, \ldots, \Psi_n$.

Remark. Notice that for a word $w = a_1 \ldots a_n$ the trivial diagram ε_w decomposes uniquely as $\varepsilon_{a_1} + \ldots + \varepsilon_{a_n}$.

Since this sum operation is defined by using a homotopy to identify vertices on the boundary of the graphs we can see, by composing the relevant maps, that it does not depend on the choice of representatives. It is associative since it corresponds to concatenation of labels of their positive paths, and it is not commutative for the same reason. With this in place we can now formalise the idea of appending the same elementary diagram at different locations. Let u, w be words over Σ and let r be a(n inverse of a) relation of $\mathcal{P}$. We call any isotopy class of the form $\varepsilon_u + \Psi_r + \varepsilon_w$ an **atomic diagram over** $\mathcal{P}$.

Definition 2.3.3. Let $\Sigma_i \Psi_i, \Sigma_j \Psi_j$ be finite sums of elementary or trivial diagrams such that the paths $\lfloor \Sigma_i \Psi_i \rfloor$ and $\lceil \Sigma_j \Psi_j \rceil$ have the same label in Σ^+. We define the **concatenation** $\Sigma_i \Psi_i \circ \Sigma_j \Psi_j$ to be isotopy class of oriented graphs labelled by Σ containing the graphs formed by identifying (via a homotopy of $\mathbb{R}^2$, assuming they are initially disjoint) the paths $\lfloor \Sigma_i \Psi_i \rfloor$ and $\lceil \Sigma_j \Psi_j \rceil$ of any representatives for $\Sigma_i \Psi_i, \Sigma_j \Psi_j$. By definition

- $\lceil \Sigma_i \Psi_i \circ \Sigma_j \Psi_j \rceil = \lceil \Sigma_i \Psi_i \rceil$;
- $\lfloor \Sigma_i \Psi_i \circ \Sigma_j \Psi_j \rfloor = \lfloor \Sigma_j \Psi_j \rfloor$;

while $i(\Sigma_i \Psi_i \circ \Sigma_j \Psi_j)$ and $t(\Sigma_i \Psi_i \circ \Sigma_j \Psi_j)$ are the vertices formed by identifying $i(\Sigma_i \Psi_i)$ with $i(\Sigma_j \Psi_j)$ and $t(\Sigma_i \Psi_i)$ with $t(\Sigma_j \Psi_j)$ respectively.

We can similarly extend this to finitely many finite sums.

As with the sum operation, we can see that concatenation is well defined on isotopy classes of these graphs. We can also see that it is associative and non-commutative since it corresponds to concatenation of labels of the positive paths in their dual graphs (see Section 2.3.2 on strand diagrams). Unlike the sum operation, concatenation is not necessarily defined for an arbitrary pair of diagrams.

Finally, we may define diagrams in general.

Definition 2.3.4. Let $\mathcal{P}$ be a semigroup presentation. We say an isotopy class of labelled oriented graphs Δ is a **diagram over** $\mathcal{P}$ if either it is equal to a trivial diagram or it decomposes as a concatenation of atomic diagrams over $\mathcal{P}$. If Δ' is a subgraph of Δ which is equal to a diagram over $\mathcal{P}$ such that Δ' contains every part of Δ within the region bounded by the path $\lceil\Delta'\rceil \circ \lfloor\Delta'\rfloor^{-1}$ then we say Δ' is a **subdiagram** of Δ. A subdiagram of Δ equal to an elementary diagram is called a **cell** of Δ. If Δ has top path labelled u and bottom path labelled w we may call it a **(u,w)-diagram**; similarly, if π is a cell of Δ such that $\lceil\pi\rceil$ is labelled by u and $\lfloor\pi\rfloor$ is labelled by w we say π is a **(u,w)-cell** and sometimes denote such a π by (u, w). If $u = w$ then we say Δ is **spherical with base** w.

We will denote the set of all diagrams over $\mathcal{P}$ by $\mathcal{D}(\mathcal{P})$, or just $\mathcal{D}$ when this will not lead to confusion. We also denote by $\mathcal{D}^+, \mathcal{D}^- \subseteq \mathcal{D}$ the set of positive and negative diagrams over $\mathcal{P}$ respectively. Given a diagram $\Delta \in \mathcal{D}(\mathcal{P})$ it may be useful to consider the sets of its incoming and outgoing positive edges from a given vertex v of Δ, which we will denote by $I(v), O(v)$ respectively. Notice that these sets are naturally totally ordered clockwise and anticlockwise rotation, respectively, around v in a planar embedding of Δ. Relatedly, we also define $I^\pi(v)$ and $O^\pi(v)$ to be the totally ordered sets of cells of Δ bounded between consecutive elements of $I(v)$ and $O(v)$ respectively.

Notice that Definition 2.3.2 extends to a binary operation on $\mathcal{D}$ while Definition 2.3.3 extends to a partial operation on $\mathcal{D}$. If we write $\mathcal{D}(\mathcal{P}, w) \subseteq \mathcal{D}(\mathcal{P})$ (or just $\mathcal{D}(w)$) to denote the set of all spherical diagrams with base w then concatenation restricts to a binary operation on $\mathcal{D}(w)$.

As an example, consider the semigroup presentation $\langle a, b \mid ab = ba\rangle$ and the derivation

$$aabb \to abab \to baab \to baba \to bbaa$$

from the word $aabb$ to $bbaa$. The diagram defined by this derivation over this presentation is shown in Figure 2.3. In fact, the derivation

$$aabb \to abab \to abba \to baba \to bbaa$$

defines precisely the same diagram. This corresponds to the fact that this dia-

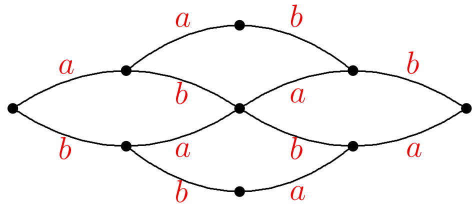

Figure 2.3: An $(aabb, bbaa)$-diagram over the semigroup presentation $\langle a, b \mid ab = ba \rangle$

gram has two distinct decompositions into atomic diagrams, which can be written algebraically as

$$(\varepsilon_a + \Psi_r + \varepsilon_b) \circ (\Psi_r + \varepsilon_{ab}) \circ (\varepsilon_{ba} + \Psi_r) \circ (\varepsilon_b + \Psi_r + \varepsilon_a)$$

and

$$(\varepsilon_a + \Psi_r + \varepsilon_b) \circ (\varepsilon_{ab} + \Psi_r) \circ (\Psi_r + \varepsilon_{ba}) \circ (\varepsilon_b + \Psi_r + \varepsilon_a)$$

respectively, where $r = (ab, ba)$.

At this point, we have that the pair $(\mathcal{D}(w), \circ)$ forms a monoid with the trivial diagram ε_w acting as the identity. We now introduce another equivalence relation, this time on $\mathcal{D}$, by which the quotient of $(\mathcal{D}(w), \circ)$ will be a group. This equivalence is analogous to the cancelling pairs equivalence on words which takes us from the free monoid $(\Sigma \sqcup \Sigma^{-1})^*$ to the free group of rank $|\Sigma|$.

Definition 2.3.5. Consider $\mathcal{D} = \mathcal{D}(\mathcal{P})$ and let $r = (u, w)$ be a relation of $\mathcal{P}$. Then $\Psi_r \circ \Psi_{r^{-1}}$ is a (u, u)-diagram while $\Psi_{r^{-1}} \circ \Psi_r$ is a (w, w)-diagram, and both are in $\mathcal{D}$ (pictures in Figure). We refer to diagrams of either of these forms as **dipoles**. If a diagram $\Delta \in \mathcal{D}$ contains a dipole as a subdiagram then we say it contains this dipole.

Suppose $\Delta \in \mathcal{D}$ contains a dipole $\Psi_r \circ \Psi_{r^{-1}}$ (equivalent for $\Psi_{r^{-1}} \circ \Psi_r$). We define a new diagram Δ' from Δ by

(a) deleting from Δ the path labelled w in this dipole;

(b) identifying the paths $\lceil \Psi_r \circ \Psi_{r^{-1}} \rceil$ and $\lfloor \Psi_r \circ \Psi_{r^{-1}} \rfloor$, both labelled u, in Δ.

We call this process **reducing a dipole**. We can also consider the inverse process - given a diagram Δ with a path p labelled u we embed a copy p' of p beginning at $i(p)$ and ending at $t(p)$ so that p' does not intersect any other part of Δ and then add another path q with the same properties as p, except it is between p and p' and it is labelled by w. We call this process **inserting a dipole**.

If $\Delta \in \mathcal{D}$ does not contain any dipoles then we say Δ is **reduced**.

Remark. To see that Δ' is a diagram, notice that there must exist a decomposition

$\Delta = \Delta_1 \circ (\varepsilon_x + (\Psi_r \circ \Psi_{r^{-1}}) + \varepsilon_y) \circ \Delta_2$ for some words x, y, since otherwise the dipole could not be a subdiagram of Δ. It follows that $\Delta' = \Delta_1 \circ \Delta_2$.

Notice that the 'reducing a dipole' process induces a relation R on $\mathcal{D}$; $(\Delta, \Delta') \in R$ if Δ' is obtained from Δ by reducing a dipole. Naturally, the 'inserting a dipole' process induces its converse relation R^C. Thus, the union of these relations $R \sqcup R^C$ is symmetric. Therefore, the reflexive, transitive closure of $R \sqcup R^C$, which we denote by $\equiv$, is an equivalence relation on $\mathcal{D}$. Notice that $\Delta_1 \equiv \Delta_2$ is equivalent to the statement 'Δ_1 can be obtained from Δ_2 by some (possibly empty) sequence of dipole reductions and insertions'.

Suppose that $\Delta_1 \equiv \Delta_1', \Delta_2 \equiv \Delta_2'$. Notice that, whatever dipoles were removed/inserted in Δ_1 to obtain Δ_1' can still be removed and inserted in $\Delta_1 \circ \Delta_2$ to obtain $\Delta_1' \circ \Delta_2$ and then, similarly, we can insert/remove dipoles in the copy of Δ_2 in $\Delta_1' \circ \Delta_2$ to obtain $\Delta_1' \circ \Delta_2'$. It follows that $\circ$ induces a partial operation (which we also denote $\circ$) on the quotient set $\mathcal{D}/\equiv$ such that $(\mathcal{D}(x)/\equiv, \circ)$ is a monoid. In fact, $(\mathcal{D}(x)/\equiv, \circ)$ forms a group; if $\mathcal{P}$ be a semigroup presentation and w be a word on Σ, then $D(\mathcal{P}, w) := (\mathcal{D}(\mathcal{P}, w)/\equiv, \circ)$ is a group, known as the **diagram group over $\mathcal{P}$ with base** w.

Remark. Given $\Delta \in D(\mathcal{P}, w)$ we can find a representative for its inverse Δ^{-1} by reflecting a plane graph representative of Δ across a horizontal line in $\mathbb{R}^2$ (i.e. flip Δ upside-down). To see this, consider an atomic decomposition of Δ - an atomic decomposition for Δ^{-1} is found by simply replacing the cell in each atomic diagram by the cell of its inverse relation and reversing the order. Every cell cancels as a dipole, leaving us with the trivial diagram ε_w.

As was originally shown in [36], every $\equiv$-equivalence class of diagrams contains a unique reduced diagram. This basic fact is used frequently when working with diagram groups in general, and can be used to deduce some important properties; in particular, it follows that diagram groups have word problem which can be solved rather efficiently, since it suffices to fully reduce a diagram to determine whether or not it represents the identity.

Diagram groups form a broad class of groups. All the diagram groups we consider will fall into a particular subclass, which we now define.

Definition 2.3.6. Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation. If, for a given $a \in \Sigma$, there is at most one relation of the form $(a, w) \in \mathcal{R}$ and, if this relation exists, we have $\text{len}(w) > 1$ then we say $\mathcal{P}$ is **tree-like**. In this case, we also say $D(\mathcal{P}, w)$ is tree-like for any base w. We say a cell π of a diagram over a tree-like presentation is **positive** if it has the form $\pi = (a, w)$ and **negative** if it has the form $\pi = (w, a)$. We call Δ **positive** if all of its cells are positive, and **negative** if all of its cells are negative.

We will give an equivalent representation for diagrams over tree-like presentations in Section 2.3.3.

We conclude this section by stating a theorem of Guba and Sapir appearing as [30, Theorem 4.1] which will be used in Chapters 4 and 5. This theorem is originally stated in terms of (the equivalent definition of) diagram groups over directed 2-complexes; we state the equivalent theorem in terms of diagram groups over semigroup presentations.

Theorem 2.3.7 (Guba-Sapir moves [30])**.** *Consider a semigroup presentation $\mathcal{P}$.*

(a) *Suppose $u_i = v_i$ and $u_j = v_j$ are distinct relations in $\mathcal{P}$. If either u_j or v_j has u_i as a subword, then replacing this u_i subword with v_i in the relation $u_j = v_j$ does not change the isomorphism type of $D(\mathcal{P}, u)$ for any word u.*

(b) *Suppose x is a generator in $\mathcal{P}$ and only appears in one relation, which has the form $x = w$ for some word w that does not contain x. In this case, removing the generator x as well as the relation $x = w$ does not change the isomorphism type of $D(\mathcal{P}, u)$ for any word u not containing x.*

2.3.2 Strand diagrams

For our purposes it will be useful to consider an alternative representation of elements of diagram groups. These are known as **strand diagrams**, a term coined by Belk in [4] (see also [2] for a more general version) although similar constructions elsewhere in other forms (e.g. [10]). Strand diagrams are closely related to the "transistor" pictures for diagram groups described by Guba and Sapir in [29, Section 4] and equivalent to the planar subgroup of the braided diagram groups they later define in [29, Section 16]. Indeed, strand diagrams are dual to diagrams. Explicitly,

Definition 2.3.8. Given a diagram Δ over a semigroup presentation $\mathcal{P}$ we define its **strand diagram** Ψ as follows

(a) There is a vertex for every cell of Δ;

(b) There is an edge from v_1 to v_2 if the corresponding cells π_1, π_2 have a shared boundary edge $e \in \lfloor \pi_1 \rfloor \cap \lceil \pi_2 \rceil$ in Δ;

(c) There is a vertex for each cell on the top (bottom) path, and an edge from (to) such a vertex to (from) another if the corresponding top (bottom) path edge forms part of the boundary of the corresponding cell.

and we label each component of Ψ with the label of the corresponding component of Δ. We refer to the vertices defined in (1) as **interior vertices** and those defined in (3) as **boundary vertices**.

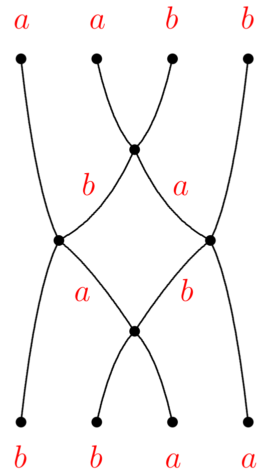

Figure 2.4: The strand diagram corresponding to the diagram shown in Figure 2.3

Remark. An example of a strand diagram is shown in Figure 2.4. Statements (1) and (2) define the dual of Δ while statement (3) provides a tweak which allows us to concatenate strand diagram as we would diagrams.

We can define dipoles and reductions in much the same way as with diagrams, and as such we can define the strand diagram group $S(\mathcal{P}, w)$ as the group of equivalence classes of (w, w)-strand diagram over $\mathcal{P}$. Indeed, the function above, which we take as defining strand diagrams, induces an isomorphism between these groups, so $S(\mathcal{P}, w) \cong D(\mathcal{P}, w)$.

We borrow an idea from [2] characterising interior vertices which we will make use of in some proofs in Chapter 4. If an interior vertex has precisely one incoming edge and at least two outgoing edges we call it a **split**. Similarly, if an interior vertex has precisely one outgoing edge and at least two incoming edges we call it a **merge**. Notice that if our semigroup presentation is tree-like then the above distinction dichotomises the interior vertices of any strand diagram over $\mathcal{P}$. Indeed, a split is equivalent to a positive cell while a merge is equivalent to a negative cell.

2.3.3 Tree diagrams

An interesting fact about reduced diagrams Δ over a tree-like presentation $\mathcal{P}$ is that, as first observed in [32], they decompose as $\Delta = \Delta_+ \circ \Delta_-$ where Δ_+ is positive and Δ_- is negative. This gives us the ‘tree diagram’ $(\Delta_+, \Delta_-^{-1})$ of Δ, so-called since the strand diagrams of Δ_+ and Δ_-^{-1} are trees. In this subsection we give this idea some formal treatment. Afterwards, with the aid of said formalism, we introduce a purely combinatorial definition of a tree diagram. This lays the groundwork for defining infinite tree diagrams in Chapter 6.

Two positive diagrams Δ_1, Δ_2 form a tree diagram if and only if $\lfloor \Delta_1 \rfloor$ and $\lfloor \Delta_2 \rfloor$ have the same label, say w. If ψ is an atomic diagram (with cell π) with $\lceil \psi \rceil$ labelled

w then $(\Delta_1 \circ \psi, \Delta_2 \circ \psi)$ is the tree diagram of the diagram obtained from $\Delta_1 \circ \Delta_2^{-1}$ by inserting a (π, π^{-1}) dipole. We therefore consider tree pairs equivalent up to adding or removing a common suffix. Formally,

Definition 2.3.9. Let $\mathcal{P}$ be a tree-like semigroup presentation. Consider the subset $\lfloor \mathcal{D}^+ \rfloor \subseteq \mathcal{D}^+ \times \mathcal{D}^+$ of pairs (Δ_1, Δ_2) such that $\lfloor \Delta_1 \rfloor, \lfloor \Delta_2 \rfloor$ have the same label, and define on it a partial operation, denoted $\circ$, by

$$(\Delta_1, \Delta_2) \circ (\Delta_2, \Delta_3) = (\Delta_1, \Delta_3)$$

We also define an equivalence relation $\equiv$ on $\lfloor \mathcal{D}^+ \rfloor$ by taking the reflexive, symmetric, transitive closure of

$$(\Delta_1, \Delta_2) \equiv (\Delta_1 \circ \psi, \Delta_2 \circ \psi)$$

for a positive atomic diagram ψ. The partial operation $\circ$ induces a binary operation on $D_t(w) = \lfloor \mathcal{D}^+ \rfloor (w)/\equiv$, where $\lfloor \mathcal{D}^+ \rfloor (w)$ is the subset of pairs (Δ_1, Δ_2) such that $\lceil \Delta_1 \rceil, \lceil \Delta_2 \rceil$ are labelled w, so that $(D_t(w), \circ)$ forms a group called the **tree diagram group** over $\mathcal{P}$ with base w. We say a pair $(\Delta_1, \Delta_2) \in \lfloor \mathcal{D}^+ \rfloor (w)$ is **reduced** if its maximal common suffix is trivial.

As indicated by the above discussion, the map $\Delta \mapsto (\Delta_+, \Delta_-^{-1})$ induces an isomorphism $D(w) \to D_t(w)$, and we can use this map to show that $D_t(w)$ is indeed a group. To see this more explicitly, we introduce an object associated with the tree-like presentation $\mathcal{P}$ which will also be relevant for later discussions.

Definition 2.3.10. Let $\mathcal{P}$ be a tree-like semigroup presentation and let w be a word. Start with $\Delta_0 = \varepsilon_w$ and inductively define Δ_n for $n > 0$ by

$$\Delta_n := \Delta_{n-1} \circ (\Sigma_{e \in \lfloor \Delta_{n-1} \rfloor} \pi_e)$$

with $\pi_e = \pi_r$ where r is the unique relation of the form $(l(e), w)$ if it exists, and $\pi_e = \varepsilon_{l(e)}$ otherwise. We define $\mathcal{M}(w) = \mathcal{M}(\mathcal{P}, w) = \bigcup_{n \in \mathbb{N}} \Delta_n$ to be the **master diagram** over $\mathcal{P}$ with base w.

Notice that, given $\Delta \in \mathcal{D}^+$ with top path labelled w, the master diagram $\mathcal{M}(w)$ contains Δ uniquely as a prefix, and every prefix of $\mathcal{M}(w)$ is a positive diagram with top path labelled w. Strictly speaking, $\mathcal{M}(w)$ is not a diagram.

Now, suppose $(\Delta_1, \Delta_2), (\Delta_3, \Delta_4) \in \lfloor \mathcal{D}^+ \rfloor (w)$ are reduced; both Δ_2, Δ_3 are contained as prefixes of $\mathcal{M}(w)$ and so their union $\Delta_2 \cup \Delta_3$ inside $\mathcal{M}(w)$ is also a positive diagram with top path labelled w and there exists $\Delta', \Delta'' \in \mathcal{D}^+$ such that $\Delta_2 \cup \Delta_3 = \Delta_2 \circ \Delta' = \Delta_3 \circ \Delta''$. We therefore have $(\Delta_1, \Delta_2) \equiv (\Delta_1 \circ \Delta', \Delta_2 \circ \Delta')$ and $(\Delta_3, \Delta_4) \equiv (\Delta_3 \circ \Delta'', \Delta_4 \circ \Delta'')$ and so we have $[(\Delta_1, \Delta_2)]_\equiv \circ [(\Delta_3, \Delta_4)]_\equiv =$

$[(\Delta_1 \circ \Delta', \Delta_4 \circ \Delta'')]_{\equiv}$. As such, $\circ$ induces a binary operation on $D_t(w)$ making it a group, where $(\Delta_1, \Delta_2)^{-1} = (\Delta_2, \Delta_1)$.

Combinatorial construction

Notice that the construction of a master diagram $\mathcal{M}$ strongly resembles that of an (irregular) infinite labelled rooted tree. We now make this precise, thereby obtaining a combinatorial construction of diagrams over a tree-like semigroup presentation. Indeed, recall that we described diagrams as two-dimensional analogue of words. Since words are nothing more than labelled totally ordered sets it is natural to suppose that an analogous construction exists for diagrams. Of course, the features of such a construction is 'baked-in' to the isotopy class definition of diagrams and so in our construction we begin by drawing these features out of this definition. This is most easily seen in the tree-like case, which is sufficient for our purposes.

Let $\mathcal{P}$ be a tree-like semigroup presentation and w a word. Define a partial order T on the positive edges E_+ of $\mathcal{M} := \mathcal{M}(\mathcal{P}, w)$ as the transitive closure of the relation defined by $(e, f) \in T$ when $e \in \lceil \pi \rceil$ and $f \in \lfloor \pi \rfloor$ for some cell π of $\mathcal{M}$. Additionally, define L_π to be the total order on $\lfloor \pi \rfloor$ defined by the positive path $\lfloor \pi \rfloor$ itself for each cell π, as well as a total order L_0 on the edges of top path defined by w. We define $M(\mathcal{P}, w) := (E_+, T, L_0, (L_\pi)_{\lceil \pi \rceil \in E_+})$ to be the **combinatorial master diagram** over $\mathcal{P}$ with base w.

Remark. Notice that $M(\mathcal{P}, w) := (E_+, T, L_0, (L_\pi)_{\lceil \pi \rceil \in E_+})$ a locally ordered labelled tree by Definition 2.1.3.

There are orderings 'baked-in' to $\mathcal{M}(\mathcal{P}, w)$ beyond those we have mentioned - in particular, each positive path of $(\mathcal{M}(\mathcal{P}, w)$ defines a total order on some subset of E_+. However, given that these paths come out of concatenating cells, one might expect that the corresponding total orders will come out of $M(\mathcal{P}, w)$. Indeed, if Lex is the lexicographical order for the locally ordered tree $M(\mathcal{P}, w)$ then the order on a maximal positive path p is given by the induced suborder of Lex on the edges of p.

Definition 2.3.11. Let $M := M(\mathcal{P}, w) = (E_+, T, L_0, (L_a)_{e \in E_+})$ be the combinatorial master diagram over $\mathcal{P}$ with base w. Suppose $(D, M|_D)$ is a finite initial subtree of M such that if $x \in D$ such that $\mathrm{is}_D(x) \neq \emptyset$ then $\mathrm{is}_D(x) = \mathrm{is}_{E_+}(x)$. Then we define $(D, M|_D)$ to be a **positive diagram**. We define the **top path** $\lceil D \rceil$ of D to be labelled totally ordered set $(L_{E_+}(0), L_0)$ while the **bottom path** is defined to be the totally ordered set $(\mathrm{end}(D), \mathrm{Lex}_M|_{\mathrm{end}(D)})$ where Lex_M is the lexicographic order on M. As before, we denote the set of positive diagram over $\mathcal{P}$ with base w by $\mathcal{D}^+(\mathcal{P}, w)$.

Notice that choosing Lex as the total order on this formalisation is not arbitrary. The point is that if we isolate a positive path in M as a subset then restricting Lex to this subset will recover the total order on the path. Indeed, we may consider the diagram from Figure 2.3 and see that the ordered subsets of its edges defined by its positive paths coincides with Lex when restricted to the same subsets. The following lemma establishes this fact in generality.

Lemma 2.3.12. *Let* $M = (E, T, L_0, (L_e)_{e\in E})$ *be the locally ordered labelled tree defined by a master diagram* $\mathcal{M}(\mathcal{P}, w)$ *and let* L_p *be the total order on the edges of* p *defined by* p *for each positive path* p *of* $\mathcal{M}(\mathcal{P}, w)$*. Then*

$$\mathrm{Lex}_M = T \sqcup \bigcup_{p\in P_+} L_p$$

where P_+ *is the set of positive paths of* $\mathcal{M}(\mathcal{P}, w)$*.*

Proof. First, notice that $T \subseteq \mathrm{Lex}_M$ by definition. As such, it suffices to show that $P = \mathrm{Lex}_M \setminus T$ where $P := \bigcup_{p\in P_+} L_p$.

Claim: Let $A \subseteq E$. Then A forms an antichain of T if and only if A forms a chain of P.

Proof: Recall from Definition 2.3.10 that we may obtain any positive diagram with top path label w as a prefix of the master diagram by starting with ε_w and concatenating cells step-by-step. Indeed, any such sequence of concatenating cells defines a unique diagram and the bottom path of this is a positive path of $\mathcal{M}(\mathcal{P}, w)$. If A is a chain of P then it is contained in $\lfloor\Delta\rfloor$ for some positive diagram Δ. If $T|_D$ is the tree part of D the initial subtree of M corresponding to Δ then each $f \in \lfloor\Delta\rfloor$ is a maximal element of $T|_D$ and $\lfloor\Delta\rfloor$ must therefore form an antichain of $T|_D$. Since $T|_D$ is induced from T we conclude A is an antichain of T.

Suppose Δ is a positive diagram whose bottom path contains $A\setminus\{e\}$ for some $e \in A$ where A is an antichain of T, then there exists some $f \in \lfloor\Delta\rfloor$ such that $e, f \in C$ for some chain C of T since D an initial subtree of M where if $x \in D$ such that $\mathrm{is}_D(x) \neq \emptyset$ then $\mathrm{is}_D(x) = \mathrm{is}_E(x)$, where D is the tree corresponding to Δ. If $(e, f) \in T$ then we may remove from Δ precisely those cells which have edges d satisfying $(e, d) \in T$; if this involves removing some $d \in A \setminus \{e\}$ then A is not an antichain of T. Similarly, if $(f, e) \in T$ then we add precisely those cells which have edges d satisfying $(f, d), (d, e) \in T$ to obtain Δ'; since T is a tree no cell is added over edges in $\lfloor\Delta\rfloor \setminus \{f\}$ so we have $A \subseteq \lfloor\Delta'\rfloor$. □

If $(x, y)\, \mathrm{Lex}_M \setminus T$ then either $(x, y) \in P$ or $(y, x) \in P$ by the claim, and we must then have $(x, y) \in P$ since positive paths move left to right and, by the definition of Lex_M, there is a maximal common prefix p of $\mathrm{pr}(x), \mathrm{pr}(y)$ so that x_p is to the left of y_p for immediate successors of p in $\mathrm{pr}(x), \mathrm{pr}(y)$ respectively; since T is a tree x must be to the left of y. If $(x, y) \in P$ then $(x, y), (y, x) \notin T$ by the claim and there is some positive path from x to y; if x, y are edges on the bottom path of the same cell then $(x, y) \in \mathrm{Lex}_M$, otherwise x, y are edges of different cells and we may trace our way up through the top paths of cells until we find edges x', y' in the same bottom path of a cell. We must have $x' < y'$ in this path since $(x, y) \in P$ and it follows that $(x, y) \in \mathrm{Lex}_M$. □

Concatenation can be defined in much the same way as before; if D_1, D_2 are positive diagrams such that $\lfloor D_1 \rfloor$ and $\lceil D_2 \rceil$ have the same label we concatenate by taking the union and identifying these paths, which we can do since there is a unique non-converse label-preserving order isomorphism $f : \lfloor D_1 \rfloor \to \lceil D_2 \rceil$.

Importing these new constructions into Definition 2.3.9 now allows us to view a tree diagram group as consisting of equivalence classes of pairs of locally ordered trees rather than equivalence classes of pairs of isotopy classes of plane graphs. As noted, the locally ordered tree is, in fact, contained in the corresponding isotopy class of plane graphs and in most cases we will still work with these objects via their planar embeddings. As mentioned at the start of this subsection, the primary purpose of this is in anticipation of Chapter 6.

