Thompson's Groups

Matthew G. Brin
Binghamton University

- - -
-
- - -

Edition 1.1
February 9, 2026

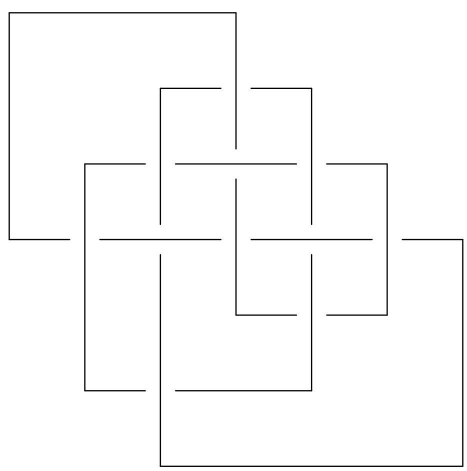

K11a266

To
Richard J. Thompson

Preface

This book is about countably infinite groups named for Richard J. Thompson, and not about finite groups associated to John G. Thompson.

Thompson's groups form a loosely defined family of countably infinite groups with several attractive features. They are easy to define. They can be both easy and difficult to work with. They have strange collections of properties. They are easy to modify and combine with other groups (and hence the looseness of the Thompson group family). They are not members of other well studied families of countably infinite groups. They exist for fairly fundamental reasons, and so interact with a number of other areas of mathematics. It is intended to at least touch on all of these features in this book.

The book was written to entice the uninterested, and help the interested. The book is heavy on basics and techniques and lighter on generality. We will give details of key examples, and references for generalizations.

The number of papers relevant to the Thompson groups is now in the high hundreds, and limitations on space and my ability make it impossible to cover, and probably impossible to even mention them all. The interactions of the Thompson groups with other areas of mathematics are covered with varying depth, mostly depending on the author's familiarity with those areas. However these interactions are an important aspect of the groups, and some attempt has been made to cover some and at least mention others.

This book was written to be readable by second year graduate students in Ph.D. programs. The accuracy of that sentence will vary greatly in different parts of the book. The book is intended as a reference, not as a textbook, and there are no formal exercises. Some work is left to the reader, and there are suggestions of extra work that can be done.

There are many people to thank and the list will grow as editions appear. The list must start with my wife Dawn Coe for infinitely many reasons, and continue with Jim Belk, Collin Bleak, Alexander Borisov, Arnaud Brothier, Ross Geoghegan, James Hyde, Yash Lodha, Roger Maddux, Cary Malkiewich, Marcin Mazur, Conchita Martínez-Pérez, Enrique Pardo, Richard Thompson, Albert Visser, Eric Wofsey, and Matt Zaremsky.

Special to the first edition

This edition contains three fairly complete chapters and fragments of others. Chapter 2, a basic introduction to the core groups, is described as fairly complete, but is the most likely to see future changes if material added in later chapters would benefit from adjustments in the more basic Chapter 2.

What will be added to the incomplete chapters and what new chapters will be added is difficult to determine now. As a first edition, there is a certain amount of experimentation with the process. Notes at the end of each chapter will give more information. A footnote at the beginning of each chapter will give a brief statement about the status of that chapter.

The figure on the title page has nothing whatsoever to do with the contents of the book and is based on a doodle I used to make in high school and college.

Matthew Brin
Vestal, New York
2025

Contents

CHAPTER 1

Introduction

[1]The Thompson groups are countably infinite groups which are almost always regarded as discrete groups. Their lack of finiteness is very much part of their identity. While uncountable groups can show up in the discussion, they are never referred to as Thompson groups.

In 1969, Thompson introduced three groups $F \le T \le V$ which form the core of the Thompson group family. Surrounding the core are layers of variations that get increasingly removed from the original three groups. Eventually groups are encountered that no one would call Thompson's groups. In spite of Sigler's law [184], the phrase "Thompson's groups" appears to be accurate for several reasons. See the outline of the early history of the groups in Section 15.

The original groups are easy to define, and we give definitions for F and T in the next paragraph. There are other definitions, but the ones we give below were chosen because they are the simplest, the most direct, and the least revealing.

The group F is the set of self homeomorphisms of the unit interval I under composition that

(1) are piecewise linear in that the graph is made of straight line segments with only finitely many discontinuities of slope,
(2) have slopes between the discontinuities that are integral powers of 2 (so all elements are orientation preserving), and
(3) have discontinuities of slope confined to the dyadic rationals $\mathbf{Z}[\frac{1}{2}]$ (those $m/2^n$ with both m and n from $\mathbf{Z}$).

The group T can be defined as those self homeomorphisms of $\mathbf{R}/\mathbf{Z}$, the circle of length one, satisfying (1)–(3) above in addition to

(4) all elements preserve the set $\mathbf{Z}[\frac{1}{2}]$.

Item (4) rules out rotations of the circle by non-dyadics.

Both F and T are not only finitely generated, but also finitely presented, and T is simple. Before 1969, there were no known examples of infinite, finitely presented, simple groups.

[1]This introduction will change as the book changes.

The group V is also finitely presented and simple, and it acts naturally on the Cantor set. A definition for V along the lines of (1)–(3) above does exist, but we offer, without details, two alternate definitions. The definitions seem different from each other, but are closely related.

For the first definition, accept that there is a variety of algebras that is universal for the property that its free algebra A on one variable has more than one element and is isomorphic to its free algebra on two variables. The group V is the automorphism group of A.

For the second definition, we simply state that V is the topological full group of the full one-sided shift on the alphabet $\{0, 1\}$. Regarding elements of the Cantor set as infinite sequences over $\{0, 1\}$, the group V specifically consists of all self homeomorphisms of the Cantor set whose germs are compositions of the germs and their inverses of the self map of the Cantor set that deletes the first entry of each sequence.

More generally, Thompson's groups arise when self similar structures arise. The Cantor set is the coproduct of its left half and its right half, and is structurally identical to each half. Similarly, if the algebra A of the first definition of V is free on x and also free on $\{y, z\}$, then the subalgebra generated by y and the subalgebra generated by z are each isomorphic to A. Now A is the coproduct of these two subalgebras. The Cantor set and the algebra A are both coproducts of two copies of themselves, and we use this as an excuse to label as self similar objects both A and the Cantor set.

Given any object X that is a coproduct (or product) of two copies of itself, there will be a homomorphic image of V, which must be all of V or (rarely) trivial, in the automorphism group of X. This accounts for many of the properties of the Thompson groups and for many of their interactions with other areas of mathematics. This also touches on the claim made in the preface that the groups exist for fairly fundamental reasons. We discuss other points brought up in the preface.

That the groups are easy to define has already been demonstrated. That the groups are easy to work with comes from the existence of several combinatorial machines that make calculations with the elements mechanical. In spite of this, there are open questions about the groups, and some questions that have been closed were closed with difficulty.

The most famous open question is whether F is amenable. Also unknown is whether F is automatic in the sense of [69].

False solutions (in both directions) for the first question were common for a period, but have abated. The group F has sat comfortably on the fence of the question with F satisfying most known consequences of amenability, and not satisfying most properties known to be stronger than amenability. It has been facetiously(?) suggested that the question is independent of ZFC.

The second question gets less attention. But some work has been done which exhibits how difficult F can be to work with. A condition necessary to having an automatic structure is to have a quadratic Dehn function. That this holds for F is shown in [98] with an intricate calculation. This followed a sequence of results showing first that the Dehn fuction is exponentially bounded, then subexponentially bounded, and then bounded by a polynomial of degree 5.

Most prominent of properties of the Thompson family is that its members tend to be simple and finitely presented while infinite. Further they enter into marriages with entire families of groups so as to embed those groups into finitely presented simple groups. As groups with calculational tools, they have solvable word problems, but some have unsolvable conjugacy problems, and some have unsolvable first order theories. An early discovery (1984 [36]) is that they supply the first examples of torsion free groups G admitting projective $\mathbf{Z}G$-resolutions finitely generated in each dimension, but having infinite cohomological dimension. It is now almost expected that each addition to the Thompson group family be proven to share this property.

A property of F that characterizes the group is that it is an initial object in the category of groups admitting an endomorphism that differs from its square, but only by an inner automorphism (is a conjugacy idempotent). This led to independent discoveries of F by two teams in the late 1970s. In fact, some of the Thompson groups were discovered independently after Thompson at least four times over three decades. It could be said that mathematics had become ready for their discovery.

The conjugacy idempotent property arose because of a question in homotopy theory. Thompson himself came upon the groups because of his interest in algebraic logic. An interest in algebraic laws led to one of the later rediscoveries of F, where it was established that, in some sense, F could be described as the structure group of the associative law.

The definitions given above for the groups are dominated by the integer 2. Changing this to an arbitrary positive integer n gives rise to an infinite family of groups having similar properties. But more can be done, and a small sample is in Chapter 6 of this edition. The fact that V acts on the Cantor set allows marriages with other groups. As naive as it seems, it can be profitable to take another group G acting faithfully on the Cantor set and look at the group generated by V and

G together. Under the right conditions the result is a finitely presented, simple group that contains G as a subgroup.

The definitions of V given above connect V to dynamical systems, and to objects having a self similar structure. The connection of V to algebraic structures exhibiting ambiguity of free rank has been useful to the Thompson family since this has given the first (and at the moment only) classification up to isomorphism of certain variations of the group V.

Strangely, the group F parametrizes (seemingly without useful consequences) that part of planar graph theory that contains the difficult core of the four color theorem. Sparking more interest is the parallel fact that F parametrizes all knots and links.

To position the Thompson groups in the landscape of countably infinite (or finitely generated, or finitely presented) groups would require a map of that landscape. While it is nonsense to try to make such a map, we will make an attempt by placing groups of rigid actions such as isometry groups to the left, and groups that act more flexibly such as homeomorphism groups to the right. In between one might put linear groups, Artin groups, mapping class groups, and the vast family of hyperbolic groups in places that one finds comfortable. The Thompson groups, being none of the above for various reasons, hang somewhere in the middle while reaching out to both ends. They exhibit almost the flexibility of the homeomorphism groups which accounts for their tendency to be simple, but in a very controlled way which accounts for their finiteness properties.

The book will attempt to lay all this out. The invitation to learn about Thompson's groups comes with the remark that groups constantly connect to other topics, and the study of the groups is not an insular study.

Most of the actual work in this book can best be described as elementary. Another word might be traditional, and a stronger term might be old fashioned. Given that the Thompson groups exist for fundamental reasons, they show up in areas both old and new. There are only minor indications of the newer apperances in this book. Limits of time, space and the knowledge of the author contribute to that. Experts in some of the areas touched on in the book will find the narrative quite unsophisticated.

As of this edition, the structure of the book is as follows. Chapter 2 covers the basics of F, T and V. This includes the various definitions, some presentations, the tools for calculation and the properties that are the most basic. The mathematical origins of the groups are covered in Chapter 3. As a bit of indulgence, that chapter starts with a history of the early literature of the groups. The author finds the flow and exchange of ideas in this period interesting. Chapter 4 (at the moment not complete) builds complexes that the groups act on. This connects to the history and to the finiteness properties. Chapter 5 has first order results, including the unsolvability of the first order theory of F and T, and also a proof that the algebra of most of these groups determines spaces that they act on and their actions. Chapter 6 (very incomplete) discusses the variants and marriages that take place among the Thompson groups. Chapter 7 (very incomplete) contains enough of the representation of some Thompson groups into rank ambiguous algebraic structures to allow for the classification of some variants of the group V. Chapter 8 (hardly begun) contains a derivation of the length function for F as a start for geometric discussions of the groups.

There is also an appendix (Chapter 9) containing a few techniques whose development is independent of that of the Thompson groups, which are needed for a discussion of Thompson groups, and which might be unfamiliar to some readers.

Very much missing is material on isomorphisms, on conjugacy, on subgroups (a large topic), and most connections to other parts of mathematics.

CHAPTER 2

A first look at the first Thompson's groups

Chapter contents

1. Introduction

[1]We define the groups F, T and V and derive some of their basic properties. By basic properties we mean those whose derivations reasonably fit into a single chapter. However, we also include topics that let us talk about the content of later chapters.

The letters F, T and V have been rather stable since the 1996 expository paper of Cannon-Floyd-Parry [43]. However, the letter G is used occasionally for T and more frequently for V. Notations used before 1996 are given in [43]. The letter F was first used in the 1979 preprint of Freyd-Heller 1993 [73] to abbreviate "Free homotopy idempotent." The releveance of this phrase and the history around it will be explained in the introduction to Chapter 3 and in Section 19 of that chapter. The letter T was used in Brown 1987 [34] in honor of Thompson. The letter V (with an additional tilde) was used in widely circulated notes that were handwritten by Thompson shortly after the 1973 appearance of his paper [189].

Most of this chapter deals with F since once the mechanics of F are laid out, discussing T and V becomes much easier. The group F can be exhibited as a group of homeomorphisms, a group of manipulations of combinatorial objects, or as specified by a presentation. This chapter follows that order starting with homeomorphisms, then extracting the combinatorial views, and lastly using the combinatorial information to extract two presentations. During this development, properties of the group are derived when it is possible to do so.

The brief Section 2 sets notation and conventions, and Section 3 gives our chosen definition of F.

Sections 4 and 5 derive properties from the structure as a group of homeomorphisms. We show that F is torsion free, locally indicable and bi-orderable. Direct sums and wreath products abound as subgroups, the center is trivial, and non-abelian subgroups cannot be free and cannot be nilpotent. The commutator subgroup is contained in every non-trivial normal subgroup of F, is simple, is not finitely generated, and every element of the commutator subgroup is a product of no more than two commutators. The free semigroup on two generators embeds in F and so the growth of F is exponential.

Sections 6 through 8 develop three combinatorial approaches to F. These make calculations with elements of the group mechanical, they each have their advantages, and they are used heavily in the literature. Much of the work in these sections is to show the equivalence of the

[1]This chapter is reasonably complete. As a foundation for the chapters that follow, changes may be made to this chapter as other chapters develop.

views and how to pass from one to another. None of this is complicated, but it makes for a lot of material. Section 9 continues the combinatorial theme and derives a finite presentation as well as a convenient infinite presentation.

Diagram groups form an interesting class of groups and the diagrams of the name give a combinatorial tool for working with them. Some but not all diagram groups are Thompson groups, and many (including F as a prime example) but not all Thompson groups are diagram groups. We will not cover diagram groups. See the end notes (Section 14) for more information.

Section 10 gives properties of F that arise from the combinatorics. Elements of the commutator subgroup are characterized, and the structure of F as an iterated HNN extension of nothing is proven. It is shown that F is universal for the property of having an endomorphism that differs from its square by an inner automorphism, and that this universality characterizes F. In Chapter 3 it is shown how this led to one of the discoveries of F.

Section 11 looks at an important submonoid F_+ of F. The properties of the monoid lie behind many of the complexes used for further analysis of F. One of the themes of the current chapter is a search for distinguished representatives and normal forms for elements of F. The monoid F_+ ties these forms together.

Sections 12 and 13 are about T and V, respectively. Both can be described by modifying a definition of F and we do this for T. However, we introduce V differently to show its inevitable appearance in certain basic dynamical systems. In Chapter 3 this is generalized to situations involving self similar objects. About T we prove that all finite subgroups are cyclic, that there are non-abelian free subgroups, and that T is finitely presentable and simple. We show that in the standard representation of T as a group of homeomorphisms on the circle, every rational in $[0, 1)$ appears as a rotation number of an element of T. We exhibit V as the topological full group of the full one-sided shift on two letters. We show that V is finitely generated and simple, and that every countable, locally finite group embeds in V. The finite presentability of V requires more work than for F and T, and we delay the proof until Chapter 4 where we discuss complexes on which the groups act.

Section 14 gives some final notes on the material in this chapter and in chapters that follow.

2. Preliminaries

2.1. Some notation. We use $\mathbf{R}$, $\mathbf{Q}$, $\mathbf{Z}$, $\mathbf{N}$ for, respectively, the reals, the rationals, the integers and the natural numbers (which include zero). We use S^1 for the circle $\mathbf{R}/\mathbf{Z}$ of length one, $\mathbf{R}_{\geq 0}$ for the non-negative reals, and I for the unit interval $[0,1]$. We use ω for the first infinite ordinal. It has the same underlying set as $\mathbf{N}$, but ω emphasizes order and $\mathbf{N}$ emphasizes arithmetic. The subring $\mathbf{Z}[\frac{1}{2}]$ of $\mathbf{R}$ consists of the dyadic rationals, those $m/2^n$ with m and n in $\mathbf{Z}$.

We will use $|S|$ for the cardinality of a set S.

2.2. Groups and actions. In this book, groups will act on the right, and actions will be composed left-to-right. Any exceptions to this will be announced loudly. So if G acts on X with g and h in G and $x \in X$, then the images of x under g and gh are xg and $xgh = (xg)h$, respectively. The *fix set* or *fixed set* $\text{Fix}(g)$ of g is $\{x \in X \mid xg = x\}$ and the more frequently referred to *support* $\text{Supp}(g)$ of g is $\{x \in X \mid xg \neq x\}$. We also have

$$\text{Fix}(G) = \{x \in X \mid \forall g \in G,\, xg = x\} = \bigcap_{g \in G} \text{Fix}(g), \text{ and}$$
$$\text{Supp}(G) = \{x \in X \mid \exists g \in G,\, xg \neq x\} = \bigcup_{g \in G} \text{Supp}(g).$$

An action of G on X is *free* if for all $1 \neq g \in G$, we have $\text{Fix}(g) = \emptyset$.

For $S \subseteq G$, we write $\langle S \rangle$ for the subgroup of G generated by S. We write $H \leq G$ to indicate that H is a subgroup of G. For $x \in X$, the *orbit* of x under G is the set $\{xg \mid g \in G\}$ and is written xG. For $g \in G$, the orbit of x under g is $x\langle g \rangle$ and is the set $\{xg^i \mid i \in \mathbf{Z}\}$. For the next comment, we refer to (xg^i, xg^{i+1}) as a consecutive pair in the orbit of x under g.

We write g^h for $h^{-1}gh$. The “fundamental triviality” (terminology of [85], p. 11) observes that the calculation $(xh)g^h = (xh)h^{-1}gh = (xg)h$ shows that h carries the consecutive pair (x, xg) in the orbit of x under g to the consecutive pair $(x, xg)h = (xh, (xg)h) = (xh, (xh)g^h)$ in the orbit of xh under g^h. From this we get $\text{Fix}(g^h) = (\text{Fix}(g))h$, and $\text{Supp}(g^h) = (\text{Supp}(g))h$. For us, the commutator $[f,g]$ of f and g is $f^{-1}g^{-1}fg$, and we will often make use of the fact that $[f,g] = (g^{-1})^f g = f^{-1}f^g$ is in the normal closure of both f and g.

In what follows, there will be functions that are not actions and there will be functions that will later be part of an action of a group that has not yet been established. It will often be convenient to put such functions to the right of their arguments and compose left-to-right. We will point out when this occurs.

2.3. Some homeomorphism groups. We will use $\mathrm{Homeo}(X)$ for the group of self homeomorphisms of a topological space X. This will be most often used when X is one of $\mathbf{R}$, I, S^1, $\mathbf{R}_{\geq 0}$ or the Cantor set $\mathfrak{C}$. The first four are orientable and for an orientable X, we use $\mathrm{Homeo}_+(X)$ for the index two subgroup of $\mathrm{Homeo}(X)$ consisting of the orientation preserving self homeomorphisms of X.

We will work with piecewise linear homeomorphisms. For X equal to one of $\mathbf{R}$, $\mathbf{R}_{\geq 0}$, I or S^1, a homeomorphism f is *piecewise linear* if there is a set of points B, discrete in X, consisting exactly of those points where the derivative f' of f does not exist so that on each component of $X \setminus B$ the derivative f' is constant. The points in B are called the *breakpoints* of f. If a piecewise linear f is defined on some $[a, b]$, then af'_+ will denote the right hand derivative of f at a and bf'_- will denote the left hand derivative of f at b.

We will use $PL(X)$ for the group of all piecewise linear self homeomorphisms of X for the usual spaces X, and $PL_+(X)$ for the index two subgroup of those elements of $PL(X)$ that preserve orientation.

The phrase "piecewise linear" not really accurate and "piecewise affine" is closer to the truth. To say that a function f is *affine* on an interval J means that there are m and b so that $tf = mt + b$ for all $t \in J$.

Note that an f in $PL(\mathbf{R})$ or $PL_+(\mathbf{R})$ can have infinitely many breakpoints. We use $PLF(\mathbf{R})$ and $PLF_+(\mathbf{R})$ for the subgroups of $PL(\mathbf{R})$ and $PL_+(\mathbf{R})$, respectively, consisting of those elements with only finitely many breakpoints. A useful subgroup of $PLF_+(\mathbf{R})$ is $BPL(\mathbf{R})$ consisting of those elements in $PL(\mathbf{R})$ of bounded support.

3. The group F, I: definition

We repeat the definition from the introduction using vocabulary from Section 2. This definition is the quickest to state, and as noted in the introduction, the definition gives a certain dramatic impact since some of the claimed properties then seem surprising. The definition drives all observations that we make about F through Section 5. However, most of what is known about the group comes from a more combinatorial view of the group, and starting in Section 6, we derive the equivalent and more frequently used, combinatorial defintions.

Definition 3.1. The group F consists of those self homeomorphisms of the unit interval I that

(1) are piecewise linear,
(2) have only slopes that are integral powers of 2 (making all slopes positive and all elements orientation preserving), and
(3) have breakpoints confined to $\mathbf{Z}[\frac{1}{2}]$.

The group operation is composition and action is on the right.

It is clear from the definition that F is countably infinite, and from the lemma below that F is closed under composition and inverse and is thus a group.

LEMMA 3.2. *Each element of F carries $\mathbf{Z}[\frac{1}{2}] \cap I$ into itself.*

PROOF. For $h \in F$ and $t \in \mathbf{Z}[\frac{1}{2}] \cap I$, induct on the number of breakpoints of h in the interval $[0, t]$. Details are left to the reader □

What is not clear from Definition 3.1 is that F is not only finitely generated, but in fact finitely presented. We will show in Section 9 that the following two elements generate F. The notation x_0 and x_1 for these elements is common.

$$tx_0 = \begin{cases} 2t, & t \in [0, \frac{1}{4}], \\ t + \frac{1}{4}, & t \in [\frac{1}{4}, \frac{1}{2}], \\ (t+1)/2, & t \in [\frac{1}{2}, 1], \end{cases} \qquad tx_1 = \begin{cases} t, & t \in [0, \frac{1}{2}], \\ 2t - \frac{1}{2}, & t \in [\frac{1}{2}, \frac{5}{8}], \\ t + \frac{1}{8}, & t \in [\frac{5}{8}, \frac{3}{4}], \\ (t+1)/2, & t \in [\frac{3}{4}, 1]. \end{cases}$$

The graphs of x_0 and x_1 are shown below.

(3.1)

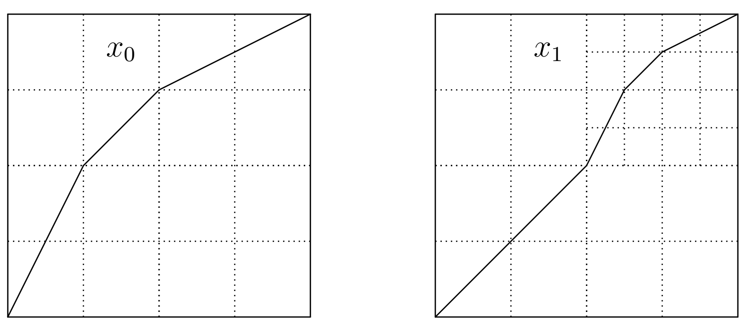

We will also show in Section 9 that a finite presentation for F using the elements above is

$$\langle x_0, x_1 \mid [x_0^2 x_1^{-1} x_0^{-1}, x_1] = [x_0^3 x_1^{-1} x_0^{-2}, x_1] = 1 \rangle. \tag{3.2}$$

3.1. Describing elements. To work with elements, we need to describe them.

DEFINITION 3.3. Let $J = [a, b]$ be a closed interval with non-empty interior and let $a = p_0 < p_1 < p_2 < \cdots p_n = b$ be a finite set of points E in J. The set E divides J into the subintervals

$[p_0, p_1], [p_1, p_2], \dots, [p_{n-1}, p_n]$, and we call this collection of intervals the *partition* of J with endpoints E. If P is a partition of J with endpoints E, we set $e(P) = E$. If P and Q with $|P| = |Q|$ are two partitions of a closed interval J, then we call (P, Q) a *partition pair*. For such a pair, the order preserving bijection from $e(P)$ to $e(Q)$ extends to a unique, orientation preserving, piecewise linear homeomorphism h from J to itself whose breakpoints are contained in $e(P)$. We say that the pair (P, Q) *determines* h.

REMARK 3.4 (Anticipating the future). The order preserving bijection from $e(P)$ to $e(Q)$ induces an order preserving bijection from P to Q where for intervals K and L with disjoint interiors, we say $K < L$ if $s < t$ for some s in the interior of K and t in the interior of L. If P and Q are partitions of I, then knowing the order preserving bijection from P to Q is enough to define an order preserving self homeomorphism of I. When we get to the groups T and V, the order preserving bijection from P to Q will be replaced by more general bijections, and it pays to get used to the practice of using the word "pair" to refer to a triple (P, σ, Q) where P and Q are sets of equal finite cardinality and σ is some bijection from P to Q. For F, the bijection is always order preserving (whatever the structues of P and Q) and will be suppressed from the notation.

3.2. The multiplication. Calculating compositions in $PL_+(I)$ can be best described as opportunistic. Compose in the obvious way when there is an obvious way to compose, and otherwise change things so that it becomes obvious. This theme will be repeated often.

Let f and g in F be determined, respectively, by partition pairs (P, Q) and (R, S). We assume that all of the endpoints of P, Q, R, and S are in $\mathbf{Z}[\frac{1}{2}]$. If by coincidence $Q = R$, then we can take advantage of that coincidence and declare that fg (remember the left-to-right order of composition) is determined by (P, S).

If the coincidence $Q = R$ does not occur, then we arrange it to do so. Set $E = e(Q) \cup e(R)$, let T be the partition with $e(T) = E$, let P' be the partition with $e(P') = Ef^{-1}$, and let S' be the partition with $e(S') = Eg$. Now f and g are determined, respectively, by (P', T) and (T, S') and fg is determined by (P', S'). The computations of P' and S' are straightforward given the affine nature of f and g on the intervals between the points of P and R.

The *rectangle diagrams* in [13, 43, 14] (attributed to William Thurston) comprise one set of visual tools for working with elements of F. If f is determined by (P, Q), then we can draw a rectangle with the top and bottom edges each representing the unit interval, with the endpoints of P marked out in proportional position on the top edge and of Q on the bottom edge. Then for each $p_i \in e(P)$, a straight line is drawn from p_i on the top edge to $p_i f \in e(Q)$ on the bottom edge. Rectangle diagrams for x_0 and x_1 are shown below.

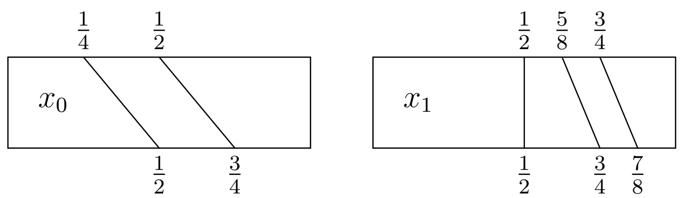

Multiplication is accomplished by stacking rectangle diagrams vertically and, if necessary, making additions to the partitions and adding a straight line from the top edge to the bottom edge for each addition. The calculation of x_0x_1 is shown below where the extra line in the rectangle for x_0 made necessary by the breakpoint at $\frac{5}{8}$ of x_1 is shown as dashed.

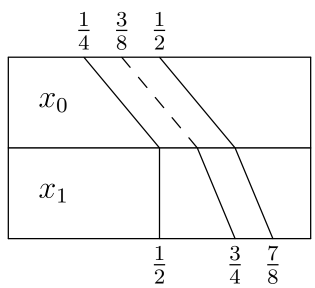

The reader might calculate x_0^2 using rectangle diagrams, and further verify the two relations in the presentation (3.2).

3.3. Normal and seminormal forms. Every element of F can be represented by a partition pair, but many partition pairs can represent the same element of F. This makes partition pairs seminormal forms rather than normal forms. However, a partition pair (P, Q) represents the identity element of F if and only if $P = Q$. A seminormal form that can pick out the identity is still a useful object. In particular, if generators for F are known, then the algorithm for multiplication of Section 3.2 makes the word problem for F solvable. This will be stated formally as Proposition 11.11.

Because we will introduce several ways to represent F, there will be several forms for elements. They all start out as seminormal forms that can pick out the identity element. Eventually, we will pick out a distinguished representative for each element of F, and have a true normal form.

4. Properties, I: from the definition

This section gives properties that have little to do with F other than the fact that F is a subgroup of $\mathrm{Homeo}_+(I)$ and of $PL_+(I)$. We need some key notions.

4.1. Orbitals. The order structure of $\mathbf{R}$ cooperates with the concepts of support and orbit. The proof of the lemma below is left to the reader.

DEFINITION 4.1. If H be a subgroup of $\mathrm{Homeo}_+(\mathbf{R})$, then an *orbital* of H is a component of $\mathrm{supp}(H)$. If $f \in \mathrm{Homeo}_+(\mathbf{R})$, then an *orbital* of f is an orbital of $\langle f \rangle$.

LEMMA 4.2. *If $H \leq \mathrm{Homeo}_+(\mathbf{R})$ and $t \in \mathbf{R}$ is in $\mathrm{supp}(H)$, then the orbital of H that contains t is the convex hull of the orbit tH.*

Not every orbital of an element of F needs to have dyadic endpoints. Non-dyadic, isolated fixed points can occur. The lemma below is an easy observation whose proof is left to the reader. If $p = \frac{m}{n} \in (0, 1)$ is a rational expressed in reduced terms with $n > 0$, then we can write $n = 2^j c$ where c is odd. Let k be the smallest positive integer so that $c|(2^k - 1)$. That such a k exists follows from Euler's generalization of Fermat's little theorem ([59, Theorem 13] or [167, Theorem 2.8]).

LEMMA 4.3. *Let p be in $(0, 1)$.*

(1) *If p is a non-dyadic rational, let k be as chosen in the paragraph above. There are then elements of F that are fixed on p but on no neighborhood of p, and futher for any such element f, the slope of f on a neighborhood of p is an integral power of 2^k.*
(2) *If p is irrational and a fixed point of $f \in F$, then f is the identity on an open neighborhood of p.*

For example, $\frac{1}{3}$ can be an isolated fixed point of an element of F, but the slope of such an element will have to be an integral power of 4 since $k = 2$ is the least k for which $3|(2^k - 1)$. For $\frac{1}{5}$, $k = 4$, and for $\frac{1}{7}$, $k = 3$, and so on. For $\frac{1}{p}$ with p a prime, k is always a divisor of $p - 1$.

4.2. Wreath products. In groups of homeomorphisms, wreath products and iterated wreath products abound. One reason for pointing this out is that wreath products have non-cyclic abelian subgroups and thus cannot be subgroups of free groups. We give definitions.

DEFINITION 4.4. If G and H are groups and H acts on a set X on the right, then the *(restricted) wreath product* $G \wr H$ (the notation neglects ingredients) is the semidirect product

$$\left(\sum_{x\in X} G\right) \rtimes H$$

where the sum ΣG on the left of copies of G can be thought of as those functions f from X to G for which only finitely many $f(x)$ are not the identity in G, and where $(fg)x = f(x)g(x)$. Then for $h \in H$, we have $f^h(x) = (h^{-1}fh)(x) := f(xh^{-1})$. Elements of $G \wr H$ can be written as fh with $(f, h) \in \Sigma G \times H$ and multiplication done as $(f_1h_1)(f_2h_2) = f_1h_1f_2h_1^{-1}h_1h_2 = f_1f_2^{h_1^{-1}}h_1h_2$. The *unrestricted wreath product* would replace the sum of copies of G by the Cartesian product of copies of G. In the case that H is acting on itself by right multiplication, then the wreath products (restricted and unrestricted) are called standard.

The paper [160], some of whose left-right conventions differ from ours, is the standard reference on standard wreath products.

DEFINITION 4.5. In the case that G is also acting on a set Y on the right, then there is a natural right action of $G \wr H$ on $Y \times X$. In this case, the resulting group and its action are referred to as a *permutation wreath product* and the action is defined as follows. Using the letters already established and with $(y, x) \in Y \times X$, we have $(y, x)(fh) = (yf(x), xh)$. The reader can check that this gives a consistent right action.

4.2.1. *Realization.* We will rarely have occasion to refer to the details of the definition of the wreath product. The following illustrates how wreath products will arise, how they can be iterated, and gives a hint to the associativity of the permutation wreath product. The reader can come up with easy examples using finite groups that show that the standard wreath product is not associative.

DEFINITION 4.6. Let H be a group acting on a set A on the right, let B be a subset of A, and let $X = \{Bh \mid h \in H\}$. Note that H also acts on X. We say that the action of H on X is *consistent* to mean that for all $h \in H$ if $Bh \cap B \neq \emptyset$, then h fixes B pointwise. We say that the action of H on X is *faithful* to mean that the only element of H that fixes all elements of X is the identity of H.

The following lemma allows us to recognize wreath products when they occur without referring to the details of Definitions 4.4 and 4.5. The proof is left to the reader.

LEMMA 4.7. *Suppose that G and H act on a set A on the right. Assume there are subsets $C \subseteq \text{supp}(G) \subseteq B \subseteq A$ such that the action* *of G on $Y = \{Cg \mid g \in G\}$ and the action of H on $X = \{Bh \mid h \in H\}$ are both consistent and faithful. Then the action of $W = \langle G, H\rangle$ on $Z = \{Cw \mid w \in W\}$ is also consistent and faithful and is isomorphic to the action of the permutation wreath product $G \wr H$ on $Y \times X$.*

4.2.2. *Existence.* The basic lemma is the following. The proof is left to the reader.

LEMMA 4.8. *Let $h \in H \leq \mathrm{Homeo}_+(\mathbf{R})$ be given with the closure of $\mathrm{supp}(h)$ a compact subset of an orbital of H. Then there is an $f \in H$ with $\langle h, f\rangle$ isomorphic to $\langle h\rangle \wr \langle f\rangle$ the standard wreath product $\mathbf{Z} \wr \mathbf{Z}$. In particular H contains a subgroup isomorphic to the direct sum of countably many copies of $\mathbf{Z}$ and H is not free.*

4.3. The shrinking commutator. The following drives results in Sections 4.5, 4.6, and 5.5.

LEMMA 4.9. *If f and g are in $PL(\mathbf{R})$ then the following hold.*

(1) *If f and g have a common fixed point t, then $[f, g]$ is fixed on an open neighborhood of t.*
(2) *If f and g are in $PLF(\mathbf{R})$, then $[f, g]$ has slope 1 outside of a compact set.*
(3) *If f and g both have slope 1 outside of a compact set, then the closure of $\mathrm{supp}([f, g])$ is a compact subset of $\mathrm{supp}(f) \cup \mathrm{supp}(g)$.*

PROOF. The first two statements follow from the chain rule. The third follows from the commutativity of translations, from the first statement, and from the fact that $\mathrm{supp}(f) \cup \mathrm{supp}(g) = \mathrm{supp}(\langle f, g\rangle)$. ☐

4.4. Absence of torsion. For an $f \in \mathrm{Homeo}_+(\mathbf{R})$, the proof of the following breaks into two symmetric cases depending on whether one starts with $t \in \mathbf{R}$ where $tf > t$ or $tf < t$. Details are left to the reader.

LEMMA 4.10. *The only element of finite order in $\mathrm{Homeo}_+(\mathbf{R})$ is the identity.*

4.5. Absence of free subgroups. That there are no non-abelian free subgroups in the groups we work with should follow immediately from Lemma 4.8 which guarantees wreath products in very general situations. The catch is that Lemma 4.8 assumes an h with support in one orbital of the ambient group. A certain amount of work with the Thompson groups revolves around cleaning up conflicts among multiple orbitals. This is an early example. The assumption that there be only finitely many breakpoints is necessary. See Corollary 12.14.1.

Theorem 4.11. *There is no subgroup of $PLF_+(\mathbf{R})$ isomorphic to the free group of rank 2 and thus no subgroup of either $PL_+(I)$ or F is isomorphic to the free group of rank 2. In particular, the following dichotomies hold. A subgroup of $PLF_+(\mathbf{R})$ is either metabelian or contains a subgroup ismorphic to $\mathbf{Z} \wr \mathbf{Z}$. A subgroup of $PL_+(I)$ is either abelian or contains a subgroup isomorphic to $\mathbf{Z} \wr \mathbf{Z}$.*

Proof. Let $PLF_1(\mathbf{R})$ be the subgroup of $PLF_+(\mathbf{R})$ consisting of all elements that have slope 1 outside of a compact set. By Lemma 4.9, the commutator subgroup $PLF'_+(\mathbf{R})$ of $PLF_+(\mathbf{R})$ is contained in $PLF_1(\mathbf{R})$. Since $F \leq PL_+(I) \leq PLF_1(\mathbf{R})$, the claims all follow if we show that a non-abelian subgroup of $PLF_1(\mathbf{R})$ contains a subgroup isomorphic to $\mathbf{Z} \wr \mathbf{Z}$.

We assume that there are elements f and g of $PLF_1(\mathbf{R})$ that do not commute. As elements of $PLF_+(\mathbf{R})$, all of f, g and $H = \langle f, g \rangle$ have finitely many orbitals. By Lemma 4.9 the closure of $\mathrm{supp}([f, g])$ is a compact subset of $\mathrm{supp}(H)$. Thus the set of elements in H whose closure of support is a compact subset of $\mathrm{supp}(H)$ is not empty. Let w be a non-trivial element of H so that the closure of $\mathrm{supp}(w)$ is a compact subset of $\mathrm{supp}(H)$ and intersects non-trivally the smallest number of orbitals of H. Let J be an orbital of H that contains points in $\mathrm{supp}(w)$ and let H_J consist of the restriction to J of the elements of H.

By Lemma 4.8, there is an $x \in H$ with $\langle w|_J, x|_J \rangle \simeq \langle w|_J \rangle \wr \langle x|_J \rangle$ the wreath product $\mathbf{Z} \wr \mathbf{Z}$. Sending each $h \in H$ to $h|_J$ is a surjective homomorphism ϕ from H to H_J. If the restriction of ϕ to $G = \langle w, x \rangle$ has trivial kernel, then G is the subgroup of H that we seek. So we assume some non-identity $y \in \langle w, x \rangle$ in the kernel of ϕ.

The abelianization of $\langle w|_J, x|_J \rangle \simeq \mathbf{Z} \wr \mathbf{Z}$ is $\mathbf{Z} \times \mathbf{Z}$ generated by the images in the abelianization of $w|_J$ and $x|_J$. It follows that the exponent sum of each of w and x in y is zero. Moving appearances of w and x that are not in commutators to the right end of y leaves behind nothing but commutators, putting y in the commutator subgroup of $\langle w, x \rangle$. Once again Lemma 4.9 has the closure of $\mathrm{supp}(y)$ a compact subset of $\mathrm{supp}(H)$. But commutators of w and x can only have support in orbitals of H that also contain points of $\mathrm{supp}(w)$. So the support of y, which misses J, lies in fewer orbitals of H than the support of w. This contradicts our choice of w. ☐

4.6. Nilpotent implies abelian.

Proposition 4.12. *Let H be a nilpotent subgroup of $PLF(\mathbf{R})$. Then H is abelian.*

Proof. We assume that H is nilpotent and not abelian. Thus in the lower central series,

$$H = H_0 \supseteq H_1 \supseteq H_2 \supseteq \cdots \supseteq H_{n-2} \supseteq H_{n-1} \supseteq H_n = \{1\}$$

we know that $H_1 = [H, H]$ is not the trivial subgroup and we know that for every $h \in H_{n-1} \subseteq H_1 = [H, H]$ and $w \in H$ that $[w, h] = 1$. Thus we are done if we show that for every non-trivial element h of $H_1 = [H, H]$ there is an element w of H for which $[w, h]$ is non-trivial.

Take $1 \neq h \in [H, H]$ and let I be an orbital of H containing a non-fixed point of h. If I is bounded to the left, then let a be the left endpoint of I. By Lemma 4.9, the support of h in I is bounded away from a. Let b be the left endpoint in I of the support of h in I. We have $b \neq a$ and $b \in I$. There is an element $w \in H$ for which $bw \neq b$. Thus the conjugate of h by w will have support in I different from that of h and we have that h and w do not commute.

Assume that I is not bounded to the left. The argument of the previous paragraph finishes the argument if the support of h is bounded to the left, so we assume that it is not. Since h is the product of finitely many commutators, Lemma 4.9 says that h must be a translation by some s near $-\infty$. We have $s \neq 0$ since the support of h is not bounded on the left. If all elements of H were translations near $-\infty$, then h would be the identity near $-\infty$. Thus some $w \in H$ behaves as $x \mapsto x\lambda + k$ near $-\infty$ with $\lambda \neq 1$. The action of w^{-1} near $-\infty$ is as $x \mapsto x\lambda^{-1} - k\lambda^{-1}$ and the action of $w^{-1}hw$ near $-\infty$ is as

$$x \mapsto (x\lambda^{-1} - k\lambda^{-1} + s)\lambda + k = x - k + s\lambda + k = x + s\lambda.$$

Since $\lambda \neq 1$ and $s \neq 0$, this is not the behavior of h near $-\infty$ and $[w, h] \neq 1$. □

Theorem 4.5 of Plante-Thurston 1976 [169] proves the above for C^2 homeomorphisms of $\mathbf{R}_{\geq 0}$. From Proposition 9.11 below, there is a smooth action of F on $\mathbf{R}_{\geq 0}$ giving another argument for Proposition 4.12.

There are non-abelian nilpotent subgroups of $PL_+(\mathbf{R})$ of every nilpotency class. Let s be the shift map $xs = x + 1$ and let f be any non-identity element of $PL_+(\mathbf{R})$ with support in $(0, 1)$. For $i \in \mathbf{Z}$, let $f_i = (f)^{s^i}$. Let $\phi : \mathbf{Z} \to \mathbf{Z}$ be any function and let f_ϕ be the product over all $i \in \mathbf{Z}$ of $f_i^{\phi(i)}$. This product makes sense and gives an element of $PL_+(\mathbf{R})$ since the supports of the f_i are pairwise disjoint. The reader can show that judicious choices of ϕ can make $G_\phi = \langle s, f_\phi \rangle$ nilpotent of any desired nilpotency class. Note that G_ϕ will be isomorphic to a subgroup of the unrestricted wreath product of $\mathbf{Z}$ by $\mathbf{Z}$.

4.7. Local indicability and order. There is much written about ordered groups. Here we introduce a bare minimum and start with the relevant definitions.

DEFINITION 4.13. If $<$ is a total order on a group G, it is called a *left order* (or a left invariant order) if for all a, b and c in G with $a < b$, we have $ca < cb$. It is called a *right order* (or right invariant order) if instead one can conclude $ac < bc$. If both conclusions hold, then the order is a *bi-order* (or bi-invariant order). A group is *locally indicable* if every finitely generated subgroup has a homomorphism to $\mathbf{Z}$ with non-trivial image.

Note that true statements about groups with a left order can be turned into true statements about groups with a right order and vice-versa by exchanging the words left and right.

In a group G with a total order $<$, an element x is *positive* if $x > 1$. If the order is any of left, right or bi-invariant, then $P = \{x \mid x > 1\}$ is closed under multiplication, and for every $1 \neq y \in G$ exactly one of $y \in P$ or $y^{-1} \in P$ will be true. If the order is a bi-order, then P is closed under conjugation by arbitrary elements of G. We have the following.

LEMMA 4.14. *Let a group G be the disjoint union $G_- \cup \{1\} \cup G_+$ where G_+ is a semigroup invariant under conjugation in G and G_- is the set of inverses of G_+. Then setting $f < g$ if and only if $gf^{-1} \in G_+$ gives a bi-order on G.*

PROOF. Trichotomy follows directly from the definition. For transitivity we have $gf^{-1} \in G_+$ and $hg^{-1} \in G_+$ gives $hg^{-1}gf^{-1} = hf^{-1} \in G_+$, so we have a total order. Now if $gf^{-1} \in G_+$ then $gh(fh)^{-1} = gf^{-1} \in G_+$ and $hg(hf)^{-1} = (gf^{-1})^{h^{-1}} \in G_+$ as required by the definition. $\square$

A left order is *Conradian* if for all positive f and g, there is an $n \in \mathbf{N}$ so that $fg^n > g$.

We have a set of implications. A bi-order is Conradian. This is immediate since $f > 1$ implies $fg > g$, and so $n = 1$ works in the definition of Conradian. A group has a Conradian order if and only if it is locally indicable. See for example Propositions 3.11 and 3.16 in Navas 2010 [158]. As a consequence, a locally indicable group has a left order. A countable group has a left order if and only if it admits a faithful left action on $\mathbf{R}$. See for example Theorem 6.8 of Ghys 2001 [81].

There are no converses to the two implications above not given as equivalences. It is straightforward to show that if G and H are locally indicable, then so is an extension of G by H. Thus $G = BS(1,-1) = \langle a, b \mid a^{-1}ba = b^{-1}\rangle$ as a semidirect product of $\mathbf{Z}$ and $\mathbf{Z}$ is locally indicable. But G cannot admit a bi-order since b and b^{-1} are conjugate. Bergman 1991 [10] has a right orderable but non locally indicable group.

Every subgroup of $BPL(\mathbf{R})$ has a bi-order and is thus locally indicable. But in the setting of $BPL(\mathbf{R})$ it is so easy to show both that we shall do so directly.

PROPOSITION 4.15. *If $G \le BPL(\mathbf{R})$, then G has a bi-order and is locally indicable.*

PROOF. We use Lemma 4.14. We let G_+ consist of those $h \in G$ so that if $t = \inf(\mathrm{supp}(h))$, then there is an $\epsilon > 0$ so that for all $s \in (t, t+\epsilon)$ we have $sh > s$. Letting G_- be the set of inverses of the elements of G_+ gives all the right properties.

For local indicability, we assume G is finitely generated and thus has bounded support. Let $t = \inf(\mathrm{supp}(G))$. For $h \in G$, let $\phi(h) = \log(th'_+)$ where th'_+ is the right hand derivative of h at t. This is a homomorphism to $(\mathbf{R}, +)$ that is not the trivial homomorphism since G is finitely generated. The image is a finitely generated, torsion free abelian group and is the direct sum of a finite number of copies of $\mathbf{Z}$. Composing ϕ by a projection to one of the summands satisfies the definition of local indicability. ☐

5. Properties, II: from transitivity

The action of F has strong transitivity properties, and many of the results about F are based on these properties.

5.1. The transitivity lemmas. We look at the level of transitivity of F. It turns out to be useful to work with points outside of the unit interval I even when discussing the action of F on I. So our statements of transitivity will cover more than the unit interval.

DEFINITION 5.1. If a group G acts on a linearly ordered set L by order preserving bijections, then we say that the action is *order n-transitive* or o-n-transitive if for every pair of subsets A and B of L with $|A| = |B| = n$, there is an element $g \in G$ with $Ag = B$. When $n = 1$, this is simply a definition of *transitive*.

To be more precise about where we can find elements in the next lemma, we define $BPL_2(\mathbf{R})$ to be the subgroup $BPL(\mathbf{R})$ consisting of those elements that are F-like in that their slopes are integral powers of 2, that have their breakpoints are in $\mathbf{Z}[\frac{1}{2}]$, and that take $\mathbf{Z}[\frac{1}{2}]$ to itself.

This will be shown in Corollary 10.1.1 to be the commutator subgroup of an isomorphic copy of F acting on $\mathbf{R}$.

LEMMA 5.2. *The action of $BPL_2(\mathbf{R})$ on $\mathbf{Z}[\frac{1}{2}]$ is o-n-transitive for every positive integer n. Specifically, if A and B are finite subsets of $\mathbf{Z}[\frac{1}{2}]$ of the same size and there are $p < q$ in $\mathbf{Z}[\frac{1}{2}]$ with $A \cup B \subseteq (p, q)$, then there is an $f \in BPL_2(\mathbf{R})$ taking A onto B that is the identity off (p, q).*

PROOF. Let A, B, p and q be as in the statement. By increasing p and decreasing q slightly, we can assume both are in $\mathbf{Z}[\frac{1}{2}]$.

We let $\{a_i \mid 0 \leq i \leq n+1\}$ be the elements of $A \cup \{p, q\}$ indexed in increasing order. This puts $a_0 = p$ and $a_{n+1} = q$. Similarly we let $\{b_i \mid 0 \leq i \leq n+1\}$ be the elements of $B \cup \{p, q\}$ putting $b_0 = p$ and $b_{n+1} = q$.

We express each element of $A \cup \{p, q\}$ as a fraction in reduced terms and let k be the largest exponent of 2 found in a denominator over all the elements of $A \cup \{p, q\}$. We choose l similarly for $B \cup \{p, q\}$.

We now replace A by A' which consists of all integer multiples of 2^{-k} in the closed interval $[p, q]$ and we let $\{a'_i \mid 0 \leq i \leq M\}$ be the elements of A' indexed in increasing order in $[p, q]$. Similarly, we replace B by B' consisting of all integer multiples of 2^{-l} in $[p, q]$ and we let $\{b'_i \mid 0 \leq i \leq N\}$ be the elements of B', also indexed in increasing order. Note that $A \subseteq A'$ and $B \subseteq B'$. Note also that intervals of the form $[a'_i, a'_{i+1}]$, $0 \leq i < M$, and intervals of the form $[b'_i, b'_{i+1}]$, $0 \leq i < N$, have lengths that are integral powers of 2.

For i with $0 \leq i < n+1$, consider the intervals $[a_i, a_{i+1}]$ defined from the elements of $A \cup \{p, q\}$, and $[b_i, b_{i+1}]$ defined from the elements of the elements of $B \cup \{[p, q\}$. Let c_i be the number of elements of A' in the interval $[a_i, a_{i+1}]$, and let d_i be the number of elements of B' in the interval $[b_i, b_{i+1}]$. If say $c_i < d_i$, then apply a sequence of $d_i - c_i$ modifications to A' consisting of adding to A' the midpoint $(a'_j + a'_{j+1})/2$ of two elements a'_j and a'_{j+1} that reside in $[a_i, a_{i+1}]$ and renumbering so that the indexing reflects the order. Each such modification increases by 1 the number of elements of A' in the interval $[a_i, a_{i+1}]$ and preserves the fact that consecutive elements of A' are an integral power of 2 apart. If the non-specificity of where to make each modification seems bothersome, then one could insist that a'_j is always chosen so that $a'_j = a_i$. The case $d_i < c_i$ has symmetric treatment.

Doing this for all i with $0 \leq i < n+1$ lets us define a homeomorphism $f : \mathbf{R} \to \mathbf{R}$ by taking each $[a'_i, a'_{i+1}]$ to $[b'_i, b'_{i+1}]$ affinely, and which is the identity on both $(-\infty, p]$ and $[q, \infty)$. Since each $[a'_i, a'_{i+1}]$ and $[b'_i, b'_{i+1}]$ has length an integral power of 2, all slopes are integral powers of 2. The a_i' contain all breaks of f and are all dyadic. The function f has been constructed to take the points in A to those of B and is the function we seek. ☐

COROLLARY 5.2.1. *For every integer $n \geq 1$, the group F is o-n-transitive on $\mathbf{Z}[\frac{1}{2}] \cap (0,1)$.*

COROLLARY 5.2.2. *The group F is C^0-dense in* $\text{Homeo}_+(I)$.

PROOF. Since I is compact and thus all elements of F are uniformly continuous, we can interpret C^0-dense as dense under the metric

$$d(f,g) = \sup_{t \in I} |tf - tg|.$$

The claim is now immediate from Lemma 5.2. ☐

Lemma 5.2 is only a warm up for the lemma we really need. The restrictions on elements of F in Definition 3.1 are local restrictions. The next lemma codifies what we need of this notion and expands on the lemma above. The reader can supply the proof.

LEMMA 5.3. *Let A and B be locally finite subsets of $\mathbf{Z}[\frac{1}{2}]$ with $|A| = |B|$, and let X be a possibly empty collection of closed intervals with pairwise disjoint interiors in I so that each $J \in X$ has its endpoints in A and no points of A in its interior. Let ϕ be an order preserving bijection from A to B, and for each $J \in X$ let g_J be an element of $PL_+(\mathbf{R})$ whose restriction to J satisfies (1) and (2) of Definition 3.1, and which agrees with ϕ on the endpoints of J. Then there is an element $f \in PL_2(\mathbf{R})$ so that f agrees with ϕ on A, and for each $J \in X$, f agrees with g_J on J. If A and B are finite so that $A \cup B \subseteq (p,q)$ with p and q in $\mathbf{Z}[\frac{1}{2}]$, then f can be chosen to be the identity off (p,q).*

Sections 5.2 through 5.5 exploit the transitivity properties of the Thompson groups.

5.2. Trivial center. In the proof of Proposition 4.12, certain elements are shown to not commute. Transitivity can be used to show that even more elements do not commute.

LEMMA 5.4. *Let P be the group of permuations on some X with a subgroup G of P so that for any $x \neq y$ in X, there is a $f \in G$ with $xf = x$ and $yf \neq y$. Then the centralizer of G in P is trivial.*

PROOF. Given $h \in P \setminus \{1\}$, we have $xh \neq x$ for some x, and h cannot commute with an element that fixes x but not xh. ☐

COROLLARY 5.4.1. *The centralizer of F in* $\text{Homeo}(I)$ *is trivial. In particular, the center of F is trivial.*

PROOF. This follows from Lemmas 5.3 and 5.4 and the density of the dyadics in $(0,1)$. $\square$

The reader can show that the set of elements in F whose support has closure in $(0,1)$ is a normal subgroup of F that intersects non-trivially each non-trivial normal subgroup of F.

5.3. Many copies of F. We will often use the notation $F_{[a,b]}$ with both a and b in $\mathbf{Z}[\frac{1}{2}]$ and $a<b$ to denote the set of all PL self homeomorphisms of $[a,b]$ that satisfy (1) and (2) in Definition 3.1. The points a and b are from $\mathbf{Z}[\frac{1}{2}]$ in $\mathbf{R}$ and not necessarily confined to the unit interval.

LEMMA 5.5. *With $a<b$ in $\mathbf{Z}[\frac{1}{2}]$ the set $F_{[a,b]}$ under composition is isomorphic to F.*

PROOF. By Lemma 5.2 there is a homeomorphism g of $BPL_2(\mathbf{R})$ that carries $[0,1]$ to $[a,b]$. By the definitions of F and $BPL_2(\mathbf{R})$, the chain rule, and the fact that g and g^{-1} carry $\mathbf{Z}[\frac{1}{2}]$ to itself, the set $F_{[a,b]}$ is exactly the set $F^g=\{f^g \mid f\in F\}$. $\square$

Note that any increasing bijection from $[0,1]$ to $[a,b]$ will conjugate F to a copy of F acting on $[a,b]$, but its elements will not necessarily resemble elements of F. Even if the bijection is affine the breakpoints will not necessarily be dyadics, although the resulting slopes will be integral powers of 2. We will take another look at $F_{[0,n]}$ for a positive integer n in Proposition 11.13.

The proofs of the next two lemmas are left to the reader.

LEMMA 5.6. *Let $a<c<b<d$ all be in $\mathbf{Z}[\frac{1}{2}]$. Then $F_{[a,d]} = \langle F_{[a,b]}\cup F_{[c,d]}\rangle$.*

LEMMA 5.7. *If $0<a<c<d<b<1$ with all of a,b,c,d in $\mathbf{Z}[\frac{1}{2}]$. then there is an isomorphism from F to $F_{[a,b]}$ that is the identity on $F_{[c,d]}$.*

5.4. Direct sums and wreath products. Theorem 4.11 shows certain groups are not subgroups of F. We briefly show a few groups that are. First a seed lemma.

LEMMA 5.8. *If a group G is isomorphic to a subgroup of F, then the standard wreath product $G\wr\mathbf{Z}$ is isomorphic to a subgroup of F. In particular the infinite direct sum $\sum_{i\in\mathbf{Z}}\oplus G$ is isomorphic to a subgroup of F.*

PROOF. The proof is almost identical to the proof of Lemma 4.8.

Let G be isomorphic to a subgroup of F. Let h be a non-trivial element of F and by replacing h by its inverse if necessary, we can assume there is $t \in \mathbf{Z}[\frac{1}{2}]$ so that $t < th$. We can use Lemma 5.5 to find an isomorphic copy G_0 of G in F with support in $[t, th]$. Now $\langle G_0, h\rangle$ is isomorphic to $G \wr \mathbf{Z}$ by Lemma 4.7. From Definition 4.4, the promised infinite direct sum is a subgroup of $G \wr \mathbf{Z}$. □

Corollary 5.8.1. *Thompson's group F has subgroups isomorphic to $\mathbf{Z} \wr \mathbf{Z}$, $F \wr \mathbf{Z}$, the infinite direct sum of copies of $\mathbf{Z}$, and the infinite direct sum of copies of F. There are also subgroups of F isomorphic to the iterated wreath products $(\cdots((\mathbf{Z}\wr\mathbf{Z})\wr\mathbf{Z})\wr\cdots\wr\mathbf{Z})$ of all finite lengths.*

5.5. On the commutator and other normal subgroups. The groups in the Thompson family tend to have "few" normal subgroups, with many being simple. In any group, a subgroup containing the commutator subgroup is normal. Here we will prove that the normal subgroups of F are exactly the subgroups containing the commutator subgroup, and that the commutator subgroup itself is simple. In particular every proper quotient of F is abelian.

The map $\phi : F \to \mathbf{Z} \times \mathbf{Z}$ defined by $\phi(f) = (\log_2(0f'_+), \log_2(1f'_-))$ is a surjective homomorphism. Since $\mathbf{Z} \times \mathbf{Z}$ is abelian, the kernel K contains the commutator subgroup F' of F, and K consists of all those $f \in F$ for which there is an open U in $[0, 1]$ about 0 and 1 on which f is the identity. From Lemma 4.9, this gives a second reason why $F' \leq K$. In fact $K = F'$, but this will have to wait until Section 10.1 when we know a presentation for F.

The transitivity properties in Lemma 5.3 give the following.

Lemma 5.9. *For $i \in \{1, 2, 3, 4\}$, let $J_i = [x_i, y_i]$ have non-empty interior, and for $i \in \{1, 2, 3\}$, assume $[x_i, y_i] \subseteq (x_{i+1}, y_{i+1})$. Then*

(i) *there is an $f \in BPL_2(\mathbf{R})$ fixed on J_1 and off J_4 with $J_3f \subseteq J_2$, and*

(ii) *there is a $g \in BPL_2(\mathbf{R})$ fixed off J_2 so that $J_1g \cap J_1 = \emptyset$.*

Lemma 5.10. *For $f \in K$ with support of f in $J = [x, y] \subseteq (0, 1)$, let N be the normal closure of f in F. Then there is a commutator $c \in F' \cap N$ with $c|_J = f|_J$.*

Proof. From Lemma 5.9(ii) there is a $g \in F$ with $Jg \cap J = \emptyset$. Now $c = [g, f] = g^{-1}f^{-1}gf = (f^{-1})^g f$ has the claimed property. □

Corollary 5.10.1. *For every integer $n \geq 1$, both K and the commutator subgroup of F are o-n-transitive on $\mathbf{Z}[\frac{1}{2}] \cap (0, 1)$.*

Proof. For $|A| = |B| = n$ in $\mathbf{Z}[\frac{1}{2}] \cap (0, 1)$, we put $A \cup B$ in the interior of a closed interval $J = [p, q] \subseteq (0, 1)$ and the claim for K

follows by adding $\{p, q\}$ to both A and B and using Lemma 5.2. The claim for F' follows now from Lemma 5.10. $\square$

LEMMA 5.11. *Every commutator of elements of F is a commutator of elements of K, and so $F' \leq K'$ giving $F' = K'$.*

PROOF. With $K \leq F$, we have $K' \leq F'$.

Let $c = [f, g]$ be a commutator in F. From Lemma 4.9, the closure of the support of c is contained in some $[a, b]$ with dyadic endpoints and with $[a, b] \subseteq (0, 1)$. We choose dyadic rationals p, q, with

$$[a, b] \subseteq (p, q) \subseteq [p, q] \subseteq (0, 1).$$

Lemma 5.9 gives an $h \in BPL_2(\mathbf{R})$ that is the identity on $[a, b]$ and taking $[0, 1]$ into $[p, q]$. Now $c = c^h = [f^h, g^h]$ is a commutator of two elements in K and we have shown $F' \leq K'$. $\square$

We prove our main claims by using techniques of Higman from [106]. One immediate consequence is that F is Hopfian. A group G is *Hopfian* if every surjection from G to itself is an automorphism.

PROPOSITION 5.12. *The following hold in F.*

1. *Every proper quotient of F is abelian.*
2. *F' is simple.*
3. *F is Hopfian.*

PROOF. Let f, g be in F and let $c = [f, g]$. We wish to show that any non-trivial normal subgroup of F contains c. That is, we wish to show that the normal closure of any non-identity element in F contains c. We have been ambiguous about where the normal closure is taken and we will play with that during the proof.

From Lemma 5.11 we know that $c = [f_1, g_1]$ with f_1 and g_1 elements of K. Let h be a non-identity element and let U be an open subset with closure in $(0, 1)$ so that $Uh \cap U = \emptyset$.

The union of the supports of f_1 and g_1 is contained in some closed interval J in $(0, 1)$ with dyadic endpoints and there is a $k \in F$ taking J into U. By Corollary 5.10.1, we may assume that k is in F'. Now $[f_1^k, g_1^{kh}] = 1$. Equivalently, $[f_1, g_1^{khk^{-1}}] = 1$. We have shown that f_1 and g_1 commute modulo a conjugate of h making $c = [f, g]$ trivial modulo h. For the faint of heart, we offer the following. Let $j = khk^{-1}$ and N the normal closure (again ambiguous) of h. We note that

$$\begin{aligned}[f_1, g_1]^{-1}[f_1, g_1^j] &= [g_1, f_1][f_1, g_1^j] \\ &= \left(g_1^{-1}\left(\left(f_1^{-1}\left((g_1 j^{-1} g_1^{-1}) j\right) f_1\right) j^{-1}\right) g_1\right) j\end{aligned}$$

where each parenthesized group is in N and thus so is the entire expression. So $[f_1, g_1]$ and $[f_1, g_1^j]$ are equal modulo N. But we have shown $[f_1, g_1^j]$ to be in N (and in fact trivial).