We now have three interconnected ways of looking at a diagram group: via diagrams, via strand diagrams and via tree diagrams. It is with the last point of view that the connection between diagram groups and Thompson's group F is most clear. Indeed, notice that the (unlabelled) strand diagrams of the tree diagram group $D_t(x)$ over the semigroup presentation $\langle x \mid x^2 = x \rangle$ are precisely the tree diagrams used by Cannon, Floyd and Parry to represent F, and which are central to their analysis of it, in [14].

2.3.4 Annular diagrams

Recall that Cannon, Floyd and Parry also use a representation for Thompson's group T very similar to that of F - indeed, it is precisely the same but with the addition of cyclic permutations on the ends of each tree diagram describing how they connect. This idea can be generalised to diagrams as a whole, from which we obtain *annular diagrams*. Just like diagrams (which, given this wider context, we might call *planar* diagrams, as suggested in [22]), annular diagrams are isotopy classes of labelled, oriented plane graphs built out of certain atomic pieces.

First, notice that, given a diagram Δ we may obtain a new (isotopy class of) plane graph(s), denoted Δ°, by identifying (via a homotopy) the vertices $i(\Delta)$ and $t(\Delta)$ so that the image of $\lceil \Delta \rceil$ bounds a region of $\mathbb{R}^2$ disjoint from Δ° while the image of $\lfloor \Delta \rfloor$ bounds a region containing Δ°. Notice that the orientation on the images of the top and bottom path must then be counter-clockwise. This new isotopy class is an annular diagram and the map $\Delta \mapsto \Delta^\circ$ embeds the set of diagrams into the set of annular diagrams. We may refer to Δ° as an **annularised planar diagram**.

Definition 2.3.13. Let $\mathcal{P}$ be a semigroup presentation. We define two kinds of diagram:

(a) Given an atomic (planar) diagram $\Psi = \varepsilon_x + \pi_r + \varepsilon_y$ over $\mathcal{P}$, we define the corresponding **atomic annular diagram** to be Ψ°. We refer to $\lceil \Psi \rceil^\circ$ and $\lfloor \Psi \rfloor^\circ$ as the **inner path** and **outer path** of Ψ and denote them by $\mathrm{inn}(\Psi^\circ)$ and $\mathrm{out}(\Psi^\circ)$ respectively. We also consider the vertex which is the image of $i(\Delta)$ and $t(\Delta)$ to be two vertices (which happen to coincide in this case) called the **inner vertex** and **outer vertex**, which we denote by $i(\Psi^\circ)$ and $o(\Psi^\circ)$ and which are the basepoints of $\mathrm{inn}(\Psi^\circ)$ and $\mathrm{out}(\Psi^\circ)$ respectively.

(b) Let $w = a_0 \dots a_{n-1}$ where $a_i \in \Sigma$ and let $0 \leq k \leq n-1$. We define the **atomic cycle diagram** $\mathcal{C} := \mathcal{C}(w, k)$ to be the isotopy class of the plane graph containing a single (up to rotation) non-self-intersecting closed positive path p with counter-clockwise orientation which has two distinguished vertices: $i(\mathcal{C})$, where p has label w if we take this vertex as its basepoint; and $o(\mathcal{C})$, where p has label $w_k := a_k \dots a_{n-1} a_0 \dots a_{k-1}$ if we take this vertex as its basepoint. The vertices $i(\mathcal{C}), o(\mathcal{C})$ are called the **inner** and **outer vertices** respectively, while the path $\mathrm{inn}(\mathcal{C})$ labelled by w is called with **inner path** and the path $\mathrm{out}(\mathcal{C})$ labelled by w_k is called the **outer path**.

We refer to both kinds collectively as **atomic annular diagrams** over $\mathcal{P}$. Naturally, we also have the **trivial annular diagrams** which take the form ε_w° for $w \in \Sigma^+$ with distinguished vertices and paths defined just as above. Notice that $\varepsilon_w^\circ = \mathcal{C}(w, 0)$.

Just as with planar diagrams, we construct annular diagrams in general out of these basic diagrams. We now define a partial operation which extends concatenation on planar diagrams.

Definition 2.3.14. Let Ψ_1, Ψ_2 be atomic annular diagrams such that $\mathrm{out}(\Psi_1)$ and $\mathrm{inn}(\Psi_2)$ have the same label. We define their **concatenation** $\Psi_1 \circ \Psi_2$ as follows: embed a copy of Ψ_1 inside the region bounded by $\mathrm{inn}(\Psi_2)$ and then identify (via a homotopy) the paths $\mathrm{out}(\Psi_1)$ and $\mathrm{inn}(\Psi_2)$ so that $o(\Psi_1)$ is identified with $i(\Psi_2)$. By definition

- $\mathrm{inn}(\Psi_1 \circ \Psi_2) = \mathrm{inn}(\Psi_1)$ and $i(\Psi_1 \circ \Psi_2) = i(\Psi_1)$;
- $\mathrm{out}(\Psi_1 \circ \Psi_2) = \mathrm{out}(\Psi_2)$ and $o(\Psi_1 \circ \Psi_2) = o(\Psi_2)$.

As before, this can be extended to finitely many atomic annular diagrams.

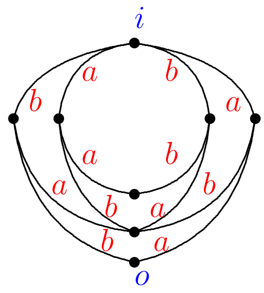

Figure 2.5: A spherical annular diagram with base $aabb$ over the semigroup presentation $\langle a, b \mid ab = ba \rangle$. Notice it is equal to $\Psi^\circ \circ \mathcal{C}(bbaa, 2)$ where Ψ is the planar diagram from Figure 2.3.

Let $\mathcal{P}$ be a semigroup presentation. We say an isotopy class of labelled oriented graphs Δ is an annular diagram over $\mathcal{P}$ if either it is equal to a trivial diagram or it decomposes as a concatenation of atomic annular diagrams over $\mathcal{P}$. Related terminology regarding subdiagrams, cells etc. all translate from planar to annular diagrams.

An example of an annular diagram is shown in Figure 2.5. We denote the set of all (base w) annular diagrams by $\mathcal{A}$ ($\mathcal{A}(w)$).Notice that $\mathcal{C}(w, k_1) \circ \mathcal{C}(w, k_2) = \mathcal{C}(w, k_1 + k_2 \mod |w|)$ for any word w and thus we may express an annular diagram as an alternating product of annularised planar diagrams and atomic cycle diagrams. It follows that the notion of dipoles, and of reducing/inserting dipoles, in planar diagrams readily translates to annular diagrams and so we define an equivalence relation $\equiv$ on $\mathcal{A}$ in exactly the same way as on $\mathcal{D}$.

Let $\mathcal{P}$ be a semigroup presentation and let w be a word on Σ. Then $D_a(\mathcal{P}, w) := (\mathcal{A}(\mathcal{P}, w)/ \equiv, \circ)$ is a group, known as the **annular diagram group over $\mathcal{P}$ with base w**.

As indicated at the opening of this section, the annular diagram group $D_a(\mathcal{P}, x)$ over $\mathcal{P} = \langle x \mid x^2 = x \rangle$ is isomorphic to Thompson's group T.

2.3.5 Symmetric diagrams

Going a step further, Cannon, Floyd and Parry show that Thompson's group V can be represented using tree diagrams with, in addition, *arbitrary* permutations in place of the cyclic permutations in the representation for T. Naturally, as before, we can generalise this idea in diagrams to obtain what Guba and Sapir in [29] called *braided pictures* from which we can similarly obtain *braided diagram groups*. Following a suggestion of Genevois in [22] we opt to call these **symmetric diagrams** and **symmetric diagram groups**.

While these objects are relevant for discussions that will follow, specifically in discussions of the Farley-Hughes theorem (Section 2.4.1) and our use of it in Chapter 5, we will not concern ourselves with them directly, and so we will not define them formally here. Suffice it to say that given a reduced symmetric diagram Δ over a tree-like presentation $\mathcal{P}$ there exists strand diagrams S_1, S_2 over $\mathcal{P}$ and a label-preserving bijection $\sigma : \lfloor S_1 \rfloor \to \lfloor S_2 \rfloor$ such that

$$\Delta = S_1 \circ \sigma \circ S_2^{-1}$$

where we understand σ as a generalisation of an atomic cycle diagram for an annular diagram (Definition 2.3.13).

2.3.6 Infinite diagrams

So far, we have only considered diagrams with finite base and consisting of finitely many cells. In the work that follows, however, it will be necessary to work with diagrams with infinite base and infinitely many cells, and so we develop the required formalisations here.

In the literature so far, the only formal definition of an infinite diagram is given in [23] where they are used to model the combinatorial boundary of the Farley complex [19] of a given diagram group. We recall the definition here and build on it significantly in Chapter 6.

Let $\mathcal{P}$ be a semigroup presentation. We define a **formal concatenation** to be a length ω sequence $(\Delta_\alpha)_{\alpha<\omega}$ of diagrams over $\mathcal{P}$ such that $\lfloor \Delta_\alpha \rfloor$ and $\lceil \Delta_{\alpha+1} \rceil$ have the same label for all $\alpha < \omega$. If there exists i such that Δ_α is a trivial diagram for all $\alpha > i$ then we say (Δ_α) is finite and otherwise we say it is infinite. We denote the set of all formal concatenations over $\mathcal{P}$ by $C(\mathcal{P})$. Given two formal concatenations $(\Delta_\alpha), (\Psi_\alpha) \in C(\mathcal{P})$ we say (Δ_α) is a **prefix** of (Ψ_α) if for all i there exists some j such that the diagram $\Pi_{\alpha<i}\Delta_\alpha = \Delta_0 \circ \ldots \circ \Delta_{i-1}$ is a prefix of the diagram $\Pi_{\alpha<j}\Psi_\alpha = \Psi_0 \circ \ldots \circ \Psi_{j-1}$.

It is clear that the prefix relation on $C(\mathcal{P})$ defined above is reflexive since we may take $j = i$, and it is transitive since the prefix relation on diagrams is. Thus it is a preorder, and so by Proposition 2.1.1 we can make the following definition.

Definition 2.3.15. Define $\mathcal{D}^\omega(\mathcal{P}) = C(\mathcal{P})/E$ where E is the canonical equivalence relation of the prefix preorder on $C(\mathcal{P})$. We call an element $\Delta \in \mathcal{D}^\omega(\mathcal{P})$ an **infinite diagram** if the formal concatenations it contains are infinite. Given $\Delta \in \mathcal{D}^\omega(\mathcal{P})$ we refer to a formal concatenation $(\Delta_\alpha) \in \Delta$ as an **atomic decomposition** of Δ if Δ_α is an atomic diagram for all $\alpha < \omega$. We say Δ is **reduced** if $\Pi_{\alpha<i}\Delta_\alpha$ is reduced for all i for some (and therefore all) $(\Delta_\alpha) \in \Delta$.

Remark. By Proposition 2.1.1 the prefix preorder on $C(\mathcal{P})$ also induces a partial order P on $\mathcal{D}^\omega(\mathcal{P})$ and we say D_1 is a prefix of D_2 if $(D_1, D_2) \in P$ for $D_1, D_2 \in \mathcal{D}^\omega(\mathcal{P})$.

2.4 Homeomorphisms locally determined by similarity structures

We now take a brief detour into yet another class of groups, a particular subclass of which will turn out to be important for understanding the connection between the two classes of groups considered in the previous two sections. This class of groups was introduced by Hughes in [34].

We begin by recalling the idea of a similarity between metric spaces. Let X, Y be metric spaces with metrics d_X, d_Y respectively and let $f : X \to Y$ be a function. If there exists a positive real number C such that $d_Y(x_1f, x_2f) = Cd_X(x_1, x_2)$ for all $x_1, x_2 \in X$ then we say f is a **similarity**.

Remark. Notice that a similarity is necessarily injective since $d(x, y) \neq 0$ if and only if $x \neq y$ for any metric d. Thus, a surjective similarity is necessarily a bijection.

Recalling that a metric d_X induces a topology, where the basic open sets are the open metric balls $B(x, r)$, on its set X we may consider the group $H(X)$ of all homeomorphisms $h : X \to X$ with respect to this topology. We are interested in a certain type of metric space whose topology has a particularly nice structure.

Let X be metric space with metric d. We say X is an **ultrametric space**, and d is an **ultrametric**, if d satisfies the strong triangle inequality:

$$d(x, z) \leq \max\{d(x, y), d(y, z)\}$$

for all $x, y, z \in X$.

An immediate consequence of this property is that if two open balls intersect $B_1 \cap B_2 \neq \emptyset$ then we either have $B_1 \subseteq B_2$ or $B_2 \subseteq B_1$. If an ultrametric space X is also compact then it follows that each open ball B (which is not an isolated point) can be partitioned into finitely many maximal open proper subballs, while B itself is properly contained in finitely many open balls.

A relevant example for us is as follows. We say a **standard dyadic interval** is a subinterval of $[0, 1]$ of the form $[\frac{a}{2^n}, \frac{a+1}{2^n}]$ for $a, n \in \mathbb{N}$. Let I_2 be the set of standard dyadic intervals; then I_2 is partially ordered by set containment and its Hasse diagram is the infinite binary rooted tree. Given this, we can define an ultrametric d on $[0, 1]$ by stipulating that $d(x, y) = 0$ if $x = y$ and $d(x, y) = \frac{1}{2^n}$ otherwise, where $[\frac{a}{2^n}, \frac{a+1}{2^n}]$ is the largest standard dyadic interval containing both x, y. Then I_2 is precisely the set of balls of the ultrametric space $([0, 1], d)$.

Since balls partition an ultrametric space and similarities map balls to balls the following definition becomes natural. Let $h \in H(X)$, we say h is a **local similarity** if for all $x \in X$ there exists a ball B_1 and a surjective similarity $S : B_1 \to B_2$ such that $x \in B_1$ and $h|_{B_1} = S$. In order to study groups of such homeomorphisms

Hughes introduced the following idea in [34].

Definition 2.4.1. Let X be a compact ultrametric space and let $\mathcal{B}$ be the set of its balls. Let Sim be a set of sets which is indexed by $\mathcal{B} \times \mathcal{B}$ such that $\mathrm{Sim}(B_1, B_2)$ only contains surjective similarities from B_1 to B_2. We call Sim a **finite similarity structure** on X if

(a) for every $B_1, B_2 \in \mathcal{B}$ the set $\mathrm{Sim}(B_1, B_2)$ is finite;

(b) for every $B_1 \in \mathcal{B}$ we have $\mathrm{id}_{B_1} \in \mathrm{Sim}(B_1, B_1)$;

(c) if $f \in \mathrm{Sim}(B_1, B_2)$ and $g \in \mathrm{Sim}(B_2, B_3)$ then $fg \in \mathrm{Sim}(B_1, B_3)$

(d) if $f \in \mathrm{Sim}(B_1, B_2)$ and $B \in \mathcal{B}$ such that $B \subseteq B_1$ then $f|_B \in \mathrm{Sim}(B, Bf)$;

(e) if $f \in \mathrm{Sim}(B_1, B_2)$ then $f^{-1} \in \mathrm{Sim}(B_2, B_1)$.

Notice that Sim induces an equivalence relation E on $\mathcal{B}$ defined by $(B_1, B_2) \in E$ if $\mathrm{Sim}(B_1, B_2) \neq \emptyset$. We refer to an equivalence class $[B]$ of E as a Sim**-class** or **similarity class**.

Remark. Notice that E is an equivalence relation by conditions $(b), (c), (e)$.

Let X be a compact ultrametric space and let Sim be a finite similarity structure. We say $h \in H(X)$ is **locally determined** by Sim if for every $x \in X$ there is a ball $B \in \mathcal{B}$ such that $h|_B = S$ for some $S \in \mathrm{Sim}(B, Bh)$. We denote the group of all homeomorphisms locally determined by Sim by $\Gamma(\mathrm{Sim})$ and refer to it as the **group locally determined by** Sim.

Remark. Notice that conditions (b)-(e) in Definition 2.4.1 are chosen precisely so that $\Gamma(\mathrm{Sim})$ is a group.

An important property of groups determined by similarity structures is the following, from [21, Proposition 3.11].

Proposition 2.4.2. *Let $\gamma \in \Gamma(\mathrm{Sim})$ for a finite similarity structure* Sim *on a compact ultrametric space X with ball set $\mathbb{B}$. There exists two finite partitions $\{B_1, \ldots, B_k\}, \mathcal{C} \subseteq \mathbb{B}$ of X into balls and a bijection $\phi : \{B_1, \ldots, B_k\} \to \mathcal{C}$ such that $\gamma|_{B_i} \in \mathrm{Sim}(B_i, \phi(B_i))$.*

2.4.1 Small similarity structures and the Farley-Hughes theorem

Our interest in groups locally determined by similarity structures lies in the subclass of those determined by *small* structures.

Let Sim be a finite similarity structure on a compact ultrametric space X with ball set $\mathcal{B}$. We say Sim is **small** if $|\operatorname{Sim}(B_1, B_2)| \leq 1$ for every $B_1, B_2 \in \mathcal{B}$. Let B be a ball which is not an isolated point and let $\mathcal{B}_B$ denote its (finite) set of maximal subballs, we call a total order L_B on $\mathcal{B}_B$ a **local ball order on B**. If L_B is defined for every $B \in \mathcal{B}$ then we say the relation $L_{\mathcal{B}} = \bigsqcup_{B \in \mathcal{B}} L_B$ is a **ball order on X**. We say $L_{\mathcal{B}}$ is **compatible** with Sim if any $h \in \operatorname{Sim}(B_1, B_2)$ induces an order isomorphism $\mathcal{B}_{B_1} \to \mathcal{B}_{B_2}$.

Farley and Hughes [21] construct a compatible ball order given a small similarity structure by assigning a local ball order on a representative B for each similarity class and letting this induce local ball orders on every other ball in a given class via the unique similarity between them. It then follows from Definition 2.4.1 that the orders agree between any two balls in a given class.

With this, we may now understand the remarkable Farley-Hughes theorem [21]. In a sentence, it states that the class of groups locally determined by small similarity structures coincides with the class of symmetric diagram groups over tree-like semigroup presentations. For our purposes, we will only require one half of this theorem, which we now state.

Theorem 2.4.3 (Farley-Hughes Theorem [21])**.** *Let X be a compact ultrametric space equipped with a small similarity structure* Sim *and a compatible ball order $L_{\mathcal{B}}$. Let Σ denote the set of* Sim*-classes of X and define*

$$\mathcal{R} = \{([B], [B_1] \dots [B_k]) \mid B \in \mathcal{B} \setminus \mathcal{I},\ B_i \in \mathcal{B}_B,\ (B_i, B_j) \in L_B \text{ for } i < j\}$$

where $\mathcal{I} = \{\{x\} \mid x \text{ is an isolated point}\}$. Then

$$\Gamma(\operatorname{Sim}) \cong D_s(\mathcal{P}_{\operatorname{Sim}}, [X])$$

where $\mathcal{P}_{\operatorname{Sim}} = \langle \Sigma \mid \mathcal{R} \rangle$.

Remark. Essentially, one ought to think of the balls of X as being labelled by their similarity classes.

We will make use of the isomorphism which witnesses this theorem, as defined in its proof in [21], which we now define for reference.

Definition 2.4.4. Let $\mathcal{C} = \{B_1, \dots, B_k\} \subseteq \mathbb{B}$ be a partition of a compact ultrametric space X with ball set $\mathbb{B}$ equipped with a small similarity structure Sim. We define $S_{\mathcal{C}}$ to be the positive strand diagram over $\mathcal{P}_{\operatorname{Sim}}$, as defined in Theorem 2.4.3, with top boundary vertex labelled by $[X]$ which contains precisely the positive paths p_i labelled by $[B_i^0], \dots, [B_i^{m_i}]$ where $X = B_i^0, \dots, B_i^{m_i} = B_i$ is the sequence on $\mathbb{B}$ such that B_i^{j+1} is a maximal proper subball of B_i^j.

Now, let $\gamma \in \Gamma(\mathrm{Sim})$ and let $\mathcal{C}_1, \mathcal{C}_2, \phi : \mathcal{C}_1 \to \mathcal{C}_2$ be partitions and a bijection between them as defined with respect to γ in Proposition 2.4.2. We define the isomorphism $\Omega : \Gamma(\mathrm{Sim}) \to D_s(\mathcal{P}_{\mathrm{Sim}}, [X])$ by

$$\Omega(\gamma) = S_{\mathcal{C}_1} \circ \sigma_\phi \circ S_{\mathcal{C}_2}^{-1}$$

understanding this image as representing its equivalence class of strand diagrams, where $\sigma_\phi : \lfloor S_{\mathcal{C}_1} \rfloor \to \lfloor S_{\mathcal{C}_2} \rfloor$ is the label-preserving bijection induced by ϕ.

Chapter 3

Trees & Cycles

The following is adapted from joint work between the author and Peter J. Cameron [13]. Upon submitting our paper for review, we learned that a number of the results were already known; this chapter represents what could be salvaged. Sections 3.1 and 3.2 are, more or less, as they appear in [13] while Sections 3.3 and 3.4 are derived from the remainder of the paper, being improved by taking the existing literature into account.

This chapter is concerned with the *Dénes cycles* of a tree T, which arise as products of the edges of T where we understand edges as transpositions of the vertices which they connect. In particular, we were interested in determining how many distinct products of the edges give rise to a given Dénes cycle, which we call the *multiplicity* of the cycle. Our primary result is described in Theorem 3.3.2, which places the Dénes cycles in one-to-one correspondence with certain partial orders on the edge set of T such that the set of distinct products of edges giving a Dénes cycle is in one-to-one correspondence with the total extensions of its corresponding partial order. This allows us to show that there exists an algorithm which determines the multiplicity of a given Dénes cycle (Theorem 3.4.2). In addition, we obtain a significantly simplified proof (Theorem 3.4.1) of the Eden-Schützenberger Theorem, originally proved in their paper [16].

3.1 Dénes cycles

Let T be a tree on the vertex set $\{1, \ldots, n\}$, with edge set $E(T)$. We regard the edge $e \in E(T)$ which joins vertices i and j as the transposition (i, j) in the symmetric group S_n. These transpositions generate the symmetric group, and form a *minimal generating set* for S_n, in the sense that no proper subset is a generating set.

Whiston [40] showed that the largest size of a minimal generating set for S_n is $n-1$, and Cameron and Cara [12] showed that, for $n \geq 7$, there are just two types of minimal generating sets of size $n-1$, both derived from trees: one consists of the set

$E(T)$ as above, while the other has the form $\{\{s\} \cup \{(st)^\epsilon : t \in E(T) \setminus \{s\}, \epsilon = \pm 1\}$, for $s \in E(T)$, where T is an arbitrary tree. The diameter of the Cayley graph of S_n with generating set $E(T)$ has been investigated by Kraft [37], and other properties of these graphs by Konstantinova and co-authors (see for example [27]).

On the other hand, the product of the transpositions corresponding to all edges of T is an n-cycle (this follows easily by induction from the fact that, if $g \in S_n$ has the property that i and j lie in different cycles, then these two cycles are 'fused' into a single cycle in the product $g(i,j)$). This correspondence was originally studied by Dénes [15] and investigated further, most notably, by Eden & Schützenberger [16] and Moszkowski [38].

According to Cayley's Theorem, the number of trees on $\{1, \ldots, n\}$ is n^{n-2}, while the number of orderings of the edges of such a tree is $(n-1)!$. On the other hand, the number of n-cycles in S_n is $(n-1)!$, and all cycles are conjugate, so each cycle can be realised in n^{n-2} ways as the product of all the edges in a tree.

One might guess that, given any tree, every cycle arises uniquely from a product of its edges; but a little thought shows that this is not so. Indeed, the only trees with this property are the stars. If a cycle arises as such a product we say it is a **Dénes cycle**. This raises the question - for a given tree, precisely how many distinct products of edges produce a given Dénes cycle?

To illustrate, Table 3.1 shows the frequencies of cycles arising from the six non-isomorphic trees on 6 vertices. (The entry (x, y) means that y cycles have frequency x.) This and other computations were performed using GAP [28]. The column labelled "Diameter" gives the diameter of the Cayley graph of S_6 with connection set $E(T)$. These values were calculated using the GAP package GRAPE [39].

Tree	Cycles	Diameter	Cycle frequencies
	16	15	$(1,2), (4,4), (6,2), (9,4), (11,2), (16,2)$
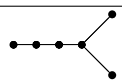	24	11	$(1,4), (3,4), (4,4), (6,4), (7,4), (9,4)$
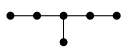	24	11	$(1,2), (2,4), (3,4), (4,4), (7,4), (8,4), (11,2)$
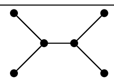	36	10	$(1,8), (3,16), (4,4), (6,8)$
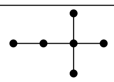	48	9	$(1,12), (2,12), (3,12), (4,12)$
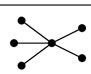	120	7	$(1,120)$

Table 3.1: Cycles from 6-vertex trees

The primary contribution of this chapter is to establish a correspondence between the Dénes cycles of a given tree and certain partial orders of its edges. This correspondence is strong enough to answer the above question as well as recover some classical results on Dénes cycles.

Let $O(T)$ denote the set of all orderings of the edges of a tree T, and let $C(T)$ denote the set of cycles arising from orderings of the edges of T; for $c \in C(T)$, the *multiplicity* of c is the number of orderings of the edges of T for which the product of the transpositions is c. We define the *evaluation map* to be the surjection $\mathcal{E} : O(T) \to C(T)$ which maps each product to the cycle which it is equal to in S_n.

3.2 Geometric correspondence

First notice that the set-up of our problem is naturally reached via a geometric, or dynamical, interpretation. We consider the edges of T as transpositions on $\{1, \ldots, n\}$: the edge $e = \{i, j\}$ corresponds to the transposition (i, j). Set-theoretically, the edge is just a pair of elements while the transposition is a function which sends the element i to j and vice versa, fixing everything else. This may seem like a bit of a leap, but this gap can be bridged by appealing to the geometry of T. If we imagine a person standing at the vertex i, they could cross the edge e and arrive at the vertex j; similarly they can cross in the opposite direction to reach i from j. Of course, if this person were at any other vertex v they cannot cross e; any attempt to do so would fail and they would remain at v.

This interpretation can naturally be extended to sequences of edges (correspondingly, products of transpositions). To take a basic example, suppose we have $e_1 = \{i, j\}$ and $e_2 = \{j, k\}$; then their corresponding transpositions are $(i, j), (j, k)$ and so the sequence $e_1 e_2$ corresponds to the 3-cycle (i, k, j). As before, imagine a person standing at the vertex i, moving along the tree with respect to the sequence $e_1 e_2$. First they cross e_1, bringing them to the vertex j, and then e_2, arriving at k. Now at k they move with respect to $e_1 e_2$ again; first they attempt to cross e_1 but cannot since e_1 is not incident to k, so they next cross e_2 and arrive at j. Finally, starting at j they cross e_1 to arrive at i since they now cannot cross e_2. Starting at any other vertices they will remain fixed, since neither e_1 nor e_2 are incident to them.

So by appealing to the geometry of a tree T we can arrive at the correspondence between edges and transpositions, since the transposition corresponding to an edge e describes the part that e plays in the dynamics of T. As such we can use the transpositions to formalise and generalise the idea described above.

Definition 3.2.1. Let T be a tree and let $s = e_1 \ldots e_m$ be a ordering of (a subset of) its edges. We define the **kth step of the traversal of T from i with respect**

to s to be the unique path p_i^k from the vertex $i(e_1 \dots e_m)^{k-1}$ to $i(e_1 \dots e_m)^k$ where the edges e_j are identified with their corresponding transpositions.

We then define the **traversal of** T **from** i **with respect to** s to be the concatenation of kth steps of the traversal from i for $k \geq 1$ in increasing order; $p_i = p_i^1 \dots p_i^r$ where r is the smallest number such that $i(e_1 \dots e_m)^r = i$.

Finally, we define the **traversal of** T **with respect to** s to be the set of circuits p_i for each i.

Remark. Notice that p_i is necessarily a circuit and if c_i is the cycle of $e_1 \dots e_m$ containing i when written in disjoint cycle form then p_i corresponds to c_i and $c_i = (i, t(p_i^1), t(p_i^2), \dots, t(p_i^{r-1}))$ which are the vertices hit by the traversal from i, where $t(p)$ denotes the terminal vertex of a path p.

We are concerned with orderings of the edges $O(T)$. These are sequences as above which are maximal without replacement. As stated in the introduction, the products of transpositions corresponding to orderings are n-cycles and so the traversal of each must be a single circuit which hits every vertex. In fact, one can use this correspondence with traversals to obtain an alternative proof that any ordering obtains an n-cycle; it can be shown independently that a traversal with respect to a given ordering of a tree must be a single circuit which lands on every vertex. Related to this is the following lemma.

Lemma 3.2.2. *Let T be a tree with vertex set $\{1, \dots, n\}$ and let $p = e_1^{\varepsilon_1} \dots e_m^{\varepsilon_m}$ be the traversal of T with respect to some ordering from $O(T)$. Then for each $e \in E(T)$ there are precisely two distinct numbers $j_1, j_2 \in \{1, \dots, m\}$ such that $e = e_{j_1}^{\varepsilon_{j_1}}$ and $e^{-1} = e_{j_2}^{\varepsilon_{j_2}}$.*

Proof. We proceed by induction on n. Considering the path of length 1 with vertices i_1, i_2, there is only one edge e and so only one ordering. The first step of the traversal from i_1 is e; the second step starts at i_2 and is also e, returning to i_1. Thus the only traversal for an ordering for this tree is ee, satisfying the claim.

Now assume the inductive hypothesis and suppose we are given an ordering $\sigma \in O(T)$, we find its traversal p step by step. Pick any vertex to start at and call it i; we denote by $e_1, \dots, e_d$ the edges incident to i and assume they are labelled so that they appear in this order in σ. Notice that each e_j leads to a subtree of T which we denote T_j and the T_j are pairwise disjoint. The first step of the traversal is the path between i and $i\sigma$; the first edge incident to i appearing in σ is e_1 and thus $i\sigma$ must be a vertex of T_1. Let $e_{1,1} \dots e_{1,k}$ be the part of the first step of the traversal which is in T_1 (this may be empty) and let v_1 be the vertex such that $e_1 = \{i, v_1\}$ (we may have $v_1 = i\sigma$).

By the inductive hypothesis each $e \in E(T_1)$ appears precisely twice in any traversal of T_1 with respect to an ordering of its edges; removing edges in $E(T) \setminus E(T_1)$

from σ, which we denote $\sigma|_{T_1}$, gives such an ordering and indeed, if $v \in V(T_1)$ then $v\sigma = v\sigma|_{T_1}$ as long as $v\sigma \in V(T_1)$. So we obtain a traversal p_{T_1} where each $e \in E(T_1)$ appears precisely twice, hits each vertex of T_1 and ends at $i\sigma$ (since that is where it started). Notice that $e_{1,1} \dots e_{1,k}$ must be a suffix of p_{T_1} as once $e_{1,1}$ is crossed the rest must immediately follow as the edges appear in that order in σ and we must then land on $i\sigma$ where the traversal ends (we know this because they are a suffix of the first step of the traversal p). Since $e_{1,1}$ is incident to v_1 the edge preceding it in p_{T_1} must be an edge e_{v_1} also incident to v_1 and this edge must be the rightmost edge incident to v_1 left of $e_{1,1}$ in $\sigma|_{T_1}$ or the rightmost edge incident to v_1 in $\sigma|_{T_1}$. Therefore, in σ, it is either the rightmost edge of T_1 left of e_1 incident to v_1 or (if there are no edges incident to v_1 left of e_1) it is the rightmost edge incident to v_1. In either case the final step of p_{T_1} differs from the corresponding step of p; after e_{v_1}, instead of $e_{1,1}$ the traversal p crosses e_1 and immediately e_2 without landing on i. So p and p_{T_1} disagree on the path $e_{1,1} \dots e_{1,k}$, but these edges are still included in p precisely twice so far since they were counted once in p_{T_1} where they do agree, and once immediately before p_{T_1} started. Thus we redefine p_{T_1} by removing the suffix $e_{1,1} \dots e_{1,k}$ and attaching it as a prefix; now p_{T_1} is the part of p containing all edges of T_1 which appear in p and each appear precisely twice. In the case where $v = i\sigma$ we have that p_{T_1} and p agree apart from landing on v_1 the second time, and so no changes need be made to p_{T_1}.

Now after landing on the final vertex of T_1, p lands on a vertex of T_2 and we have precisely the same situation as before. Thus this continues inductively until we reach T_d; in the final step of p we start at a vertex in T_d and cross e_d. Since e_d is the rightmost edge incident to i in σ, p must land on i and the traversal is complete. We now see that $p = e_1 p_{T_1} e_1 e_2 \dots e_{d-1} e_d p_{T_d} e_d$ where each edge of T_j appears precisely twice in p_{T_j}, as required. □

Using similar ideas to the above we can prove the following lemma which, while a simple idea, turns out to be crucial.

Lemma 3.2.3. *Let $\sigma, \tau \in O(T)$. Then σ, τ give the same n-cycle $c \in C(T)$ if and only if they differ only by some sequence of commuting non-adjacent edges.*

Proof. Since non-adjacent edges correspond to disjoint transpositions, it is trivial that if the orderings differ only by a sequence of commuting non-adjacent edges then they give the same cycle. We show that if they differ by a sequence of commutes which involves at least one pair of adjacent edges then they must give different cycles.

Indeed, suppose $\sigma = e_1 \dots e_j e_{j+1} \dots e_{n-1} \in O(T)$ and consider

$$\sigma' = e_1 \dots e_{j+1} e_j \dots e_{n-1}$$

where $e_j, e_{j+1}$are incident to a common vertex denoted i. Listing the edges incident to i by the order they appear in each ordering, we have

$$e_{i_1}, \ldots, e_{i_k}, e_j, e_{j+1}, e_{i_{k+3}}, \ldots, e_{i_d}$$

for σ and

$$e_{i_1}, \ldots, e_{i_k}, e_{j+1}, e_j, e_{i_{k+3}}, \ldots, e_{i_d}$$

for σ'. We denote the subtree that each e_m leads to by T_m and the traversal of T with respect to σ (resp. σ') by p (resp. p'). As illustrated in the proof of Lemma 3.2.2 above, p and p' hits each vertex of T_m in some order then each vertex of T_{m+1} according to the order of their respective listing of edges incident to i. The order in which a traversal of T with respect to an ordering in $O(T)$ hits the vertices fully determines the cycle in $C(T)$ obtained from that ordering. Thus we can see that p hits the vertices in T_j then those in T_{j+1} while p' hits the vertices in T_{j+1} then those in T_j and hence they must correspond to distinct n-cycles. This argument can be applied inductively so we see that a sequence of commutes involving any number of pairs of adjacent edges will result in a necessarily different cycle. □

Having established this basic fact we now proceed with investigating the underlying order theory.

3.3 Order-theoretic correspondence

Recall that, given a tree T, we are interested in products - that is, sequences - of its edges $e_1 \ldots e_n$ which are maximal without replacement. Notice, then, that such a sequence uniquely determines a total order on the set E of edges of T, and vice versa. In other words, we are interested in equivalence classes of total orders on the set E. The question is, what are the equivalence classes?

The answer lies in Lemma 3.2.3. Put in order-theoretic terms, it says that if L is a total order on E then another total order L' is equivalent to L precisely when the induced suborders $L|_{\{e_1,e_2\}} = L'|_{\{e_1,e_2\}}$ are equal for every pair of adjacent edges e_1, e_2. With this in mind we introduce some definitions.

Definition 3.3.1. Let T be a tree and let c be a Dénes cycle of T. Let σ be a total order on its set of edges E giving the cycle c. If v is a vertex of T then we define $\sigma|_v$ to be the induced suborder $\sigma|_{E(v)}$ where $E(v) \subseteq E$ denotes the set of edges incident to v; we call $\sigma|_v$ the **local suborder of σ at v**. Thus we also define the **Dénes partial order** corresponding to c, denoted P_c, to be the transitive closure of the relation

$$\bigcup_{v \in V} \sigma|_v$$

where V is the set of vertices of T. We denote the set of Dénes partial orders by $\mathcal{P}(T)$.

Remark. As discussed above, it is a consequence of Lemma 3.2.3 that the Dénes partial order is independent of the choice of total order which gives the cycle in question.

With these ideas in place we may now establish the primary result of this chapter.

Theorem 3.3.2. *Let T be a tree. Then the map $C(T) \to \mathcal{P}(T)$ defined by $c \mapsto P_c$ is a bijection and satisfies $L(P_c) = \{\sigma \in O(T) \mid \mathcal{E}(\sigma) = c\}$.*

Proof. By Lemma 3.2.3 two products $\sigma, \tau \in O(T)$ satisfy $\mathcal{E}(\sigma) = \mathcal{E}(\tau) = c \in C(T)$ if and only if $\sigma|_{\{e_1,e_2\}} = \tau|_{\{e_1,e_2\}}$ for every pair of adjacent edges e_1, e_2 of T, which is true if and only $\sigma|_v = \tau|_v$ for every vertex v of T. By Definition 3.3.1 this confirms that the map is well-defined and injective. Since $P(T)$ is defined precisely as the partial orders obtained from this process it is clear that the map is surjective.

If $\mathcal{E}(\sigma) = c$ it is clear that σ is a total extension of P_c since P_c is defined as a transitive subset of σ. On the other hand, if l is a total extension of P_c then it must be a product satisfying $\mathcal{E}(l) = d$ for some cycle d, and since l is a total extension we have $P_c|_v = l|_v$ for every v. Thus, by Lemma 3.2.3 we must have $d = c$. □

We will make use of this result in the following section. For now, let us consider what these Dénes partial orders look like.

Proposition 3.3.3. *Let T be a tree. Then the Hasse diagram of $P \in \mathcal{P}(T)$ is a tree.*

Proof. By definition, P is formed as the transitive closure of a union of local suborders at each vertex of T. Finding the Hasse diagram H of P, then, is simply a matter of taking the union of the Hasse diagrams for each local suborder, since transitivity is encoded in the path relation of a Hasse diagram by definition.

The Hasse diagram of a local suborder, being a total order, is a path. Any pair of such suborders may or may not have an element in common; if the corresponding vertices of T are joined by an edge and if they are not, respectively. It follows that, in H, two local suborders corresponding to vertices v, w are in the same connected component if and only if there is a path between v and w in T. For the same reason, there is precisely one path in H between such local suborders if and only if there is precisely one path between v and w in T. Thus, H must be a tree. □

3.4 Consequences

We may now exploit the connection described in Theorem 3.3.2 to answer the question posed at the beginning of this chapter, as well as quickly recovering a classical result.

The classical result in question is known as the Eden-Schützenberger Theorem, proved in [16]. Their argument is somewhat involved and careful, dealing directly with the trees and cycles. The following proof is considerably simpler (given Theorem 3.3.2, of course) and, thus, more transparent.

Theorem 3.4.1 (Eden-Schützenberger Theorem [16]). *Let T be a tree with vertex set V. Then*

$$|C(T)| = \prod_{v \in V} d(v)!$$

where $d(v)$ is the degree of $v \in V$.

Proof. By Theorem 3.3.2 it suffices to determine $|\mathcal{P}(T)|$. By definition, an element of $\mathcal{P}(T)$ is uniquely determined by choosing a local suborder at each vertex of T, since otherwise we simply take the union and transitive closure. At a given vertex v, the local suborder is a total order of $E(v)$ and so there are $d(v)!$ distinct choices. The choice of local suborders are independent between distinct vertices, so we obtain the counting formula simply by taking the product of $d(v)!$ for every $v \in V$. □

We may also use Theorem 3.3.2 to help us determine the multiplicities of the Dénes cycles of a given tree. It tells us that the multiplicity of a cycle c is equal to $|L(P_c)|$ and so determining the multiplicity of a cycle is equivalent to counting the linear extensions of P_c. By Proposition 3.3.3 we know the Hasse diagram of P_c is a tree. This leads us to the following theorem.

Theorem 3.4.2. *Let T be a tree. There exists an algorithm which determines the multiplicity of a given Dénes cycle $c \in C(T)$ in quadratic time.*

Proof. The multiplicity of c is equal to $|L(P_c)|$, and the Hasse diagrams of P_c is a tree. In [1], Atkinson gives an algorithm with complexity $O(n^2)$ for counting the total extensions of a partial order whose Hasse diagram is a tree. □

Chapter 4

Fast groups acting on the interval

From here we begin with the primary focus of this thesis; analysis of *fast bump groups*. The cornerstone of this analysis lies in Theorem 4.2.7 which states that each member of the class of fast bump groups is isomorphic to a member of a particular subclass of diagram groups, which we term *bumpy diagram groups*. Naturally, this opens up a plethora of new techniques in the study of these groups. In particular, having obtained these representations we may apply a theorem of Guba and Sapir (Theorem 2.3.7) to conclude that two particular fast bump groups are isomorphic, thereby answering a question of Matthew Brin first publicly posed at a meeting of Thompson group theorists in Oberwolfach in 2018 [11].

The following chapter contains proofs of each of these claims, as well as discussion where appropriate. This chapter is adapted from the author's paper jointly authored with James Belk [3] which is currently under review. Since submitting, this author became aware of a more general proof which obtains Theorem 4.2.7 which uses the Farley-Hughes theorem. We detail this more general argument in Chapter 5 where some of its wider implications are relevant and, for now, consider a more direct argument; the ideas and technology developed through this argument will be relevant for Chapter 7.

4.1 The diagram representation

Consider a geometrically fast set of bumps B with dynamical diagram D. Then there exists a marking of B which witnesses its sources and destinations being disjoint. Taking the sources and destinations along with the connected subsets of their complement in the support of B (i.e. the gaps between the feet) we obtain a general partition of the support of B.

We can improve on this partition by assuming that B is canonical (Definition 2.2.3). In particular, for any dynamical diagram without isolated bumps there exists a geometrically fast set with that diagram such that there exists a marking

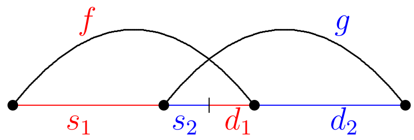

Figure 4.1: The canonical partition for the dynamical diagram generating F.

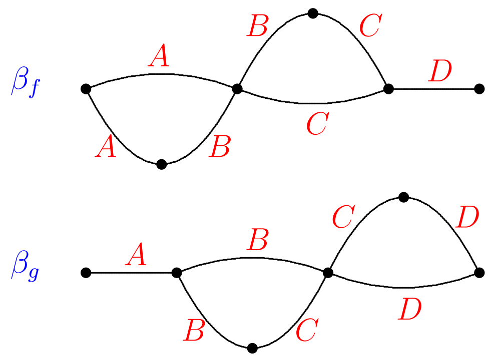

Figure 4.2: The diagrams representing the bumps f, g from Figure 4.1

which defines disjoint feet that cover the support of the set (modulo finitely many isolated points). An example is shown in Figure 4.1. For dynamical diagrams with isolated bumps this is nearly possible—in this case, we just require a 'gap' (i.e. a fundamental domain) between the source and destination of each isolated bump. We may call this the **canonical partition** of a dynamical diagram.

Having obtained such a partition for a dynamical diagram we may use it to define diagrams (in the sense of Guba and Sapir) representing each bump in B. This is best described through a simple example. Consider the partition of the dynamical diagram shown in Figure 4.1. By the definition of this partition we have

$$(s_1)f = s_1 s_2, \ (s_2 d_1)f = d_1$$

and

$$(s_2)g = s_2 d_1, \ (d_1 d_2)g = d_2$$

where in each case restrictions of the maps witness the intervals as homeomorphic. This observation makes it natural to consider the diagrams

$$\beta_f = \pi_{A,AB} + \pi_{BC,C} + \varepsilon_D$$

$$\beta_g = \varepsilon_A + \pi_{B,BC} + \pi_{CD,D}$$

over the semigroup presentation $\mathcal{P} = \langle A, B, C, D \mid A = AB, B = BC, C = BC, D = CD \rangle$ as candidates for representing the bumps f and g. Pictures of these diagrams are shown in Figure 4.2.

We now describe the semigroup presentation and diagrams defined by a dynamical diagram of a geometrically fast set of bumps B in full generality. Suppose B contains n bumps, k of which are isolated, and let $A_1, \dots, A_{2n+k}$ be the canonical partition of its dynamical diagram. Consider $b \in B$ and let $i(b), j(b)$ be such that $\mathrm{src}(b) = A_{i(b)}$ and $\mathrm{dest}(b) = A_{j(b)}$. Then $A_{i(b)}, A_{i(b)+1}, \dots, A_{j(b)-1}, A_{j(b)}$ is the partition of $\mathrm{supt}(b)$ contained in the canonical partition and we define $G(b) = A_{i(b)+1} \dots A_{j(b)-1}$, which is a fundamental domain of b. By definition

$$(A_i(b))b = A_i(b)G(b), (G(b)A_{j(b)})b = A_{j(b)}$$

and so the diagram representing b is

$$\beta_b = \varepsilon_{A_1 \dots A_{i(b)-1}} + \pi_{A_i(b), A_i(b)G(b)} + \pi_{G(b)A_{j(b)}, A_{j(b)}} + \varepsilon_{A_{j(b)+1} \dots A_{2n+k}}$$

over the semigroup presentation $\mathcal{P} = \langle A_1, \dots, A_{2n+k} \mid A_{i(b)} = A_{i(b)}G(b), G(b)A_{j(b)} = A_{j(b)}$ for each $b \in B\rangle$. We may refer to the β_b as **generator diagrams**.

Definition 4.1.1. Let B be a fast set of n bumps, k of which are isolated, and let $A_1 \dots, A_{2n+k}$ be the canonical partition of D_B. For $b \in B$ define $i(b)$ and $j(b)$ to be the integers such that $\mathrm{src}(b) = A_{i(b)}$ and $\mathrm{dest}(b) = A_{j(b)}$. We define

$$\mathcal{P} = \langle A_1, \dots, A_{2n+k} \mid A_{i(b)} = A_{i(b)}G(b), G(b)A_{j(b)} = A_{j(b)} \text{ for each } b \in B\rangle$$

where $G(b) = A_{i(b)+1} \dots A_{j(b)-1}$. We say $\mathcal{P}$ is a **bumpy presentation**, and call $D := D(\mathcal{P}, A_1 \dots A_{2n+k})$ a **bumpy diagram group**. We also define $\delta : B \to D$ by $b \mapsto \beta_b$ where

$$\beta_b = \epsilon_{A_1 \dots A_{i(b)-1}} + \pi_{A_i(b), A_i(b)G(b)} + \pi_{G(b)A_{j(b)}, A_{j(b)}} + \epsilon_{A_{j(b)+1} \dots A_{2n+k}}$$

referring to δ as the **representation** of B, and to β_b as **generator diagrams**.

We claim that $\{\beta_b \mid b \in B\}$ is a generating set for the bumpy diagram group D of B, and D is isomorphic to the geometrically fast group $\langle B\rangle$.

4.2 The isomorphism

Let B be a finite geometrically fast set of bumps and let $a_1, \dots, a_n$ and $\mathcal{P}$ be the canonical partition and bumpy presentation obtained from B from Definition 4.1.1. We prove that the representation $\delta : B \to D(\mathcal{P}, a_1 \dots a_n)$ extends to an isomorphism between $\langle B\rangle$ and $D(\mathcal{P}, a_1 \dots a_n)$, which we also denote δ.

For the purposes of the following proposition we add an additional set of labels to our strand diagrams. Notice from how $\mathcal{P}$ is constructed there is a two-to-one correspondence between its relations and the set of bumps B. Thus, we will label each interior vertex of a strand diagram S with the bump corresponding to the relation it is labelled by (recalling that cells of diagrams, and therefore interior vertices of strand diagrams, are labelled by relations). We denote by $BP(S)$ the set of bump vertex labels of maximal paths in S; that is, the paths starting at a top boundary vertex and ending at a bottom boundary vertex.

Proposition 4.2.1. *Let $w \in B^+$ where B is a fast set of bumps and let S be the reduced strand diagram of the image $w\delta$. There is an interval partition $\mathscr{P}_1, \ldots, \mathscr{P}_k$ of the support of B in left-to-right, one-to-one correspondence with the set of maximal paths $p_1, \ldots, p_k$ of S such that given any $x \in \mathscr{P}_i$ the simply local reduction w_x is the bump vertex label of p_i. Furthermore, if p_i is a path with boundary labels a_1 and a_2 respectively then $\mathscr{P}_i \subseteq a_1$ and $(\mathscr{P}_i)w \subseteq a_2$.*

Proof. We proceed by induction on the length of w. Let $b \in B$ be supported on $A = a_i \ldots a_j$ and set $\beta = b\delta$ and, without loss, suppose $S \circ \beta$ does not contain dipoles. If a maximal path on S ends at a boundary vertex labelled $a \notin \{a_i, \ldots, a_j\}$ then it is preserved in $S \circ \beta$ and, similarly, if we take a point x such that $xw \notin A$ then $(wb)_x = w_x$.

Suppose p is a maximal path on S ending at a vertex labelled by $a \in \{a_i, \ldots, a_j\}$ and let $\mathscr{P}$ be its corresponding part. In $S \circ \beta$, consider those maximal paths which contain p. If there is only one, call it q, its final interior vertex must be a merge and we take $\mathscr{P}$ as corresponding to q. Since the bump vertex label of p is equal to the simply local reduction w_x of w for any $x \in \mathscr{P}$ and $(\mathscr{P})w \subseteq a$ we see that $(wb)_x = w_x b$ is the bump vertex label of q for any $x \in \mathscr{P}$. Further, the final vertex of q being a merge corresponds to the action of b sending the interval a into the destination of b, which labels the bottom boundary vertex of q.

If p is contained in d distinct maximal paths $q_1, \ldots, q_d$ of $S \circ \beta$ then they must share a final interior vertex and it must be a split. First notice that each q_l must only contain splits since otherwise somewhere a merge would immediately precede a split but it follows from $\mathcal{P}$ being tree-like that it would form a dipole in $S \circ \beta$.

Claim: If p is a path in S from a_1 to a_2 consisting only of splits then its corresponding part $\mathscr{P}$ satisfies $\mathscr{P}w = a_2$.

Proof: By the inductive hypothesis we have that $\mathscr{P}w \subseteq a_2$. If $\mathscr{P}w \subset a_2$ then there would have to be some other part $\mathscr{P}'$ which is also mapped into a_2 by w—however, since p consists only of splits, any other path must lead to a boundary vertex labelled by a distinct part of the canonical partition and so, by the hypothesis, would not map into a_2. Thus, $\mathscr{P}w = a_2$.

The final vertex of the paths $q_1, \ldots, q_d$ being a split corresponds to the action of b mapping the interval a, which is its source, onto the complement of its destination inside A. By the claim, this means that $\mathscr{P}w$ is mapped onto a subset of the canonical partition such that we may partition $\mathscr{P}$ into intervals $\mathscr{Q}_1, \ldots, \mathscr{Q}_d$ with $\mathscr{Q}_l w = a_l$ where a_l is the label of the boundary vertex of q_l. We take $\mathscr{Q}_l$ to be the part corresponding to q_l and note that if $x \in \mathscr{Q}_l$ we have that w_x is the bump vertex label of p, from which it follows that $w_x b$ is the bump vertex label of q_l and it is clear that $w_x b$ is simply locally reduced with respect to x. □

One useful upshot of this proposition is the following corollary. Recall the definition of local reduction and related ideas from Definition 2.2.4. In particular, recall that $L(w)$ denotes the set of simply local reductions of a word $w \in B^+$ while $L^\vee(w)$ denotes the set of free reductions of elements of $L(w)$.

Corollary 4.2.2. *Let $w \in B^+$ and let S be a strand diagram of the image $w\delta$. Then*

$$L^\vee(w) = \{u^\vee \mid u \in BP(S)\}.$$

To illustrate, let's recall the fast bump set from Figure 4.1 and consider $w = fg^{-1}f^{-1}$. The bump vertex labelled strand diagram S of $w\delta$ is shown in Figure 4.3. We can thus see that

$$BP(S) = \{ff^{-1}, fg^{-1}f^{-1}, g^{-1}f^{-1}\}$$

and we can also see that $L(w) = \{ff^{-1}, fg^{-1}f^{-1}, g^{-1}f^{-1}\}$ since, letting m be the left endpoint of the support of g, if $x \in (mf^{-1}, mfg)$ we have $x \neq xf \neq xfg^{-1} \neq xfg^{-1}f^{-1}$, whereas $x \in (0, mf^{-1}]$ implies $x \neq xf = xfg^{-1} \neq xff^{-1}$ and $x \in [mfg, 1)$ implies $x = xf \neq xg^{-1} \neq xg^{-1}f^{-1}$. Taking free reductions of the elements of both sets we see that $L^\vee(w) = \{u^\vee \mid u \in BP(S)\}$ as claimed.

If we consider the details of last proof we can be more precise regarding the structure of the partition $\mathscr{P}_1, \ldots, \mathscr{P}_k$ for a given w. Essentially, the partition is constructed by inductively refining the canonical partition—following a path starting at the top boundary of its strand diagram, each time we cross a split labelled by $b \in B$ we partition the current part $\mathscr{P}$ in question according to which part of the canonical partition the action of b sends each point to; that is, we break $\mathscr{P}$ into disjoint intervals where the breakpoints are the images of the breaks in the canonical partition under b^{-1}. Since the canonical partition is defined by markers and images of markers we can then see that the partition $\mathscr{P}_1, \ldots, \mathscr{P}_k$ is defined by a finite subset of the set of orbits of markers $M\langle B\rangle$ where M is the canonical marking. We may refer to such a partition as the **canonical refinement** for the word w.

Corollary 4.2.3. *Let $w \in B^*$ where B is a geometrically fast set of bumps and let*

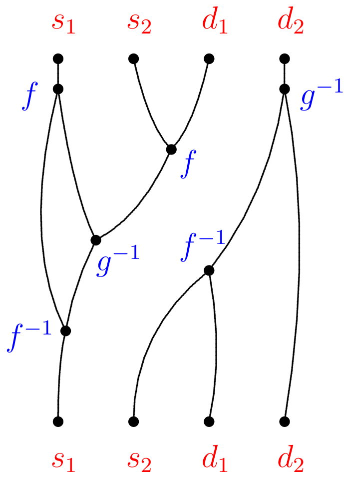

Figure 4.3: The bump vertex labelled strand diagram of $w = fg^{-1}f^{-1}$ where $\{f, g\}$ is the fast bump set from Figure 4.1.

M be a marking which witnesses that B is fast. Then

$$L(w) = \{w_x \mid x \in M\langle B\rangle\}.$$

Proof. The set of simply local reductions is independent of choice of marking so we may choose M to be canonical. Consider the canonical refinement $\mathscr{R}$ of w and its reduced strand diagram S. The refinement is defined by some finite subset $R \subset M\langle B\rangle$ and by Proposition 4.2.1 there is a surjection $\mathscr{R} \to L(w)$ such that if $x \in \mathscr{P} \in \mathscr{R}$ then $\mathscr{P} \mapsto w_x$. If one of the parts $\mathscr{P}$ does not contain its boundary then we can see from how canonical refinements are constructed that we may further subdivide this part via a split so that the breakpoints, contained in $\mathscr{P}$, are elements of $M\langle B\rangle$. $\square$

Proposition 4.2.4. *The map $\delta : \langle B\rangle \to D(\mathcal{P}, a_1 \ldots a_n)$ is a homomorphism.*

Proof. Let w be a word on B such that $w \equiv 1$ and consider the image $w\delta$. We want to show that the reduced strand diagram S of $w\delta$ is trivial. By Corollary 4.2.2 it suffices to show that $L^\vee(w) = \{1\}$ and, by Corollary 4.2.3, it then suffices to show $w_x^\vee = 1$ for $x \in M\langle B\rangle$.

Consider $x \in M\langle B\rangle$. Since markers have trivial history x must have finite history. By [7, Lemma 5.7] if y is a point in the orbit of x then there must be precisely one word u locally reduced at x such that $xu = y$. By assumption we have $xw_x^\vee = xw = x$ and we therefore conclude that $w_x^\vee = 1$. $\square$

The following lemma, which will be useful for the proof of δ being surjective, describes in detail the possible labellings of the edges and cells incident to a given vertex in a diagram over $\mathcal{P}$. Recall the notation defined in the paragraph following Definition 2.3.4. For such a vertex v notice that the order on the sets $I^\Pi(v)$ and

$O^\Pi(v)$ must consist of a (possibly empty) sequence of positive cells followed by a (possibly empty) sequence of negative cells. As such there exists an edge $e \in I(v)$, and similarly for $O(v)$, which marks the point where the sequence $I^\Pi(v)$ changes from positive to negative, and we will refer to e as the **inflection edge** of $I^\Pi(v)$.