At this point, we have shown that every proper quotient of F is abelian.

Now we take into account that the element k can be taken from F' and we could have started with a non-identity h from F'. This shows that every element of F' is in the normal closure in F' of any non-identity element of F'. This shows that F' is simple.

That F is Hopfian follows from (1) and the fact that F is not abelian. $\square$

We do not know at this point the true fact that $F' = K$ and that the abelianization of F is $\mathbf{Z} \times \mathbf{Z}$. This needs the fact that x_0 and x_1 generate F and more combinatorial arguments.

It might be tempting to argue that $F' = K$ by noting that every element of K has support bounded away from the endpoints of $[0, 1]$ and guessing that each such might be a commutator. But an open question about F is whether all elements of F' are single commutators. However, the following is known. The argument goes back to [57].

Proposition 5.13. *(I) Let G act in an order preserving way on $I = [0, 1]$, so that every element of G has support bounded away from 0 and 1, and so that for all $p < q$ in $(0, 1)$, there is an $f \in G$ with $pf > q$. Then every element of G' is a proeuct of two or fewer commutators.*

(II) Every element of F' is a product of two or fewer commutators.

Proof. By Lemma 5.11, (II) follows from (I). We focus on (I).

We wish to show that every product of three commutators is equal to a product of two commutators. To that end let $c_i = [a_i, b_i]$, $i \in \{1, 2, 3\}$, be in G'. Some $[p, q] \subseteq (0, 1)$ contains the supports of the a_i and b_i (and thus the c_i). Choose $f \in G$ so that $pf > q$ and we have

$$0 < qf^{-1} < p < q < pf < 1.$$

The supports of $\{a_1, b_1\}$, $\{a_2^f, b_2^f\}$ and $\{a_3^{f^{-1}}, b_3^{f^{-1}}\}$ are now in pairwise disjoint intervals and so c_1, c_2^f, $c_3^{f^{-1}}$ commute and their product is a commutator. These two properties give

$$\begin{aligned} c_1c_2c_3 &= c_1c_2^f c_3^{f^{-1}} \left((c_3^{-1})^{f^{-1}} c_2\right)\left((c_2^{-1})^f c_3\right) \\ &= \left(c_1c_2^f c_3^{f^{-1}}\right)\left(fc_3^{-1}f^{-1}c_2f^{-1}c_2^{-1}fc_3\right) \\ &= \left(c_1c_2^f c_3^{f^{-1}}\right)[c_2^{-1}c_3^{f^{-1}}, f], \end{aligned}$$

which is a product of two commutators. □

Lastly we show that F' is not finitely generated. This would be slightly easier if we knew $F' = K$, but transitivity makes it fairly direct.

PROPOSITION 5.14. *The commutator subgroup F' of F is not finitely generated.*

PROOF. Let c be any non-identity commutator in F. Let (p, q) be an orbital of c. We know $0 < p < q < 1$ and for each $\epsilon > 0$ there is an element of F taking p into $(0, \epsilon)$ and q into $(1-\epsilon, 1)$. Conjugating c by these elements shows that $\mathrm{supp}(F') = (0, 1)$. But every element of F' has support bounded away from 0 and 1. So every finitely generated subgroup of F' has support bounded away from 0 and 1, and no finite set in F' can generate F'. □

5.6. Centralizers. For a subset S of a group G it will be convenient to use $C_G(S)$ (or $C(S)$ if the group G is clear) for $\{f \in G \mid \forall h \in S, fh = hf\}$, the centralizer of S in G. For a single $f \in G$, we write $C_G(f)$ (or $C(f)$) for the centralizer of f in G.

The structure of the centralizer of an element is determined by the "true" orbitals of the element. From Lemma 4.3, a non-dyadic rational can be an isolated fixed point of an element of F. Such fixed points are "unimportant" in certain ways.

DEFINITION 5.15. For $f \in F$, let $R(f)$ be the set of isolated fixed points of f that are non-dyadic rationals. A *dyadic orbital* of f is a component of $\mathrm{supp}(f) \cup R(f)$.

LEMMA 5.16. *If $1 \neq f \in F$ has one dyadic orbital $J = (a, b)$, then the centralizer in $F_{[a,b]}$ of $f|_{[a,b]}$ is the maximal cyclic subgroup of $F_{[a,b]}$ that contains $f|_{[a,b]}$.*

PROOF. Note that the statement contains the claim that the maximal cyclic subgroup exists.

Let C be the centralizer in $F_{[a,b]}$ of $f|_{[a,b]}$. Let $\phi : F_{[a,b]} \to \mathbf{Z}$ be defined by $\phi(g) = \log_2(ag'_+)$ for $g \in F_{[a,b]}$. Assume $1 \neq g \in F_{[a,b]}$ is in the kernel of the homomorphism ϕ. Then g is the identity on some $[a, s]$ with $s \in (a, b)$, but $xg \neq x$ for every $x \in (s, t)$ for some $t \in (s, b)$. The point s must be dyadic which implies that $sf \neq s$. Now $g^f \neq g$ since g^f is the identity on $[a, sf]$ and $xg^f \neq x$ for every $x \in (sf, tf)$. So ϕ is one-to-one on C, and C is cyclic. A cyclic subgroup of $F_{[a,b]}$ containing $f|_{[a,b]}$ is abelian and is thus contained in C. □

PROPOSITION 5.17. *Let S be a finite subset of F whose elements have pairwise disjoint supports, and let A be the set of dyadic orbitals of the elements in S. Let B be the set of components with non-empty interior of* $\operatorname{fix}(S)$. *For each $(a,b) \in A$ and $f \in S$ whose support contains (a,b), let $g_{[a,b]}$ be the identity off $[a,b]$ and be the generator of the maximal cyclic subgroup of $F_{[a,b]}$ that contains $f|_{[a,b]}$. Then with $m = |A|$ and $n = |B|$, the centralizer $C(S)$ of S in F is a direct sum of m copies of $\mathbf{Z}$ and n copies of F as follows:*

$$C(S) = \sum_{(a,b)\in A} \langle g_{[a,b]} \rangle + \sum_{[a,b]\in B} F_{[a,b]}. \tag{5.1}$$

PROOF. Because of Lemma 5.16, all we need to argue is that if g centralizes S, then g fixes the end points of the intervals in $A \cup B$. There are only finitely many such endpoints and if g moves one of them, then the least such will demonstrate that $f^g \neq f$ for some $f \in S$. □

The following eliminates the second summand from (5.1).

COROLLARY 5.17.1. *Let S be a finite subset of F whose elements have pairwise disjoint supports, and let A be the set of dyadic orbitals of the elements in S. Then*

$$C(C(S)) = \sum_{(a,b)\in A} \langle g_{[a,b]} \rangle$$

where $g_{[a,b]}$ is as in Proposition 5.17.

PROOF. This follows from the fact that $S \subseteq C(S)$, from Proposition 5.17, from Lemma 5.5 which says that each $F_{[a,b]}$, $(a,b) \in B$, in (5.1) is isomorphic to F, and from Corollary 5.4.1 which says that the center of F is trivial. □

5.7. Growth. We take our first look at a metric aspect of F.

DEFINITION 5.18. If X is a generating set for a semigroup S, then the *norm* $\|s\|$ of an element $s \in S$ is the minimum n so that s is represented by a word in X of length n. For $n \in \mathbf{N}$, the n-ball $B_n(S,X)$ in S is the set $\{s \in S \mid \|s\| \leq n\}$. The *growth function* of S with respect to X is the function $\gamma_{S,X} : \mathbf{N} \to \mathbf{N}$ where $\gamma_{S,X}(n) = |B_n(S,X)|$. The growth function of a group G with respect to a finite set X of group generators for G is the growth function of G with respect to $X \cup X^{-1}$.

When X is finite, the growth function exhibits some level of independence of X. If X_1 and X_2 are two finite generating sets of a semigroup S, then for some m and n in $\mathbf{N}$, X_2 is contained in $B_m(S,X_1)$, and X_1 is contained in $B_n(S,X_2)$. Now $B_{mn}(S,X_1)$ contains all words of length n in X_2 and $B_{mn}(S, X_2)$ contains all words of length m in X_1. This gives

$$\gamma_{S,X_1}(mn) \geq \gamma_{S,X_2}(n), \text{ and}$$
$$\gamma_{S,X_2}(mn) \geq \gamma_{S,X_1}(m).$$

Letting $K = \max(m, n)$, we get

$$\gamma_{S,X_1}(Kn) \geq \gamma_{S,X_2}(n) \geq \gamma_{S,X_1}(K^{-1}n).$$

This leads naturally to the next definition.

DEFINITION 5.19. For $f, g : \mathbf{N} \to \mathbf{N}$, we write $f \preceq g$ to mean there is a $K > 1$ so that

$$f(n) \leq g(Kn)$$

for all sufficiently large n, and $f \sim g$ to mean $f \preceq g$ and $g \preceq f$.

The discussion preceeding Definition 5.19 gives the following.

LEMMA 5.20. *If X_1 and X_2 are two finite sets generating a group G, then the growth functions with respect to the two generating sets are equivalent.*

For various reasons that will not concern us, the definition commonly used varies from Definition 5.19. For example, different constant functions are not equivalent under Definition 5.19. We have already done some accomodating with the phrase "for sufficiently large n" since it is generally accepted that the concerns with a growth function are its long term behavior. See the Chapter VI of the book [100] for considerably more information and discussion.

Let $\exp_k(n) = k^n$. It is a consequence of Definition 5.19 that for any real $c > 1$ and $k > 1$ we have $\exp_c \sim \exp_k$.

DEFINITION 5.21. Let S be a semigroup with finite generating set X. We say that S has *exponential growth* if $\gamma_{S,X} \sim \exp_2$.

No growth function exceeds $\exp_2$. We have the following.

LEMMA 5.22. *Let S be a semigroup generated by a set X with k elements. Then $\gamma_{S,X} \preceq \exp_k$.*

PROOF. Let T be the free semigroup on X. Then for all $n \in N$, we have $|B_n(S, X)| \leq |B_n(T, X)|$. ☐

The next lemma shows that Definition 5.19 is not too liberal.

LEMMA 5.23. *Let S be an infinite semigroup with finite generating set X, and let K be a real number greater than 1. Then there is a finite generating set Y for S so that for all $n \in \mathbf{N}$, $\gamma_{S,Y}(n) \geq \gamma_{S,X}(Kn)$.*

PROOF. Let $Y = B_K(S, X)$ and let w be a word in X of length no more than Kn. We can write $w = w_1 w_2 \cdots w_n$ where each w_i has length no more than K. This puts w in $B_n(S, Y)$. □

Lemma 5.23 shows that any function in the class of a growth function can be exceeded by changing the generating set.

LEMMA 5.24. *Let $Q \subseteq P \subseteq S$ be semigroups with finite generating set Z, Y and X, respectively with $Y \subseteq X$. Then the following hold.*

(1) *For all n, $\gamma_{P,Y}(n) \leq \gamma_{S,X}(n)$.*
(2) *We have $\gamma_{Q,Z} \preceq \gamma_{S,X}$.*

PROOF. The first is clear. For the second, let $W = X \cup Z$ and we get $\gamma_{Q,Z} \preceq \gamma_{S,W} \sim \gamma_{S,X}$ from the first claim and Lemma 5.20. □

We can now apply all of the above to F.

LEMMA 5.25. *The free semigroup of rank 2 embeds in F. It follows that F has exponential growth.*

PROOF. From the transitivity lemmas of Section 5 it follows that there is an interval J and elements l and r so that Jl is the left half of J and Jr is the right half of J. If something specific is desired, $J = [\frac{1}{4}, \frac{3}{4}]$, $l = x_0 x_1^{-1} x_0^{-1}$ and $r = x_0 x_1^{-1}$ will do. Let $u = u_1 u_2 \cdots u_m$ and $v = v_1 v_2 \cdots v_n$ be different products of the elements of $\{l, r\}$. We want to show that $Jl \neq Jr$, and it suffices to assume that $u_m \neq v_n$. But under this assumption, Jl and Jr reside in different halves of J. □

A bit more work gets the following.

LEMMA 5.26. *[43, Theorem 4.6] The elements x_0^{-1} and $x_1^{\pm 1}$ of F generate a monoid isomorphic to the free product of $\mathbf{N}$ with $\mathbf{Z}$.*

6. Combinatorics, I: binary partitions

6.1. Words, the free monoid, and the Cantor set. We briefly introduce the tools we will work with. The free monoid of words over an alphabet organizes much of the combinatorics of the Thompson groups. The Cantor set has two uses. The compactness of the Cantor set simplifies an argument, and the Cantor set is a superb model of what the Thompson groups act on.

6.1.1. *Words.* In our applications, words are sequences of symbols from a finite, totally ordered alphabet. The order will be used to order the words over the alphabet.

DEFINITION 6.1. An *alphabet* is a finite totally ordered set A. The elements of A will be called letters or symbols. A *word* over A is a sequence in A either with domain some $n = \{0, 1, \dots, n-1\}$ (including $0 = \emptyset$, giving the empty word $\emptyset$), or with domain ω, the first infinite ordinal. A word u with domain n is finite and its length $|u|$ is n. A word that is not finite is infinite.

For an alphabet A, we use A^* to denote the set of all finite words (including the empty word $\emptyset$) over A and A^ω to denote the set of all infinite words over A.

Symbols in a word will be written without separators, so examples of finite words over $\{x, y\}$ would be $\emptyset$, x, y, xy, $xyxx$. We use subscripts (as in a_0) to denote the independent variable, and words are written from left to right with increasing subscript (as in $a_0a_1a_2 \dots a_{n-1}$ for a word of length n). We use adjacency to denote concatenation, and the concatenation uv (for u followed by v) makes sense for a finite word u and any word v. We leave formal definition of concatenation to the reader. We can also concatenate over sets. If u is a finite word, U a set of finite words, v a finite or infinite word, and V a set of finite or infinite words, then $uV = \{uv \mid v \in V\}$, $Uv = \{uv \mid v \in U\}$, and $UV = \{uv \mid u \in U, v \in V\}$ all make sense. We will often use the notation u^n with $u \in A$ or A^* which represents the concatenation of n copies of u, and u^ω which is the concatenation of infinitely many copies of u.

Definition 6.2. In uv, u is a *prefix* of uv and v is a *suffix* of uv. We will write $u \preceq uv$ and note that $\preceq$ is a partial order on A^*. As usual $u \prec v$ means $u \preceq v$ and $u \neq v$. We will write $u \perp v$ to mean that both $u \preceq v$ and $v \preceq u$ are false, and will say that u and v are *orthogonal*. If $u \perp v$, then there is a longest common prefix $u \wedge v$ of u and v. We extend $\preceq$ to a total order $\leq$ by writing $u < v$ if $u \prec v$, or if $u \perp v$ with $w = u \wedge v$, then there are a, b in A with $wa \preceq u$ and $wb \preceq v$, and we declare $u < v$ if and only if $a < b$. The resulting order $\leq$ is the usual prefix order defined for trees. See Definition 8.1.

6.1.2. *The free monoid.* Concatenation makes A^* a monoid with identity $\emptyset$, and we denote it $\mathfrak{M}_A$ when we wish to emphasize its structure as a monoid. The alphabet of heaviest use in this chapter will be $\{0, 1\}$, and for $\{0, 1\}$ we simply write $\mathfrak{M}$. The monoid $\mathfrak{M}_A$ is the free monoid over A since different words are different elements in $\mathfrak{M}_A$. We use A^+ to denote the set of all non-empty finite words over A and this forms the free semigroup over A to which we give no separate notation.

We will deal with more monoids than the free monoid and we will work with several properties of monoids. Here we list those properties that apply to the free monoid. More properties will be added when more monoids are added.

In the following, M will represent an arbitrary monoid.

DEFINITION 6.3. We say that M is *left cancellative* if $ab = ac$ always implies $b = c$. We say that M is *right cancellative* if $ab = cb$ always implies $a = c$. We say that M is *cancellative* if M is both right and left cancellative.

DEFINITION 6.4. If $a = bc$ in M, then b is a *left factor* for a. We say that d is a *common left factor* for e and f if d is a left factor of both e and f. We say that g is a *greatest common left factor* of h and k if it is a left factor of h and k and every common left factor of h and k is a left factor of g. We say that M *has greatest common left factors* if every pair of elements in M has a greatest common left factor.

DEFINITION 6.5. Right factors, (greatest) common right factors, and having greatest common right factors are defined by replacing "left" by "right" everywhere in Definition 6.4.

In $\mathfrak{M}_A$, left and right factor are synonymous, respectively, to prefix and suffix. The latter terms will often be used instead.

DEFINITION 6.6. A *unit* in a monoid M is an $x \in M$ with a $y \in M$ so that $xy = yx = 1$. We say M has *trivial units* if the only unit is 1. A *length function* for M is a homomorphism to $(\mathbf{N}, +)$ whose preimage of 0 is contained in the units of M.

LEMMA 6.7. *The free monoid $\mathfrak{M}_A = A^*$ is cancellative, has greatest common left and right factors, has trivial units, and the length of a word is a length function on $\mathfrak{M}_A$.*

6.1.3. *The Cantor set.* If A is finite, then the collection $\{uA^\omega \mid u \in A^*\}$ forms a basis for a topology on A^ω that is homeomorphic to the Cantor set. We use $\mathfrak{C}_A$ to denote A^ω with this topology and use $\mathfrak{C}$ when $A = \{0, 1\}$. The notation $\mathfrak{C}_n$ uses the set convention $n = \{0, 1, \ldots, n-1\}$. With $u \in A^*$, we will often use $u\mathfrak{C}$ as an equivalent to uA^ω to denote a typical basic open set in $\mathfrak{C}$. We will refer to $u\mathfrak{C}$ as the *cone* at u.

6.2. Binary partitions.

6.2.1. *Definitions and basics.*

DEFINITION 6.8. For integers $n \geq 0$ and $0 \leq i < 2^n$, a *dyadic interval* in $I = [0, 1]$ is an interval of the form

$$J = \left[\frac{i}{2^n}, \frac{i+1}{2^n}\right] \text{ with midpoint } m(J) = \frac{2i+1}{2^{n+1}}. \tag{6.1}$$

With J as in (6.1), we set

$$J0 = \left[\frac{i}{2^n}, m(J)\right], \text{ and } J1 = \left[m(J), \frac{i+1}{2^n}\right],$$

and are each said to be obtained from J by a *binary splitting*. In the above, we regard 0 and 1 as operators on the right and extend this to elements of $\mathfrak{M} = \{0,1\}^*$ recursively by setting $J\emptyset = J$, $J(w0) = (Jw)0$ and $J(w1) = (Jw)1$.

The next lemma is left to the reader.

LEMMA 6.9. *The following hold.*

(1) *For J a dyadic interval, $J = J0 \cup J1$, the intersection $J0 \cap J1$ contains only $m(J)$, and $J < J0 < J1$ under the order on intervals from Remark 3.4 augmented (anti-intuitively) by declaring $A < B$ for a closed interval A that properly contains a closed interval B.*
(2) *For intervals J and K in I, an increasing, continuous function takes J onto K affinely, if and only if it takes $J0$ affinely onto $K0$ and $J1$ affinely onto $K1$.*
(3) *The function that sends $w \in \mathfrak{M}$ to Iw is a bijection from $\mathfrak{M}$ to the set of dyadic intervals of I for which $Iw \subseteq Iv$ if and only if v is a left factor of w, and for which Iw and Iv have disjoint interiors if and only if $w \perp v$, in which case $w < v$ implies $Iw < Iv$.*
(4) *Two dyadic intervals in I either have disjoint interiors or are nested.*
(5) *The midpoint function m of (6.1) is a bijection from the set of dyadic intervals of I to the set of dyadics in $(0,1)$.*
(6) *If $J \subseteq L$ are dyadic intervals with $J \neq L$, then J is a binary splitting of a unique dyadic interval K. Further $K \subseteq L$.*

DEFINITION 6.10. We call a partition P of J a *binary partition* of J if all its elements are dyadic intervals. If P and Q are partitions of J, then we say that Q *refines* P if every interval in Q is contained in some interval in P. If (P,Q) is a partition pair as in Definition 3.3 with both P and Q binary partitions of J, then we call the pair a *binary partition pair*. As in Remark 3.4, a bijection σ from P to Q will eventually be inserted into the pair.

DEFINITION 6.11. If P is a partition of a closed interval J and K is in P, then a *binary splitting* of P at K is the partition $P' = (P \setminus \{K\}) \cup \{K0, K1\}$. We say that P' is obtained from P by a binary splitting. We say that Q is a *binary refinement* of P if Q is obtained from P by a sequence of binary splittings.

In I, notions come together.

LEMMA 6.12. *Let P be a binary partition of I and let Q be a refinement of P. Then Q is a binary partition of I if and only if Q is a binary refinement of P.*

PROOF. The claim is clear if Q is a binary refinement of P. The argument assuming that Q is a binary partition of I is inductive. If J is shortest so that J is in Q but not P, then Q is obtained as a binary splitting of some binary partition R of I at some K. From Lemma 6.9(6), K is contained in an interval of P, and R is a refinement of P. $\square$

COROLLARY 6.12.1. *A partition of I is a binary partition of I if and only if it is a binary refinement of the partition $\{I\}$ of I.*

COROLLARY 6.12.2. *If P and Q are binary partitions of I, then there is a common binary refinement T of P and Q so that every common binary refinement of P and Q is a binary refinement of T.*

PROOF. The desired T consists of all the intervals in $P \cup Q$ that do not properly contain any other interval in $P \cup Q$. $\square$

The following is immediate.

LEMMA 6.13. *If P is a binary parition of a dyadic interval J and $f \in F$ is affine on J, then f carries P to a binary partition of Jf.*

6.2.2. *Representations by binary partition pairs.* If (P, Q) is a binary partition pair for I, then the PL homeomorphism of I determined by (P, Q) as described in Definition 3.3 is clearly an element of F. The following shows that all elements of F can be obtained this way.

PROPOSITION 6.14. *Every element of F can be represented by a binary partition pair of I.*

PROOF. Let f be in F. There are finitely many breakpoints of f, all dyadic, and there is a $j \in \mathbf{N}$ so that every breakpoint of f occurs at an integral multiple of 2^{-j}. Let P_1 be the binary partition of I whose endpoints are all the integral multiples of 2^{-j}. Now let $k \in \mathbf{N}$ be such that the endpoints of all the Jf with $J \in P_1$ are integral multiples of 2^{-k}. The set of all integral multiples of 2^{-k} are the endpoints of a binary partition Q of I.

For $J \in P_1$, the length of Jf is an integral power of 2 and thus a union of m intervals in the partition Q where m is an integral power of

2. Thus this set of m intervals from Q forms a binary partition of Jf and its image under f^{-1} is a binary partition P_J of J, and all intervals in P_J are dyadic. The partition $P = \cup\{P_j \mid J \in P_1\}$ is now a binary partition of I whose image under f is Q, and (P, Q) represents f. $\square$

The binary partition pair representing an element of f is not unique To discuss this, we elevate splitting partitions to splitting partition pairs. We introduce some flexibility to be able to discuss more than F by the end of this chapter. Recall that a binary pair (P, Q) is generalized in Remark 3.4 to a triple (P, σ, Q) where σ is a bijection from P to Q. For elements of F, the bijection σ is order preserving.

Definition 6.15. If (P, σ, Q) is a binary partition pair of a closed interval J, and K is in P, then (P', σ', Q') is the result of a *matched binary splitting* of (P, σ, Q) at K if P' is obtained from P by a binary splitting at K, Q' is obtained from Q by a binary splitting at $K\sigma$, and $\sigma' = \sigma$ off K and takes $K0$ to $(K\sigma)0$, and $K1$ to $(K\sigma)1$. The notation is as in Definition 6.8. We say that (P', σ', Q') is a *matched binary refinement* of (P, σ, Q) if (P', σ', Q') is obtained from (P, σ, Q) by a sequence of matched binary splittings. The reverse of a matched binary splitting is a matched binary reduction. If (P, σ, Q) (or (P, Q) if σ is order preserving) is a binary partition pair, then $[P, \sigma, Q]$ (or $[P, Q]$) will denote the equivalence class of (P, σ, Q) (or (P, Q)) under the equivalence relation generated by matched binary splittings.

The following is clear.

Lemma 6.16. *For a binary partition pair (P, Q), all elements in $[P, Q]$ represent the same element of F.*

If (P, Q), (P', Q) and (P, Q') are binary partition pairs representing the same element of F, then $P = P'$ and $Q = Q'$.

Proposition 6.17. *Every f in F is represented by a unique binary partition pair (P, Q) that is irreducible with respect to matched binary reduction, and the elements of $[P, Q]$ are exactly the pairs that represent f. Further any binary partition pair representing f is a matched binary refinement of (P, Q).*

Proof. If (P, Q) and (R, S) are binary partition pairs that both represent $f \in F$, then there is a common binary refinement T of P and R. There are then corresponding matched binary partitions (T, Q') and (T, S'), respectively, of (P, Q) and (R, S). Now $Q' = S'$ by Lemma 6.16, and we get that (R, S) is in $[P, Q]$.

Since reductions lower the number of elements in a partion, it follows from Proposition 6.14 that at least one irreducible representative exists for each element of F. Assume that a partition pair (P, Q) allows for two reductions. One reduction must replace some $J0$ and $J1$ in P with J, and the other reduction must replace some $K0$ and $K1$ in P with K (and corresponding replacements in Q). Since $J0$ and $K0$ are left halves of intervals, then either $J = K$ or they are intervals with disjoint interiors. If $J \neq K$, then doing both reductions in either order gives the same result. This verifies the diamond condition of Section 47 and the claims follow from Lemma 47.2 and Corollary 47.2.1. $\square$

The irreducible representative of Proposition 6.17 is a normal (and not just seminormal) form. It is the model for all normal forms that will occur in this chapter. However, even the non-reduced representatives give a workable seminormal form since the only partition pairs that represent the identity are of the form (P, P).

6.2.3. *Multiplication and Thompson's group F.* Multiplication as in Section 3.2 is opportunistic. We have the special case $(P, Q)(Q, R) = (P, R)$. Given $(P, Q)(R, S)$, we let T be the common binary refinement of Q and R. Matched binary refinements produce (P', T) and (T, S') for (P, Q) and (R, S) to give (P', S') for the product. This can be reduced to an irreducible if desired.

The product formula $(P, Q)(Q, R) = (P, R)$ agrees with composition of the functions represented by the pairs. From Lemma 6.16 the calculation of the product $(P, Q)(R, S)$ also agrees with the composition of the functions represented. If we let $F_{\mathscr{P}}$ the set of all classes $[P, Q]$ with (P, Q) a binary partition pair, and give $F_{\mathscr{P}}$ the multiplication just defined, then we have the following.

Proposition 6.18. *The multiplication above is well defined and makes $F_{\mathscr{P}}$ a group that is isomorphic to Thompson's group F.*

Remark 6.19. One can think of F as the set of irreducible, binary partition pairs where multiplication consists of the procedure given above followed by a reduction of the result to an irreducible pair.

7. Combinatorics, II: Prefix sets

We use Propositions 6.14 and 6.17 along with Lemma 6.9(3) to turn working with elements of F into combinatorial exercises.

7.1. Generalizing Section 6. This topic is large and must be approached with caution. Generalizations offer the following.

(1) Agreement with the literature.
(2) Easier calculations.
(3) Easier access to different aspects of an object.

(4) Easier access to variations of an object.
(5) Make clear general and universal aspects of an object.
(6) Give unified proofs of some favorite properties of a family of objects.

We will be motived less by Item (6) and more by Items (1) through (5). While certain properties such as simplicity or finite presentability are proven many times over in the family of Thompson groups, this book will give only one or two examples, and will give references for more. There have been several successful constructs that unify different portions of the Thompson family and these will be discussed in the end notes to this chapter. We will only generalize slightly here, and will do so in small steps. Deeper generalization will occur in other chapters.

Section 6.2 starts with a set of intervals and builds a structure $F_{\mathscr{P}}$ that is established in Proposition 6.18 to be a group isomorphic to Thompson's group F. The fact that the set of intervals is parametrized by the free monoid $\mathfrak{M}$ on $\{0,1\}$ in a way that cooperates with much of the structure is the key to the first generalization that we build. We will use the parametrization to build a structure on $\mathfrak{M}$ that will be isomorphic to $F_{\mathscr{P}}$ and thus F.

This has two consequences. It gives a combinatorial tool for working with F, and it gives an easy way to show that other structures (here we focus on constructs using intervals) are isomorphic to F. Much of the work is imitation of material in Section 6.2 and will be left to the reader.

There are three levels of structure to model.

7.2. The interval level. The order preserving bijection from $\mathfrak{M}$ to the set of dyadic intervals in I has already been established by definition ($\emptyset \mapsto I$, $w0 \mapsto (Jw)0$ and $w1 \mapsto (Jw)1$) and Lemma 6.9(3).

7.3. The binary partition level: prefix sets. We imitate binary partitions.

DEFINITION 7.1. If P is a subset of $\mathfrak{M}$ and p is in P, then a *binary splitting* of P at p yields the subset $Q = (P \setminus \{p\}) \cup \{p0, p1\}$. A *prefix set* for $\mathfrak{M}$ is either $\{\emptyset\}$ or one derived from $\{\emptyset\}$ by a sequence of binary splittings.

The term "prefix set" is justified by the following.

LEMMA 7.2. *A subset P of $\mathfrak{M}$ is a prefix set if and only if for every word $w \in \{0,1\}^\omega$ there is a unique prefix for w in P.*

Proof. The direction assuming that P is a prefix set is an easy induction starting with the fact that every word in $\{0,1\}^\omega$ has the unique prefix $\emptyset$ in $\{\emptyset\}$.

Assume now that P has a unique prefix for every word in $\{0,1\}^\omega$. Recall that we view of the Cantor set $\mathfrak{C}$ as $\{0,1\}^\omega$, and note that P is finite since $\{p\mathfrak{C} \mid p \in P\}$ is a cover of $\mathfrak{C}$ by pairwise disjoint open sets. We can now imitate the proof of Lemma 6.12. Assume $P \neq \{\emptyset\}$ and let p be longest in P. If $p = p'0$, then an infinite word beginning with $p'1$ must have a unique prefix in P which must be $p1$ by our choice of p. So P is the binary splitting of $(P \setminus \{p'0, p'1\}) \cup \{p'\}$ which is easily shown to have a unique prefix for every word in $\{0,1\}^\omega$. If $p = p'1$, the argument is similar. The claim follows inductively. □

Corollary 7.2.1. *A subset P of $\mathfrak{M}$ is a prefix set if and only if $IP = \{Iw \mid w \in P\}$ is a binary partition of I, in which case $w \mapsto Iw$ is an order preserving bijection from P to IP. The map $P \mapsto IP$ gives a bijection that commutes with binary splittings from the set of prefix sets in $\mathfrak{M}$ to the set of binary partitions of I.*

Proof. The fact that if Q is a binary splitting of the prefix set P at p, then $IQ = \{Iw \mid w \in Q\}$ is a binary splitting of $IP = \{Iw \mid w \in P\}$ at Ip is straightfoward. The if and only if is an induction based on the parallel Corollary 6.12.1 and Definition 7.1. The other claims are straightforward. □

7.4. The binary partition pair level: prefix set pairs. We continue to mirror the structures in Section 6.2.

Definition 7.3. In the following P and Q are prefix sets.

(1) If every $q \in Q$ has a prefix in P, then Q *refines* P.
(2) If Q is obtained from P by a sequence of binary splittings, then Q is a *binary refinement* of P.
(3) If $\sigma : P \to Q$ is a bijection, then (P, σ, Q) is a *prefix set pair*. If σ is omitted, it is assumed to be the order preserving bijection.

Definition 7.4. In the following (P, σ, Q) and (P', σ', Q') are prefix set pairs.

(1) If p is in P, then (P', σ', Q') is the result of a *matched binary splitting* of (P, σ, Q) at p if P' is obtained from P by a binary splitting at p, Q' is obtained from Q by a binary splitting at $p\sigma$, and $\sigma' = \sigma$ off p and takes $p0$ to $(p\sigma)0$ and $p1$ to $(p\sigma)1$.
(2) We say that (P', σ', Q') is a *matched binary refinement* of (P, σ, Q) if (P', σ', Q') is obtained from (P, σ, Q) by a sequence of matched binary splittings,

(3) The reverse of a matched binary splitting is a matched binary reduction. A prefix set pair permitting no matched binary reduction is said to be *irreducible*.
(4) We use $[P, \sigma, Q]$ to denote the equivalence class of (P, σ, Q) under the equivalence relation generated by matched binary splittings.

The items in the next proposition are either proven by imitating the proof of the corresponding fact from Section 6.2, or by using the proven properties of the parametrizaton of dyadic intervals by $\mathfrak{M}$.

PROPOSITION 7.5. *For prefix sets P and Q, Q is a refinement of P if and only if Q is a binary refinement of P.*

Any two prefix sets have a common (binary) refinement.

Setting $(P, \sigma, Q) \mapsto (IP, \sigma', IQ)$ where σ' is induced by σ gives a bijection from the set of prefix set pairs to the set of binary partition pairs of I that commutes with matched binary splittings and that carries $[P, \sigma, Q]$ onto $[IP, \sigma', IQ]$.

Every $[P, \sigma, Q]$ has a unique irreducible element from which all other elements in $[P, \sigma, Q]$ are obtainable as matched binary refinements.

7.5. Multiplication. At this point, a multiplication on the classes $[P, Q]$ of prefix set pairs can be copied word for word and letter for letter from Section 6.2.3 to create a set with multiplication that we denote $F_{\mathfrak{M}}$. We have the following.

PROPOSITION 7.6. *The multiplication just discussed makes $F_{\mathfrak{M}}$ a group isomorphic to $F_{\mathcal{P}}$ and thus also isomorphic to F.*

The prefix set pairs that represent the identity are those pairs of the form (P, P). This follows from the fact that the only binary partition pairs that represent the identity are those pairs of the form (P, P).

Multiplication using prefix set pairs is practical, but often space inefficient. We will introduce tree pairs in Section 8 which tend to be more popular since multiplication with tree pairs is also practical and perhaps more space efficient.

We give an example. The generators x_0 and x_1 of (3.1) can be represented as prefix set pairs as follows.

$$x_0 = \begin{pmatrix} 00 & \to 0 \\ 01 & \to 10 \\ 1 & \to 11 \end{pmatrix} \qquad x_1 = \begin{pmatrix} 0 & \to 0 \\ 100 & \to 10 \\ 101 & \to 110 \\ 11 & \to 111 \end{pmatrix}$$

Listing the elements of a prefix set horizontally, and placing one prefix set over the other is another option. We calculate x_0^2 as follows.

$$
\begin{aligned}
x_0^2 &= \begin{pmatrix} 00 & \to 0 \\ 01 & \to 10 \\ 1 & \to 11 \end{pmatrix} \begin{pmatrix} 00 & \to 0 \\ 01 & \to 10 \\ 1 & \to 11 \end{pmatrix} = \begin{pmatrix} 000 & \to 00 \\ 001 & \to 01 \\ 01 & \to 10 \\ 1 & \to 11 \end{pmatrix} \begin{pmatrix} 00 & \to 0 \\ 01 & \to 10 \\ 1 & \to 11 \end{pmatrix} \\
&= \begin{pmatrix} 000 & \to 00 \\ 001 & \to 01 \\ 01 & \to 10 \\ 1 & \to 11 \end{pmatrix} \begin{pmatrix} 00 & \to 0 \\ 01 & \to 10 \\ 10 & \to 110 \\ 11 & \to 111 \end{pmatrix} = \begin{pmatrix} 000 & \to 0 \\ 001 & \to 10 \\ 01 & \to 110 \\ 1 & \to 111 \end{pmatrix}
\end{aligned} \tag{7.1}
$$

The first change is a matched binary splitting at the first line of the first factor. The second change is a matched binary splitting at the last line of the second factor. After the second change, opportunistic multiplication can take place.

Remark 7.7. In parallel to the point in Remark 6.19, one can think of F as the set of irreducible, prefix set pairs where multiplication of two pairs is followed by a reduction of the result to an irreducible pair.

7.6. Examples. When the properties proven about prefix set pairs and a system of intervals having the properties in Items (1) and (2) from Lemma 6.9 are combined we get other versions of F. We will give several examples. We also immediately get a faithful action of F on the Cantor set.

7.6.1. *F on the line.* Let $\mathcal{J}$ be the set of intervals in $\mathbf{R}$ consisting of $\mathbf{R}$, all $[k,\infty)$ and $(-\infty,-k]$ for $k\in\mathbf{N}$, and all integer translates of the dyadic intervals from Definition 6.8. We tell how to "binary split" these intervals. Rules for the unbounded intervals differ from the rules for the bounded intervals. For a bounded interval $J\in\mathcal{J}$ we let $J0$ be the left half of J and $J1$ be the right half of J exactly as in Definition 6.8. We set $\mathbf{R}0=(-\infty,0]$, $\mathbf{R}1=[0,\infty)$ and for all other unbounded $J\in\mathcal{J}$, and for $i\in\mathbf{N}$ we set

$$
\begin{aligned}
(-\infty,-k]0 &= (-\infty,-k-1], & [k,\infty)0 &= [k,k+1], \\
(-\infty,-k]1 &= [-k-1,-k], & [k,\infty)1 &= [k+1,\infty).
\end{aligned}
$$

We pick out certain select maps between the intervals in $\mathcal{J}$. For J and K in $\mathcal{J}$ both bounded, our select map is the unique increasing affine map from J onto K. For $J=(-\infty,-i]$ and $K=(-\infty,-k]$, our select map is the unique map of slope 1 from J onto K. Similarly for $J=[i,\infty)$ and $K=[k,\infty)$, our select map is also the unique map of slope 1 from J onto K. Our select map from $\mathbf{R}$ to $\mathbf{R}$ is the identity.

There are no other maps that need to be defined since any partition by elements of $\mathcal{J}$ of $\mathbf{R}$ with more than one element must have an interval of the form $(-\infty, j]$ as the leftmost interval and an interval of the form $[k, \infty)$ as the rightmost interval. The only partition of $\mathbf{R}$ with one element is $\{\mathbf{R}\}$. Thus a pair of partitions P and Q of $\mathbf{R}$ by elements of $\mathcal{J}$ with the same number of elements induces a homeomorphism from $\mathbf{R}$ to itself using the select maps that we have defined to take the elements of P to those of R so as to preserve the order.

We claim that we have reproduced the properties in Items (1) and (2) of Lemma 6.9 which we now repeat so as to fit our current situation. The following hold.

(1) For $J \in \mathcal{J}$, $J = J0 \cup J1$, the intersection $J0 \cap J1$ contains only the common endpoint of $J0$ and $J1$, and $J < J0 < J1$ under the order from Remark 3.4 as augmented by Lemma 6.9(1).
(2) For intervals J and K in $\mathcal{J}$, a select function takes J onto K, if and only if its restrictions $J0$ and $J1$ are select functions onto $K0$ and $K1$, respectively.

We leave it to the reader to define the bijection from $\mathfrak{M}$ to $\mathcal{J}$ that starts wtih $\emptyset \mapsto \mathbf{R}$, and to verify that properties corresponding to Items (3) and (4) from Lemma 6.9 hold. The rest of the behavior of binary partitions of I can be imitated either from what has just been estasblished or by using the known properties of prefix sets, prefix set pairs and the structure of $F_{\mathfrak{M}}$. We then get equivalence classes of partition pairs that use intervals from $\mathcal{J}$ to produce a multiplicative structure on those classes that we will denote $F_{\mathbf{R}}$. The reader can verify the truth of the following.

Proposition 7.8. *The mulitplicative structure of $F_{\mathbf{R}}$ is a group isomorphic to $F_{\mathfrak{M}}$ and thus also isomorphic to F.*

The generators corresponding to x_0 and x_1 of F are easy to describe. The generators x_0 and x_1 are represented by the prefix set pairs

$$\begin{aligned} x_0 =&(\{00, 01, 1\}, \{0, 10, 11\}) \\ x_1 =&(\{0, 10, 101, 11\}, \{0, 10, 110, 111\}) \end{aligned}$$

Using the intervals from $\mathcal{J}$, the first gives the homeomorphism s with $ts = t + 1$ on $\mathbf{R}$, and the second gives the homeomorphism y_0 on $\mathbf{R}$ defined by

$$ty_0 = \begin{cases} t, & t \le 0, \\ 2t, & 0 \le t \le 1, \\ t+1, & t \ge 1. \end{cases} \tag{7.2}$$

The notation y_0 will be explained and extended in Section 9.4.

The graphs of s and y_0 are shown below.

(7.3)

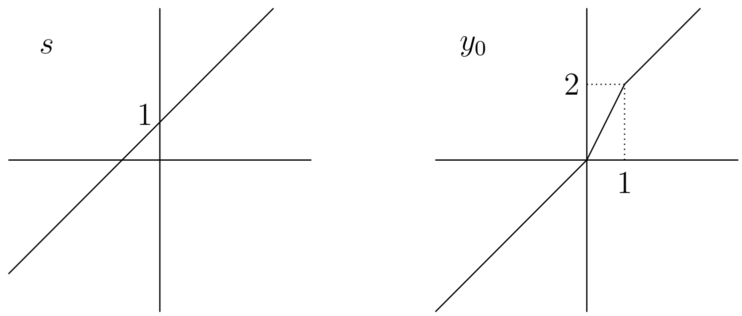

The reader can also verify the following.

Proposition 7.9. *The elements of $F_{\mathbf{R}}$ are those PL homeomorphisms f of $\mathbf{R}$ whose slopes are integral powers of 2, whose break points are dyadic rationals and for which there are integers N, P, n and p that depend on f so that the restriction of f to $(-\infty, N]$ agrees with $x \mapsto x+n$ and the restriction to $[P,\infty)$ agrees with $x \mapsto x+p$.*

Remark 7.10. The group $F_{\mathbf{R}}$ is conjugate to the action of F restricted to $(0,1)$. Let h map

$$\begin{aligned} \frac{1}{2^{n+1}} &\mapsto -n, \quad n \in \mathbf{N}, \\ 1-\frac{1}{2^{n+2}} &\mapsto n+1, \quad n \in \mathbf{N}. \end{aligned}$$

So $\ldots$, $\frac{1}{8}h = -2$, $\frac{1}{4}h = -1$, $\frac{1}{2}h = 0$, $\frac{3}{4}h = 1$, $\frac{7}{8}h = 2 \ldots$, etc.

We extend h to be piecewise linear between the values defined. The reader can show that h conjugates F to $F_{\mathbf{R}}$. It is entertaining to observe the cancellation of all but finitely many breakpoints of h during the conjugation of an $f \in F$ by h.

7.6.2. *F on a half line.* Let $F_{\geq 0}$ denote the subset of $F_{\mathbf{R}}$ consisting of those elements that are the identity on $(-\infty, 0]$. This set is clearly closed under product and inverse, and so is a subgroup of $F_{\mathbf{R}}$. The reader can show in either of two ways that $F_{\geq 0}$ is isomorphic to F. One way is to build a conjugator taking the action of F on $(0,1)$ to $F_{\geq 0}$. The other way is to note that the intervals in $\mathcal{J}$ from Section 7.6.1 that are contained in $[0,\infty)$ form a system of intervals parametrized by $\mathfrak{M}$ starting with $\emptyset \mapsto [0,\infty)$ that has all the properties of $\mathcal{J}$ needed to establish the isomorphism. If we conjugate $F_{\geq 0}$ by appropriate elements of $F_{\mathbf{R}}$, we get isomorphisms from F to every $F_{\geq a}$ with $a \in \mathbf{Z}[\frac{1}{2}]$ consisting of those elements in $F_{\mathbf{R}}$ that are the identity on $(-\infty, a]$.

7.6.3. *A piecewise projective version F.* We consider maps of closed intervals to closed intervals, but instead of maps that are affine on each interval, we consider maps that are projective on each interval. We start with a description of the intervals involved.

Let $A = [\frac{m}{n}, \frac{p}{q}]$ be an interval with rational endpoints. From this point all rationals will be represented as ratios of integers in lowest terms with positive denominator. As an extension of this, we allow the endpoints $\infty = \frac{1}{0}$ and $-\infty = \frac{-1}{0}$, and we think of $\mathbf{R}$ as $[-\infty, \infty]$.

We use the matrix $\begin{pmatrix} m & p \\ n & q \end{pmatrix}$ which we also call A to represent the interval $A = [\frac{m}{n}, \frac{p}{q}]$. We will not distinguish between the matrix and the interval. It is easy to show that, except for $A = \mathbf{R}$, we have $\det(A) < 0$ if and only if $\frac{m}{n} < \frac{p}{q}$. The set $\mathcal{P}$ of intervals that we will work with will consist of $\mathbf{R}$ and those intervals A with rational endpoints with $\det(A) = -1$. Note that for each integer i, the interval $[i, i+1]$ is in $\mathcal{P}$.

To describe the parametrization of $\mathcal{P}$ by $\mathfrak{M}$, we note that for $A = \begin{pmatrix} m & p \\ n & q \end{pmatrix}$ in $\mathcal{P} \setminus \{\mathbf{R}\}$, both $A0 = \begin{pmatrix} m & m+p \\ n & n+q \end{pmatrix}$ and $A1 = \begin{pmatrix} m+p & p \\ n+q & q \end{pmatrix}$ are in $\mathcal{P}$. In particular, as intervals, $A0$ can be viewed as a "left half" of A and $A1$ can be viewed as a "right half" of A. We set $\mathbf{R}0 = [-\infty, 0]$ and $\mathbf{R}1 = [0, \infty]$, as expected.

Given A and B in $\mathcal{P} \setminus \{\mathbf{R}\}$, we can form the matrix BA^{-1} which, acting on the left, carries A to B. This will be an integer matrix with determinant 1 and thus an element of $SL(2, \mathbf{Z})$. We will think of BA^{-1} as an element of $PSL(2, \mathbf{Z})$ since $-I$ takes $A = [\frac{m}{n}, \frac{p}{q}]$ to itself. Now, Given $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in PSL(2, \mathbf{Z})$, and $A = \begin{pmatrix} m & p \\ n & q \end{pmatrix} \in \mathcal{P} \setminus \{\mathbf{R}\}$ we have $MA = \begin{pmatrix} ma+nb & pa+qb \\ mc+nd & pc+qd \end{pmatrix}$. This agrees with the projective action of M on $\mathbf{R}$ defined by

$$t \mapsto \frac{at+b}{ct+d} \tag{7.4}$$

because if t is the rational $\frac{m}{n}$, then the effect of M on t is

$$\frac{m}{n} \mapsto \frac{am+bn}{cm+dn}.$$

So our select map between intervals in A and B in $\mathcal{P} \setminus \{\mathbf{R}\}$ is the action by the matrix $BA^{-1} \in PSL(2, \mathbf{Z})$ acting on the left. The only special case is that $\mathbf{R}$ only maps to $\mathbf{R}$ and our select map there is the identity. The following is now a set of exercises for the reader. The first item is the bulk of the work.

LEMMA 7.11. *The following hold*

(1) *Starting with* $\mathbf{R}\emptyset = \mathbf{R}$*, every interval in* $\mathcal{P}$ *is of the form* $\mathbf{R}w$ *for some* $w \in \mathfrak{M}$*.*
(2) *The "action" of* $PSL(2, \mathbf{Z})$ *is truly an action in that* $M_1(M_2t) = (M_1M_2)t$*.*
(3) *Properties corresponding to (1) and (2) of Section 7.6.1 hold. Specifically for (2) and* A *and* B *in* $\mathcal{P} \setminus \{\mathbf{R}\}$ *we have that the restriction of* BA^{-1} *to* $A0$ *is* $(B0)(A0)^{-1}$ *and to* $A1$ *is* $(B1)(A1)^{-1}$*.*

If the process in Section 7.6.1 is followed, if equivalence classes of pairs of partitions of $\mathbf{R}$ using elements of $\mathcal{P}$ are defined, and projective maps are used to to take the intervals in the first partition of the pair to the intervals of the second, then we have a our set of piecewise projecive elements of Homeo($\mathbf{R}$). Opportunistic multiplication gives a structure $F_{\mathcal{P}}$ for which the following holds. It will be necessary to note that every group is isomorphic to its opposite.

PROPOSITION 7.12. *The multiplicative structure* $F_{\mathcal{P}}$ *is a group isomorphic to* $F_{\mathfrak{M}}$ *and thus also isomorphic to* F*.*

As bits of extra interest, the reader can show that the element corresponding to x_0 is $t \mapsto t+1$, that the element corresponding to x_1 agrees with the element y_0 of (7.2) and (7.3) on the complement of $[0,1]$ but not on $[0,1]$, and that, unlike F, every element of $F_{\mathcal{P}}$ has continuous first derivative (but not necessarily continuous second derivative).

7.6.4. *Action on the Cantor set.* Prefix set pairs give a very direct way to describe a natural action of F on the Cantor set $\mathfrak{C}$. Recall that we identify the Cantor set with $\{0,1\}^\omega$, and that we use $u\mathfrak{C} = \{u\alpha \mid \alpha \in \mathfrak{C}\}$ to denote the clopen set in $\mathfrak{C}$ of all words that start with u. Given a prefix set pair (P, σ, W) representing $f \in F$, we set $(u\alpha)f = (u\sigma)\alpha$ for all $u \in P$. With σ the order preseving bijection from P to Q, this gives an action of F on $\mathfrak{C}$.

We can interpret each word $\alpha \in \mathfrak{C}$ as a real number in $I = [0,1]$ by thinking of it as a binary expansion of a real number with "binary point" immediately to the left of the first symbol in α. This creates a map $\pi : \mathfrak{C} \to I$ that is two-to-one to the dyadics in I and one-to-one everywhere else. If θ is the homomorphism from F to the homeomorphisms of the Cantor set given above, then for $f \in F$ we have $\pi f = \theta(f)\pi$, composing left to right.

8. Combinatorics III: Trees

Binary trees mirror prefix sets over $\{0,1\}$. The shift from prefix sets to binary trees is mostly an exercise in changing the vocabulary, but trees give a convenient way to represent elements of F and to calculate with them. They are one of the most popular ways to work with Thompson's groups. We adopt a restricted definition of a tree.

8.1. The definitions.

Definition 8.1. Given a finite, totally ordered alphabet A with $|A| = n$, the *complete, rooted n-ary tree* $\mathcal{T}_A$ has as vertices (nodes) the elements of the free monoid $\mathfrak{M}_A = A^*$ over A. Each $v \in T$ has n *children*, consisting of all the va with $a \in A$. The *parent* of each va is v. The *root* of $\mathcal{T}_A$ is $\emptyset$ and is the only node of $\mathcal{T}_A$ with no parent. The transitive closure of "child of" is *descendant*, and the transitive closure of "parent of" is *ancestor*. The root is the ancestor of all nodes in $\mathcal{T}_A$ other than itself. We order the nodes of $\mathcal{T}_A$ using the order of Definition 6.2. This is the usual prefix order given to trees where each node has a fixed total order on its children.

Definition 8.2. All n-ary trees other than the complete n-ary tree will subtrees of $\mathcal{T}_A$. An n-ary tree is finite if it has finitely many nodes. For a subset T of $\mathcal{T}_A$ to be a subtree of $\mathcal{T}_A$, all of the following must be satisfied.

(1) There must be a root r of T that is either specified or understood to be $\emptyset$, the root of $\mathcal{T}_A$.
(2) A node v of a T will either have no children or will have as children all the children that v has in $\mathcal{T}_A$.
(3) Every node w other than the root r will have as parent the parent that w has in $\mathcal{T}_A$ and r will be the ancestor of every node of T other than r.

The nodes of T with no children will be the *leaves* of T. We will often use $\Lambda(T)$ to denote the set of leaves of the tree T. All nodes of T that are not leaves of T are *internal nodes* of T. The restricton of the prefix order on $\mathcal{T}_A$ to T is used as the order on the nodes of T.

The *trivial tree* has only one node, which is then both its root and a leaf.

We will constantly exploit the fact the nodes of a tree are members of a monoid.

Definition 8.3. If T is an n-ary tree rooted at $r \in \mathcal{T}_A$ with $v \in T$, then T_v, the subtree of T rooted at v, is the set of nodes $\{vu \mid vu \in T\}$.

The tree T/r is the set $\{u \in \mathcal{M}_A \mid ru \in T\}$ and is the tree T *rerooted* to $\emptyset$. Note that T_v/v is the subtree of T rooted at v rerooted to $\emptyset$. If T is a tree rooted at r, if v is a leaf of T, and U is a tree rooted at $\emptyset$, then the result of *attaching* U to T at (the leaf) v is the tree $T \cup vU = T \cup \{vu \mid u \in U\}$.

LEMMA 8.4. *If S and T are finite n-ary trees with $S \subseteq T$, and $\Lambda(S)$ is the set of leaves of S, then*

$$T = S \cup \left(\bigcup_{v \in \lambda(S)} T_v \right) = S \cup \left(\bigcup_{v \in \lambda(S)} v(T_v/v) \right).$$

There are several places in the definitions above where it is implied that a certain collection of nodes forms a tree. The reader can verify that these implications are correct. The following can be useful.

LEMMA 8.5. *A subset of $\mathcal{T}_A$ is a tree rooted at $\emptyset$ if it is closed under parent and sibling, where the latter means that if va is in T for an $a \in A$, then vb is in T for all $b \in A$.*

At this point it is clear that ancestor is the same as prefix or left factor and descendant is the same as right multiple.

8.2. Connections to prefix sets and F. We now replace the general alphabet A by $\{0, 1\}$ and drop the subscript A. So all trees are binary and $\mathcal{T}$ is the complete, rooted, binary tree. To continue the connection between prefix sets and trees, we need more definitions.

DEFINITION 8.6. A *caret* $\wedge$ is a triple $\{v, v0, v1\}$ in $\mathcal{T}$. A *binary splitting* at v of a finite binary tree T with leaf v is the tree $T \cup \{v0, v1\}$. This is also called *attaching a caret* to T at v. A *binary refinement* is obtained by a sequence of binary splittings.

A *tree pair* is a triple (T, σ, S) where T and S are finite, binary trees, and σ is a bijection from the leaves of T to the leaves of S. If σ is omitted, it is understood to be the bijection that preserves the prefix order.

REMARK 8.7. For binary trees, an *infix order* on the nodes is also available. This is defined recusively by saying that for each node u, $u0$ and its descendants come before u in the order, and $u1$ and its descendants come after u in the order. For a finite tree T, the orders on its leaves inherited from the prefix order and the infix order agree. So there is never a question as to what is meant by saying that σ in (T, σ, S) is order preserving. The infix order is sometimes useful when dealing recursively with all nodes in a binary tree.

The key connection to prefix sets is the following which is left to the reader.

LEMMA 8.8. *If T is a finite binary tree rooted at $\emptyset$, then $\Lambda(T)$, the set of leaves of T is a prefix set for $\mathfrak{M}$. Conversely, if P is a prefix set for $\mathfrak{M}$ then there is a unique finite binary tree T_P rooted at $\emptyset$ with $\Lambda(T_P) = P$. Further, a prefix set Q is a binary splitting of P at p if and only if T_Q is a binary splitting of T_P at P.*

At this point, the reader can follow the outline of Section 7 and define matched binary splittings, irreducible tree pairs with respect to the reverse of matched binary splittings, matched binary refinements, the equivlance classes on binary tree pairs, and the opportunistic multiplication on the set of such pairs. The resulting multipicative structure can be denoted F_T, and the reader can show the following.

PROPOSITION 8.9. *Opportunistic multiplication makes F_T a group isomorphic to $F_{\mathfrak{M}}$ and thus also isomorphic to F.*

Each element of F_T has a unique representative that is irreducible and from which all other representatives of that element can be obtained by sequences of matched binary splittings.

The representatives of the identity are those tree pairs of the form (T, T).

8.3. Working with trees.

8.3.1. *More infrastructure.*

DEFINITION 8.10. If T and S are binary trees with $T \subseteq S$, then S is a *refinement* of T. The reader can show that for binary trees, refinement and binary refinement (Definition 8.6) are synonymous.

The leaves of a finite binary tree inherit a total order from the prefix order on the nodes of the tree. We number the leaves of a finite binary tree in order starting at 0.

Since binary trees are subtrees of the complete, binary tree $\mathcal{T}$, we can take unions and intersections of binary trees. We can also take set differences, but when we do so we regard the trees as unions of carets rather than collections of vertices.

A *binary forest* is a sequence of binary trees. If the sequence is finite and the trees in the sequence are all finite, then the forest is a finite forest.

An *exposed caret* in a binary tree T is a caret in T so that both leaves of the caret are also leaves of T.

The next lemma is a collection of trivial but useful observations.

LEMMA 8.11. *The following hold.*

(1) *Both the union and the intersection of a family of binary trees are binary trees, and the union of a family of binary trees is the smallest common refinement of the family.*
(2) *Given a finite set V of vertices in $\mathcal{T}$, there is a smallest finite tree T rooted at $\emptyset$ containing V. If the vertices in V are pairwise orthogonal, then V will be contained among the leaves of T.*
(3) *If S and T are finite, binary trees and S refines T, then $S \setminus T$ is a finite binary forest, so that if v_i is the i-th leaf of T, then S_{v_i} is the i-th tree of the forest.*
(4) *Every finite binary tree has at least one exposed caret.*
(5) *If T' is the result of attaching a caret to a finite binary tree T at the i-th leaf, then leaves i and $i+1$ of T' are the leaves of an exposed caret of T'.*
(6) *A binary tree pair (T, S) is the result of a matched binary splitting (and is thus reducible) if and only if for some i, the leaves i and $i+1$ are leaves of a single exposed caret both in T and in S.*

We can use this to say more about multiplication. To multiply tree pairs $(R, S)(T, U)$ when $S \neq T$, we can let $V = S \cup T$, produce R' by attaching, for each relevant i, the i-th tree of $V \setminus S$ to the i-th leaf of R, produce U' by attaching, for each relevant j, the j-th tree of $V \setminus T$ to the j-th leaf of U, to obtain the result $(R, S)(T, V) = (R', U')$.

We give a sample calculation. In (7.1), a prefix set pair was calcuated for x_0^2 which we can turn into a binary tree pair.

$$x_0^2 = \begin{pmatrix} 000 & \to 0 \\ 001 & \to 10 \\ 01 & \to 110 \\ 1 & \to 111 \end{pmatrix} =$$

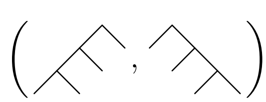

We then use this to calculate x_0^3 as follows.

$$x_0^3 = x_0^2 x_0 =$$

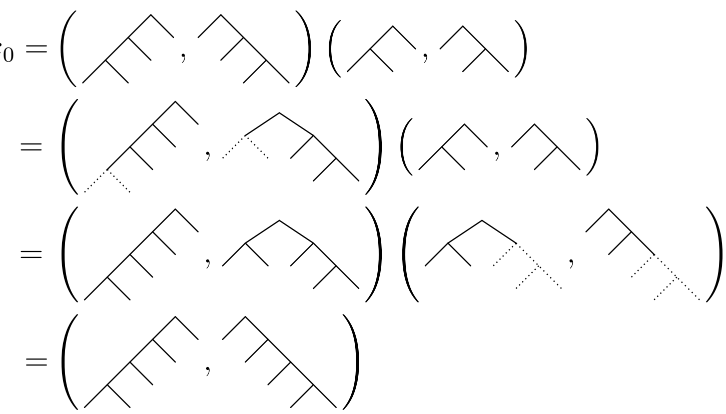

The first change attaches a single caret (drawn with dotted lines) to the leftmost leaf of both trees in the first binary tree pair. The second change attaches the subtree of two carets  to the rightmost leaf of both trees in the second binary tree pair. Finally the "middle trees" match, and opportunistic multiplication can take place.

Remark 8.12. Finally, in parallel to Remarks 6.19 and 7.7, one can think of F as the set of irreducible, tree pairs where multiplication is followed by a reduction of the result to an irreducible pair.

It is trivial to build examples with (R, S) and (T, U) both irreducible but for which the process above produces a reducible result. The inverse of (T, S) is (S, T) so if (T, S) is not the identity, then $(T, T) = (T, S)(S, T)$ represents the identity and is not irreducible. One can build less trivial examples.

8.4. Actions on a tree. There are several ways to think of an element of f as acting on the complete binary tree $\mathcal{T}$. The element f can be given as a tree pair (T, S) or the corresponding prefix set pair $(\Lambda(T), \Lambda(S)$ consisting of the leaf sets of T and S. As in Section 7.6.4, and with $\sigma : \Lambda(T) \to \Lambda(S)$, we have f acting on the Cantor set $\mathfrak{C}$ by $(u\alpha)f = (u\sigma)\alpha$ for $(u, \alpha) \in \Lambda(T) \times \mathfrak{C}$. Here we can think of $\mathfrak{C}$ as the ends of $\mathcal{T}$. But if we restrict α to $\mathfrak{M} = \{0, 1\}^*$, then we have an action of f on all but finitely many nodes of the complete binary tree $\mathcal{T}$.

We can extend the action of f to all the nodes of $\mathcal{T}$ by using the bijection of Lemma 6.9(5) from $\{0, 1\}^*$ to the dyadics in $(0, 1)$. If we use the bijection to regard an element of $\{0, 1\}*$ as both a dyadic and a node of $\mathcal{T}$, it is easy to see that the action of F given above on all but finitely many nodes of $\mathcal{T}$ agrees with the action of f on $\mathbf{Z}[\frac{1}{2}] \cap (0, 1)$. Thus action of f can be extended to all of $\mathcal{T}$. Note that this action fails to preserve the parent-child relation in finitely many places.

8.5. Parenthesized expressions. We briefly mention a way to view F as acting on fully parenthesized expressions. Actions on such expressions will be covered more fully in Section 16.

If A is a set of variable symbols, then for $a \in A$, the expression a is a fully parenthesized expression, and if e_1 and e_2 are fully parenthesized expressions, so is $(e_1 e_2)$. A correspondence between finite binary trees and fully parenthesized expressions is hinted at below.

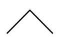 $\longleftrightarrow (ab)$ 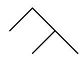 $\longleftrightarrow (a(bc))$ 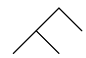 $\longleftrightarrow ((ab)c)$

Clearly pairs, splittings (plain and matched) and so forth can be defined. Changes of the variable names should make no difference, and so extra care is needed in defining F as operations on classes of pairs of expressions. However, it takes very little space to calculate with expressions. See Section 16 for many examples of this.