Lemma 4.2.5. *Let Δ be a reduced $(a_1 \dots a_n, a_1 \dots a_n)$-diagram over $\mathcal{P}$ and consider $v \neq i(\Delta), t(\Delta)$ one of its vertices of degrees at least three. Then there exists k such that a_{k-1} is the label of the first edge in $I(v)$ and a_k is the label of the first edge in $O(v)$. Further,*

(a) *if a_k is neither a source nor destination then a_{k-1} must be a source and the sequence of labels on $I(v)$ is a_{k-1}, a_k while the sequence of labels on $O(v)$ is a_k, a_{k+1} and the singletons $I^\Pi(v), O^\Pi(v)$ are positive and negative respectively;*

(b) *if a_{k-1} is neither a source nor destination then a_k must be a destination and the sequence of labels on $O(v)$ is a_k, a_{k-1} while the sequence of labels on $I(v)$ is a_{k-1}, a_{k-2} and the singletons $I^\Pi(v), O^\Pi(v)$ are negative and positive respectively;*

(c) *if a_{k-1} is a destination and a_k is a source then every $e \in I(v)$ is labelled a_{k-1} while every $e \in O(v)$ is labelled a_k and for both $I^\Pi(v)$ and $O^\Pi(v)$ the inflection edge occurs at either the first edge or the last;*

(d) *if a_{k-1} is a source then let $\mathcal{T}_m$ be the maximal stretched transition such that $a_{k-1} \in S(\mathcal{T}_m)$. Then there exists two stretched transition chains $\mathcal{T}_1, \mathcal{T}_2$ such that $\mathcal{T}_1$ starts at a_{k-1} and the sequence of labels on $I(v)$ in order is $\overrightarrow{S(\mathcal{T}_1)}$, followed by some number $m_1 \geq 0$ of a_{m-1} if $a_m = d(\mathcal{T}_1) = d(\mathcal{T}_m)$, where the first or the last may be the inflection edge, and then $\overleftarrow{S(\mathcal{T}_2)}$ where $\mathcal{T}_2$ ends at $d(\mathcal{T}_1)$, the first of which is the inflection edge if it has not yet occurred, while*

- *if $\mathcal{T}_2 \subseteq \mathcal{T}_1$ the sequence of labels on $O(v)$ is some number $m_2 \geq 1$ of a_k (if a_k is a destination then $m_2 = 1$), either the first or the last one being the inflection edge, followed by $\overrightarrow{D(\mathcal{T}_1 \setminus \mathcal{T}_2)}$ but;*
- *if $\mathcal{T}_1 \subseteq \mathcal{T}_2$ then $O(v)$ is some number $m_2 \geq 1$ of a_k (if a_k is a destination then $m_2 = 1$) followed by $\overleftarrow{D(\mathcal{T}_2 \setminus \mathcal{T}_1)}$ where the last may be the inflection edge and, finally, some number $m_3 \geq 0$ of a_{m+1} if $a_m = s(\mathcal{T}_2) = s(\mathcal{T}_m)$, the first or the last being the inflection edge if it has not yet occurred;*

(e) *if a_k is a destination then let $\mathcal{T}_m$ be the stretched transition chain such that $a_k \in D(\mathcal{T}_m)$. Then there exists two stretched transition chains $\mathcal{T}_1, \mathcal{T}_2$ such that $\mathcal{T}_1$ ends at a_k and the sequence of labels on $O(v)$ in order is $\overleftarrow{D(\mathcal{T}_1)}$, followed by some number $m_1 \geq 0$ of a_{m+1} if $a_m = s(\mathcal{T}_1) = s(\mathcal{T}_m)$, where the first or the* *last may be the inflection edge, and then* $\overrightarrow{D(\mathcal{T}_2)}$ *where* $\mathcal{T}_2$ *begins at* $s(\mathcal{T}_1)$*, the first of which is the inflection edge if it has not yet occurred, while*

- *if* $\mathcal{T}_2 \subseteq \mathcal{T}_1$ *the sequence of labels on* $I(v)$ *is some number* $m_2 \geq 1$ *of* a_{k-1} *(if* a_{k-1} *is a source then* $m_2 = 1$*), either the first or the last one being the inflection edge, followed by* $\overleftarrow{S(\mathcal{T}_1 \setminus \mathcal{T}_2)}$ *but;*
- *if* $\mathcal{T}_1 \subseteq \mathcal{T}_2$ *then* $I(v)$ *is some number* $m_2 \geq 1$ *of* a_{k-1} *(if* a_{k-1} *is a source then* $m_2 = 1$*) followed by* $\overrightarrow{S(\mathcal{T}_2 \setminus \mathcal{T}_1)}$ *where the last may be the inflection edge and, finally, some number* $m_3 \geq 0$ *of* a_{m-1} *if* $a_m = d(\mathcal{T}_2) = d(\mathcal{T}_m)$*, the first or the last being the inflection edge if it has not yet occurred.*

Proof. Recalling that, since $\mathcal{P}$ is tree-like, $\Delta = \Delta_1 \circ \Delta_2$ for Δ_1 consisting only of positive cells and Δ_2 consisting only of negative, notice that every vertex of Δ is contained in Δ_1. So, consider a vertex v of Δ_1 of degree at least three which is neither the initial nor terminal vertex. Then v must be either be initial or terminal for some positive cell and since the top path has label $a_1 \dots a_n$ it follows by a recursive argument that there exists k such that the first edge in $I(v)$ is labelled a_{k-1} and the first edge in $O(v)$ is labelled a_k. A fact we will use implicitly throughout this proof is that for a given a_i there is at most one relation of the form $(a_h, a_h \dots a_i)$ and at most one of the form $(a_j, a_i \dots a_j)$. For each case that follows we consider what diagram could be built beginning from a vertex with one incoming edge and one outgoing edge.

(a) Since a_k is not a foot it must be in the support of an isolated bump and so the only relations it appears in which could add an edge incident to v are $(a_{k-1}, a_{k-1}a_k)$ and $(a_{k+1}, a_k a_{k+1})$, and we achieve this by adding a positive cell along a_{k-1} and a negative cell along $a_k a_{k+1}$. Any way of adding further edges incident to v will produce dipoles.

(b) Symmetric to (2).

(c) Since a_{k-1} is a destination, the relation here has the form $(a_{k-1}, a_j \dots a_{k-1})$ for some j and so we may attach a positive or negative cell of this form, which introduces a second incoming edge labelled a_{k-1}, and now we may repeat this to add an arbitrary number of such edges. Similarly, we can add an arbitrary number of outgoing edges labelled a_k. In order to add edges with a different label incident to v, we would require a relation of the form $(a_i, a_i \dots a_{k-1})$ or $(a_i, a_k \dots a_i)$ for some i, but the first implies that a_k is a destination and the second implies that a_{k-1} is a source.

(d) If positive cells are to be attached so as to add edges incident to v they must come before any negative cells. Since a_{k-1} is a source we have a relation

$(a_{k-1}, a_{k-1} \dots a_j)$ where a_j is the successor of a_{k-1} in $\overrightarrow{S(\mathcal{T}_m)}$; if a_{k-1} is maximal then $a_j = a_{m-1}$ is a destination, where $a_m = d(\mathcal{T}_m)$. As such we may attach a positive cell at a_{k-1} and then at a_j and so on, as far through $\overrightarrow{S(\mathcal{T}_m)}$ as we like. If we reach the end then a_{m-1} is the label of the last edge added to $I(v)$ and it is then possible to add arbitrarily many edges labelled a_{m-1} to $I(v)$ by attaching either that many positive cells or that many negative cells at a_{m-1} (it cannot be a combination of positive and negative since they would cancel as dipoles). Whether we add those cells or not, we may then attach a sequence negative cells to add incoming edges to v labelled by an interval of $\overleftarrow{S(\mathcal{T}_m)}$, depending where on $\mathcal{T}_m$ we started and where we stop (if we did not reach the end of $\mathcal{T}_m$ or did not add edges labelled a_{m-1} to $I(v)$ this is still possible since there may be a cell added along the bottom path of any positive cell with top path from $S(\mathcal{T}_m)$ which does not add an edge to $I(v)$ whence attaching the inverse of such a positive cell would not form a dipole).

In order for this diagram to be a $(a_1 \dots a_n, a_1 \dots a_n)$-diagram the last edge of $O(v)$ must be labelled a_{l+1} if the last edge of $I(v)$ is labelled a_l - if this pair of edges doesn't form part of the bottom path of the diagram then there must be a cell that contains it as a subpath of its top path. We look at each possibility for a_l mentioned in the previous paragraph and consider the cells which must therefore be added. First notice that if a_k is a source we have a relation of the form $(a_k, a_k \dots a_j)$ for some j we may add arbitrarily many positive cells or negative cells of this form. Suppose that $a_l \in S(\mathcal{T}_m)$ and this edge is in the bottom path of a positive cell, then this cell is of the form $(a_s, a_s \dots a_l)$ where a_s is the predecessor of a_m in $\overrightarrow{S(\mathcal{T}_m)}$ and we therefore have a relation of the form $(a_{l+1}, a_{s+1} \dots a_{l+1})$ - indeed, we have the relation $(a_{s+1}, a_{t+1} \dots a_{s+1})$ wherever a_t is the predecessor of a_s in $\overrightarrow{S(\mathcal{T}_m)}$. If the edge a_{l+1} was in the bottom path of a positive cell then it must have the form $(a_j, a_{m+1} \dots a_j)$ for some $j > k$ and a_j would be the label of the predecessor in $O(v)$ and must also be in the bottom path of a positive cell of the form $(a_{j'}, a_j \dots a_{j'})$ for $j' > j > k$. Continuing this argument indefinitely we can see that this is not possible since the first element of $O(v)$ is a_k and, so, the edge labelled a_{l+1} must be the bottom path of a negative cell and it must have the form $(a_{l+1}, a_{s+1} \dots a_{l+1})$. We can now apply the same argument to a_{s+1} to see it must be the bottom path of a negative cell of the form $(a_{s+1}, a_{t+1} \dots a_{s+1})$ and so on until we reach the edge labelled a_k. Noticing that $a_{s+1} \in D(\mathcal{T}_m)$ when $a_s \in S(\mathcal{T}_m)$ and $a_s \neq s(\mathcal{T}_m)$ completes the proof for this case. If $a_l = a_{m-1}$ where $a_m = d(\mathcal{T}_m)$ then $a_{l+1} = d(\mathcal{T}_m)$ and the argument follows the same way. If $a_l \in S(\mathcal{T}_m)$ and this edge is the bottom path of a negative cell such that $m \geq k - 1$ then the argument is the same; otherwise (i.e. if $m < k - 1$) a symmetric argument shows that we must add a sequence of positive cells which add edges to $O(v)$ labelled by an interval from $\overleftarrow{D}(\mathcal{T}_m)$ in order to get the final edge to have label a_{l+1}. Finally, if $a_l = a_m = s(\mathcal{T}_m)$ then a_{l+1} is a source and we may attach arbitrarily many cells of the form $(a_{l+1}, a_{l+1} \dots a_j)$ for some j (either all positive or all negative).

(e) Symmetric to (4)

□

Proposition 4.2.6. *The map $\delta : \langle B \rangle \to D(\mathcal{P}, a_1 \dots a_n)$ is an isomorphism.*

Proof. Let $g = b_{i_1}^{\epsilon_1} \dots b_{i_m}^{\epsilon_m} \in \langle B \rangle \setminus \{1\}$ for $\epsilon_i = \pm 1$ be a freely reduced word. Then consider $\Delta = g\delta = \beta_{i_1}^{\epsilon_1} \circ \dots \circ \beta_{i_m}^{\epsilon_m}$ and choose one of its dipoles. This dipole consists of two cells π_j, π_k such that π_j is a cell of a factor $\beta_{i_j}^{\epsilon_j}$ and π_k is a cell of a factor $\beta_{i_k}^{\epsilon_k}$ for some $j < k$. Since each relation of $\mathcal{P}$ occurs in precisely one of the diagrams β_i we must have $\beta_{i_j} = \beta_{i_k}$ and $\epsilon_j = -\epsilon_k$.

Our aim is to reduce Δ and see that it is non-trivial. Consider a factor $\beta_{i_k}^{\epsilon_k}$ in the unreduced product above. If either of its cells remain after reducing then we're done, so suppose both are reduced. Notice that if the two dipoles are both formed with the cells of a single factor equal to $\beta_{i_k}^{-\epsilon_k}$ then there is a word equivalent to g containing $b_{i_j}^{\epsilon_j} b_{i_k}^{\epsilon_k}$ as a subword and this forms a cancellable pair. Thus there must exist a factor where this does not happen; otherwise, g would be equivalent to the empty word. So let $\beta_{i_k}^{\epsilon_k}$ be such a factor; there must be two distinct factors equal to $\beta_{i_k}^{-\epsilon_k}$ which each reduce with one of the cells of $\beta_{i_k}^{\epsilon_k}$ and we are thus left with a cell from each factor $\beta_{i_k}^{-\epsilon_k}$. These remaining cells can't reduce with each other since they are labelled by different relations. Since the factor we chose with this property was arbitrary we can see there must be at two least cells remaining once Δ is reduced. Thus $\ker \delta$ is trivial and δ is injective.

Now, let Δ be a reduced $(a_1 \dots a_n, a_1 \dots a_n)$-diagram over $\mathcal{P}$ - we want to show that a diagram equivalent to Δ can be decomposed as a product of generator diagrams. Consider the cells $\pi_1, \dots, \pi_m$ of Δ such that $\lfloor \pi_l \rfloor$ is a subpath of $\lfloor \Delta \rfloor$ in ascending order.

First notice that, for a given π_l, if there does not exist a cell π such that π_l and π form a generator diagram in Δ then π_l must be negative. To see this, suppose some π_l is a positive cell of the form $(a_i, a_i \dots a_{j-1})$ without loss and consider its terminal vertex v. We can then see from Lemma 4.2.5 that v must be the initial vertex of a negative cell of the form $(a_{i+1} \dots a_j, a_j)$ with its bottom path being a subpath of $\lfloor \Delta \rfloor$, which means this cell must be π_{l+1} and thus π_l and π_{l+1} form a generator diagram β such that $\Delta = \Delta' \circ \beta$.

Now suppose π_l is a cell for which there does not exist a cell π such that π_l and π form a generator diagram in Δ; as noted, it must be negative. Suppose without loss that it has the form $(a_i \dots a_{j-1}, a_i)$ consider its terminal vertex v. By Lemma 4.2.5 we can see the label of the first edge in $I(v)$ must be a source and the label of the first edge in $O(v)$ must be a destination, and since π_l is negative there, then, must be a positive cell π of the form $(a_i, a_i \dots a_{j-1})$ whose terminal vertex is v. We can now introduce a dipole (π', π'') along the subpath of the bottom path of Δ labelled $a_{i+1}, \dots, a_j$ to obtain an equivalent diagram $\bar{\Delta}$ so that π and π' form a generator diagram while π_l and π'' form a generator diagram β such that $\bar{\Delta} = \bar{\Delta}' \circ \beta$, and notice $\bar{\Delta}'$ is reduced.

Now, as we have seen, if any of the cells $\pi_1, \dots, \pi_m$ are positive then we can decompose $\Delta = \Delta' \circ \beta$ and Δ' has fewer cells than Δ. If they are all negative then by Lemma 4.2.5 at least one of them must not be part of a generator diagram in Δ whence we obtain an equivalent diagram $\bar{\Delta}$ which decomposes $\bar{\Delta} = \bar{\Delta}' \circ \beta$ such that $\bar{\Delta}'$ has the same number of cells as Δ but has fewer cells which do not immediately form a generator diagram. This proves the claim and it follows that δ is surjective. □

We have now proved the following Theorem.

Theorem 4.2.7. *Let $B = \{b_1, \dots, b_n\}$ be a geometrically fast set of bumps, let k be the number of isolated bumps in B and let $A_1, \dots, A_{2n+k}$ be the canonical partition of the support of B. Then the fast bump group $\langle B \rangle$ is isomorphic to the diagram group $D(\mathcal{P}, A_1 \dots A_{2n+k})$ over the presentation*

$$\mathcal{P} = \langle A_1, \dots, A_{2n+k} \mid A_i = A_i A_{i+1} \dots A_{j-1}, A_j = A_{i+1} \dots A_{j-1} A_j \text{ for each } b \in B \rangle$$

where A_i and A_j is the source and destination of each $b \in B$ respectively.

4.3 PF_4 is isomorphic to F_4

We now turn to the $|B| = 4$ case of the isomorphism type question for these groups. Consider a geometrically fast set $B = \{b_1, b_2, b_3, b_4\}$. By Theorem 2.2.2 the group $\langle B \rangle$ is invariant under isomorphism of its dynamical diagram so we may study isomorphism types of geometrically fast bump groups via these objects. Dynamical diagrams are combinatorially nice and easily enumerable; there are $\prod_{i=0}^{n-1}(2i+1) = 105$ distinct dynamical diagrams of $n = 4$ bumps, and so they generate at most 105 distinct isomorphism classes of groups.

4.3.1 Irreducible systems and ideal form

We can, however, do much better than this since many of these diagrams are decomposable into dynamical diagrams with fewer bumps in a way which respects the isomorphism type of the groups they generate. For example, if B_1, B_2 are geometrically fast sets with disjoint support then the dynamical diagram for $B_1 \sqcup B_2$ decomposes naturally into the dynamical diagrams for B_1 and B_2, while $\langle B_1 \sqcup B_2 \rangle \cong \langle B_1 \rangle \times \langle B_2 \rangle$.

At this point it is worth introducing some formalisms which will be useful both here and in Chapter 7.

Definition 4.3.1. Let G be a fast bump group generated by B. Define an oriented graph Γ on B by saying (b_1, b_2) is a positive edge if $x_1 < x_2 < y_1 < y_2$ where $\text{supt}(b_1) = (x_1, y_1)$, $\text{supt}(b_2) = (x_2, y_2)$. We refer to $\mathbf{\Gamma}$ as the **determinant of B** and we say B is **irreducible** if Γ is connected, in which case we also say G is **irreducible**.

Remark. Notice that the positive paths of the determinant Γ are precisely the transition chains of B.

We know that a given dynamical diagram will always generate a single group isomorphism type; in addition, we can see that many distinct dynamical diagrams will generate a single isomorphism type. The idea is to sequentially replace a generator $b \in B$ with its conjugate by some other generator, thereby preserving isomorphism type but changing the supports of the generators, such that fastness is preserved. This lets us see that two different arrangements of bumps in fact generate the same group. As such, we need to understand the sorts of replacements we can do.

To this end, consider the totally ordered set $E_b := \{f_1, \ldots, f_m\}$ of feet of $B \setminus \{b\}$ contained in the support of $b \in B$. We define an equivalence relation S on E_b as the reflexive, transitive closure of the relation $\{(f_i, f_j) \in E_b \mid f_i, f_j \subseteq \text{supt}(c)$ for some $c \in B \setminus \{b\}$ such that $\text{supt}(c) \subset \text{supt}(b)\}$. We may then consider the quotient set E_b/S, noticing that each equivalence class of S forms an interval of E_b.

Lemma 4.3.2. *Let B be a fast set of bumps and suppose $E_b := \{f_1, \ldots, f_m\}$ is the totally ordered set of feet of $B \setminus \{b\}$ contained in the support of $b \in B$. Suppose B' is obtained by a cyclic shift of E_b/S while keeping the remaining feet fixed, then B' generates a group isomorphic to $\langle B \rangle$. Furthermore, if B is irreducible then B' is irreducible.*

Let us consider an example before we proceed to the proof. Figure 4.4 depicts the feet of some fast bump set B contained in $b \in B$ (not including the feet of b itself). We have $E_b = \{f_1, f_2, f_3, f_4\}$ as the ordered set of feet in question. Consider the

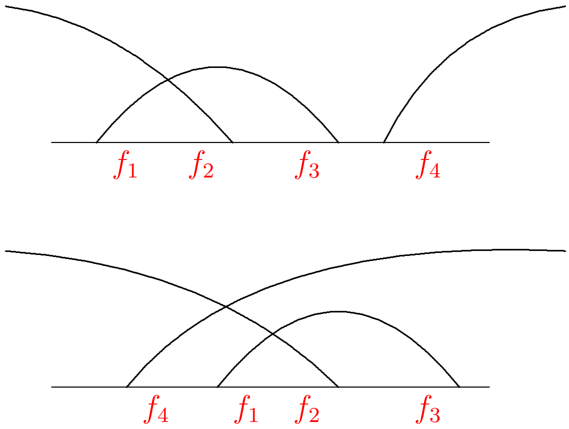

Figure 4.4: The feet of a fast bump system contained in the fundamental domain of a particular bump and a cyclic shift as described in Lemma 4.3.2.

equivalence relation S, we see that f_1, f_2, f_3 are all equivalent since they are all contained in some bump c such that $\operatorname{supt}(c) \subset \operatorname{supt}(b)$. Thus $E_b/S = \{\{f_1, f_2, f_3\}, f_4\}$ is the totally ordered quotient set. In accordance with Lemma 4.3.2, the fast bump set obtained by cyclically shifting E_b/S is shown in Figure 4.4.

Proof. Consider the minimal element of $f \in E_b/B$ and suppose it is a destination. Then f is a foot of some bump $c \in B$ and we replace c with c^b in B to obtain B' and $\langle B \rangle = \langle B' \rangle$ since $B \subseteq \langle B' \rangle$ and vice versa. As is well known, $\operatorname{supt}(c^b) = \operatorname{supt}(c)b$ for any group action and since the source of c is not in the support of b we can see that $\operatorname{supt}(c^b) = (x, yb)$ if $\operatorname{supt}(c) = (x, y)$. The feet f and $\operatorname{src}(b)$ must be disjoint and so $fb \subseteq \operatorname{dest}(b)$ and we take the marker of c^b to be mb if m is the marker of c. To see that B' is fast, then, we must choose a new marker for b. Indeed, we may take y as the new marker of b since f (or any subset of it) is not part of a foot of B' and the left endpoint of $\operatorname{dest}(b)$ is thus yb. Otherwise keeping the same markers from B therefore witnesses that B' is fast and E_b/B' is equal to the 1-left shift of E_b/B'. This argument also shows that the same holds if f is a source since it applies to the converse arrangements of feet starting with $(B')^C$ and reaching B^C.

Now suppose f is an interval of feet and let $c_1, \ldots, c_k$ be the bumps whose feet are contained in f. We replace c_i with c_i^b for all i and, since f is an interval, we may choose markers precisely as before to witness the set B' is fast.

Finally, suppose B is irreducible. This means the oriented graph Γ defined on B with edges $\{b, c\}$ if $\operatorname{supt}(b) \cap \operatorname{supt}(c) \neq \emptyset$ while $\operatorname{supt}(b) \not\subseteq \operatorname{supt}(c)$ and $\operatorname{supt}(c) \not\subseteq \operatorname{supt}(b)$ is connected by definition. It is clear any edges $\{c_1, c_2\}$ of Γ for $c_1, c_2 \in B \cap B'$ are also edges of the graph Γ' of B'. If $\{c_1, c_2\}$ is an edge of Γ for $c_1, c_2 \in B \setminus B'$ then $\{c_1^b, c_2^b\}$ is an edge of Γ' since b is an order isomorphism and thus preserves the arrangement of feet. If $\{c_1, c_2\}$ is an edge of Γ for $c_1 \in B \cap B'$ and $c_2 \in B' \setminus B$ then it is possible that $\{c_1, c_2^b\}$ is not an edge of Γ', but in this case c_1 and c_2 must each have a single foot in $\text{supt}(b)$ and so $\{c_1, b\}$ and $\{c_2^b, b\}$ are both edges of Γ'. Thus Γ' is connected and B' is irreducible. □

Definition 4.3.3. Let B' and B be as in the statement of Lemma 4.3.2. Then we say B' and B are **conjugate equivalent**.

In effect, we are able to consider a group isomorphism type as represented by an equivalence class of dynamical diagrams, the equivalence being the transitive closure of the relation (B_1, B_2) where B_1, B_2 are both fast and $B_2 = (B_1 \setminus \{b\}) \cup \{b^c\}$ for some $b, c \in B_1$. It would be useful to have some canonical representative of this equivalence class. We now define

Definition 4.3.4. Let B be a geometrically fast set of bumps. We say the form of B is **ideal**, or B has **ideal form**, if the following conditions hold:

- every source is contained in $\text{supt}(f)$ and every destination is contained in $\text{supt}(l)$;
- the only destination in $\text{supt}(f)$ is $\text{dest}(f)$ and the only source in $\text{supt}(l)$ is $\text{src}(l)$.

where f is the bump to which the leftmost foot belongs and l is the bump to which the rightmost foot belongs.

The following lemma establishes that every equivalence class of dynamical diagrams contains one of ideal form. Its statement was first communicated to the author by Matthew Brin via private correspondence. The proof is original.

Lemma 4.3.5. *Let G be an irreducible fast bump group. Then there exists a fast set of bumps B with ideal form such that $\langle B\rangle \cong G$.*

Remark. Notice that, in general, there may be multiple fast bump systems with ideal form in a single conjugate equivalence class.

Proof. Let B be an arbitrary (irreducible) fast generating set B of G and let f, l be the first and last bumps of B respectively. We apply Lemma 4.3.2 repeatedly to obtain a fast set B with ideal form. For convenience we relabel B' with B according to the marked isomorphism at each step so that we always use the same labelling.

First, we see that there is a transition chain from f to l. Since B is irreducible there must be a reduced path without loops p in the determinant Γ from f to l, meaning there exists a sequence $f =: b_0, \dots, b_m := l$ such that $\text{supt}(b_i) \cap \text{supt}(b_{i+1})$ while $\text{supt}(b_i) \not\subseteq \text{supt}(b_{i+1})$ and $\text{supt}(b_{i+1}) \not\subseteq \text{supt}(b_i)$ for each i. We assume $f, l \in B$ do not appear as vertices of p without loss. If, in addition, this path is positive then we have $\text{src}(b_i) < \text{src}(b_{i+1}) < \text{dest}(b_i) < \text{dest}(b_{i+1})$ for all i and so this forms a transition chain. Suppose p is not positive; then $p = p_1q_1 \dots p_nq_np_{n+1}$ for positive paths p_i and negative paths q_i. Consider p_1; there must be some p_j containing an edge e so that $\mathrm{dest}(t(p_1)) \subseteq \mathrm{supt}(i(e))$; otherwise it would be impossible to reach l. Thus we see that $p_1es_jq_j \dots p_{n+1}$ where $p_j = p'es_j$ is also a path of Γ and we may recursively apply this argument to see that there exists a positive path from f to l.

Notice, also, that we may assume $\mathrm{supt}(f) \cap \mathrm{supt}(l) = \emptyset$. If this is not the case but there is a bump b such that $\mathrm{dest}(f) \cup \mathrm{src}(l) \subseteq \mathrm{supt}(b)$ then we may cyclically shift the feet contained in $\mathrm{supt}(b)$ to make it so. Otherwise, assuming $|B| > 2$, there must be a bump b such that either $\mathrm{src}(l) \subseteq \mathrm{supt}(b) \subseteq \mathrm{supt}(f)$ or $\mathrm{dest}(f) \subseteq \mathrm{supt}(b) \subseteq \mathrm{supt}(l)$. In either case we may cyclically shift the feet contained in either $\mathrm{supt}(f)$ or $\mathrm{supt}(l)$ so that $\mathrm{src}(l) \cup \mathrm{dest}(f) \subseteq \mathrm{supt}(b)$.

From here, we show how we may repeatedly apply Lemma 4.3.2 to obtain a set B of bumps generating the same group isomorphism type such that every source of $B \setminus \{l\}$ is contained in $\mathrm{supt}(f)$ while every destination of $B \setminus \{f\}$ is contained in $\mathrm{supt}(l)$. First, suppose that there a destination d of $B \setminus \{f, l\}$ such that there is a positive path from b to l in Γ where $d = \mathrm{dest}(b)$. Then there is a transition chain $\mathcal{T} = (b =: t_0, \dots, t_k := l)$ where $\mathrm{supt}(t_{i+1})$ must contain $\mathrm{dest}(t_i)$ and $\mathrm{src}(t_{i+2})$ but not $\mathrm{src}(t_i)$ and $\mathrm{dest}(t_{i+2})$ for each i. Thus we may cyclically shift the feet contained in $\mathrm{supt}(t_{i+1})$ until $\mathrm{dest}(t_i)$ is contained in $\mathrm{supt}(t_{i+2})$ for any i so that $(t_0, \dots, t_i, t_{i+2}, \dots, t_k)$ forms a transition chain and so we may eventually form (b, l) as a transition chain. Furthermore, we can see that no destinations were removed from $\mathrm{supt}(l)$ nor were any sources removed from $\mathrm{supt}(f)$ in this process. Similarly, neither f nor l has been moved. This argument also shows that we may move a source $\mathrm{src}(b)$ into $\mathrm{supt}(f)$ without removing sources from $\mathrm{supt}(f)$ or removing destinations from $\mathrm{supt}(l)$ if there is a transition chain from f to b.

Now, we show that B can be transformed into a system such that for every bump b, if $\mathrm{supt}(b) \not\subseteq \mathrm{supt}(l)$ then there is a transition chain from b to l and if $\mathrm{supt}(b) \not\subseteq \mathrm{supt}(f)$ then there is a transition chain from f to b. It then follows that an arbitrary B can be transformed into a system such that every source is contained in $\mathrm{supt}(f)$ and every destination is contained in $\mathrm{supt}(l)$ by the argument of the previous paragraph. So, suppose without loss that we have a bump b such that $\mathrm{supt}(b) \not\subseteq \mathrm{supt}(l)$ but there is not a transition chain from b to l; let $p_1q_1 \dots q_{k-1}p_k$ be a shortest reduced path without loops from b to l in the determinant Γ, where p_i is positive and q_i is negative, and without loss suppose that $p_1 \neq \emptyset$. Notice that $\mathrm{supt}(B(p_i) \setminus \{i(q_i)\}) \subseteq \mathrm{supt}(b(q_i)) \setminus \mathrm{supt}(c(q_i))$ and $\mathrm{supt}(B(q_i) \setminus \{i(p_{i+1})\}) \subseteq \mathrm{supt}(b(p_{i+1})) \setminus \mathrm{supt}(c(p_{i+1}))$ for all i, where $B(p) \subseteq B$ is the set of bumps used in the path p of Γ while $b(p), c(p)$ are the bumps such that $(i(p), b(p))(b(p), c(p))$ is a prefix of p. This is because if this is not the case for some i then, say, $\mathrm{supt}(B(p_i) \setminus \{i(q_i)\}) \not\subseteq \mathrm{supt}(b(q_i)) \setminus \mathrm{supt}(c(q_i))$ and this would imply there is a bump in $B(p_i) \setminus \{i(q_i)\}$ whose support intersects that of $c(q_i)$, contradicting the path from b to l being minimal. There is a transition chain from b to $t(p_1) = i(q_1)$ and we may repeatedly cyclically shift feet so that $(b, i(q_1))$ is a positive edge of Γ. Notice, then, that we must have $\text{dest}(b), \text{dest}(b(q_1)) \subseteq \text{supt}(i(q_1))$ so we may cyclically shift the feet here until $\text{dest}(b(q_1)) < \text{dest}(b)$ and, since we know $\text{src}(b) \subseteq \text{supt}(b(q_1))$, we thus have $(b, b(q_1))$ as an edge of Γ. Now the path from b to l has the form $q_1 p_2 \ldots q_{k-1} p_k$ and so we may apply this argument repeatedly until there is a positive path from b to l.

At last, we have obtained B as a system such that every source of $B \setminus \{l\}$ is contained in $\text{supt}(f)$ while every destination of $B \setminus \{f\}$ is contained in $\text{supt}(l)$ and $\text{supt}(f) \cap \text{supt}(l) = \emptyset$. Thus it only remains to make (f, l) a stretched transition chain. If there is a bump b so that the only feet contained in its support (besides its own) are $\text{dest}(f)$ and $\text{src}(l)$ - that is, (f, b, l) forms a stretched transition chain - then we simply swap these feet. Otherwise, there must be two bumps $b_f, b_l \in B$ such that $(f, b_f), (b_l, l)$ each form stretched transition chains and so we cyclically shift the feet contained in $\text{supt}(b_l)$ one place to the right, so that $\text{src}(l)$ becomes the leftmost foot of $B \setminus \{b_l\}$ contained in $\text{supt}(b_l)$. Notice it is still the case that the feet contained in $\text{supt}(f)$ are precisely the sources of B (besides $\text{dest}(f)$). Thus we shift the feet in $\text{supt}(f)$ until (f, l) forms a stretched transition chain, and now B is in ideal form. □

Returning to our discussion of the isomorphism problem for groups generated by geometrically fast sets of bumps B, we may apply this lemma to simplify the problem. Indeed, if $|B| < 4$ then it follows that there is only one irreducible system of bumps with ideal form; if $|B| = 3$ then two bumps must form the first f and the last l, and the positions of these are fully determined, leaving only one bump b whose position is also fully determined by the definition of ideal. Similarly, if $|B| = 4$ then there are precisely two irreducible bump systems of ideal form; one of them is known to generate F_4 (see Section 8 of [7], for instance), while the other generates an unknown isomorphism type. The dynamical diagram of the unknown isomorphism type is shown in Figure 4.5.

This unknown isomorphism type has been dubbed pseudo-F_4, sometimes called PF_4. In what follows we will show that $PF_4 \cong F_4$.

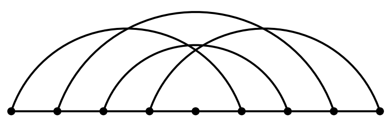

Figure 4.5: The four bumps in ideal form that generate the group pseudo-F_4

4.3.2 Proof

The proof begins by establishing a particular diagram group representation for F_4 over a semigroup presentation $\mathcal{P}_1$ which is, essentially, a refinement of the standard representation over $\langle x \mid x^4 = x \rangle$. We then take the representation produced by Theorem 4.2.7 from the ideal form bump system generating PF_4, over the semigroup presentation $\mathcal{P}_2$, and repeatedly apply moves from Theorem 2.3.7 to obtain $\mathcal{P}_1$ from $\mathcal{P}_2$ whence it follows that the diagram groups in question are isomorphic.

Standard

Recall that a **standard 4-adic interval** in $[0, 1]$ is any interval of the form

$$\left[\frac{i}{4^n}, \frac{i+1}{4^n}\right]$$

where $0 \leq i < 4^n$. The standard 4-adic intervals in $[0, 1]$ form a rooted 4-ary tree under inclusion. There is a natural partial action of F_4 on these intervals, where $f \in F_4$ maps an interval J to an interval J' if f is linear on J and $f(J) = J'$. Thus we may represent the elements of F_4 using pairs of 4-ary rooted trees (just as F can be represented using pairs of binary rooted trees [14]).

This is closely related to the standard representation of F_4 as a diagram group. Guba and Sapir show in [29] that F_4 is isomorphic to the diagram group over $\langle x \mid x^4 = x \rangle$ with base x (in fact any base word will do). As discussed in [32, p. 1111] any reduced diagram Δ over this presentation (since it is tree-like) can be decomposed as $\Delta = \Delta_1 \circ \Delta_2$ where Δ_1 only contains cells of the form (x, x^4) and Δ_2 only contains cells of the form (x^4, x). If $f \in F_4$ is represented by a reduced diagram Δ then the strand diagrams of Δ_1 and Δ_2^{-1} (omitting the top boundary vertex of each) give a pair of binary rooted trees representing f. In the other direction, given a pair of binary trees we can horizontally reflect the codomain tree and join the two along their leaves according to the action of f to obtain the interior of a strand diagram where every edge is unlabelled, which is equivalent to being labelled by the same letter x, and the diagram of this represents f.

Intervals Types

The **type** of a standard 4-adic interval is the value of i modulo 3 (either 0, 1, or 2). It is easy to check that if a 4-adic interval has type τ, then its four children have types τ, $\tau + 1$, $\tau + 2$, and τ, respectively, where the addition is modulo 3.

Note that if T is any finite, rooted subtree of the tree of 4-adic intervals, then the types of the leaves of T are precisely $0, 1, 2, 0, 1, 2, 0, 1, 2, \ldots$, with the last leaf having type 0. It follows that if $f \in F_4$ has domain partition $J_1, \ldots, J_n$ and range partition $J_1', \ldots, J_n'$, then each pair of intervals J_i, J_i' have the same type. That is, the partial action of F_4 on the set of standard 4-adic intervals is type-preserving.

Given this, we may use the connection between pairs of rooted trees and diagrams to find a new semigroup presentation over which F_4 is a diagram group. Indeed, given a pair of 4-ary rooted trees representing an element f we can form the interior of a strand diagram representing f as before, except that instead of having unlabelled edges we label each according to the type of the interval which that edge led to in the rooted tree.

It follows from this analysis that F_4 is isomorphic to the diagram group over the presentation

$$\langle x_0, x_1, x_2 \mid x_0 = x_0x_1x_2x_0,\ x_1 = x_1x_2x_0x_1,\ x_2 = x_2x_0x_1x_2\rangle$$

with base word x_0. Indeed, any base word of the form $(x_0x_1x_2)^n x_0$ gives F_4.

Even more letters

One way of looking at the introduction of the letters x_0, x_1, and x_2 above is that they represent orbits of intervals under the partial action of F_4 on the standard 4-adic intervals. Of course, this isn't quite true—although the intervals of type 1 form a single orbit and the intervals of type 2 form a single orbit, there are actually *four* orbits of intervals of type 0. Specifically, the interval $[0, 1]$ of type 0 is in its own orbit, the remaining intervals that contain 0 form an orbit, the remaining intervals that contain 1 form an orbit, and the last orbit consists of all intervals of type 0 that lie in $(0, 1)$.

This leads to another diagram group representation of F_4, where we break the letter x_0 into four letters u_0, v_0, w_0, x_0 such that the letter u_0 represents the whole interval $[0, 1]$, the letter v_0 represents other intervals that contain 0, the letter w_0 represents other intervals that contain 1, and the letter x_0 represents intervals of type 0 that lie in $(0, 1)$. The resulting semigroup presentation is

$$\left\langle u_0, v_0, w_0, x_0, x_1, x_2 \;\middle|\; \begin{array}{l} u_0 = v_0x_1x_2w_0,\ v_0 = v_0x_1x_2x_0,\ w_0 = x_0x_1x_2w_0, \\ x_0 = x_0x_1x_2x_0,\ x_1 = x_1x_2x_0x_1,\ x_2 = x_2x_0x_1x_2 \end{array} \right\rangle$$

It is not hard to show that the diagram group over this presentation with base word u_0 is isomorphic to F_4. The base word $v_0x_1x_2w_0$ also suffices, in which case we can remove the generator u_0 and the relation $u_0 = v_0x_1x_2w_0$ from the presentation without affecting the isomorphism type by Theorem 2.3.7.

We now prove that pseudo-F_4 is isomorphic to F_4. First we label the eight feet of the bumps using the eight letters A, B, C, D, $\overline{A}$, $\overline{B}$, $\overline{C}$, and $\overline{D}$, as shown in

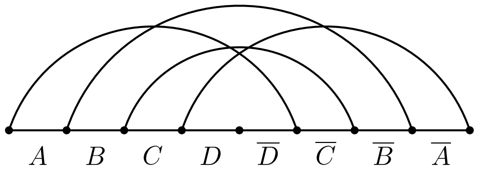

Figure 4.6: Labels for the eight feet of the bumps that generate pseudo-F_4.

Figure 4.6. According to Theorem 4.2.7, pseudo-F_4 is isomorphic to the diagram group for the presentation

$$\left\langle \begin{array}{l} A, B, C, D, \\ \overline{A}, \overline{B}, \overline{C}, \overline{D} \end{array} \;\middle|\; \begin{array}{l} A = ABCD,\ B = BCD\overline{DC},\ C = CD\overline{D},\ D = D\overline{DCB} \\ \overline{A} = \overline{DCBA},\ \overline{B} = CD\overline{DCB},\ \overline{C} = D\overline{DC},\ \overline{D} = BCD\overline{D} \end{array} \right\rangle$$

with base word $ABCD\overline{DCBA}$. Note that this presentation is symmetric with respect to the operation of:

(a) Switching the pairs $(A, \overline{A})$, $(B, \overline{B})$, $(C, \overline{C})$, and $(D, \overline{D})$ and then

(b) Reversing the order of the generators in every word.

We now apply a sequence of Guba and Sapir's moves from Theorem 2.3.7 to this presentation to obtain further diagram groups that are isomorphic to pseudo-F_4. First, since the last relation is $\overline{D} = BCD\overline{D}$, we can replace the $BCD\overline{D}$ in the relation $B = BCD\overline{DC}$ by a $\overline{D}$, and similarly we can replace the $D\overline{DCB}$ in the relation $\overline{B} = CD\overline{DCB}$ by a D. This yields the semigroup presentation

$$\left\langle \begin{array}{l} A, B, C, D, \\ \overline{A}, \overline{B}, \overline{C}, \overline{D} \end{array} \;\middle|\; \begin{array}{l} A = ABCD,\ B = \overline{DC},\ C = CD\overline{D},\ D = D\overline{DCB} \\ \overline{A} = \overline{DCBA},\ \overline{B} = CD,\ \overline{C} = D\overline{DC},\ \overline{D} = BCD\overline{D} \end{array} \right\rangle$$

Next we replace $D\overline{DC}$ by $\overline{C}$ in the relation $D = D\overline{DCB}$, and we replace $CD\overline{D}$ by C in the relation $\overline{D} = BCD\overline{D}$, yielding the presentation

$$\left\langle \begin{array}{l} A, B, C, D, \\ \overline{A}, \overline{B}, \overline{C}, \overline{D} \end{array} \;\middle|\; \begin{array}{l} A = ABCD,\ B = \overline{DC},\ C = CD\overline{D},\ D = \overline{CB} \\ \overline{A} = \overline{DCBA},\ \overline{B} = CD,\ \overline{C} = D\overline{DC},\ \overline{D} = BC \end{array} \right\rangle$$

Next we replace CD by $\overline{B}$ in the relation $C = CD\overline{D}$, and we replace $\overline{DC}$ by B in the relation $\overline{C} = D\overline{DC}$, yielding the presentation

$$\left\langle \begin{array}{l} A, B, C, D, \\ \overline{A}, \overline{B}, \overline{C}, \overline{D} \end{array} \;\middle|\; \begin{array}{l} A = ABCD,\ B = \overline{DC},\ C = \overline{BD},\ D = \overline{CB} \\ \overline{A} = \overline{DCBA},\ \overline{B} = CD,\ \overline{C} = DB,\ \overline{D} = BC \end{array} \right\rangle.$$

By Guba and Sapir's theorem, the diagram group over this presentation with base word $ABCD\overline{DCBA}$ remains isomorphic to pseudo-F_4.

Next we break the symmetry by substituting CD for $\overline{B}$, substituting DB for $\overline{C}$, and substituting BC for $\overline{D}$ in all of the other relations as well as in the base word. This yields the semigroup presentation

$$\left\langle \begin{array}{l} A, B, C, D, \\ \overline{A}, \overline{B}, \overline{C}, \overline{D} \end{array} \middle| \begin{array}{l} A = ABCD,\ B = BCDB,\ C = CDBC,\ D = DBCD \\ \overline{A} = BCDBCD\overline{A},\ \overline{B} = CD,\ \overline{C} = DB,\ \overline{D} = BC \end{array} \right\rangle$$

with base word $ABCDBCDBCD\overline{A}$. We can now eliminate the generators $\overline{B}$, $\overline{C}$, and $\overline{D}$ as well as the corresponding relations to get the semigroup presentation

$$\left\langle A, B, C, D, \overline{A} \middle| \begin{array}{l} A = ABCD,\ B = BCDB,\ C = CDBC,\ D = DBCD \\ \overline{A} = BCDBCD\overline{A} \end{array} \right\rangle$$

with base word $ABCDBCDBCD\overline{A}$. If we replace $BCDB$ by B in the relation for $\overline{A}$ as well as twice successively in the base word, we obtain the semigroup presentation

$$\left\langle A, B, C, D, \overline{A} \middle| \begin{array}{l} A = ABCD,\ B = BCDB,\ C = CDBC,\ D = DBCD \\ \overline{A} = BCD\overline{A} \end{array} \right\rangle$$

with base word $ABCD\overline{A}$, and the corresponding diagram group remains isomorphic to pseudo-F_4.

This is almost the same as the diagram group representation of F_4 discussed above. To finish, we introduce a new generator E to the presentation:

$$\left\langle A, B, C, D, \overline{A}, E \middle| \begin{array}{l} A = ABCD,\ B = BCDB,\ C = CDBC,\ D = DBCD \\ \overline{A} = BCD\overline{A},\ E = D\overline{A} \end{array} \right\rangle$$

Next we substitute E for $D\overline{A}$ in the relation for $\overline{A}$:

$$\left\langle A, B, C, D, \overline{A}, E \middle| \begin{array}{l} A = ABCD,\ B = BCDB,\ C = CDBC,\ D = DBCD \\ \overline{A} = BCE,\ E = D\overline{A} \end{array} \right\rangle$$

and similarly in the base word to get the new base word $ABCE$. Next we substitute BCE for $\overline{A}$ in the relation for E:

$$\left\langle A, B, C, D, \overline{A}, E \middle| \begin{array}{l} A = ABCD,\ B = BCDB,\ C = CDBC,\ D = DBCD \\ \overline{A} = BCE,\ E = DBCE \end{array} \right\rangle$$

which allows use to eliminate the generator $\overline{A}$ and its corresponding relation. The result is the presentation

$$\left\langle A, B, C, D, E \;\middle|\; \begin{array}{l} A = ABCD,\ B = BCDB,\ C = CDBC,\ D = DBCD \\ E = DBCE \end{array} \right\rangle$$

with base word $ABCE$. This is exactly the diagram group representation for F_4 discussed above, with $v_0 = A$, $w_0 = E$, $x_0 = D$, $x_1 = B$, and $x_2 = C$, so we conclude that pseudo-F_4 is isomorphic to F_4.

Chapter 5

Fast groups acting on the circle

Having established Theorem 4.2.7, in this chapter we now turn to the analogous question for homeomorphism groups of the circle - given a fast set (of a particular form) does the group it generates have a faithful representation as a diagram group? We find a partial answer, detailed in Theorem 5.2.5, which says such a group is isomorphic to a particular annular diagram group precisely when a condition on the fast set is satisfied. It is unclear whether those fast sets which do not satisfy the condition generate groups isomorphic to some other diagram group. Along the way, we introduce an analogue of dynamical diagrams for a subclass fast sets on the circle.

We begin by defining a small similarity structure S on the space of histories of an abstract ping-pong system on a set of permutations A. We show that, in general, the group generated by the ping-pong representation $\hat{A}$ embeds in the group determined by S. By the Farley-Hughes theorem, this means $\langle \hat{A} \rangle$ embeds in the symmetric diagram group D over the semigroup presentation determined by the blueprint of A, which generalises the semigroup presentation determined by the canonical partition of a fast set (Corollary 5.1.5). By restricting (Proposition 5.1.7) to cases where the blueprint is orderable we can show that this embedding induces an isomorphism of the planar subgroup of D, recovering Theorem 4.2.7 (as mentioned in the introduction of Chapter 4). Restricting to cases where the blueprint is cyclically orderable produces what will turn out to be the class of fast sets acting on the circle mentioned above.

5.1 A similarity structure on the space of histories

Recall the definitions from Section 2.2.1 and consider a finite set of permutations A equipped with a ping-pong system with space of histories K_A. Our first step is to show that K_A is an ultrametric space.

Lemma 5.1.1. *The space of histories K_A for a ping-pong system on a set of permutations A is metrisable. More specifically, there is an ultrametric d on K_A which induces the topology on K_A.*

Proof. We define the ultrametric $d : K_A \times K_A \to K_A$ by

$$d(\eta, \mu) = \begin{cases} 0 & \text{if } \eta = \mu \\ \frac{1}{e^u} & \text{otherwise, where } u = \sup\{\text{len}(w) \mid w \in \eta \cap \mu\} \end{cases}$$

where $\text{len}(w)$ denotes the length of the word w. To see that this satisfies the strong triangle inequality, consider $\eta, \mu, \nu \in K_A$ and notice that if $d(\eta, \nu) > d(\eta, \mu)$ we must have $\sup\{\text{len}(w) \mid w \in \eta \cap \nu\} < \sup\{\text{len}(w) \mid w \in \eta \cap \mu\}$ and let $u \in \eta \cap \mu$ be the word which attains this bound. Since any history is suffix closed we have $v \in \eta \cap \mu$ for any suffix v of u and so $v \in \mu \cap \nu$ if and only if $v \in \eta \cap \nu$, and since any history contains at most one word of a given length we may now conclude that $\eta \cap \nu = \mu \cap \nu$ and it follows that $d(\eta, \nu) = d(\mu, \nu)$. It is otherwise clear that d is an ultrametric on K_A.

Consider a metric ball $B_d(\eta, r)$ for some $\eta \in K_A$, $r > 0$. We have $\mu \in B_d(\eta, r) \setminus \{\eta\}$ if and only if $d(\mu, \eta) \leq r$ if and only if $-\log(r) \leq \sup\{\text{len}(w) \mid w \in \eta \cap \mu\}$. Let n be the smallest non-negative integer such that $-\log(r) \leq n$; if η does not contain a word of length n then no μ can satisfy this condition and so $B_d(\eta, r) = \{\eta\} = [w]$ where w is the maximal length word in η. Otherwise, let $w \in \eta$ be the unique word of length n; any μ which satisfies $-\log(r) \leq \sup\{\text{len}(w) \mid w \in \eta \cap \mu\}$ must contain a word $u \in \eta$ of length greater than or equal to n. Since η and μ are both suffix closed and contain at most one word of a given length we must have $\{v \in \eta \mid \text{len}(v) \leq \text{len}(u)\} = \{v \in \mu \mid \text{len}(v) \leq \text{len}(u)\}$ and in particular $w \in \mu$. We conclude $B_d(\eta, r) = [w]$ as required. □

Remark. Effectively, we are showing that K_A is isometric to the endspace of a rooted tree; the prototypical ultrametric space. For more information on the correspondence between ultrametric spaces and endspaces of rooted trees see Hughes [35].

From here, we will consider the space of histories K_A to be a compact ultrametric space equipped with d. Thus, we may define a small similarity structure on it.

Definition 5.1.2. Let K_A be the space of histories for a ping-pong system on a set of permutations A, and let $\mathcal{B}$ be its ball set. For each pair (B_1, B_2) where $B_1, B_2 \in \mathcal{B}$ we define a set $\text{Sim}_A(B_1, B_2)$ as follows. If w_1, w_2 are words on $A^{\pm} \sqcup M_A$ such that $B_1 = [aw_1], B_2 = [aw_2]$ for some $a \in A^{\pm} \sqcup M_A$ then we define the function $S_{aw_1,aw_2} : [aw_1] \to [aw_2]$ by

$$\eta S_{aw_1,aw_2} = \{uaw_2 \mid uaw_1 \in \eta\} \sqcup \{s \mid s \text{ is a suffix of } w_2\}$$

referring to it as the **suffix similarity** for aw_1, aw_2, and we define

$$\mathrm{Sim}_A(B_1, B_2) = \{S_{aw_1,aw_2}\}$$

in this case. Otherwise, we define

$$\mathrm{Sim}_A(B_1, B_2) = \emptyset$$

and, finally, we define $\mathrm{Sim}_A = \{\mathrm{Sim}_A(B_1, B_2) \mid B_1, B_2 \in \mathcal{B}\}$.

Remark. Notice that $S_{aw_1,aw_1} = \mathrm{id}_{[aw_1]}$ for any w_1 such that $aw_1 \in L_A$.

Proposition 5.1.3. *Let K_A be the space of histories for a ping-pong system on a set of permutations A. Then Sim_A is a small similarity structure on K_A.*

Proof. First, we show that a suffix similarity is, indeed, a surjective similarity. Let $S_{aw_1,aw_2} : [aw_1] \to [aw_2]$ be a suffix similarity and consider the set of local reductions L_A of A. Notice that if uaw_1 is a word on $A^{\pm} \sqcup M_A$ then it follows from Definition 2.2.7 that $uaw_1 \in L_A$ if and only if $uaw_2 \in L_A$. Thus we can see that $\eta S_{aw_1,aw_2} \subseteq L_A$ if $\eta \in [aw_1]$. By definition, $uaw_2 \in \eta S_{aw_1,aw_2}$ for every $uaw_1 \in \eta$ and $s \in \eta S_{aw_1,aw_2}$ for every suffix of w_2, so we can also see $\eta S_{aw_1,aw_2}$ is suffix closed, contains at most one word of a given length and every word v on $A^{\pm}$ such that $v \in \eta S_{aw_1,aw_2}$ is a suffix of some $v' \in \eta S_{aw_1,aw_2}$ since η itself satisfies these conditions. Hence, $\eta S_{aw_1,aw_2} \in K_A$ and $aw_2 \in \eta S_{aw_1,aw_2}$, so $\eta S_{aw_1,aw_2} \in [aw_2]$.

Consider $\eta, \mu \in [aw_1]$ with $d(\eta, \mu) = \frac{1}{e^{\mathrm{len}(v)}}$ where v is the longest word in $\eta \cap \mu$. Noticing that v must contain aw_1 as a suffix, say $v = v'aw_1$, consider $d(\eta S_{aw_1,aw_2}, \mu S_{aw_1,aw_2})$. By definition of S_{aw_1,aw_2} we can see that $v'aw_2 \in \eta S_{aw_1,aw_2} \cap \mu S_{w_1,w_2}$ and it is the longest such word, since if $w' \in \eta S_{aw_1,aw_2} \cap \mu S_{aw_1,aw_2}$ it is either a suffix of w_2 or $w' = pw_2$ such that $pw_1 \in \eta \cap \mu$. Thus we have

$$d(\eta S_{aw_1,aw_2}, \mu S_{aw_1,aw_2}) = \frac{1}{e^{\mathrm{len}(v'w_2)}} = \frac{1}{e^{\mathrm{len}(v)} e^{\mathrm{len}(w_2)-\mathrm{len}(w_1)}} = \frac{1}{e^{\mathrm{len}(w_2)-\mathrm{len}(w_1)}} d(\eta, \mu)$$

since $\mathrm{len}(v') = \mathrm{len}(v) - \mathrm{len}(w_1)$ implies $\mathrm{len}(v'w_2) = \mathrm{len}(v) - \mathrm{len}(w_1) + \mathrm{len}(w_2)$, and so S_{aw_1,aw_2} is a similarity. This implies that S_{aw_1,aw_2} is injective and we may conclude it is bijective by noting that $S_{aw_1,aw_2} \circ S_{aw_2,aw_1} = S_{aw_2,aw_1} \circ S_{aw_1,aw_2} = \mathrm{id}_{[aw_1]}$.

Finally, let us confirm that Sim_A satisfies Definition 2.4.1. It is clear that Sim contains identities, is inverse closed (by the last sentence of the previous paragraph) and every element is finite (indeed, it is small). Let w_1, w_2, w_3 be words such that $aw_1, aw_2, aw_3 \in L_A$ and consider $S_{aw_1,aw_2} \circ S_{aw_2,aw_3}$. We have

$$\begin{aligned}\eta S_{aw_1,aw_2} \circ S_{aw_2,w_3} &= (\{uaw_2 \mid uaw_1 \in \eta\} \sqcup \{s \mid \ \text{is a suffix of } w_2\}) S_{aw_2,aw_3} \\ &= \{uaw_3 \mid uaw_1 \in \eta\} \sqcup \{s \mid s \text{ is a suffix of } w_3\} = \eta S_{aw_1,aw_3}\end{aligned}$$

for every $\eta \in [aw_1]$. Now, let B be a subball of $[aw_1]$ - then $B = [vaw_1]$ for some $vaw_1 \in L_A$ - and consider $S_{aw_1,aw_2}|_B$. We have

$$\eta S_{aw_1,aw_2}|_B = \{uvaw_2 \mid uvaw_1 \in \eta\} \sqcup \{s \mid \text{ is a proper suffix of } vaw_2\} = \eta S_{vaw_1,vaw_2}$$

for every $\eta \in B$. $\square$

With this in place, we can now see a group generated by a ping-pong representation embeds in the group locally determined by this similarity structure.

Proposition 5.1.4. *Let K_A be the space of histories for a set of permutations A with ping-pong representation $\hat{A}$. Then $\langle \hat{A} \rangle$ is a subgroup of $\Gamma(\mathrm{Sim}_A)$.*

Proof. It suffices to show that $\hat{a} \in \hat{A}$ is locally determined by Sim_A. Notice that

$$\eta\hat{a} = \begin{cases} \eta S_{b,ba} & \text{if } \eta \in [b] \text{ such that } \mathrm{dest}(b) \subseteq \mathrm{supt}(a) \setminus \mathrm{src}(a) \\ \eta S_{ba^{-1},b} & \text{if } \eta \in [ba^{-1}] \text{ where } ba^{-1} \in L_A \\ \eta & \text{if } \eta \in [c] \text{ for some } c \in A \text{ such that } \mathrm{dest}(c) \cap \mathrm{supt}(a) = \emptyset \\ \eta & \text{if } \eta \in [\bar{x}] \ \&\ a \notin \bar{x} \\ \eta S_{\bar{x},\bar{x}a} & \text{if } \eta \in [\bar{x}] \ \&\ a \in \bar{x} \end{cases}$$

where we have simply translated the definition of $\hat{a}$ from Definition 2.2.8 in terms of similarities from Sim_A. Thus $\hat{A} \subseteq \Gamma(\mathrm{Sim}_A)$ as required. $\square$

Now applying the Farley-Hughes theorem (Theorem 2.4.3) we obtain the following corollary.

Corollary 5.1.5. *Let $\hat{A}$ be a ping-pong representation of some set of permutations A, and let $\mathcal{B}$ be the blueprint of the ping-pong system on A. Then $\langle \hat{A} \rangle$ embeds in $D_s(\mathcal{P}_{\mathrm{Sim}_A}, \varepsilon)$ where*

$$\mathcal{P}_{\mathrm{Sim}_A} = \langle A^{\pm} \sqcup M_A \sqcup \{\varepsilon\} \mid (a, (\mathrm{supt}_{\mathcal{B}}(a) \setminus \{a^{-1}\}, L_a)) \text{ for } a \in A^{\pm} \rangle$$

where L_a is some total order on $\mathrm{supt}_{\mathcal{B}}(a) \setminus \{a^{-1}\}$.

Remark. We leave the total orders L_a, which would be determined given a ball order on K_A compatible with Sim_A, unspecified since this does not effect the isomorphism type of $D_s(\mathcal{P}, \varepsilon)$.

Proof. Here, we use the identification $[[a]] \mapsto a$ for $a \in A^{\pm} \sqcup M_A \sqcup \{\varepsilon\}$ where $[[a]] = \{[aw] \mid aw \in L_A\}$ is the Sim_A-class containing $[a]$. The corollary then follows from Theorem 2.4.3, noting that $\mathrm{supt}_{\mathcal{B}}(a) \setminus \{a^{-1}\}$ is precisely the set of letters b such that $[baw]$ is a maximal subball of $[aw]$ for any w such that $aw \in L_A$. $\square$

Definition 5.1.6. We will refer to a symmetric diagram group $D_s(\mathcal{P}_{\mathrm{Sim}_A}, \varepsilon)$ with a specified total order L_a for each $a \in A^\pm$ defined in Corollary 5.1.5 as a **diagram group of $\hat{A}$**.

This is all we will say about ping-pong representations in general. We can, however, say something stronger for appropriate subclasses of ping-pong representations.