8.6. Vines, deferments and wedges. The following useful notions are easy to define with trees.

Definition 8.13. A tree with exactly one exposed caret is a *vine*. Given $v \in \{0,1\}^*$, we use V_v to denote the unique minimal tree containing v. It is immediate that V_v is a vine and that v is one of the leaves of its only exposed caret. We refer to V_v as the *vine determined by* v. A *left vine* is a vine determined by 0^n for some n and a *right vine* is a vine determined by 1^n for some n. We will use V_n to denote the right vine of n carets.

There are many ways to approach the next definition. Given $f \in F$ and a dyadic interval J, we want a new function that "acts on J as f acts on $I = [0,1]$." We can extend f to all of $\mathbf{R}$ by declaring that it acts as the identity off I, and then conjugate by an element in $F_{\mathbf{R}}$ that takes I affinely to J. The following alternative definition gives useful notation.

Definition 8.14. Let $f \in F$ be represented by the tree pair (T,S) and let u be in $\mathfrak{M}$ (equivalently a node of $\mathcal{T}$). Then f_u, the *deferment of* f *to* u, is represented by $(V_u \cup uT, V_u \cup uS)$.

Note that any tree with u as a leaf can be used instead of V_u.

Lemma 8.15. *The group F is closed under deferments.*

Definition 8.16. If S and T are binary trees, we define

$$S \wedge T = \emptyset \cup 0S \cup 1T.$$

This is the result of attaching S and T in left-right order to the two leaves of a caret.

For f and g in F, we set $f \wedge g = f_0 g_1 = g_1 f_0$ where f_0 and g_1 are deferments.

REMARK 8.17. For a binary tree T, we have

$$T = (T_0/0) \wedge (T_1/1)$$

as a special case of Lemma 8.4. For f and g in F represented by tree pairs (P, Q) and (R, S), respectively, we have $f \wedge g$ represented by $(P, Q) \wedge (R, S) = (P \wedge R, Q \wedge S)$. Note that $\wedge : F \times F \to F$ is an injective homomorphism whose image is the maximal subgroup of F consisting of those elements that fix $\frac{1}{2}$.

9. The group F, II: Presentations

We can now derive presentations for F. The only obstruction to doing the derivation earlier has been the problem of establishing a generating set. We do so by using Proposition 6.14.

9.1. Generators.

PROPOSITION 9.1. *The group F is generated by $\{x_0, x_1\}$.*

PROOF. Let $f \in F$ be represented by a binary partition pair (P, Q) of I where P and Q have n elements each. We will induct on n. If $n = 1$ or $n = 2$, there is only one possible partition and f is the identity. So we assume $n \geq 3$. We will show that after multiplying f on the right and left by well chosen combinations of x_0, x_1 and their inverses, the resulting function can be represented by a binary partition pair having $n - 1$ elements each.

If $\frac{1}{4}$ is an endpoint of an interval in P (i.e., the interval $[0, \frac{1}{2}]$ was divided in the creation of P from the trivial partition $\{I\}$), then $x_0^{-1} f$ is represented by a pair of binary partitions with n elements each and the number of endpoints of the new first partition in $[0, \frac{1}{2}]$ is fewer than in P. This is illustrated below using rectangle diagrams.

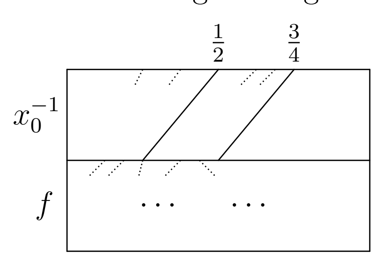

Thus after multiplying f on the left by a sufficient power (which might be zero) of x_0^{-1} we arrive at an element (which we will continue to call f) represented by a pair of partitions (which we will continue to denote (P, Q)) with n elements each, and where $[0, \frac{1}{2}]$ is the leftmost element of P.

Similarly, after multiplying f on the right by a sufficient power (which might be zero) of x_0 we arrive at an element (which we will continue to call f) represented by a pair of partitions (which we will continue to denote (P, Q)) with n elements each, and where $[0, \frac{1}{2}]$ is the leftmost element of Q.

We then repeat the two paragraphs above but now multiplying f by a power (possibly zero) of x_1^{-1} on the left and a power (also possibly zero) of x_1 on the right so that the resulting function (still called f) is represented by a pair of partitions (still denoted (P, Q)) with n elements each, where now $[0, \frac{1}{2}]$ and $[\frac{1}{2}, \frac{3}{4}]$ are the two leftmost elements of P and of Q. That is, $\frac{5}{8}$ is no longer an endpoint of either P or Q.

Lastly, we conjugate the latest version of f by x_0^{-1} which results in our final version of f and our final version of (P, Q) where now the two leftmost elements of both P and Q are the intervals $[0, \frac{1}{4}]$ and $[\frac{1}{4}, \frac{1}{2}]$. But this last (P, Q) is a matched binary splitting at the leftmost interval of a pair of binary partitions with $n-1$ elements each, and with leftmost interval $[0, \frac{1}{2}]$ in each. $\square$

9.2. An infinite presentation. We have claimed that the finite presentation (3.2) (copied just below) presents F.

$$\langle x_0, x_1 \mid [x_0^2 x_1^{-1} x_0^{-1}, x_1] = [x_0^3 x_1^{-1} x_0^{-2}, x_1] = 1 \rangle. \tag{3.2}$$

We also claim that F can be presented by the infinite presentation

$$\langle x_0, x_1, x_2 \ldots \mid x_j x_i = x_i x_{j+1} \text{ whenever } i < j \rangle. \tag{9.1}$$

We show that (3.2) and (9.1) present isomorphic groups. Thus finding a valid presentation for F can be done by showing that either (3.2) or (9.1) presents F.

PROPOSITION 9.2. *Sending x_0 and x_1 in the finite presentation (3.2) to the generators of the same name in the infinite presentation (9.1) extends to an isomorphism of the presented groups.*

PROOF. The relations in (9.1) can be rewritten as $x_j^{x_i} = x_{j+1}$ whenever $i < j$, or equivalently

$$x_i^{-1} x_{i+j} x_i = x_{i+j+1}, \text{ for all } j \geq 1. \tag{9.2}$$

For all $i \geq 2$ we have

$$x_i = x_1^{x_0^{i-1}} \tag{9.3}$$

and we can turn the relations in (9.3) into definitions of the x_i for $i \geq 2$. If we add the x_i with $i \geq 2$ with these definitions to the finite presentation (3.2), then the two relations in (3.2) can be given different appearances. The relation

$$[x_0^2 x_1^{-1} x_0^{-1}, x_1] = 1$$

can become

$$x_1 = x_1^{(x_0^2 x_1^{-1} x_0^{-1})} \quad \text{or} \quad x_1^{x_0 x_1} = x_1^{x_0 x_0} \quad \text{or} \quad x_2^{x_1} = x_3. \tag{9.4}$$

Similarly $[x_0^3 x_1^{-1} x_0^{-2}, x_1] = 1$ becomes

$$x_3^{x_1} = x_4. \tag{9.5}$$

These replacements are relations in (9.1) so we have an extension to a homomorphism.

We return to the relations in (9.2). The relations (9.4) and (9.5) can be conjugated by powers of x_0 to give all the relations in (9.2) for $j = 1$ and $j = 2$. For $k \geq 3$, if we assume (9.2) for all $j < k$, then we have

$$x_i^{-1} x_{i+k} x_i = x_i^{-1} x_{i+k-2}^{-1} x_{i+k-1} x_{i+k-2} x_i = x_{i+k-1}^{-1} x_{i+k} x_{i+k-1} = x_{i+k+1}.$$

Inductively, we have that all in (9.2) hold in the presentation (3.2) with (9.3) as definitions. Thus the homomorphism is an isomorphism. □

9.3. Final arguments. Let $\widehat{F}$ be the group presented by (9.1) which we repeat here for convenience.

$$\widehat{F} = \langle x_0, x_1, x_2 \ldots \mid x_j x_i = x_i x_{j+1} \text{ whenever } i < j \rangle \tag{9.1}$$

LEMMA 9.3. *There is a homomorphism from $\widehat{F}$ onto F taking x_0 and x_1 to the elements of F with the same name.*

PROOF. We look at some behavior of the x_i using the graphs in (3.1) for $i = 0, 1$ and the definitions in (9.3) for $i \geq 2$. For $i = 0, 1$, we have that the support of x_i is $(1 - 2^{-i}, 1)$ and x_i has slope $\frac{1}{2}$ on $(1 - 2^{-i-1}, 1)$. From (9.3) we see that this holds for all $i \geq 0$. For $i < j$, we have the support of x_j contained in an interval where the behavior of x_0 and x_i agree. So $x_j^{x_i} = x_j^{x_0} = x_{j+1}$. The image of the homomorphism is all of F since x_0 and x_1 generate F. □

LEMMA 9.4. *The function σ on $\{x_i \mid i \in \mathbf{N}\}$ defined by $\sigma(x_i) = x_{i+1}$ extends to an endomorphism of $\widehat{F}$.*

PROOF. The relation set in (9.1) is preserved by σ. □

We will refer to the endomorphism σ of Lemma 9.4 as the shift endomorphism. At this point it is an endomorphism of $\widehat{F}$, but once it is shown that the homomorphism of Lemma 9.3 is an isomorphism, we can also think of σ as an endomorphism of F.

We will give two proofs that Lemma 9.3 is about an isomorphism. One proof refers to the action of F on I, while the other only considers the structure of $\widehat{F}$. Both make use of a seminormal form for elements expressed in the x_i.

Lemma 9.5. *Every element of $\widehat{F}$ can be expressed in the following form*

$$x_{i_0}x_{i_1}\cdots x_{i_k}x_{j_l}^{-1}x_{j_{l-1}}^{-1}\cdots x_{j_0}^{-1}, \quad i_0 \le \cdots \le i_k \ne j_l \ge \cdots \ge j_0. \tag{9.6}$$

Note that the expression in (9.6) is not unique since $x_1 = x_0x_2x_0^{-1}$ and both sides are in the form of (9.6).

Proof. Each relation in (9.1) can also be written in the following three forms. In each we have $i < j$.

$$x_j^{-1}x_i = x_ix_{j+1}^{-1}, \qquad x_i^{-1}x_j = x_{j+1}x_i^{-1}, \qquad x_i^{-1}x_j^{-1} = x_{j+1}^{-1}x_i^{-1}.$$

The original together with the first above say that positive powers of x_i with low subscripts can move to the left of arbitrary powers of x_j with high subscripts at the expense of raising the higher subscript. The last two above say negative powers of x_i with low i can move to the right of powers of x_j with similar comments.

If w is a word of length n in the x_i and their inverses using powers in $\{\pm 1\}$, then let m be the lowest subscript used in w and let p be the number of appearances of x_m or its inverse in w. We can push all appearances of x_m^{+1} to the left and all appearances of x_m^{-1} to the right without increasing the length of the word. The length may go down if a generator at some point ends up next to its inverse. The result is a word in the form $x_m^q w_1 x_m^{-r}$ for some $q + r \le p$ with w_1 a word of length no greater than $n - p$ in the x_i and their inverses using powers in $\{\pm 1\}$ and with all $i > m$. We can continue the process with w_1 and eventually achieve the form (9.6). ☐

As discussed in Section 3.3 some seminormal forms are useful. The seminormal form (9.6) has its power, and we will use it in both proofs that (9.1) is a valid presentation for F. The first proof establishes that the form (9.6) can pick out the identity element.

Proposition 9.6. *The only word in the x_i in the form (9.6) that represents the identity in F is the empty word. The surjection of Lemma 9.3 from $\widehat{F}$ to F is an injection and (9.1) is a presentation for F.*

Proof. The second sentence will follow from the first since a word in the form (9.6) represents the identity of F if and only if it is in the kernel of the surjection from $\widehat{F}$ to F.

Let w be a non-empty word in the form (9.6) that represents the identity in F. We know that w is the identity in $\widehat{F}$ if any conjugate or an inverse of a conjugate of w is the identity. So we may assume that w is of the form $x_m^p w_1$ where $p > 0$ and w_1 is a word in x_i with all $i > m$. This says that $x_m^p = w_1^{-1}$. We argue that $p > 0$ leads to a contradiction.

Let n be the lowest subscript of an x_i that appears in w_1. From the discussion in the proof of Lemma 9.3 we know that the support of x_m is $(1 - 2^{-m}, 1)$ and the support of w_1 is contained in $(1 - 2^{-n}, 1)$. But we can also deduce from the definitions $x_i = x_1^{x_0^{i-1}}$ for $i > 1$ that the slope of x_i at $1 - 2^{-i}$ is 2 and so the slope of x_m^p at $1 - 2^{-m}$ is 2^p. With $m < n$ the equality $x_m^p = w_1^{-1}$ with $p > 0$ is not possible, even if w_1 is empty. □

Our second proof lifts the first result of Proposition 5.12 for F to $\widehat{F}$. This proposition will be useful in the future because the first conclusion of its statement has already absorbed the work of the seminormal form argument that is used in the proof of Proposition 9.6.

PROPOSITION 9.7. *Every proper quotient of $\widehat{F}$ is abelian and (9.1) is a presentation for F.*

PROOF. Again we start with a word w in the x_i in the form (9.6), but this time we assume that it is not the identity in $\widehat{F}$. Let N be the normal closure of w in $\widehat{F}$. We want to show that x_0 and x_1 commute modulo N. We will use $u \equiv v$ to say that $u = v$ holds in $\widehat{F}$ modulo N.

Again we can assume that w is in the form $x_m^p w_1$ where $p > 0$ and w_1 is a word in the x_i with all $i > m$. We will use the shift endomorphism σ from Lemma 9.4 of $\widehat{F}$ that takes each x_i to x_{i+1}. We know that $x_m^p \equiv w_1^{-1}$ and conjugating both sides by x_m^k with $k > 0$ gives $x_m^p \equiv \sigma^k(w_1^{-1})$. For any $i > m$ and $k > i - m$, we have that $x_i^{-1}\sigma^k(w_1^{-1})x_i = \sigma^{k+1}(w_1^{-1})$. But the $\sigma^k(w_1^{-1})$ are all equivalent to x_m^p. Thus $[x_m^p, x_i] \equiv 1$ for all $i > m$ and in fact for all $i \geq m$.

The previous argument will be reapplied so we take note that we have proven that if some x_m^p is equivalent to a word in generators of subscript greater than m, then x_m^p commutes modulo N with all x_i with $i \geq m$.

If $j > m$, then conjugating x_j by x_m^p yields x_{j+p} in $\widehat{F}$. But with $[x_m^p, x_j] \equiv 1$, we have $x_j \equiv x_{j+p}$ and our previous argument shows that $[x_j, x_i] \equiv 1$ for all $i \geq j$.

Now if $k > j > m$, we have $x_j^{-1}x_k x_j = x_{k+1}$ in $\widehat{F}$, but with $[x_j, x_k] \equiv 1$, we have $x_k \equiv x_{k+1}$. From the definitions we have $x_1^{x_0^{k-1}} \equiv x_1^{x_0^k}$. But conjugating this by x_0^{1-k} gives $x_1 \equiv x_1^{x_0}$ and we have proven the first claim.

For the second we note that the homomorphism of Lemma 9.3 takes x_0 and x_1 of $\widehat{F}$ to the elements of the same name in F and these generate F and do not commute. □

9.4. PL and smooth actions of F on the line. Lemma 9.5 and Proposition 9.7 give a representation free way to recognize F. We make the following a lemma because we will refer to it frequently.

LEMMA 9.8. *For a non-abelian group generated by a sequence x_i of elements that satisfy the relations of (9.1), taking each x_i to the element of the same name in F extends to an isomorphism.*

We give two applications. The first gives another (and easier?) proof of Proposition 7.8 and a slight extension. The second gives smooth actions of F.

Consider the following family of homeomorphisms from $\mathbf{R}$ to itself indexed by $i \in \mathbf{Z}$.

$$ty_i = \begin{cases} t, & t < i, \\ i + 2(t - i), & i \le t < i + 1, \\ t + 1, & i + 1 \le t. \end{cases}$$

Note that y_0 agrees with the map of the same name from (7.2). And again let $s : \mathbf{R} \to \mathbf{R}$ be such that $ts = t + 1$.

PROPOSITION 9.9. *The groups*

$$F_{\mathbf{R}} = \langle s, y_0 \rangle,$$
$$F_{\mathbf{R},\ge i} = \langle s, y_j \mid j \in \mathbf{Z},\, j \ge i \rangle$$

with $i \in \mathbf{Z}$, are all equal and are isomorphic to F, and the groups

$$F_{\ge i} = \langle y_j \mid j \in \mathbf{Z},\, j \ge i \rangle,$$

with $i \in \mathbf{Z}$ are also isomorphic to F.

PROOF. All the groups are not abelian. The relations $y_k y_j = y_j y_{k+1}$ hold for all $j < k$ as well as $y_k s = s y_{k+1}$ for all k. Thus the groups $F_{\mathbf{R}}$ and $F_{\mathbf{R},\ge i}$ are all equal. From Lemma 9.8, an isomorphism from $F_{\mathbf{R},\ge 1}$ to F as presented by (9.1) is built by taking s to x_0 and y_j to x_j for $j \ge 1$. An isomorphism from $F_{\ge i}$ to F is built by taking y_j to x_{j-i} for $j \ge i$. □

And we have a slight extension of Proposition 7.9.

PROPOSITION 9.10. *The elements of $F_{\ge i}$ are the elements of $F_{\mathbf{R}}$ that are the identity on $(-\infty, i]$.*

Now we can build smooth actions of F.

PROPOSITION 9.11. *There are infinitely differentiable actions of F on $\mathbf{R}$, $\mathbf{R}_{\geq 0}$ and $(0,1)$.*

PROOF. The nature of the functions y_i in Proposition 9.9 can be summarized as saying that y_i is the identity on $(-\infty, i]$, has slope 2 on $[i, i+1]$ and has slope 1 on $[i+1, \infty)$. The relation $y_k y_j = y_j y_{k+1}$ does not depend on the nature of y_i on $[i, i+1]$ as long as it successfully carries $[i, i+1]$ homeomorphically to $[i, i+2]$ preserving order. If y_0 is replaced by a smooth homeomorphism z_0 that equals y_0 off $[0,1]$ and is infinitely differentiable (see the techniques in the beginning of Chapter 2, Section 2 of [110]), then defining $z_i = s^{-i} z_0 s^i$ for $i \in \mathbf{Z} \setminus \{0\}$ gives a family of functions so that Proposition 9.9 holds if y_j is everywhere replaced by z_j. If $F^\infty_{\mathbf{R}}$ is what is obtained by that substitution on $F_{\mathbf{R}}$, then $F^\infty_{\mathbf{R}}$ can be conjugated to a smooth action on $(0,1)$ by a suitable modification of the arctan function. □

REMARK 9.12. The following representation of F appears in Dydak 1977 [63] and is attributed there to Minc. For each $n \in \mathbf{N}$ the permutation g_n on $\mathbf{Z} \times \mathbf{N}$ is defined by

$$g_n(j,k) = \begin{cases} (j,k), & j < n, \\ (n, 2k) & j = n, \\ (n, 2k+1) & j = n+1, \\ (j-1, k) & j > n+1. \end{cases}$$

The reader can show that $x_i \mapsto g_i$, $i \in \mathbf{N}$, extends to an isomorphism from F to the group generated by the g_n.

10. Properties, III: from the combinatorics

10.1. The commutator subgroup revisited. We can now describe the commutator subgroup. It was observed in Section 5.5 that the set of elements of F with support bounded away from 0 and 1 is exactly the kernel K of the "end point slope" homomorphism from F to $\mathbf{Z} \times \mathbf{Z}$ defined by $f \mapsto (\log_2(0f'_+), \log_2(1f'_-))$. We can now prove the following.

PROPOSITION 10.1. *The commutator subgroup F' of F equals K.*

PROOF. We know $F' \leq K$. The abelianization applied to the presentation (9.1) makes all the x_i with $i \geq 1$ equal to each other and trivializes the relations in the presentation. So the abelianization of F is $\mathbf{Z} \times \mathbf{Z}$, and modulo F' all elements of F are equivalent to words of the form $x_0^m x_1^n$. Under the end point slope homomorphism ϕ of Section

5.5, the image of x_1 is $(1,-1)$, the image of x_1 is $(0,-1)$, and the image of $x_0^m x_1^n$ is $(m, -(m+n))$. So any element of F not in F' is not in K, the kernel of ϕ. $\square$

COROLLARY 10.1.1. *The commutator subgroup of $F_{\mathbf{R}}$ is $BPL_2(\mathbf{R})$.*

PROOF. This follows from Proposition 10.1 and Remark 7.10. $\square$

As mentioned immediately before Proposition 5.13, it is not known if every non-identity element of F' is a single commutator. Equivalently, it is an open question if every non-identity element of F with support bounded away from 0 and 1 is a single commutator.

10.2. A universal property. We recall the endomorphism σ defined by $\sigma(x_i) = x_{i+1}$ of Lemma 9.4. For every x_i, we have $\sigma^2(x_i) = x_{i+2} = x_{i+1}^{x_0}$. Composing endomorphisms right to left, we have $\sigma^2 = C_{x_0}\sigma$ where C_{x_0} is conjugation by x_0. In order to argue that F is universal with respect to this behavior, we become more formal.

If a group G has an endomorphism ϕ and an element g so that $\phi^2 = C_g\phi$, then we will say that the triple (G, ϕ, g) forms a *group with a conjugacy idempotent*. We can make a category whose objects are groups with a conjugacy idempotent. A morphism from (G, ϕ, g) to (G', ϕ', g') is a homomorphism $\eta : G \to G'$ so that $\eta(g) = g'$ and for all $h \in G$, we have that $\eta\phi(h) = \phi'\eta(h)$.

Recall [144], Page 20, that an initial object in a category is an object a so that for each other object b there is a unique morphism from a to b. It is clear that any two initial objects in a category are joined by a unique isomorphism. We have the following.

PROPOSITION 10.2. *The triple (F, σ, x_0) with notation as above is an initial object in the category of groups with a conjugacy idempotent. Further if η is a morphism in that category from (F, σ, x_0) to some (G, ϕ, g), then either η is an injection, or the image of η is abelian and $\phi^i(g) = \phi^j(g)$ for all $1 \le i < j$.*

PROOF. Let (G, ϕ, g) be a group with conjugacy idempotent ϕ.

For $i \ge 0$, let $g_i = \phi^i(g)$ where we interpret $\phi^0(g)$ as g. For $i < j$ we have

$$\begin{aligned} g_i^{-1} g_j g_i &= \phi^i(g^{-1}\phi^{j-i}(g)g) \\ &= \phi^i(g^{-1}\phi(\phi^{j-(i+1)}(g))g) \\ &= \phi^i(\phi^2(\phi^{j-(i+1)}(g))) \\ &= \phi^{j+1}(g) = g_{j+1}. \end{aligned}$$

So the relations of the presentation (9.1) for F are satisfied and there is homomorphism $\eta_1 : F \to G$ taking x_i to g_i for all $i \geq 0$. If η is a morphism in the category from (F, σ, x_0) to (G, ϕ, g), then we must have $\eta(x_0) = g_0$, and the requirement that $\eta\phi(x_0) = \phi\eta(x_0)$ says that $\eta(x_1) = \phi(g_0) = g_1$. So $\eta = \eta_1$ since $\{x_0, x_1\}$ generates F.

The last claim follows from Proposition 5.12. □

Note that the uniqueness up to isomorphism of an initial object means that F can be discovered by studying groups with a conjugacy idempotent and noticing that there is an initial one. In fact this happened, and we will discuss the details in Section 19.

10.3. HNN extensions. See Chapter IV of [141] for basics on HNN extensions.

Given an isomorphism $\theta : A \to B$ between subgroups of a group G, we can form an HNN extension G_θ which is most easily defined if G is given as a presentation $G = \langle X \mid R \rangle$. A new generator t (an element not in G) is added to G and a relation set $S = \{a^t = \theta(a) \mid a \in A\}$ is formed so that G_θ is given as the presentation

$$G_\theta = \langle X \cup \{t\} \mid R \cup S \rangle.$$

It is easy to argue that S can be restricted to $\{a^t = \theta(a) \mid a \in Y\}$ if Y generates A. The element t in the above is usually called the *stable letter* of the extension (as in [141]). If one of A or B is G, the resulting HNN extension is usually called an *ascending* HNN extension.

The shift endomorphism $\sigma : F \to F$ is an isomorphism between F and the image of σ. If we form the HNN extension F_σ, we get the presentation

$$F_\sigma = \langle x_0, x_1, \ldots, t \mid x_j^{x_i} = x_{j+1},\ x_i^t = x_{i+1},\ \text{whenever } 0 \leq i < j \rangle.$$

Now sending $x_0 \in F$ to $t \in F_\sigma$ and $x_j \in F$ to $x_{j-1} \in F_\sigma$ for $j \geq 1$ preserves the relations of F. Thus we have a homomorphism from F to F_σ which is surjective since all the generators of F_σ are in the image and injective since the image is not abelian.

Thus F is isomorphic to an ascending HNN extension of itself using σ. Another view is that F is the HNN extension of the subgroup $\langle x_i \mid i \geq 1 \rangle$ with stable letter x_0. Now $\langle x_i \mid i \geq 1 \rangle$ is the HNN extension of the subgroup $\langle x_i \mid i \geq 2 \rangle$ with stable letter x_1, etc. Ken Brown once commented that if stripping off stable letters in search of the "core" of a group was an act of stripping off fluff, then F is all fluff and no core.

11. Combinatorics IV: Forests and the positive monoid

11.1. Introduction. The most important substructure of F is the monoid of all products of positive powers of the x_i from the presentation (9.1). We call this the *positive monoid* of F and denote it by F_+. The seminormal form (9.6) shows that every element f of F has the form $f = pn^{-1}$ where both p and n are from F_+. Note that F_+ is not related to the bi-order discussed in Section 4.7.

If a group G has a submonoid M so that every $g \in G$ is of the form $g = pq^{-1}$ with both p and q in M, then we say that G is the group of fractions of M. Various properties of M follow from this situation. It is immediate that M is cancellative and for every a and b in M, we have that $a^{-1}b = pq^{-1}$ for some p and q from M. So $bq = ap$ is a common right multiple of a and b and the monoid M is said to have common right multiples. This tells us that F_+ has common right multiples.

A celebrated theorem of Ore gives a converse in that every cancellative monoid with common right multiples embeds in and generates a group of fractions of M that is unique up to isomorphism. See Theorems 1.23 and 1.25 of [50] where the vocabulary is somewhat different. The existence of a group of fractions will be proven in a more general setting as Theorem 24.2.

If M has common right multiples, then for every a and b in M, we have the common right multiple $ap = bq$ and $a^{-1}b = pq^{-1}$. Now if M generates a group G, every $g \in G$ is a word in elements of M and their inverses which by the calculation just done can be arranged to have only positive powers of elements of M at the beginning of the word followed by only negative powers of elements of M to finish the word.

The positive monoid F_+ is isomorphic to a monoid of forests, and so F is isomorphic to a group of fractions of the monoid of forests. This will be important in Chapter 4 when complexes are built for the Thompson groups to act on. Here we will use forests to show that certain variations in the definition of F do not change F even though those variations will be shown in Chapter 6 to change the groups T and V.

A first goal of this section is to gain knowledge of the structure of F_+. Another is to use this knowledge to improve on the seminormal form (9.6) for elements of F and get a true normal form. This will involve relating the view of elements of F as words in a generating set, to the view of elements as pairs of partitions, prefix sets, or trees. Even more analysis of F_+ will occur when we build the complexes of Chapter 4.

We will get the improvement over the seminormal form (9.6) by exploiting what we know about irreducible tree pairs. Since the seminormal form (9.6) is about words in the $x_i^{\pm 1}$, we need to relate the structure of trees in a tree pair to words in a generating set. Uncovering this relation starts with the elements of F_+.

11.2. Turning F_+ into forests. We start with generators.

11.2.1. *Generators as tree pairs.* The tree pairs for x_0 and x_1 are

(11.1) $x_0 =$ 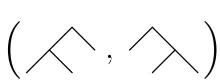 , $x_1 =$ 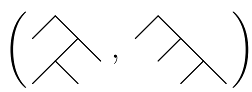 .

We defined $x_i = x_1^{x_0^{i-1}}$ in (9.3), and we calculate the tree pair for x_2 below. We show as dotted the added carets needed to construct the common binary refinements.

$x_2 = x_0^{-1}x_1x_0 =$ 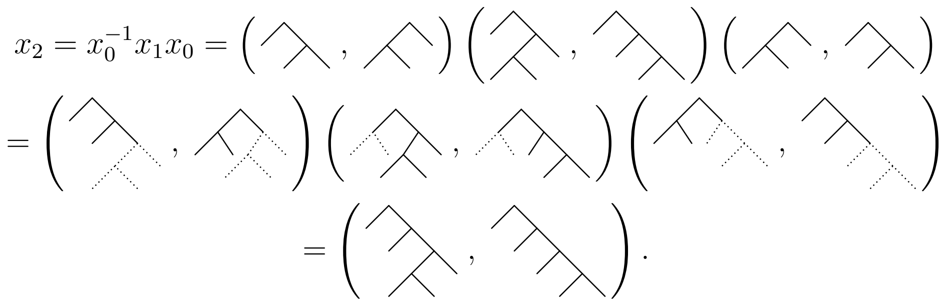

That the right trees of in x_0, x_1 and x_2 are all right vines (Definition 8.13) is not a coincidence. The leaves of a right vine V_n that are not the rightmost are all of the form 1^i0 for $0 \le i < n$. These vertices will be of particular importance, and in $\mathcal{T}$ we let $v_i = 1^i0$ for $i \ge 0$.

In x_0, x_1 and x_2, the left trees are almost right vines. To accomodate the extra carets needed to depict the left trees of the x_i, we will use $V_{n,m_1,\ldots,m_k}$ with $m_1 < m_2 < \cdots < m_k < n$ to denote the right vine V_n to which carets have been added at v_{m_i} for $1 \le i \le k$. Now $x_0 = (V_{1,0}, V_2)$, $x_1 = (V_{2,1}, V_3)$ and $x_2 = (V_{3,2}, V_4)$.

Since a binary splitting at the rightmost leaf of V_n produces V_{n+1}, we can repeatedly apply matched binary splittings at the rightmost leaves of the binary tree pair for x_0 to get $x_0 = (V_{n,0}, V_{n+1})$ for all $n > 0$. Similar statements apply to x_1 and x_2.

LEMMA 11.1. *For $i \ge 0$, we have $x_i = (V_{i+1,i}, V_{i+2}) = (V_{n,i}, V_{n+1})$ for all $n > i$.*

PROOF. The second form follows from the first by matched binary splittings at the rightmost leaf. In the following inductive calculation

$$\begin{aligned} x_{i+1} &= x_0^{-1} x_i x_0 \\ &= (V_2, V_{1,0})(V_{i+1,i}, V_{i+2})(V_{1,0}, V_2) \\ &= (V_{i+2}, V_{i+1,0})(V_{i+1,i}, V_{i+2})(V_{i+2,0}, V_{i+3}) \\ &= (V_{i+2,i+1}, V_{i+1,0,i})(V_{i+1,0,i}, V_{i+2,0})(V_{i+2,0}, V_{i+3}) \\ &= (V_{i+2,i+1}, V_{i+3}), \end{aligned}$$

the fourth line has had matched binary splittings done on the first pair at leaf $i+1$ and on the second pair at leaf 0. Since $i > 0$, leaf $i+1$ is at $1^{i+1}0$ in V_{i+2} and at $1^i 0$ in $V_{i+1,0}$. The reader is invited to draw corresponding pictures. □

11.2.2. *Special subtrees of a tree.* From Lemma 11.1 we know that those elements of F representable in the form (T, V_n) where T is a finite binary tree and n is the number of carets in T include all the generators x_i. Let F_+° temporarily denote the subset of those elements of F that can be represented in the form (T, V_n). We will show $F_+^\circ = F_+$. To do so, we look at the anatomy of the tree T in the pair (T, V_n).

We will make use of figures shown below.


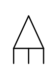

The figure on the left is a stylized representation of a typical tree, and the figure on the right is decorated to emphasize the presence of leaves. The picture is not meant to imply that there are exactly three leaves.

Recall that $v_i = 1^i 0$. Given a finite binary tree T, consider the infinite sequence $\Phi(T)$ where the i-th element of the sequence Φ_i is defined as

$$\Phi_i = \begin{cases} T_{v_i}/v_i & v_i \in T, \\ \cdot & \text{otherwise}, \end{cases}$$

where T_{v_i}/v_i is as in Definition 8.3 and $\cdot$ represents the trivial tree. Since only finitely many v_i are nodes of T and each T_i is finite, this is an infinite sequence of finite trees where only finitely many in the sequence are non-trivial. The picture below gives a picture of the beginning of the sequence (Φ_i).

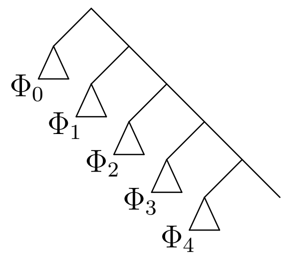

LEMMA 11.2. *Consider $f = (T_1, V_m)$ and $g = (T_2, V_n)$ in F_+°. The following are equivalent.*

(1) *$f = g$ in F.*
(2) *One of T_1 or T_2 can be obtained from the other by a sequence of binary splittings all of which are at the rightmost leaf.*
(3) *$\Phi(T_1) = \Phi(T_2)$.*

PROOF. Assume $f = g$. If $m = n$, then the last provision of Lemma 6.16 reinterpreted for tree pairs says that $T_1 = T_2$. If $m < n$, the matched binary splitings at the rightmost leaf of the pair for f changes the V_m of that pair to V_n and Lemma 6.16 now shows the altered T_1 equals T_2. Thus (1) $\Rightarrow$ (2).

If (2) holds, then $f = g$ since the splittings can be matched by splittings at the rightmost leaf V_m or V_n as appropriate, and the resulting right vines will have the same number of carets by Lemma 6.16 and be equal.

If (1) and (2) hold, then $\Phi(T_1) = \Phi(T_2)$ follows by noting that adding a caret to a tree T at its rightmost leaf does not change $\Phi(T)$.

If (3) holds and $m = n$, then $T_1 = T_2$ since $\Phi(T)$ and the number of carets in T completely determine the struture of a finite tree T. If $m < n$, then preforming $n - m$ matched splittings on (T_1, V_m) at the rightmost leaf produces the pair (T_1', V_n) where $\Phi(T_1') = \Phi(T_1) = \Phi(T_2)$ and T_1' and T_2 both have n carets. Thus $T_1' = T_2$ and $f = g$. □

11.2.3. *Forests.* The sequences $\Phi(T)$ motivate the next definition.

DEFINITION 11.3. A *finitary forest* is a sequence indexed over $\mathbf{N}$ of finite binary trees so that all but finitely many trees in the sequence are trivial. After this definition we will only include the word finitary in formal settings. Finite forests (finite sequences) will come later in Section 11.6. We let $\mathcal{F}$ be the set of finitary forests. Given a forest $\Psi \in \mathcal{F}$, we order the leaves of Ψ so that the leaves of each Ψ_i are given the prefix order and so that all leaves of Ψ_i come before all leaves of Ψ_j if $i < j$. The leaves of Ψ will be numbered in order from $\mathbf{N}$.

If $f = (T, V_n)$ is in F_+°, then from Lemma 11.2 we can write $\Phi(f)$ for $\Phi(T)$ and obtain a function from F_+° to $\mathcal{F}$ which is well defined and injective by Lemma 11.2, and which is clearly surjective.

11.2.4. *Multiplication.* If we want to multiply two elements of F_+°, then we start with a pair $(T_1, V_m)(T_2, V_n)$ of binary tree pairs as shown below.

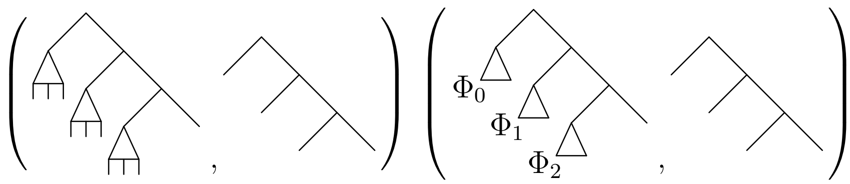

To proceed with the multiplication, a common refinement of V_m and T_2 must be found. To do this the trees in the sequence $\Phi(T_2)$ must be attached to leaves $v_i = 1^i0$ of V_m. If there are not enough leaves of the form v_i in V_m, then matched binary splittings of the pair (T_1, V_m) at the rightmost leaf can be done until there are enough such leaves. Once that is done, then the Φ_i can be attached simultenously to the v_i of V_m and the leaves of T_1. Specifically, if λ_i is the i-th leaf of T_1, counting from the leftmost leaf, then Φ_i must be attached to v_i in V_m and λ_i in T_1.

There is also the possibility that V_m has more carets than T_2 has along its right edge. In that case, matched binary splittings of (T_2, V_n) along the rightmost leaf can be done until there are enough carets in the right edge of T_2.

The result of the multiplication will be a pair (T_3, V_k) where $\Phi(T_3)$ is obtained from $\Phi(T_1)$ and $\Phi(T_2)$ by attaching the i-th tree of $\Phi(T_2)$ to the i-th leaf of $\Phi(T_1)$. We have the following.

LEMMA 11.4. *The set F_+° is closed under multiplication and is a monoid.*

We turn the multiplication in F_+° into a multiplication on $\mathcal{F}$ which we give a more formal treatment. We make use of the fact that nodes of a tree are elements of the monoid $\mathfrak{M} = \{0,1\}^*$. We define the product $\Psi\Theta$ of $\Psi \in \mathcal{F}$ and $\Theta \in \mathcal{F}$ by giving $(\Psi\Theta)_i$, the i-th tree in the sequence. We set

$$(\Psi\Theta)_i = \Psi_i \cup \bigcup_{v_j \in \Lambda(\Psi_i)} v_j\Theta_j \tag{11.2}$$

where the v_j are subscripted by the numbering of the leaves in Ψ and $v_j\Theta_j = \{v_j u \mid u \in \Theta_j\}$. In words, we obtain $\Psi\Theta$ by hanging, for each $j \in \mathbf{N}$, a copy of Θ_j on the j-th leaf of Ψ. That each $(\Psi\Theta)_i$ in (11.2) is finite is clear, and that it is a tree follows from Lemma 8.5.

We illustrate a product below. We have numbered the leaves of Ψ and the roots of both Θ and $\Psi\Theta$ and have used dotted lines for the edges in Θ. An infinite sequence of trivial trees is represented by $\cdots$.

(11.3)

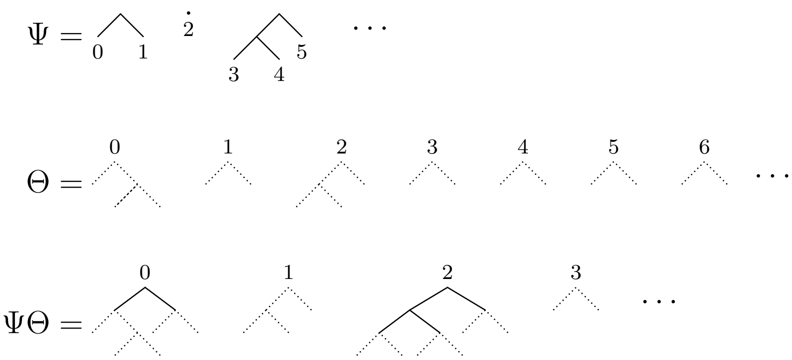

The multiplication given in (11.2) and the fact that the infinite sequence of trivial trees is an identity makes $\mathcal{F}$ a monoid that we refer to as the monoid of forests. Comparing the multiplication of (11.2) to the multiplication in F_+ shows that taking $(T, V_n) \in F_+$ to $\Phi(T) \in \mathcal{F}$ is a homomorphism. From Lemma 11.2 and the fact that any element of $\mathcal{F}$ is some $\Phi(T)$, we get the following.

Lemma 11.5. *The assignment $(T, V_n) \mapsto \Phi(T)$ is an isomorphism from F_+° to $\mathcal{F}$.*

For each $i \in \mathbf{N}$, we let $\nu_i = \Phi(x_i)$. From Lemma 11.1, each ν_i is the forest whose only non-trivial tree is the i-th tree which has only one caret. It is now obvious that for each $\Psi \in \mathcal{F}$ the product $\Psi\nu_i$ is obtained from Ψ by attaching a caret to the i-th leaf of Ψ. This and Lemma 11.5 give the following.

Lemma 11.6. *The ν_i generate $\mathcal{F}$, and the x_i generate F_+°.*

Corollary 11.6.1. *The monoids F_+ and F_+° are equal.*

Proof. The x_i generate F_+. □

Corollary 11.6.2. *For $f \in F_+^\circ = F_+$, we have that $\Phi(fx_i)$ is obtained from $\Phi(f)$ by adding a caret at the i-th leaf of $\Phi(f)$.*

Consequences of Corollary 11.6.2 will be explored in Section 11.4 and then in more detail in Section 23.3.

The observation made in Section 11.1 that F is a group of fractions of the monoid F_+ can be done by a cosmetically different argument.

Corollary 11.6.3. *Each $f \in F$ can be put in the form pn^{-1} with p and n in the positive monoid F_+ of F.*

PROOF. Represent f by a tree pair (D, R). Now with n the number of carets of both D and R, we have

$$f = (D, R) = (D, V_n)(V_n, R) = (D, V_n)(R, V_n)^{-1}$$

as required. □

11.3. Presentations. We have a parallel to the presentation (9.1) of F.

PROPOSITION 11.7. *The relations $x_j x_i = x_i x_{j+1}$ hold in F_+ for all $i < j$, and*

$$\langle x_0, x_1, x_2 \ldots \mid x_j x_i = x_i x_{j+1} \text{ whenever } i < j \rangle \tag{11.4}$$

is a monoid presentation for F_+. A normal form for elements of F_+ in the generators is

$$x_{i_0} x_{i_1} \cdots x_{i_k}, \quad i_0 \le \cdots \le i_k. \tag{11.5}$$

Replacing each x_i by ν_i in (11.4) and (11.5) gives a corresponding presentation and normal form for $\mathcal{F}$.

PROOF. The last point follows from the isomorphism between F_+ and $\mathcal{F}$ given by Lemma 11.5 and Corollary 11.6.1. So we work with whichever is more convenient.

The relations hold in F_+ because they hold in F. That the relations allow any word in the x_i to be put in the form (11.5) follows from the argument in Lemma 9.5 while ignoring inverses. We argue uniqueness in $\mathcal{F}$.

If words w and w' in the form (11.5) using the ν_i differ, then either they have different numbers of generators (in which case the corresponding forests have different numbers of carets), or they differ first from the left where one (say w) uses some ν_i and the other w' uses some ν_j with $i < j$. The longest common prefix of w and w' is some p giving a forest Φ. Now building the forest for w attaches a caret to Φ at the i-th leaf of Φ and building the forest for w' never does. □

REMARK 11.8. It is not hard to show that if a group G is a group of fractions of a monoid M, and $\langle X \mid R \rangle$ is a monoid presentation for M, then it is also a group presentation for G. Thus if one carefully avoids deriving the presentation (9.1) for F before deriving the presentation (11.4) for F_+ and showing that F is a group of fractions of F_+, then one could derive (9.1) from (11.4).

11.4. Words, trees and a normal form. We add a condition to the seminormal form (9.6) of words in the $x_i^{\pm 1}$ to arrive at a true normal form. We already have a normal form for elements of F represented as pairs of either binary partitions, prefix sets, or binary trees. The last of these will give a corresponding form for words. Corollary 11.6.2 tells us how to go from a word in positive powers of the x_i to the corresponding tree pair (T, V_n). Then given a word

$$w = x_{i_0}x_{i_1}\cdots x_{i_k}x_{j_l}^{-1}x_{j_{l-1}}^{-1}\cdots x_{j_0}^{-1}, \quad i_0 \le \cdots \le i_k \ne j_l \ge \cdots \ge j_0,$$

in seminormal form, we will turn the positive part $p = x_{i_0}x_{i_1}\cdots x_{i_k}$ and negative part $q^{-1} = (x_{j_0}x_{j_1}\cdots x_{j_l})^{-1}$ into a tree pair (D, R) and declare that w is in normal form if the tree pair is irreducible. We will turn this criterion into a condition on the entries in the sequences $(i_0, i_1, \cdots, i_k)$ and $(j_0, j_1, \cdots, j_l)$ that is easy to state and practical to detect.

11.4.1. *Single word to single tree.* Given $p = x_{i_0}x_{i_1}\cdots x_{i_k}$ in the positive powers of the x_i, Corollary 11.6.2 builds a forest $\Phi(p)$ caret by caret. Then a tree pair $(T(p), V_m)$ for p is obtained by building $T(p)$ with $\Phi(T(p)) = \Phi(p)$ and letting m be the number of carets in $T(p)$. Building $T(p)$ is done by taking some V_n with n sufficiently large, and hanging the tree $(\Phi(p))_i$ on leaf $v_i = 0^i1$ of V_n. We pause to discuss the exposed carets of $T(p)$ since exposed carets are key in a discussion of irreducible tree pairs.

One possible exposed caret in $T(p)$ will involve the rightmost leaf of $T(p)$ which will occur if n is chosen larger than necessary. If r is the largest so that $(\Phi(p))_r$ is not trivial, then any v_i in $T(p)$ with $i > r$ will be a leaf of $T(p)$. This will occur if the n used above has $n > r$. If $n > r$, then $i = n$ is the largest with v_i in $T(p)$, and v_i and the rightmost leaf of $T(p)$ will be the leaves of an exposed caret of $T(p)$. The leaves of this caret will have the same leaf numbers as the only exposed caret in V_m and the pair $(T(p), V_m)$ will not be reduced. From this point, we assume that $n = r$, and the rightmost leaf of $T(p)$ is not a leaf of an exposed caret. This makes $T(p)$ refer to a fixed tree and not a family of trees.

To discuss other exposed carets of $T(p)$, we take advantage of the normal form for words in F_+ and insist that in $p = x_{i_0}x_{i_1}\cdots x_{i_k}$, the sequence of subscripts satisfies $i_0 \le i_1 \le \cdots \le i_k$. We can argue the following.

LEMMA 11.9. *Let* $p = x_{i_0}x_{i_1}\cdots x_{i_k}$ *with* $i_0 \le i_1 \le \cdots \le i_k$, *and let* $T(p)$ *be as described above. Then the leaf numbers of the exposed carets in* $T(p)$ *are exactly the pairs* $(i_j, i_j + 1)$ *for all* i_j *where* $j = k$ *or*

$i_{j+1} \geq i_j + 2$. *Further, the pair* $(T(p), V_m)$ *where* m *is the number of carets in* $T(p)$ *is an irreducible tree pair representative of the element* p *in* F_+.

PROOF. With $n = r$ as discussed above, we build $T(p)$ from V_n by hanging the trees of $\Phi(p)$ on the v_i of V_n. Excepting the rightmost leaf of V_n, the leaves of $T(p)$ are exactly those gotten from the leaves of $\Phi(p)$. This is true even for the trivial trees of $\Phi(p)$. Since our choice of n prevents V_n from contributing to the exposed carets of $T(p)$, the exposed carets of $T(p)$ are those gotten from the exposed carets of $\Phi(p)$. At the time a caret is added corresponding to x_{i_j}, leaves i_j and $i_j + 1$ are exposed. If now $i_{j+1} = i_j$ or $i_{j+1} = i_j + 1$, then one of those previously exposed leaves has been "covered" by a new caret, is no longer a leaf, and no longer a leaf of an exposed caret. However if $i_{j+1} \geq i_j + 2$ (the only remaining possibility), then the exposed pair is left uncovered and will remain so for the rest of the word because of the non-decreasing order of the subscripts. The last caret to be added inevitably contributes an exposed caret. This covers all the cases in the claim related to exposed carets.

For the last sentence of the lemma, we note that our choice of $T(p)$ avoids an exposed caret that uses the rightmost leaf. Since V_m only has one exposed caret, and that caret uses the rightmost leaf, the pair $(T(p), V_m)$ is irreducible. ☐

11.4.2. *Word pairs to tree pairs.* We now return to

$$w = x_{i_0} x_{i_1} \cdots x_{i_k} x_{j_l}^{-1} x_{j_{l-1}}^{-1} \cdots x_{j_0}^{-1}, \quad i_0 \leq \cdots \leq i_k \neq j_l \geq \cdots \geq j_0,$$

in seminormal form, with positive part $p = x_{i_0} x_{i_1} \cdots x_{i_k}$ and negative part $q^{-1} = (x_{j_0} x_{j_1} \cdots x_{j_l})^{-1}$. From p, we get $(T(p), V_s)$, and from q, we get $(T(q), V_t)$. Since we may have $s \neq t$, let us assume $s < t$ and apply $t - s$ matched binary splittings to the rightmost leaves of $(T(p), V_s)$ to get the reducible $(T'(p), V_t)$ that still represents the element $p \in F_+$. Now

$$w = pq^{-1} = (T'(p), V_t)(V_t, T(q)) = (T'(p), T(q))$$

which might be reducible. However any reducibilty will not involve the rightmost leaves of $T'(p)$ and $T(q)$ since $T(q)$ has no exposed caret that uses its rightmost leaf. We can now state the following.

THEOREM 11.10. *Each element* f *of* F *is represented uniquely as a word in the form*

$$x_{i_0} x_{i_1} \cdots x_{i_k} x_{j_l}^{-1} x_{j_{l-1}}^{-1} \cdots x_{j_0}^{-1}, \quad i_0 \leq \cdots \leq i_k \neq j_l \geq \cdots \geq j_0 \qquad (9.6)$$

that satsifies the additional requirement

(†) *if* x_i *and* x_i^{-1} *are present in (9.6), then* x_{i+1} *or* x_{i+1}^{-1} *is present.*

Proof. That every $f \in F$ is representable by a word in the form (9.6) plus (†) follows because we already know that a word in the form (9.6) represents f. If a word in (9.6) violates (†), then in the notation of the statement, x_i and x_i^{-1} are present but not x_{i+1}. The requirement $i_k \neq j_l$ prevents the adjacency of x_i and x_i^{-1} in the word and so the subscripts of the non-empty subword u between them are all at least $i+2$. Now applying a relation from the presentation (9.1) removes x_i and x_i^{-1} and lowers all the subscripts of the subword u by one. This preserves (9.6), shortens the length of the word, and repeated applications must ultimately end in a word that also satisfies (†)..

It follows from Lemma 11.9 that a word in the form (9.6) with positive part p and negative part q^{-1} will have that $(T(p), T(q))$ as described above, after adjustment to guarantee an equal number of carets in each tree, is irreducible if and only if (†) holds. Now the uniquness of an irreducible tree pair from Proposition 8.9 for a given element of F completes the proof. □

11.5. Algorithms.

11.5.1. *Words to trees and back.* The discussion in Section 11.4 contains algorithms for going back and forth between two forms of representation of an element of F: words in the $x_i^{\pm 1}$ in the seminormal form (9.6) on the one hand, and binary tree pairs on the other. Also, reducing a word in the seminormal form (9.6) to the normal form (9.6) plus (†) is also algorithmic. Thus there is an algorithm behind the bijection between words in the normal form (9.6) plus (†) and reduced binary tree pairs that matches those objects that represent the same element of F.

The three equalities below illustrate points made above.

$$x_0^2 x_2 x_4 x_5 x_2^{-2} x_{-1} =$$

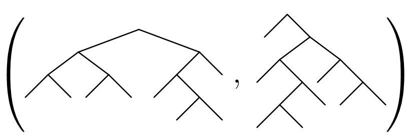

$$= x_0^2 x_3 x_4 x_2^{-1} x_{-1} =$$

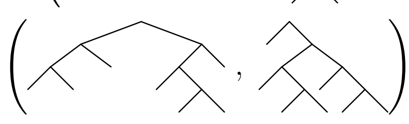

11.5.2. *Multiplying words.* From Section 8, we can view F as a set of irreducible binary tree pairs with an algorithm to perform multiplication that ends with a reduction of a tree pair to an irreducible pair. Similarly F can be viewed as a set of words in the normal form (9.6) plus (†) with an algorithm for multiplication that consists of concatenation followed by a reduction to the normal form.

As an example, we compute the square of $x_0x_2x_4x_5^{-1}x_3^{-1}x_1^{-1}$. Note that this is in the normal form (9.6) plus (†). We have

$$\begin{aligned}&(x_0x_2x_4x_5^{-1}x_3^{-1}x_1^{-1})^2\\=&(x_0x_2x_4x_5^{-1}x_3^{-1}x_1^{-1})(x_0x_2x_4x_5^{-1}x_3^{-1}x_1^{-1})\\=&x_0^2x_3x_5x_6^{-1}x_4^{-1}x_2^{-1}x_2x_4x_5^{-1}x_3^{-1}x_1^{-1}\\=&x_0^2x_3x_5x_6^{-1}x_5^{-1}x_3^{-1}x_1^{-1}\\=&x_0^2x_4x_5^{-1}x_4^{-1}x_1^{-1}\end{aligned}$$

The last transformation eliminates a violation of (†). The reader can do this calculation using tree pairs. Calculating with words can take less space, but not always less time. It is also interesting to see that intermediate steps in the calculation do not have to correspond precisely.

We take this opportunity to formally and belatedly announce the following. The point could have been made many times before this.

PROPOSITION 11.11. *The word problem for F is solvable.*

11.6. More fractions for F. We know that F is the group of fractions of F_+ and thus also of the monoid of finitary forests $\mathcal{F}$. This view of F can be turned into a partition pair view of $F_{\geq 0}$ by taking a finitary forest $\Phi \in \mathcal{F}$ as a set of instructions to partition $[0, \infty)$. The i-th tree Φ_i tells how to partition the interval $[i, i+1]$.

The group F can also be represented by pairs of finite forests. In some sense this has already covered in Lemma 5.5. We review this here because this variation of an approach to F does not change the isomorphism type of the resulting group, but we will see in Chapter 6 that a parallel variation changes the isomorphism types when applied to T and V. All the work below will be left to the reader.

DEFINITION 11.12. A *finite forest* is a finite sequence of finite trees with the leaves ordered as for a finitary forest (Definition 11.3). The forest is binary if all the trees in it are binary.

From this point all trees and forests will be binary.

Let $\psi = (T_0, \ldots, T_k)$ be a finite binary forest. Then the finitary forest Ψ associated to ψ is the concatenation of sequences $(T_0, \ldots, T_{k-1})$ followed by the infinite but finitary forest $\Phi(T_k)$ where $\Phi(T_k)$ is as described in Section 11.2.2.

Let an integer $r \geq 1$ be fixed. We consider finite binary forests of length r. That is, r is the number of trees (or the number of roots) in the forest. On such objects, one can define binary splittings and binary refinements, and on pairs (ψ_1, σ, ψ_2) with σ a bijection from the leaves of ψ_1 to those of ψ_2, and one can define matched binary splitings, and corresponding equivalence classes. As usual, if σ is omitted it is assumed to be the unique order preserving bijection. We denote the set of equivalence classes by $F_{2,r}$. The notation anticipates Chapter 6 where the 2 emphasizes the binary aspect.

The reader can define the usual opportunistic multiplication to turn $F_{2,r}$ into a group and use the finitary forest associated to a finite forest to prove the first claim in the proposition below. For the second claim, a group $F_{\mathcal{F}}$ can be defined on classes of pairs of finitary, binary forests in the usual way. The second claim summarizes the comments at the beginning of the section.

PROPOSITION 11.13. *The group $F_{2,r}$ is isomorphic to F, and also naturally isomorphic to the group $F_{[0,r]}$. The group $F_{\mathcal{F}}$ is isomorphic to F, and also naturally isomorphic to the group $F_{\geq 0}$.*

The extraction described above of a finitary forest from a finite forest shows how to conjugate $F_{[0,r]}$ to $F_{\geq 0}$ in a manner similar to that in Remark 7.10.

We illustrate below the forest pair for the element corresponding to x_0 when $r = 2$.

$$\left(\wedge\ \cdot\ ,\ \cdot\ \wedge\right)$$

The reader can work out the pairs for x_0 for $F_{2,r}$ when $r \geq 3$. The forest pair in $F_{\geq 0}$ is made obvious by the element y_0 of (7.2) and the element ν_0 of Lemma 11.6.

12. The group T

As for F there is a definition of T as a group of piecewise linear maps, and also definitions that are more combinatorial. We will start with the PL definition.

12.1. The PL definition. We first need to discuss some structure on the circle.

DEFINITION 12.1. We use S^1 to denote $\mathbf{R}/\mathbf{Z}$ and refer to it as the *circle of length one*. We use $\mathrm{Homeo}_+(S^1)$ to denote the group of orientation preserving self homeomorphisms of S^1. Let $p : \mathbf{R} \to S^1$ be the corresponding projection map. A lift of $h \in \mathrm{Homeo}_+(S^1)$ to $\mathbf{R}$ is some $\widetilde{h} \in \mathrm{Homeo}_+(\mathbf{R})$ so that for all $t \in \mathbf{R}$, we have $t\widetilde{h}p = tph$. The group $PL_+(S^1)$ is the subgroup of $\mathrm{Homeo}(S^1)$ whose lifts are in $PL_+(\mathbf{R})$. We can refer to an element of S^1 as being *dyadic* (being in $\mathbf{Z}[\frac{1}{2}]$) if its representatives are dyadic.

Note that all lifts of the identity in $\mathrm{Homeo}_+(S^1)$ are integral powers of the translation $s \in \mathrm{Homeo}_+(\mathbf{R})$ defined by $ts = t+1$. Given two different lifts of an $h \in \mathrm{Homeo}_+(S^1)$, one will equal the other composed with an integral power of s. Every lift of an element of $\mathrm{Homeo}_+(S^1)$ commutes with s.

DEFINITION 12.2. The group T is the subgroup of $PL_+(S^1)$ satisfying the following.

(1) At every point where $h \in T$ has a derivative, it acts on an open neighborhood of the point as $t \mapsto 2^n t + b$ with $n \in \mathbf{Z}$ and $b \in \mathbf{Z}[\frac{1}{2}]$.
(2) The *breakpoints* of $h \in T$ are those points where h' does not exist, are finite in number, and are elements of $\mathbf{Z}[\frac{1}{2}]$.

Note that the set of breakpoints can be empty since T contains all rotations (translations) of S^1 by dyadic rationals. The careful wording of (1) in the definition is to rule out rotations by amounts that are not dyadic rationals. We could have given a less careful wording of (1) and added a third requirement that the dyadic rationals be preserved as a set. It is left to the reader to verify that elements of T do preserve the dyadics and that T is a group under composition.

12.2. Elements and multiplication. Below we show the graph of two typical elements of T where we treat S^1 as the interval $[0, 1]$ with the points 0 and 1 identified. The first has breakpoints at 0, $\frac{1}{2}$ and $\frac{3}{4}$. The second is rotation by $\frac{1}{4}$ and has no breakpoints.

(12.1)

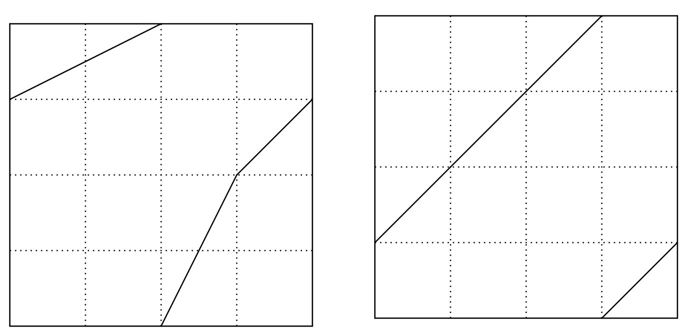

The fact that T contains all rotations by dyadic rationals gives a quick look at some of its structure.

PROPOSITION 12.3. *The following hold.*

(1) *For each dyadic rational d, the subgroup T_d of T that fixes d is isomorphic to F, and all these T_d are conjugate in T.*
(2) *Every element of T is a product fg where $f \in T_0$ and g is a rotation by a dyadic rational.*

(3) *Every element of T is a product hj with $h \in T_0$ and $j \in T_d$ for some dyadic rational d.*
(4) *Every element of T is a product of no more than two elements each of which has fixed set with non-empty interior.*
(5) *For each dyadic $d \in S^1$, the group T_d is maximal in T.*

PROOF. (1) That T_0 is isomorphic to F is immediate from the definitions of T and F. Now T_d is a conjugate of T_0 by a rotation of S^1 by d.

For (2–4), let $k \in T$ not be the identity.

(2) If $0k \neq 0$, then $0k$ is a dyadic rational d. With g rotation by d, we have $f = kg^{-1} \in T_0$.

(3) Again if $0k \neq 0$, we now pick some dyadic rational d different from both 0 and $0k$. Using the transitivity properties of T_d, and the fact that all of 0, $0k$ and d are dyadic rationals, there is an element $j \in T_d$ that carries 0 to $0k$, and we have $h = kj^{-1} \in T_0$.

(4) With k not the identity, there is an open set U with Uk disjoint from U. By choosing U small enough we may assume that the closure of $U \cup Uk$ misses an open set W in S^1. Let d be a dyadic in W. By the transitivity properties of T_d, there is some g in T_d that is fixed on a neighborhood of d in S^1 and that agrees with k on a non-degenerate interval J in U. Now $h = kg^{-1}$ is the identity on J and $k = hg$.

(5) Since all T_d are conjugate, so we can work with T_0. If $G \leq T$ contains T_0 and an element h with $d = 0h \neq 0$, then G contains T_d and by the transitivity of T_0 on the dyadics, G contains all T_s with s dyadic in S^1. Now we are done by (3). □

We skip several ways to represent elements of T and go directly to pairs of trees. The reader can use the material of Sections 6 and 7 to derive other representations.

Definition 8.6 gives a tree pair as a triple (T, σ, S), but this overuses the letter T in a section on the group T. So we will consistently switch to (D, σ, R) where the binary trees D and R are to suggest "domain" and "range." For the group T, there are more possibilities for the bijection σ from $\Lambda(D)$ to $\Lambda(R)$ than there are for the group F, but the choices are still somewhat limited. We give the specifics.