Proposition 5.1.7. *Let $\hat{A}$ be a ping-pong representation of a set A with blueprint $\mathcal{B}$. If $\mathcal{B}$ is totally orderable then let T be a total order on $A^\pm \sqcup M_A$ which witnesses this and define L_a be the total order on $\mathrm{supt}_{\mathcal{B}}(a) \setminus \{a^{-1}\}$ induced by T. If $D_s(\mathcal{P}_{\mathrm{Sim}_A}, \varepsilon)$ is the diagram group of $\hat{A}$ with the specified total orders L_a then*

$$\langle \hat{A} \rangle \cong D_p(\mathcal{P}_{\mathrm{Sim}_A}, \varepsilon)$$

and this isomorphism is induced by the embedding $\langle \hat{A} \rangle \hookrightarrow D_s(\mathcal{P}_{\mathrm{Sim}_A}, \varepsilon)$ from Corollary 5.1.5.

Meanwhile, if $\mathcal{B}$ is cyclically orderable then

$$\langle \hat{A} \rangle \hookrightarrow D_a(\mathcal{P}_{\mathrm{Sim}_A}, \varepsilon)$$

which is induced by the embedding from Corollary 5.1.5, where $D_s(\mathcal{P}_{\mathrm{Sim}_A}, \varepsilon)$ is the diagram group of $\hat{A}$ with specified total orders induced by a cyclic order on $A^\pm \sqcup M_A$ which witnesses $\mathcal{B}$ being cyclically orderable. Moreover, we have

$$\langle \hat{A} \rangle \cong D_a(\mathcal{P}_{\mathrm{Sim}_A}, \varepsilon)$$

precisely when the label of $\lfloor \Delta \rfloor$ is aperiodic for every positive planar diagram Δ over $\mathcal{P}$ with top path label ε.

Proof. For the first claim, it suffices to show that $\Omega(\langle \hat{A} \rangle) = S(\mathcal{P}_{\mathrm{Sim}_A}, \varepsilon)$ where Ω is as in Definition 2.4.4. In other words, we want to show that if $\Omega(\gamma) = S_{\mathcal{C}_1} \circ \sigma_\phi \circ S_{\mathcal{C}_2}^{-1}$ for $\gamma \in \Gamma(\mathrm{Sim}_A)$ then $\sigma_\phi : \lfloor S_{\mathcal{C}_1} \rfloor \to \lfloor S_{\mathcal{C}_2} \rfloor$ is a non-converse order isomorphism if and only if $\gamma \in \langle \hat{A} \rangle$.

Let $\hat{a} \in \hat{A}$. Recalling the definition of $\hat{a}$, and in particular the equivalent definition in the proof of Proposition 5.1.4, we define partitions of K_A by $\mathcal{C}_1 = \{[ba^{-1}] \mid ba^{-1} \in L_A\} \sqcup \{[c] \mid c \in A^\pm \sqcup M_A \setminus \{a^{-1}\}\}$ and $\mathcal{C}_2 = \{[ba] \mid b \in \mathrm{supt}_{\mathcal{B}}(a) \setminus \{a^{-1}\}\} \sqcup \{[b] \mid b \in A^\pm \sqcup M_A \setminus \{a\}\}$. Let $a_1 \dots a_n$ be the word corresponding to the totally ordered set $(A^\pm \sqcup M_A, T)$ where $a_i = a^{-1}, a_j = a$ assuming $i < j$ without loss. We can then see that

$$l(\lfloor S_{\mathcal{C}_1} \rfloor) = a_1 \dots a_i \dots a_{j-1} a_{i+1} \dots a_n$$

and

$$l(\lfloor S_{\mathcal{C}_2} \rfloor) = a_1 \dots a_{j-1} a_{i+1} \dots a_j a_{j+1} \dots a_n$$

since we know supt(a) forms an interval of T. Notice $a_k \in \text{supt}_{\mathcal{B}}(a) \setminus \{a^{-1}\}$ precisely when $i+1 \leq k \leq j$, $a_k a^{-1} \in L_A$ precisely when $i \leq k \leq j-1$ and otherwise $a_k \notin \text{supt}_{\mathcal{B}}(a)$. We can then see that the map $\phi : \mathcal{C}_1 \to \mathcal{C}_2$ is defined by

$$\begin{aligned} \phi([a_k a^{-1}]) &= [a_k] \text{ where } i+1 \leq k \leq j \\ \phi([a_k]) &= [a_k a] \text{ where } i \leq k \leq j-1 \\ \phi([a_k]) &= [a_k] \text{ where } k \in \{1, \dots, i-1\} \sqcup \{j+1 \dots n\} \end{aligned}$$

and, noticing that $[a_k a^{-1}]$ corresponds to the first occurrence of the letter a_k in $l(\lfloor S_{\mathcal{C}_1} \rfloor)$ while $[a_k a]$ corresponds to the second occurrence of a_k in $l(\lfloor S_{\mathcal{C}_2} \rfloor)$, we see that σ_ϕ is the non-converse order isomorphism, as required. We can now see that

$$\begin{aligned} \Omega(\hat{a}) = (\varepsilon, a_1 \dots a_n) \circ (\epsilon_{a_1 \dots a_{i-1}} + (a, a_i \dots a_{j-1}) + \epsilon_{a_{i+1} \dots a_n}) \\ \circ (\epsilon_{a_1 \dots a_{j-1}} + (a_j, a_{i+1} \dots a_j) + \epsilon_{a_{j+1} \dots a_n})^{-1} \circ (\varepsilon, a_1 \dots a_n)^{-1} \end{aligned}$$

in $D_p(\mathcal{P}, \varepsilon)$ where $\mathcal{P} = \langle A^{\pm} \sqcup M_A \sqcup \{\varepsilon\} \mid \mathcal{R} \rangle$ and

$$\begin{aligned} \mathcal{R} = \{(\varepsilon, a_1 \dots a_n)\} \\ \sqcup \{(a_i, a_i \dots a_{j-1}), (a_j, a_{i+1} \dots a_j) \mid a \in A \text{ where } a_i = a^{-1}, a_j = a\} \end{aligned}$$

Notice that conjugating by $(\varepsilon, a_1 \dots a_n)$ shows $D_p(\mathcal{P}, \varepsilon) \cong D_p(\mathcal{P}, a_1 \dots a_n)$ and we may apply Theorem 2.3.7 to remove the letter ε and the relation $(\varepsilon, a_1 \dots a_n)$ from $\mathcal{P}$ without changing the isomorphism type of $D_p(\mathcal{P}, a_1 \dots a_n)$. Comparing this with the definition of a bumpy diagram group (Definition 4.1.1) we can see they are identical apart from the letters M_A. Recall that, when defining bumpy diagram groups, we chose the canonical partition to be as efficient as possible; this was simply for tidiness, and if we included letters representing the gaps between feet in a bumpy presentation we would find this will not affect the validity of the proofs from Chapter 4; these letters would then play the same role as M_A. In particular, all of this means that it follows from a minor modification to these proofs, and in particular the proof of Proposition 4.2.6, that $\langle \Omega(\hat{A}) \rangle = D_p(\mathcal{P}_{\text{Sim}_A}, \varepsilon)$.

The second claim follows similarly to the first half of the above proof - when the relations of $\mathcal{P}_{\text{Sim}_A}$ have total orders induced by the cyclic order on $A^{\pm} \sqcup M_A$ we will find $\Omega(\hat{a}) = S_{\mathcal{C}_1} \circ \sigma_\phi \circ S_{\mathcal{C}_2}^{-1}$ where σ_ϕ is a non-converse cyclic order isomorphism. For the third, notice that if u, w are words on $\hat{A}^{\pm}$ such that $\Omega(u) = S_{\mathcal{C}_1} \circ \sigma_{\phi_u} \circ S_{\mathcal{C}_2}$ and

$\Omega(w) = S_{\mathcal{D}_1} \circ \sigma_{\phi_w} \circ S_{\mathcal{D}_2}$ then

$$S_{\mathcal{C}_i} \equiv S_{\mathcal{D}_i} \text{ for } i = 1,2 \implies \Omega(u) = \Omega(w)$$

since the antecedent implies that $L^{\vee}(u) = L^{\vee}(w)$ which, in turn, implies that $u = w$. It is also true that, for any positive strand diagram S over $\mathcal{P}_{\mathrm{Sim}_A}$ with top boundary vertex label ε such that $S \circ \sigma \circ S'^{-1}$ is an annular diagram over $\mathcal{P}_{\mathrm{Sim}_A}$, there is some word w on $\hat{A}^{\pm}$ such that $S_{\mathcal{C}_1} \equiv S$ where $\Omega(w) = S_{\mathcal{C}_1} \circ \sigma_{\phi} \circ S_{\mathcal{C}_2}$. Thus, if $l(\lfloor S \rfloor)$ is aperiodic for every positive strand diagram S over $\mathcal{P}_{\mathrm{Sim}_A}$ then there is at most one annular diagram of the form $S \circ \sigma \circ S'^{-1}$ it must be equivalent to $\Omega(w)$. On the other hand, if $l(\lfloor S \rfloor)$ is periodic for some S there is some σ such that $S \circ \sigma \circ S^{-1}$ is a non-trivial annular diagram and we see $S \circ \sigma \circ S^{-1} \notin \Omega(\langle \hat{A} \rangle)$ since we already have $S \circ S^{-1} \equiv \epsilon_{\varepsilon} \in \Omega(\langle \hat{A} \rangle)$. □

To understand the scope of this result, it is natural to wonder whether there are groups generated by faithful ping-pong representations with orderable blueprints which do not arise as fast sets of the interval. It turns out there are not.

Proposition 5.1.8. *Let $\hat{A}$ be a faithful ping-pong representation with orderable blueprint $\mathcal{B}$. Then there exists a geometrically fast set of bumps B such that $\hat{B}$ has blueprint isomorphic to $\mathcal{B}$.*

Proof. Let $a_1 \ldots a_n$ be the word corresponding to the total order on $A^{\pm} \sqcup M_A$ which witnesses $\mathcal{B}$ being orderable. Then the word corresponding to the suborder on $\mathrm{supt}_{\mathcal{B}}(a)$ is $a_i a_{i+1} \ldots a_j$ where $a_i = a^{-1}, a_j = a$. We may construct a graph on $a_1 \ldots a_n$ by defining an oriented edge from $a_i = a^{-1}$ to $a_j = a$ for each $a \in A$. This graph is isomorphic to a dynamical diagram for some fast set of bumps B and a choice of marking for B witnessing this defines an abstract ping-pong system on B which is isomorphic to $\mathcal{B}$. □

It follows from this that the first part of Proposition 5.1.7 is a recovery of Theorem 4.2.7.

The next question is: is there an analogous result for faithful representations with cyclically orderable blueprints? We tackle in the next section.

5.2 Fast groups on the circle

Let $S^1 = \{x \in \mathbb{R}^2 \mid \|x\| = 1\}$ be the unit circle and consider $\mathrm{Homeo}_+(S^1)$ the group of all increasing homeomorphisms of S^1. Just as we have considered geometrically fast sets of homeomorphisms of the unit interval I, we may do so with S^1. Just as we investigated groups generated by fast bump sets on I previously, we will be

(a) (b) (c) (d)

Figure 5.1: Non-isomorphic cyclic dynamical diagrams

investigating fast sets on S^1 consisting of bumps and another kind of closely related type of function.

Definition 5.2.1. Let $f \in \text{Homeo}_+(S^1)$. If f has two orbitals C_1, C_2 such that $f|_{C_1}$ is a positive bump and $f|_{C_2}$ is a negative bump and $C_1 \sqcup C_2 = S^1 \setminus \{x, y\}$ for some $x, y \in S^1$ then we call f a **2-bump**. If $A \subseteq \text{Homeo}_+(S^1)$ contains only bumps and 2-bumps we call it a **bump set**.

The majority of definitions from Section 2.2 adapt readily to this context. One exception is that, for a 2-bump f, we define $\text{src}(f) = \text{src}(f|_{C_1}) \sqcup \text{src}(f|_{C_2})$ and $\text{dest}(f) = \text{dest}(f|_{C_1}) \sqcup \text{dest}(f|_{C_2})$ where the sources and destinations of $f|_{C_i}$ are defined as before.

Another definition from Section 2.2 which does not fully adapt is that of dynamical diagrams. We address this now for our context.

Definition 5.2.2. Let $A \subseteq \text{Homeo}_+(S^1)$ be a bump set. We define the **cyclic dynamical diagram** C_A of A to be the (partially) oriented chord diagram defined as follows:

- Draw a unit circle (with counterclockwise orientation);
- Place a vertex on the circle for each foot of A in order;
- For each bump $b \in A$ draw an oriented edge from $\text{src}(b)$ to $\text{dest}(b)$;
- For each 2-bump function f draw an unoriented edge between $\text{src}(f)$ and $\text{dest}(f)$.

We consider two cyclic dynamical diagrams to be **isomorphic** if they are isomorphic as unlabelled (partially) oriented chord diagrams.

Remark. Given a bump b in a cyclic dynamical diagram C we understand the support of b to be the section of the circle between $\text{src}(b)$ and $\text{dest}(b)$ which has the same orientation as the edge b.

There are four examples of similar yet pairwise non-isomorphic cyclic dynamical diagrams shown in Figure 5.1. All except (b) are equivalent to a (linear) dynamical diagram.

We are now in a position to establish some facts about groups generated by bump sets. The first is a corollary of Theorem 2.2.10.

Corollary 5.2.3. *Let $A_1, A_2 \subseteq \mathrm{Homeo}_+(S^1)$ be fast bump sets with cyclic dynamical diagrams C_1, C_2. If C_1, C_2 are isomorphic then $\langle A_1\rangle \cong \langle A_2\rangle$.*

Proof. If C_1, C_2 are isomorphic then the blueprints of the abstract ping-pong systems determined by the markings on A_1, A_2 witnessing fastness must also be isomorphic. The result now follows from Theorem 2.2.10. □

Proposition 5.2.4. *Let $\hat{A}$ be a faithful ping-pong representation with cyclically orderable blueprint $\mathcal{B}$. Then there exists a geometrically fast bump set $A \subseteq \mathrm{Homeo}_+(S^1)$ such that $\hat{A}$ has blueprint isomorphic to $\mathcal{B}$.*

Proof. The proof is similar to that of Proposition 5.1.8. Let $a_1 \ldots a_n$ be the word corresponding to a total order on $A^{\pm} \sqcup M_A$ obtained by cutting the cyclic order C C which witnesses $\mathcal{B}$ being cyclically orderable and arrange $a_1 \ldots a_n$ counter-clockwise on a unit circle. We know $\mathrm{supt}_{\mathcal{B}}(a)$ for $a \in A$ is an interval in C. If a is an endpoint of the interval on $\mathrm{supt}_{\mathcal{B}}(a) \setminus \{a^{-1}\}$ then define an oriented edge from a^{-1} to a; if a is not an endpoint of $\mathrm{supt}_{\mathcal{B}}(a) \setminus \{a^{-1}\}$ then define an unoriented edge between a^{-1} and a. The defined graph is isomorphic to a cyclic dynamical diagram for some fast bump set A and a choice of marking which witness this defines an abstract ping-pong system on A with blueprint isomorphic to $\mathcal{B}$. □

This proposition tells us, by Theorem 2.2.10, that the class of groups generated by faithful ping-pong representations with cyclically orderable blueprints coincides with the class of groups generated by fast bump sets. It makes sense, then, to bring things 'down to earth' and translate what we learned in the previous section from ping-pong representation to fast bump sets.

As with the linear case, we may take the feet of A to cover S^1 modulo finitely many isolated points and a connected interval between the source and destination of each isolated bump. We take this canonical partition $a_1, \ldots, a_n$ of S^1 as the generating set of a semigroup presentation $\mathcal{P}$. The partition is cyclically ordered, and as such we define $a_i = a_{i (\bmod n)}$. For a bump $b \in A$ with $\mathrm{src}(b) = a_i$ and $\mathrm{dest}(b) = a_j$ we add to $\mathcal{P}$ the relations

$$(a_i, a_i \ldots a_{j-1})$$

$$(a_j, a_{i+1} \ldots a_j)$$

since the interval $a_{i+1} \ldots a_{j-1}$ is the fundamental domain of b defined by its marker. For a 2-bump $f \in A$ with $\mathrm{src}(f) = a_i$ and $\mathrm{dest}(f) = a_j$ we add to $\mathcal{P}$ the relations

$$(a_i, a_{j+1} \ldots a_{i-1} a_i a_{i+1} \ldots a_{j-1})$$

$$(a_j, a_{i+1} \dots a_{j-1} a_j a_{j+1} \dots a_{i-1})$$

since the interval $a_{i+1} \dots a_{j-1}$ is the fundamental domain of the positive bump while the interval $a_{j+1} \dots a_{i-1}$ is the fundamental domain of the negative bump, both defined by their respective markers.

Finally, we define a map $\delta : \langle A \rangle \to D_a(\mathcal{P}, a_1 \dots a_n)$ extending the map $a \mapsto \alpha^\circ$ where

$$\alpha = \pi_{a_i, a_i \dots a_{j-1}} + \pi_{a_{i+1} \dots a_j, a_j} + \varepsilon_{a_{j+1} \dots a_{i-1}}$$

if a is a bump and

$$\alpha = \pi_{a_i, a_{j+1} \dots a_{i-1} a_i a_{i+1} \dots a_{j-1}} + \pi_{a_{i+1} \dots a_{j-1} a_j a_{j+1} \dots a_{i-1}, a_j}$$

if a is a 2-bump. In both cases the inner vertex and the outer vertex of α° are, respectively, the vertex separating the edges labelled a_n and a_1 on the inner path and the outer path so that α° is an annular $(a_1 \dots a_n, a_1 \dots a_n)$-diagram

Theorem 5.2.5. *Let $A = B_1 \sqcup B_2$ be a geometrically fast bump set where B_1 is its subset of bumps and B_2 is its subset of 2-bumps. Then $\langle A \rangle \hookrightarrow D_a(\mathcal{P}, a_1 \dots a_n)$ and $\langle A \rangle \cong D_a(\mathcal{P}, a_1 \dots a_n)$ if and only if B_1 does not contain any closed stretched transition chains.*

Proof. This is a translation of the second part of Proposition 5.1.7, so it suffices to show that B_1 does not contain any closed stretched transition chains if and only if the bottom path label of every positive planar diagram over $\mathcal{P}$ is aperiodic. If B_1 does contain a closed stretched transition chain starting at a_i then

$$a_1 \dots a_i \dots a_n \to a_1 \dots a_i a_{i+1} \dots a_j a_{i+1} \dots a_n \to \dots$$
$$\dots \to a_1 \dots a_i \dots a_n a_1 \dots a_{i-1} a_i \dots a_n$$

is a positive derivation in the semigroup $\mathcal{P}$ starting at the base word and ending at a periodic word. On the other hand, if B_1 does not contain a closed stretched transition chain then consider a positive derivation from the base word using relations of the form $(a_i, a_i \dots a_j), (a_j, a_{i+1} \dots a_j)$. If we try to produce a periodic word as in the previous case, by following a stretched transition chain, we will eventually produce a subword of the form $a_k a_j$ for $j < k$; precisely which word will be only be the same if we choose the same starting letter. If we expand letters from two different stretched transition chains, we meet a similar problem. So, such a derivation cannot produce a periodic word. On the other hand, applying relations coming from B_2 will never produce a periodic word either; since these expand to both sides of the letter simultaneously and the new subword produced is everything in the support minus one other letter we can see that the same issues will arise. □

Chapter 6

Infinite diagrams

Recall that in Definition 2.3.15 we introduced infinite diagrams over $\mathcal{P}$ as formal concatenations of infinitely many diagrams from $\mathcal{D}(\mathcal{P})$. In this chapter we push this idea further, primarily with the purpose of establishing overgroups of the groups we are interested in which we can use in future analysis.

Returning to the idea of diagrams as two-dimensional analogues of words, so far infinity in a diagram is only defined along one dimension. Our next step is to formalise infinity in a diagram along the orthogonal dimension.

6.1 Diagrams with infinite base

Put simply, the idea of diagrams with infinite base, which we will soon define formally, is to generalise the forest diagrams defined by Belk and Brown [5] for representing elements of F; diagrams with infinite base are two forest diagrams as tree diagrams in general are to the classical binary tree pair representation of F [14].

For a letter $x \in \Sigma$ the trivial diagram ε_x consists of an initial and a terminal vertex with an edge labelled by x between them. Then, of course, for any word $w = x_1 \dots x_n \in \Sigma^+$ we have $\varepsilon_w = \varepsilon_{x_1} + \dots + \varepsilon_{x_n}$. This motivates the following definition.

Definition 6.1.1. Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a semigroup presentation and let γ be a cordinal number. We define a **formal sum of trivial diagrams** to be a sequence $(\varepsilon_{x_\alpha})_{\alpha\in\gamma}$ indexed by γ for $x_\alpha \in \Sigma$. If $u_\alpha \in \Sigma^+$ then we understand the sequence $(\varepsilon_{u_\alpha})_{\alpha\in\gamma}$ to be the formal concatenation $((\varepsilon_{x_i})_{i<n_\alpha})_{\alpha\in\gamma}$ where $u_\alpha = x_0 \dots x_{n_\alpha - 1}$. We also define the **sum of two formal sums of trivial diagrams** $(\varepsilon_{x_\alpha})_{\alpha\in\gamma} + (\varepsilon_{x\alpha})_{\alpha\in\delta}$ to be the formal sum of trivial diagrams $(\varepsilon_\alpha)_{\alpha\in\zeta}$ where $\zeta = \gamma + \delta$. An **embedding** of a formal sum of trivial diagrams is a plane graph with precisely one maximal positive path, which is labelled by $(x_\alpha)_{\alpha\in\gamma}$. We consider two formal sums of trivial diagrams $(\varepsilon_{x_\alpha})_{\alpha\in\gamma}, (\varepsilon_{x_\alpha})_{\alpha\in\delta}$ to be **isotopic** if there exists an isotopy of $\mathbb{R}^2$ between embeddings

of them. We will often denote a particular formal sum of trivial diagrams $(\varepsilon_{u_\alpha})_{\alpha\in\gamma}$ by $\Sigma_{\alpha\in\gamma}\varepsilon_{u_\alpha}$ and refer to its isotopy class of formal sums of trivial diagrams as a **transfinite trivial diagram**.

Remark. As has been mentioned previously, a trivial diagram is essentially a word. Thus, this translates to an analogous theory of words indexed by cordinals. Theories of transfinite words i.e. words indexed by ordinals are well established e.g. see [8].

Having a formal definition for infinite trivial diagrams now allows us to define infinite atomic diagrams.

Definition 6.1.2. Let $\Sigma_{\alpha\in\gamma}\varepsilon_{x_\alpha}, \Sigma_{\alpha\in\delta}\varepsilon_{x_\alpha}$ be transfinite trivial diagrams and let $\Delta \in \mathcal{D}(\mathcal{P})$ be an elementary diagram. We may embed the formal sum

$$\Psi := \Sigma_{\alpha\in\gamma}\varepsilon_{x_\alpha} + \Delta + \Sigma_{\alpha\in\delta}\varepsilon_{x_\alpha}$$

as a plane graph so that it has precisely two maximal positive paths, which are the top and bottom paths, and we refer to its isotopy class as a **transfinite atomic diagram**. Formally, the top path and bottom path of Ψ are the transfinite trivial diagrams $\lceil\Psi\rceil := \Sigma_{\alpha\in\gamma}\varepsilon_{x_\alpha} + \lceil\Delta\rceil + \Sigma_{\alpha\in\delta}\varepsilon_{x_\alpha}$ and $\lfloor\Psi\rfloor := \Sigma_{\alpha\in\gamma}\varepsilon_{x_\alpha} + \lfloor\Delta\rfloor + \Sigma_{\alpha\in\delta}\varepsilon_{x_\alpha}$ respectively.

We may then extend concatenation to transfinite atomic diagrams. In general, this is slightly more complicated than the finite base case. More specifically, given two finite total orders there is at most one order isomorphism between them, meaning that if two finite trivial diagrams have the same label there is only one way to identify them. However, there may be (infinitely) many order isomorphisms between two given cordinals, and so in general we will need to specify by which isomorphism we are identifying two transfinite trivial diagrams.

Definition 6.1.3. Let Ψ_1, Ψ_2 be transfinite atomic diagrams and let $f : \lfloor\Psi_1\rfloor \to \lceil\Psi_2\rceil$ be a label-preserving order isomorphism. We define the concatenation $\Psi_1 \circ_f \Psi_2$ with respect to f to be the isotopy class of the plane graph formed from embeddings of Ψ_1, Ψ_2 by identifying e with $f(e)$ for every $e \in \lfloor\Psi_1\rfloor$ via a homotopy. We have $\lceil\Psi_1 \circ_f \Psi_2\rceil = \lceil\Psi_1\rceil$ and $\lfloor\Psi_1 \circ_f \Psi_2\rfloor = \lfloor\Psi_2\rfloor$. This naturally extends to finite sequences of concatenation $\Psi_1 \circ_{f_1} \ldots \circ_{f_n} \Psi_n$.

This ambiguity, as it stands, would prevent us from defining a group structure on this class of objects, even with respect to a fixed base. In order to overcome this we augment these diagrams with additional data which specifies which automorphism we should concatenate by.

Definition 6.1.4. Let $\Sigma_{\alpha\in\gamma}\varepsilon_{x_\alpha}$ be a transfinite trivial diagram and let $\zeta_1, \ldots, \zeta_n$ be the ζ summands of the cordinal number γ defined in Definition 2.1.2. Suppose the labellings of ζ_i are periodic (relabel summands if necessary) and choose a period w_i and tiling t_i for each. We define an **augmentation** of $\Sigma_{\alpha\in\gamma}\varepsilon_{x_\alpha}$ by distinguishing two n-tuples $(o_1,\dots,o_n),(b_1,\dots,b_n)$ of edges such that o_i, b_i are edges of the subdiagram equal to $\Sigma_{\alpha\in\zeta_i}\varepsilon_{x_\alpha}$. We consider two augmentations $(o_1,\dots,o_n),(b_1,\dots,b_n)$ and $(o'_1,\dots,o'_n),(b'_1,\dots,b'_n)$ to be **equivalent** if either;

(a) $o_i, o'_i \in t_i^{-1}(w_{i,j_1})$ and $b_i, b'_i \in t_i^{-1}(w_{i,j_2})$ for all i where $t_i^{-1}(w_{i,j_1}), t_i^{-1}(w_{i,j_2})$ are tiles of ζ_i, or;

(b) there exist non-converse label-preserving automorphisms $f_i : \zeta_i \to \zeta_i$ such that $f_i(o_i) = o'_i$ and $f_i(b_i) = b'_i$ for all i.

We refer to an object obtained as such as an **augmented trivial diagram**.

Similarly, we define an augmentation of an atomic diagram Ψ by distinguishing two n-tuples $(o_1,\dots,o_n),(b_1,\dots,b_n)$ of edges such that $o_i \in \lceil\Psi\rceil$ and $b_i \in \lfloor\Psi\rfloor$ for all i, and we consider two augmentations of Ψ to be equivalent by condition (1) above with the caveat that $\lceil\Psi\rceil$ and $\lfloor\Psi\rfloor$ may have different labels. We call these objects **augmented atomic diagrams**.

We now adapt Definition 6.1.3 to reflect this change.

Definition 6.1.5. Let Ψ_1, Ψ_2 be augmented atomic diagrams such that $\lfloor\Psi_1\rfloor$ and $\lceil\Psi_2\rceil$ have the same label and let $\Sigma_{\alpha\in\gamma}\varepsilon_{x_\alpha}$ be the (non-augmented) transfinite trivial diagram with that label. If $(b_1,\dots,b_n)$ and $(b'_1,\dots,b'_n)$ are the distinguished tuples of $\lfloor\Psi_1\rfloor$ and $\lceil\Psi_2\rceil$ respectively then define $f : \lfloor\Psi_1\rfloor \to \lceil\Psi_2\rceil$ to be the unique non-converse label-preserving order isomorphism defined by $t_i^{-1}(w_{i,j_1}) \mapsto t_i^{-1}(w_{i,j_2})$ for each i where $b_i \in t_i^{-1}(w_{i,j_1})$ and $b'_i \in t_i^{-1}(w_{i,j_2})$ for tiles $t_i^{-1}(w_{i,j_1}), t_i^{-1}(w_{i,j_2})$ where t_i is a tiling with period w_i for the zeta summand ζ_i of γ.

We thus define the **concatenation** $\Psi_1 \circ \Psi_2$ of augmented atomic diagrams by

$$\Psi_1 \circ \Psi_2 := \overline{\Psi}_1 \circ_f \overline{\Psi}_2$$

where $\overline{\Psi}_i$ is the non-augmented counterpart of Ψ_i and $\circ_f$ is defined as in Definition 6.1.3. As before, this can be extended to finite sequences of concatenations $\Psi_1 \circ \ldots \circ \Psi_n$.

Remark. Defining this concatenation and equivalence between augmentations in general requires us to do so with respect to a particular period and tiling. However, (following the discussion around periodic labellings in Chapter 2) we can see that the objects - and, in what follows, the group structure - defined are actually independent of this choice. Indeed, when working with such augmentations we will never make explicit reference to a period or tiling but leave this implicit; usually the diagrams we consider will have at most one ζ summand and we specify an augmentation by way of 'pointers' *a la* Belk and Brown [5].