DEFINITION 12.4. A *marked binary tree pair* is a binary tree pair (D, σ, R) as in Definition 8.6 with the bijection $\sigma : \Lambda(D) \to \Lambda(R)$ restricted as follows. With $n = |\Lambda(D)| = |\Lambda(R)|$, and u_i and w_i denoting the i-th leaves, respectively, of D and R, there is a k with $0 \leq k < n$ so that for all i with $0 \leq i < n$, we have $u_i\sigma = w_{i+k}$ where the addition in the subscript of w is modulo n.

In words, the bijection σ is rotation by k. It is useful to see this as a piecewise order preserving map that interchanges the relative positions of the two intervals consisting of the first $n-k$ elements and the last k elements.

The only information beyond D and R in a marked binary tree pair that is needed is the value of k. This will be reflected in our drawings of pairs of trees that represent elements of T. In the notation of Definition 12.4, it suffices to know which leaf of $\Lambda(R)$ is $u_0\sigma = w_k$. In a drawing, we can indicate this with a "bullet" $\bullet$ at the leaf w_k. Below we show marked binary tree pairs for the two elements shown in (12.1).

(12.2) 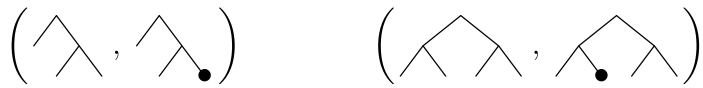

Note that both pairs in (12.2) are irreducible.

The figures in (12.2) should justify the terminology "marked binary tree pair." Note that elements of F can be represented this way by putting the mark on the leftmost leaf of R. This repeats the fact that the subgroup of T that fixes 0 is isomorphic to F.

We describe a matched binary splitting of a marked binary tree pair (D, σ, R) with u_i, w_i and k as in Definition 12.4 to yield (D', σ', R'). If the splitting of D is at the leaf u_i of D, then the splitting of R is at $u_i\sigma = w_{i+k}$. The bijection σ' agrees with σ off u_i and it takes u_i0 to $w_{i+k}0$ and u_i1 to $w_{i+k}1$. The marked vertex of R' depends on whether the splitting of D is at the leftmost leaf u_0 of D or not. If the splitting of D is at the leftmost leaf u_0, then the marking of R' is at w_k0. Otherwise the marking of R' is at the same leaf as the marked leaf of R.

Starting with the marked tree pair on the right in (12.2) for rotation of S^1 by $\frac{1}{4}$, the pair below on the left is a mathched binary splitting at u_0 and the pair below on the right is a matched binary splitting at u_3.

(12.3) 

Multiplication is opportunistic, as expected. Given two marked binary tree pairs $(D_1, \sigma_1, R_1)(D_2, \sigma_2, R_2)$, the result is $(D_1, \sigma_1\sigma_2, R_2)$ if $R_1 = D_2$. Otherwise multiplication is done by doing matched binary splittings to get (D_1', σ_1', R_1') and (D_2', σ_2', R_2') where $R_1' = D_2'$ (ignoring the marking) to give the result of the multiplication as $(D_1', \sigma_1'\sigma_2', R_2')$.

If (12.3) is taken as a multiplication problem then the result is

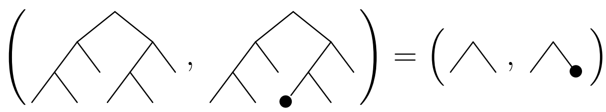

after reduction. Note that during the reduction the leftmost exposed caret of the left tree "cancels" the rightmost exposed caret of the right tree, and vice versa. Of course this particular calculation simply verfies that the square of rotation by $\frac{1}{4}$ is rotation by $\frac{1}{2}$.

We observe that if (D, σ, R) represents an element $f \in T$, then (R, σ^{-1}, D) represents f^{-1}. The inverses of the pairs in (12.2) are as below. The identical trees in each pair do not illustrate the interchange of D and R, but that should by now be less important than the change in the marking.

(12.4) 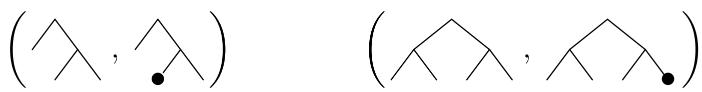

LEMMA 12.5. *Every element of T can be represented by a marked binary tree pair, every marked binary tree pair represents an element of T, and two marked binary tree pairs represent the same element of T if and only if they are related by the equivalence relation generated by matched binary splittings.*

PROOF. For the first claim both elements of F and rotations by dyadic rationals are representable by marked binary tree pairs, and so their product can also be so represented. The second claim follows because the resulting action fits the definition of elements of T. The last claim is clear in one direction, and the details of the other direction are left to the reader guided by the material in Sections 6 through 8. $\square$

Details for the following parallels to statements in Sections 6 through 8 can be supplied by the reader.

LEMMA 12.6. *Every element of T has an irreducible representative (that is minimal as measured by numbers of carets) by marked binary tree pairs from which all other representatives are obtained by matched binary refinements. The only representatives of the identity of T are of the form (D, σ, D) where σ is the identity and the irreducible representative of the identity uses the trivial tree.*

REMARK 12.7. There is a smooth, faithful action of T on the circle. See [82], Theorem A.

12.3. Finite subgroups.

PROPOSITION 12.8. *Every finite subgroup of T is cyclic, and every cyclic group is isomorphic to a subgroup of T.*

PROOF. The first claim has little to do with T and applies to $\mathrm{Homeo}_+(S^1)$. Let H be a finite subgroup of $\mathrm{Homeo}_+(S^1)$ and let t_0 be a point in S^1. The orbit $t_0 \cdot H$ is finite. Let $t_0, t_1, \ldots, t_{n-1}$ be a listing of $t \cdot H$ in counterclockwise order around S^1. There is an $h \in H$ with $t_0 h = t_1$. Since h takes the interval (t_0, t_1) in S^1 that contains no other points of $t \cdot H$ to an interval with similar behavior, it follows that $t_1 h = t_2$. Inductively $t_i h = t_{i+1}$ treating arithmetic on the subscripts modulo n.

If h does not generate H, then let g be some element of H that is not a power of h. We have $t_0 g = t_i$ for some i and $t_0 g h^{-i} = t_0$. But gh^{-i} is a non-identity element of $\mathrm{Homeo}_+(S^1)$ with a fixed point and must have infinite order. This contradicts the fact that H is finite.

For the second claim, with V_{n-1} the right vine of $n-1$ carets and n leaves, the marked binary tree pair $(V_{n-1}, \sigma, V_{n-1})$ as depicted by

(12.5)

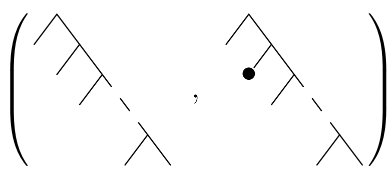

generates a copy of the cyclic group of order n in T. □

12.4. Rotation numbers.

DEFINITION 12.9. Given $h \in \mathrm{Homeo}_+(S^1)$, a lift $\tilde{h}$ in $\mathrm{Homeo}_+(\mathbf{R})$ and $t \in \mathbf{R}$, then the *rotation number* of h is

$$\tau(h) = \left[\lim_{|n|\to\infty} \frac{1}{n}(t\tilde{h}^n - t)\right] \mod \mathbf{Z}. \tag{12.6}$$

The defined term makes no mention of t because the quantity defined does not depend on t. See Chapter 11 of [121] for a proof that $\tau(h)$ exists for all $t \in \mathbf{R}$, is independent of t and of the choice of lift of h, and is an invariant of conjugacy.

In a later edition we will prove that all rotation numbers of elements of T are rational. For now we only establish a converse.

LEMMA 12.10. *For every $r \in \mathbf{Q} \cap [0,1)$, there is an element of T whose rotation number is r.*

PROOF. Let $r = \frac{m}{n}$ in reduced terms with $0 < m < n$ in $\mathbf{N}$ and consider $f = (V_{n-1}, \sigma, V_{n-1})$ as depicted in (12.5) of Proposition 12.8.

The orbit S of 0 in S^1 under $h = f^m$ contains n points. The elements s_i of the preimage $\widetilde{S}$ in $\mathbf{R}$ of S under the projection $p : \mathbf{R} \to \mathbf{R}/\mathbf{Z} = S^1$ can be indexed by $\mathbf{Z}$ in increasing order so that $s_0 = 0$. Each half open interval of length one in $\mathbf{R}$ contains exactly n points of $\widetilde{S}$.

There is a lift $\tilde{h}$ of h so that the action of $\tilde{h}$ on the indexes of the elements of $\widetilde{S}$ is $i \mapsto i + m$. Now the action of $\tilde{h}^{nk}$ on the indexes is as $i \mapsto i + mnk$. But $s_{mnk} - s_0 = mk$ and the limit in (12.6) is $\frac{m}{n} = r$. $\square$

12.5. Transitivity. The transitivity properties of F are given by Lemma 5.3. We give a straightfoward adaptation for T.

LEMMA 12.11. *Let $I_0 < I_1 < \cdots < I_{n-1}$ be a sequence of single points and closed intervals with non-empty, pairwise disjoint interiors in $[0, 1]$ with the single points and the endpoints of the intervals in $\mathbf{Z}[\frac{1}{2}]$. Let π act on $k = \{0, 1, \ldots, n-1\}$ as addition of a constsant modulo n. Let $J_0 < J_1 < \cdots < J_{n-1}$ be similarly described so that $J_{i\pi}$ is a point if and only if I_i is a point and intervals $J_{i\pi}$ and $J_{i\pi+1}$ share an endpoint if and only if I_i and I_{i+1} share an endpoint. For each interval I_i, let g_i be an element of $PL_+(I)$ whose restriction to I_i has image $J_{i\pi}$ and which satisfies Definition 12.2. Then there is an element of T taking each I_i to $J_{i\pi}$ and whose restriction to each interval I_i equals g_i.*

12.6. Simplicity.

PROPOSITION 12.12. *The group T is simple.*

PROOF. Let $f \in T$ not be the identity in a normal subgroup N of T. There is an open U in S^1 so that Uf is disjoint from U and there is a dyadic $d \in U$. In T_d there is an element g whose support is in U so $c = [g, f] = g^{-1}g^f$ is in N and is not trivial since g and g^f have disjoint support. Also d is not in the support of g by choice and not in the support of g^f which is in Uf. So c is in T_d. By Proposition 5.12, N contains the commutator subgroup of T_d which by Proposition 10.1 consists of all elements of T that fix an open neighborhood of d.

Since T is transitive on the dyadics in S^1, we know that N contains all elements of T that fix an open set containing a dyadic. But every open set contains a dyadic, so N contains all elements of T whose fixed set has non-empty interior. We are done by (4) of Proposition 12.3. $\square$

12.7. Free subgroups. We will find a non-cyclic free subgroup of T. The argument used, known as the "ping-pong" argument, is usually attributed to Klein [125, §III, 16] or Fricke and Klein [74, §II, 3.8].

There is an element $x \in T$ fixed at 0 and $\frac{1}{2}$ and no other point, and that takes $\frac{1}{16}$ to $\frac{7}{16}$ and $\frac{15}{16}$ to $\frac{9}{16}$. Conjugating x by rotation by $\frac{1}{4}$ gives

$y \in T$ fixed on $\frac{1}{4}$ and $\frac{3}{4}$ and no other point, and that takes $\frac{5}{16}$ to $\frac{11}{16}$ and $\frac{3}{16}$ to $\frac{13}{16}$.

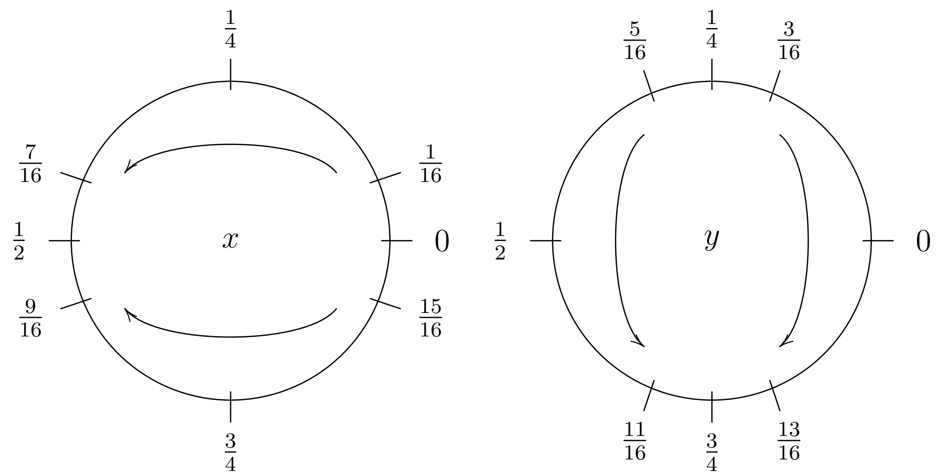

Proposition 12.13. *The subgroup $\langle x, y\rangle \leq T$ is freely generated by x and y.*

Proof. Let w be a non-trivial, freely reduced word in x, y, x^{-1} and y^{-1}. We will show that $\frac{1}{8}w \neq \frac{1}{8}$.

Let $X = (\frac{7}{16}, \frac{9}{16})$, $X^{-1} - = (\frac{15}{16}, \frac{1}{16})$, $Y = (\frac{11}{16}, \frac{13}{16})$ and $Y^{-1} = (\frac{3}{16}, \frac{5}{16})$ where the exponent is purely formal. The interval notation (a, b) is to be interpreted as those t with $a < t < b$ in counterclockwise order on the circle. None of these four sets contains $\frac{1}{8}$. Note that x takes all points not in X^{-1} into $\overline{X}$, the closure of X, x^{-1} takes all points not in X into $\overline{X^{-1}}$ and similar statements for $y^{\pm 1}$ and the intervals Y and Y^{-1}.

If a is the last letter of w, then the claim is that $\frac{1}{8}w$ is in X if $a = x$, in X^{-1} if $a = x^{-1}$, in Y if $a = y$, and in Y^{-1} if $a = y^{-1}$. That is $\frac{1}{8}w$ is in the capitalized version of a. It is true if $|w| = 1$ and if $w = w'a$ for non-empty w', then a corresponding statement is true for $\frac{1}{8}w'$ and the last letter of w'. The last letter of w' is not a^{-1} and so $\frac{1}{8}w'$ is not in the capitalized version of a^{-1}. By the previous paragraph $\frac{1}{8}w = \frac{1}{8}w'a$ is in the claimed interval. □

Remark 12.14. The method used in the proof of Proposition 12.13 is standard and worth demonstrating. It is even easier to show that T contains an isomorphic copy of $PSL(2, \mathbf{Z})$. This is isomorphic to the free product of $\mathbf{Z}/3\mathbf{Z}$ and $\mathbf{Z}/2\mathbf{Z}$ [4] and contains non-abelian free groups [162]. The generators in T are

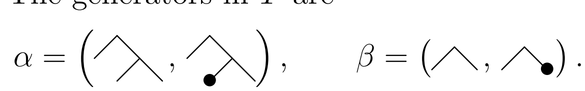

A word w in $\{\alpha, \beta\}$ can be reduced to an alternation of $\alpha^{\pm 1}$ and β. The reader can show that if such a w starts with $\alpha^{\pm 1}$, then $[0, \frac{1}{2}]$ is taken to a strictly shorter interval by w, and that if w starts with β, then $w = \beta$ and is not the identity, or $[\frac{1}{2}, 1]$ is taken to a strictly shorter interval.

Theorem 4.11 shows that $PLF_+(\mathbf{R})$, the group of orientation preserving PL self homeomorphisms of $\mathbf{R}$ with only finitely many breakpoints, contains no subgroup isomorphic to the free group on two generators. We now have the importance of the finiteness assumption.

COROLLARY 12.14.1. *There is a subgroup of $PL_+(\mathbf{R})$ isomorphic to the free group of rank 2.*

PROOF. Let x and y be as in Proposition 12.13, and let $\tilde{x}$ and $\tilde{y}$ be any lifts of x and y, repsectively, to $\mathbf{R}$. The elements $\tilde{x}$ and $\tilde{y}$ both commute with the shift map s where $ts = t + 1$, so every element in $\langle \tilde{x}, \tilde{y} \rangle$ commutes with s. It follows that $\tilde{x} \mapsto x$ and $\tilde{y} \mapsto y$ extends to a well defined homomorphism from $\langle \tilde{x}, \tilde{y} \rangle$ onto $\langle x, y \rangle$. Since x and y freely generate the latter, $\tilde{x}$ and $\tilde{y}$ must freely generate the former. □

12.7.1. *An application to F.* A group G *satisfies a law* if there is a non-identity element w in some group F_S free on a finite set S so that every homomorphism from F_S to G has w in its kernel. Using the standard embedding of a finitely generated free group into the free group of rank 2 (see Proposition 3.1 in [141]) we can assume S has two elements.

PROPOSITION 12.15. *The Thompson group F satisfies no laws.*

PROOF. We will work within the isomorphic copy $F_{\mathbf{R}}$ of F from Section 7.6.1.

Let w be a non-identity element in the free group $\langle \tilde{x}, \tilde{y} \rangle$ from the proof of Corollary 12.14.1, and note that we can insist that the lifts $\tilde{x}$ and $\tilde{y}$ have infinitely many fixed points each. Let n be the length of w in $\{\tilde{x}^{\pm 1}, \tilde{y}^{\pm 1}\}$ and let p_i, $0 \le i \le n$, be the prefixes of w with p_0 the empty word, and $p_n = w$. Since w is not the identity, there is a $t \in \mathbf{R}$ for which $tw \ne t$ and let $t_i = tp_i$, $0 \le i \le n$.

Let K be a compact interval that contains all the t_i. There is a compact interval L containing K, and elements $\tilde{x}'$ and $\tilde{y}'$ in F_∞ that agree with $\tilde{x}$ and $\tilde{y}$ on K and are the identity off L. Sending $\tilde{x}$ to $\tilde{x}'$ and $\tilde{y}$ to $\tilde{y}'$ is a homomorphism h from $\langle \tilde{x}, \tilde{y} \rangle$ into F_∞, and we let w' be the image of w under h. Our construction gives $tw' = tw \ne t$ and w is not in the kernel of h. □

12.8. A finite presentation. As with F we find a seminormal form for elements of T. And like Proposition 9.6 for F, we show that the seminormal form is powerful enough because it only has one representative for the identity element.

We identify F with T_0 and regard the usual generators x_0 and x_1 of F as elements of $T_0 \subseteq T$. We give the name $\pi_{0,1}$ to the rotation of S^1 by $\frac{1}{2}$,

$$\pi_{0,1} =$$ 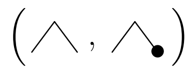 ,

since it switches (permutes) the two halves of S^1.

Since $\{x_0, x_1\}$ generates T_0, Item (5) of Proposition 12.3 gives that $\{x_0, x_1, \pi_{0,1}\}$ generates T.

Our seminormal form will be based on the following.

LEMMA 12.16. *The group T is the disjoint union of F and $F\pi_{0,1}F$.*

PROOF. The sets are disjoint since every element of F fixes 0 and no element of $F\pi_{0,1}F$ does. Let $g \in T$ not be in F. Since $0 \neq 0g \in \mathbf{Z}[\frac{1}{2}]$, there is $f \in F$ with $0gf = \frac{1}{2}$. So $0gf\pi_{0,1} = 0$ and $gf\pi_{0,1} = f' \in F$. Now $g = f'\pi_{0,1}f^{-1} \in F\pi_{0,1}F$. □

We say that an element of T is in seminormal form if it is expressed either as f or $f\pi_{0,1}f'$ with f and f' in F. From the lemma, the only seminormal form for the identity of T is the identity in F.

LEMMA 12.17. *The group F is the disjoint union of the three sets*

$$S = \{f \in F \mid \tfrac{1}{2}f = \tfrac{1}{2}\},$$
$$Sx_0S = \{f \in F \mid \tfrac{1}{2}f > \tfrac{1}{2}\},$$
$$Sx_0^{-1}S = \{f \in F \mid \tfrac{1}{2}f < \tfrac{1}{2}\},$$

PROOF. The sets on the right clearly partition F, so we need to show that they are equal to the sets on the left. The first is a definition and the argument for the third is similar to that of the second. In the second, the left side is clearly contained in the right. For the other containment, if $f \in F$ has $\frac{1}{2}f > \frac{1}{2}$, then some $g \in S$ has $\frac{1}{2}fg = \frac{3}{4}$ making $g' = fgx_0^{-1}$ an element of S. So $f = g'x_0g^{-1} \in Sx_0S$. □

To define the relations for T, we will use the operation $\wedge$ on F from Definition 8.16 where $f \wedge g$ is the composition of f_0g_1 of deferments of f and g. Note that $x_1 = 1 \wedge x_0$, and that the image of $\wedge$ is S.

With $X = \{x_0, x_1, \pi_{0,1}\}$, we let R consist of the following relations.

(1-2) The two defining relations for F from (3.2).

(3) $\pi_{0,1}^2 = 1$.

(4) $\pi_{0,1}(1 \wedge x_0)\pi_{0,1} = x_0 \wedge 1$.

(5) $\pi_{0,1}(1 \wedge x_1)\pi_{0,1} = x_1 \wedge 1$.
(6) $(\pi_{0,1}x_0)^3 = 1$.

PROPOSITION 12.18. *A presentation for T is $\langle X \mid R\rangle$.*

PROOF. The validity of the relations (1–5) in T is clear and (6) is a straightforward check using tree pairs. The more useful forms of relation (6) for us are $\pi_{0,1}x_0\pi_{0,1} = x_0^{-1}\pi_{0,1}x_0^{-1}$, and $\pi_{0,1}x_0^{-1}\pi_{0,1} = x_0\pi_{0,1}x_0$ since in each, replacing the left side by the right side reduces the number of appearances of $\pi_{0,1}$.

In the following f, f' and f'' will be understood to represent elements of F, and for various i, s_i will be understood to represent elements of S.

Since X generates T, relation (3) says that every element of T reduces to a word consisting of alternations of words in $\{x_0, x_1\}$ and single copies of $\pi_{0,1}$. We are done if we can reduce such a word to our seminormal form of f or $f\pi_{0,1}f'$. This will follow if we can reduce an arbitrary word to one having no more than one appearance of $\pi_{0,1}$. To do this we show that every $\pi_{0,1}f\pi_{0,1}$ can be replaced by some $f'\pi_{0,1}f''$.

Since x_0 and x_1 generate F, that $\pi_{0,1}$ normalizes S follows formally from (4) and (5). So we only need to consider $f \notin S$. If $f = s_1x_0s_2$, then from the first alternate form for (6), we get

$$\pi_{0,1}(s_1x_0s_2)\pi_{0,1} = (s_1'x_0^{-1})\pi_{0,1}(x_0^{-1}s_2')$$

for some $s_1', s_2' \in S$ which contains fewer appearances of $\pi_{0,1}$ than $\pi_{0,1}f\pi_{0,1}$. A similar argument holds if $f \in Sx_0^{-1}S$ using the second alternate form of (6). □

Relation (6) is often called the hexagon relation and an efficient verification using parenthesized expressions is given in (16.6).

13. The group V

Section 7.6.4 presents F as acting on the Cantor set $\mathfrak{C}$. This also applies to express T as a group acting on $\mathfrak{C}$. It is easy to tweak the machinery and the definitions and define V as a subgroup of the self homeomorphisms of $\mathfrak{C}$. In spite of this, we will work a bit harder to give a different definition of V, and then work back to a structure similar to that of F and T. We do this because we wish to emphasize the strong relationship between Thompson's groups and self similar structures.

We will never give a strict definition of a self similar structure, but for now the Cantor set $\mathfrak{C}$ will serve as a typical example. It has two halves, each of which is "identical" to the whole. From this we will arrive at V.

13.1. Building V. Recall that the set $\{u\mathfrak{C} \mid u \in \{0,1\}^*\}$ is a basis of clopen sets for the topology on the Cantor set $\mathfrak{C}$, and that we refer to $u\mathfrak{C}$ with $u \in \{0,1\}^*$ as the cone at u. The following is mostly a translation of Lemma 7.2 and Corollary 7.2.1. Details are left to the reader.

LEMMA 13.1. *Let C be a cover of $\mathfrak{C}$ by cones. Then the maximal elements of C under containment form the unique minimal subcover of C, and this subcover consists of finitely many pairwise disjoint cones.*

A set $\{u_1\mathfrak{C}, u_2\mathfrak{C}, \ldots, u_k\mathfrak{C}\}$ is a cover of $\mathfrak{C}$ by pairwise disjoint cones if and only if $\{u_1, u_2, \ldots, u_k\}$ is a prefix set.

13.1.1. *The one-sided shift.* Consider the map $\sigma : \mathfrak{C} \to \mathfrak{C}$ acting on the right that deletes the first letter of each $w \in \mathfrak{C}$. Specifically $(0u)\sigma = (1u)\sigma = u$ for all $u \in \mathfrak{C}$. The map σ is two-to-one and is usually referred to as the *one-sided shift* on $\{0,1\}^\omega$. The adjective *full* is often applied since the domain is all of $\{0,1\}^\omega$ rather than some invariant subspace. We can also refer to σ as the *doubling map*. The map σ is interesting dynamically since any element of $\mathfrak{C}$ of the form u^ω, $u \in \{0,1\}^*$, (i.e., is periodic) has a finite orbit under σ, and the periodic elements of $\mathfrak{C}$ are dense in $\mathfrak{C}$.

After defining terms, we will state that the Thompson group V is the unique group of homeomorphisms from $\mathfrak{C}$ to itself that are determined by the local behaviors generated by σ. In the terminology (to be defined below) of symbolic dynamics, the Thompson group V is the topological full group of the full one-sided shift on $\{0,1\}^\omega$.

13.1.2. *Groupoids of germs.* In what follows, we only give the generality that we need.

DEFINITION 13.2. A *groupoid* is a category in which all morphisms are isomorphisms. A *groupoid of germs of a topological space X* will have the elements of X as objects and the morphisms from $x \in X$ to $y \in X$ will be a set of invertible germs of local homeomorphisms from x to y. An *invertible germ* from x to y is an equivalence class of homeomorphisms h with $xh = y$, each with domain some open set about x and codomain some open set about y. Two such homeomorphisms are equivalent if they agree on some open set about x. The groupoid $\mathcal{G}_X$ is the groupoid of all invertible germs on X. The groupoid $\mathcal{I}_X$ is the groupoid of all germs on X represented by identity homeomorphisms. If A is a set of morphisms in $\mathcal{G}_X$, then $\mathcal{G}_A$, called the *groupoid generated by A*, is the smallest category in $\mathcal{G}_X$ containing A.

LEMMA 13.3. *If X is a topological space and A is a set of morphisms in $\mathcal{G}_X$, then $\mathcal{G}_A$ is $\mathcal{I}_X$ together with all compositions of elements of A and their inverses.*

PROOF. The nominated collection of morphisms contains all identities, and is closed under composition and inversion. □

DEFINITION 13.4. A *local homeomorphism* on a topological space X is a continuous function $f : X \to X$ so that for each $x \in X$, the germ of f at x is an invertible germ. If S is a set of local homeomorphisms on X, then $\mathcal{G}_S$ is the groupoid generated by the set of germs of the elements of S. If S consists of one local homeomorphism f, we will write $\mathcal{G}_f$ for $\mathcal{G}_S$.

13.1.3. *Topological full groups.*

DEFINITION 13.5. If $\mathcal{G}$ is a groupoid of germs on a topological space X (i.e., $\mathcal{I}_X \subseteq \mathcal{G} \subseteq G_X$), then the *topological full group* of $\mathcal{G}$ is the set of those $h \in \mathrm{Homeo}(X)$ all of whose germs are in $\mathcal{G}$.

13.1.4. *A definition of V.*

DEFINITION 13.6. The group V is the topological full group of $\mathcal{G}_\sigma$ where $\sigma : \mathcal{C} \to \mathcal{C}$ is the one-sided full shift on on $\{0,1\}^\omega$ (the doubling map on $\mathcal{C}$).

We must now reveal some of the structure of V and relate it to the structures of F and T. In particular, we will work to prove the following parallel to Lemma 12.5 which will give a mechanism for multiplication. The proof will be in Section 13.1.6. The development contains no surprises.

We continue to use D and R for domain and range as we did for the group T.

PROPOSITION 13.7. *Every element of V can be represented by a binary tree (prefix set) pair (D, σ, R), every binary tree (prefix set) pair represents an element of V, and two binary tree (prefix set) pairs represent the same element of V if and only if they are related by the equivalence relation generated by matched binary splittings. Conversely, every binary tree (prefix set) pair (D, σ, R) with σ an arbitrary bijection represents an element of V.*

In the above statement, partitions of intervals are not mentioned as they are rarely used. Instead of binary partions of $I = [0,1]$ into dyadic closed intervals $[a,b]$, one would use binary partitions of $[0,1)$ into dyadic half open intervals $[a,b)$.

Before we start the process of proving Proposition 13.7, we state a straightforward corollary.

COROLLARY 13.7.1. *The Thompson groups F and T embed in V.*

13.1.5. *The germs of $\mathcal{G}_\sigma$.*

DEFINITION 13.8. Given u and w in $\{0,1\}^*$, the *rigid cone map* $\sigma_{uw} : u\mathfrak{C} \to w\mathfrak{C}$ is defined by $(u\alpha)\sigma_{uw} = w\alpha$ for all $\alpha \in \mathfrak{C}$.

The following lemma is immediate. The second provision follows from the fact that given $u \in \{0,1\}^*$, the cone $u\mathfrak{C}$ is the disjoint union (coproduct) of $(u0)\mathfrak{C}$ and $(u1)\mathfrak{C}$. This provision supplies the crucial fact that we need from the behavior of a coproduct.

LEMMA 13.9. *(I) Given u, w and s in $\{0,1\}^*$, the restriction of the rigid cone map $u\mathfrak{C} \to w\mathfrak{C}$ to $(us)\mathfrak{C}$ is the rigid cone map $(us)\mathfrak{C} \to (ws)\mathfrak{C}$.*

(II) Given a function $f : \mathfrak{C} \to \mathfrak{C}$, the restriction of f to $u\mathfrak{C}$ is the rigid cone map to some $v\mathfrak{C}$ if and only if for each $i \in \{0,1\}$, the restriction of f to $(ui)\mathfrak{C}$ is the rigid cone map to $(vi)\mathfrak{C}$.

The relevance of the cone maps is the following.

PROPOSITION 13.10. *The set of germs in the groupoid of germs generated by the full one sided shift σ on $\{0,1\}^\omega$ is exactly the set of germs represented by rigid cone maps $\sigma_{uv} : u\mathfrak{C} \to v\mathfrak{C}$ with u and v in $\{0,1\}^*$. Specifically, the cone map σ_{uv} represents a germ in $\mathcal{G}_\sigma$ at at each element of $u\mathcal{C}$, and a germ in $\mathcal{G}_\sigma$ at $w \in \mathfrak{C}$ is represented by a rigid cone map on some $u\mathfrak{C}$ for some prefix u of w.*

PROOF. Given $v \in \mathfrak{C}$ and a finite prefix of v in the form au with $a \in \{0,1\}$, then the germ of σ at v can be represented by the rigid cone map from $(au)\mathfrak{C}$ to $u\mathfrak{C}$. The inverse of such a germ is represented by a rigid cone map taking $u\mathfrak{C}$ to $(au)\mathfrak{C}$ for some $a \in \{0,1\}$. Thus a composition of k germs of σ and their inverses rigidly takes some $u\mathfrak{C}$ with the length of u at least k to some $u'\mathfrak{C}$ where u' is obtainable from u by k deletions and/or insertions of single symbols at the beginning of the word. Thus given a $v \in \mathfrak{C}$, a germ in $\mathcal{G}_\sigma$ at v is represented by some rigid cone map defined on $u\mathfrak{C}$ where u is a sufficiently long prefix of v.

Conversely, given any u and u' in $\{0,1\}^*$, we can build u' from u by the blunt technique of deleting all the symbols from u by germs of σ and then building u' by the properly chosen inverses of germs of σ. $\square$

13.1.6. *Proof of Proposition 13.7.*

PROOF. For $f \in V$ and $v \in \mathfrak{C}$, from Proposition 13.10 there is a prefix u of v so that $u\mathfrak{C}$ is the domain of a rigid cone map that represents the germ of f at v. Choosing such a domain for each $v \in \mathfrak{C}$ gives an open cover of $\mathfrak{C}$ and from Lemma 13.1 there is a finite subcover $\{u_1\mathfrak{C}, u_2\mathfrak{C}, \ldots, u_k\mathfrak{C}\}$ of pairwise disjoint cones. Since each is a domain of a rigid cone map that represents a germ of f we have that for $1 \leq i \leq k$, there is a w_i so that f takes $u_i\mathfrak{C}$ rigidly to $w_i\mathfrak{C}$, and further that $\{w_1\mathfrak{C}, w_2\mathfrak{C}, \ldots, w_k\mathfrak{C}\}$ is a cover of $\mathfrak{C}$ by pairwise disjoint cones.

The two sets $U = \{u_1, \ldots, u_k\}$ and $W = \{w_1, \ldots, w_k\}$ are both prefix sets by Lemma 13.1. However, the indexing does not necessarily reflect the prefix order. Thus f is represented by a prefix set pair (U, σ, W) where $u_i\sigma = w_i$ and the bijection σ does not necessarily carry the prefix order of U to that of W. Now there is a corresponding binary tree pair $(D, \sigma, R) = (T_U, \sigma, T_W)$ where we have overused the symbol σ. (The notation is from Lemma 8.8.)

The claim about the pairs that represent the same element is argued as in the first part of the proof of Proposition 6.17 with appropriate minor modifications. Finally the converse follows because the germs are seen to behave correctly from the first provision of Lemma 13.9. $\square$

Since the bijection σ in (D, σ, R) can be arbitrary, the typical way to represent an element graphically as a pair of trees is to number the leaves of D and R so as to show the action of σ. It is somewhat customary to number the leaves of D consecutively, but that is not necessary. It is also not necessary to use numbers. Two examples follow that represent the same element.

(13.1)

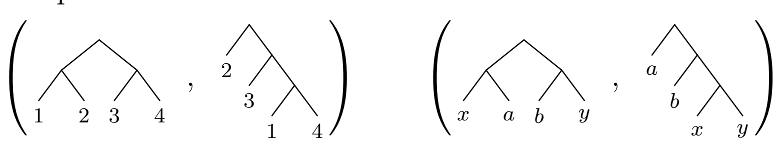

This example is fixed on an open subset of $\mathfrak{C}$.

The machinery of binary splittings and refinements, matched binary splittings and refinements is as in Section 6.2 for binary partitions and later sections for prefix sets, trees, and finally Section 12.2 for T. We copy Lemma 12.6 changing T to V.

LEMMA 13.11. *Every element of V has an irreducible representative (that is minimal as measured by numbers of carets) by binary tree pairs from which all other representatives are obtained by matched binary refinements. The only representatives of the identity of V are of the form (D, σ, D) where σ is the identity and the irreducible representative of the identity uses the trivial tree.*

13.1.7. *Multiplication and inverse.* As in F and T, multiplication is opportunistic. Details are left to the reader based on discussions in Sections 3.2, 6.2.3, and 12.2. The reader can work through the changes to show that the square of the element in (13.1) is as shown below.

(13.2)

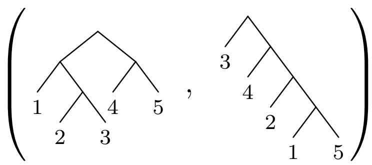

The inverse of (D, σ, R) is (R, σ^{-1}, D).

13.2. Transitivity, generators, and simplicity. Our goal in this section is to show that V is simple. The argument is rather easy. We leave more difficult properties to later sections and later chapters.

13.2.1. *Transitivity.* The next lemma has many consequences.

LEMMA 13.12. *Let* $A = \{u_1\mathfrak{C}, \ldots, u_k\mathfrak{C}\}$ *and* $B = \{v_1\mathfrak{C}, \ldots, v_k\mathfrak{C}\}$ *each be a set of pairwise disjoint cones that do not cover* $\mathfrak{C}$*. Then there is an element* $f \in V$ *so that for each* $i \in \{1, \ldots, k\}$*, the restriction of* f *to* $u_i\mathfrak{C}$ *is the rigid cone map to* $v_i\mathfrak{C}$*.*

PROOF. The lemma follows immediately if we can put A in A' and B in B' where A' and B' have the same number of elements and are each a cover of $\mathfrak{C}$ by pairwise disjoint cones.

The portion of $\mathfrak{C}$ not covered by A can be covered by finitely many pairwise disjoint cones. The same is true replacing A by B. Adding the appropriate cones to A and B results in two finite covers A' and B' that might have different numbers of elements. Assume A' has fewer elements than B'. Replacing some $u\mathfrak{C}$ in $A' - A$ by the two sets $(u0)\mathfrak{C}$ and $(u1)\mathfrak{C}$ raises the number of elements of A' by one. This can be repeated as needed to make the number of elements in the two covers the same while preserving A and B as subsets of the two covers. $\square$

The following parallel to Lemma 8.15 is an immediate consequence of Lemma 13.12. Note there can be no such parallel for the group T.

COROLLARY 13.12.1. *The group* V *is closed under deferment.*

13.2.2. *Generators.* In V an element in the form (P, β, P) will be called a *permutation.* A permutation (P, β, P) in V for which β fixes all but two elements of P will be called a *transposition.* If u and v are the two elements of P not fixed by β, then the pair (u, v) completely determines (P, β, P). We will use $\pi_{u,v}$ to denote this transposition.

A transposition (P, β, P) where P has more than two elements will be called a *proper transposition.* The following is immediate from

Lemma 13.12 and will be used in the argument for the simplicity of V. It is the reason for forcing on the proper transpositions.

LEMMA 13.13. *Any two proper transpositions are conjugate in V.*

The only non-proper transposition is $\pi_{0,1}$ which has no fixed points in $\mathfrak{C}$ and is not conjugate to any proper transposition.

Before we show that the proper transpositions generate V, we need the next observation.

LEMMA 13.14. *If P is a complete prefix set with more than one element, then for some $u \in \{0,1\}^*$, both $u0$ and $u1$ are in P.*

PROOF. The force of having more than one element is that P cannot contain the empty word. Letting v be a longest element in P, the argument proceeds exactly as in the proof of Lemma 7.2. □

PROPOSITION 13.15. *The proper transpositions generate V. The normal closure of any proper transposition is all of V.*

PROOF. The second sentence follows from the first and Lemma 13.13.

If P_r is the subgroup of V generated by proper transpositions, then we want to show that $V = P_r$. The only non-proper transposition $\pi_{0,1}$ has $\pi_{0,1} = \pi_{00,10}\pi_{01,11}$ and is thus in P_r. Permutations on a finite set are generated by transpositions, so all the permutations are in P_r.

To end the argument, we show that given an element f of V, we can multiply f on the right by permutations so that the result is the identity element. Representing f as (P, β, Q), we know that if P has only one element, then so does Q and that element is the empty word. Thus f is the identity element.

If P (and thus also Q) has more than one element, then P contains a pair of the form $u0$ and $u1$ and Q contains a pair of the form $v0$ and $v1$. By choosing the right permutation γ, we can insure that $\beta\gamma$ in $(P, \beta\gamma, Q) = (P\beta, Q)(Q, \gamma, Q)$ takes $u0$ to $v0$ and $u1$ to $v1$. Now $(P, \beta\gamma, Q)$ represents the same element as (P', δ, Q') where P' is obtained from P by replacing $u0$ and $u1$ by the single element u, Q' is obtained from Q by replacing $v0$ and $v1$ by the single element v and δ sends u to v and equals $\beta\gamma$ otherwise. Now P' has one fewer element than P, and this process can be repeated until the identity element is achieved. □

PROPOSITION 13.16. *The group V is generated by the four elements in $X = \{x_0, x_1, \pi_{0,1}, \pi_{10,11}\}$.*

Proof. From Proposition 12.18 the subgroup $\widehat{V}$ of V generated by X contains all of T. We must show that $\widehat{V}$ contains all proper transpositions. We separate the proper transpositions into several classes. If $\pi_{u,v}$ is a proper transposition, then we know that $u \perp v$ and we can assume that $\pi_{u,v}$ has been written so that $u < v$ under the prefix order. Our classes will depend on the locations of u and v in the complete binary tree $\mathcal{T}$.

The "left edge" of $\mathcal{T}$ are nodes of the form 0^n with $n > 0$. The "right edge" are nodes of the form 1^n with $n > 0$. All other nodes are "inner." We say that a transposition $\pi_{u,v}$ is inner if both u and v are inner. We don't have to say "proper" since an inner transposition is not $\pi_{0,1}$. By the transitivity properties of F and the fact that F is contained in $\widehat{V}$, we know that all inner transpositions are conjugate in $\widehat{V}$. Thus we need to show that $\widehat{V}$ contains at least one inner transposition, and that the inner transpositions with X generate all proper transpositions. We offer

$$((\pi_{10,11})^{x_0})^{\pi_{10,11}} = (\pi_{110,111})^{\pi_{10,11}} = \pi_{100,101}$$

as evidence that $\widehat{V}$ contains at least one inner transposition.

If the proper $\pi_{u,v}$ is not inner, then we say it is of class 2 if both u and v are on an edge of $\mathcal{T}$. All other proper transpositions that are not inner are of class 1. The depth of $\pi_{u,v}$ is the smaller of the depths of u and v.

For any transposiion $\pi_{u,v}$, we have $\pi_{u,v} = \pi_{u0,v0}\pi_{u1,v1}$. Let $\pi_{u',v'}$ be one of the factors. We note that $\pi_{u',v'}$ has depth at least 2 and cannot be of class 2, and that u' and v' end in the same symbol. So it suffices to show that all proper transpositions of this description are in $\widehat{V}$.

Let $\pi_{u,v}$ be of class 1 and depth at least 2 where u and v end in the same symbol. We will show that $\pi_{u,v}$ is conjugate to an inner transposition. Our conjugators will be from among $\pi_{0,1}$, $\pi_{10,11}$ and $\pi_{00,01} = (\pi_{10,11})^{\pi_{0,1}}$. The symmetry of our tools and the fact that $\pi_{u,v} = \pi_{v,u}$ allows us to assume that $u = 0^n$ for some $n \geq 2$. We know that v ends with 0, and because $u \perp v$ must be true, we know that v contains at least one 1.

If v starts with 1, then conjugating $\pi_{u,v}$ by $\pi_{00,01}$ produces an inner transposition. If v starts with 0, then conjugating $\pi_{u,v}$ by $\pi_{0,1}$ produces an inner transposition. This completes the proof. □

In fact V is finitely presented. The work involved is greater than for F and T and will be deferred to Section 27 where we use an action on a complex to get a very geometric presentation. A more direct computation is to be found in [43].

13.2.3. *The simplicity of V.*

THEOREM 13.17. *The group V is simple.*

PROOF. From Proposition 13.15, we must show that any non-identity element $f = (P, \beta, Q)$ has a proper transposition in its normal closure. Since f is not the identity, the empty word is in neither P nor Q.

We first argue that there is a cone $v\mathfrak{C}$ with the length of v greater than one on which f is a rigid cone map and which is disjoint from its image under f. There is some $u \in P$ with $u\beta \neq u$. If $u \perp u\beta$, then we take $v = u0$ to make sure v is as long as desired. Otherwise one of u and $u\beta$ is a proper prefix of the other. Different appeals to symmetry allow us to assume that not only is u a proper prefix of $u\beta$, but also that $u0$ is a prefix of $u\beta$. Now f takes $u1$ to $(u\beta)1$, $(u1)\mathfrak{C}$ is disjoint from $((u\beta)1)\mathfrak{C}$, and we set $v = u1$.

Let $w\mathfrak{C} = (v\mathfrak{C})f$. Now $g = [f, \pi_{v0,v1}] = \pi_{v0,v1}\pi_{w0,w1}$ is in the normal closure of f. And $[g, \pi_{v0,w0}] = \pi_{v0,w0}\pi_{v1,w1}$ is in the normal closure of g and thus of f. But by Lemma 13.9, $\pi_{v0,w0}\pi_{v1,w1} = \pi_{v,w}$, a transposition. Since the length of v is at least two, $\pi_{v,w}$ is a proper transposition. □

The proof of Theorem 13.17 works by showing that ultimately, every permutation in V is even. This can be compared to the variations of V in Section 38.8 where this does not hold.

13.3. Subgroups. For the following, we note that for a group G, a set X and a bijection $\beta : G \to X$, we get a right action of G on X induced from β in which $(\beta g)h = \beta(gh)$ for $g, h \in G$.

PROPOSITION 13.18. *The free group on two generators is isomorphic to a subgroup of V. Every countable, locally finite group is isomorphic to a subgroup of V.*

PROOF. Expressed as a group of homeomorphisms of $\mathrm{Homeo}(\mathfrak{C})$, the group T is a subgroup of V. So the first sentence follows from Proposition 12.13.

For the second sentence, let G be finite and $n = |G|$. Let T be a subtree of $\mathcal{T}$ rooted at $\emptyset$ with strictly more than n leaves and let β be an injection of G into $\Lambda(T)$. We note that the action of G on the image of β induces an embedding of G into V as a group of permutations (in the sense of Section 13.2.2) by declaring that G fixes the leaves of T that are not in the image of β.

Now assume that G is a subgroup of index i in the finite group H and let K be a left transversal for G in H containing the identity. Let V_{i-1} be the right vine of $i-1$ carets and i leaves and attach a copy of V_{i-1} to each leaf of T that is in the image of β. Let T' denote the resulting tree. Let γ be a bijection from K to the leaves of V_{i-1}. Note that the action of G within V permutes the attached copies of the V_{i-1}. We build a bijection $\hat{\gamma}$ from H to the leaves of the n attached copies of V_{i-1} in T' by setting for each $k \in K$ and $g \in G$, $\hat{\gamma}(kg) = (\beta g)(\gamma k)$ viewing the vertices of $\mathcal{T}$ as elements of the monoid $\mathfrak{M}$. Now the bijection $\hat{\gamma}$ induces an action of H on the leaves of the attached copies of V_{i-1} which, in turn, induces an embedding of H into V by fixing the leaves of T' that are not in the image of $\hat{\gamma}$. These are just the leaves of T not in the image of β.

The action of G on the leaves $(\beta g)(\gamma 1)$ is conjugate to the original action of G induced by β. This uses the fact that the image of β is not all the leaves of T. Since the action of H before the conjugation was induced by a bijection $\hat{\gamma}$ from H to some of the leaves of T', we can insist that after the conjugation the action of H is induced by some bijection from H to some but not all the leaves of a tree T''. Since any countable, locally finite group is a countable union of an ascending nested sequence of finite subgroups, this allows us to embed any countable, locally finite group into V. □

Remark 13.19. Both T and V can be exhibited as a group of fractions of a positive monoid. The resulting structures are slightly less well behaved than for F. As of this time, there is no plan to discuss this in the book.

14. End notes

The standard reference since 1996 for the basics on F, T and V has been Cannon-Floyd-Parry [43]. Notes on F by Burillo [39] are available online. There are two nice introductory books with chapters on F. One is Meier 2008 [153] (Chapter 10 with an appropriate epigraph), and the other is Bonanome-Dean-Dean 2018 [23] (Chapter 2). The 1992 paper [173] is a survey of the finitely presented, infinite, simple groups known as of that date and covers the generalizations of V by Higman [107] and parallel generalizations of T from Brown 1987 [34].

Most of the material in this chapter comes under the heading of well known, and shows up in several places. Initial knowledge of the Thompson groups comes from a set of notes handwritten by Thompson some time after 1973 and very widely copied and circulated. The lower limit on the date is apparent from the fact that the notes refer to the published version of [152]. Much material was adapted from [43] which expanded on and added to the material in Thompson's notes, and from [34] and [36]. The piecewise projective version of F is attributed to Thurston. The universal property of Section 10.2 has multiple sources which are discussed in Section 19. Remark 12.14 is from [179]. The derivation of a finite presentation for T in Section 12.8 is simpler than ones found in the literature and I learned of it from Jim Belk. I learned the simple proof of the simplicity of V from Matti Rubin. The second provision of Proposition 13.18 is from Higman [107].

It is mentioned in the preface that Thompson's groups tend to be "none of the above." They are not linear because they are not residually finite. They are not hyperbolic because they contain infinitely generated free abelian subgroups. They are not Kähler by [156] and [157]. They are not right angled Artin groups (RAAGs) since RAAGs embed in linear groups. They are not groups that satisfy the Tits alternative which includes mapping class groups of hyperbolic surfaces of finite type, certain general Artin groups, and outer automorphism groups of free groups. Groups in the Thompson family do show up as subgroups of big mapping class groups (BMCGs), those mapping class groups of surfaces of infinite type.

See Guba-Sapir 1997 [99] and Genevois 2025 [78] for information on diagram groups. Diagram groups overlap with Thompson's groups (in particular, both contain F) but neither class contains the other. For groups such as F, the diagrams are basically dual to tree pairs when stripped of certain labels, and (with some work) facts about F derived with the use of diagram group diagrams can be derived in other ways. An example of this is our treatment of the word length function for F in Section 44.

Exercises such as seen in Theorem 4.11 to resolve behavior on different orbitals lie behind the main result of Brin 1999 [28], which is reworked from a more universal viewpoint in Hyde-Moore 2023 [115], and they appear in Bleak 2008–9 [18, 19, 20]

Chapter 3 discusses the mathematical structures from which the Thompson groups arise. Section 15 of that chapter gives some of the early history of the groups. Since the groups were discovered several times, the history has some wrinkles.

The definitions of F, T and V and the mechanics for the manipulation of their elements are clearly subject to modification. There is an entire world of modifications of the three original groups, and some of this (all would be impossible) is discussed in Chapter 6. Subgroups of F, T and V can be surprisingly varied, and at the same time there are surprising restrictions. Subgroups are not covered in this edition.

The high level of transitivity of the groups makes their algebra very expressive. To be more precise, the first order theory of the groups is powerful and complex. This is discussed in Chapter 5.

The rich combinatorial structure of both the groups and the machinery that goes with them, along with their structures as groups of fractions allows the construction of simplicial complexes with good properties on which the groups act. Some of these complexes, the actions and their consequences are covered in Chapter 4. That chapter also covers a complex that arises from the status of F as an initial object in the category of groups with a conjugacy idempotent.

That the word problem is solvable in F, T and V is very straighforward. Less straightforward is the conjugacy problem and other algorithmic problems. This will be covered in a planned later chapter.

Subgroups are inclusions and thus a special type of morphism. Isomorphisms between the variants of Chapter 6 can be subtle, and the automorphisms of the groups themselves just as subtle. Necessary conditions for isomorphisms between variants of V are covered in Chapter 6, and are shown to be sufficient in Chapter 7. Other morphisms among the Thompson groups will be covered in a planned later chapter.

The heavy use of words opens the door to applications of automata and there are strong connections there. This topic will be included in planned later chapters.

The interaction between the Thompson groups and self similar structures will show up regularly in later chapters. The key idea is that the behavior of the groups is determined by their action on pieces of a space that resemble the entire space. This idea lies behind the material on the origins of Thompson's groups in Sections 16, 17, and 18 of Chapter 3. This is discussed further in the end notes (Section 21) to that chapter. Algebraic structures that abstract the local actions on self similar spaces interact strongly with Thompson's groups. Some of this interaction is introduced in Chapter 7 where it is shown that the representations of the Thompson groups into these structures proves useful in classifying a family of variants of V.

The simplicity of the groups (or as will be seen later, simplicity of subgroups in the derived series) and generalizatons of finite presentability are two constant themes of groups in the Thompson family. Simplicity of some of the variants is discussed in Chapter 6. Generalizations of finite presentability are discussed in Chapter 4.

The piecewise projective representation of Thompson's groups have proven useful in building important examples. See Lodha-Moore 2016 [139].

The many flavors of the Thompson groups and the similarity among them of some properties motivates a search for generalizations. We will not discuss these, but we mention a few here. To some extent, diagram groups [99, 78] are a generalization. An approach using semigroups due to Lawson and at times with coauthors perhaps starts with 2007 [131] and extends to Lawson-Vdovina 2024 [134]. There are strong connections with these structures and those considered in Sections 16 and 40. See the end notes in those chapters. For another approach, see Martínez-Pérez-Nucinkis in 2013 [150], and together with Matucci in 2016 [149]. A categorical approach is in Thumann 2017 [191]. Locally symmetric spaces are used in Hughes 2009 [113] with an appendix by Farley, and several succeeding papers by the two of them. The connection of V and other groups to shifts as described in Paragraphs 13.1.1 open the possibility to generalizations and unifications from dynamical systems and related algebraic structures. These may be sorted out in later editions.

CHAPTER 3

Origins of Thompson's groups

Chapter contents

[1]The Thompson groups have been discovered several times for several reasons. We give details in this chapter. Hindsight gives alternate ways that the groups might have been discovered, and we include one because of its generality and because it adds weight to one of the approaches. Relating these stories allows us to exhibit some mathematical reasons for the existence of the groups.

The first section of this chapter is both an introduction to the rest of the chapter and a brief history of the early literature of the groups, from Thompson's first introduction in the late 1960s to the expository 1996 paper [43] of Cannon-Floyd-Parry. Only a few dozen papers relating to Thompson's groups appeared in that period. Since [43], there have been hundreds.

The remaining sections cover a monoid of Thompson that preceeded the groups (Section 16), a section of hindsight on self similar objects in a category (Section 17), an algebraic structure whose automorphism group is V (Section 18), the intimate connection beteween F and homotopy idempotents (Section 19), and the groups as structure groups of common algebraic laws (Section 20).

15. Introduction and some early history

15.1. Late 1960s. The Thompson groups F and V first appeared in print in McKenzie-Thompson 1973 [152]. The paper was presented at a September, 1969, conference in Irvine on decision problems in groups. The groups F and V were not the main point of the paper, and some of their basic properties were used to build an example, more elementary than previous examples, of a finitely presented group with unsolvable word problem. The groups F (denoted $\mathfrak{P}$) and V (denoted $\mathfrak{C}'$) are mentioned in the concluding remarks of [152] where it is stated that V is infinite, finitely presented and simple. These facts about V were announced by Thompson at the conference.

Earlier in 1969, Thompson had described the structure of V and its properties to Fred Galvin, a vistor at Berkeley. Galvin, a student of Bjarni Jónsson, was familiar with a variety of algebras $\mathfrak{A}_{1,2}$ from Jónsson-Tarski 1961 [119] which was universal for the property of having its free algebra on one generator also free on two generators. Galvin suggested that V was the automorphism group of this algebra which Thompson then verified. In May of 1969 Galvin wrote about this to Jónsson who forwarded the information to Graham Higman.

[1]This chapter is reasonably complete. Minor changes may be made in the future.

Thompson's discovery came from an attempt to apply facts about a monoid containing V to algebraic logic. Modifying a formal statement can alter whether the statement is true, motivating a sytematic study of how formal expressions are modified. As an example, the focus of Jónsson 1962 [117] is on substitution of variables, and generators and relations are found for the monoid of finitely supported endofunctions on a set S (those $f : S \to S$ for which $\{x \in S \mid xf \neq x\}$ is finite), where elements of S are interpreted as variables, and the endofunctions as substitutions. The analysis is used in [117] to simplify the proof of a result in cylindric and polyadic algebras. Further analysis of this monoid appears in Thompson's 1979 thesis [190].

Separate from his thesis, Thompson had considered a monoid $\mathscr{M}$ of transformations of formal expressions that also take into account the movement of parentheses. This monoid never saw application to algebraic logic, but it led directly to the groups we call Thompson groups. The group V is the group of invertible elements of $\mathscr{M}$.

Thompson found a finite presentation for $\mathscr{M}$, and observed that the group of invertible elements was infinite and also finitely presented. Thompson asked John Rhodes, the closest algebraist he knew, what to do with it, and Rhodes advised him to look for normal subgroups.[2] Proving that there were none, Thompson had built the first example of an infinite, simple, finitely presented group.

Thompson never published his results about $\mathscr{M}$, but did describe its structure and the proof of its finite presentability in talks at Palo Alto in 2004 and Luminy in 2008. We give the details in Section 16.

Thompson surmised (as he put it "On reading Hewitt and Ross") that F might be a counterexample to what is oftern referred to as the von Neumann-Day conjecture, that a finitely generated group with no non-abelian free subgroups is amenable. News of this question did not spread rapidly until the 1980s. See Section 15.4.

15.2. Early to mid 1970s. The introductory paragraphs to Boone and Higman 1974 [25] mention that in his presentation of [152] at the Irvine conference, Thompson pointed out that all subgroups of a finitely presented simple group have a solvable word problem. While easy to demonstrate, this was not generally known, and the information led Boone and Higman to prove in [25] that a finitely generated group G has a solvable word problem if and only if it embeds in a simple subgroup S of a finitely presented group P, and further to ask in (4) of Boone 1974 [24] whether S and P could be made the same. An intermediate question was also raised in [24] as to whether S could be

[2]Both Thompson and Rhodes tell this story.

made finitely generated. Earlier sources for Thompson's observation were noted in a late addition to the end of [25].

Thompson answered the intermediate question in his only other published paper on the Thompson groups 1980 [189] (announced 1976 at a conference in Oxford) by showing that the simple group S could be made finitely generated. This was done by having an arbitrary group G given as an action on a set L and a modified version of V acting on the Cantor set $\mathfrak{C}$ act together on $L \times \mathfrak{C}$, and then looking at the group that they generate. For G countable with solvable word problem, enlarging G and shrinking the combination preserves the solvability of the word problem and forces simplicity and finite generation of the combination. The Higman embedding theorem then embeds the result in a finitely presented group. A modification of the techniques in [189] was picked up later and carried farther by E. A. Scott.

Progress on the full question (now known as the Boone-Higman conjecture) has focused on finding more finitely presented, infinite, simple groups, understanding their subgroups, and putting classes of groups known to have solvable word problems into them. All infinite, finitely presented, simple groups discovered so far have been closely related to Thompson's groups with the exception of the groups of Burger-Mozes 2000 [37] and Caprace-Remy 2009 [45]. We will not cover these exceptions. The groups of [37] and [45] are somewhat related to each other, and are quite different from groups in the Thompson family.

Building on the news from Jónsson, Higman 1974 [107] generalized the Jónsson-Tarski variety $\mathfrak{A}_{1,2}$ to a family $\mathfrak{A}_{1,n}$ where the free algebra $J_{n,r}$, $1 \le r < n$, on r variables in $\mathfrak{A}_{1,n}$ is isomorphic to exactly the algebras in $\mathfrak{A}_{1,n}$ that are free on $r + k(n-1)$ variables as k ranges over $\mathbf{N}$. Higman then showed that all the $V_{n,r} = \mathrm{Aut}(J_{n,r})$ are finitely presented, are simple if n is even, and have a simple subgroup of index two if n is odd. Further Higman partially classified the $V_{n,r}$ to the extent of showing that $n \neq m$ implied that $V_{n,r}$ and $V_{m,s}$ cannot be isomorphic, giving infinitely many isomorphism classes of infinite, finitely presented, simple groups. Higman also showed, Proposition 13.18, that every countable, locally finite group embeds into $V = V_{2,1}$. By contrast, Higman shows that $GL(3, \mathbf{Z})$ does not embed in any $V_{n,r}$, noting that $GL(3, \mathbf{Z})$ has a solvable word problem. Much of this material is covered in Section 38.8.

The Jónsson-Tarski paper [119] and its earlier announcement 1956 [118] appeared during a flurry of interest in the notion of independent generating sets and possible variations from the predictable behavior of rank in finite dimensional vector spaces. Varieties more general than the $\mathfrak{A}_{1,n}$ of Higman were independently constructed and explored by others in the late 1950s and early 1960s. The paper Świerczkowski 1961/62 [185] shows that in any variety of algebras, if the free algebra on k generators is isomorphic to the free algebra on $l > k$ generators, with k the smallest for which this holds, and with $l = k+d$ the smallest for which this holds for that k, then for any $n > m$, the free algebras on m and n generators are isomorphic if and only if $m \geq k$ and $n = m+jd$ for some $j \geq 0$. Further Świerczkowski shows in [185] that for each $0 < k < l$ there is a variety $\mathfrak{A}_{k,l}$ exhibiting exactly that behavior. It is also shown in [185] that the varieties $\mathfrak{A}_{k,l}$ are in a sense universal for the stated properties.

This setting produced an independent discovery of a Thompson group. The paper Smirnov 1974 [179] considers the automorphism group of $J_{n,1}$, the free algebra on one generator in the variety $\mathfrak{A}_{1,n}$ of [185]. Generators are found for the group, and it is shown that the group is generated by permutations analogous to the permutations of Section 13.2.2. Smirnov also proves that the groups contain an isomorphic copiy of $PSL(2, \mathbf{Z})$, and that their subgroups are closed under taking finite products. In an end note added prior to publication, Smirnov states that, after obtaining the results in [179], he learned of Higman's and Thompson's groups, having received notes of lectures given by Higman in Australia in August of 1973.

No one has found a practical way to analyze the automorphism group of the free algebra on k generators in the variety $\mathfrak{A}_{k,l}$ of [185] when k is greater than one.

We give the details of the simplest variety $\mathfrak{A}_{1,2}$ of Jónsson and Tarski in Section 18 and show that the automorphism group of its free algebra on one generator is isomorphic to V. But before that, we show how V appears as a subgroup of automorphism groups in a more general setting that is outside the history of Thompson's groups. We explain the connection.

If an algebra A is free on $\{a\}$ and is also free on $\{b, c\}$ with $b \neq c$, then with A_x denoting the subalgebra generated by x, we have A isomorphic to both A_b and A_c and further A is the coproduct of A_b and A_c. Speaking loosely, A is the coproduct of two copies of itself. The action of V from Section 13 on the Cantor set $\mathfrak{C}$ essentially takes advantage of the fact that $\mathfrak{C}$ is the coproduct of two copies of itself. In Section 17 we show that in any category, if an object X is a coproduct (or product) of two objects each isomorphic to X, then there is a canonical homomorphism from V into $\mathrm{Aut}(X)$. Since V is simple, this homomorphism is either trivial or an embedding.

The copy of V in $\mathrm{Aut}(X)$ can be well hidden. In the category of sets, the set of integers is the coproduct of the set of odd integers and the set of even integers, making V a subgroup of the group $\Sigma(\mathbf{Z})$ of all permutations of the integers. But $\Sigma(\mathbf{Z})$ contains an isomorphic copy of every countable group, so this is hardly a surprise. However this shows both that the Thompson groups can appear frequently, and that the Jónsson-Tarski example is special in that it gives an example where the automorphism group is exactly V.