Finally, we have;

Definition 6.1.6. We say a plane graph Δ is an **augmented diagram** if it is equal to an augmented trivial diagram or a finite concatenation of augmented atomic diagrams. We denote the set of augmented diagrams over a semigroup presentation $\mathcal{P}$ by $\mathcal{D}^*(\mathcal{P})$. If $\mathcal{E} = \Sigma_{\alpha\in\gamma}\varepsilon_{x_\alpha}$ is a (non-augmented) transfinite trivial diagram then we denote the set of augmented diagrams over $\mathcal{P}$ whose top and bottom paths are equal to $\mathcal{E}$ by $\mathcal{D}^*(\mathcal{P}, (x_\alpha)_{\alpha\in\gamma})$.

Remark. Notice that, in many cases, a given diagram will have only one augmentation. This is equivalent to no augmentation, since there is then only one non-converse label-preserving automorphism of the top (bottom) path.

Notice that the concepts of dipoles, dipole reduction and dipole equivalence $\equiv$ (Definition 2.3.5) readily carry over to augmented diagrams and so we have everything we need to define a group structure on this set of objects.

Definition 6.1.7. Let $\mathcal{P}$ be a semigroup presentation. We define the (augmented) diagram group to be $D(\mathcal{P}, (x_\alpha)_{\alpha\in\gamma}) := (\mathcal{D}^*(\mathcal{P}, (x_\alpha)_{\alpha\in\gamma}), \circ)$.

Remark. Notice that, in general, an augmented trivial diagram may represent a non-trivial group element. For example, consider $\Sigma_{\alpha\in(\omega^{-1}+\omega)}\varepsilon_w$ for some $w \in \Sigma^+$. An (equivalence class of) augmentation(s) is defined by distinguishing two (not necessarily distinct) ε_w in $\Sigma_{\alpha\in(\omega^{-1}+\omega)}\varepsilon_w$ for some $\alpha_1 < \alpha_2$ such that $\alpha_2 - \alpha_1$ is constant. Notice that if we compose the equivalence class such that $\alpha_2 - \alpha_1 = 1$ with itself we obtain the equivalence class such that $\alpha_2 - \alpha_1 = 2$. If we compose it with itself $n-1$ times we obtain the equivalence class such that $\alpha_2 - \alpha_1 = n$. Every augmentation can be obtained as such and so we see that the subgroup of $D(w^{(\omega^{-1}+\omega)})$ consisting of augmented trivial diagrams is the infinite cyclic group (and, in general, if the cordinal number γ has n ζ-summands then the subgroup of all augmented trivial diagrams is the free Abelian group of rank n). The identity element of $D(w^{(\omega^{-1}+\omega)})$ is, therefore, the augmented trivial diagram where the distinguished tiles coincide (which we may call the *trivial* trivial diagram, if one likes).

Remark. Augmented diagrams are not really different to diagrams; they are just diagrams with extra data specifying how to concatenate them when extra data is necessary. Augmented diagram groups, then, are just a generalisation of the typical diagram group to allow transfinite bases. As such, there is not really a distinction to be made, and we will simply refer augmented diagrams as diagrams and augmented diagram groups as diagram groups (understanding that, when a diagram group has base containing a periodically labelled ζ summand, we will need to use augmented diagrams for things to make sense).

Although (to the best of the author's knowledge) these diagram groups have not been defined in the literature, some important examples have appeared. Most notably, the strand diagram group corresponding to $D(\langle x \mid x^2 = x\rangle, x^{(\omega^{-1}+\omega)})$ is precisely the forest diagram representation of Thompson's group F studied in [5].

6.2 Diagrams with infinitely many cells

Having established a precise definition of transfinite bases for diagrams and diagram groups we would like to define concatenation on diagrams with infinitely many cells to naturally extend the finite case. But we will quickly see that we are still missing some pieces of the puzzle.

Basically, the issue lies in defining the bottom path of an infinite diagram. So what is the bottom path of a diagram? In the finite case, if S is the strand diagram of a diagram Δ, the bottom path is the totally ordered set of labelled termini of all maximal positive paths of S. However, in the infinite case, there will be at least one maximal positive path which does not terminate.

Let's consider the semigroup presentation obtained from a fast pair of bumps generating F, namely $\mathcal{P} = \langle A, B, C, D \mid A = AB, B = BC, C = BC, D = CD\rangle$, and consider the infinite diagram

$$\Delta := [\varepsilon_A + (B, BC) + \varepsilon_{CD}] \circ [\varepsilon_{AB} + (C, BC) + \varepsilon_{CD}] \circ \ldots \circ [\varepsilon_{AB^n} + (C, BC) + \varepsilon_{CD}] \circ \ldots$$

over $\mathcal{P}$. We can see that the top path is finite, and has label $ABCD$, but the bottom path is infinite. We would like there to be a transfinite trivial diagram which we can identify with this bottom path; a natural candidate is $\varepsilon_{AB^\omega CD}$. Indeed, we can see that there is a label-preserving order isomorphism between the set of termini of maximal positive paths in S and the edges of $\varepsilon_{AB^\omega CD}$, where S is the strand diagram of Δ. However, there is a maximal positive path of S which does not terminate (its path label is BC^ω) which is to the left of the path with terminus labelled C and to the right of every path with terminus labelled B. We can see that this 'gap' will produce a point of discontinuity when attempting to identify this 'bottom path' with $\varepsilon_{AB^\omega CD}$.

So, we need some way to 'plug' this 'gap'. To this end, we will need to introduce a new type of cell which will attached to the bottom of an infinite sequence of cells. Attaching this cell will require an additional operation after countably many have been added; in other words, we need to allow our concatenations to have length greater than ω. However, this causes our idea of equality of infinite diagrams to break, as demonstrated in the following example.

Consider the semigroup presentation $\mathcal{P}$ above and define infinite diagrams

$$\Delta(X) = (X, BC) \circ [(B, BC) + \varepsilon_C] \circ \ldots \circ [(B, BC) + \varepsilon_{C^n}] \circ \ldots$$

and

$$\Psi(X) = (X, BC) \circ [\varepsilon_B + (C, BC)] \circ \ldots \circ [\varepsilon_{B^n} + (C, BC)] \circ \ldots$$

over $\mathcal{P}$ for $X \in \{B, C\}$. These are both represented by length ω formal concatenations. We may then consider the diagram

$$\Delta(B)\circ[\varepsilon_{BC^{\omega-1}}+\Delta(C)]\circ[\varepsilon_{BC^{\omega-1}}+\Delta(C)+\varepsilon_{BC^{\omega-1}}]\circ\ldots\circ[\varepsilon_{BC^{\omega-1}}+\Delta(C)+\varepsilon_{(BC^{\omega-1})^n}]\circ\ldots$$

that is, we adjoin a copy of $\Delta(C)$ to every leaf of $\Delta(B)$ labelled by C. Repeating this ω times we obtain a formal concatenation with length ω^{ω^ω} defining a diagram which we call Δ. Now, starting with $\Psi(B)$ and repeatedly adjoining copies of $\Psi(B)$ to each leaf indefinitely, just as with Δ, we obtain another length ω^{ω^ω} formal concatenation called Ψ.

Notice that Δ and Ψ ought to be equal as diagrams. However, no prefix of the concatenation defining Δ contains $\Psi(B)$ as a prefix since the nth term of $\Psi(B)$ only appears as the $\omega^\omega.(n-1)+1$th term in Δ. This means that, as well as introducing a new type of cell, we will have to use a different notion of equality. To do this, we now restrict our attention to diagram groups over tree-like presentations (since this is sufficient our the purposes).

6.3 Transfinite tree-like diagrams

Take a tree-like semigroup presentation $\mathcal{P}$, a word w over $\mathcal{P}$ and consider its corresponding master diagram $\mathcal{M}(\mathcal{P}, w) = (\Delta_n)_n$ which is defined as an increasing (in the prefix relation) sequence of finite diagrams. Now that we have defined transfinite bases we would like to extend $\mathcal{M}$ to a transfinite sequence $(\Delta_\alpha)_\alpha$.

To this end, recall Definition 2.3.11 which defines the master diagram as a locally ordered labelled tree $M = (E_+, P, L_0, (L_\pi)_{\pi\in\mathcal{C}})$. We now extend this definition.

Definition 6.3.1. Let $\mathcal{P} = \langle \Sigma \mid \mathcal{R} \rangle$ be a tree-like semigroup presentation. Given a word $w \in \Sigma^+$ we understand it to be a labelled totally ordered set and define $M_0 := w$. For an ordinal α and a limit ordinal β, we denote $M_\alpha := (E_\alpha, P_\alpha, L_0, (L_c)_c)$ and inductively define:

- $E_{\alpha+1}$ to be the disjoint union of E_α and a labelled totally ordered set $(w_{l(x)}, L_x)$ defined by the relation $(l(x), w_{l(x)}) \in \mathcal{R}$ for each $x \in \text{end}(M_\alpha)$. Additionally, we add pairs (x,y) for $y \in w_{l(x)}$ for each $x \in \mathrm{end}(M_\alpha)$ to P_α and take the transitive closure to define $P_{\alpha+1}$. Then $M_{\alpha+1} := (E_{\alpha+1}, P_{\alpha+1}, L_0, (L_c)_c, (L_x)_{x\in \mathrm{end}\, M_\alpha})$.

- E_β to be the disjoint union of $\bigcup_{\alpha<\beta} E_\alpha$ and a singleton $\{x\}$ labelled by $a \in \Sigma$ for every maximal chain C_x of $(\bigcup_{\alpha<\beta} M_\alpha, \bigcup_{\alpha<\beta} P_\alpha)$ such that C_x has an infinite suffix S_x whose every element is labelled by a, for each $a \in \Sigma$. We define P_β as the transitive closure of the disjoint union of $\bigcup_{\alpha<\beta} P_\alpha$ and (s,x) for each added singleton $\{x\}$ and $s \in S_x$. Then $M_\beta := (E_\beta, P_\beta, L_0, (L_c)_c)$.

We refer to M_α as the **transfinite master diagram of height** $\alpha+1$ and $\bigcup_{\alpha<\beta} M_\alpha$ as the **transfinite master diagram of height** β.

This lets us define:

Definition 6.3.2. Let M_α be a transfinite master diagram over some tree-like semigroup presentation $\mathcal{P}$ and top path label w. Suppose D is an initial subtree of M_α such that if $x \in D$ where $\mathrm{is}_D(x) \neq \emptyset$ then $\mathrm{is}_D(x) = \mathrm{is}_{M_\alpha}(x)$. Then we say D is a **positive diagram** over $\mathcal{P}$ with top path label w, and we denote its top path by $\lceil D \rceil$.

This definition lets us recover, and extend, a suitable idea of equality for our infinite diagrams - two positive diagrams are equal exactly when they are relation isomorphic (to the same initial subtree of the master diagram).

This brings us to our other problem - the bottom path. As we discussed, in order to concatenate a diagram we need its bottom path to be 'complete' - that is, it should not have any 'gaps'.

Definition 6.3.3. Let D be a positive diagram with top path label w over $\mathcal{P}$. We say D is **well-formed** if for every maximal chain C of D we have $C \cap \mathrm{end}(D) \neq \emptyset$. In other words, every maximal chain has a maximal element. In this case we refer to the totally ordered labelled set $(\mathrm{end}(D), \mathrm{Lex}_M |_{\mathrm{end}(D)})$ as the **bottom path** of D and denote it by $\lfloor D \rfloor$. We denote by $\mathcal{D}^+_\beta(w)$ the set of well-formed positive diagrams over $\mathcal{P}$ with top path w of height less than β for a limit ordinal β, and by $\mathcal{D}^+_\beta(w)$ the subset of diagrams with top path.

Remark. Notice, then, that no positive diagram of limit ordinal height is well-formed, since it must have a maximal chain with no maximal element by definition.

As with Definition 2.3.11 we may concatenate two transfinite positive diagrams Δ_1, Δ_2 such that $\lfloor \Delta_1 \rfloor$ and $\lceil \Delta_2 \rceil$ have the same label. As discussed in a previous section, when this common label is transfinite there may be multiple ways to concatenate and we may specify this by way of augmenting Δ_1, Δ_2 as in Definition 6.1.4.

At last, we can define the infinite diagram groups we seek.

Definition 6.3.4. Let α be an ordinal, $\mathcal{P}$ be a tree-like semigroup presentation and consider the set $\lfloor\mathcal{D}^*_\beta\rfloor(w)$ of pairs (Δ_1, Δ_2) of positive (augmented) diagrams over $\mathcal{P}$ with top path label w such that $\lfloor\Delta_1\rfloor$ and $\lfloor\Delta_2\rfloor$ have the same label, where we define $\beta := \omega^\alpha$. We define a partial operation $\circ$, extending that of Definition 2.3.9, by

$$(\Delta_1, \Delta_2) \circ (\Delta_2, \Delta_3)$$

as well as defining an equivalence relation $\equiv$ on $\lfloor\mathcal{D}^*_\beta\rfloor(w)$ by taking the reflexive, symmetric, transitive closure of

$$(\Delta_1, \Delta_2) \equiv (\Delta_1 \circ \psi, \Delta_2 \circ \psi)$$

for a positive atomic diagram ψ over $\mathcal{P}$. The partial operation $\circ$ induces a binary operation on $\mathbf{D}_\alpha(\mathcal{P}, \mathbf{w}) = \lfloor\mathcal{D}^*_\beta\rfloor / \equiv$ so that $(D_\alpha(\mathcal{P}, w), \circ)$ forms a group called the **αth tree diagram group** over $\mathcal{P}$ with base w.

Of course, we should establish that this is, indeed, a group. We also show that a term in the sequence of αth tree diagram groups over $\mathcal{P}$ with base w as α increases contains every previous term as a normal subgroup.

Proposition 6.3.5. *The pair $(D_\alpha(\mathcal{P}, w), \circ)$ from Definition 6.3.4 forms a group.*

Proof. Just as with the finite tree diagram group we can see that the operation is associative, has a trivial element $(\varepsilon_w, \varepsilon_w)$ and inverses $(\Delta_1, \Delta_2)^{-1} = (\Delta_2, \Delta_1)$.

To see that $D_\alpha(w)$ is closed under $\circ$ notice that, in general, if we have two reduced pairs of well-formed trees $(\Delta_1, \Delta_2), (\Delta_3, \Delta_4)$ with height at most β, δ respectively then their (reduced) product is a pair of trees with height at most $\beta + \delta$. This is because, in the worst case, Δ_2 and Δ_3 are disjoint and the suffix added to the pair (Δ_1, Δ_2) is adjoined at a point $x \in \text{end}(\Delta_1)$ such that $\rho(x)$ is maximal. Therefore, the ordinal γ which (properly) bounds the height that the trees may take must satisfy the property that $\beta + \delta < \gamma$ if $\beta, \delta < \gamma$. This property of γ is known as additive indecomposability and the ordinals satisfying it are known as gamma numbers, which are precisely the ordinals of the form ω^α for an ordinal α.

In spirit, we now want to show that products are well-formed. Strictly speaking, since we defined our group on the set of well formed diagrams, this is already satisfied. With this framing, then, what we need to show is that the operation is binary. So, let ψ, θ be such that $\Delta_2 \circ \psi = \Delta_3 \circ \theta$ is the union of Δ_2 and Δ_3 in $M := \bigcup_{\beta<\omega^\alpha} M_\beta(\mathcal{P}, w)$ and consider $\Delta_1 \circ \psi$ and $\Delta_4 \circ \theta$. Since Δ_3 is well-formed and ψ is precisely those edges of Δ_3 which are not edges of Δ_2 in M we can see that ψ is also a well-formed diagram. The same argument shows that θ is well-formed too, and since the definition of well-formed is only concerned with the ends of a diagram we can see that $\Delta_1 \circ \psi$ and $\Delta_4 \circ \theta$ are also well-formed, and so the pair $(\Delta_1 \circ \psi, \Delta_4 \circ \theta)$ represents an element of $D_\alpha(w)$. □

Proposition 6.3.6. *Let $\alpha < \beta$ be ordinals. Then $D_\alpha(\mathcal{P}, w)$ is a normal subgroup of $D_\beta(\mathcal{P}, w)$.*

Proof. We want to show that

$$(\Delta_2^\beta, \Delta_1^\beta) \circ (\Delta_1^\alpha, \Delta_2^\alpha) \circ (\Delta_1^\beta, \Delta_2^\beta) \in D_\alpha$$

if $(\Delta_1^\alpha, \Delta_2^\alpha) \in D_\alpha$ and $(\Delta_1^\beta, \Delta_2^\beta) \in D_\beta$ are reduced. In order to evaluate the first product we must consider equivalent pairs such that Δ_1^β and Δ_1^α are replaced with equal diagrams. If Ψ is the union of $\Delta_1^\beta, \Delta_1^\alpha$ in $M := \bigcup_{\gamma<\omega^\beta} M_\gamma(\mathcal{P}, w)$ then let ψ be the diagram satisfying $\Psi = \Delta_1^\beta \circ \psi$. Notice that each edge e of ψ as a subdiagram of Ψ satisfies $\rho_M(e) < \omega^\alpha$ since ψ must also be a subdiagram of Δ_1^α. Thus we have

$$(\Delta_2^\beta \circ \psi, \Psi) \circ (\Psi, \Delta_2^\alpha \circ \theta) \circ (\Delta_1^\beta, \Delta_2^\beta) = (\Delta_2^\beta \circ \psi, \Delta_2^\alpha \circ \theta) \circ (\Delta_1^\beta, \Delta_2^\beta)$$

where θ satisfies $\Psi = \Delta_1^\alpha \circ \theta$. Notice that θ is also a suffix of Δ_1^β and, in particular, every edge e of Δ_1^β such that $\rho_M(e) > \omega^\alpha$ is an edge of the suffix θ since Δ_1^α has height at most ω^α; on the other hand, every edge e of θ with $\rho_M(e) > \omega^\alpha$ is an edge of Δ_1^β. Repeating this for the remaining product we obtain

$$(\Delta_2^\beta \circ \psi \circ \lambda, \Theta) \circ (\Theta, \Delta_2^\beta \circ \sigma) = (\Delta_2^\beta \circ \psi \circ \lambda, \Delta_2^\beta \circ \sigma)$$

where Θ is the union of $\Delta_2^\alpha \circ \theta$ and Δ_1^β, and λ, σ are the diagrams satisfying $\Theta = \Delta_2^\alpha \circ \theta \circ \lambda = \Delta_1^\beta \circ \sigma$. Since e such that $\rho_M(e) > \omega^\alpha$ is an edge of θ if and only if it is an edge of Δ_1^β we can see that if e is an edge of λ or σ as a suffix of Θ then $\rho_M(e) < \omega^\alpha$. Thus, considering the tree diagram we are left with

$$(A, B) := (\Delta_2^\beta \circ \psi \circ \lambda, \Delta_2^\beta \circ \sigma)$$

we can see that an edge e such that $\rho_M(e) > \omega^\alpha$ is an edge of either A or B if and only if it is an edge of Δ_2^β, and it then follows that such an e is an edge of A if and only if it is an edge of B. Neither A nor B contain edges at level greater than ω^β so their common subdiagram consisting of all edges at level greater than ω^α must be a suffix of them both. Thus the reduced form of (A, B) consists of two positive diagrams of height less than ω^α, as required. □

Chapter 7

Properties of fast groups

We now investigate fast groups in more detail. Specifically, we develop a method for deducing information on the relations of a fast group, via a combinatorial argument on the diagrams which make up its diagram group representation. This method has potential scope for further work then is completed in this thesis, in particular in concert with the ideas developed in Chapter 6. For now, we reach the conclusion that the relations of a fast group are all consequences of commutators (Theorem 7.1.14). This falls in line with strong results regarding commutators in diagram groups more generally. Specifically, a classical result due to Guba and Sapir [29] states that the abelianisation of a diagram group is always free, which is to say that when we factor a diagram group by its commutator subgroup we obtain a free abelian group. To the author's knowledge, neither of these results imply the other in themselves.

7.1 Group presentations

As previously noted, diagram groups are, in an important sense, a natural generalisation of free groups. Indeed, we can obtain any free group of rank n as a diagram group $D(\mathcal{P}_n, x_1)$ over the semigroup presentation

$$\mathcal{P}_n = \langle x_1, \ldots, x_n, y_1, \ldots, y_n, z_1, \ldots, z_n \mid x_i = y_i, y_i = z_i, z_i = x_i, x_i = x_{i+1} \rangle$$

and $D(\mathcal{P}_\infty, x_1)$ is isomorphic to the countable rank free group where $\mathcal{P}_\infty = \bigcup_{n=1}^{\infty} \mathcal{P}_n$. We can also obtain any rank n free abelian group as $D(\mathcal{P}_1, x_1^n)$, and the countable rank free abelian group is $D(\mathcal{P}_1, x_1^\omega)$.

In each of these cases, and in diagram groups generally, we may represent the generators as reduced diagrams over the relevant presentation. When we form a word w on these generators we can represent this by concatenating the corresponding diagrams in the same order - the diagram Δ obtained may have dipoles which can then be reduced iteratively to reach a reduced form Δ_r, which is the canonical representative for this element of the diagram group. Indeed, any word u on the generators represented by diagrams in this way will reduce to Δ_r if and only if u and w are equal in the group. As such, given a reduced diagram Δ_r, if we introduce dipoles to find an equivalent diagram Δ_e such that Δ_e can be decomposed into generator diagrams then all words on the generators represented by Δ_e are equal in the group and, indeed, equal to the words obtained in the same way from any other equivalent diagram Δ'_e. Doing this for all diagrams equivalent to Δ_r (modulo free reduction of generators) gives us all the (reduced) words on the generators which represent the same group element as Δ_r. Finally, performing this procedure for all reduced diagrams over the relevant presentation with the relevant base gives us all the relations of the group, and we therefore obtain a presentation. We call this set of relations $\mathcal{D}$.

Of course, there are many problems with this in general. We may be able to introduce dipoles in a non-trivial way indefinitely, and it will often be difficult to recognise when a diagram is decomposable. Even besides that, this presentation is a priori much too large to be useful. We will find, however, that these problems can be overcome when we restrict our attention to bumpy diagram groups $D(\mathcal{P}, a_1 \dots a_n)$ over bumpy presentations $\mathcal{P}$, which are isomorphic to groups generated by a fast set of bumps B on the interval by Theorem 4.2.7.

Definition 7.1.1. Let $D(\mathcal{P}, a_1 \dots a_n)$ be a bumpy diagram group. Given a spherical diagram Δ with this base we say it is **supported on a subword** $u = a_i \dots a_j$ of $a_1 \dots a_n$ if $\Delta = \varepsilon_{a_1 \dots a_{i-1}} + \Delta' + \varepsilon_{a_{j+1} \dots a_n}$ where Δ' is a $(a_i \dots a_j, a_i \dots a_j)$-diagram. We say a word w on B is supported on u if $w\delta$ is supported on u. For a word $w = a_1 \dots a_k$ with $a_i \in B$ we call the diagram $\Delta_w = a_1\delta \circ \dots \circ a_k\delta$ the **unreduced diagram of** w. We say a diagram Δ over $\mathcal{P}$ is **B-decomposable** if there exists a word $w \in B^+$ such that $\Delta = \Delta_w$.

Remark. We denote $G = \langle B \rangle$ and abuse notation slightly by considering $w\delta$ to be the reduced form of Δ_w (this is coherent enough since the reduced form of Δ_w is the canonical normal form of the equivalence class $w\delta$).

We have the following presentation

$$G = \langle B \mid \mathcal{D} \rangle$$

where $\mathcal{D} = \{r_\Delta \mid \Delta \text{ reduced}\}$ for r_Δ the set of relations (w_1, w_2) whenever $\Delta_1 \equiv \Delta_2 \equiv \Delta$ and Δ_1, Δ_2 are B-decomposable such that $\Delta_1 = \Delta_{w_1}$ and $\Delta_2 = \Delta_{w_2}$.

We now define a procedure.

Definition 7.1.2. Let Δ be a reduced diagram and let $\pi_1, \dots, \pi_k$ be the list of cells whose bottom path is a subpath of $\lfloor \Delta \rfloor$. We say a subdiagram Γ of Δ is a **choice**

of Δ if either

- Γ is a generator diagram consisting of π_i and π_{i+1} for some i, or;
- Γ is equal to π_i for some i such that π_i does not form a generator diagram with any other cell of Δ.

and we denote the set of such subdiagrams by $C(\Delta_0)$.

Remark. Notice that a choice always corresponds to a unique generator diagram β_b for $b \in B^{\pm}$.

Definition 7.1.3. Let Δ be a reduced diagram with set of choices $C(\Delta)$, and let w be some word on B. We call the pair (Δ, w) a **step** in our procedure. Let $c \in C(\Delta)$ with corresponding generator diagram β_b and define $\Delta(c) = (\Delta \circ \beta_b^{-1})^r$ and redefine $w := wb$. We refer to the pair $(\Delta(c), wb)$ as a **next step** of our procedure. If a pair is a next step of a next step and so on, we say it is a **subsequent step**.

Remark. We extend this notation so that $\Delta(c_2, c_1) = \Delta(c_1)(c_2)$ for $c_1 \in C(\Delta)$ and $c_2 \in C(\Delta(c_1))$ and, in general, $\Delta(c_k, \ldots, c_1) = \Delta(c_{k-1}, \ldots, c_1)(c_k)$ for $c_k \in C(\Delta(c_{k-1}, \ldots, c_1))$.

Suppose we have a diagram Δ, a choice $c \in C(\Delta)$ and a word $w := \varepsilon$; we obtain the next step $(\Delta(c), b)$ and the subdiagram c forms dipoles in $\Delta \circ \beta^{-1}$ and so is not a subdiagram of $\Delta(c)$. We may then find a next step of $\Delta(c)$, choosing a subdiagram in $C(\Delta(c))$ corresponding to a bump b. At each step we redefine $w := bw$ and notice that $\Delta(c_k, \ldots, c_1 = c)$ either has fewer cells than $\Delta(c_{k-1}, \ldots, c_1 = c)$ or has fewer cells which do not immediately form a generator diagram by Lemma 4.2.5. As such, we will eventually reach $\Delta(c_m, \ldots, c_1) = \varepsilon$ whence this procedure will terminate and we have obtained a length m word w such that $\Delta_w \equiv \Delta_0 = w\delta$.

Since we assume Δ is reduced, every subdiagram c equal to a generator diagram will be preserved in every subsequent diagram in the procedure until c itself is chosen, after which it will not be a subdiagram in any subsequent step. Similarly, negative cells will persist until they are chosen, and positive cells will persist until its corresponding negative cell is chosen, whence it forms part a generator diagram. Thus the cells of any given reduced diagram Δ are partitioned by its set of future choices (understanding choices as sets of either one or two cells) and the set of sequences s of choices such that $\Delta(s) = \varepsilon$ is in one to one correspondence with $W(\Delta)$. We note that the natural partial order on the cells of a given reduced diagram quotients to a partial order on its future choices.

Notice that, as suggested by the name, at a given step in the procedure we may have $|C(\Delta_j)| > 1$, in which case the word obtained by the procedure is not unique; it depends on the sequence of choices made. If $w_1, \ldots, w_t$ are all words obtained from Δ_0 via this procedure then $\Delta_{w_1} \equiv \ldots \equiv \Delta_{w_t} \equiv \Delta_0$ and therefore $w_1 \equiv_G \ldots \equiv_G w_t$.

Definition 7.1.4. Let Δ be reduced and consider the step (Δ, ε). Let $c_1, \ldots, c_m$ be a sequence such that $c_{i+1} \in C(\Delta(c_i \ldots c_1))$ for all i and $\Delta(c_m, \ldots, c_1) = \varepsilon_{a_1 \ldots a_n}$. If $b_i \in B$ corresponds to c_i then we have $(\varepsilon_{a_1 \ldots a_n}, b_1 \ldots b_m)$ as a subsequent step of (Δ, ε). We denote $W(\Delta) = \{w \mid (\varepsilon_{a_1 \ldots a_n}, w) \text{ is a subsequent step of } (\Delta, \varepsilon)\}$. In addition, we define

$$R_\Delta = \{(u, w) \mid u, w \in W(\Delta)\}$$

and $\mathcal{R} = \{R_\Delta \mid \Delta \text{ reduced}\}$.

We claim that this subset of $\mathcal{D}$ is sufficient to describe all the relations of G.