15.3. Late 1970s. The third and fourth discoveries of a Thompson group were almost simultaneous, and were for almost the same reasons which were quite different from the first two discoveries. The setting is of homotopy classes of basepoint preserving maps of pointed, connected CW complexes. We use $\sim$ to denote "homotopic to," and to save typing we will not distinguish between a loop at a basepoint and the element of the fundamental group it represents.

If a CW complex X with basepoint p has a self map f that is homotopic to f^2, then f is a *homotopy idempotent*. But whether or not that homotopy keeps p fixed during the course of the homotopy is an important distinction. If p is kept fixed, then f is called a *pointed homotopy idempotent* and it can be shown that there is a CW complex Y, maps $d : X \to Y$ and $u : Y \to X$ so that (composing right-to-left) we have $f \sim ud$ and $du \sim \mathbf{1}_Y$. In this situation, it is said that the homotopy idempotent f *splits*.

On the other hand, if the homotopy drags p along a loop α, then f is called an *unpointed homotopy idempotent*, and passing to the fundamental group, we have that $f_* = C_\alpha \circ f_*^2$ on $\pi_1(X, p)$ where C_α is conjugation by α. In the terminology of Section 10.2, the triple $(\pi_1(X, p), f_*, \alpha)$ is a conjugacy idempotent, and since (F, σ, x_0) is an initial object for groups with conjugacy idempotent, there is a canonical homomorphism from F to $\pi_1(X, p)$ taking x_0 to α and x_1 to $f_*(\alpha)$. In this situation, f splits if and only if this homomorphism is not injective.

In the early 1970s, it was known that pointed homotopy idempotents split, and nothing was known about unpointed homotopy idempotents. This question was related to other questions in shape and homotopy theory. See Geoghegan 1978 [79] and the end of Hastings-Heller 1981 [102]. In the second half of the 1970s, the group F was discovered twice by realizing that in the situation above, if α and $f_*(\alpha)$ do not commute, then they have to generate a subgroup of $\pi_1(X, p)$ whose presentation is the infinite presentation (9.1) of F. This was done in Dydak 1977 [63] where the author credits Minc with finding a faithful represenation for the presentation. It was also done in Freyd-Heller 1993 [73] which had circulated as a prepreprint since the late 1970s. In Section 19 we give the details showing that the canonical homomorphism from F discussed above characterizes when unpointed homotopy idemptotents split. The idempotent splits if and only if the canonical homomorphism has non-trivial kernel.

15.4. Early 1980s. At a geometric topology converence in Warsaw in the summer of 1978, the pairs that discovered F via homotopy idempotents, Dydak-Minc and Freyd-Heller, were represented, respectively, by Dydak and by Hastings, a student of Heller. During an afternoon meeting with the two, Ross Geoghegan learned of the details of the group F. At the time, Geoghegan was interested in the application of the techniques of shape theory to groups via the ends of universal covers of their classifying spaces, and he recognized the potential of F to exhibit interesting behavior.

Geoghan conjectured, independently of Thompson, that F might supply a counterexample to the von Neumann-Day conjecture. Secondly, he conjectured that F would answer question F11 of the problem list of the proceedings [193] of the 1977 Durham conference on homological group theory. This asked whether there could exist a torsion free group with property FP_∞ (see Section 22.1 for the definition and the related F_n and F_∞) having a free abelian subgroup of infinite rank. Lastly, Geoghegan conjectured that F would have trivial cohomology at infinity. Specifically $H^n(F, \mathbf{Z}F) = 0$ for all n. Over the next few years many people, including the author of this book, learned of Thompson's groups from Geoghegan.

During the Warsaw conference Dydak and Hastings used the properties of F to show that homotopy idempotents on 2-dimensional complexes must split, and the paper [64] appeared in the 1980 proceedings of the conference. Hastings and Heller then proved, Theorem 28.9, that homotopy idempotents on any finite dimensional complex must split. The proof used the fact that F must be a subgroup of a counterexample Y together with a spectral sequence argument to give facts about the homology of Y. The result appeared in an early 1981 summary [102] in the proceedings of a Dubrovnik conference that year on shape theory and geometric topology, and also in 1982 [103].

The question from [193] was settled by Brown-Geoghegan 1984 [36] where they show that F is of type F_∞ by building a classifying space for F based on its status as an initial object in the category of groups with a conjugacy idempotent that has exactly two cells in each dimension.

They derive from that a more elementary proof that homotopy idempotents on finite dimensional complexes split, and also the fact that $H_n(F, \mathbf{Z}) = \mathbf{Z} \oplus \mathbf{Z}$ for all $n \geq 1$. These results are covered in Section 28. Also in [36] is a proof that $H^n(F, \mathbf{Z}F) = 0$ for all n. The paper was also the first among those working on the homotopy idempotents queston to show awareness of Thompson.

We note that homology calculations for Thompson's groups have been difficult to obtain and slow to arrive. There will be a few more homology results mentioned in this early history. However, there have been recent results, starting in 2019 with [186], using techniques from homotopy theory.

The von Neumann-Day question first requires showing, Theorem 4.11, that F contains no non-abelian free subgroups. This was shown independently in Brin-Squier 1985 [31], Freyd-Heller 1993 (but written much earlier) [73], and Cannon-Floyd-Parry 1996 [43]. The question of whether F is amenable remains open, and the question became well known from talks given by Geoghegan.

In unpublished notes written after the 1976 Oxford conference, Higman suggested modifications to Thompson's method in [189] for combining families of groups with V, and further suggested several goals. The goals included refining the method further to create larger classes of infinite, finitely presented, simple groups, finding groups in these classes that contained subgroups isomorphic to $GL(n, \mathbf{Z})$ for $n \geq 3$, and finding groups in the classes that had some undecidable properties. For the last, Higman suggested an undecidable conjugacy problem.

In a series of three papers appearing in the same 1984 volume of the Journal of Algebra, E. A. Scott fulfilled all three suggestions. The first [174] gives a general construction for building infinite, finitely presented, simple groups. The construction differs from that of Thompson [189] in that it restricts the class of groups it combines with V, but has each such group act on the Cantor set simultaneously with V, rather than act on enlargement of the Cantor set. Enough connection to [189] exists for many lemmas from [189] to apply. In [176], the technique is applied to several groups including the $GL(n, \mathbf{Z})$, the semidirect products $\mathbf{Z}^n \ltimes GL(n, \mathbf{Z})$, the groups $\mathbf{Z}[\frac{1}{n}]$, for fixed $n \geq 2$, and certain other abelian groups to put them each in a finitely, presented group. In [175] a group of Miller [154] with an unsolvable conjugacy problem is used to build a finitely presented, simple group with the same property.

15.5. Mid to late 1980s. Some important properties of F come from its status as a group of PL homeomorphism of $\mathbf{R}$. The paper [31] is primarily about $PL_+(\mathbf{R})$. Bieri and Strebel learned of the group F

and [31] from Geoghegan and wrote an extensive set of notes 1985 [13] that included general results on presentations and finiteness properties of subgroups of $PL_+(\mathbf{R})$, as well as the first results on automorphisms of these groups. The groups $G(K, A, P)$ considered were determined by a finitely generated multiplicative group P of allowed slopes, and a P-invariant ring A of allowed breakpoints, all on a closed (possibly unbounded) interval K in $\mathbf{R}$. The group F equals $G([0,1], \mathbf{Z}[\frac{1}{2}], \langle 2\rangle)$. The notes were amended and published in book form in 2016 [15].

Brown 1987 [34] derives a general technique for proving property F_n for a group and applies this to show the $V_{n,r}$ of Higman and corresponding generalizations of F and T are all of type F_∞. Some of this is covered in Sections 25 and 26. The paper [34] established a general outline for proving facts about the properties F_n and F_∞. This outline is discussed briefly in Section 27.3. The paper also analyzes the normal subgroups of these generalizations, showing that the commutator subgroups of the variants of F and V and the second commutator subgroups of the variants of T are all simple. The paper [34] along with Higman's notes [107] served as the earliest reference for facts about the Thompson groups.

Ghys and Sergiescu learned of the Thompson group family from Geoghegan. Their paper 1987 [82] derives a number of results about the groups, mostly about T (denoted G in [82]) and the dynamics of its actions on the circle $S^1 = \mathbf{R}/\mathbf{Z}$. Among the results are that every element of T has rational rotation number, and (Lemma 12.10) every rational in $(0, 1)$ is a rotation number of some element of T. Also shown is that T is conjugate to a smooth action on S^1 and facts about the stablity and invariants of the various C^r actions are given. The cohomology ring structure is described for F, F', T and $\widetilde{T}$ where the latter is the lift of T to $\mathbf{R}$. See [82] for more.

15.6. Early to mid 1990s. The braid group B_n on n strands includes naturally into B_{n+1} by adding an unbraided strand at the right, and the homology of the direct limit B_∞ of the B_n is understood. The analysis of the homology of F' in [82] and the known homology of B_∞ hinted to Greenberg and Sergiescu that there might be an extension of F' by B_∞ that is acyclic. In 1991 [96], they construct such an extension with the predicted property. Sergiescu also promoted Thompson's groups widely, and Greenberg put out a rapid series of papers [90, 89, 94, 93, 91, 92, 95] on the geometry, combinatorics, and projective aspects of Thompson and related groups before his untimely death in 1993.

Stein, a student of Brown, in 1992 [183] worked with variants of the Thompson groups of the form $G(K, A, P)$ studied in [13] in which the group P of allowable slopes has rank greater than one. In [183], Stein builds complexes similar to those in [34] using more sophisticated collapsing arguments that include some of the type used in [36]. With these complexes, property F_∞ is proven for many of the groups, simplicity results were obtained paralleling those of [34] for corresponding variants of T and V, presentations are obtained for some, and abelianizations are calculated for the non-simple variants. More intricate are homology calculations done for the F variants where the slope group has rank no greater than two.

A 1992 survey paper [173] by E. A. Scott in the proceedings of a 1989 conference in Berkeley gave a survey of the known finitely presented, infinite, simple groups including those from [107], [174], [176], [175], and [34].

In the same proceedings the paper [35] by Brown uses the complexes from [34] with the collapsings from [183] to present V as an amalgamation of three finite groups. Brown also uses the complex to prove that V is rationally acyclic and asks if V might be integrally acyclic.

In the 1995 paper [47] Cleary considers subgroups of $PL_+(\mathbf{R})$ with irrational slopes. In the notation of [13], groups such as $G = G(K, A, P)$ are considered with P the cyclic, multiplicative subgroup of $\mathbf{R}$ generated by an irrational α, A, the set of breakpoints of elements of G, a P-invariant $\mathbf{Z}P$ module in $\mathbf{R}$, and K an interval with endpoints in A. The irrationality of α allows for multiple choices of how to subdivide intervals in the construction of elements of G. A specific example with $\alpha = \sqrt{2} + 1$ is studied. With α satisfying $\alpha^{-2} + 2\alpha^{-1} = 1$, an interval can be divided into three intervals of unequal length and still be used to create PL homeomorphisms with slopes that are powers of α. Because of this, there are several choices as to how to subdivide a given interval. It is shown in [47] that the patterns created allow a complex along the lines of [34] and [183] to be built to prove that $G([0, 1], \mathbf{Z}[\sqrt{2}], \langle\alpha\rangle)$ has type F_∞. Indications are given to show how the analysis applies to certain other algebraic numbers α.

Fordham's 1995 thesis [71] starts the geometric studies of the group F by calculating the word length metric for F. The thesis circulated widely before the published version [72] appeared in 2003.

In Brin 1996 [27] the automorphism groups of F and T are completely analyzed. In particular, it is shown that all elements of an index 2 subgroup of $\mathrm{Aut}(F)$ come from conjugations by elements of $PL_+([0, 1])$, and that the outer automorphism group of T has order 2.

The only writing by Thompson on the Thompson groups beyond what has been mentioned above was a set of handwritten notes written some time after the 1973 apperance of [152]. These were copied, recopied and passed around many times until the appearance of Cannon-Floyd-Parry 1996 [43]. The paper [43] fleshed out Thompson's notes by adding numerous missing details and extra results. The paper also includes a projective version by Thurston that generalizes to higher dimensions. Since its appearance, [43] has been the standard introduction to Thompson's groups.

15.7. The last (known) independent discovery. A discovery by Dehornoy of F and V was for reasons that are somewhat closer to Thompson's original discovery. The LD or *left (self) distributive* identity $x\cdot(y\cdot z) \sim (x\cdot y)\cdot(x\cdot z)$ had been known to be of importance in the study of large cardinals in set theory. See the introduction to Dehornoy 1989 [53]. Conjugation in groups (with a reversal of order) satisfies this identity since $(z^y)^x = (z^x)^{(y^x)}$, but is not the freest example. Questions in free algebras defined by LD are quite difficult to answer, and in [53], Dehornoy studied a monoid of expression modifications generated by the LD identity.

In Dehornoy 1993 [54] the approach of [53] was generalized to build monoids for certain algebraic identites which could then be embedded in a group of fractions. This was applied to both the associative law and the commutative law. In each case an infinite presentation is extracted that is explicit in the case of the commutative law, and implicit in the case of the associative law. In the latter case, a normal form and a solution of the word problem is derived. The group obtained from the associative law is considered again in Dehornoy 1996 [55], and also in the Appendix to Chapter IX of the book Dehornoy 2000 [56] where it is noted that this group is isomorphic to the Thompson group F. Motivated by this material, Section 20 gives an argument to justify that F is "the structure group of the associative law." The structure group combining the associative and commutative laws is V.

16. The Thompson monoid

We describe Thompson's monoid $\mathscr{M}$ and give the argument that $\mathscr{M}$ is finitely presented. We do so for historical interest, because the technique for showing that $\mathscr{M}$ has a certain presentation is different from arguments that we have used before, and because the move from group to monoid introduces ideas that will be repeated later. See the end notes of this chapter.

We work with two representations of $\mathscr{M}$ and do not distinguish between them. One acts as modifications of formal expressions, and the other acts as endomorphism of the Cantor set $\mathfrak{C}$. The identification of the two represenations factors through functions between leaf sets of finite, binary trees. While trees tie the two representations together and help with certain arguments, we will not use them for calculations. In spite of the fact that a monoid of expression modifications and a monoid of endofunctions of $\mathfrak{C}$ are different monoids, we will recklessly use $\mathscr{M}$ to denote both.

16.1. Expressions and their modifications. We assume an infinite supply of variable symbols. However, we reserve the letter e to represent an unknown complex expression of one or more variables. We will use e with subscripts or primes such as e_0 or e'' to denote separate complex expressions rather than use different letters.

DEFINITION 16.1. Recursively, a *fully parenthesized expression* is either a single variable, or an expression (e_1e_2) where e_1 and e_2 are fully parenthesized expressions. We usually omit "fully parenthesized," and also omit the outer matched parentheses in expressions having more than one variable. For example, we will write $(ab)(cd)$, ab, and $a(bc)$ instead of $((ab)(cd))$, (ab), and $(a(bc))$.

An *expression pair* is a pair (e_0, e_1) usually denoted $e_0 \to e_1$ of fully parenthesized expressions that satisfy the following. No variable appears twice in e_0. Every variable that appears in e_1 also appears in e_0.

The monoid $\mathscr{M}$ is built on expression pairs, and we give five important elements of $\mathscr{M}$ which illustrate the points in definition above.

$$\begin{aligned} &P : ab \to ba && R : a(bc) \to (ab)c \\ &U : a \to aa && K : ab \to a \\ &L : ab \to b && \end{aligned} \tag{16.1}$$

In particular, the two expressions in a pair do not have to have the same number of positions, not all the variables used in the first expression need show up in the second expression, and a variable may show up more than once in the second expression.

Variables are regarded as placeholders, and the second expression is the modification of the first expression. There are no needs beyond the fact that the modification needs to be able to specify unambiguously each position of the first expression, and be able to say where the information in each position in the second position comes from.

Different expression pairs can represent the same element of $\mathscr{M}$. Replacing all appearances of one variable in both expressions by a fixed expression (which could simply be another single variable) gives another representation of the same element of $\mathscr{M}$ as long as no variable ends up duplicated in the first expression. Thus $L : (pq)((rs)t) \to (rs)t$ and $U : (xy) \to (xy)(xy)$ are alternate representations of L and U. Consequently, a bijective substitution of variables also does not change an element of $\mathscr{M}$. So $U : b \to bb$ and $U : z \to zz$ also represent U.

16.2. Labeled trees and the Cantor set.

Definition 16.2. If T is a finite binary tree, then a *labeling* of T is a function $\lambda : \Lambda(T) \to A$ from the leaves of T to a set of variables A.

The *expression given by* λ is the label of the root of T defined by recursively extending λ to all the nodes of T as follows. If $\{u, u0, u1\}$ is a caret in T then $\lambda(u) = (\lambda(u0)\lambda(u1))$. An easy induction shows that every fully parenthesized expression is given by some labeling of a tree.

Expressions used in (16.1) correspond to labeled trees as follows.

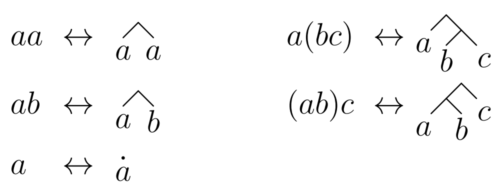

This lets us write the elements of $\mathscr{M}$ given in (16.1) as follows.

(16.2)

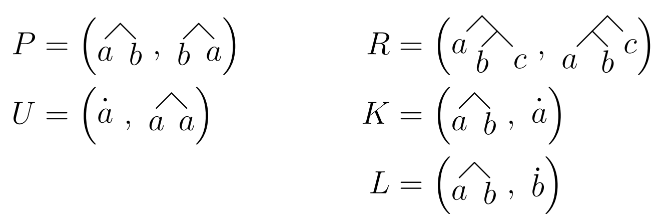

We regard the pairs in (16.2) as specifying a function from the leaves of the second tree to the leaves of the first tree. The function is determined by the labels in that it takes a leaf of the second tree with a given label to the unique leaf of the first tree with the same label. With our conventions, the function is well defined, need not be injective, and need not be surjective.

If we regard an expression modification as a morphism between labelings, and we regard a map from the leaves of one tree to the leaves of another as a morphism between trees, then the relation between the two kinds of morphisms is contravariant since a labeling is a map from leaves of trees to a fixed set of variables. This requires a careful discussion of conventions of order of writing and order of composition.

As in $K : ab \to a$, we choose to write expression modifications with an arrow to the right. Using the same symbol K for the corresponding map between tree leaves, we will write $K : \overset{\frown}{a\ b} \leftarrow \dot{a}$ with an arrow to the left. We continue to compose maps between trees from left to right, and we will compose expression modifications from right to left. We anticipate the discussion of multiplication with the simple example

$$ab \xrightarrow{P} ba \xrightarrow{K} b \tag{16.3}$$

which allows us to write $L = KP$.

On the Cantor set, we continue to compose from left to right and write the action on the right. Given a function $\rho : \Lambda(T_1) \leftarrow \Lambda(T_2)$, the induced continuous function that we still denote by ρ from the Cantor set $\mathfrak{C}$ to itelf has the usual definition

$$(u\alpha)\rho = (u\rho)\alpha, \quad u \in \Lambda(T_2), \quad \alpha \in \mathfrak{C}. \tag{16.4}$$

As usual, we can omit mentioning the suffix α as predictable and write the effects of the elements in (16.2) on $\mathfrak{C}$ as follows.

$$\begin{aligned} &P : \begin{cases} 1 & \leftarrow 0 \\ 0 & \leftarrow 1 \end{cases} \qquad && R : \begin{cases} 0 & \leftarrow 00 \\ 10 & \leftarrow 01 \\ 11 & \leftarrow 1 \end{cases} \\ &U : \begin{cases} \emptyset & \leftarrow 0 \\ \emptyset & \leftarrow 1 \end{cases} && K : 0 \leftarrow \emptyset \\ & && L : 1 \leftarrow \emptyset \end{aligned} \tag{16.5}$$

In words, P affinely interchanges the two halves of $\mathfrak{C}$ and is the generator $\pi_{0,1}$ of T and V from Sections 12.8 and 13.2.2, R is the generator x_0 of F, T and V, U is the one-sided full shift (doubling map) of Section 13.1.1 on $\mathfrak{C}$, K maps $\mathfrak{C}$ affinely onto its left half, and L maps $\mathfrak{C}$ affinely onto its right half.

16.3. The monoid. We base the monoid $\mathscr{M}$ on the Cantor set. Expression modifications clearly map to Cantor set endomorphisms, but we do not discuss the possible kernel in going from expression modifications to Cantor set endomorphisms. So calculations using expression modifications will be limited to showing that certain elements are equal, and functions on the Cantor set will be used when we want to show that certain elements are not equal.

Definition 16.3. The Thompson monoid $\mathscr{M}$ is the topological full monoid of the full one sided shift σ of Section 13.1.1 on the Cantor set $\mathfrak{C}$ in that the elements of $\mathscr{M}$ are the endofunctions of $\mathfrak{C}$ whose germs are compositions of germs of σ and their inverses.

The definition makes it clear that $\mathscr{M}$ is a monoid under composition, and from Definition 13.6 it is clear that the group of invertible elements of $\mathscr{M}$ is the group V. The following is a parallel of Proposition 13.7. It ties the definition above to the discussion in Sections 16.1 and 16.2. The vocabulary is from Section 13.1.5.

Proposition 16.4. *For each $f \in \mathscr{M}$, there are finite binary trees T_1 and T_2 and a function $\rho : \Lambda(T_1) \leftarrow \Lambda(T_2)$ that determines f by the formula in (16.4). Further, if T is a tree so that for each $u \in \Lambda(T)$, f is a rigid cone map on $u\mathfrak{C}$, then T_2 can be chosen to contain T.*

Proof. If T is not given as in the assumption of the second sentence of the statement, then we find one. As in Section 13.1.5 germs of f are represented by rigid cone maps whose domains are a cover of $\mathfrak{C}$ by clopen sets, and as in Section 13.1.6 there is a finite cover $\{u_1\mathfrak{C}, u_2\mathfrak{C}, \ldots, u_k\mathfrak{C}\}$ of $\mathfrak{C}$ by pairwise disjoint cones on each of which f is a rigid cone map. From Lemmas 8.5 and 8.8, the u_i are the leaves of a finite binary tree T.

The image of each $u_i\mathfrak{C}$ is some $v_i\mathfrak{C}$, but the collection $\{v_1, \ldots, v_k\}$ need not be of pairwise orthogonal elements of $\{0,1\}^*$. It is not a problem if $v_i = v_j$ for some $i \neq j$ since the ρ that we seek need not be injective, but we might have $v_i \prec v_j$ which prevents the v_i from being the leaves of a tree.

From Lemma 8.11 there is a smallest finite tree T_1 that contains all the v_i. For each v_i, we form $S_i = (T_1)_{v_i}/v_i$, the tree rooted at $\emptyset$ isomorphic to the subtree of T_1 rooted at v_i. The tree S_i is trivial for each v_i that is a leaf of T_1, but keeping these v_i makes the next discussion more uniform. Let

$$T_2 = T \bigcup_{u_i \in \Lambda(T)} u_i S_i.$$

In words, for each u_i we hang the tree S_i on T at u_i.

A leaf of T_2 is of the form $u_i\lambda$ for some $\lambda \in \Lambda(S_i)$ and $v_i\lambda$ is in $\Lambda(T_1)$. We let ρ take each such $u_i\lambda$ to $v_i\lambda$, and we note that f takes $u_i\lambda\mathfrak{C}$ rigidly to $v_i\lambda\mathfrak{C}$. Now T_1, T_2 and ρ are the items sought. $\square$

16.4. The multiplication. Proposition 16.4 makes multiplication in $\mathscr{M}$ reasonably straightforward, if a bit more complicated than in F, T and V. Doubts of correctness might arise when we use the representation as expression modifications, but they can be dispelled by checking calculations in $\mathfrak{C}$. We will not do such checks and leave them to the reader.

Consider $(R_2, \sigma_2, D_2)(R_1, \sigma_1, D_1)$ to be two elements to multiply where $\sigma_i : \Lambda(R_i) \leftarrow \Lambda(D_i)$. As usual the multiplication is opportunistic, and if $D_2 = R_1$, then the product is $(R_2, \sigma_1\sigma_2, D_1)$. Remember we compose the σ_i from left to right. If the equality is not available, we get it in two steps. The first to replace (R_2, σ_2, D_2) by (R_2', σ_2', D_2') representing the same element as (R_2, σ_2, D_2) and with D_2 containing R_1. This is done in the usual way with binary splittings, but now any splitting of a leaf v of R_2 must be matched by binary splittings at all leaves of D_2 that map to v under σ_2. Now (R_1, σ_1, D_1) is replaced by (R_1', σ_1', D_1') where R_1 is enlarged (if necessary) to R_1' to get $R_1' = D_2'$ using the same care with matching all splittings correctly. The result is $(R_1', \sigma_1'\sigma_2', D_1')$.

In spite of the careful description above, we will almost always justify calculations by composing modifications of expressions as illustrated in (16.3). The reason is that the calculations take very little space on the page.

The calculations can become intricate in resolving domains and ranges. Recall that R is the same as x_0. We have that $a(b(cd)) \to a((bc)d)$ is the same as x_1, and $a(b(c(df))) \to a(b((cd)f))$ is the same as x_2. The following verifies that $x_2 = x_0^{-1}x_1x_0$. It is only one line, but it requires many substitutions to set up.

$$a(b(c(df))) \xrightarrow{x_0} (ab)(c(df)) \xrightarrow{x_1} (ab)((cd)f) \xrightarrow{x_0^{-1}} a(b((cd)f))$$

Sometimes no substitutions are needed. Below is an efficient verification of the hexagon relation that appeared as (6) in Proposition 12.18. This will be a relation among elements of $\mathscr{M}$, so we use P for the equivalent $\pi_{0,1}$ and R for the equivalent x_0.

(16.6)
$$a(bc) \xrightarrow{R} (ab)c \xrightarrow{P} c(ab) \xrightarrow{R} (ca)b \xrightarrow{P} b(ca) \xrightarrow{R} (bc)a \xrightarrow{P} a(bc)$$

This gives $(PR)^3 = 1$ or $(\pi_{0,1}x_0)^3 = 1$ as in Proposition 12.18.

The fact that a variable can show up more than once in the second expression of a modification can introduce problems in stringing together a set of compositions. The following calculation will be useful. It shows the problem, and also a solution. After the calculation we will explain why we will systematically ignore the solution.

In the calculation

$$
(16.7) \quad \begin{aligned} & ab \xrightarrow{U} (a_1b_1)(a_2b_2) \xrightarrow{R} ((a_1b_1)a_2)b_2 \xrightarrow{K} \\ & (a_1b_1)a_2 \xrightarrow{P} a_2(a_1b_1) \xrightarrow{R} (a_2a_1)b_1, \end{aligned}
$$

the introduction of two copies of each of a and b in the first modification creates problems for the modifications that follow. We resolve this by giving the two appearances different subscripts so that they uniquely determine positions for the modifications that come after the first one. The convention will be that as maps of leaf sets, all leaves labeled a_i, for example, map to a in the leftmost expression.

However, (16.7) can be written less carefully as

$$
(16.8) \quad ab \xrightarrow{U} (ab)(ab) \xrightarrow{R} ((ab)a)b \xrightarrow{K} (ab)a \xrightarrow{P} a(ab) \xrightarrow{R} (aa)b
$$

and the net effect is unambiguous in the end, showing that that $RPKRU$ is the modification $ab \to (aa)b$. In all of the calculations that we will encounter, the result of the calculation will be clear without the use of extra subscripts on variables.

16.5. Deferments. Before we start work on the presentation of $\mathscr{M}$, we pause for some infrastructure. The calculation (16.8) shows that a "deferment" of U can be gotten from the elements in (16.1). Definition 8.14 describes the deferment of an element f of F to some $v \in \{0,1\}^*$. We give an equivalent definition for $\mathcal{M}$.

DEFINITION 16.5. Let $g \in \mathscr{M}$ be given as an endomorphism of $\mathfrak{C}$ and let v be in $\{0,1\}^*$. Then the *deferment* $g_v : \mathfrak{C} \to \mathfrak{C}$ of g to v is defined by $wg_v = w$ if v is not a prefix of w and $(vw)g_v = v(wg)$.

For examples, given as expression manipulations, we have the following.

$$
\begin{aligned}
P_0 &: (ab)c && \to && (ba)c \\
R_1 &: a(b(cd)) && \to && a((bc)d) \\
U_0 &: ab && \to && (aa)b \\
K_1 &: a(bc) && \to && ab \\
L_{01} &: (a(bc))d && \to && (ac)d
\end{aligned}
$$

This gives us more to work with. The calculation

$$
ab \xrightarrow{U} (ab)(ab) \xrightarrow{P_0} (ba)(ab) \xrightarrow{K} ba
$$

gives $P = KP_0U$. More generally if $g : e_1 \to e_2$ is an expression modfication, then

$$e_1 \xrightarrow{U} (e_1)(e_1) \xrightarrow{g_0} (e_2)(e_1) \xrightarrow{K} e_2 \tag{16.9}$$

shows that $g = Kg_0U$. Similarly $g = Lg_1U$.

We have a variation that will be more important later. Note that by replacing a variable by an expression if needed, we can guarantee that in $g : e_1 \to e_2$, the expression e_2 is not a single variable and thus splits into $e_2 = e_2'e_2''$. Now we have

(16.10)

$$e_1 \xrightarrow{U} e_1e_1 \xrightarrow{g_1} e_1(e_2'e_2'') \xrightarrow{g_0} (e_2'e_2'')(e_2'e_2'') \xrightarrow{K_0} e_2'(e_2'e_2'') \xrightarrow{L_1} e_2'e_2''$$

showing $g = L_1K_0g_0g_1U$.

We will later use the fact that $U_0 = RPKRU$ follows from (16.8).

16.6. Generators. Our generating set will be $X = \{P_0, R_0, K, U\}$. It is possible to start with $Y = \{P, R, K, U\}$ but it is easier to work with X. The reader can take as an exercise to show that P_0 and R_0 can be obtained as combinations of the elements of Y. We will use $\langle X\rangle$ to denote the monoid generated by X.

From X we get Y since $R = KR_0U$, and from (16.9) we get $P = KP_0U$. We also have $L = KP$ from (16.3).

We note that P and P_0 are their own inverses and we get $R^{-1} = PRPRP$ from P and R using (16.6), and thus get R_0^{-1} from R_0 and P_0. So $\{P, P_0, R, R_0\}$ generates a group. Trivially $R_1 = PR_0P$ and $P_1 = PP_0P$, and in the language of Proposition 13.16 we have $\{R, R_1, P, P_1\} = \{x_0, x_1, \pi_{0,1}.\pi_{10,11}\}$. We have shown the following.

LEMMA 16.6. *The group V is in $\langle X\rangle$.*

Those who would rather deal with semigroups than monoids might be bothered by inverses, subgroups and discussions of conjugations that will come later. It is not at all hard to treat $\mathscr{M}$ formally as a semigroup and we will indicate how do so in Section 16.8. However, until then we will be liberal with our use of inverses.

We return to deferments. We have

$$(ab)c \xrightarrow{P_0} (ba)c \xrightarrow{R^{-1}} b(ac) \xrightarrow{L} ac$$

giving $K_0 = LR^{-1}P_0$. And we know that (16.8) gives $U_0 = RPKRU$.

LEMMA 16.7. *All deferments of elements of Y are in $\langle X\rangle$.*

PROOF. The generators of V and their inverses are in $\langle X\rangle$ and so V is in $\langle X\rangle$. From Proposition 13.15, we know that for every u and v

in $\{0,1\}^*$ that are not the empty word, there is an element of V that carries $u\mathfrak{C}$ rigidly to $v\mathfrak{C}$. We know that the deferments P_0, R_0, K_0 and U_0 are in $\langle X\rangle$, and the conclusion follows. $\square$

LEMMA 16.8. *The set* $X = \{P_0, R_0, K, U\}$ *generates* $\mathscr{M}$.

PROOF. Let $e_1 \to e_2$ be an expression manipulation. Let n be the length of e_2. Every variable from e_1 appears no more than n times in e_2. Applying U^n to e_1 gives an expression e_3 with 2^n repetitions of every variable. Judicious applications of deferments of K and L can cut e_3 to an expression e_4 which has the same number of repetitions of each variable as are found in e_2. Thus e_4 reads the same as e_2 except that the order and the paretheses might be incorrect. Now order and parentheses can be rearranged at will by elements of V. But V is contained in $\langle X\rangle$, and we have succeeded in transforming e_1 to e_2. $\square$

16.7. Deferments from the generators. Our arguments concerning a presentation for $\mathscr{M}$ will make use of a systematic method of defining deferments of the generators. Note that P_0 and R_0 are already generators and we have given definitions of K_0 and U_0 in Section 16.6. Consider $g : e_1 \to e_2$ in $\mathscr{M}$ and consider the following calculations.

$$ae_1 \xrightarrow{P} e_1a \xrightarrow{g_0} e_2a \xrightarrow{P} ae_2,$$

$$(e_1b)c \xrightarrow{R^{-1}} e_1(bc) \xrightarrow{g_0} e_2(bc) \xrightarrow{R} (e_2b)c, \text{ and}$$

$$(ae_1)c \xrightarrow{P} c(ae_1) \xrightarrow{R} (ca)e_1 \xrightarrow{g_1} (ca)e_2 \xrightarrow{R^{-1}} c(ae_2) \xrightarrow{P} (ae_2)c$$

Now let x be in X. Motivated by the above, if x_{0v} is defined in $\mathscr{M}_X$, we define $x_{1v} = Px_{0v}P$, and $x_{00v} = Rx_{0v}R^{-1}$, and if x_{1v} is defined in $\mathscr{M}_X$, we define $x_{01v} = PR^{-1}x_{1v}RP$. Note that the word in the subscript is being modified or added to at its left end and not the right.

We already have definitions of K_0 and U_0 in terms of X and since P_0 and R_0 are of the form x_{0v} with $v = \emptyset$, we can use the scheme above to acquire definitions of P_{00} and K_{00}. Thus we have x_0 defined for every $x \in X$. The scheme now gives x_u for each $x \in X$ and $u \in \{0,1\}^*$, but we can say some additional nice things about the scheme.

Let $X_0 = \{x_0 \mid x \in X\}$. The scheme above gives for each $u \in \{0,1\}^*$ a word w in the invertible elements P and R so that "conjugation" of each $x_0 \in X_0$ by that w gives x_u. If we let $X_u = \{x_u \mid x \in X\}$, then the equality $X_u = X_0^w := \{w^{-1}x_0w \mid x_0 \in X_0\}$ makes sense. Further if v is a word in X, then we can write $v_u = w^{-1}v_0w$ where v_0 is obtained from v by replacing each appearance of $x \in X$ in v by x_0. The resulting v_u is what is obtained by replacing each appearance of $x \in X$ in v by x_u.

16.8. Relations. Our aim is to find a presentation for the monoid $\mathscr{M}$. Relations will take the form $w_1 = w_2$ where the w_i are words in the generators. The power of such a relation will be that for all words u and v in the generators, we can conclude $uw_1v = uw_2v$.

We build a relation set slowly. Instead of the process of previous chapters where an infinite set of relations was shown to reduce to a finite set, we work with a finite set of relations and show that they have infinitely many consequences. We will build a finite relation set $\mathscr{R}$ in stages by adding new relations at each stage. After each addition, we will continue to use $\mathscr{R}$ to denote the set of relations accumulated so far.

As is expected, we only introduce relations that hold in the monoid $\mathscr{M}$. But we have to be careful to distinguish between relations that are known to hold in $\mathscr{M}$ and relations that are derivable from the relations that we have introduced. To notate this distinction, we will label a relation with $\models$ if the relation is known to hold in $\mathscr{M}$. We will label a relation with $\vdash$ if the relation follows from our relation set $\mathscr{R}$. The set $\mathscr{R}$ grows throughout our discussion, but a late change will never invalidate any early claim that a relation follows from $\mathscr{R}$.

In spite of the fact that the relation set $\mathscr{R}$ takes a while to stabilize, we will use $\mathscr{M}_X$ to denote the monoid presented by $\langle X \mid \mathscr{R}\rangle$. Since every relation in $\mathscr{R}$ holds in $\mathscr{M}$ and X generates $\mathscr{M}$, we will have a homomorphism from $\mathscr{M}_X$ onto $\mathscr{M}$. The goal is to get enough into $\mathscr{R}$ to make this homomorphism a monomorphism.

Our usual process has been to derive a normal form for a word in terms of the generators, or of elements derived from the generators. The interest in the following is that this outline is not followed literally. There is a normal form in the argument but it is hidden and comes in pieces. For each element of $\mathscr{M}$, there will be no single expression that captures the behavior of that element.

16.8.1. *Invertiblity relations.* From previous observations, we have $PP = \mathbf{1}$, and $(PR)^3 = (RP)^3 = \mathbf{1}$. This gives inverses for P and R. We set

$$\mathscr{R} := \mathscr{R}_I := \{PP = \mathbf{1}, (RP)^3 = \mathbf{1}, (PR)^3 = \mathbf{1}\}.$$

For those who would prefer semigroup relations we offer the alternative

$$\begin{aligned}\mathscr{R} := \mathscr{R}_I := \{PPx = xPP = x(RP)^3 = (RP)^3x \\ = x(PR)^3 = (PR)^3x = x \mid x \in X\},\end{aligned} \tag{16.11}$$

which is still finite. We set $P^{-1} = P$ and $R^{-1} = PRPRP$ as formal definitions so that the symbols P^{-1} and R^{-1} can be used in what follows.

16.8.2. *Commutativity relations.* The following is clear.

$$\models: \forall g \in \mathscr{M}, \forall h \in \mathscr{M}, g_0 h_1 = h_1 g_0.$$

To select from these relations we use from Section 16.7 the sets $X_0 = \{P_{00}, R_{00}, K_0, U_0\}$ and $X_1 = \{P_{10}, R_{10}, K_1, U_1\}$. We set

$$\begin{aligned} \mathscr{R}_C &:= \{x_0 y_1 = y_1 x_0 \mid (x_0, y_1) \in X_0 \times X_1\}, \\ \mathscr{R} &:= \mathscr{R}_I \cup \mathscr{R}_C. \end{aligned}$$

Now with g and h representing words in X, we have notations such as g_0 or h_1 interpreted as having each appearance of $x \in X$ replaced by x_0 or x_1 as appropriate. The following is immediate from $\mathscr{R}_C$.

$$\vdash: \forall g \in \mathscr{M}_X, \forall h \in \mathscr{M}_X, g_0 h_1 = h_1 g_0.$$

16.8.3. *Splitting relations.* We first argue that

$$\models: \forall g \in \mathscr{M}, Ug = g_0 g_1 U.$$

Let $g : e_1 \to e_2$. We have

$$e_1 \xrightarrow{g} e_2 \xrightarrow{U} (e_2)(e_2),$$

$$e_1 \xrightarrow{U} (e_1)(e_1) \xrightarrow{g_1} (e_1)(e_2) \xrightarrow{g_0} (e_2)(e_2).$$

We set

$$\begin{aligned} \mathscr{R}_S &:= \{Ux = x_0 x_1 U \mid x \in X\}, \\ \mathscr{R} &:= \mathscr{R}_I \cup \mathscr{R}_C \cup \mathscr{R}_S. \end{aligned}$$

To argue

$$\vdash: \forall g \in \mathscr{M}_X, Ug = g_0 g_1 U$$

we note that if $\vdash: Ug = g_0 g_1 U$ for some word g in X, then

$$Ugx = g_0 g_1 Ux = g_0 g_1 x_0 x_1 U = g_0 x_0 g_1 x_1 U$$

by use of the commutativity relations.

16.8.4. *Reconstruction relations.* From our previous calculation (16.10) we have that

$$\models: \forall g \in \mathscr{M}, g = L_1 K_0 g_0 g_1 U.$$

We set

$$\begin{aligned} \mathscr{R}_R &:= \{x = L_1 K_0 x_0 x_1 U \mid x \in X\}, \\ \mathscr{R} &:= \mathscr{R}_I \cup \mathscr{R}_C \cup \mathscr{R}_S \cup \mathscr{R}_R. \end{aligned}$$

Now we get

$$\vdash: \forall g \in \mathscr{M}_X, g = L_1 K_0 g_0 g_1 U$$

by noting

$$\begin{aligned} gx &= L_1K_0g_0g_1Ux \\ &= L_1K_0g_0g_1x_0x_1U \\ &= L_1K_0g_0x_0g_1x_1U. \end{aligned}$$

16.8.5. *Rewriting relations.* From this point, no more calculations with expression modifications will appear. Arguments will be based on previous calculations and the view of elements of $\mathscr{M}$ as endofunctions of $\mathfrak{C}$.

We have arrived at the core of the argument. For motivation, consider $g \in \mathscr{M}$ and an expression Wg with W a word in $\{K, L\}$. We illustrate with an example of sufficient lack of symmetry, such as $Wg = KLKKg$. For $v \in \mathfrak{C}$, we have

$$(v)KLKKg = (0v)LKKg = (10v)KKg = (010v)Kg = (0010v)g$$

or, in other words, $(u^*v)g$ where u^* is the word in $\{0, 1\}$ created from the reverse $KKLK$ of $W = KLKK$ by the replacements $K \to 0$ and $L \to 1$. If now g is determined by a function $f : A \to B$ between prefix sets A and B for $\{0, 1\}^*$ so that u^* is in A, then u^*f is in B and $(u^*v)g = (u^*f)v$ for all v. Thus we would have $Wg = W'$ as elements of $\mathcal{M}$ where W' is the reverse of the word in $\{K, L\}$ that translates to u^*f under the replacements discussed above.

The reader can verify that the following all hold in $\mathscr{M}$.

$$\begin{aligned} KU &= \mathbf{1}, & KKKR_0 &= KK, \\ LU &= \mathbf{1}, & LKKR_0 &= KLK, \\ KKP_0 &= LK, & LKR_0 &= LLK, \\ LKP_0 &= KK, & LR_0 &= L. \\ LP_0 &= L, \end{aligned} \tag{16.12}$$

Note that for each $x \in X \setminus \{K\}$, the reverse of the words W used in front of x in the above form a prefix set for $\{K, L\}^*$. If we note that $K = K$ always holds, then we can say that $W = \emptyset$ works for $x = K$, noting that $\emptyset$ is also a prefix set for $\{K, L\}^*$.

We set $\mathscr{R}_W$ to be the set of relations given in (16.12) above. A relation $K = K$ is not needed. As usual, we grow $\mathscr{R}$ and set it to

$$\mathscr{R} := \mathscr{R}_I \cup \mathscr{R}_C \cup \mathscr{R}_S \cup \mathscr{R}_R \cup \mathscr{R}_W.$$

Once again, to view $\mathscr{M}$ as a semigroup instead of a monoid, we could deal with UK and UL as was done in (16.11).

We claim that a family of relations follow from $\mathscr{R}$. Specifically for each $g \in \mathscr{M}_X$, there is an integer $N(g) \geq 0$ so that if W is a word in $\{K, L\}$ with $|W| \geq N(g)$, then $\vdash$: $Wg = W'$ for some word W' in $\{K, L\}$.

The argument starts with those g with $|g| = 1$ by taking $N(K) = 0$, $N(U) = 1$, $N(P_0) = 2$ and $N(R_0) = 3$. Any word W in $\{K, L\}$ of length greater than $N(x)$ as just given is of the form $W = W_1W_2$ where W_2x appears as a left side of a relation in $\mathscr{R}_W$. Then $Wx = W_1W_2x = W_1W'$ as desired.

To induct on $|g|$ we consider gx and let $N(gx) = N(g)+N(x)$ where it is assumed that $N(g)$ has been defined. Now W with $|W| \geq N(gx)$ is of the form $W = W_1W_2$ where $|W_1| \geq N(x)$ and $|W_2| \geq N(g)$. We get $Wgx = W_1W_2gx = W_1W'x$ and we know that $|W_1W'| \geq N(x)$ as needed.

It will be clear when the presentation of $\mathscr{M}$ is fully developed, that the argument for it also solves the word problem. Since the number of words in $\{K, L\}$ of length n is 2^n, this only gives an exponential upper bound on the complexity of the solution.

16.8.6. *Implied relations, I.* This brings in the core machinery of the Thompson groups that deduces infinitely many relations from finitely many. This machinery alone suffices for Thompson's group F. It needs a bit of help in the setting of $\mathscr{M}$ which is supplied by the sections surrounding this one.

We currently have a set of relations that reads as

$$\mathscr{R} := \mathscr{R}_I \cup \mathscr{R}_C \cup \mathscr{R}_S \cup \mathscr{R}_R \cup \mathscr{R}_W.$$

Let us give this set another name

$$\mathscr{R}_\emptyset := \mathscr{R}_I \cup \mathscr{R}_C \cup \mathscr{R}_S \cup \mathscr{R}_R \cup \mathscr{R}_W.$$

We want to define a "deferment" $\mathscr{R}_v$ of $\mathscr{R}_\emptyset$ for each $v \in \{0,1\}^*$. We start with $\mathscr{R}_0$ and the rest will follow easily.

For each $g = h$ in $\mathscr{R}_\emptyset$, we can insist that each of g and h be written as fixed words in the elements of X. We now define for each $v \in \{0,1\}^*$ with $|v| \geq 1$, the set $\mathscr{R}_v$ to be derived from $\mathscr{R}_\emptyset$ by replacing each $g = h$ in $\mathscr{R}_\emptyset$ by $g_v = h_v$, where g_v and h_v are formed by replacing the appearancs of $x \in X$ with x_v as defined in Section 16.7.

We augment $\mathscr{R}$ for the last time and set

$$\mathscr{R} := \mathscr{R}_\emptyset \cup \mathscr{R}_0.$$

We have two claims to make at this point. First we claim that for all $v \in \{0,1\}^*$, and all $g_v = h_v$ in $\mathscr{R}_v$, that $\vdash$: $g_v = h_v$. This is immediate if $v \in \{\emptyset, 0\}$. But it follows inductively on $|v|$ from Section 16.7 that since every $g_0 = h_0$ in $\mathscr{R}_0$ is a word in symbols that are deferments to a location in $\{0,1\}^*$ that starts with 0, then for $|v| \geq 2$ or $v = 1$, the elements of $\mathscr{R}_v$ are conjugates by a single invertible element of $\mathscr{M}_X$ (that depends on v alone) of the elements of $\mathscr{R}_0$. This holds whether the monoid relations or the semigroup relations are being used. The remarks at the beginning of Section 16.8 now apply. This completes the argument for the first claim.

The second claim justifies the title of this section. We claim

$$\forall g \in \mathscr{M}_X, \forall h \in \mathscr{M}_X, \forall v \in \{0,1\}^*, \Big((\vdash: g = h) \Longrightarrow (\vdash: g_v = h_v)\Big).$$

This follows immediately from the first claim and the fact that every step in the chain of alterations connecting g to h using relations in $\mathscr{R}$ can have everything in it defered to v, giving a chain of alterations connecting g_v to h_v. By the first claim all such deferments are consequences of $\mathscr{R}$.

In fact we are going to use nothing of this claim except the following two special cases.

$$\begin{aligned}&\forall g \in \mathscr{M}_X, \forall h \in \mathscr{M}_X,\\ &\quad\Big(\big(\vdash: g = h\big) \Longrightarrow \big((\vdash: g_0 = h_0) \wedge (\vdash: g_1 = h_1)\big)\Big).\end{aligned}$$

16.8.7. *Implied relations, II.* We do not need to do this, but as motivation for a calculation that we will do immediately after, we will argue that in $\mathscr{M}$, we have

$$\begin{aligned}&\forall g \in \mathscr{M}, \forall h \in \mathscr{M},\\ &\quad\Big(\big((\models: Kg = Kh) \wedge (\models Lg = Lh)\big) \Longrightarrow \big(\models: g = h\big)\Big).\end{aligned}$$

It is easiest to argue from the point of view of actions on $\mathfrak{C}$. We have $(v)Kg = (0v)g$ so $Kg = Kh$ says that g and h agree on all infinite words that begin with 0. Similarly $Lg = Lh$ says that g and h agree on all infinite words that begin with 1. The implication follows.

Without adding to $\mathscr{R}$ we argue that

$$\begin{aligned}&\forall g \in \mathscr{M}_X, \forall h \in \mathscr{M}_X,\\ &\quad\Big(\big((\vdash: Kg = Kh) \wedge (\vdash Lg = Lh)\big) \Longrightarrow \big(\vdash: g = h\big)\Big).\end{aligned} \tag{16.13}$$

We have

$$\begin{aligned}(\vdash: Kg = Kh)\ &\wedge (\vdash: Lg = Lh),\\ (\vdash: K_0g_0 = K_0h_0)\ &\wedge (\vdash L_1g_1 = L_1h_1),\\ \vdash: K_0g_0L_1g_1 &= K_0h_0L_1h_1,\\ \vdash: K_0L_1g_0g_1 &= K_0L_1h_0h_1,\\ \vdash: g = K_0L_1g_0g_1U &= K_0L_1h_0h_1U = h.\end{aligned}$$

We now give a more complicated and weaker statement for convenience. We claim that given an integer $n>0$ we have

$$\begin{aligned} &\forall g \in \mathscr{M}_X, \forall h \in \mathscr{M}_X, \\ &\Big(\big(\forall W \in \{K,L\}^*, \\ &\qquad (|W| = n \implies \vdash : Wg = Wh) \big) \implies \big(\vdash : g = h \big) \Big). \end{aligned} \tag{16.14}$$

This follows by noting that the inner portion

$$\forall W \in \{K,L\}^*, (|W| = n \implies \vdash : Wg = Wh)$$

of (16.14) implies

$$\forall W \in \{K,L\}^*, (|W| = n-1 \implies \vdash : Wg = Wh)$$

by using (16.13) to eliminate the first letter of W.

16.9. Non-relations. In $\mathscr{M}$, the set $\{K,L\}^*$ is the free monoid on K and L. That is two different words in $\{K,L\}$ are different elements of $\mathscr{M}$. This will be used immediately. It follows by seeing what prefix the two words add to an arbitrary $v \in \mathfrak{C}$ using the coding of the first paragraph of Section 16.8.5.

16.10. The presentation. We now claim that $\langle X \mid \mathscr{R} \rangle$ presents $\mathscr{M}$. There is a homomorphism from $\mathscr{M}_X$ to $\mathscr{M}$ and we need to look at some g and h with the same image in $\mathscr{M}$. But we use the introductory remarks to Section 16.8.5 and claim that the identical behavior of g and h in their actions on $\mathfrak{C}$ say that there is an integer $N \geq 0$ (the larger of the depths of their domain trees, for example) so that for each $W \in \{K,L\}^*$ with $|W| \geq N$, we have $Wg = W'$ and $Wh = W''$ with W' and W'' words in $\{K,L\}$. But W' and W'' must represent the same element of $\mathscr{M}$, so by the non-relations of Section 16.9 we have $W' = W''$ as words in $\{K,L\}$. Thus $Wg = Wh$ in $\mathscr{M}_X$, and the fact that $g = h$ in $\mathscr{M}_X$ follows from (16.14).

17. Categorical squares

This section contains hindsight rather than history.

A very simple assumption on an object X in a category can cause Thompson's group V to be isomorphic to a subgroup of $\mathrm{Aut}(X)$. The assumption is so simple and the result so general, that as remarked in Section 15, the conclusion can sometimes be of little interest. But the result does have consequences, and we get a more focused example in Section 18.

We will show that in any category, if an object X is either the product or coproduct of two copies of itself, then there is a homomorphism from V to $\mathrm{Aut}(X)$. Since V is simple, this homomorphism is either trivial or injective.

We give the details in the case of a coproduct and leave the almost identical argument for products to the reader. From [144], we have the following.

Definition 17.1. A coproduct of objects A and B in a category $\mathcal{C}$ is an object C and a pair of arrows $\alpha : A \to C$ and $\beta : B \to C$ so that for any pair of arrows $f : A \to D$ and $g : B \to D$, there is a unique arrow $h : C \to D$ so that the following diagram commutes.

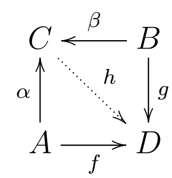

17.1. The construction. Let $X \xrightarrow{\alpha} X \xleftarrow{\beta} X$ be a coproduct. Let (D, σ, R) be a binary tree pair as in Definition 8.6. We wish to interpret D and R as a pair of diagrams. To treat them both, we let T be a finite binary tree in $\{0,1\}^*$.

At each vertex v of T, we place a copy of the object X that we will denote by X_v. For an internal vertex v, we insert the morphism α from X_{v0} to X_v and we insert the morphism β from X_{v1} to X_v.

We then do this for both D and R. To distinguish between the copies of X in D from those in R, we use $\widetilde{X}$ for the copies in R. For each leaf v in T, we put an identity morphism from X_v to $\widetilde{X}_{v\sigma}$. For example the element of V represented as

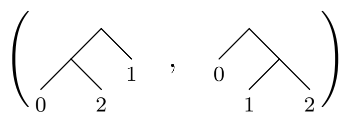

becomes the following diagram where we have flipped the second tree over a horizontal line for visibility.

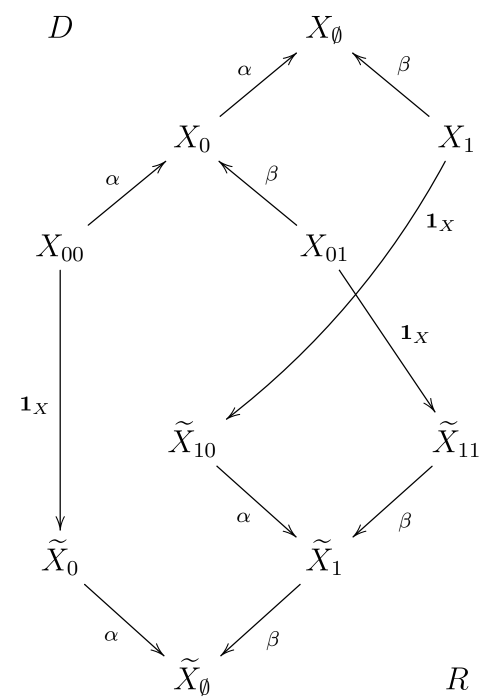

The treatment of the trees D and R now becomes assymmetric. Each $\widetilde{X}_v$ in R has a morphism μ_v from $\widetilde{X}_v$ to $\widetilde{X}_\emptyset$ at the root obtained by composing the arrows in the path from $\widetilde{X}_v$ to $\widetilde{X}_\emptyset$. Now for each leaf u of D there is a morphism γ_u from X_u in D to $\widetilde{X}_\emptyset$ in R by precomposing $\mu_{u\sigma}$ with $\mathbf{1}_X$.

For each vertex u in D, we inductively build a morphism γ_u from X_u in D to $\widetilde{X}_\emptyset$ in R. We already have γ_u if u is a leaf. If u is not a leaf of D, and both γ_{u0} and γ_{u1} have been defined, then we have the following square

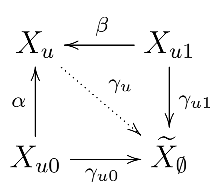

which induces the unique γ_u from the fact that $X_{u0} \xrightarrow{\alpha} X_u \xleftarrow{\beta} X_{u1}$ is a coproduct. At the root of D we have a morphism $\gamma_\emptyset$ from $X_\emptyset$ to $\widetilde{X}_\emptyset$. Remebering that the two objects are two copies of the same object, we have the following.

Proposition 17.2. *If* $X \xrightarrow{\alpha} X \xleftarrow{\beta} X$ *is a coproduct in some category, then the above construction creates a well defined homomorphism from* V *to* $\mathrm{Aut}(X)$*.*

If $X \xleftarrow[\alpha]{} X \xrightarrow[\beta]{} X$ *is a product in some category, then a dual of the above construction creates a well defined homomorphism from* V *to* $\mathrm{Aut}(X)$*.*

Proof. Left to the reader where most of the work has been done for the first claim. Checking well definedness and composition is needed first, and preservation of identity and inverses follows easily. For the second claim the reader has to modify the argument that gives a morphism from the top of one tree to the other. □

There is no guarantee that the homomorphism from V to $\mathrm{Aut}(X)$ is non-trivial. Since V is simple, the homomorphism is either trivial or injective. The homomorphism is injective if and only if the image of the element $\pi_{0,1} = (\underset{0\ \ 1}{\wedge}, \underset{1\ \ 0}{\wedge})$ maps to a non-trivial automorphism of X. In the first case, a typical coproduct is a "disjoint union," and examples of non-trivial homomorphisms abound. In the second case, products based on sets rarely have switching coordinates result in an identity morphism.

It is straightforward to verify that K and L in (16.5) from $\mathfrak{C}$ into itself expresses $\mathfrak{C}$ as the coproduct of two copies of itself and that the corresponding image of V in $\mathrm{Aut}(\mathfrak{C})$ is the usual representation of V as a group of homeomorphisms of $\mathfrak{C}$.

18. Jónsson-Tarski algebras

We return to history and present the construction from Jónsson-Tarski 1961 [119] of an algebra A that is free on both a set of size one and also on a set of size two. As mentioned in Section 15, this makes A a coproduct of two copies of itself. From Proposition 17.2, there is a homomorphism from V into $\mathrm{Aut}(A)$ which is easy to show is an injection. But V is actually isomorphic to $\mathrm{Aut}(A)$, and not just a subgroup.

We give the construction of the algebra following [119], and then go on to understand its automorphism group. Oddly, the construct starts with an object that is a product of two copies of itself in the extremely general setting of sets and functions between them.

18.1. Definitions. Let X be a set with a bijection $\alpha : X \to X \times X$ written on the right. Thus for every $x \in X$, we have $x\alpha = (x\alpha_0, x\alpha_1)$

where the α_i are surjections from X to X. Since α is a bijection, it has an inverse $\beta : X \times X \to X$ satisfying

$$\begin{aligned}(x_0, x_1)\beta\alpha_0 &= x_0,\\ (x_0, x_1)\beta\alpha_1 &= x_1,\\ x\alpha\beta = (x\alpha_0, x\alpha_1)\beta &= x,\end{aligned} \tag{18.1}$$

for all x_0, x_1 and x in X.

We can think of α_0, α_1 and β as operations on X where the α_i are unary operations and β is binary. The equalities in (18.1) are then identities that these operations satisfy.

Definition 18.1. A *Jónsson-Tarski algebra* (JT-algebra) is a quadruple $(X, \alpha_0, \alpha_1, \beta)$ where X is a set, and where α_0, α_1 are unary operations on X, β is a binary operation on X, all acting on the right, which satisfy the identities in (18.1)

A homomorphism from a JT-algebra $(X, \alpha_0, \alpha_1, \beta)$ to a JT-algebra $(X', \alpha'_0, \alpha'_1, \beta')$ is a function $h : X \to X'$ so that for each x and y in X we have $h(x\alpha_i) = (h(x))\alpha'_i$, $i \in \{1, 2\}$, and $h((x, y)\beta) = (h(x), h(y))\beta'$.

A JT-algebra $(X, \alpha_0, \alpha_1, \beta)$ is *free on* a set $S \subseteq X$ if for every JT-algebra $(X', \alpha'_0, \alpha'_1, \beta')$ and function $f : S \to X'$ there is a unique homomorphism $h : X \to X'$ that extends f.

The following is clear.

Lemma 18.2. *Let $m : X \times X \to X$ be a bijection and let $m^{-1} : X \to X \times X$ be denoted $xm^{-1} = (xa, xb)$. Then (X, a, b, m) is a J-T algebra.*

Conversely, let $(X, \alpha_0, \alpha_1, \beta)$ be a J-T algebra. Then $\beta : X \times X \to X$ is a bijection and $x\beta^{-1} = (x\alpha_0, x\alpha_1)$.

Lastly, the constructions just given connecting the two structures are mutually inverse to each other.

18.2. A free algebra. It follows from very general principles that for any set S, a free JT-algebra on S exists. The general principles can be learned from [42, Ch. II]. We will benefit from the details of the construction, so we will build a free JT-algreba on one generator from scratch. To show freeness, we will use a normal form, and for normal forms, we need terms.

Definition 18.3. If $(X, \alpha_0, \alpha_1, \beta)$ is a JT-algebra and V is a set of variables disjoint from $X \cup \{\alpha_0, \alpha_1, \beta\}$, then a *term* over V is defined recursively as follows. An element of V is a term. If x, s and t are terms, then so are $x\alpha_0$, $x\alpha_1$ and $(s, t)\beta$. If $v : V \to X$ is a fuction and T is a term over V, then v is called an *assignment* and the *value* of $T(v)$ in X is defined recursively in the obvious way.

Some terms can be simplified using the identities in (18.1). If a term T contains a subterm of the form $(t_0, t_1)\beta\alpha_i$ where t_0 and t_1 are terms, then this pair of operations can be replaced by t_i reducing the number of operations in T. If a term T contains a subterm of the form $(t\alpha_0, t\alpha_1)\beta$ where t is some term, then this triple of operations can be replaced by t reducing the number of operations in T. We say that a term T is *irreducible* if T contains no subterm of either of these two forms.

A term can be viewed as a tree of mixed arity in which a node labeled β can have two children and a node labeled either α_0 or α_1 can have one child. The following configurations in the tree can be eliminated and replaced by a single node, reflecting applications of the equalities in (18.1). The parenthsized numbers (0), (1), (2) over the arrows give a name to each type of replacement.

(18.2)

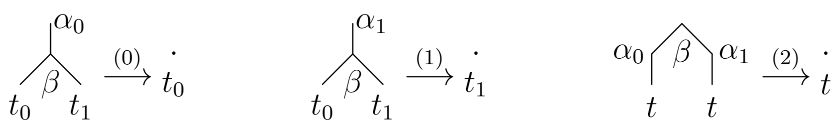

Because the term has only finitely many operations (the tree has finitely many nodes), it is clear that a term can be reduced to an irreducible. We will show that there is only one irreducible that a given term can reduce to. We use the diamond condition of Section 47, Lemma 47.2 and Corollary 47.2.1.

If a term allows for a choice between two replacements, then the two locations in the tree share no edges or they share an edge. If the two locations share no edge, then "doing both" brings the two choices to a single configuration as required by the diamond condition. The two possible arrangements where the locations share an edge are shown below.

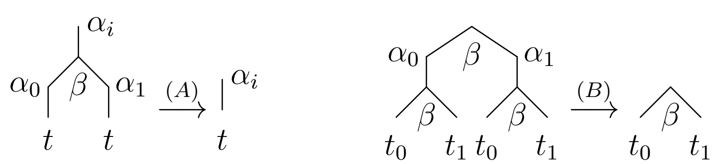

The reduction (A) on the left can be accomplished by doing a replacement of type (0) or (1) (depending on i), or by doing a type (2) replacement. So the two methods of reduction give the same result. The configuration on the right offers three places a replacement can be done, one each of type (0), (1) and (2). The final reduction (B) can be accomplished by a replacement of type (2) alone, or one each of types (0) and (1) in either order. Here, two of the three choices for starting the reduction process can be combined with another reduction to give the same result as the third. It follows from the diamond condition that every term reduces to a unique irreducible for that term.

The free JT-algebra JT_x on one variable is now defined to be the set of irreducible terms in which the only variable is x. If t and t' are terms, then $t\alpha_0$, $t\alpha_1$, and $(t, t')\beta$ can be formed and reduced to an irreducible, thus defining the expected operations on JT_x. The operations satisfy equations in (18.1) because of the reductions done. The algebra is free because sending x to any element of a JT-algebra J specifies a unique place to send every element of JT_x, and this is a homomorphism since the relations in (18.1) hold on J.

There is an alternate way to deduce uniqueness of reduced terms and freeness of the algebra. Starting from the fact that every term over x can be reduced to a (not yet known to be unique) irreducible, one carefully defines the operations α_i and β on irreducible terms covering all cases so that the result is again reduced. For example defining $(t_0, t_1)\beta$ breaks into two cases depending on whether or not there is a t' with $t_i = t'\alpha_i$, $i = 1, 2$. The result is a JT-algebra that the free algebra JT_x must map to, showing that different reduced terms in JT_x must be different elements in JT_x.

18.3. Structure of the free algebra. We look at the structure of an irreducible term in JT_x. From reduction types (0) and (1) in (18.2), we know that an irreducible term has no operation α_i that follows (is above in the tree) an operation β. Thus the appearance of the term tree is that of a binary tree whose root contains the value of the term, and whose leaves are roots of unary trees that are sequences of applications of the α_i. A simple example is the following where the nodes with two children correspond to the operaton β and the nodes with one child correspond to α_0 or α_1.

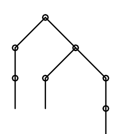

LEMMA 18.4. *The function that sends each $w \in \{\alpha_0, \alpha_1\}^*$ to xw in JT_x is an injection.*

PROOF. The term for xw is a tree of arity one. All such are irreducible, and thus give different elements of JT_x. □

LEMMA 18.5. *Let $S \subseteq JT_x$ be such that JT_x is free on S with $a \in S$. Then JT_x is free on $S' = (S \setminus \{a\}) \cup \{a\alpha_0, a\alpha_1\}$.*

PROOF. If given $f : S' \to J$ where $J = (X, \alpha_0', \alpha_1', \beta')$ is a JT-algebra, then let $a' = (a\alpha_0, a\alpha_1)f\beta'$ and let $f : S \to J$ have $f' = f$ on

$(S \setminus \{a\})$ and $af' = a'$. Then the unique homomorphic extension f'' of f' to JT_x agrees with f on S', and any homomorphic extension of f to JT_x must take a to a' and thus agree with f''. □

COROLLARY 18.5.1. *For any integer $n \geq 1$, there is a set $S \subseteq JT_x$ of n elements where JT_x is free on S.*

DEFINITION 18.6. A set S in JT_x for which JT_x is free on S will be called a *base* or *basis* for JT_x.

LEMMA 18.7. *A basis for JT_x is a minimal generating set in that removing any one element no longer generates. The converse does not hold.*

PROOF. If S is a basis with $a \in S$ and $S' = S \setminus \{a\}$ generates, then a is a term over S'. Now the function from S into JT_x that is the identity on S' and not on $\{a\}$ cannot extend to a homomorphism from JT_x to itself. For the lack of a converse, consider $y = (x, x\alpha_0)\beta$. The set $\{y\}$ generates JT_x, but not freely. □

It will be convenient to refer to the change from S to S' in Lemma 18.5 as a binary splitting of S at a. A sequence of binary splittings can be referred to as a binary refinement. It will also be convenient to identify words in $\{\alpha_0, \alpha_1\}^*$ with the vertices of the complete binary tree $\mathcal{T}$. A finite sequence $w = w_0 w_1 \cdots w_{n-1}$ over $\{0, 1\}^*$ corresponds to $\alpha_{w_0}\alpha_{w_1} \cdots \alpha_{w_{n-1}}$. Finite subtrees of $\mathcal{T}$ will be thought of as collections of words in $\{\alpha_0, \alpha_1\}^*$.

PROPOSITION 18.8. *If $U \neq \emptyset$ is a finite subset of $\{\alpha_0, \alpha_1\}^*$, then JT_x is free on $xU = \{xw \mid w \in U\}$ if and only if U is a prefix set for $\{\alpha_0, \alpha_1\}^*$ and thus the leaf set of a finite tree T_U.*

PROOF. The correspondence between prefix sets and leaf sets of finite trees is covered in Lemma 8.8. Assume that U is a prefix set for $\{\alpha_0, \alpha_1\}^*$, and the leaf set of the finite tree T_U. Each caret in T_U is a triple of the form $(u, u\alpha_0, u\alpha_1)$ where u is a proper prefix of some element of U. The tree T_U can be built from the trivial tree by a sequence of binary splittings. By inducting on the number of carets in T_U, it can be seen that xU is obtained from $\{x\}$ by a sequence of binary splittings, and that Lemma 18.5 implies that JT_x is free on xU.

For the inverse, the characterization of Lemma 7.2 says that if U is not a prefix set, then for some $w \in \{\alpha_0, \alpha_1\}^\omega$, either w has no prefixes in U or at least two prefixes in U. If $p \neq q$ from U are prefixes of w, then (say) $p \prec q$ and $q = pv$ for some non-empty v in $\{\alpha_0, \alpha_1\}^*$. Now in JT_x there are y and z with $z \neq yv$. So a function from xU to JT_x taking xp to y and xq to z cannot extend to a homomorphism.

If no $w \in \{\alpha_0, \alpha_1\}^\omega$ has more than one prefix in U, but some $w \in \{\alpha_0, \alpha_1\}^\omega$ has no prefixes in U, then the elements of U are pairwise orthogonal and there is a finite prefix p of w orthognal to all elements of U. By Lemma 8.11 there is a tree T with $U \cup \{p\}$ in its leaf set L. From the forward direction of this proposition, xL is a basis for JT_x. But it is not a minimal generating set since $U \subsetneq L$ generates JT_x, violating Lemma 18.7. □

18.4. Automorphisms.

Proposition 18.9. *For every automorphism η of JT_x there is a triple (U, σ, V) where $\sigma : U \to V$ is a bijection of prefix sets for $\{\alpha_0, \alpha_1\}^*$ and η is the unique homomorphic extension to JT_x of $\sigma' : xU \to xV$ defined by $(xu)\sigma' = x(u\sigma)$. Conversely, every such bijection $\sigma : U \to V$ induces an automorphism of JT_x.*

Proof. Let η be an automorphism of JT_x. Then JT_x is free on $\{x\eta\}$, and Lemma 18.4 and Proposition 18.8 apply in that sending $w \in \{\alpha_0, \alpha_1\}^*$ to $x\eta w$ is an injection, and if U is a finite subset of $\{\alpha_0, \alpha_1\}^*$, then JT_x is free on $x\eta U = \{x\eta w \mid w \in U\}$ if and only if U is a prefix set for $\{\alpha_0, \alpha_1\}^*$.

The element $x\eta$ is a term over $\{x\}$ and we let T be the binary tree part of the term from whose leaves the unary parts of the term are hung. The paths in T from the root to the leaves are words in $\{\alpha_0, \alpha_1\}$ that form a prefix set U for $\{\alpha_0, \alpha_1\}^*$, and $x\eta U = (xU)\eta$ is a basis for JT_x. The unary trees attached to the leaves of T express $(xU)\eta$ as xV where V must be a prefix set for $\{\alpha_0, \alpha_1\}^*$ by Proposition 18.8 because $(xU)\eta$ is a basis for JT_x.

We have $(xU)\eta = xV$, and if we set $\sigma : U \to V$ to be such that for all $u \in U$ we have $x(u\sigma) = (xu)\eta$, then setting $\sigma' : xU \to xV$ to have $(xu)\sigma' = x(u\sigma) = (xu)\eta$ makes η agree with σ' on xU. Since JT_x is free on xU, η must be the unique extension of σ' to JT_x.

The claimed converse follows from Proposition 18.8. □

The techniques in the proof of Proposition 18.9 also give the following in which (I) is a stepping stone to (II).

Proposition 18.10. *(I) If B is a finite basis of JT_x, and A is a binary refinement of $\{x\}$ with $|A| = |B|$, then there is a binary refinement of B that is also a binary refinement of A. (II) Any two finite bases of JT_x have a common binary refinement.*

Proof. (I) Let β be a bijection from A to B. There is a prefix set S for $\{\alpha_0, \alpha_1\}^*$ with $A = xS$. The automorphism η taking A to B with bijection β shows that $B = A\eta = (xS)\eta = x\eta S$.

The proof of Proposition 18.9 applied to $x\eta$ lets us conclude that for some prefix sets U and V for $\{\alpha_0, \alpha_1\}^*$ we have $x\eta U = xV$. Now S and U are prefix sets applied to $x\eta$ and S and V are prefix sets applied to x, and these prefix sets can be viewed as leaf sets of trees. If we form binary refinements of $x\eta U = xV$ we are simultaneously enlarging the trees for U and V. We can make a refinement deep enough so that the corresponding enlargements of the trees for U and V each contain the tree for S. This refinement of $x\eta U = xV$ will then be a common refinement of $A = xS$ and $B = x\eta S$. This proves (I).

(II) Enlarge the smaller basis by Lemma 18.5 if the bases are not of the same size. Then an automorphism applied to both carries one of them to a refinement of $\{x\}$. This reduces (II) to (I). $\square$

In the following, we take matched binary splittings of bijections between prefix sets for $\{\alpha_0, \alpha_1\}^*$ to be defined in parallel to Definition 7.4. We also take the following to be clear.

Lemma 18.11. *Triples of the type (U, σ, V) as in Proposition 18.9 that are related by matched binary splittings induce identical automorphisms of JT_x.*

Theorem 18.12. *The group $\text{Aut}(JT_x)$ is isomorphic to V.*

Proof. The bijection from $\{0,1\}^*$ to $\{\alpha_0, \alpha_1\}^*$ induced by sending $i \in \{0,1\}$ to α_i, in turn induces a bijection from triples (U, σ, V) with $\sigma : U \to V$ a bijection of prefix sets for $\{0,1\}^*$ to triples (U', σ', V') with $\sigma' : U' \to V'$ a bijection of prefix sets for $\{\sigma_0, \sigma_1\}^*$. This takes matched binary splittings to matched binary splittings, and so induces a well defined function from V to $\text{Aut}(JT_x)$. It is clear that this is a homomorphism which must be surjective by Proposition 18.9. From Lemma 18.4, one can show that the homomorphism is not trivial, and thus must be injective since V is simple. $\square$

Remark 18.13. As mentioned in Section 15, the autmorphisms of JT_x and generalizations were considered by Smirnov in 1973 [179]. In earlier papers Smirnov gave algebras such as JT_x and some generalizations the name "Cantor algebras" because JT_x was based on a bijection from an infinite set to a direct product of two copies of itself. A particular (and recursively computable) example of such a function is Cantor's function from $\mathbf{N} \times \mathbf{N}$ to $\mathbf{N}$ that "counts" the elements in $\mathbf{N} \times \mathbf{N}$ by running diagonally through the pairs starting at $(0,0)$ and counting the pairs (m,n) with $m + n = k$ before counting the pairs with $m + n = k + 1$. See Cantor 1895 in [44] Section 6, proof of (8) on Pages 106–107.

Not only is V isomorphic to $\mathrm{Aut}(JT_x)$, but the Thompson monoid $\mathscr{M}$ is isomorphic to the monoid of endomorphisms of JT_x. See Theorem 1.1 of De Witt-Elliott 2023 [194].

19. Splitting homotopy idempotents

We will show how Thompson's group F separates those unpointed homotopy idempotents that split from those that do not. As in Section 15, we let $\sim$ be the relation "homotopic to" and in the category of pointed, connected CW complexes we call a morphism $f : X \to X$ a homotopy idempotent if $f^2 \sim f$. Such an f is said to split if there are morphisms $d : X \to Y$ and $u : Y \to X$ so that, composing right to left, $ud \sim f$ and $du \sim \mathbf{1}_Y$. Note that in any situation where maps $d : X \to Y$ and $u : Y \to X$ are given where $du \sim \mathbf{1}_Y$, then setting $f = ud$ gives $f^2 = (ud)(ud) = u(du)d \sim ud = f$.

If the homotopy from f^2 to f keeps the basepoint of X fixed throughout the homotopy, then f is a pointed homotopy idempotent and it splits. This fact has nothing to do with F and we refer the reader to Theorem (2.1) of Hastings-Heller 1981 [102] for one proof. The reader should note that one of the mentions there of ud should read du. A different proof that factors through the Brown representability theorem is hinted at in the paragraphs following the statement of Theorem 1.3 in Heller 1981 [105].

We will be concerned with the case where the homotopy $H : X \times [0,1] \to X$ from f to f^2 does not preserve the basepoint. That is, the homotopy connects morphisms in the category, but is not a path through morphisms in the category. To make the statement of the main result precise, we expand on the discussion in Section 15. We assume for all $x \in X$ that $H(x,0) = f(x)$ and $H(x,1) = f^2(x)$. We let p be the basepoint of X, and we let $\alpha : [0,1] \to X$ be the loop at p defined by $\alpha(t) = H(p,t)$. If $\beta : [0,1] \to X$ is another loop at p, then with $\bar{\beta} : [0,1]^2 \to X \times [0,1]$ acting as $\bar{\beta}(s,t) = (\beta(s),t)$, the composition of $\bar{\beta}$ with H shows that we have $f^2\beta \sim \alpha^{-1}(f\beta)\alpha$. With C_α the inner automorphism of $\pi_1(X,p)$ given by conjugation by α, we have that $f_*^2 = C_\alpha f_*$. In the language of Section 10.2, we have that $(\pi_1(X,p), f_*, \alpha)$ is a group with a conjugacy idempotent.

With σ the shift homomorphism on F taking each x_i to x_{i+1}, Proposition 10.2 gives a unique homomorphism $\eta : F \to \pi_1(X,p)$ which acts as a morphism of groups with conjugacy idempotent from (F, σ, x_0) to $(\pi_1(X,p), f_*, \alpha)$. We can now state the following whose proof shows the entanglements of F in the situation.

THEOREM 19.1. *The following hold.*

(1) *With the notation as above, and where f is an unpointed homotopy idempotent, we have the following alternatives.*
  (a) *In the category of connected, pointed topological spaces, if η is an injection, then f does not split.*
  (b) *In the category of connected, pointed CW complexes, if η is not an injection, then f is unpointed homotopic to a pointed homotopy idempotent and thus splits.*

(2) *In the category of connected, pointed CW complexes, there is an unpointed homotopy idempotent that does not split.*

(3) *Every homotopy idempotent (pointed or not) on a finite dimensional CW complex splits.*

Item (3) will be proven later as Theorem 28.9.

Proof. (1a) With the notation as above, and with $\{x_0, x_1\}$ the usual two element generating set for F, we have

$$\eta(x_i) = \eta\sigma^i(x_0) = f_*^i\eta(x_0) \sim f^i\alpha \quad \text{for all } i \geq 0.$$

Since η is assumed to be injective, no two $f^i\alpha$ are homotopic. But if f splits, then there is a space with basepoint (Y,q) and maps $d : (X,p) \to (Y,q)$ and $u : (Y,q) \to (X,p)$ so that $du \sim \mathbf{1}_Y$ and $ud \sim f$. Thus the following diagram homotopy commutes.

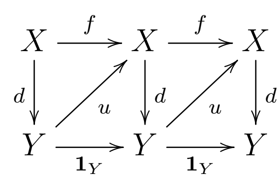

If we let $G = \pi_1(X,p)$ and $H = \pi_1(Y,q)$, then the following diagram commutes.

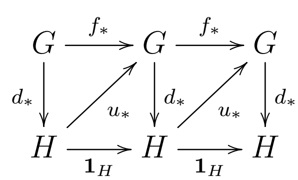

From this it follows that

$$f^2\alpha \sim f_*^2[\alpha] = u_*\mathbf{1}_H d_*[\alpha] = u_*d_*[\alpha] = f_*[\alpha] \sim f\alpha,$$

contradicting the injectivity of η.

(1b) The following argument is adapted from the proof of $5 \Rightarrow 1$ in Theorem 9.2.2 of Dydak-Segal 1978 [65].

If η is not an injection, then the image of η is abelian and all the x_i with $i \geq 1$ have the same image under η. In particular, $f\alpha \sim f^2\alpha$. We will essentially show that by dragging the base point around the reverse of α, we can homotop f to a pointed homotopy equivalence. The argument takes several steps.

We will use $f \xrightarrow{\alpha} f^2$ to sybolize the fact that f is homotopic to f^2 under a homotopty that drags the basepoint around the path α. If we follow the homotopy by f then we get a homotopy from f^2 to f^3 that drags the basepoint around the path $f\alpha$. This is symbolized by $f^2 \xrightarrow{f\alpha} f^3$. Doing this once more lets us write down the following string of symbols.

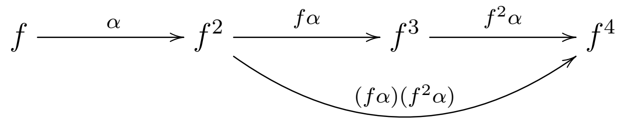

However, $f\alpha \sim f^2\alpha$ and we make the seemingly odd choice to point out that this implies $(f\alpha)(f^2\alpha) \sim (f^2\alpha)^2 = f^2\alpha^2$ where α^2 refers to the concatenation of two copies of the loop α and f^2 is simply the composition of f with itself.

By dragging the base point around the reverse of α^2, we build a homotopy J from f^2 to a map $g : (X,p) \to (X,p)$ and can denote this situation by $g \xrightarrow{\alpha^2} f^2$. If we apply J to the left factor of $gg = g^2$, we get $g^2 \xrightarrow{\alpha^2} f^2g$, and if we follow J by f^2, we get $f^2g \xrightarrow{f^2\alpha^2} f^4$. We now have the following.

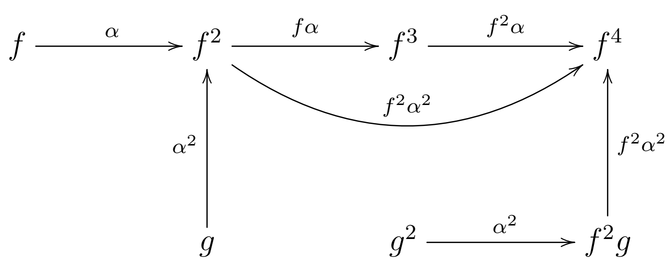

From this it is seen that that there is a homotopy from g to g^2 that drags the basepoint around a loop that is homotopic to a constant. It follows that g is a pointed homotopy idempotent. Since $f \sim f^2 \sim g$, we have the first conclusion that f is homotopic to a pointed homotopy idempotent. We know that pointed homotopy idempotents split, and so f splits as well.

(2) We build our example from F to fit the situation of (1a). Chapter 8, Section 1 of Spanier 1966 [181] covers the material needed. Let (X,p) be an Eilenberg-McLane space of type $(F,1)$. There is a self map $f : X \to X$ with $f_* = \sigma$, and since f_* and f_*^2 differ by conjugation by x_0, we have that f and f^2 are homotopic under a homotopy dragging

the basepoint around a loop that represents x_0. We may as well refer to this loop as x_0. Now $(\pi_1(X,p), f_*, x_0)$ is a group with conjugacy idempotent, and the canonical homomorphism η from (F, σ, x_0) is an isomorphism. By (1a), f does not split. $\square$

20. The associative law

The purpose of this section is to demonstrate the validity of the statement "Thompson's group F is the structure group of the associative law." This is not a well defined target, so we will do it twice. For the first, we will assume that the associative law gives a congruence relation, and for the second we will assume that it is not.

DEFINITION 20.1. If X is a set, if $\sim$ is an equivalence relation on X and if $\mathcal{O}$ is a set of operations on X, then $\sim$ is a *congruence relation* if for every operation $f \in \mathcal{O}$, if k is the arity of f, if (a_i) and (b_i), $0 \le i < k$, are two k-tuples in X, and if $a_i \sim b_i$ for $0 \le i < k$, then $f(a_0, \ldots, a_{k-1}) \sim f(b_0, \ldots, b_{k-1})$.

If H is a subgroup of G, then "living in the same left coset of H" is an equivalence relation on G which is a congruence relation for the group multiplication if and only if H is normal in G. The relation $=$ is always taken to be a congruence relation.

20.1. If the associative law is a congruence relation. The material below is extracted from Dehornoy's papers 1989 [53], 1993 [54], and 1996 [55]. The discussion will allow us to give an interesting infinite presentation for F.

If the associative law is given a direction, then it "acts" on fully parenthesized expressions. As in Section 16, we can associate a fully parenthesized expression with a partition of the Cantor set, and so an application of the associative law can be turned into an action on the Cantor set. This action is invertible, reflecting the fact that the associative law works both ways. Since we are taking the associative law to be a congruence relation, the law can also act on subexpressions. We now give formal details for building a group from the associative law.

We already have a symbol for one direction of the associative law from Secction 16, namely $R : a(bc) \to (ab)c$. As noted in that section, this corresponds to $x_0 \in F$ whose action on the Cantor set is determined by the prefix set map $00 \to 0$, $01 \to 10$, $1 \to 11$. Thus R (and its inverse) belong to the group that we are building.

The associative law can act on subexpressions, so given any position in an expression, if $a(bc)$ appears in that position, we can replace that appearance by $(ab)c$ while leaving the rest of the expression fixed. This translates to putting in our group every deferment R_u of R acting as $u00 \to u0$, $u01 \to u10$, $u1 \to u11$ and as the identity on every v with $v \perp u$.

We call the set $S = \{R_u \mid u \in \{0,1\}^*\}$ the *symmetric generating set* of our group, and we declare that the group generated by S to be the structure group of the associative law. Whether this term is acceptable to the reader is entirely up to the reader. Of course since S consists of x_0 and its deferments, we have that $S \subseteq F$. Since x_0 and x_1 are in S, we have $\langle S \rangle = F$. We can pretend not to know about F, but to do so completely will be confining. The argument just given depends on knowing about the action of F on the interval or Cantor set. We will give an argument based only on presentations. We do so because the presentation based on the symmetric generating set is interesting.

The elements of S satisfy the following relations.

$$(C_u) \qquad R_v R_u = R_u R_v \qquad v \perp u$$

$$(C_{u00}) \qquad R_{00v} R_u = R_u R_{0v} \qquad u \preceq 00v$$

$$(C_{u01}) \qquad R_{01v} R_u = R_u R_{10v} \qquad u \preceq 01v$$

$$(C_{u1}) \qquad R_{1v} R_u = R_u R_{11v} \qquad u \preceq 1v$$

$$(\Delta_u) \qquad R_{u0} R_u R_{u1} = R_u^2 \qquad u \in \{0,1\}^*$$

The relations (C_u) through (C_{u1}) follow because they express how R_u conguates or commutes with the actions of other generators where the fundamental triviality of Section 2.2 applies. The last relation can be verified any number of ways, but the following is efficient. It verifies the relation $\Delta_\emptyset$, but it translates to all other locations.

$$a(b(cd)) \xrightarrow{R_1} a((bc)d) \xrightarrow{R} (a(bc)d) \xrightarrow{R_0} ((ab)c)d$$

$$a(b(cd)) \xrightarrow{R} (ab)(cd) \xrightarrow{R} ((ab)c)d$$

Remember that under the order conventions of Section 16.4, the calculation above verifies $R_0RR_1 = RR$.

Note that all five ways to fully parenthesize a string of four variables are used in the calculation above. These correspond to the five finite binary trees with four leaves. Arranging them appropriately at the vertices of a pentagon and giving each edge a well chosen direction also verifies the relation (Δ_u). For this reason, (Δ_u) is often called the pentagon relation. Its fame goes back to (3.5) of MacLane 1963 [143], and the complex K_4 of Stasheff 1963 [182].

Proposition 20.2. *With S the symmetric generating set and $\mathcal{R}$ the relations (C_u) through (C_{u1}) and (Δ_u), the group presented by $\langle S \mid \mathcal{R}\rangle$ is isomorphic to F.*

Proof. From the actions of the R_u on expressions, we know that the group presented by $\langle S \mid \mathcal{R}\rangle$ is not abelian. We attempt to interpret each R_{1^i} as x_i of the presentation (9.1). The relations (C_{u1}) with $v = 1^j$, $j \geq 0$, as applied to the elements R_{1^i}, $i \in \mathbf{N}$, give all the relations of the presentation (9.1) for F under this interpretation. So by Lemma 9.8, sending each R_{1^i} to x_i extends to an isomorphism from $\langle S \mid \mathcal{R}\rangle$ to F if we can show that the R_{1^i} generate the rest of S modulo the relations in $\mathcal{R}$.

We induct on the depth of $u \in \{0,1\}^*$. Let A be the subset of S in the subgroup generated by the R_{1^i}, $i \in \mathbf{N}$. All the R_u with u of depth 0 are in A. Assume that all the R_u with depth u less than k are in A. We will induct backwards on position in the prefix order of Definition 8.1 to get all R_u with depth u equal to k in A. For u of depth k, assume that all v of depth k with $u < v$ in prefix order are in A. Note that 1^k is maximal in prefix order among elements of $\{0,1\}^*$ of depth k, and we already have R_{1^k} in A.

Let u have depth k, $u \neq 1^k$. Let p be the longest prefix of u ending in 0, making $p = q0$. If $p = u$, then $u = q0$, and $R_u = R_{q0} = R_q^2 R_{q1}^{-1} R_q^{-1}$ by (Δ_u). If $p \prec u$, then $u = q011^m$ for some $m \geq 0$ and $R_u^{R_q} = R_{q101^m}$ or $R_u = R_{q011^m} = R_q R_{q101^m} R_q^{-1}$ by (C_{u01}). In both cases R_u is a composition of elements of the form R_w where either w has depth less than k, or w has depth k and $u < w$ in the prefix order. ☐

20.2. If the associative law is not a congruence relation. We assume a fixed countably infinite set of variables. We consider the set E of pairs (e_1, e_2) where e_1 and e_2 are fully parenthesized expressions on the same finite string of variables where there the string uses no variable twice. Specifically, the order of the variables in e_1 and e_2 are the same.

If a is a variable used in e_1 and e_2 and we replace a in both e_1 and e_2 by the same fully parenthesized expression e_3 so that no duplication of variables is introduced, then we get new expressions e_1' from e_1 and e_2' from e_2 so that (e_1', e_2') is in E. We write $(e_1, e_2) \to (e_1', e_2')$ and let $[E]$ be the set of equivalence classes in E under the equivalence relation generated by $\to$. Under the relationship between expressions and trees given in Definition 16.2, together with a relation between the classes just described and representatives of elements of F, it is legitimate to identify $[E]$ with F as sets.

If we write $e_1 \sim e_2$ for a pair (e_1, e_2), then we can say that $e_1 \sim e_2$ is an associative law. Note that the symbol $\sim$ is not the equivalence relation of the previous paragraph. Under the usual interpretation of an associative law, it is accepted that $e_1 \sim e_2$ would imply $e_1' \sim e_2'$ for all (e_1', e_2') in the same class as (e_1, e_2) in $[E]$. So we can think of $\sim$ as an attribute that might or might not be possessed by elements of $[E]$, and that certain elements of $[E]$ are associative laws. We can ask under what assumptions all in $[E]$ are associative laws.

If A is the subset of $[E]$ for which $\sim$ holds, then with no assumptions on $\sim$, A is just an arbitrary subset. But if $\sim$ is assumed to be an equivalence relation, then A becomes a subgroup of F. For the reflexive hypothesis puts all (e, e) in A, the symmetric hypothesis puts the "inverse" (e_2, e_1) in A whenever (e_1, e_2) is in A and the transitive hypothesis puts (e_1, e_3) in A whenever both (e_1, e_2) and (e_2, e_3) are in A.

Since x_0 and x_1 generate F, we see that the assumption that $\sim$ is an equivalence relation and an assumption that both $(ab)c \sim a(bc)$ and $a((bc)d) \sim a(b(cd))$ hold, implies that all associative laws hold. The usual argument given in courses on basic algebra that the associative law $(ab)c = a(bc)$ implies all associative laws uses the fact that $=$ is a congruence relation. Whether this discussion justifies calling F the structure group of the associative law is left to the reader's judgement.

21. End notes

In Section 38, generalizations $V_{n,r}$ of V from Higman 1974 [107] are built. In 2009 [17], Birget defines four monoids $M_{n,1}$, $totM_{n,1}$, $surM_{n,1}$, and $invM_{n,1}$ closely related to $\mathscr{M}$, of which $M_{n,1}$ contains the other three. The monoid $totM_{2,1}$ is the Thompson monoid $\mathscr{M}$. Further, Proposition 2.1 of [17] states that $V_{n,1}$ is the group of units of all four.

All of these monoids generalize the group V by emphasizing the actions on the cones of the Cantor set rather than the entire Cantor set. But important power is lost if the focus is restricted to actions on only single cones. Hidden in all structures mentioned so far is the "coproduct condition" that (when $n = 2$) if a map f on a cone $u\mathfrak{C}$ is restricted to its two subcones $(u0)\mathfrak{C}$ and $(u1)\mathfrak{C}$, then those two restrictions determine f. This condition is captured in the monoid $\mathscr{M}$ by the reconstruction relations in Section 16.8.4.

Certain inverse semigroups ([129] or Pages 28 ff. of [50]), or the associated inverse monoids, focus entirely on actions on individual cones. A semigroup S is an inverse semigroup if for every $a \in S$, there is a unique $b \in S$ (called the inverse of a) for which $aba = a$ and $bab = b$. The archtypical example of an inverse semigroup is a symmetric inverse semigroup S for which there is a set A in which each element $a \in S$ has a domain $\mathbf{d}(a)$ and a range $\mathbf{r}(a)$ so that a is defined only on $\mathbf{d}(a)$ which is then taken bijectively by a onto $\mathbf{r}(a)$. A Cayley type theorem (Theorem 1.20 of [50]) embeds every inverse semigroup into a symmetric inverse semigroup.

Inverse monoids relevant to the Thompson groups are the polycyclic monoids of Nivat-Perrot 1970 [166] which generalize (to $n \geq 2$) the bicyclic monoid (Example 2 in Pages 43 ff. in [50]). The polycyclic monoid of rank n acts naturally on the n-ary Cantor set, but its domains (if non-empty) are always single cones. The polycyclic monoids cannot combine several domains into a larger domain and cannot express the coproduct condition.

But the $\mathbf{Z}$-linear ring over a polycyclic monoid can combine domains, and the coboundary condition can be captured by a single linear relation. If this is done, rings of Leavitt 1952–62 [136, 137, 138] are discovered in hindsight. In fact this was done independently of Leavitt and of each other in 2004 by Birget [16] and by Nekrashevych [159] (in a more analytic setting initiated by Dixmier 1964 [60] (Example 2.1) and developed by Cuntz 1977 [51]). It was observed in [16] and [159] that the $V_{n,r}$ represent into the resulting structures.

This discussion has nothing to do with the motivations of Leavitt, Dixmier and Cuntz. Relevant to us are the motivations of Leavitt which were to find rings over which the rank of free modules is ambiguous. The Leavitt ring $L_n = L_{\mathbf{Z}}(n, 1)$ parallels the Jónsson-Tarski algebra JT_x in that L_n is universal for the property that its free modules of ranks 1 and n are isomorphic as L_n-modules. The representations of the $V_{n,r}$ into the endomorphism rings of the L_n-modules allowed for the first (and thus far only) complete classification of the $V_{n,r}$ up to isomorphism. This is the subject of Section 40 and 41.

To round out this discussion, we note that the connections between the polycyclic monoids and the Thompson groups have been investigated further by Lawson 2007 [130, 131] and 2021 [132]. The Cuntz algebra $\mathcal{O}_n$ of [51] can be obtained from a complex-linear version $L_{\mathbf{C}}(n, 1)$ of the Leavitt ring by adding a suitable norm and taking a completion. Theorem 1.8 of [17] points out that the monoid $M_{n,1}$ of [17] is a submonoid of the multiplicative part of $\mathcal{O}_n$.

In parallel to the word problem for groups, Glass had been working on the word problem for lattice-ordered groups (see Section 1.3 of [85] for a definition) since 1975 [84]. In 1981 Glass-Gurevich (Chapter 13 of [85] with a more polished version in 1983 [86]) a finitely presented lattice-ordered group is built with an unsolvable word problem. The construction is motivated partly by McKenzie-Thompson [152], and uses the representation of arbitrary computable functions found there. But the construction in [152] has torsion elements which cannot exist in a lattice-ordered group, so the construction departs from [152] considerably. Interestingly, a corollary is given that there must be a finitely generated lattice-ordered group with an unsolvable word problem that is 1-related as a lattice ordered group.

Another set of irrational slopes, the integral powers of the golden ratio, was introduced in Cleary 2000 [48] where finiteness properties are proven. This group (with its T and V counterparts) are further studied in Burillo-Nucinkis-Reeves 2021 [40] and 2022 [41].

The exact lengths from Fordham's thesis [71] are not always needed to make geometric observations, and easily obtained approximations sometimes suffice. See for example Burillo 1999 [38].

It took time for various groups of people working with Thompson's groups to become aware of each other. The group F is written up in detail in Section 6.2 of the book Dydak-Segal 1978 [66] with no mention of Thompson, and there is also no mention in the 1981 and 1982 papers of Hastings-Heller [102] and [103]. Thompson is mentioned prominently in the introduction to Brown-Geoghegan 1984 [36].

The paper [36] also verifies Geoghegan's conjecture that $H^n(F, \mathbf{Z}F) = 0$ for all n, and we note, but do not prove this in Chapter 4.

CHAPTER 4

Actions on complexes

Chapter contents

22. Introduction

[1]Much of this chapter will be concerned with finiteness properties of groups. See the introduction to Chapter VIII of Geoghegan 2008 [80] for some useful references. The unpublished Bieri 1976 [12] can easily be found online. Remarks in the introduction to Bestvina-Brady 1997 [11] give some history. The material in these papers and Brown 1987 [34] covers the notions we will discuss.

Information about members of the Thompson family can be obtained from their actions on complexes. This chapter is an introduction to the topic. We give three examples of how such complexes can be built. The first is built for F from a detailed knowledge of the algebraic structure of its positive monoid. The second is built for V from its action on the free Jónsson-Tarski algebra. The third is the most specialized and is also built for F, but it is built from the knowledge that F is an initial object in the category of groups with a conjugacy idempotent. The complexes will be built as simplicial complexes, but some can also be shown to have a cubical structure. See Section 48 for basics on both kinds of complexes, and comments on the CAT(0) property that we mention below.

The most common use of the actions of Thompson groups on complexes is to prove finiteness properties. This activity tends to follow a certain outline that has become quite standard, and we illustrate that outline with the first complex and give a brief summary in Section 27.3. The finiteness properties will be defined in Section 22.1.

The first complex is built in Section 25. In Section 26 the complex is shown to have a CAT(0), cubical structure, and is used to prove the finiteness properties for F. Using results that lie outside the scope of this book, we will point out that this also shows that F has the Haagerup property and satisfies the Baum-Connes and Novikov conjectures.

The second complex is built in Section 27. The complex is used to derive a finite presentation for V. The calculation of a presentation for V by the techniques of Chapter 2 can be quite elaborate. In addition, the presentation we obtain has a certain simple appeal. The complex for V can also be used with the "standard outline" to show that V has type F_∞, and we point out what must be strengthened in our description of the outline and give a reference for the details.

The third complex is built in Section 28. The complex was historically the first complex built for a Thompson group, gives an alternate

[1]This chapter is not complete. More material will be added in the future.

proof of the main results in Sections 25 and 26, and in addition allows for a direct calculation of the integral homology groups of F. The construction uses the fact that F is an initial object in the category of groups with a conjugacy idempotent to build a complex that is an initial object for path connected CW complexes that have a homotopy idempotent. All the facts gathered lead to a direct proof of the last, unproven part of Theorem 19.1, that homotopy idempotents on finite dimensional complexes split.

22.1. Definitions. The finiteness properties that will be shown generalize the group properties of being finitely generated and finitely presented. Section 7.2 of [80] gives full backbground for the vocabulary and concepts used. We will give some definitions, but they will use terms that also need definitions and any gaps can be filled from [80]. Finiteness properties have been used to separate groups in various classes.

For us, a *classifying space* or $K(G,1)$ for a group G is a connected CW complex X so that $\pi_1(X) = G$ and $\pi_i(X) = 0$ for all $i > 1$. The spaces we will work with will be simplicial complexes within which a cubical structure can be detected. Any two classifying spaces for a group are homotopy equivalent ([80, Corollary 7.1.7]).

A group being finitely generated corresponds to having a classifying space with finite 1-skeleton, and being finitely presented corresponds to having a classifying space with finite 2-skeleton. Having a classifying space with finite n-skeleton is referred to as having type F_n. Having a classifying space with finitely many cells in every dimension is referred to as having type F_∞.

There is a homological version of type F_n. A group G is of type FP_n if the $\mathbf{Z}G$ module $\mathbf{Z}$ has a projective resolution that is finitely generated in dimensions no greater than n. It is of type FP_∞ if it is of type FP_n for all n. A finitely presented group is of type FP_n if and only if it is of type F_n. This follows from the material in Chapter VIII, Section 7 of Brown 1982 [32]. Non-implications, including groups groups of type FP_2 that are not finitely presented are given in Section 6.3 of Bestvina-Brady 1997 [11]. See also Section 8.3 of [80].

For a group G, we say that G has *geometric dimension* ∞ if there does not exist a finite dimensional $K(G,1)$. Otherwise the smallest integer d for which there exists a d-dimensional $K(G,1)$ is the geometric dimension of G.

22.2. Facts. If H is a subgroup of G, and G has a classifying space X of dimension d, then the cover $\widetilde{X}$ of X corresponding to H is a classifying space for H that is also of dimension d. Thus geometric dimension is "monotone" in the sense that $H \leq G$ implies that the geometric dimension of H is no greater than the geometric dimension of G. From this we have the following.

PROPOSITION 22.1. *The Thompson groups F, T and V all have geometric dimension ∞.*

PROOF. For all positive integers n, Lemma 5.8 and Corollary 13.7.1 tell us that $\mathbf{Z}^n \leq F \leq T \leq V$. The n dimensional torus $T^n = (S^1)^n$ has fundamental group $\mathbf{Z}^n$, has all other homotopy groups trivial, and has non-trivial homology in dimension n. So any classifying space for $\mathbf{Z}^n$ must have non-trivial homology in dimension n and be at least n dimensional. □

Thus the claim that F is of type F_∞ is the best one can hope for. The following two results show what we must do to prove that F is of type F_∞.

THEOREM 22.2. *A group G has type F_∞ if and only if G has type F_n for all $n \geq 1$.*

PROOF. One direction requires quite a bit of care. See the proof of Proposition 7.2.2 in [80]. □

For the next statement a non-empty space X is n-connected if $\pi_i(X)$ is trivial for all non-negative $i \leq n$. A non-empty, path connected space X is k-aspherical if $\pi_i(X) = 0$ for $2 \leq i \leq k$.

THEOREM 22.3. *A group G has type F_n if and only if there is a finite n-dimensional $(n-1)$-aspherical simplicial complex whose fundamental group is isomorphic to G. This in turn is equivalent to the existence of an $(n-1)$-connected simplicial complex on which G acts freely and simplicially with compact quotient.*

PROOF. See the proof of Proposition 7.2.1 in [80] and the remark immediately following the proof. □

In this chapter we will show (twice) that F is of type F_∞. But T and V also are of type F_∞. This will be discussed in Section 27.3.

Sections 23 and 24 are preliminaries to the constructions. Aspects of the discussion will be applicable to all the complexes built. Section 23 gives more detail about the positive monoid F_+ of F. Section 24 prepares for the complexes by building categorical structures from F_+. It also locates the cubical structures within the associated simplicial complexes. Section 25 builds the first complex for F and Section 26 completes the argument that F has type F_∞. Section 27 builds a complex for V and derives a finite presentation. Section 28 builds the third complex. Section 29 contains some final comments.

23. The positive monoid revisited

The positive monoid F_+ of F was introduced in Section 11. Here we consider F_+ separately from its existence as a submonoid of F and generalize it to a category.

A category has both objects and morphisms. The algebraic structure is in the morphisms. For us the "elements" of the algebraic structure of a category will be the morphisms.

23.1. Categories of finite and finitary forests. The finitary forests were defined in Section 11.2.3 as infinite sequences of finite, binary trees, almost all trivial, and the finite forests were defined in Section 11.6 as finite sequences of finite, binary trees. The product $\Phi\Theta$ of two finitary forests was defined by the formula (11.2), turning the set $\mathcal{F}$ of finitary forests into a monoid. The product $\Phi\Theta$, illustrated in (11.3), is obtained by hanging the tree Θ_i on the i-th leaf of Φ for each $i \in \mathbf{N}$.

23.1.1. *The category of finite forests.* If P denotes the set of finite forests, then the definition (11.2) can be used to define a partial multiplication on P. If Φ and Θ are now finite forests, then we can form $\Phi\Theta$ using (11.2) only if the number of leaves of Φ equals the number of trees (number of roots) in Θ.

As with any monoid, we can regard $\mathcal{F}$ as a category with one object ω whose morphisms are the elements of of $\mathcal{F}$. We can turn P into a category whose objects are the positive integers, and where each $P(m,n)$, the morphisms in P from m to n, consists of those forests with m trees and n leaves. To be consistent with remarks about about when the product $\Phi\Theta$ can be formed, we will compose morphisms left-to-right. That is, we compose "the wrong way" for a category.

23.1.2. *Going between the finite and the infinite.* If $\Phi \in P(m,n)$, then we write $\Phi\cdots$ to denote the element of $\mathcal{F}$ obtained by following Φ by an infinite sequence of trivial trees. We say that $\Psi \in \mathcal{F}$ is of type (m,n) if all Ψ_i are trivial for $i \geq m$ and the total number of leaves on the Φ_i, $0 \leq i < m$ is n. Clearly if $\Psi \in \mathcal{F}$ is of type (m,n) it is of type $(m+k, n+k)$ for all $k \geq 0$. And if $\Phi \in P(m,n)$, then $\Phi\cdots$ is of type (m,n).

In inverse to taking $\Phi \in P(m,n)$ to $\Phi\cdots$, we can restrict a $\Psi \in \mathcal{F}$ of type (m,n) to the first m trees and obtain the element $\Psi|_m$ of $P(m,n)$. As a function on the elements of $\mathcal{F}$ of type (m,n), taking Ψ to $\Psi|_m$ is a bijection onto $P(m,n)$. As a function on $P(m,n)$, taking $\Phi \in P(m,n)$

to $\Phi\cdots \in \mathcal{F}$ is an injection. These are homomorphisms in that if $\Phi\Theta$ is an allowable composition in P, then $(\Phi\cdots)(\Theta\cdots) = (\Phi\Theta)\cdots$ in $\mathcal{F}$.

From the previous paragraph, we know that if $\Psi = \Phi\Theta$ in P with $\Phi \in P(m,n)$ and $\Theta \in P(n,p)$, then $\Psi\cdots$ is of type (m,p) in $\mathcal{F}$ and it factors in $\mathcal{F}$ as a product of forests of types (m,n) and (n,p). We will need some reverse information, and this requires facts about the numbers.

For $\Phi \in \mathcal{F}$, let us write $\Phi \sim (m,n)$ to mean that Φ is of type (m,n) and m is the minimum for which this holds. Note that $\Phi \sim (m,n)$ implies that m is one plus the largest i so that the i-th root is not a leaf, and n is one plus the largest j so that the j-th leaf is not a root.

LEMMA 23.1. *Let Φ, Θ and $\Psi = \Phi\Theta$ be in $\mathcal{F}$ with $1 \neq \Phi \sim (m,n)$, $1 \neq \Theta \sim (p,q)$ and $\Psi \sim (r,s)$. Then for some t with $r < t < s$, we have Φ of type (r,t) and Θ of type (t,s).*

PROOF. Because none of Φ, Θ or Ψ is an identity, we have $m < n$, $p < q$ and $r < s$.

The three cases to consider are $n < p$, $n = p$ and $n > p$. If $n = p$, then $(r,s) = (m,q)$ and $t = n = p$ works.

If $n < p$, then with $k = p-n$, we have that Φ is of type $(m+k, n+k = p)$ and $\Psi = \Phi\Theta$ is of type $(m+k, q)$. But q is one more than the index of the rightmost non-root leaf of Ψ, so $m+k$ is one more than the index of the rightmost non-leaf root. So $(r,s) = (m+k,q)$ and $t = n+k = p$ works.

If $n > p$, then with $k = n-p$, we have that Θ is of type $(n = p+k, q+k)$ and $\Psi = \Phi\Theta$ is of type $(m, q+k)$. But m is one more than the index of the rightmost non-leaf root of Ψ, so $q+k$ is one more than the index of the rightmost non-root leaf. So $(r,s) = (m, q+k)$ and $t = n = p+k$ works. $\square$

23.2. Properties. The categories $\mathcal{F}$ and P have many properties that we will use. We list the definitions first. We are interested in the multiplicative (compositional) structure of the morphisms. We use M to represent an arbitrary category where all variables mentioned are morphisms. This is partly because morphism starts with M and partly because some properties are repeats of properties defined in Section 6.1.2 for monoids.

Once the properties are proven, we will be able to build groupoids of fractions from P and from $\mathcal{F}$. For $\mathcal{F}$, the result will be no suprise and will be a groupoid with one object (a group) which will be isomorphic to F.

There is a natural partial order $\leq$ on P and $\mathcal{F}$ which will carry over later to their groupoids of fractions. The order $\leq$ will dictate some of the vocabulary that we use in the properties.

DEFINITION 23.2. If $fg = h$ in P or $\mathcal{F}$, we will declare $f \leq h$.

Since P and $\mathcal{F}$ have identities, we know that $\leq$ is reflexive and since P and $\mathcal{F}$ are closed under composition, we know that $\leq$ is transitive. Since the relations in the presentation from Proposition 11.7 for $\mathcal{F}$ are length preserving, it follows that the only invertible element in $\mathcal{F}$ is the identity. Thus $\leq$ is anti-symmetric on $\mathcal{F}$ and a partial order. Since $fg = h$ in P implies that the codomain of h is strictly greater than the codomain of f unless g is a trivial forest, we have that $\leq$ is anti-symmetric on P as well and thus a partial order.

The mix of categories, order and algebra will cause certain concepts to naturally have more than one name. Each will emphasize a different aspect of the structure. We will never use the categorical terms after mentioning them once.

In the following there are occasional requirements that certain pairs of morphisms have the same domain or the same codomain. When applied to $\mathcal{F}$, this will automatically be true for all pairs. Some of the definitions that follow are identical to definitions in Section 6.1.2. We copy them here for the convenience of the reader, and also because a few of them have extra comments, adjustments, or assumptions.

DEFINITION 23.3. We say that M is *left cancellative* if $ab = ac$ always implies $b = c$. We say that M is *right cancellative* if $ab = cb$ always implies $a = c$. We say that M is *cancellative* if M is both right and left cancellative. With our left-to-right composition convention in a category, left cancellative is equivalent to all morphisms being epimorphisms, and right cancellative is equivalent to all morphisms being monomorphisms.

DEFINITION 23.4. If $c = ab$, then c is a *right multiple* of a, and a is a *left factor* of c. These are equivlent and are equivalent to $a \leq c$. Furher c is a *left multiple* of b, and b is a *right factor* of c.

DEFINITION 23.5. Two elements a and b in M have a *common right multiple* if $ac = bd$ for some c and d in M. We say that M has *common right multiples* if every pair of elements in M with the same domain has a common right multiple. Common right multiples are upper bounds under $\leq$.

DEFINITION 23.6. If elements a and b in M have a common right multiple, then we say they have a *least common right multiple* if they have a common right multiple c so that any other common right multiple d of a and b is also a right multiple of c (that is, $c \leq d$). We say that M has *conditional least common right multiples* if every a and b in M with a common right multiple also has a least common right multiple. We simply say that M has *least common right multiples* if every a and b in M with the same domain has a least common right multiple. A least common right multiple is a least upper bound, and in a cancellative category, it is a pushout.

Definition 23.7. Left multiples, (shortest) common left multiples, and having (conditional) greatest common left multiples are defined by replacing "right" by "left" and replacing "least" by "shortest" in (23.5) and (23.6), where common codomains are assumed when needed. In a cancellative category, a shortest common left multiple is a pullback. The term shortest is used since q being a left multiple of p does not translate to a relation between p and q under $\leq$.

Definition 23.8. We say that c is a *common left factor* for a and b if c is a left factor of both a and b. We say that c is a *greatest common left factor* of a and b if it is a left factor of a and b and every common left factor d of a and b is a left factor of c (that is, $d \leq c$). We say that *M has greatest common left factors* if every pair of elements in M with the same domain has a greatest common left factor. A common left factor is a lower bound, and a greatest common left factor is a greatest lower bound.

Definition 23.9. Right factors, common right factors, longest right factors and having longest common right factors are defined by replacing "left" by "right" and replacing "greatest" by "longest" in (23.8) where common codomains are assumed when needed. The term longest is used since q being a right factor of p does not translate to a relation between p and q under $\leq$.

Definition 23.10. A *unit* in M is an $x \in M$ with a $y \in M$ so that xy and yx are identities (x and y are isomorphisms in a category). We say M has *trivial units* if the only units are identities (i.e., all automorphism groups are trivial). A *length function* for M is a homomorphism to $\mathbf{N}$ whose preimage of 0 is contained in the units of M.

The strange sounding conditional properties have uses.

We now start the work of proving properties for $\mathcal{F}$ and P.

Definition 23.11. For a forest Φ (finite or finitary), we let $l(\Phi)$, the *length* of Φ, be the number of carets in Φ. In $\mathcal{F}$ this is the length of Φ expressed as a word in the generators ν_i (well defined since the relations in (9.1) are length preserving). In P this is the codomain minus the domain.

DEFINITION 23.12. For two forests Φ and Θ with the same domain, we define the union by $(\Phi\cup\Theta)_i = \Phi_i\cup\Theta_i$, the intersection by $(\Phi\cap\Theta)_i = \Phi_i\cap\Theta_i$, and containment by declaring that $\Phi \subseteq \Theta$ if and only if $\Phi_i \subseteq \Theta_i$ for all i in the domain of Φ and Θ. The operations and comparisons are operations and comparisons of finite binary trees regarded as subtrees rooted at $\emptyset$ of the complete binary tree $\mathcal{T}$.

LEMMA 23.13. *For forests, $\Phi \subseteq \Psi$ holds if and only if Ψ is a right multiple of Φ (i.e., Φ is a left factor for Ψ, and equivalently $\Phi \leq \Psi$).*

PROOF. The formula (11.2) makes it clear that a forest is contained in all its right multiples.

If $\Phi \subseteq \Psi$, then the concept is that of Lemma 8.4. Let (u_i) be the sequence of leaves of Φ indexed in left to right to right order and let $\phi(u_i)$ be such that $\Phi_{\phi(i)}$ is the tree in Φ to which u_i belongs. For each u_i let $\Theta_i = (\Psi_{\phi(i)})_{u_i}/u_i$ as given in Definition 8.3. Now $\Psi = \Phi\Theta$. □

The following gives the properties defined above that apply to $\mathcal{F}$ and P, and gives the structure of the promised items.

PROPOSITION 23.14. *Under the product/composition formula (11.2), both $\mathcal{F}$ and P have length function l as given in Definition 23.11 , trivial units, least common right multiples, longest common right factors, greatest common left factors, and conditional shortest common left multiples. The least common right multiple for Φ and Θ is $\Phi \cup \Theta$. The greatest common left factor for Φ and Θ is $\Phi \cap \Theta$. If $\Phi\Psi = \Theta\Pi$ is a common left multiple of Ψ and Π, then the shortst common left mutiple of Ψ and Π is $\Phi'\Psi = \Theta'\Pi$ where $\Phi = (\Phi \cap \Theta)\Phi'$ and $\Theta = (\Phi \cap \Theta)\Theta'$.*

PROOF. That l is a length function and that all units are trivial follow from (11.2) and from the definitions of length function and of l.

Cancellativity follows for $\mathcal{F}$ because the isomorphic F_+ is cancellative, and it follows for P because $\Phi \mapsto \Phi\cdots$ is a homomorphic injection that takes trivial forests to the trivial forest.

The claims about $\Phi \cap \Theta$ and $\Phi \cup \Theta$ follow from Lemma 23.13. Left to show are longest common right factors and conditional shortest common left multiples. Neither notion cooperates with $\leq$. The longest common right factor will have the largest number of carets. In P, the shortest common left multiple will have the highest domain.

We first consider longest common right factors. We note that the greatest common left factor $\Phi \cap \Theta$ of Φ and Θ finds the largest structure common to Φ and Θ "starting at the roots." Finding the longest common right factor does the same "starting at the leaves." The basic common structure at the leaves is a matched pair of exposed carets. Distinct exposed carets in a single forest must be disjoint. This will be used with the diamond condition to find the desired longest factor.

We start with $\mathcal{F}$. If Φ and Θ have a common right factor p, then $\Phi = ap$ and $\Theta = bp$. If now a and b have a common right factor $q \neq 1$, then qp is a "longer" common right factor of Φ and Θ. We can keep making the common right factor "longer" until it is no longer possible. The process must stop at some point because Φ and Θ have finitely many carets. If the diamond condition of Section 47 of the appendix holds, then there will be a unique longest common right factor.

So now assume that Φ and Θ have two different common right factors r and s expressed as words in the ν_i. With t the longest common suffix of r and s, We can get rid of t and assume that the last generators of r and s are, respectively, some ν_i and ν_j with $i \neq j$. So we can assume that $\Phi = f\nu_i$ and $\Theta = g\nu_i$ as well as $\Phi = f'\nu_j$ and $\Theta = g'\nu_j$.

The caret that ν_i attaches to f and g is exposed and has leaves numbered i and $i+1$ in both Φ and Θ. The caret that ν_j attaches to f' and g' is exposed and has leaves numbered j and $j+1$ in both Φ and Θ. Since $i \neq j$ and since leaves i and j are left leaves of their respective exposed carets, we must have that the two exposed carets just discussed are disjoint in both Φ and Θ. Assuming $i < j$, we have that $\nu_{j-1}\nu_i = \nu_i\nu_j$ is common right factor of Φ and Θ longer than both ν_i and ν_j. This verifies the diamond condition and shows that there is a unique longest common right factor.

That longest common right factors exist for pairs in P with common codomain follows from the existence of longest common right factors in $\mathcal{F}$ using the homomorphic injections and their inverses from the morphisms of P into $\mathcal{F}$ together with Lemma 23.1.

The proof of last point, conditional shortest common left multiples, will be deferred until after the next lemma. $\square$

The monoid $\mathcal{F}$ does not have unconditional common left multiples. The reader can take it as an exercise that ν_i and ν_j have a common left multiple if and only if $|i-j| \geq 2$.

LEMMA 23.15. *In $\mathcal{F}$ and P, if $f = pa = qb$ and $g = ra = sb$ are common left multiples of a and b, then there is a common left multiple h of a and b for which f and g are both left multiples.*

PROOF. The diagram below illustrates the hypothesis.

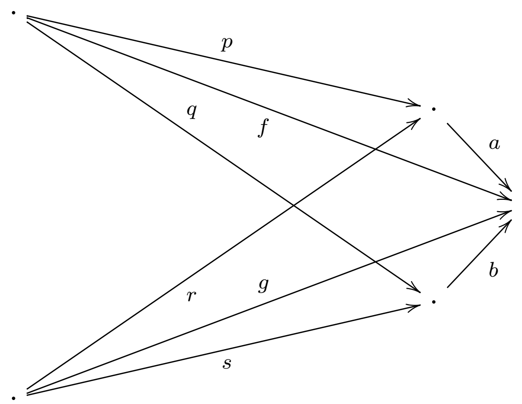

Both a and b are common right factors of f and g and we let h be the longest common right factor of f and g. For some x, y, w and z we have all of $h = xa = yb$, $f = wh$ and $g = zh$ as shown below.

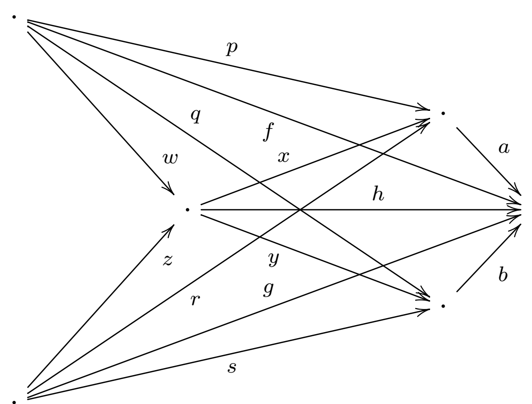

Thus $h = xa = yb$ is the desired common left multiple. ☐

PROOF OF PROPOSITION 23.14. *(Continued.)* We use the notation $\Phi\Psi = \Theta\Pi$ as well as $\Phi = (\Phi \cap \Theta)\Phi'$ and $\Theta = (\Phi \cap \Theta)\Theta'$ from the end of the statement of the proposition. That $\Phi'\Psi = \Theta'\Pi$ holds follows from

$$(\Phi \cap \Theta)\Phi'\Psi = \Phi\Psi = \Theta\Pi = (\Phi \cap \Theta)\Theta'\Pi$$

and cancellativity.

We claim that $f = \Phi'\Psi = \Theta'\Pi$ is a shortest common left multiple of Ψ and Π.

If $f = wg$ where $g = x\Psi = y\Pi$, then

$$\begin{aligned} \Phi'\Psi &= f = wg = wx\Psi, \text{ and} \\ \Theta'\Pi &= f = wg = wy\Pi, \end{aligned}$$

giving $\Phi' = wx$ and $\Theta' = wy$ by cancellation. This makes $\Phi = (\Phi \cap \Theta)\Phi' = (\Phi \cap \Theta)wx$ and $\Theta = (\Phi \cap \Theta)\Theta' = (\Phi \cap \Theta)wy$. But $(\Phi \cap \Theta)$ is the greatest common left factor of Φ and Θ, so $w = 1$ and $f = g$.

On the other hand, if $g = x\Psi = y\Pi$ and it is not assumed that f is a left multiple of g, then by Lemma 23.15 there is a common left multiple h of Ψ and Π for which f and g are both left multiples. By the above argument $f = h$ and f is a left multiple of g. □

One power of Proposition 23.14 is that the shortest common left multiple of a pair can be calculated from any common left multiple of that pair.

The following lemma is a warmup, and is not strong enough for our needs. However, it is brief and shows what we are looking for. The proof only uses the existence of least common right multiples. A more general version, Lemma 24.6, will come later and its proof will use more properties.

Lemma 23.16. *In $(\mathcal{F}, \leq)$ and in $(P, \leq)$, the intersection of two closed intervals is empty or a closed interval.*

Proof. In the following, new symbols are assumed to exist as used.

Let $[a, b]$ and $[c, d]$ be closed intervals, and assume $S = [a, b] \cap [c, d]$ is not empty. Given $e \in S$, we have $af = e = cg$, $b = eh$ and $d = ej$. Thus e is a common right multiple of a and c, and in the case of $(P, \leq)$ the domains of a and c coincide. Let $l = a \cup c$ be the least common right multiple of a and c, giving $l \leq e$. Since $e \in S$ was arbitrary, we have $S \subseteq [l, b] \cap [l, d]$ and b and d are right multiples of every element of S.

Let $m = \bigcup_{e \in S} e$ be the least common right multiple of the elements in S. Now for all $e \in S$ we have $a \leq l \leq e \leq m \leq b$ and $c \leq l \leq e \leq m \leq d$ giving both $S \subseteq [l, m]$ and $[l, m] \subseteq S$. □

23.3. Equivalent words. We discuss the set of words in the ν_i that represent a single element of $\mathcal{F}$. There will be much talk of the carets in a tree, and we use the vocabulary of trees from Section 8. For convenience we will let κ_u denote the caret in the complete binary tree $\mathcal{T}$ whose root is at the vertex u. There is a one-to-one correspondence between the carets in a tree and the internal nodes of the tree, and we could refer to the internal nodes. But carets are the main focus here.

Let Ψ be a forest in $\mathcal{F}$ and let κ_u and κ_v be carets in Ψ. We write $\kappa_u \geq \kappa_v$ if u is an ancestor of v. Specifically, κ_u and κ_v are in the same tree Ψ_i of Ψ and the unique simple path in Ψ_i from the root of Ψ_i to v passes through u. The order $\geq$ is a partial order. A linear ordering (most likely different from the order $\geq$) of the carets in Ψ is said to be consistent with $\geq$ if $\kappa_u \geq \kappa_v$ implies that κ_u comes before κ_v in the linear ordering.

Given a word $w = \nu_{i_0} \cdots \nu_{i_k}$ in the ν_i that represents a forest $\Psi \in \mathcal{F}$, we will extract from w a linear ordering of the carets of Ψ that is consistent with the partial order $\geq$. We will do this by setting up a bijection between the letters in w and the carets in Ψ and using the linear ordering of the letters given by the sequence w itself.

For $0 \leq j < k$, let $w|_j = \nu_{i_0} \cdots \nu_{i_{j-1}}$ be the prefix of w of length j and let $\Psi|_j \subseteq \Psi$ be the forest corresponding to $w|_j$. The forest $\Psi|_{j+1}$ given by $(w|_j)\nu_j$ differs from $\Psi|_j$ by a single caret κ and we associate κ to ν_j. This is a bijection from the letters in w to the carets in Ψ and we carry the order on the letters in w over to the carets in Ψ by this bijection. The linear ordering of the carets in Ψ induced from the linear order of the symbols in w and the bijection just described is consistent with $\geq$ on the carets of Ψ. This is because for each j, all the ancestors of the caret added by ν_j to $\Psi|_j$ must already be in $\Psi|_j$ and must have been introduced by generators in $w|_j$.

On the other hand, given a linear ordering of the carets in Ψ consistent with $\geq$, we can build a word in the ν_i that leads to that ordering. If we redefine $\Psi|_j$ to now be the set of carets given as the first j carets in the given linear ordering, we can claim that $\Psi|_j$ is a forest with $\Psi|_j \subseteq \Psi$. This is because the requirements of consistency with $\geq$ force $\Psi|_j$ to be prefix closed. We have the following which is left to the reader to check.