Lemma 7.1.5. *Let $w \in B^+$ be freely reduced and consider Δ_w. If there is a sequence of dipole reductions to obtain a diagram Δ from Δ_w containing a pair of dipoles together forming two subdiagrams, one equal to β and the other equal to β^{-1} where $\beta = b\delta$ for some $b \in B$, then there exists a word u and a relation $(u, v) \in \langle\langle \mathcal{R} \rangle\rangle$ such that $w = w'uw''$ and $w'vw''$ contains either bb^{-1} or $b^{-1}b$ as a subword.*

Proof. Suppose Δ_w has a pair of dipoles with the stated property, formed by cells π_1, π_2 and ρ_1, ρ_2 respectively. If the two dipoles in fact form a subdiagram equal to either $\beta \circ \beta^{-1}$ or $\beta^{-1} \circ \beta$ (that is, π_1 and ρ_1 form one generator diagram while π_2 and ρ_2 form the other) then there is a subword u such that $w = w'b^{-1}ubw''$ where the letter b^{-1} corresponds to the generator diagram β^{-1} formed by π_1 and ρ_1 while b corresponds to the other. Since the bottom path of β^{-1} is a subpath of $\lfloor \Delta_{w'b^{-1}u} \rfloor$ we have $\Delta_{w'b^{-1}u} = \Delta_{w'u} \circ \Delta_{b^{-1}}$ and by definition we have $\Delta_{w'b^{-1}u} = \Delta_{w'b^{-1}} \circ \Delta_u$. It follows that Δ_u and $\Delta_{b^{-1}} = \beta^{-1}$ have disjoint support and since Δ_u is B-decomposable each generator diagram in Δ_u has support disjoint from that of β^{-1}. Clearly $(ab, ba) \in \mathcal{R}$ for every bump a with support disjoint from b, thus $(ub^{-1}, b^{-1}u) \in \langle\langle \mathcal{R} \rangle\rangle$ and this witnesses $w'b^{-1}ubw'' \equiv_G w'ub^{-1}bw''$.

Otherwise, π_1 and ρ_2 form one diagram while π_2 and ρ_1 form the other such that $t(\pi_1) = t(\pi_2) = i(\rho_1) = i(\rho_2)$ while $\lfloor \pi_1 \rfloor = \lceil \pi_2 \rceil$ and $\lfloor \rho_1 \rfloor = \lceil \rho_2 \rceil$. Since $\Delta_w = a_1\delta \circ \ldots \circ a_k\delta$ where $w = a_1 \ldots a_k$ it cannot be the case that both π_1 and ρ_2 form $a_i\delta$ and, at the same time, π_2 and ρ_1 form $a_j\delta$ since this would imply $a_i\delta \circ a_j\delta = \beta^{-1} \circ \beta$ or $\beta \circ \beta^{-1}$. Therefore, without loss, there must exist cells π_3 and ρ_3 in Δ_w such that π_3 and ρ_1 form the subdiagram $a_i\delta$ while π_2 and ρ_3 form the subdiagram $a_l\delta$ and π_1 and ρ_2 form a subdiagram $a_j\delta$ where $a_i\delta \circ a_j\delta \circ a_l\delta = \beta \circ \beta^{-1} \circ \beta$ or $\beta^{-1} \circ \beta \circ \beta^{-1}$ for $i < j < l$. It follows that $w = w'a_iu_1a_ju_2a_kw''$ for some words u_1, u_2 and $t(\pi_3) = i(\rho_1) = t(\pi_2) = i(\rho_3)$ and so all six cells we consider share a common vertex v. Notice that if $a_j = b$ then π_1 is positive and it follows that π_3 is negative, but the only cells that have the source s of b as the rightmost label of their top or bottom path is the relation labelled by s itself and, if a_{t+1} is a destination where $a_t = s$, the relation labelled by the source paired with a_{t+1}. Considering

$\Delta_{a_i u_1 a_j u_2 a_k} = \Delta_{a_i} \circ \Delta_{u_1} \circ \Delta_{a_j} \circ \Delta_{u_2} \circ \Delta_{a_k}$ we can see that, since the cell in Δ_{a_i} with bottom path label $a_{t+1} \ldots d$ (where d is the destination of b) must form a dipole with the cell in Δ_{a_j} with the same label on its top path, every generator diagram of Δ_{u_1} must be supported to the left of a_{t+1}, which is a contradiction, and we thus have $a_j = b^{-1}$ while $a_i = a_j = b$.

Now considering $\Delta_b \circ \Delta_{u_1} \circ \Delta_{b^{-1}} \circ \Delta_{u_2} \circ \Delta_b$ it follows by the same argument that every generator diagram of Δ_{u_1} must be supported to the left of d and, similarly, Δ_{u_2} must be supported to the right of s. It is easy to see that, if c is a bump supported to the right of s, we have $(ac, ca), (ab, ba) \in \mathcal{R}$ if a is a bump supported to the left of s and $(bab^{-1}c, cbab^{-1}) \in \mathcal{R}$ otherwise. It follows that $(bu_1b^{-1}u_2, u_2bu_1b^{-1}) \in \langle\langle \mathcal{R} \rangle\rangle$.

Finally, suppose Δ is obtained from Δ_w by a sequence of dipoles reductions such that Δ contains a pair of dipoles with the stated property while Δ_w and every diagram in between do not. Further suppose the sequence is as short as possible i.e. if there exists a Δ' obtained from Δ_w by a sequence of dipole reductions containing, as dipoles, the pair of dipoles of Δ then any dipole of Δ' which is a dipole of Δ_w must be a dipole of Δ. It follows that at least one cell of every term in Δ_w is in Δ, and since the pair of dipoles is not a pair of dipoles in Δ_w there must be subword bub^{-1} of w such that dipoles formed in Δ_u must reduce so that a cell of $b\delta$ and a cell of $b^{-1}\delta$ together form a dipole in Δ where they did not in Δ_w. Therefore, there must be a (maximal) subword u' of u such that $\Delta_{u'}$ is a subdiagram of Δ and $u = cu'c^{-1}$ for some word c. Replacing bumps a with cac^{-1} wherever necessary in the relations of $\mathcal{R}$ the result now follows from the previous cases. □

Lemma 7.1.6. *Let w be a word on B. Then $w \notin W(w\delta)$ if and only if there exists a diagram Δ obtained from Δ_w by a sequence of elementary reductions which contains a pair of dipoles together forming two subdiagrams, one equal to a generator diagram β and the other equal to β^{-1}.*

Proof. If $w \in W(w\delta)$ then there is a sequence of choices of $w\delta$ which obtains w. At a given step of the procedure, if we choose a generator diagram then we do not introduce any dipoles. On the other hand, if we choose a negative cell then we do introduce a dipole, the bottom cell of which forms a generator diagram with the choice in question and, then, we obtain the diagram to consider for the next step by removing this generator diagram. This leaves the top cell π of the dipole and the only way to add a dipole via the procedure to introduce a cell that would pair with π is for π itself to be a choice; however, by Lemma 4.2.5 this cell already forms a generator diagram which will become a choice after removing all choices above it in the partial order. Making these choices according to w therefore defines a sequence of dipole introductions and it is clear that at no point in this sequence is there a pair of dipoles with the stated property. This suffices for any other order of these introductions obtaining Δ_w since each introduction still only occurs at cells which did not form part of a generator diagram.

Now suppose Δ_w is such that no diagram Δ obtained from Δ_w by a sequence of reductions contains a pair of dipoles with the stated property - we proceed by induction on the length of w. Let w' satisfy $w = w'b$, then by the inductive hypothesis $w' \in W(w'\delta)$ and if both cells of this term $\beta = b\delta$ are contained in $w\delta$ then it is clear that $w \in W(w\delta)$. Otherwise the positive cell of β is not in $w\delta$ while the negative one is, so let v be the vertex connecting the cells of β and suppose without loss that its negative cell π is in $I^{\Pi}(v)$. Since the positive cell of β is not in $w\delta$ it must form part of a dipole in some Δ obtained from Δ_w by a sequence of reductions and therefore there is a positive cell π^{-1} in $I^{\Pi}(v)$ labelled by the same relation as π. Notice that π^{-1} must be contained in $w\delta$, and therefore in $w'\delta$, by the assumption. Since $w'\delta$ is the reduced form of $w\delta \circ \beta^{-1}$ it follows that π^{-1} forms part of a generator diagram in $w'\delta$ where the bottom path of the other cell is a subpath of $\lfloor w'\delta \rfloor$ and it is now clear that $w \in W(w\delta)$. □

Lemma 7.1.7. *A word w satisfies $w \equiv_G 1$ if and only if $w \in \langle\langle \mathcal{R} \rangle\rangle$.*

Proof. Since $\mathcal{R}$ is a subset of $\mathcal{D}$ it is clear that $r \in \langle\langle R \rangle\rangle$ implies $r \equiv_G 1$. Now consider $(u, w) \in \mathcal{D} \setminus \mathcal{R}$ - this implies that u, w are words such that $\Delta_u \equiv \Delta_w$ and either $u \notin W(w\delta)$ or $w \notin W(w\delta)$. Suppose for now that $u \in W(w\delta)$, then by Lemma 7.1.6 there is a Δ obtained from Δ_w by a sequence of dipole reductions containing a pair of dipoles together forming two subdiagrams, one equal to a generator diagram β and the other its inverse. By Lemma 7.1.5 there is a word w_1 containing a freely reducible pair such that $(w_1, w) \in \langle\langle \mathcal{R} \rangle\rangle$. If $w_1^{\vee} \notin W(w\delta)$ then we may reapply the two lemmas to obtain w_2 such that $(w_2, w) \in \langle\langle \mathcal{R} \rangle\rangle$ and, continuing as far as possible, we eventually reach a word w' such that $w' \in W(w\delta)$ and $(w, w') \in \langle\langle \mathcal{R} \rangle\rangle$ and it follows that $(u, w) \in \langle\langle \mathcal{R} \rangle\rangle$ since $u, w' \in W(w\delta)$. If both $u, w \notin W(w\delta)$ then we may simply apply the same argument to each of them in order to obtain $(u, u'), (w, w') \in \langle\langle \mathcal{R} \rangle\rangle$ where $u', w' \in W(w\delta)$. □

So far we have restricted the relations we need to consider by analysing the structural similarities between equivalent diagrams - now we do the same between non-equivalent diagrams.

Definition 7.1.8. Let Δ be reduced such that $|C(\Delta)| = 1$ and $\Delta \equiv \beta \circ w\delta \circ \beta^{-1}$ contains a single dipole for some generator diagram β and word w then we say it is a **conjugate diagram**. By Lemma 4.2.5 the vertex v which the dipole of $\beta \circ w\delta \circ \beta^{-1}$ shares with Δ has edges of Δ incident to it labelled by feet such that reading from top to bottom while writing the corresponding generator at each edge obtains a word $w_p a^k w_p^{-1}$ for a word w_p, generator a and integer k. We call v the **principal vertex** of Δ, w_p the **principal conjugator** of Δ and β the **first conjugator** of Δ.

In general, if Δ is reduced then the presence of a negative cell that does not form part of a generator diagram implies there is a subdiagram of Δ which is a conjugate diagram whose first conjugator is the negative cell in question. More explicitly, this subdiagram consists of this negative cell, its paired positive cell and every cell of Δ bounded between them in the partial order. As such, Δ may be partitioned into maximal conjugate diagrams and generator diagrams.

Definition 7.1.9. Let Δ be reduced such that $|W(\Delta)| = 1$; we say Δ is **basic**. On the other hand, if Δ is reduced such that $W(\Delta) = \{w_1, w_2\}$ and w_1 and w_2 share no common non-trivial prefix or suffix then we say Δ is a **basic relator diagram**. We denote $\mathcal{R}_{\mathcal{B}} = \{r \in R_\Delta \mid \Delta \text{ is a basic relator diagram}\}$.

Remark. Notice that, if $|W(\Delta)| = 1$ for all reduced Δ then G must be free of rank $|B|$.

Lemma 7.1.10. *A reduced diagram Δ with $|W(\Delta)| = 2$ is a basic relator diagram if and only if Δ and Δ^{-1} have precisely two choices each and there exists Δ_1 and Δ_2 such that $\Delta = \Delta_1 \circ \Delta_2 = \Delta_2 \circ \Delta_1$ where one of the following holds:*

(a) *Δ_1 and Δ_2 are generator diagrams with disjoint support;*

(b) *Δ_1 is a generator diagram β and Δ_2 is a basic conjugate diagram whose first conjugator has support overlapping with β;*

(c) *Δ_1 and Δ_2 are both basic conjugate diagrams where the support of Δ_2 overlaps the support of the first conjugator of Δ_1 and vice versa.*

Proof. Let Δ be a basic relator diagram. There cannot be more than two choices of Δ since no two distinct choices of a diagram can correspond to the same generator. If there is only a single choice of Δ then the elements of $W(\Delta)$ would have a common suffix - similarly, there must be precisely two choices of Δ^{-1} since otherwise we have a common prefix.

If both choices are generator diagrams β_1, β_2 then they must have disjoint support and after choosing one, the other is immediately available as a choice. It follows that $\Delta = \beta_1 \circ \beta_2 = \beta_2 \circ \beta_1$. If one choice is a negative cell (labelled by a source, without loss) and the other is a generator diagram β, call them c_1 and c_2 respectively, then c_1 must be the only choice after choosing c_2, meaning that $\Delta = \Delta_1 \circ \beta = \beta \circ \Delta_1$ where Δ_1 is the diagram obtained after choosing c_2. Since c_1 is a negative cell which does not form part of a generator diagram and $t(c_1)$ has only one outgoing edge in Δ_1, it follows by Lemma 4.2.5 that there exists a generator diagram γ such that $\Delta_1 \equiv \gamma \circ \Delta' \circ \gamma$ for a non-trivial subdiagram Δ' of Δ_1. This means choosing c_1 in Δ must block c_2 from being an immediate choice since the choice of Δ' must now be an available choice, and therefore the support of γ overlaps the support of β. This argument extends to the case where both choices are negative cells, thereby proving the lemma. □

Lemma 7.1.11. *Let Δ be a reduced diagram such that $\Delta = \Delta_1 \circ \Delta_2$ where Δ_1, Δ_2 are conjugate diagrams with disjoint support such that the support of the first conjugator of Δ_1 contains the label of the choice of Δ_2 and vice versa. Then Δ_i is a product of basic diagrams $\Psi_{i,j}$ such that $\Psi_{1,j} \circ \Psi_{2,k}$ is a basic relator diagram for all j, k.*

Proof. Since Δ_1 is a conjugate diagram the choice of Δ_2 will not be a future choice of $\Delta(c_1)$ until after all future choices of Δ_1 have been removed, and vice versa. It follows that $W(\Delta) = W(\Delta_1)W(\Delta_2) \sqcup W(\Delta_2)W(\Delta_1)$ and so we treat Δ_1 and Δ_2 separately. Without loss, suppose c_1 is labelled by a source s and c_2 is labelled by a destination.

We claim that Δ_1 is a product of basic diagrams each supported to the left of a_{m+1} where $a_m = s$ (and symmetrically for Δ_2) and we proceed by induction on the cardinality of $W(\Delta_1)$. We choose the only choice so long as there is only one choice available to obtain $\Delta' = \Delta(c_1, c_1^2, \ldots, c_1^k)$ and a word w_k which must be the length k suffix of the principal conjugator w. Notice that the subdiagram Γ of Δ' consisting of all cells below π and above π^{-1} in the partial order, where (π, π^{-1}) is the next term in the principal conjugator, is itself a conjugate diagram whose principal conjugator is r where $w = rw_k$. Consider choices in the subdiagram $\Delta' \setminus \Gamma$ - if a choice is a negative cell (other than π) then the conjugate diagram Θ it belongs to is a subdiagram of $\Delta' \setminus (w_k\delta^{-1} \circ \Gamma)$ and since this is also a subdiagram of Δ it follows that Θ is supported to the left of the destination d of the first letter of w_k. By the inductive hypothesis Θ is a product of basic diagrams each supported to the left of d and it follows that the reduced form of $\Theta^{w_k\delta}$ is a product of basic diagrams each supported to the left of a_{m+1}. On the other hand, if there is a choice which is a generator diagram β then it similarly follows that the reduced form of $\beta^{w_k\delta}$ is a product of basic diagrams each supported to the left of a_{m+1} and, indeed, it then follows that the subdiagram $\Delta' \setminus (w_k\delta^{-1} \circ \Gamma)$ satisfies this property. Again - by applying the inductive hypothesis to Γ we therefore see that the reduced form of $\Gamma^{w_k\delta}$ also satisfies this property, and the claim follows since $\Delta \equiv (\Gamma \circ (\Delta' \setminus (w_k\delta^{-1} \circ \Gamma)))^{w_k\delta}$.

We have shown that Δ_1 is a product of basic diagrams each supported to the left of a_{m+1} and that Δ_2 is a product of basic diagrams each supported to the right of a_{l-1} where a_l is the label of c_2. The result now follows since if Ψ_1 and Ψ_2 are basic conjugate diagrams with disjoint support such that the first conjugator of Ψ_1 is supported on the label of the choice of Ψ_2 and vice versa then $\Psi_1 \circ \Psi_2$ is a basic relator diagram. □

Proposition 7.1.12. $\langle\langle \mathcal{R}_{\mathcal{B}} \rangle\rangle = \langle\langle \mathcal{R} \rangle\rangle$.

Proof. Notice that it suffices to show that, for a given reduced diagram Δ, each relation in R_Δ is a consequence of the relations in $\mathcal{R}_\mathcal{B}$. We proceed by induction on the cardinality of $W(\Delta)$.

If $|W(\Delta)| = 1$ then R_Δ contains no non-trivial relations and there is nothing to show. If $|W(\Delta)| > 1$ then choose the sole choice of Δ repeatedly until we reach a Δ which has at least two choices, so far obtaining a word s. Let $c_1, \ldots, c_k$ be the choices of Δ - notice that the relations in $R_{\Delta(c_i)}$ are consequences of $\mathcal{R}_\mathcal{B}$ by the induction hypothesis. As such we see that for $u_i, w_i \in W(\Delta(c_i))$ we have $(u_i c_i s, w_i c_i s) \in \langle\langle \mathcal{R}_\mathcal{B} \rangle\rangle$. Noticing that for any $w \in W(\Delta)$ there exists an i such that $w = u_i c_i s$, it now suffices to show for each i, j there exists $u_i \in W(\Delta(c_i))$ and $u_j \in W(\Delta(c_j))$ such that $(u_i c_i s, u_j c_j s) \in \langle\langle \mathcal{R}_\mathcal{B} \rangle\rangle$. We split our argument into four cases by considering the form these choices may take. Without loss we assume c_i is left of c_j and let β_i, β_j be their corresponding generator diagrams. In each case we find two sequences of choices s_i and s_j, starting with c_i and c_j respectively, such that the diagrams reached by the sequences coincide i.e. $\Delta(s_i, s_j) = \Delta(s_j, s_i)$, since we can then pick any $p \in W(\Delta(s_i, s_j)) = W(\Delta(s_j, s_i))$ and see that $ps_j s_i s, ps_i s_j s \in W(\Delta)$, and finally observe that $(s_j s_i, s_i s_j) \in \langle\langle \mathcal{R}_\mathcal{B} \rangle\rangle$.

In the first case, both c_i and c_j are generator diagrams and it follows that $\Delta(c_i, c_j) = \Delta(c_j, c_i)$ and $(c_i c_j, c_j c_i) \in \mathcal{R}_\mathcal{B}$. In the second (and, symmetrically, the third) case, c_i must be negative cell and c_j must be a generator diagram. Let Ψ_i be the conjugate diagram which c_i belongs to - since c_i is part of the bottom path Ψ_i must be maximal - and suppose for now there are no cells above Ψ_i in the partial order. Consider $\Delta(c_i)$ - if $c_j \in C(\Delta(c_i))$ then $\Delta(c_i, c_j) = \Delta(c_j, c_i)$ and $\beta_i \circ \beta_j$ is a basic relator diagram. On the other hand, if $c_j \notin C(\Delta(c_i))$ then Lemma 7.1.11 tells us that Ψ_i is a product of basic diagrams $\Psi_{i,k}$ such that $\Psi_{i,k} \circ \beta_j$ is a basic relator diagram for all k and it follows that $s = s_1 \ldots s_n \in W(\Psi_i)$ where $s_k \in W(\Psi_{i,k})$ whence $(sc_j, c_j s) \in \langle\langle \mathcal{R}_\mathcal{B} \rangle\rangle$ and it is clear that $\Delta(s, c_j) = \Delta(c_j, s)$.

Now suppose there are cells above Ψ_i in the partial order and consider the sub-diagram Γ of Δ formed by these cells. It may be partitioned into maximal conjugate diagrams and generator diagrams supported to the left of the label of c_i. For any such generator diagram β, we have that $\beta \circ \beta_j$ is a basic relator diagram. Let Ψ be such a maximal conjugate diagram and let w_p be its principal conjugator; if l is a prefix of w_p such that $c_j \in C((\Psi \circ \beta_j)(l^{-1}))$ then $b\delta \circ \beta_j$ is a basic relator diagram for any letter b of l. Suppose l is the maximal such prefix, then we may apply Lemma 7.1.11 to the reduced form of $\Psi^{l\delta} \circ \beta_j$ to see that there is a sequence s of future choices of $\Psi^{l\delta}$ such that $(sc_j, c_j s) \in \langle\langle \mathcal{R}_\mathcal{B} \rangle\rangle$ and it follows that lsl^{-1} is a sequence of future choices of Ψ such that $(lsl^{-1} c_j, c_j lsl^{-1}) \in \langle\langle \mathcal{R}_\mathcal{B} \rangle\rangle$. Making choices in a way which respects the partition of Γ, followed by Ψ_i, it is now clear that we may obtain sequences of future choices of Δ with the desired properties.

The fourth and final case, where both c_i and c_j are negative cells, may be seen applying the above arguments to each where necessary, depending on the forms that $\Delta(c_i), \Delta(c_j)$ take. □

We have shown that

$$G \cong \langle B \mid \mathcal{R}_B \rangle$$

and we can also show that the relations in $\mathcal{R}_{\mathcal{B}}$ have the following form.

Lemma 7.1.13. *Let Δ be a basic relator diagram. Then the relation $r \in R_\Delta$ is a commutator.*

Proof. By Lemma 7.1.10 we know that $\Delta = \Delta_1 \circ \Delta_2 = \Delta_2 \circ \Delta_1$ where Δ_1, Δ_2 are basic diagrams. If w_1, w_2 are the unique words in $W(\Delta_1), W(\Delta_2)$ respectively then it follows that $W(\Delta) = \{w_1 w_2, w_2 w_1\}$ and so $R_\Delta = \{(w_1 w_2, w_2 w_1)\}$ and this relation is equivalent to $(w_1 w_2 w_1^{-1} w_2^{-1}, \varepsilon)$. □

By consolidating the results of this section we conclude the following theorem.

Theorem 7.1.14. *Let $G = \langle B \rangle$ be a fast group on the interval. Then the relations of G are all consequences of commutators.*

Proof. By Theorem 4.2.7 G is isomorphic to a bumpy diagram group. By Lemma 7.1.7 and Proposition 7.1.12 it then follows that the relations of G are all consequences of the relations in $\mathcal{R}_{\mathcal{B}}$. By Lemma 7.1.13 all relations in $\mathcal{R}_{\mathcal{B}}$ are commutators. □

Bibliography

[1] M.D. Atkinson. “On counting the number of linear extensions of a tree”. In: *Order* 7 (1990), pp. 23–25.

[2] J. Belk and F. Matucci. “Conjugacy and Dynamics in Thompson’s Groups”. In: *Geometriae Dedicata* 169.1 (2014), pp. 239–261.

[3] J. Belk and L. Stott. “PF_4 is isomorphic to F_4”. In: *arXiv:2303.16868* (2023).

[4] J.M. Belk. “Thompson’s Group F”. PhD thesis. arXiv:0708.3609: Cornell University, 2004.

[5] J.M. Belk and K.S. Brown. “Forest diagrams for elements of Thompson’s group F”. In: *Int. J. Alg. and Comp.* 15.5-6 (2005), pp. 815–850.

[6] P. Bellingeri, A. Genevois, and N Nanda. “Right-angled Artin groups are symmetric diagram groups”. In: *arXiv:2305.11810* (2023).

[7] C. Bleak et al. “Groups of fast homeomorphisms of the interval and the ping-pong argument”. In: *Journal of Combinatorial Algebra* 3.1 (2019), pp. 1–40.

[8] L. Boasson and O. Carton. “Transfinite Lyndon words”. In: *International Conference on the Developments in langauge Theory, Lecture Notes in Comput. Sci.* 9168 (2015), pp. 179–190.

[9] M. G. Brin. “The ubiquity of Thompson’s groups F in groups of piecewise linear homeomorphisms of the unit interval”. In: *J. London Math. Soc.* 60.2 (1999), pp. 449–460.

[10] M.G. Brin. “Higher dimensional Thompson’s groups”. In: *Geom. Dedicata* 108 (2004), pp. 163–192.

[11] J. Burillo, K. Bux, and B. Nucinkis. “Cohomological and metric properties of groups of homeomorphisms of $\mathbb{R}$”. In: *Oberwolfach Reports* 15 (2018), pp. 1579–1633.

[12] P.J. Cameron and P. Cara. “Independent generating sets and geometries for symmetric groups”. In: *J. Algebra* 258 (2002), pp. 641–650.

[13] P.J. Cameron and L. Stott. “Trees and cycles”. In: *arXiv:2010.14902* (2020).

[14] J. Cannon, W. Floyd, and W. Parry. "Introductory notes on Richard Thompson's groups". In: *L'Enseignement Mathematique* 42 (1996), pp. 215–256.

[15] J. Dénes. "The representation of a permutation as the product of a minimal number of transpositions and its connection with the theory of graphs". In: *Publ. Math. Inst. Hungar. Acad. Sci.* 4 (1959), pp. 63–71.

[16] M. Eden and M.P. Schützenberger. "Remark on a theorem of Dénes". In: *Magyar Tud. Akad. Mat. Kutató Int. Közl.* 7 (1962), pp. 353–355.

[17] P. Ehrlich. "Number Systems with Simplicity Hierarchies: A Generalization of Conway's Theory of Surreal Numbers". In: *The Journal of Symbolic Logic* 66.3 (2001), pp. 1231–1258.

[18] D.S. Farley. "Actions of Picture Groups on CAT(0) Cubical Complexes". In: *Geom. Dedicata* 110 (2005), pp. 221–242.

[19] D.S. Farley. "Finiteness and CAT(0) properties of diagram groups". In: *Topology* 42 (2003).

[20] D.S. Farley. "The planar pure braid group is a diagram group". In: *arXiv:2109.02815* (2021).

[21] D.S. Farley and B. Hughes. "Braided diagram groups and local similarity groups". In: *Geometric and cohomological group theory, London Math. Soc. Lecture Note Ser.* 444 (2018), pp. 15–33.

[22] A. Genevois. "An introduction to diagram groups". In: *arXiv:2211.12068* (2022).

[23] A. Genevois. "Contracting isometries of CAT(0) cube complexes and acylindrical hyperbolicity of diagram groups". In: *Algebr. Geom. Topol.* 20.2 (2020), pp. 49–134.

[24] G. Golan and Sapir M.V. "On Jones' subgroup of R. Thompson's group F". In: *J. Algebra* 470 (2017), pp. 122–159.

[25] G. Golan and Sapir M.V. "On subgroups of R. Thompson's group F". In: *Transactions of the American Mathematical Society* 369.12 (2017), pp. 8857–8878.

[26] G. Golan-Polak. "The generation problem in Thompson's group F". In: *arXiv:1608.02572* (2016).

[27] S. Goryainov et al. "PI-eigenfunctions of the Star graphs". In: *Linear Algebra Appl.* 586 (2020), pp. 7–27.

[28] The GAP group. "GAP – Groups, Algorithms, and Programming, Version 4.10.2". In: `https://www.gap-system.org` (2019).

[29] V. Guba and M. Sapir. “Diagram groups”. In: *Memoirs of the Amer. Math. Soc.* 130.620 (1997), pp. 1–117.

[30] V. Guba and M. Sapir. “Diagram groups and directed 2-complexes: homotopy and homology”. In: *Journal of Pure and Applied Algebra* 205.1 (2006), pp. 1–47.

[31] V. Guba and M. Sapir. “Diagram groups are totally orderable”. In: *J. Pure Appl. Algebra* 205.1 (2006), pp. 48–73.

[32] V. Guba and M. Sapir. “On subgroups of R. Thompson’s group F and other diagram groups”. In: *Sb. Math.* 190.8 (1999), pp. 1077–1130.

[33] V.S. Guba and M.V. Sapir. “Rigidity properties of diagram groups”. In: *Internat. J. Algebra Comput.* 205.1 (2006), pp. 9–17.

[34] B. Hughes. “Local similarites and the Haagerup property”. In: *Groups Geom. Dyn.* 3.2 (2009), pp. 299–315.

[35] B. Hughes. “Trees and ultrametric spaces: a categorical equivalence”. In: *Adv. Math.* 189 (2003), pp. 148–191.

[36] V. Kilibarda. “On the algebra of semigroup diagrams”. In: *Internat. J. Algebraic Comput.* 7 (1997), pp. 313–338.

[37] B. Kraft. “Diameters of Cayley graphs generated by transposition trees”. In: *Discrete Appl. Math.* 184 (2015), pp. 178–188.

[38] P. Moszkowski. “A Solution to a Problem of Dénes: a Bijection Between trees and Factorizations of Cyclic Permutations”. In: *Europ. J. Combin.* 101 (1989), pp. 13–16.

[39] L.H. Soicher. “GRAPE, GRaph Algorithms using PErmutation groups, Version 4.8.1”. In: *https://gap-packages.github.io/grape* (2018).

[40] J. Whiston. “Maximal independent generating sets of the symmetric group”. In: *J. Algebra* 232 (2000), pp. 255–268.