LEMMA 23.17. *For $\Psi \in \mathcal{F}$, the following hold.*

1. *Taking a word w in the ν_i representing Ψ to the linear order on the carets of Ψ induced by w gives a one-to-one correspondence from the set of words in the ν_i that represent Ψ to the set of linear orders on the carets in Ψ that are consistent with $\geq$.*
2. *The set of forests with a unique linear order on the carets consistent with $\geq$ (and thus a unique word in the ν_i representing the forest) are the forests whose only non-trivial tree is a vine. That is, those forests for which $\geq$ is already a linear order (i.e., a chain).*
3. *For each $n \geq 1$, the set of forests with exactly n carets and $n!$ linear orders consistent with $\geq$ are those forests with n carets where every non-trivial tree has exactly one caret. That is, those forests for which $\geq$ is an anti-chain.*

The forests described by (3) in Lemma 23.17 will be the subject of Sections 23.4 and 24.3.

The relations $\nu_j\nu_i = \nu_i\nu_{j+1}$ whenever $i < j$ from the presentation from Proposition 11.7 for $\mathcal{F}$ connect all the words in the ν_i representing a given element of $\mathcal{F}$. If the relation is taken as giving two possible changes $\nu_j\nu_i \to \nu_i\nu_{j+1}$ and $\nu_i\nu_{j+1} \to \nu_j\nu_i$ to adjacent symbols in a word in the generators, then these induce two possible changes in a linear ordering of the carets in a forest that are compatible with $\geq$. It is seen that such changes are reversals of the order of two carets that are adjacent in the linear ordering. We get the following consequence.

Lemma 23.18. *For $\Psi \in \mathcal{F}$ any two linear orderings of the carets in Ψ that are compatible with $\geq$ are connected by a chain of linear orderings compatible with $\geq$ in which passage from one ordering to the next is accomplished by a transposition of consecutive carets in the linear ordering, induced by a relation from the presentation from Proposition 11.7 for $\mathcal{F}$.*

23.4. Elementary elements.

Definition 23.19. In $\mathcal{F}$ or P an elementary element is a forest in which each tree has zero or one carets. In P, we use $\Delta_n : n \to 2n$ to denote the elementary forest of n trees in which every tree has exactly one caret. In $\mathcal{F}$, $\Delta_n \cdots$ (using the notation $\Phi \cdots$ of Section 23.1) is then the finitary forest in which only the first n trees are non-trivial and each of these has exactly one caret.

Lemma 23.20. *Factors of elementary elements are elementary. There are exactly 2^n left factors of $\Delta_n \cdots$ and Δ_n.*

Proof. The first sentence follows from Lemma 23.13. The second follows from the fact that $\Phi \subseteq \Delta_n$ is determined entirely by which of the n trees of Φ are non-trivial. □

The next lemma adds detail to (3) of Lemma 23.17.

Lemma 23.21. *In $\mathcal{F}$, the element $\Delta_n \cdots$ is represented by exactly $n!$ different words in the ν_i of which one is the normal form $\nu_0\nu_2\nu_4 \cdots \nu_{2n-2}$ and another of which is $\nu_n\nu_{n-1} \cdots \nu_1\nu_0$.*

Proof. The count $n!$ comes from Lemma 23.17, and it is easy to check that the claimed forms represent Δ_n. □

24. Two groupoids of fractions

We will use the conclusions of Proposition 23.14 to embed P and $\mathcal{F}$ in groupoids of fractions. The conclusions in Proposition 23.14 mimic the assumptions needed in Ore's theorem which embeds a cancellative semigroup with common right multiples into a group of fractions.

The embedding of $\mathcal{F}$ into a groupoid and the properties of that embedding will duplicate some of the results about F that were obtained in Chapter 2.

This exercise can also be viewed as applying the constructs in Gabriel-Zisman 1967 [76]. The restrictions (a–d) from Page 12 of [76] are all satsified. Restrictions (a) and (b) are satisfied because we will invert all morphisms, the existence of common right multiples is exactly (c), and cancellativity is stronger than (d). After we prove the validity of the construction we will explain the difference.

The Ore construction has only to do with the properties in the conclusions of Proposition 23.14 and little to do with the exact nature of the morphisms involved, so we axiomatize the situation. We continue the practice of composing in left-to-right order.

Definition 24.1. An *Ore category* is a category where every morphism can be cancelled on the right and the left and where any two morphisms f and g with the same domain have morphisms p and q so that $fp = gq$. A *groupoid* is a category where every morphism is an isomorphism.

The setting of the theorem below is mildly more abstract than the setting (semigroups) of the classical Ore theorem [50, Theorem 1.23], but in spite of that the proofs can be made identical. The proof below differs considerably in style from the proof in [50] but is quite traditional. Further we assume common right multiples, while [50] assumes common left multiples (there called *right reversible*).

Theorem 24.2 (Ore). *Given an Ore category C, there is a groupoid $C^{\pm}$ with the same objects as C and a functor $K : C \to C^{\pm}$ with the following properties.*

1. *K is the identity on the objects.*
2. *Every morphism in $C^{\pm}$ is of the form $K(f)(K(g))^{-1}$ for some morphisms f and g in C having the same codomain.*
3. *K is injective on the morphisms.*

Proof. We set objects of $C^{\pm}$ to be the objects of C giving (1) by definition. Morphisms of $C^{\pm}$ will be equivalence classes of pairs (f, g) with f and g morphisms of C with the same codomain. We need to define the equivalence relation $\sim$ that will be put on the pairs. If (f, g) is a pair of morphisms from C and j is a morphism of C so that the composition fj exists, then gj exists as well, and we write $(f, g) \to (fj, gj)$. We let $\sim$ be the equivalence relation generated by $\to$. To keep control of notation we will not invent a notation to distinguish between a pair and its equivalence class.

We define the domain of a pair (f,g) to be the domain of f, and the codomain of the pair to be the domain of g. Composition is opportunistic based on the special case: $(f,g)(g,h) = (f,h)$. The general case uses one of the hypotheses of an Ore category and repeats an argument used many times with F. If we want the product $(f,g)(h,k)$, we note that g and h must have the same domain, and we can take p and q with $gp = hq$ and set

$$(f,g)(h,k) = (fp, gp)(hq, kq) = (fp, kq).$$

If we use r and s with $gr = hs$ to obtain (fr, ks) for the composition, then we need to show $(fp, kq) \sim (fr, ks)$. But p and r have the same domain and there are t and u with $pt = ru$. Now

$$hqt = gpt = gru = hsu, \tag{24.1}$$

and h can be cancelled from the left to give $qt = su$. Now we get

$$(fp, kq) \to (fpt, kqt) = (fru, ksu) \leftarrow (fr, ks),$$

and composition is seen to be well defined. Associativity of composition is straighforward now that well definedness is established. It is also immediate that for an object X, the pair $(\mathbf{1}_X, \mathbf{1}_X)$ is an identity on X in $C^{\pm}$, and thus equally immediate that the inverse of (f,g) is (g,f). So $C^{\pm}$ is a category and a groupoid.

Before defining the functor K, we prove the following.

Claim 1. *If $(f,g) \sim (h,k)$, then there are morphisms p and q in P so that*

$$(f,g) \to (fp, gp) = (hq, kq) \leftarrow (h,k).$$

We discuss the relation between $\sim$ and $\to$. It is clear that $\to$ is reflexive by letting j be an appropriate identity in the definition given for $\to$, and it is straightforward that $\to$ is also transitive. Thus $(f,g) \sim (h,k)$ occurs when the two pairs are connected by a string of alternations of $\to$ and $\leftarrow$. It suffices to show that any $\leftarrow \cdot \to$ can be replaced by $\to \cdot \leftarrow$. To that end we start with $(fa, ga) \leftarrow (f,g) \to (fb, gb)$. Now there are p and q so that $ap = bq$, and we have

$$(fa, ga) \to (fap, gap) = (fbq, gbq) \leftarrow (fb, gb),$$

establishing the claim.

Now that Claim 1 is established, we point out that not only is it clear that all pairs (k,k) act as identities, but the claim shows that they are the only pairs to do so.

We start defining our functor K by setting it to be the identity on the objects as promised. Now given a morphism f in C, we let

$K(f) = (f, \mathbf{1}_X)$ with X the codomain of f. This cooperates with our definition of K on objects.

Now given $f : X \to Y$ and $g : Y \to Z$ in C, then we have

$$(f, \mathbf{1}_Y)(g, \mathbf{1}_Z) = (fg, g)(g, \mathbf{1}_Z) = (fg, \mathbf{1}_Z).$$

Item (2) is seen by noting that since a morphism (f, g) in $C^{\pm}$ has some common codomain Y for f and g, we can write $(f, g) = (f, \mathbf{1}_Y)(\mathbf{1}_Y, g) = K(f)(K(g))^{-1}$.

For the claim of injectivity, if both f and g are from X to Y and $(f, \mathbf{1}_Y) \sim (g, \mathbf{1}_Y)$, then from Claim 1 there are p and q with

$$(f, \mathbf{1}_Y) \to (fp, p) = (gq, q) \leftarrow (g, \mathbf{1}_Y)$$

forcing $p = q$, $fp = gq = gp$ and $f = g$ by canceling p on the right. □

Note that the above argument used cancellation on the left for well definedness of composition and cancellation on the right for the injectivity of K. Cancellativity on the left can be weakened and is done so in [76]. The assumption (d) on Page 12 of [76] is (after adjusting for the reversed order of composition) that if f, g, h are morphisms of C with $hf = hg$, then there is a morphism w so that $fw = gw$. Now after (24.1), we can write $qtw = suw$ and with the given $pt = ru$ we can write

$$(fp, kq) \to (fptw, kqtw) = (fruw, ksuw) \leftarrow (fr, ks).$$

This wasn't incorporated into the statement of the theorem since we don't need the generality.

Theorem 24.2 did not use several properties of P. If we take them into account, we get the following. In the following we will drop the symbol K and assume that in $C^{\pm}$ that morphiusm symbols such as f or g^{-1} indicate that f and g come from C. We also use $|f|$ for the length of a morphism f when there is a length function.

Proposition 24.3. *Assume that an Ore category C has a length function, conditional shortest common left multiples, and trivial automorphism groups. Then every morphism in the groupoid $C^{\pm}$ has a unique representation pr^{-1} with p and r in C and $|p| + |r|$ minimal. For this p and r, if $qs^{-1} = pr^{-1}$ also has q and s in C, then there is a unique morphism y in C so that $q = py$ and $s = ry$.*

Proof. Let a morphism in $C^{\pm}$ be represented as pr^{-1} with p and r from C. Since $|p|+|r|$ is finite, we can assume it is minimal as specified in the statement of the proposition. We need to show that it has the rest of the claimed properties.

Assume that $qs^{-1} = pr^{-1}$ with q and s from C. Then p and q have the same domain and r and s have the same domain. There are a and b with $pa = qb$. We have almost arrived at the situation of Lemma 23.15 and the following simplification of its first diagram which needs $ra = sb$ if we are to claim it commutes.

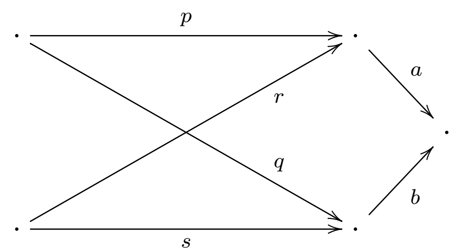

But $pr^{-1} = qs^{-1}$ gives $r = sq^{-1}p$ so $ra = sq^{-1}pa = sq^{-1}qb = sb$. So a and b have $pa = qb$ and $ra = sb$ as common left multiples, and Lemma 23.15 gives $h = xa = yb$, as well as w and z for which $p = wx$, $q = wy$, $r = zx$ and $s = zy$ all hold. This is illustrated below as a simplification of the second diagram in the proof of Lemma 23.15.

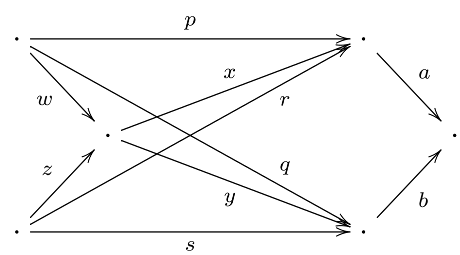

Now $wz^{-1} = (wx)(zx)^{-1} = pr^{-1}$ is another representation of pr^{-1} with w and z in C. But unless $x = 1$, we have $|w| + |z| < |p| + |r|$. So $w = p$ and $r = z$. Now $q = wy$ and $s = zy$ get rewritten as $q = py$ and $s = ry$ as desired. The uniquenss of pr^{-1} follows from assuming that $|q| + |s| = |p| + |r|$ which gives $y = 1$. □

We have arrived at the groupoids we want.

COROLLARY 24.3.1. *There is a groupoid $P^{\pm}$ whose objects are the non-negative integers and whose morphisms from m to n are equivalence classes of pairs of forests (Φ, Θ) where Φ has m trees and k leaves and Θ is a forest with n trees and k leaves. Sending a finite forest Φ with m trees and k leaves to $(\Phi, \mathbf{1}_k)$ is an injective functor that is the identity on the objects from P into $P^{\pm}$.*

There is a groupoid $\mathcal{F}^{\pm}$ with one object ω and whose morphisms are equivalence classes of pairs of finitary forests (Φ, Θ). Sending a finitary forest Φ to $(\Phi, \mathbf{1})$ is an injective functor from $\mathcal{F}$ into $\mathcal{F}^{\pm}$.

The rest applies to both groupoids.

The equivalence relation is generated by $(\Phi, \Theta) \to (\Phi\Psi, \Theta\Psi)$*. Composition is opportunistic as discussed in the proof of Theorem 24.2. For a given morphism there is a unique pair* (Φ, Θ) *characterized by having the minimun total number of carets representing the morphism so that* $(\Phi, \Theta) \to (\Phi', \Theta')$ *for any other pair* (Φ, Θ') *representing the same morphism.*

Proof. The fact that the length of a forest (finite or finitary) is the number of its carets combined with Proposition 23.14 gives all that is needed to apply Theorem 24.2 and Proposition 24.3. □

Corollary 24.3.2. *In* $P^{\pm}$*, the automorphism group of each object is isomorphic to* F*. The groupoid* $\mathcal{F}^{\pm}$ *is isomorphic as a group to* F*.*

Proof. In $P^{\pm}$, we have that $\mathrm{Aut}(1)$ consists of all pairs (Φ, Θ) where Φ and Θ are single trees with the same number of leaves under the equivalence relation generated by $\to$ as specified in Corollary 24.3.1, and with multiplication as given in that corollary. But this is F. Now for $n \geq 1$ we have $\mathrm{Aut}(n)$ isomorphic to $\mathrm{Aut}(1)$ since $P^{\pm}$ is connected. The claim about $\mathcal{F}^{\pm}$ is covered by the isomorphism between $\mathcal{F}$ and F_+ and the remarks in Section 11.2. □

From Corollary 24.3.2 we could try to treat F as a substructure of $P^{\pm}$. But in Section 24.2 we will extend $\leq$ to $P^{\pm}$, and to $\mathcal{F}^{\pm}$ regarded as the group F. It turns out that with this extra structure it is easier to keep F separate from $P^{\pm}$. The groupoid $P^{\pm}$ will see much use in spite of this separation.

24.1. Composition of finite forest pairs. If we have two composable morphisms in $P^{\pm}$, then each morphism is represented by a forest pair, and we are faced with the composition $(\Phi, \Theta)(\Psi, \Xi)$, where each entry is a finite forest. We repeat the formalities of the multiplication from Theorem 24.2.

The number of leaves of Φ and Θ must be the same and the number of leaves of Ψ and Ξ must be the same. To be composable the codomain of (Φ, Θ) or the number of trees in Θ must be the same as the domain of (Ψ, Ξ) or number of trees in Ψ. Thus $\Theta \cup \Psi$ can be formed which is a right multiple of Θ and Ψ in P. Thus for some Γ and Ω, we have $\Theta\Gamma = \Psi\Omega$. The number of trees in Γ is the number of leaves of Θ and thus also the number of leaves of Φ. The number of trees in Ω is the number of leaves of Ψ and thus also the number of leaves of Ξ. Thus we can form $(\Phi\Gamma, \Theta\Gamma)$ and $(\Psi\Omega, \Xi\Omega)$ which represent the same morphisms in $P^{\pm}$ as (Φ, Θ) and (Ψ, Ξ). Now the composition is $(\Phi\Gamma, \Xi\Omega)$.

An example follows.

(24.2) 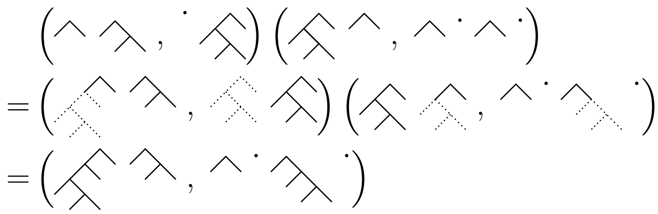

Extra carets corresponding to the forests Γ and Ω of the general discussion were temporarily shown as dotted.

There is more graphical way to compose elements of $P^{\pm}$. We make no attempt to make this rigorous. A forest pair (Φ, Θ) where the number of leaves of the two forests are the same can be regarded as a code for a PL map between intervals of various integral lengths. If the common number of leaves is k and the number of trees of Φ and Θ are m and n respectively, then Φ can be viewed as a way to split the m intervals of length 1 with integer endpoints in $[0, m]$ into k subintervals whose lengths are integral powers of 2, and Θ can be similarly viewed as splitting the interval $[0, n]$. The pair (Φ, Θ) then represents the PL homeomorphism from $[0, m]$ to $[0, n]$ that takes the k intervals determined by Φ affinely, in order, to the k intervals determined by Θ. We represent this pictorally by inverting Θ and drawing it below Φ with the leaves matched.

The result for the two factors in (24.2) is as follows. The horizontal dotted lines represent PL functions between partitions and not an identification of intervals.

(24.3) 

Reading from top to bottom, the left figure represents a PL map from $[0, 2]$ to $[0, 2]$ and the second a PL map from $[0, 2]$ to $[0, 4]$.

Since we read from top to bottom, we can compose the two figures by placing the right below the left, joining the bottom roots of the left figure with the top roots of the right figure.

(24.4) 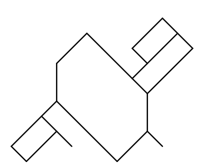

Any occurrence of 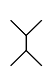 can be replaced by  since the first figure represents a merge of two intervals followed by a split of the merged interval. The composition has had no effect on the two intervals.

If this is applied to (24.4), we get the left figure below. A slight rearrangement gives the right figure below which is what is obtained if, in the final pair in (24.2), the right forest is inverted and placed below the left forest with the leaves matched.

(24.5) 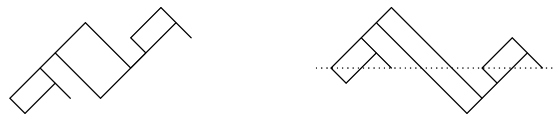

There is no occurrence of 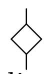 in either (24.4) or (24.5). Such a figure represents a split immediately followed by a merge of the two resulting intervals. Were the figure to appear, it could be replaced by |. This move is usually associated to a reduction of a forest pair to the minimum representative pair promised by Corollary 24.3.1.

24.2. Partial orders on the groupoids. We lift $\le$ from P and $\mathcal{F}$ to their groupoids of fractions. From this point, we want to treat the groupoid of fractions of $\mathcal{F}$ as the group F and it makes more sense notationally to refer to F as the group of fractions of its positive monoid F_+ rather than of the isomorphic $\mathcal{F}$. This practice starts in the definition of $\le$ on $P^\pm$ and F below.

Definition 24.4. For morphisms f and g in $P^\pm$, we say $f \le g$ if for some $h \in P$ we have $fh = g$. That is, $f \le g$ if $f^{-1}g$ is in P. For f and g in F, we say $f \le g$ if $f^{-1}g$ is in F_+.

Since P and $\mathcal{F}$ have only trivial units, $\le$ is a partial order. We have the following. Again the requirement of a common domain is always satisfied in F.

Lemma 24.5. *If f and g have a common domain in $(P^\pm, \le)$ or $(F, \le)$ then they have an upper bound. If f and g also have a lower bound h, then they have a least upper bound M and a greatest lower bound m. Specifically, if $\Omega_1 = h^{-1}f$ and $\Omega_2 = h^{-1}g$, then Ω_1 and Ω_2 have a common domain, and $M = h(\Omega_1 \cup \Omega_2)$ and $m = h(\Omega_1 \cap \Omega_2)$.*

Proof. For the first sentence, we have $f = \Phi_1\Theta_1^{-1}$ and $g = \Phi_2\Theta_2^{-1}$ with the Φ_i and Θ_i in F_+, or in P so that the Φ_i have a common domain. Now Proposition 23.14 gives $\Phi_1 \cup \Phi_2 = f\Theta_1 \cup g\Theta_2$ as a common upper bound for $f\Theta_1$ and $g\Theta_2$ and thus for f and g.

Let $f\Psi_1 = g\Psi_2$ with the Ψ_i in F_+ or P be an upper bound for f and g and assume h is a lower bound for f and g. Then $f = h\Omega_1$, $g = h\Omega_2$ with the Ω_i in F_+ or P. Cancellativity and $h\Omega_1\Psi_1 = h\Omega_2\Psi_2$ give $\Omega_1\Psi_1 = \Omega_2\Psi_2$. If j is another lower bound for f and g with $f = j\Omega_3$, $g = j\Omega_4$, then we will also have $\Omega_3\Psi_1 = \Omega_4\Psi_2$ with all Ω_i in

F_+ or P. In other words, all lower bounds for f and g create common left multiples of Ψ_1 and Ψ_2 in F_+ or P.

Proposition 23.14 now gives Ψ_1 and Ψ_2 a shortest common left multiple $\Omega_1'\Psi_1 = \Omega_2'\Psi_2$ in F_+ or P where $\Omega_1 = (\Omega_1 \cap \Omega_2)\Omega_1'$ and $\Omega_2 = (\Omega_1 \cap \Omega_2)\Omega_2'$. We claim that $m = h(\Omega_1 \cap \Omega_2)$ is the greatest lower bound for f and g. Taking j as analyzed in the previous paragraph as a typical other lower bound for f and g, the situation creates another common left multiple $\Omega_3\Psi_1 = \Omega_4\Psi_2$ for Ψ_1 and Ψ_2, so for some x in F_+ or P we have that $\Omega_3\Psi_1 = x\Omega_1'\Psi_1$ giving $\Omega_3 = x\Omega_1'$. Now

$$h(\Omega_1 \cap \Omega_2)\Omega_1' = h\Omega_1 = f = j\Omega_3 = jx\Omega_1'$$

so $jx = h(\Omega_1 \cap \Omega_2) = m$ and $j \le m$.

Similar arguments make $M = h(\Omega_1 \cup \Omega_2)$ the least upper bound for f and g. □

Corollary 24.5.1. *If S is a finite set of morphisms in F or $P^\pm$ that all have the same domain, then S has an upper bound in $(F, \le)$ or $(P^\pm, \le)$. If every pair of elements in S also has a lower bound in $(F, \le)$ or $(P^\pm, \le)$, then S has least upper bound and a greatest lower bound in $(F, \le)$ or $(P^\pm, \le)$.*

Proof. This follows inductively from Lemma 24.5. □

We now get the version of Lemma 23.16 that we need. Only slight changes are needed in the wording of the proof.

Lemma 24.6. *In $(F, \le)$ or $(P^\pm, \le)$, the intersection of two closed intervals is empty or a closed interval.*

Proof. Let $[a, b]$ and $[c, d]$ be closed intervals, and assume $S = [a, b] \cap [c, d]$ is not empty. Given $e \in S$, we have $af = e = cg$ with f and g in F_+ or P. This gives all of a, e and c the same domain which then must be the domain of all in S. Now with M the least upper bound of S and m the greatest lower bound, we get $S = [m, M]$ exactly as in the proof of Lemma 23.16. □

24.3. Cubes. The simplicial complexes (see Section 48 in the appendix) associated to the posets $(F, \le)$ and $(P^\pm, \le)$ have subcomplexes that organize themselves neatly into cubes. These arise from a restriction of the order $\le$ that is based on the elementary forests of Section 23.4.

Definition 24.7. For morphisms f and g in F or in $P^\pm$ with common domain, we write $f \preceq g$ if $f^{-1}g$ is an elementary forest.

Since $f \le g$ holds if $f^{-1}g$ is in F_+ or P, the relation $\preceq$ is a restriction of $\le$ and is thus anti-symmetric. It is reflexive since the trivial forest is elementary. However, $\preceq$ is not transitive.

LEMMA 24.8. *Let $f \preceq g$ in F or $P^{\pm}$ and let $h = f^{-1}g$. Then the closed interval $[f,g] = f[1,h] = \{fk \mid 1 \preceq k \preceq h\}$ under $\preceq$ equals the closed interval $[f,g]$ under $\le$. In particular $f \preceq j \preceq g$ if and only if $f \le j \le g$. If h has n carets, then $[f,g]$ has the structure of an n-cube whose faces are the closed intervals in $[f,g]$.*

PROOF. The equality of $[f,g]$ under $\le$ and $\preceq$ follows from Lemma 23.20. Let Supp(h), the support of h, be the set of those $i \in \mathbf{N}$ for which h_i is a non-trivial tree. Taking $j \in [1,h]$ to Supp(j) is a bijection from $[1,h]$ to the subsets of Supp(h). From Lemma 23.13, this bijection is an order isomorphism from $[1,h]$ to the subsets of Supp(h) ordered by inclusion. As remarked in Section 48.9 of the appendix, this gives $[1,h]$ the structure of an n-cube. The remark about faces also comes from Section 48.9. □

LEMMA 24.9. *The set of closed intervals under $\preceq$ gives $(F,\preceq)$ and $(P^{\pm},\preceq)$ the structure of a cubical complex.*

PROOF. The union of the closed intervals is all of $P^{\pm}$. It follows from Lemma 24.6 and from the equality of $\le$ and $\preceq$ inside a closed interval $[f,g]$ with $f \preceq g$ that the intersection of two cubes is empty or a face of each. □

24.4. Faces and generators. We give some motivation for the discussion in this section. The material here will not see use until Section 28. In Section 28, we will claim that the cubical complex associated to $(F,\preceq)$, regarded as a subcomplex of the complex associated to $(F,\le)$, is the universal cover of a classifying space Y for F. Each cube $f[1,h]$ in the complex of $(F,\preceq)$ will be identified, under the action of F, to the cube $[1,h]$ in Y. We are thus interested in the structures of cubes of the form $[1,h]$ and their faces. However, each face $[j,k]$ of a cube $[1,h]$ must be understood in the form $j[1,j^{-1}k]$ to see how it fits into the complex Y.

A cube of the form $[1,h]$ with $1 \preceq h$ has h elementary. Since we are working in F we should think of h as an element of the positive monoid F_+. However we want to discuss placements of carets that correspond to specific generators, and it is easiest to think of h as a finitary forest. Thus we will work with the generators ν_i of the presentation from Proposition 11.7 for the monoid $\mathcal{F}$.

We consider an elementary element h. If h has k carets, then f can be given as a word w in the ν_i of length k in normal form. It is easy to characterize words in normal form that correspond to elementary forests. What follows adds information to Lemma 23.21.

If $w = \nu_{i_0}\nu_{i_1}\cdots\nu_{i_{k-1}}$ represents h, it is in normal form if the subscripts are in non-decreasing order. If $p = \nu_{i_0}\nu_{i_1}\cdots\nu_{i_j}$, $j < k-1$, is a proper prefix of w, then we can look at the forest h_p corresponding to p and discuss the placement on h_p of the caret corresponding to $\nu_{i_{j+1}}$. The numbers of the leaves of the caret placed by ν_{i_j} are i_j and i_j+1. With $i_{j+1} \ge i_j$, the caret placed by $\nu_{i_{j+1}}$ is disjoint from the caret placed by ν_{i_j} if and only if $i_{j+1} \ge i_j+2$. The caret corresponding to $\nu_{i_{j+1}}$ will disjoint from all carets in h_p if $i_{j+1} \ge i_j+2$ because i_j is the largest subscript in p. We have the following.

Lemma 24.10. *A word $w = \nu_{i_0}\nu_{i_1}\cdots\nu_{i_{k-1}}$ in normal form in the generators of $\mathcal{F}$ represents an elementary forest if and only if for all j with $0 \le j < k-1$ we have $i_{j+1} \ge i_j+2$.*

We look at some of the faces of $[1,h]$ and start with the codimension 1 faces. If h has k carets, then $[1,h]$ is a k-dimensional cube and has $2k$ codimension 1 faces. Let j be such that $0 \le j < k$. If $w = \nu_{i_0}\nu_{i_1}\cdots\nu_{i_{k-1}}$ in normal form represents h, then the appearance of ν_{i_j} in w corresponds to the j-th caret in h ordered from left to right with the count starting at 0. One of the faces of $[1,h]$ "orthogonal to" coordinate j consists of all left factors of h that omit the j-th caret (the *lower face*), and the other consists of all left factors of h that include the j-th caret (the *upper face*).

The lower face has as maximal element h_j° consisting of h with the j-th caret omitted and has 1 as minimal element. So the lower face has the form $[1,h_j^\circ]$. The upper face has as maximal element h and has the forest containing only the j-th caret as minimal element. This forest is some ν_p so that with $h = \nu_p h_j'$ and the upper face has the form $[\nu_p, h] = \nu_p[1, h_j']$. We need to obtain the value of p and to express h_j° and h_j' in terms of the ν_i.

We bring in ideas surrounding Lemma 23.18 to change when the j-th caret is added to the forest h. We are interested in two words that are not necessarily in normal form that represent h. We start with $w = \nu_{i_0}\nu_{i_1}\cdots\nu_{i_j}\cdots\nu_{i_{k-1}}$ in normal form and apply relations from the presentation from Proposition 11.7 for $\mathcal{F}$.

First we use the relations from Proposition 11.7 in the form of rewriting rules $\nu_m\nu_{n+1} \to \nu_n\nu_m$ when $m < n$ to "move" ν_{i_j} to the end of the word. This can be done because each i_n with $n > j$ exceeds i_j by at least two. We get that

$$\nu_{i_0}\nu_{i_1}\cdots\nu_{i_{j-1}}\nu_{i_{j+1}-1}\nu_{i_{j+2}-1}\cdots\nu_{i_{k-1}-1}\nu_{i_j} = w_j^\circ\nu_{i_j}$$

also represents h. The subword w_j° represents h_j° since the forest it represents has all the carets in h except the j-th caret.

Similarly, we can use rewriting rules $\nu_n \nu_m \to \nu_m \nu_{n+1}$ when $m < n$ to "move" ν_{i_j} to the beginning of the word. Again this can be done because the differences between consecutive subscripts in the normal form are at least two. We get that

$$\nu_{i_j - j}\nu_{i_0}\nu_{i_1}\cdots\nu_{i_{j-1}}\nu_{i_{j+1}}\cdots\nu_{i_{k-1}} = \nu_{i_j - j} w_j'$$

also represents h.

The lower face is $[1, h_j^\circ]$ where w_j° represents h_j°. The upper face is $\nu_{i_j-j}[1, h_j']$ where w_j' represents h_j'. The upper face is located by its least vertex v_{i_j-j} and its structure is given by $[1, h_j']$.

We now simplify the notation and invent two "face operators." The word $w = \nu_{i_0}\nu_{i_1}\cdots\nu_{i_{k-1}}$ in normal form representing an elementary forest, can be coded by the k-tuple $\mathbf{i} = (i_0, \cdots, i_{k-1})$ where $i_{j+1} \ge i_j + 2$ for $0 \le j < k-1$ is assumed. We call such a tuple a *cubical* tuple. We define "face operators" as follows.

$$\begin{aligned} A_j(\mathbf{i}) &= A_j(i_0, \cdots, i_{k-1}) = (i_0, \cdots, i_{j-1}, i_{j+1}, \cdots, i_{k-1}), \\ B_j(\mathbf{i}) &= B_j(i_0, \cdots, i_{k-1}) = (i_0, \cdots, i_{j-1}, i_{j+1} - 1, \cdots, i_{k-1} - 1) \end{aligned} \tag{24.6}$$

Now A_j codes the structure (but not the location) of the upper face, and B_j codes the structure of the lower face.

There are other useful modifications to cubical tuples that will be seen to arise naturally. Given a cubical k-tuple $\mathbf{i}$, we let $\widehat{\sigma}(\mathbf{i})$ be the k-tuple obtained from $\mathbf{i}$ by raising each entry in $\mathbf{i}$ by 1.

We define another modification on cubical k-tuples in two steps. We will use $0 \wedge \mathbf{i}$ to denote the $(k+1)$-tuple whose first entry is 0 and whose remaining entries are copied in order from the entries from $\mathbf{i}$ and we let $\widehat{H}(\mathbf{i}) = 0 \wedge \widehat{\sigma}^2(\mathbf{i})$.

LEMMA 24.11. *If* $\mathbf{i}$ *is a cubical tuple, then all of* $A_j(\mathbf{i})$, $B_j(\mathbf{i})$, $\widehat{\sigma}(\mathbf{i})$, *and* $\widehat{H}(\mathbf{i}) = 0 \wedge \widehat{\sigma}^2(\mathbf{i})$ *are cubical tuples.*

The modifications $\widehat{\sigma}$ and $\widehat{H}$ give a way to systematically run through all the cubical tuples. Let $C_1 = \{(0)\}$. If C_i is defined let $C_i' = \{\widehat{\sigma}^i(C_i) \mid i \ge 1\}$. If C_i' is defined let $C_{i+1} = \widehat{H}(C_i') = \{\widehat{H}(\mathbf{i}) \mid \mathbf{i} \in C_i'\}$. Inductively, we have the following.

LEMMA 24.12. *For each* $n \ge 1$, C_n *is the set of cubical* n*-tuples that have 0 as the first entry, and* C_n' *is the set of cubical* n*-tuples that have a first entry greater than 0.*

25. The Brown, Stein and Farley complexes

25.1. The Brown complex. In this section, we build a complex for F. It is built in stages, starting with a simplicial complex, then moving to a smaller simplicial complex having the same homotopy type, and then transforming the second complex into a cubical complex. In Section 26, the last complex will be analyzed to show that F is of type F_∞. The first complex will be a simplicial complex derived from a partial order on a category. See Section 48.4.

We will work with the groupoid $P^\pm$. We regard P as a subcategory of $P^\pm$ identifying it with the image of the functor K from Theorem 24.2. If we want to refer to a morphism in $P^\pm$ we will use traditional lower case Roman letters such as f, g, etc., but we will continue to use upper case Greek letters such as Φ, Θ, etc., for elements of P. When we want the details of an element of $P^\pm$ as a "fraction" of elements of P, we will write $f = \Phi\Theta^{-1}$ instead of $f = (\Phi, \Theta)$ or $f = K(\Phi)(K(\Theta))^{-1}$.

For each positive integer n, we have a group $\mathrm{Aut}(n)$, and since the groupoid is connected there are (non-canonical) isomorphisms between these automorphism groups. From Corollary 24.3.2, we know that all these groups are isomorphic to F.

For positive integers m and n we let $P^\pm(m, n)$ be the set of morphisms in $P^\pm$ from m to n. Let $P^\pm_1$ be all the morphisms in $P^\pm$ with domain 1. Identifying F with $\mathrm{Aut}(1)$, we have a left action of F on $P^\pm_1$ by $f \cdot g = fg$ for which each $P^\pm(1, j)$ is invariant.

We turn our attention to the restriction to $P^\pm_1$ of the partial order $\leq$ from Definition 24.4 in which $f \leq g$ if $f^{-1}g$ (which must exist since both f and g have domain 1) is in P. That is $f \leq g$ if $f^{-1}g$ is a finite forest. Note that if for f and g in $P^\pm_1$ we have $f \leq g$, then $\mathrm{codomain}(g) \geq \mathrm{codomain}(f)$. The next lemma and corollary restate Lemma 24.5 and its corollary for $P^\pm_1$.

Lemma 25.1. *If f and g are in $(P^\pm_1, \leq)$ then they have an upper bound. If f and g also have a lower bound h, then they have a least upper bound M and a greatest lower bound m. Specifically, if Θ_1 and Θ_2 are the unique elements in P where $f = h\Theta_1$ and $g = h\Theta_2$, then $M = h(\Theta_1 \cup \Theta_2)$ and $m = h(\Theta_1 \cap \Theta_2)$.*

Corollary 25.1.1. *If $f_1, f_2, \ldots, f_n$ are in $(P^\pm_1, \leq)$, then they have a common upper bound. If they also have a common lower bound, then they have a common least upper bound and a common greatest lower bound.*

As remarked just before Lemma 23.15, there are pairs of elements in $(P^\pm_1 \leq)$ that have no lower bound.

The first version of a complex that we are interested in is the simplicial complex $(P_1^{\pm}, \leq)$. It has the following properties.

LEMMA 25.2. *The action of F by left multiplication on $(P_1^{\pm}, \leq)$ is simplicial, free, keeps invariant each $P^{\pm}(1, j)$ and is transitive on each $P^{\pm}(1, j)$. Further $(P_1^{\pm} \leq)$ is contractible.*

PROOF. The action of F is on the left and the order structure is given by an "action" of P on $P_1^{\pm}$ on the right. The left and right actions commute and the left action of F is thus order preserving and simplicial. From now on, a reference to an action is to the left action unless stated otherwise.

Since all elements in $P^{\pm}$ are invertible and the action involves multiplication in $P^{\pm}$, the action is free on the vertices. Since order is preserved, the vertices of a simplex cannot be permuted non-trivially by the action. So the action is free on the entire topological realization of the complex. Multiplication by F on the left preserves codomain and so each $P^{\pm}(1, j)$ is kept invariant. And if f and g come from the same $P^{\pm}(1, j)$, then $gf^{-1} = h \in F$ and $g = hf$.

To show that $(P_1^{\pm}, \leq)$ is contractible, it suffices to show that every finite subcomplex is contained in a cone. This follows if we show that every finite set of vertices has an upper bound. But elements in $P_1^{\pm}$ all have common domain and Corollary 24.5.1 gives an upper bound. □

The quotient of $(P_1^{\pm}, \leq)$ under the action of F is not compact. In fact it has infinitely many vertices. Our task is to find subcomplexes with compact quotient with sufficient connectivity properties to satisfy the reqirements laid out in Theorem 22.3. It is possible to do so starting with $(P_1^{\pm}, \leq)$, but it is easier to do so if we start with a smaller complex that is also contractible.

25.2. The Stein complex. We will cut down on the simplices in $(P_1^{\pm}, \leq)$ by cutting down on the partial order. We restrict the relation $\preceq$ from Definition 24.7 to $P_1^{\pm}$. This means that for f and g in $P_1^{\pm}$ we have $f \preceq g$ if $f^{-1}g$ is in P and is elementary (Definition 23.19). The relation $\preceq$ is reflexive and anti-symmetric, but not transitive. We give the complex below its own letter since it will be the complex we work with.

DEFINITION 25.3. The complex X has $P_1^{\pm}$ as its vertex set and for $0 \leq k$, a chain $f_0 < f_1 < \cdots < f_k$ will be a k-simplex in X if $f_0 \prec f_k$.

Since factors of elementary forests are elementary by (2) of Lemma 23.20, we have $f_i \prec f_j$ for all $0 \leq i < j \leq k$ in the above definition. The set of simplices in X is closed under taking faces and forms a simplicial complex.

We are now working with two relations $\le$ and $\preceq$ with only the first a partial order. We will use interval notation such as $[f,g]$, $(f,g]$, (f,g), etc., frequently and in all cases the order referred to will be $\le$. However, the properties of $\preceq$ let us say the following.

Lemma 25.4. *If $f \preceq g$ in $P_1^{\pm}$, then for all h and j in $[f,g]$ if $h \le j$, then $h \preceq j$.*

We need the following concept in order to show that X is contractible.

Definition 25.5. A *greatest elementary left factor* of a forest Ψ is an elementary left factor $E(\Psi)$ of Ψ so that every elementary left factor of Ψ is a left factor of $E(\Psi)$.

The reader can verify that for every finite or finitary forest Ψ, there is a unique greatest elementary left factor $E(\Psi)$ and that for each i, $(E(\Psi))_i$ is trivial when Ψ_i is trivial, and $(E(\Psi))_i$ consists of a single caret when Ψ_i is non-trivial.

Proposition 25.6. *The inclusion of X into $(P_1^{\pm}, \le)$ is a homotopy equivalence, and thus the complex X is contractible.*

Proof. The proof that follows is taken directly from Brown 1992 [35] where certain key elements are attributed to Stein 1992 [183].

We will rebuild $(P_1^{\pm}, \le)$ from X by adding to X simplices that are eliminated from $(P_1^{\pm}, \le)$ by the defining restrictions of $\preceq$, and will do so in a way that does not change the homotopy type of X. Note that if $f \preceq g$, then $[f,g]$ lies in X. So we are interested in adding to X intervals of the form $[f,g]$ where $f < g$ but $f \not\preceq g$.

Recall that $E(\Psi)$, the greatest elementary left factor of Ψ, has as many carets as Ψ has non-trivial trees.

Fix $f \in P_1^{\pm}$. For any h in $P_1^{\pm}$ with $f < h$, we have $h = f\Psi$ for some $\Psi \in P$. We let $h' = fE(\Psi)$ giving $f \preceq h'$. We note that $h' \le h$ and that $[f,h']$ is in X. Also with $f < h$, we have Ψ not trivial. So $E(\Psi)$ is not trivial and $f \prec h'$.

Next we show that the open interval (f,g) is contractible as a complex if $f < g$ but $f \not\preceq g$. First we map each $h \in (f,g]$ to h'. Let $(f,g)'$ denoted the image of restriction of this map to (f,g). Since $h' \le h$ for all h, this restriction to (f,g) is homotopic to the identity on (f,g). See Section 48.4. For $h \in (f,g)$, we have $g = h\Theta = f\Psi\Theta$. Since $E(\Psi)$ is an elementary left factor of $\Psi\Theta$ it is a left factor of the greatest elementary left factor of $\Psi\Theta$ and we see that $h' \le g'$. But $g' < g$ since $f \not\preceq g$, and

$(f,g'] \subseteq (f,g)$. So the image $(f,g)'$ of (f,g) lies in $(f,g'] \subseteq (f,g)$, and $(f,g']$ has a maximal element and is contractible.

We now attach intervals in $(P_1^{\pm}, \leq)$ that are missing from X.

If an interval $[f,g]$ is not in X, then $f \not\preceq g$. Let the codomains of f and g be, respectively, p and q. We know that $f^{-1}g$ is in P, and we must have $q-p>1$ since $f \not\preceq g$. We call $d=q-p$ the height of $[f,g]$ and we induct on d. Thus we assume that d is at least 2 and that the addition to X of all closed intervals of height less than d has already been done and has not changed the homotopy type of X. When we add the interval $[f,g]$ regarded as a complex, the part that has already been added to X by our inductive assumption is the complex $[f,g) \cup (f,g]$. But this is contractible since it is the suspension of (f,g) and we have shown (f,g) to be contractible. Thus the contractible $[f,g]$ is being attached along a contractible complex which results in no change in the homotopy type. This completes the proof. □

The left action of F preserves $\preceq$ for the same reason that it preserves $\leq$, and so F acts on the complex X. The quotient is still not compact since X contains all the vertices of $(P_1^{\pm}, \leq)$. We obtain complexes with finite quotients by restricting to subcomplexes of X. At this point we lose contractibility, but from Theorem 22.3, contractibility is not necessary. However, the contractibilty of X will be used to establish the connectivity conditions of its subcomplexes. The following proposition shows how.

Lemma 25.7. *Let $X_0 \subseteq X_1 \subseteq \cdots$ be a sequence of complexes whose union is a contractible complex X. Fix $n \geq 0$. If for some i and all $j \geq i$ the inclusions $X_j \hookrightarrow X_{j+1}$ induce injections on π_n, then $\pi_n(X_i)=0$.*

Proof. Every singular n-sphere in X_i contracts in X and thus in X_j for some $j>i$. But $\pi_n(X_i) \to \pi_n(X_j)$ is an injection. □

We filter X into subcomplexes X_i where X_i consists of all morphisms in X with codomain no bigger than i. We have the following which satisfies part of the requirements of Theorem 22.3.

Lemma 25.8. *For each $i \geq 1$, the action of F on X preserves X_i and the restriction of the action to X_i is free, simplicial and with compact quotient. As extra information, the dimension of the quotient $F\backslash X_i$ is no more than $\lfloor i/2 \rfloor$.*

Proof. From Lemma 25.2 the action of F on $(P_1^{\pm}, \leq)$ is free, simplicial, preserves codomains, and is transitive on each set of morphisms with the same codomain. So the X_i are preserved, the action on each

X_i is free and simplicial, and the orbit space $F\backslash X_i$ has only finitely many vertices. We identify the vertices of $F\backslash X_i$ with the elements in $[1, i]$. Because the action is free and simplicial, we know that for each $f \in X_i$ with codomain n, there is a bijection between the simplices in X_i that have f as a vertex and the simplices in $F\backslash X_i$ that have n as a vertex. But in X_i there is a severe restriction on the simplices that a vertex can belong to.

Let f be some f_i in a chain $f_0 \prec \cdots \prec f_k$ from Definition 25.3, and assume f has codomain n. Then $f_0 = f\Phi^{-1}$ and $f\Psi = f_k$ for some elementary Φ and Ψ. There are only finitely many elementary elements with codomain n and only finitely many with domain n. Thus f can be a vertex in only finitely many simplices. But $F\backslash X_i$ has only finitely many vertices, so $F\backslash X_i$ is a finite complex and compact.

A simplex σ in X_i with least vertex f_0 of codomain m must be contained in some $[f_0, f_0\Psi]$ with Ψ elementary. The dimension d of σ will be no more than the number of carets k in Ψ. The number of carets k in the elementary Ψ can be no greater than the codomain m of f_0. So $d \le k \le m$. The codomain of the largest vertex in σ must have codomain no larger than i. So $m + k \le i$. Whether one assumes that $m \le \lfloor i/2 \rfloor$ or $m > \lfloor i/2 \rfloor$, it follows that $k \le \lfloor i/2 \rfloor$. $\square$

We will satisfy the remaining requirement of Theorem 22.3 if we show that for each $n \ge 0$ there is an i so that X_i is $(n-1)$-connected. Lemma 25.7 will be the key to controlling the homotopy groups of the X_i. Since Lemma 25.7 asks to understand the relationship between X_j and X_{j+1}, simplicial Morse theory will be useful. This is discussed in Section 48.8 where it is pointed out that an understanding of the descending links of vertices with respect to a Morse function will be important. However, before we build a Morse function and investigate descending links, we transform X from a simplicial complex to a cubical complex. This makes the analysis of descending links easier. The analysis will be done in Section 26.

25.3. The Stein-Farley complex. The complex X of the previous section also has the structure of a cubical complex. Recall from Definition 23.19 that Δ_n denotes the elementary forest with n carets, domain n, and codomain $2n$. We start with the following.

LEMMA 25.9. *If f is in $P_1^{\pm}$ with codomain n, then for all $g \in P_1^{\pm}$ with $f \preceq g$, we have $g \preceq f\Delta_n$. Further, if the codomain of g is $n + k$, then $k \le n$, and the interval $[f, g]$ has the structure of a k-cube.*

PROOF. Every g with $f \preceq g$ has the form $g = f\Psi$ with $\Psi = f^{-1}g$ elementary of domain n. Thus Ψ is a left factor of Δ_n. From Lemma

23.20, $g \preceq f\Delta_n$. The second claim follows from $\Psi \preceq \Delta_n$, and from the last provision of Lemma 24.8. $\square$

LEMMA 25.10. *The collection* $C = \{[f, g] \mid f, g \in P_1^{\pm},\ f \preceq g\}$ *gives* X *the structure of a cube complex which is preserved by the action of* F.

PROOF. That the collection C gives X the structure of a cubical complex is argued as for Lemma 24.9. It has already been observed that the action of F preserves $\preceq$ and so the action of F preserves the cube structure. $\square$

26. Analyzing links

At the end of Section 25.2, subcomplexes X_i of X are defined and it is pointed out that to arrive at the F_∞ properties of F, we must understand the connectivity properties of descending links. From Theorem 48.3 we also see that to establish that a complex is CAT(0), it suffices to understand combinatorial properties of links. We start by descibing the structure of the full links of vertices in X. As mentioned at the end of Section 25.2, the links will be viewed within the cubical structure of X.

26.1. The links. Let f be a vertex of X. It is a morphism in $P^{\pm}$ from 1 to n. Specifically f is a pair (T, Θ) where T is a tree and Θ is a forest with n roots. The roots are ordered and we think of them as numbered from 0 through $n-1$. Other than the ordering and numbering of the roots of Θ, the internal structure of f has little to do with the following discussion. In fact, it will be less confusing if we simply refer to the n roots of Θ as the *ends* of f and forget about Θ.

The vertices of the link of f in X correspond to those vertices $g \neq f$ of X for which there is a cube C having both g and f as vertices where g and f are the vertices of a 1-face of C. To save words in the discussion that follows, we will refer to such a g as a *neighbor* of f . A set A of neighbors of f is the vertex set of a simplex in the link of f in X if a single cube verifies that all the elements of A are neighbors of f. We give the details of the relationships just mentioned.

Assume that f belongs to a cube C which we may as well assume is maximal and has the form $[f_0, f_0\Delta_k]$ where f_0 has codomain $k \leq n$. That makes $f = f_0\Phi$ for some $\Phi \preceq \Delta_k$ with $n-k$ carets. That is, Φ is an elementary forest with k roots and n leaves where the n leaves correspond to the ends of f.

A g in C that is a neighbor of f is of the form $g = f_0\Psi$ where Φ and Ψ differ from each other by a single caret. Either Ψ is obtained from Φ

by removing one caret of Φ, or by adding one caret to Φ from Δ_k that is missing from Φ. Since the leaves of carets in Φ are at consecutive positions in the ends of f, we see that a neighbor of f in C is identified by a consecutive pair of ends of f occupied by leaves of Φ or by a single end of f not occupied by a leaf of Φ.

Thus the set of neighbors of f in C correspond to the elements of a partition of the n ends of f where each element of the partition is either a consecutive pair of ends of f or a single end of f. We argue in the other direction that every such partition of the ends of f corresponds to a cube having f as a vertex.

If S is a partition of the ends of f into singletons and consecutive pairs, then we build a cube that contains f corresponding to S. Let S_1 be the set of singletons in S and S_2 be the set of pairs in S.

Let h be morphism formed from f by attaching the leaves of a caret to each pair in S_2. If $|S_2| = p$, then p pairs of ends of f have been combined so that h has $n-p$ ends and is a morphism in $P^{\pm}$ with codomain $n-p$. Note that $2p \leq n$. Specifically, if Φ is the elementary forest with $n-p$ trees and p carets arranged so that the positions of the leaves of the carets correspond to the elements of S_2, then $f = h\Phi$ or $h = f\Phi^{-1}$. The discussion and pictures in Section 24.1 should help visualize the product.

Let j be the morphism formed from f by attaching the root of a caret to each element of S_1. We must have $q = |S_1| = n - 2p$ and j has $n+(n-2p) = 2(n-p)$ ends and is a morphism whose codomain is twice that of h. If Ψ is the elementary forest with n trees and q carets whose roots are located at the positions given by S_1, then $j = f\Psi$. The discussion and pictures in Section 24.1 should help visualize the product.

The arrangement guarantees that $\Phi\Psi = \Delta_{n-p}$ and $[h,j] = [h, h\Delta_{n-p}]$ is the cube containing f corresponding to the partition S.

Thus a neighbor of f corresponds to an "attaching site" that is either a pair of consecutive ends of f or a single end of f. A set of neighbors bounds a simplex in the link of f in X if and only if their attaching sites are pairwise disjoint. This leads immediately to the CAT(0) discussion so we will take that up next.

26.2. The CAT(0) property. We continue the notation and discusion of the previous section.

Theorem 48.3 says that a simply connected cubical complex is CAT(0) if and only if its vertex links are flag complexes. If A is a set of neighbors of f and every pair u and v of vertices in A bounds a 1-simplex in the link at f, then the attaching sites for u and v must be disjoint. Since the attachiing sites associate to all the elements of A are pairwise disjoint, the vertices in A are the vertices of a single simplex in the link. This shows that all links are flag complexes.

From Theorem 48.3 we know that X is CAT(0), and from Proposition 25.6 (or Theorem 48.4) we know that X is contractible. We have shown the following.

THEOREM 26.1. *Thompson's group F acts freely and cellularly on a contractible, cubical CAT(0) complex.*

26.3. Finiteness properties. Recall that the subcomplex X_i of X is defined at the end of Section 25.2 as all morphisms in X with codomain no bigger than i. As mentioned X acts on each X_i with compact quotient. We wish to show that the X_i satisfy the hypotheses of Lemma 25.7. Specifically, for each $n \geq 0$ we want to show that for some i and all $j \geq i$ the inclusions $X_j \to X_{j+1}$ induce isomorphisms on π_n.

The function ρ on the morphisms of $P_1^{\pm}$ that takes a morphism f to the codomain of f is a Morse function on the set of vertices of the simplicial complex given by $(P_1^{\pm}. \leq)$ since if $\rho(f) = \rho(g) = n$ with (say) $f < g$, then $f^{-1}g$ is in P with domain and codomain both equal to n. This makes $f^{-1}g = \mathbf{1}_n$ and $f = g$. So ρ cannot take on the same value at the two ends of a 1-simplex.

From Lemma 48.2, we know that the structure of X_{j+1} is that of X_j with a cone attached at the descending link $Lk_{\downarrow}(f, X)$ for each morphism $f \in P_1^{\pm}$ with codomain $j+1$. We are thus interested in the properties of the descending links.

From our discussion in Section 26.1, the full link of an f in X with codomain $j+1$ consists of all partitions of the ends of f into pairs containing consecutive ends and singletons. The elements of the partition are attaching sites for carets. Each element of a partition gives a vertex in the link and a set of vertices in the link belongs to a simplex if their set of attaching sites are pairwise disjoint.

A pair of consecutive ends of f corresponds to multiplying f by the inverse of a forest with one caret, and so such an attaching site gives a neighbor g of f with $\rho(g) = \rho(f) - 1$. A singleton corresponds to multiplying f by a forest with one caret, and so such an attaching site gives a neighbor g with $\rho(g) = \rho(f) + 1$. Thus a vertex in the descending link corresponds to a pair of consecutive ends of f and vertices in the descending link belong to a simplex in the link if their attaching sites are pairwise disjoint.

It is clear that the structure of $Lk_{\downarrow}(f, X)$ depends only on $\rho(f)$ and not f itself. If $\rho(f) = n$, then the structure of $Lk_{\downarrow}(f, X)$ has the structure of the simplicial complex M_n whose vertices are the consecutive pairs $(i, i+1)$ with $0 \leq i < n-1$ in $\{0, 1, \ldots, n-1\}$ and where a set of vertices bounds a simplex if the corresponding pairs are pairwise disjoint.

There is a natural inclusion of M_{n-1} into M_n whose image is all sets of pairs where the element $n-1$ is involved in no pair, and there is a natural copy of the cone on M_{n-2} in M_n with cone point the vertex corresponding to the pair $(n-2, n-1)$ in $\{0, 1, \ldots, n-1\}$ and with base the image of M_{n-2} under the composition of the natural inclusions $M_{n-2} \to M_{n-1} \to M_n$. Thus M_n is the union of M_{n-1} and the cone on M_{n-2}, where the intersection of these two subsets is M_{n-2}.

In particular, if the natural inclusion of $M_{n-2} \to M_{n-1}$ is a homotopy equivalence, then M_n is contractible. If M_{n-1} is contractible, then M_n has the homotopy type of the suspension of M_{n-2}, and the connectivity of M_n is one more than the connectivity[2] of M_{n-2} And if M_{n-2} is contractible, then the inclusion of M_{n-1} into M_n is a homotopy equivalence. This gives enough to induct on if some facts about M_n for small values of n are gathered.

By inspection M_1 is empty, M_2 is a single vertex and thus contractible, M_3 is a pair of isolated vertices one of which is the image of the inclusion of M_2 into M_3, and M_4 is the disjoint union of an isolated vertex and a 1-simplex for which the inclusion $M_3 \to M_4$ is a homotopy equivalence. If the reader wishes, a further check shows that the pattern continues with M_5 contractible and the inclusion $M_6 \to M_7$ is a homotopy equivalence where both spaces have the homotopy type of a circle. The complexes M_n are known as *matching complexes* whose homotopy properties are well known. The homotopy properties of the M_n exhibit periodic behavior modulo 3, and it is useful to define

$$\nu(n) = \left\lfloor \frac{n-2}{3} \right\rfloor .$$

LEMMA 26.2. *With M_n as defined above, M_n is contractible if $n \equiv 2 \pmod 3$, and $M_n \to M_{n+1}$ is a homotopy equivalence if $n \equiv 0 \pmod 3$. For all $n \geq 2$, M_n is $(\nu(n) - 1)$-connected.*

PROOF. The specific descriptions of M_n given above for $n < 5$ show that the claims hold for these values. The first unclaimed value

[2]A suspension is a union of two intersecting contractible sets with the suspended set as the intersection. Now the van Kampen, Meyer-Vietoris, and Hurewicz isomorphism theorems give the claimed connectivity.

5 is equivalent to 2 modulo 3, so we start the induction with $n \equiv 2 \pmod 3$, and assume the claims for all values less than n. We make use of the observations about the natural inclusions $M_{n-2} \to M_{n-1} \to M_n$ made above.

For $n \equiv 2 \pmod 3$, the assumption that $M_{n-2} \to M_{n-1}$ is a homotopy equivalence shows that M_n is contractible.

For $n \equiv 0 \pmod 3$, we know that M_{n-1} is contractible and the connectivity of M_n is one more than that of M_{n-2}. But for $n \equiv 0 \pmod 3$, $\nu(n-2) = \nu(n)-1$ and the connectivity for M_n is as claimed.

For $n \equiv 1 \pmod 3$, we know that M_{n-2} is contractible and the inclusion $M_{n-1} \to M_n$ is a homotopy equivalence which verifies that specific claim, and $\nu(n-1) = \nu(n)$ verifies the connectivity claim. $\square$

COROLLARY 26.2.1. *If the codomain of f in X is n, then $Lk_{\downarrow}(f, X)$ is $(\nu(n)-1)$-connected.*

PROPOSITION 26.3. *For each $n \geq 0$, there is an i so that for all $j \geq i$, the inclusions $X_j \hookrightarrow X_{j+1}$ induce isomorphisms on π_q for $0 \leq q \leq n$.*

PROOF. Corollary 26.2.1 tells us that, eventually, going from X_j to X_{j+1} is done by coning over n-connected subsets. The trio of the van Kampen, Meyer-Vietoris, and Hurewicz isomorphism theorems says that $X_j \to X_{j+1}$ induces an isomorphism on all π_q with $q \leq n$. $\square$

We now have the following.

THEOREM 26.4. *Thompson's group F has type F_∞.*

PROOF. This is a gathering of results. From Theorem 22.3 we must show that for each $n \geq 0$ there is an $(n-1)$-connected simplicial complex on which F acts freely and simplicially with compact quotient. From Lemma 25.8, we have such complexes among the X_i if for each n there is an i for which X_i is $(n-1)$-connected. Combining Propositions 25.6, 26.3 and Lemma 25.7 gives the needed connectivity of the X_i. $\square$

Consequences that are beyond the scope of this book are as follows.

THEOREM 26.5. *Thompson's group F has the Haagerup property and satisfies the Baum-Connes conjecture and the Novikov conjecture.*

The claims follow from Theorem 26.1. The connection of Theorem 26.1 to the Haagerup property is due to Niblo and Reeves [163] (Theorem B), and the connection of the Haagerup property to the Baum-Connes conjecture (not stated that way) is due to Higson and Kasparov [108] (Theorem 1.2 and corollary). Julg 1998 [120] is an Asterix elaboration on [108] and its main theorem (Haagerup implies Baum-Connes) is stated that way. Theorem 26.1 is a special case of a more general theorem of Farley in [70], and more details are given in [70] about how Theorem 26.1 links to the hypotheses used in [163]. The introductory chapter to [46] puts all of this material in a much larger context. It is noted in [108] that their results imply that amenable groups satisfy Baum-Connes and thus the Novikov conjecture. This is an example where something holds for F that holds for all amenable groups.

27. A presentation for V

Presentations for V are harder to come by than for F and T, and there are several approaches in the literature. Very direct is the approach in Section 6 of Cannon-Floyd-Parry 1996 [43]. An attempt to make V look like a Coxeter group (with a corresponding braided Artin version) is in Brin 2006 [30]. Smaller and more elegant presentations are given in Quick-Bleak 2017 [22]. Here we give not the smallest presentation but the presentation with the most elegant packaging. In spite of its elegant appearance, it is no more elegant than any of the others when completely unpacked.

The construction is based on the following.

Theorem 27.1. *Let X be a connected and simply connected simplicial complex and let G act simplicially on X. Assume that there is a connected subcomplex T of X so that for every simplex σ of X there is a unique simplex σ' in T in the orbit of σ. Then with G_v denoting the subgroup of G fixing a vertex v, we have that G is the free product of the elements $\{G_v \mid v \in T\}$ amalgamated over their intersections.*

We offer some interpretation, background and cautions about Theorem 27.1.

For background, the theorem appears as the main theorem in Soulé 1973 [180]. Its proof will not be covered here, and the theorem will be taken as a black box. A precursor appears in Macbeath 1964 [142]. A varation appears in Brown 1984 [33]. Further, Theorem 3 of [33] states that the amalgamations only have to be performed over the intersections of those pairs of the vertex stabilizers where the two vertices span an edge in T.

For interpretation, if we use Theorem 3 of [33], then we have that G can be presented with generators the disjoint union of generating sets for the G_v, $v \in T$, and with the relations consisting of the union of the defining relations for the G_v, $v \in T$, together with a relation for each a in a generating set for $G_v \cap G_u$, where u and v span an edge in T, of the form $w_v(a) = w'_u(a)$ where $w_v(a)$ expresses a in terms of the generators of G_v and $w'_u(a)$ expresses a in terms of the generators of G_u.

By way of caution, we start by mentioning that the subcomplex T is often referred to as a *fundamental domain* of the action of G on T. The uniqueness in the statement must be taken seriously. It implies that the vertices of T are representatives of the partition of the vertices of X into pairwise disjoint orbits under G. If X is the complex often associated with $\mathbf{R}$ whose vertices are the integers and whose 1-simplexes are the intervals $[i, i+1]$ for all integers i, then $[0, 1]$ is not a fundamental domain for the action of $\mathbf{Z}$ on X where n acts as the translation t_n in which $xt_n = x + n$. This is because 0 and 1 are in the same orbit of the action. Were $[0, 1]$ taken to be a fundamental domain then one might conclude that $\mathbf{Z}$ was the free product of two trivial groups.

On the other hand, it is easy to show that the reflections r_i, $i \in \mathbf{Z}$, where $tr_i = 2i - t$, generate a group G consisting of the r_i and the translations t_{2i} for $i \in \mathbf{Z}$. Now $[0, 1]$ is a fundamental domain and G is isomorphic to the free product $\mathbf{Z}/2\mathbf{Z} * \mathbf{Z}/2\mathbf{Z}$.

We will apply Theorem 27.1 to a simply connected 2-complex on which V acts with a 2-simplex as fundamental domain. This will exhibit V as a free product of three groups amalgamated along their intersections. The complex will essentially be a subcomplex of the complex of Section 25 adapted to V. However, we give the adapted complex a completely different description by using the representation of V as the automorphism group of the algebra JT_x of Section 18. We will need to take a subcomplex so that a single 2-simplex becomes a fundamental domain of the action.

27.1. The complex. Let JT_x be the free Jónsson-Tarski algebra of Section 18. Let $\mathcal{B}$ be the set of bases of JT_x. For A and B bases of JT_x, we write $A \le B$ if B is a binary refinement of A. This is a partial order on the bases. For $A \le B$, we note that we have $|A| \le |B|$. As in Section 25.2 we restrict $\le$ to $\preceq$ by declaring that $A \preceq B$ if $A \le B$ and if all the binary splittings that take A to B are only applied to elements of A and not to any of their "descendants." In parallel to Section 25.2, we can refer to B as an elementary refinement of A. The relation $\preceq$ is reflexive and antisymmetric and not transitive.

For a positive integer i, we let $\mathcal{B}_{\ge i}$ be those $A \in \mathcal{B}$ with $|A| \ge i$. For an $A \in \mathcal{B}_{\ge i}$, every $B \in \mathcal{B}$ with $A \le B$ has $B \in \mathcal{B}_{\ge i}$. We let Y_i be the simplicial complex associated to $(\mathcal{B}_{\ge i}, \le)$. From Proposition 18.10, every pair of bases in $\mathcal{B}_{\ge i}$ has a common binary refinement in $\mathcal{B}_{\ge i}$, and so Y_i is contractible. We let X_i be the simplicial complex associated to

$(\mathcal{B}_{\geq i}, \preceq)$. The proof of the following can be copied from the proof of Proposition 25.6 by making the appropriate translations.

PROPOSITION 27.2. *The inclusion of X_i into Y_i is a homotopy equivalence, and thus X_i is contractible.*

For our purposes, we only need a simply connected complex. For positive integers $i < j$, we let $X_{i,j}$ be the subcomplex of X_i spanned by vertices (bases) $A \in \mathcal{B}$ with $i \leq |A| \leq j$. We have $X_i = \bigcup_{j>i} X_{i,j}$, so if there is a k so that the inclusion induced $\pi_1(X_{i,j}) \to \pi_1(X_{i,j+1})$ is an injection for all $j \geq k$, then $\pi_1(X_{i,k})$ must be trivial.

LEMMA 27.3. *For all $i \geq 5$ and $k \geq i+2$, we have $\pi_1(X_{i,k})$ is trivial.*

PROOF. As in Section 26.3, the discussion in Section 48.8 says that the complex $X_{i,j+1}$ is obtained from $X_{i,j}$ by attaching cones. That is, each vertex B in $X_{i,j+1} \setminus X_{i,j}$ has as its link in $X_{i,j+1}$ a complex $L(B)$ in $X_{i,j}$. A vertex A in $L(B)$ has $A \preceq B$, and B is obtained from A by a sequence of binary splittings of A with the splittings restricted to only be at elements of A. Or, A is obtained from B by applying β to each pair in a set $S(A)$ of pairs of elements of B where the pairs in $S(A)$ are pairwise disjoint. By two pairs being disjoint, we mean the underlying sets of size two are disjoint. Since we must have $|A| \geq i$, the number of pairs in $S(A)$ is no more than $j+1-i$. A difference between the current situation and that of Section 26.3 is that in Section 26.3 pairs are consecutive in some total order, where here there is no such requirement.

Call a set of no more than $j+1-i$ pairwise disjoint pairs from B an allowable set. Thus the vertices of $L(B)$ are allowable sets, and a simplex in $L(B)$ is a chain of allowable sets. If we show that each such $L(B)$ is simply connected, then attaching the cones that form $X_{i,j+1}$ from $X_{i,j}$ will cause no change in $\pi_1(X_{i,j})$.

A companion complex to $L(B)$ is $K(B)$ where an n-simplex is an allowable set of n pairs. Chains of allowable sets are simplices of $L(B)$, and are also chains of faces and thus barycenters of simplices of $K(B)$. So $L(B)$ is the barycentric subdivision of $K(B)$. We show that $L(B)$ is simply connected by showing that $K(B)$ is simply connected.

A closed edge path in the 1-skeleton of $K(B)$ is trivial if it has less than 3 edges. We will show that every edge path in $K(B)$ with $m \geq 3$ edges is homotopic in $K(B)$ to a path with no more than $m-1$ edges. Our hypotheses $i \geq 5$ and $k \geq i+2$ give $|B| \geq 8$, and all allowable sets with up to 3 pairs are in $K(B)$.

If $A_0 \to A_1 \to A_2 \to A_3$ are the vertices, in order, of a path of length 3 in $K(B)$, then 4 pairs are involved. If the pairs for A_0 and A_2 are disjoint, then the pairs for all of A_0, A_1, A_2 are pairwise disjoint, and the path $A_0 \to A_1 \to A_2$ is homotopic in across a 2-cell to $A_0 \to A_2$. The pairs for A_1 and A_2 use four elements of B, and we can thus assume that the pair for A_0 adds no more than one element to the elements used. Similarly, we can assume that the pair for A_2 adds no more than one element to the elements used, and the number of elements of B used in the path $A_0 \to A_1 \to A_2 \to A_3$ is no more than six. With at least 8 elements in B, a vertex A_4 exists in $K(B)$ using a pair disjoint from the elements used in the path. So the cone on the path exists in $K(B)$ and the path is homotopic in $K(B)$ to the path $A_0 \to A_4 \to A_3$ in $K(B)$. This completes the proof. □

27.2. A presentation. We will use Σ_n to denote the full permutation group on a set of size n.

THEOREM 27.4. *The group V acts on the simply connected 2-complex $X = X_{5,7}$ with fundamental domain a single 2-simplex. The vertices of the fundamental domain are bases $A \prec B \prec C$ of sizes 5, 6 and 7, respectively. Their stablizers are, respectively, $G_a \simeq \Sigma_5$, $G_b \simeq \Sigma_6$, and $G_c \simeq \Sigma_7$. Their pairwise intersections and the stablizers of the corresponding edges are $G_a \cap G_b \simeq \Sigma_4$, $G_a \cap G_c \simeq \Sigma_3 \times \Sigma_2$, and $G_b \cap G_c \simeq \Sigma_5$.*

The statement and its consequences need interpretation which we give before a proof. The vertex stabilizers and edge stabilizers are most easily summarized in the triangle below. Justification and details of the inclusions will follow.

$$\begin{array}{ccccc} & & \Sigma_7 & & \\ & \nearrow & & \nwarrow & \\ \Sigma_3 \times \Sigma_2 & & & & \Sigma_5 \\ \swarrow & & & & \searrow \\ \Sigma_5 & \longleftarrow & \Sigma_4 & \longrightarrow & \Sigma_6 \end{array} \tag{27.1}$$

With $A \prec B \prec C$, the elements of A can be given as $\{a, b, c, d, e\}$. The basis B is a splitting of A at one of the elements of A, say e, and the basis C is a splitting of A at e and one other element of A, say d. Thus we have $B = \{a, b, c, d, e\alpha_0, e\alpha_1\}$ and $C = \{a, b, c, d\alpha_0, d\alpha_1, e\alpha_0, \alpha_1\}$. A permutation of C giving an element of V that restricts to a permutation of A must preserve $\{a, b, c\}$ as a set, and either fix pointwise $\{d\alpha_0, d\alpha_1, e\alpha_0, e\alpha_1\}$ or simultaneously transpose $d\alpha_0$ with $e\alpha_0$ and $d\alpha_1$ with $e\alpha_1$. This explains $G_a \cap G_c \simeq \Sigma_3 \times \Sigma_2$. Simpler discussions cover the other two intersections. The restrictions given might be made more clear by looking at the arbitrarily chosen figures below for A, B and C.

Placement of the splittings was chosen more for graphical convenience than to fit well with the letters a through e used above.

$A =$ 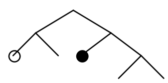 $B =$ 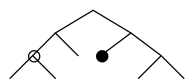 $C =$ 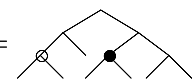

Using standard presentations for Σ_n consisting of $n-1$ transpositions as generators, $n-1$ relations establishing that the generators are of order 2, and $\frac{1}{2}(n-1)(n-2)$ relations giving the interactions of pairs of generators, the above triangle yields a presentation of V with 15 generators and 53 relations. We will not write out this presentation. The inclusions of the edge groups can be done more efficiently since two generators (a transposition and an n-cycle) can be used for each Σ_n with $n > 2$. These result in 7 of the 53 relations.

The theorem is attractive more for the visual appeal and compactness of (27.1), then for the resulting presentation. No presentation of V is particularly pleasant. In Cannon-Floyd-Parry 1996 [43] is a presentation (Lemma 6.1 and following discussion) with 4 generators and 14 relations. In Quick-Bleak 2017 [22] are presentations with 3 generators, all transpositions of elements of JT_x, and 8 relations (Theorem 1.2), and with 2 generators and 7 relations (Theorem 1.3).

None of this discusses the lengths of the relations. When fully expanded, the lengths can be huge. The 7 relations in Theorem 1.3 of [22] have lengths, when exponents are expanded, that sum to close to 200. The presentation in Theorem 1.2 of [22] is not fully expanded and a full expansion is displayed there as (2.4). The compressed form in Theorem 1.2 of [22] is about as simple a presentation for V as has been found.

Proof of Theorem 27.4. That X is a 2-complex comes from $7-5=2$. The simple connectivity comes from Lemma 27.3. That the fundamental domain is a single 2-cell comes from Proposition 18.9 and the fact that an automorphism of JT_x must take each basis to another of the same size. Proposition 18.9 also gives the stabilizer of each of A, B and C. $\square$

The discussion of Theorem 27.4 and the analysis of the complex is from Brown 1992 [35]. That paper gives further discussions about groups with triangular presentations such as given in (27.1).

27.3. A general process. The complex $(\mathcal{B}, \leq)$ is a disguised form of a generalization of the complex $(P_1^{\pm}, \leq)$ of Section 25. The generalization incorporates extra structures that allow for arbitrary bijections between the leaves of two forest. The approach involving bases of JT_x is used here and in both [35] and Brown 1987 [34] to make the action of V more obvious. The analysis in Section 26.3 that culminates in Proposition 26.3 forms a special case of general approaches to finiteness properties that are given in Sections 1–3 of [34]. See also Part II of Geoghegan 2008 [80]. The finiteness properties considered in [34] are the homological finiteness properties FP_n. Recall that a finitely presented group is of type FP_n if and only if it is of type F_n.

Combining the remarks above with the main results, Theorems 2.2 and 3.2 of [34], one can prove that V is of type F_∞. This is done for V and T and an infinite family of generalizations as Theorem 4.17 of [34].

A standard outline stretches through all of the above. Finiteness properties of a group G are established by first finding a good (contractible) complex X that G acts on. Then filter X by subcomplexes and using arguments similar to those in Section 26, prove injectivity of inclusions sufficiently far out in the filtration. This is usually referred to as analyzing the descending link. Then quote the main results of [34].

It should be noted that the finiteness criteria in Theorem 2.2 of [34] are both necessary and sufficient. So these criteria can be used to show that a certain group is FP_n but not FP_{n+1}.

28. The Brown-Geoghegan complex

In this section we describe a complex built in Brown-Geoghegan 1984 [36] where it gave the first proof that F is of type F_∞ but not of type F_n for any finite n. This answered question F11 from the section about finiteness properties in problem list in Wall 1979 [193]. The question was whether there could be a torsion free group with those properties. The space we build below and the outline that we follow are from [36], but with some modification of the arguments.

In Section 10.2 it is shown that F is an initial object in the category of groups with a conjugacy idempotent. In Section 19, this fact is used to show that F characterizes those homotopy idempotents on connected CW-complexes that split. In this section, we continue this thread and build a classifying space X for F, and show that X is an initial object in the category of connected CW complexes with a homotopy idempotent. To be specific, a an object in this category is a tuple $(Z, z_0, \rho, \alpha, K)$ with (Z, z_0) a connected CW complex with basepoint z_0, self map ρ, loop α at z_0, and homotopy K from ρ to ρ^2 that takes z_0 along α. A morphism in this category should preserve and commute with all listed structures.

Further we will show that X has finitely many cells (two, in fact) in each positive dimension, and the integral homology of F is non-trivial ($\mathbf{Z}+\mathbf{Z}$, in fact) in each positive dimension. That is, F is of type F_∞ and its integral homology groups can be calculated.

28.1. A first classifying space. We start with a classifying space Y for F that has infinitely many cells in each dimension. We will get the promised space X as a strong deformation retract of Y by collapsing certain cells.

The space Y builds itself inevitably and automatically from the data that (F,σ,x_0), with $x_i\sigma = x_{i+1}$ the shift endomorphism, is an initial object in the category of groups with a conjugacy idempotent. However, the universal cover $\widetilde{Y}$ of Y is so easy to describe that the fastest way to arrive at Y is as the quotient $F\backslash\widetilde{Y}$. Proving that $\widetilde{Y}$ is contractible and that Y arises inevitably from the stated properties of F will then proceed more easily with both Y and $\widetilde{Y}$ available.

The complex $\widetilde{Y}$ is the cubical complex $(F,\preceq)$ of Lemma 24.9. The vertices of $\widetilde{Y}$ are the elements of F, and each cube in $\widetilde{Y}$ is of the form $[f,fh] = f[1,h]$ with h an elementary element. We will make use of the analysis of Section 24.4 which codes cubes and their faces in terms of cubical tuples.

The group F acts on $\widetilde{Y}$ on the left by left multiplication and the quotient Y under the action consists of the cubes $[1,h]$ with h an elementary element. The cubes in Y are singular. In particular Y has only one vertex, but the cubes are more singular than that.

We outline our immediate goals. We will show that $\widetilde{Y}$ is contractible by showing that it is CAT(0) (see Section 48.10 in the Appendix) and quoting Theorem 48.4. We will use Theorem 48.3 to show that $\widetilde{Y}$ is CAT(0) which requires that we first show that $\widetilde{Y}$ is simply connected. We also want to show that Y is an initial object for connected CW complexes with a homotopy idempotent. These goals are interconnected and we get them by studying Y and $\widetilde{Y}$ simultaneously. Our first step is the following.

Lemma 28.1. *The space $\widetilde{Y}$ is cubical complex on which F acts freely and cellularly. The projection p from $\widetilde{Y}$ to $Y = F\backslash\widetilde{Y}$ induced by the action of F is a covering projection and F is isomorphic to the quotient $\pi_1(Y)/p_*(\pi_1(\widetilde{Y}))$.*

Proof. We know that the complex of the poset $(F,\preceq)$ is cubical by Lemma 24.9, and its underlying simplicial complex is a subcomplex of the simplicial complex given by $(F,\le)$. The action of F is on the left and the detection of $\leq$ and $\preceq$ is done by understanding multiplication on the right. Thus both $\leq$ and $\preceq$ are preserved under the action of F, and the action preserves simplices and cubes.

A simplex is a finite chain under $\leq$ and can only be carried to itself by a bijection. Since the action is order preserving, the bijection must be the identity. Since the vertices of $\widetilde{Y}$ are the elements of F and the action is by left multiplication, only the identity of F fixes a vertex. Thus the only element of F that fixes a simplex is the identity and by the remarks in Section 48.6 of the appendix, the rest of the claims follow. $\square$

The discussion that follows makes heavy use of the machinery developed for cubes and their faces in Section 24.4.

PROPOSITION 28.2. *The group $\pi_1(Y)$ is isomorphic to F and $\widetilde{Y}$ is simply connected. Further, the space Y can be given a structure $(Y, *, \sigma, x_0, H)$ as object in the category of connected CW complexes with a homotopy idempotent, and is an initial object in that category.*

PROOF. The structure of Y is given as $F\backslash\widetilde{Y}$, but to better understand Y and $\widetilde{Y}$, we will rebuild Y as a CW complex by adding cells (singular cubes) in the order implied by Lemma 24.12 and parametrized by the cubical tuples as discussed in Section 24.4. The cubical tuples correspond to words in normal form of the generators ν_i of $\mathcal{F}$, but since we are working in F with its positive submonoid F_+ that is isomorphic to $\mathcal{F}$, we will refer to the generators x_i of F_+ rather than the ν_i.

To show that the promised structure $(Y, *, \sigma, x_0, H)$ is the claimed initial object, we start by letting (Z, z_0, ρ, β, K) be a connected CW complex with homotopy idempotent ρ. As we build Y we will build a function $\phi : Y \to Z$ that commutes with all of the structure.

We start Y with a single vertex $*$ to be the image under the projection $p : \widetilde{Y} \to F\backslash\widetilde{Y} = Y$ of the vertices of $\widetilde{Y}$. We let $\phi(*) = z_0$.

We will add cells to Y in a sequence of steps in which each step has two parts. The first part (nH) of step n will add n-cells required by the existence of the homotopy H. The second part $(n\sigma)$ will add n-cells required by the existence of the idempotet σ. The 1-cells will be labeled with generators of F_+. In the narrative below, the sets of tuples C_i and C'_i are as described in the paragraph before Lemma 24.12. We will also use the face operators A_i and B_i of (24.6).

$(1H)$ Add the 1-cell given by the cubical tuple (0). This is the image of the cube $[1, x_0]$. It will be oriented from its minimal element 1 to its maximal element x_0, and labeled x_0. This is the path required by H to be the image of H restricted to $\{*\} \times [0, 1]$. We have added cells for all the tuples in C_1. We define ϕ on x_0 to carry its parametrization to β.

(1σ) For each $i \geq 1$, add the 1-cell given by the tuple $(i) = \widehat{\sigma}^i(0)$ and label it x_i to be $\sigma^i(x_0)$ and copy the orientation of x_0. The map ϕ on each x_i is defined inductively on i to preserve $\phi\sigma = \rho\phi$. We have added cells for all the tuples in C_1'.

We go through dimension two in detail before making the process inductive because dimension two will give us $\pi_1(Y)$.

($2H$) For each $i \geq 0$, add the 2-cell $\widehat{H}(i) = 0 \wedge \widehat{\sigma}^2(i) = (0, i+2)$ which carries the homotopy H that connects $\sigma(i) = \sigma(x_i) = x_{i+1}$ to $\sigma^2(i) = \sigma^2(x_i) = x_{i+2}$. Stated differently, $(0, i+2)$ connects the lower face $B_0(0, i+2) = (i+1)$ to the upper face $A_0(0, i+2) = (i+2)$. The remaining faces are $B_1(0, i+2) = (0)$ and $A_1(0, i+2) = (0)$. Recall that our cells are singular cubes. The labels on the faces of the 2-cell $(0, i+2)$ are as follows.

(28.1)

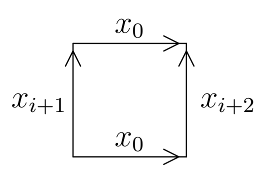

The map ϕ on the 2-cell is dictated by the homotopy K from $\rho^i(\beta)$ to $\rho^{i+1}(\beta)$. We have added cells for all the tuples in C_2.

(2σ) For each $i \geq 1$ and $(0, j+2) \in C_2$ add the 2-cell given by $\widehat{\sigma}^i(0, j+2) = (i, i+j+2)$ to be the image of $(0, j+2)$ under σ^i. The labels on the faces are as follows.

(28.2)

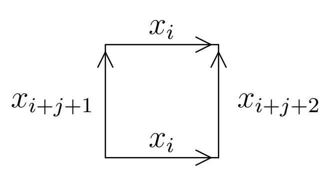

The map ϕ is defined to preserve $\phi\sigma = \rho\phi$. We have added cells for all the tuples in C_2'.

All cells added beyond this will be of dimension 3 or higher and will not affect the fundamental group of Y. The generators of $\pi_1(Y)$ are the x_i, $i \geq 0$, and the relations given by the 2-cells are all $x_j x_i = x_i x_{j+1}$ for all $0 \leq i < j$. But these are exactly the relations in the presentation (9.1) for F. So we have that $\pi_1(Y)$ is isomorphic to F. From Lemma 28.1, we have that F is isomorphic to $\pi_1(Y)/p_*(\pi_1(\widetilde{Y})) \simeq F/p_*(\pi_1(\widetilde{Y}))$. Since p is a covering projection and by Proposition 5.12 the non-abelian F cannot be a proper quotient of F, we have that $\widetilde{Y}$ is simply connected.

Adding the rest of the cells to Y continues with the pattern established above including the extension of the map ϕ to Z. Since the full two part step n adds all cells for the cubical n-tuples, when the $(n+1)$-cells are added, all their faces are already part of the n-skeleton of Y. Thus the attaching maps are all established. This completes the proof. □

PROPOSITION 28.3. *The space $\widetilde{Y}$ is a contractible, cubical CAT(0) complex on which F acts freely and cellularly, and Y is a classifying space for F.*

PROOF. Lemma 28.1 says that the action is free and cellular, and Proposition 28.2 says that $\widetilde{Y}$ is simply connected and $\pi_1(F\backslash\widetilde{Y}) = \pi_1(Y)$ is isomorphic to F. By Theorem 48.4 in the appendix, we will have that $\widetilde{Y}$ is contractible if it is CAT(0), and since $\widetilde{Y}$ is simply connected, we will have by Theorem 48.3 that $\widetilde{Y}$ is CAT(0) if the link of every vertex in $\widetilde{Y}$ is flag. A link of a vertex is flag if every finite set S of vertices in the link spans a simplex whenever every pair in S spans a simplex. In this discussion, we return to the topic of carets in a forest and the argument will finish by noting that a forest is elementary if every pair of carets is disjoint.

Let v be a vertex in $(F, \preceq)$ and let S be a finite set of vertices in the link of v. That is $w \in S$ when $\{v, w\}$ forms a 1-cube which happens when either $v \preceq w$ or $w \preceq v$ so that the difference is a single caret. We let L be the set of $w \in S$ with $w \preceq v$ and U be the set of $w \in S$ with $v \preceq w$. We first look at pairs in L, then pairs in U and lastly pairs with one element in each of L and U.

If $w_1 \neq w_2$ in L bound a 1-simplex in the link, then $\{w_1, w_2, v\}$ are vertices of a 2-cube. The vertex v can only be the maximum element of the 2-cube which must have the form $[l, v = lh]$ where h is an elementary forest of two carets. Thus every pair in L has a lower bound and by Corollary 24.5.1, L has a greatest lower bound b in $(F, \leq)$. We claim $b \preceq v$.

We bring in w_1 and w_2 as above. We know $v = bf$ for some forest f and that $b \leq l \preceq w_1 \preceq v$ and $b \leq l \preceq w_2 \preceq v$ are chains and that $[l, v = lh]$ is a 2-cube. Thus $l = bg$ for some $g \in F_+$, so $bgh = lh = v = bf$ and $f = gh$. Thus f has two exposed carets κ_1 and κ_2 obtained from h whose removal gives g. The removal of κ_1 from f leaves f_1 with $bf_1 = w_1$ and the removal of κ_2 from f leaves f_2 with $bf_2 = w_2$. If this is done for an arbitrary pair $w_i \neq w_j$ in L, we obtain a pair of exposed carets κ_i and κ_j in f where the removal of κ_i leaves f_i with $bf_i = w_i$ and removal of κ_j leaves f_j with $bf_j = w_j$. The κ_i are pairwise disjoint exposed carets in f and the removal of all of them leaves f' with $b \leq bf'$ and bf' a lower bound for all of L. Thus $b = bf'$, f' is trivial and f is elementary. Thus $[b, v = bf]$ is a cube containing L, and the vertices in L bound a simplex in the link.

To discuss U, we note that F is regarded as a groupoid on the single object ω which we can think of as the set of roots of a trivial forest to which carets may be attached to build other forests. Now for $u_1 \neq u_2$ in U, we have $u_i = v\lambda_i$, $i \in \{1, 2\}$, where each λ_i is a forest of one caret. Since $u_1 \neq u_2$, we have that $\lambda_1 \neq \lambda_2$ and we can form the forest $f_{1,2} = \lambda_1 \cup \lambda_2$. The forest $f_{1,2}$ is elementary since it consists of carets attached to different leaves of v, and we have a 2-cube $[v, vf_{1,2}]$ that contains all of v, u_1 and u_2. Thus the union of all the λ_i for which $v\lambda_i = u_i \in U$ is an elementary forest Λ and $[v, v\Lambda]$ is a cube containing all the elements of U. In the discussion of U, we did not need to use the assumption that every pair in U lies in a 2-cube.

Now we consider $w \in L$ and $u \in U$. We know that w is a vertex in $[b, bf = v]$, and u is a vertex in $[v, v\Lambda]$. So we are looking at the interaction of f and Λ. We have $w \preceq v \preceq u$ where each difference is a single caret. Our assumption says that $[w, u]$ is a 2-cell. Thus $u = wg$ with g an elementary forest of two carets and thus the union $g_1 \cup g_2$ with each g_i a forest of one caret. For one of the g_i (g_1, say) we have $v = wg_1$. Going back to $v = bf$, we have that g_1 is one of the carets in f. Since $g = g_1 \cup g_2$ is elementary, the caret attached to v to obtain u cannot be attached at a leaf of the exposed caret attached by g_1 in $v = wg_1$. Since w and u were arbitrary we have shown that in $v\Lambda = bf\Lambda$, the carets attached by f and the carets attached by Λ are pairwise disjoint. Thus $f\Lambda$ is elementary and $[b, v\Lambda] = [b.bf\Lambda]$ is a cube containing all of $S = L \cup U$.

Now $\widetilde{Y}$ is CAT(0) and thus contractible, and Y is a classifying space for F. □

Theorems 28.6 and 28.8 below give that there is a classifying CW complex X for F so that for each $n \geq 1$, X has two n-cells, and $H_n(Z; \mathbf{Z}) = \mathbf{Z} + \mathbf{Z}$. Both arguments rely on a more detailed understanding on how the cells of Y fit together, and we give those details first.

The space X will be a strong deformation retract of Y obtained by a sequence of collapses of various n-cells. To obtain the sequence we need to know which cells collapse, how they collapse and in what order. To understand the homology, we need to understand boundaries. To do all this, we classify cubes by the nature of their entries in their corresponding tuples, and our language will not distinguish between the cube and its corresponding tuple.

Consider an n-cube $\mathbf{i} = (i_0, \ldots, i_{n-1})$. For each n there are two *essential* cubes of dimension n given by $(0, 3, \ldots, 3n-3)$ and $(1, 4, \ldots, 3n-2)$. If $\mathbf{i}$ is not essential, then there is a largest j for which $i_j - i_{j-1} \neq 3$ or $i_0 \geq 2$. The cube $\mathbf{i}$ is *collapsible* if there is a largest j for which $i_j - i_{j-1} \neq 3$, and for that particular j we have $i_j - i_{j-1} = 2$. All cubes that are not essential and not collapsible are *redundant*. The set of n-cubes is thus the disjoint union of the essential n-cubes, the collapsible n-cubes, and the redundant n-cubes. Note that a cube with $i_0 \geq 2$ and all $i_j - i_{j-1} = 3$ is redundant.

In the following, we use the face operators (24.6). If $\mathbf{i}$ is an n-tuple that is collapsible with j largest so that $i_j - i_{j-1} \neq 3$ (and thus equal to 2), we call the face $A_{j-1}(\mathbf{i})$ the *free face* of $\mathbf{i}$ and the face $B_{j-1}(\mathbf{i})$ the *base face* of $\mathbf{i}$. The free face of a collapsible n-tuple is redundamt.

If $\mathbf{i}$ is a redundant n-tuple, and there is a largest j with $i_j - i_{j-1} \neq 3$, then for that j we must have $i_j - i_{j-1} \geq 4$ and we let

$$C_j(\mathbf{i}) = (i_0, \ldots, i_{j-1}, i_j - 2, i_j, \ldots, i_{n-1})$$

which is a collapsible $(n+1)$-cube whose free face $A_j(C_j(\mathbf{i}))$ is $\mathbf{i}$. Note that in this case $j \geq 1$, and that the value i_j is now found in position $(j+1)$ of $Cj(\mathbf{i})$.

If $\mathbf{i}$ is a redundant n-tuple with no j for which $i_{j+1} - i_j \neq 3$, then $i_0 \geq 2$ and we let

$$C_0(\mathbf{i}) = (i_0 - 2, i_0, \ldots, i_{n-1})$$

which is a collapsible $(n+1)$-cube whose free face $A_0(C_0(\mathbf{i}))$ is $\mathbf{i}$. We have the following.

Lemma 28.4. *Taking each collapsible $(n+1)$-cube to its free face gives a bijection from the set of collapsible $(n+1)$-cubes to the set of redundant n-cubes.*

We order cubical tuples with the short-lex order. Specifically, n-tuples are ordered lexicographically with the leftmost entry the most significant and all $(n-1)$-tuples come before all n-tuples.

Lemma 28.5. *If $\mathbf{i}$ is a redundant n-cube and $C_j(\mathbf{i})$ is the collapsible $(n+1)$-cube whose free face is $\mathbf{i}$, then for all $k \in \{0, \ldots, n\}$, we have the following possibilities.*

(1) *For $k = j$, we have $A_k(C_j(\mathbf{i})) = A_j(C_j(\mathbf{i})) = \mathbf{i}$ is redundant. The face $B_k(C_j(\mathbf{i})) = B_j(C_j(\mathbf{i}))$ can be of any type and has $B_k(C_j(\mathbf{i})) < A_k(C_j(\mathbf{i}))$ in the lexicographic order on n-tuples.*
(2) *For $k < j$, both $A_k(C_j(\mathbf{i}))$ and $B_k(C_j(\mathbf{i}))$ are collapsible.*

(3) *For $k > j$, we have $B_k(C_j(\mathbf{i})) \leq A_k(C_j(\mathbf{i})) < \mathbf{i}$ in the lexicographic order on n-tuples, and $A_k(C_j(\mathbf{i}))$ and $B_k(C_j(\mathbf{i}))$ are either equal or are the free and base faces of a collapsible $(n+1)$-cube.*

Proof. When $k = j$, the claim about $A_k(C_j(\mathbf{i})) = A_j(C_j(\mathbf{i})) = \mathbf{i}$ is by construction, the claim about the order is from the definition of the A_i and B_i, and examples that verify the claim about the type of $B_k(C_j(\mathbf{i}))$ can be constructed by the reader.

If $k < j$, then the gap of 2 in $C_j(\mathbf{i})$ between $i_j - 2$ and i_j survives as the last gap not equal to 3 in both $A_k(C_j(\mathbf{i}))$ and $B_k(C_j(\mathbf{i}))$ and they are both collapsible.

If $k > j$, then the first place from the left where $A_k(C_j(\mathbf{i}))$ and $\mathbf{i}$ disagree is in the j-th place at which the entry for $A_k(C_j(\mathbf{i}))$ is the lower of the two. The comparison between $A_k(C_j(\mathbf{i}))$ and $B_k(C_j(\mathbf{i}))$ comes from the definitions of the face operators.

If $j < k < n$, then both $A_k(C_j(\mathbf{i}))$ and $B_k(C_j(\mathbf{i}))$ are redundant since the gap in $A_k(C_j(\mathbf{i}))$ formed by removing the k-th entry from $C_j(\mathbf{i})$ is 5 or 6. That a collapsible $(n+1)$-cube can be formed for which $A_k(C_j(\mathbf{i}))$ and $B_k(C_j(\mathbf{i}))$ are the free and base faces is easily seen.

If $j < k = n$, then $A_k(C_j(\mathbf{i})) = B_k(C_j(\mathbf{i}))$. Specifically, the projection from $\widetilde{Y}$ to Y identifies these two faces (among other singularities). The type of this common face is not predictable when $j+1 = k = n$. $\square$

Theorem 28.6. *There is a subcomplex X of Y consisting of two singular cubes in each dimension so that X is a strong deformation retract of Y. It follows that F is of type F_∞.*

Proof. The complex X is the union of the unique 0-cell and all of the essential n-cells. It is easier to show how to expand X to Y, and we will do so in steps. Step 0 is to start with X. We point out that all the 0-cells are present in X as are all the essential n-cells, and all the collapsible 1-cells because there are no collapsible 1-cells. We will show inductively that after step n, all of the n-skeleton of Y and all the collapsible $(n+1)$-cells will be present. Thus after step n, the cells missing from the $(n+1)$-skeleton of Y are all redundant.

The short-lex order is a well order. At some point in the expansion process we will have a complex Z with $X \subseteq Z \subseteq Y$ and a redundant n-cell $\mathbf{i}$ that is least in the short-lex order among those cells that are not in Z. By an inductive assumption and Lemma 28.5, all faces of $C_j(\mathbf{i})$ except $\mathbf{i}$ are either collapsible n-cells, or come before $\mathbf{i}$ in the short-lex order. Thus all faces of $C_j(\mathbf{i})$ except $\mathbf{i}$ are in Z. A collapsible $(n+1)$-cell can be collapsed to the union of its non-free faces. The reverse of this collapse expands Z to Z' which includes $C_j(\mathbf{i})$ and its free face $\mathbf{i}$. ☐

Lemma 28.5 addresses faces of collapsible cubes. To compute the homology of X, we need to add information about the faces of essential cubes. We leave the verification of the following to the reader.

LEMMA 28.7. *If* $\mathbf{i}$ *is an essential* n*-cube with* $n > 1$*, then* $A_{n-1}(\mathbf{i}) = B_{n-1}(\mathbf{i})$ *is essential. For* $j < n-1$*,* $A_j(\mathbf{i})$ *and* $B_j(\mathbf{i})$ *are both redundant, and* $C_j(A_j(\mathbf{i}))$ *is the collapsible* n*-cube whose free face is* $A_j(\mathbf{i})$*, and whose base face is* $B_j(\mathbf{i})$*.*

THEOREM 28.8. *The classifying space* X *for* F *has* $H_n(X;\mathbf{Z}) = \mathbf{Z}+\mathbf{Z}$ *for each* $n \geq 1$*.*

PROOF. For $n \geq 1$, the n-chains of X form the free abelian group on the two essential n-cells. The result follows if we show that the boundary map is zero. While it would suffice to show this for the essential n-cells, the argument will need to show more. We will use cellular homology (see Chapter 2 of [80]), but will treat cells as cubes to compute boundaries.

As chains, every n-cell is a sum of essential n-cells. We note that the boundary of an n-cube $\mathbf{i}$ is the sum over $j \in \{0,\ldots,n-1\}$ of $(-1)^j(A_j(\mathbf{i}) - B_j(\mathbf{i}))$. Compare to 8.3.1 of [109]. We will be done if we show that for every essential n-cube $\mathbf{i}$, we have $A_j(\mathbf{i}) = B_j(\mathbf{i})$ as $(n-1)$-chains for all $0 \leq j < n$.

For an essential n-cube $\mathbf{i}$, Lemma 28.7 tells us that it suffices to show that for $j < n-1$, the free and base faces of the collapsible n-cube $C_j(A_j(\mathbf{i}))$ are equal as $(n-1)$-chains. We claim that the free and base faces of all collapsible n-cubes are equal as $(n-1)$-chains. We will argue this inductively.

There are no collapsible 1-cubes, so we start with $n = 2$. All 2-cubes are pictured in (28.1) and (28.2) and the claim is clear. We now consider a collapsible n-cube $\mathbf{i}$ and assume the truth of the claim for all collapsible cubes whose free face precedes the free face of $\mathbf{i}$ in the short-lex order.

The structure of a collapsible n-cube $\mathbf{i}$ shows that its free and base faces are equal as $(n-1)$-chains modulo the "side faces" (those faces of $\mathbf{i}$ other than the free and base faces). If $A_i(\mathbf{i})$ and $B_i(\mathbf{i})$ are a pair of opposite side faces of $\mathbf{i}$, then by Items (2) and (3) of Lemma 28.5, they are either collapsible, equal, or a face-base pair of some other collapsible n-cube. In the last two cases, both faces precede the free face of $\mathbf{i}$ in the lexicographic order on the $(n-1)$-cubes. By the inductive assumption, all the side faces of $\mathbf{i}$ are either collapsible or cancel in pairs, and the proof is complete. □

We can now finish Theorem 19.1 by proving its third item.

THEOREM 28.9. *Every homotopy idempotent (pointed or not) on a finite dimensional, connected CW complex splits.*

PROOF. We work with Y which is homotopy equivalent to X.

Assume $(Z, z_0, \rho, \alpha, K)$ is a conncected, CW complex with a homotopy idempotent ρ that does not split. From Proposition 28.2, we have a morphism ϕ from $(Y, *, \sigma, x_0, H)$ to $(Z, z_0, \rho, \alpha, K)$, and from Item (1) of Theorem 19.1, we know that $\phi_* : F = \pi_1(Y, *) \to \pi_1(Z, z_0)$ is an injection. Let $G \simeq F$ be the image of ϕ_* and let $\widetilde{Z}$ be the cover of Z corresponding to G. Consider the following diagram.

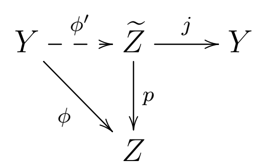

The map p is the covering projection. The lift ϕ' exists because the image of ϕ_* is contained in G (Proposition 1.33 of [104]), and induces an isomorphism on fundamental groups. The map ϕ takes x_0 to α and x_1 to $\rho(\alpha)$, and so if y_0 and y_1 are the lifts, respectively, of α and $\rho(\alpha)$ to $\widetilde{Z}$, then ϕ' takes x_i to y_i, $i \in \{0, 1\}$. Because Y is a classifying space for F, there is a map $j : \widetilde{Z} \to Y$ taking y_i to x_i, $i \in \{0, 1\}$, so that j_* is an isomorphism. It follows that $\zeta = j\phi'$ is a homotopy equivalence and $\zeta_\#$ is an isomorphism on all homology groups. From Theorem 28.8, $H_n(Y)$ is non-trivial for all n, so ζ cannot factor through a finite dimensional space, and $\widetilde{Z}$ and Z are infinite dimensional. □

The following is Theorem 7.2 of [36] whose proof is based on the HNN structure of F (Section 10.3) and its FP_∞ property. The proof while short is outside the scope of the book, and we just give the statement. This establishes one of the three properties of F mentioned in Section 15.4 that had been conjectured by Geoghegan.

THEOREM 28.10. $H^n(F, \mathbf{Z}F) = 0$ *for all* n.

29. End notes

More history around finiteness properties can be gleaned from the following: Eilenberg-Ganea 1957 [67], Wall 1965 [192] and 1979 [193], Serre 1971 [178] and 1979 [177], Baumslag-Dyer-Heller 1980 [7], Brown 1982 [32].

There are many ways to build a complex for the Thompson groups to act on. The approach in Section 23 reflects the author's fondness for the positive monoid of F.

At this point no variants of the Thompson groups have been brought up. There are some in Chapter 6 and there will be more. References for the complexes that they act on (if known) and what conclusions can be reached will be given. Most of the time, F_∞ can be proven.

Theorem 26.1 is a special case of the main theorem of Farley 2003 [70]. The cubical structures associated to Thompson's groups had been noted before, but [70] was the first to generalize and exploit them. Actions on CAT(0) spaces are proven in [70] for a large class of groups.

See Brown [34], Stein [183], and Farley [70] for the original analysis of the complexes that we have associated to those names.

The category P of finite forests is also considered in detail in Appendix A of Luo-Wan 2024 [140].

Greenberg 1992 [91] builds a classifying space for F whose cells are the associahedra of Stasheff 1963 [182].

CHAPTER 5

First order theory

Chapter contents

30. Introduction

[1]First order logic uses a restricted language which gives the first step for dealing with infinite structures. The scope of all quantifiers must be the same (possibly infinite) set. We will show that in spite of these restrictions first order techniques have considerable power in the setting of Thompson's groups. Some results are from purely first order techniques, and some are arguments with first order techniques mixed in.

The chapter is meant to illustrate, and we give details for a sampling of results. The end matter of the chapter will discuss other results that we do not cover in detail.

The main results covered are that the elementary theory of the group F (those first order sentences that hold for F) is undecidable

[1]This chapter is reasonably complete. Minor changes may be made in the future.

(in fact hereditarily undecidable), and that the theory determines the isomorphism type of F among all finitely generated groups and also largely determines the action of F on the unit interval.

Section 31 will be a brief and not entirely self contained review of the terms just used and of first order logic and model theory. What we give should be enough for the reader who is not too skeptical, and for the others we will give references for more complete background. The sections after Section 31 will contain the results mentioned in the previous paragraph.

31. Background

Two related objects coexist in this topic: structures and sets of first order statements that are made about structures. A structure $\mathscr{A} = (A, \Sigma)$ consists of a set A and a set Σ of relations, operations and constants on A. Sometimes A is allowed to be empty, but we will forbid that here. Usually Σ is allowed to be uncountable, but we will not deal with such cases. Each relation and each operation in Σ has a non-zero arity where a relation in Σ of arity n is a subset of A^n, and an operation of arity n is a function from A^n to A. Constants are elements of A.

At times it will be useful to view operations and constants as relations, with an operation of arity n viewed as a relation of arity $n+1$, and a constant as a relation of arity 1. But there are reasons for keeping them separate at other times.

Statements about a structure such as $\mathscr{A} = (A, \Sigma)$ are built from a set of symbols. Some are symbols representing the elements of Σ, some symbols are variables, and the rest are the logical connectives $\wedge$ and $\neg$, the logical quantifier $\forall$, the equal sign $=$ with its usual meaning and taken as an extra relation of arity 2, and grammatical symbols (technically not needed) such as the parentheses and comma. Symbols such as $\vee$, $\Rightarrow$, $\Leftrightarrow$, and $\exists$ can be defined in terms of $\wedge$, $\neg$ and $\forall$.

It is important that the symbols representing the elements of Σ be kept separate from Σ and thus there is a set L of such symbols and a bijection (which in our brevity we will not give a notation to) from L to Σ. Each symbol in L mapped to a relation or operation will have a non-zero arity that must equal the arity of the corresponding relation or operation. We use the letter L since this set is often called a language. This displeases us and we will use the word "signature" for the set L instead. We prefer to think of the language as the set of possible statements about a structure that are built from L with the other symbols mentioned above. In the sitation with a bijection from L to Σ as just described, the structure $\mathscr{A}$ is called an L-structure.

Keeping the symbols for Σ separate from Σ allows a single signature L to have many L-structures, hopefully somewhat related. The only requirement on two L structures about the nature of the relations (or operations) corresponding to the same symbol in L is that they be of the same arity.

Statements are built in the usual way from the symbols in L, from an agreed upon countable set of variable symbols, from the logical symbols $\wedge$, $\neg$, $\forall$, and grammatical symbols. Operations recursively applied to variable symbols and constant symbols create terms, and relation symbols applied recursively to terms while combined with logical symbols create statements. It is important to note that all statements are finite. The scope of every quantifier applied to an L-structure $\mathscr{A} = (A, \Sigma)$ will always be the set A.

There are ways of avoiding grammatical symbols. However no one actually works that way, and grammatical symbols will be used here. The usual classification of variables used in a statement into free and bound variables is made. Details can be found in the first few pages of any text on logic or model theory. Possible sources are Marker 2002 [148] and Hodges 1993 [112].

Statements formed as described are called formulas, or L-formulas to emphasize that they are based on the symbols in a signature L. Formulas without free variables (all variables used are bound) are called sentences. Sentences play very different roles than formulas with free variables. Much of the discussion of theory involves sentences, but much of the actual work done involves formulas with free variables.

31.1. Substitution, truth, and definability. It is necessary to distinguish between what is true and what can be shown to be true. Here we discuss what is true. The discussion requires a signature L and an L-structure $\mathscr{A} = (A, \Sigma)$.

Given an L-formula ϕ using n free variables $\mathbf{v} = (v_1, \ldots, v_n)$, we can ask which elements $\mathbf{a} = (a_1, \ldots, a_n)$ from A^n make ϕ true when each a_i is substituted for v_i. The subset S_ϕ of A^n consisting of those n-tuples that make ϕ true is said to be (first order) *defined by* ϕ. Sets defined by formulas dominate this chapter, and we sketch one way (elementary set theory) to arrive at the set defined by a formula.

If an L-formula ϕ consists of a single n-ary relation symbol all of whose terms are variables, then S_ϕ is the subset of A^n associated to the corresponding relation in Σ. The set defined by $\phi \wedge \theta$ would be the intersection S_ϕ and S_θ except that ϕ and θ might not use identical sets of free variables. In that case, it will be necessary to first take the Cartesian product of each of S_ϕ and S_θ with appropriate powers of A. The set defined by $\forall x, \phi$ where x is free in ϕ (we refuse to deal with the case where x is not free in ϕ) is an intersection of projections of cross sections of S_ϕ. When operations or constants are involved, then the binary relation $=$ between well chosen variable symbols is brought in to equate the values of operations or constants (each regarded as a relation here) with other parts of the formula. Thus in the L-structure $\mathscr{A} = (A, \Sigma)$ the sets defined by L-formulas having free variables are the sets associated with the elements in Σ (and $=$) combined under cartesian products with A, intersection, and intersections of projections of cross sections.

If $\mathscr{A} = (A, \Sigma)$ is an L-structure, and ϕ is an L-formula of arity n, we say that S_ϕ is (first order) definable. Note that it is possible to quantify over S_ϕ by the ruse of using $\forall \mathbf{x} \in S_\phi, \theta$ as an abbreviation for $\forall \mathbf{x}, (\phi(\mathbf{x}) \Rightarrow \theta)$. Thus one can effectively quantify over definable sets.

The sketch just given is not the usual way to arrive at a set defined by a formula, and [148] and [112] can be consulted for the usual method. There, a recursive definition is given for when an n-tuple $\mathbf{a} \in A^n$ makes an n-ary L-formula ϕ true when substituted for the free variables in ϕ. When this happens we say $\mathscr{A}$ *satisfies*, or *models*, or *is a model for* $\phi(\mathbf{a})$ and write $\mathscr{A} \models \phi(\mathbf{a})$. When ϕ is a sentence (no free variables), we write $\mathscr{A} \models \phi$.

The elementary set theory approach to defined sets fails when it comes to sentences. This is because there is no set defined by a sentence. It is either "true" for a structure or not. However, we can come close and say that if the last symbol applied in building a sentence is $\forall$, as in $\forall x, \phi$ where x is the only free variable in ϕ, then $\mathscr{A} \models \forall x, \phi$ if and only if the set defined by ϕ is exactly A. With this agreement, the usual interpretations of $\neg$ as "not" and $\wedge$ as "and" give what is needed to deal with sentences whose last symbol applied is not $\forall$.

31.2. Theories. The *elementary theory* of an L-structure $\mathscr{A}$, or $\mathrm{Th}(\mathscr{A})$, is the set of L-sentences ϕ for which $\mathscr{A}$ is a model. The elementary set theory approach makes the following clear.

(1) For every L-sentence ϕ at least one of ϕ or $\neg\phi$ is in $\mathrm{Th}(\mathscr{A})$.
(2) For every L-sentence ϕ, no more than one of ϕ and $\neg\phi$ is in $\mathrm{Th}(\mathscr{A})$.

An L-theory T is a set of L-sentences. An L-theory is *complete* if it satisfies (1) above. An L-theory is *consistent* if it satisfies (2) above. An L-theory T has a model $\mathscr{A}$ if $\mathscr{A}$ is an L-structure so that for every $\phi \in T$, $\mathscr{A} \models \phi$. Equivalently, $T \subseteq \mathrm{Th}(\mathscr{A})$.

Every theory that has a model is consistent. One result that follows of Gödel's completeness theorem for first order logic 1930 [87] is the converse and a theory is thus consistent if and only if it has a model. We will discuss another result that follows from the completeness theorem.

If S is a set of L-sentences and ϕ is another L-sentence, then we can say that ϕ is *derivable* from S if there is a formal proof of ϕ from S (i.e., a proof of $S \Rightarrow \phi$) using a fixed set of (hopefully reasonable and standard) rules of inference. We will discuss neither formal proofs nor rules of inference. On the other hand we say that ϕ is a *consequence* of S if every model for S also is a model for ϕ. The second result that follows from the completeness theorem is that ϕ is derivable from S if and only if ϕ is a consequence of S. We will use the word consequence rather than derivable, and ϕ is a consequence of S will be denoted $S \models \phi$.

We will be concerned mostly with theories of specific structures, such as the elementary theory of the Thompson group F. A theory of a class of structures comes up as we make our next point.

If T is an L-theory, we will let T^c denote its set of consequences. We have $T \subseteq T^c$. If T is consistent, then T^c is consistent since T has a model which must then be a model for T^c.

But if T is not complete, T^c might not be complete as well. There may be a pair of sentences ϕ and $\neg\phi$ not in T^c because there is a model for T that satisfies ϕ and a different model for T that does not satisfy ϕ. As an example, we can use the elementary theory of the class of groups. We let the language L consist of one binary operation that we will suggestively write by adjacency, and we let the L-theory $\mathcal{G}$ consist of the three sentences:

(i) $\forall x, \forall y, \forall z, x(yz) = (xy)z$,
(ii) $\forall x, \forall y, \exists z, x = yz$,
(iii) $\forall x, \forall y, \exists w, x = wy$.

It is a straightforward exercise that a structure is a model for $\mathcal{G}$ if and only if it has a unique, global two-sided identity and each element has a unique, two-sided inverse. Thus sentences expressing these facts are consequences of $\mathcal{G}$, and a structure is a model for $\mathcal{G}$ if and only if it is a group. However both abelian groups and non-abelian groups satisfy $\mathcal{G}$, and so neither the sentence $\forall x, \forall y, xy = yx$ nor its negation is a consequence of $\mathcal{G}$.

31.3. Undecidable theories. Let T be an L-theory. We say that T is decidable if there is a finitely describable algorithm that, given an L-sentence ϕ, decides whether or not ϕ is in T^c. Note that for a decidable, incomplete theory T, the algorithm might decide that neither ϕ nor $\neg\phi$ are in T^c. The algorithm need not "explain itself" by proving by some technique that its decisions are correct.

We are not going to discuss the details of what goes into an algorthm since that will not concern us here. However, it will be important that over a countable signature, there are only countably many possible alogorithms. This is the reason that we restrict our attencion to countable signaures. This restriction is implicit in the discussion of undecidable theories on Page 8 of Tarski 1953 [187], and explicit in Definition 10.2(i) in Monk 1976 [155] of those languages considered in the chapters of [155] on decidable and undecidable theories. Note that the vocabulary in [187] and [155] is of a certain time and differs from ours in places.

That there are undecidable theories was first established by using techniques from one of the equivalent forms of recursive function theory. Recursive function theory saw intense development after its use in the incompleteness theorem of Gödel 1931 [88], and was used in early results on undecidable theories. Shortly after, other theories were proven undecidable by the alternate technique of reducing the argument to a theory already known to be undecidable. This was first referred to as "defining" one theory in another, and eventually evolved to a more general notion of "interpreting" one theory in another. Interpretations will be discussed in Section 31.4. During this evolution, the book [187] on undecidable theories summarized and added to the alternate technique, but the concept and the vocabulary continued to evolve over time.

There are two modifiers relevant to the term "undecidable." Given two L-theories $T_1 \subseteq T_2$, we say that T_2 is an extension of T_1, and that T_1 is a subtheory of T_2. Let T be a consistent L-theory. We say that T is *essentially undecidable* if T and every consistent extension of T is undecidable. We say that T is *hereditarily undecidable* if T and every subtheory of T is undecidable.

The theory $\mathcal{G}$ of groups from Section 31.2 is an example of a hereditarily undecidable theory. That the theory $\mathcal{G}$ is undecidable is shown in (Chapter III of [187]). That $\mathcal{G}$ is hereditarily undecidable follows from the basic facts about derivability including the deduction theorem ($S \cup \{\phi\} \models \theta$ if and only if $S \models (\phi \Rightarrow \theta)$). We can regard (i)–(iii) of Section 31.2 as axioms for $\mathcal{G}$ and note that $\phi =$(i)$\wedge$(ii)$\wedge$(iii) is a single sentence from which all sentences in $\mathcal{G}^c$ are consequences. That is $\theta \in \mathcal{G}^c$ if and only $(\phi \Rightarrow \theta) \in \mathcal{G}^c$. If T_1 is a subtheory of $\mathcal{G}$, then so is $T_2 = T_1 \cup \{\phi\}$, and $T_2^c = \mathcal{G}^c$. Now $\theta \in T_2^c$ if and only if $(\phi \Rightarrow \theta) \in T_1^c$. If T_1 is decidable, then the truth of $(\phi \Rightarrow \theta) \in T_1^c$ is decidable. By our chain of equivalences and equalities, this makes $\theta \in \mathcal{G}^c$ decidable, a contradiction.

More generally, we have the following given as Lemma 2.13 in Altinel-Muranov 2009 [5]

LEMMA 31.1. *If a theory T has finite signature and has a finitely axiomatizable subtheory that is essentially undecidable, then T is hereditarily undecidable.*

For an example of a theory that fits the hypotheses of Lemma 31.1, we have the following from Chapter II, Theorem 9 of [187].

THEOREM 31.2. *The theory $\mathcal{N}$ of the semiring $(\mathbf{N}, +, \times)$ has a finitely axiomizable subtheory that is essentially undecidable.*

That the cardinality of an axiom set is important lets us explain why operations and constants should not always be thought of as relations. If an n-ary operation f on a structure $\mathscr{A} = (A, \Sigma)$ is viewed as an $(n+1)$-ary relation R where the last coordinate a_n is functionally determined from the first n coordinates $\mathbf{a} = (a_0, \ldots, a_{n-1})$, then

$$\begin{aligned}&\forall \mathbf{a}, \exists x, R(\mathbf{a}, x), \quad \text{and}\\ &\forall \mathbf{a}, \forall x, \forall y, (R(\mathbf{a}, x) \wedge R(\mathbf{a}, y)) \Rightarrow x = y\end{aligned}$$

are sentences in $\mathrm{Th}(\mathscr{A})$. Unless these sentences follow from something else, they must be regarded as axioms of $\mathrm{Th}(\mathscr{A})$. Having infinitely many operations in Σ that are regarded as relations can change whether or not $\mathrm{Th}(\mathscr{A})$ is seen as having finitely many axioms. A similar discussion holds for constants.

That the theory $\mathcal{G}$ of groups is hereditarily undecidable has a consequence worth noting. As given in Section 31.2, $\mathcal{G}$ is an L-theory where the language L has as its only element a binary operation. Thus the L-theory consisting of consequences of the empty set of sentences is undecidable. That is, the theory of the class of structures whose signature is a single binary operation is undecidable. This is noted as 1(c) in the table on P. 279 in [155] of some undecidable theories.

We can also use the theory $\mathcal{G}$ to show that neither modifier "hereditarily" nor "esentially" is automatic for undecidable theories. The L-theory $\mathcal{AG}$ of abelian groups is decidable (Janiczak 1950 [116]), and it is obtainable from the theory $\mathcal{G}$ by adding to the sentences (i)–(iii) of Section 31.2, the sentence

(iv) $\forall x, \forall y, xy = yx$.

This makes $\mathcal{AG}$ a decidable, consistent extension of the undecidable theory $\mathcal{G}$, and so $\mathcal{G}$ is not essentially undecidable.

We can guarantee the existence of an L-theory that is an undecidable extension of $\mathcal{AG}$ which will then be an undecidable, but not hereditarily undecidable theory. We will build an uncountable collection of extensions $\mathcal{AG}(P)$ of $\mathcal{AG}$ where the $\mathcal{AG}(P)^c$ are pairwise unequal. Since there can only be countably many algorithms, only countably many of the $\mathcal{AG}(P)$ can be decidable. We let P be an infinite collection of primes in $\mathbf{N}$, and we build $\mathcal{AG}(P)$ by adding, for each $p \in P$, the following sentence to the sentences (i)–(iv) of $\mathcal{AG}$:

(v_p) $\forall x, x^{p+1} \neq x$.

Since the identity is not mentioned in L, the sentence (v_p) is a cheap way to say that no x has order p. These theories are all consistent because the additive group $\mathbf{Z}$ is a model for each $\mathcal{AG}(P)$. Because of the existence of the cyclic groups of prime order, it is clear that the $\mathcal{AG}(P)^c$ are all different. Uncountably many must be undecidable and not hereditarily undecidable.

31.4. Interpretations and bi-interpretations. One way to prove that a structure $\mathscr{A}$ has an undecidable theory is to interpret within $\mathscr{A}$ another structure $\mathscr{B}$ that has an undecidable theory. The technique involves definable sets as discussed in Section 31.1.

Definition 31.3. Let $\mathscr{A} = (A, \Sigma_1)$ be an L_1-structure and $\mathscr{B} = (B, \Sigma_2)$ be an L_2-structure. Here, we regard operations and constants in Σ_2 as relations and thus each $R \in \Sigma_2$ is a subset of some B^j. An *interpretation* of $\mathscr{B}$ in $\mathscr{A}$ is a surjection f from a definable subset S of A^k to B so that for each $R \in \Sigma_2$ (as well as $=$) identified with its associated subset of some B^j, the inverse image $f^{-1}(R)$ is a definable subset of A^{kj}.

If it looks like f goes in the wrong direction, this can be corrected by replacing f by its inverse, a set valued function going from B to the definable subset S of the definition.

To give an illustration of Definition 31.3, we consider the usual interpretation of $(\mathbf{Q}, +, \times)$ in $(\mathbf{Z}, +, \times)$. Here $Q = \{(a, b) \in \mathbf{Z}^2 \mid b + b \neq b\}$ is a definable subset of $\mathbf{Z}^2$, and we set $f : Q \to \mathbf{Q}$ as $f(a, b) = \frac{a}{b}$. Now $E = \{(a_1, b_1, a_2, b_2) \in \mathbf{Z}^4 \mid a_1 b_2 = a_2 b_1\}$ defines the inverse image of "=," $P = \{(a_1, b_1, a_2, b_2, a_3, b_3) \in \mathbf{Z}^6 \mid a_3 b_1 b_2 = b_3(a_1 b_2 + a_2 b_1)\}$ defines the inverse image of $+$ regarded as a relation, and $M = \{(a_1, b_1, a_2, b_2, a_3, b_3) \in \mathbf{Z}^6 \mid a_3 b_1 b_2 = a_1 a_2 b_3\}$ defines the inverse image of $\times$ regarded as a relation.

Interpretations sometimes have limitations that can be overcome by the use of parameters (temporary constants). We discuss the use of parameters.

If ϕ is an L-formula and $\mathscr{A} = (A, \Sigma)$ an L-structure, and if $\mathbf{v} = (v_1, \ldots, v_j)$ and $\mathbf{w} = (w_1, \ldots, w_k)$ partition the free variables of ϕ into tuples that share no entries, then with $\mathbf{a} = (a_1, \ldots, a_k) \in A^k$, substituting each a_i, $1 \leq i \leq k$, for w_i in ϕ creates a formula $\phi(\mathbf{v}, \mathbf{a})$ of arity j which defines a subset of A^j *with parameters* $\mathbf{a}$. If sets defined with parameters are used in an interpretation, then the interpretation is *with parameters*. The following can be found as Lemma 6.2 of Altinel-Muranov 2009 [5] where a proof is given.

LEMMA 31.4. *Let M and N be two structures of finite signatures such that* $\mathrm{Th}(M)$ *is hereditairily undecidable, and N interprets M with parameters. Then* $\mathrm{Th}(N)$ *is hereditarily undecidable as well.*

The inclusion of parameters in a discussion can have drastic effects. We showed in Section 12.7.1 that F satisfies no laws. Consider the two parameters $a = x_1$ and $b = x_0^2(x_0x_1)^{-1}$ with their usual action on $[0,1]$. The support of a is $[\frac{1}{2}, 1]$ and the support of b is $[0, \frac{1}{2}]$. For any $f \in F$, the point $\frac{1}{2}f$ is either in $[0, \frac{1}{2}]$ or in $[\frac{1}{2}, 1]$. So either $[a, b^f] = 1$ or $[a^f, b] = 1$, and for all $f \in F$, we have $[[a, b^f], [a^f, b]] = 1$. So F satisfies a law with parameters.

Structures that are interpretable within each other are said to be bi-interpretable under a special condition. First it is noted that a composition of interpretations makes sense in the usual situations and is an interpretation.

DEFINITION 31.5. If there is an interpretation of a structure $\mathscr{B}$ in a structure $\mathscr{A}$ and vice versa, then they form a pair of bi-interpretations if the two compositions, regarded as subsets of Cartesian powers of the respective universes, are definable sets. If any definitions involve parameters, then the bi-interpretations are with parameters.

Interpretations induce homomorphisms of automorphism groups. If $\mathscr{A} = (A, \Sigma)$ is an L-structure, then $\mathrm{Aut}(\mathscr{A})$ is the group of permutations of A that preserve (or commute with) the elements of Σ. In terms of sets, if σ is a permutation of A, then σ acts on A^n diagonally, $(a_1, \ldots, a_n)\sigma = (a_1\sigma, \ldots, a_n\sigma)$ and to be in $\mathrm{Aut}(\mathscr{A})$, we must have $R\sigma = R$ for every $R \in \Sigma$ where every R in Σ is regarded as a subset of a Cartesian power of A. That is each $R \in \Sigma$ is preserved as a set by σ.

It follows that every definable set in $\mathscr{A}$ is preserved as a set by $\mathrm{Aut}(\mathscr{A})$. If a set S is definable using parameters, then S is only preserved as a set by those elements of $\mathrm{Aut}(\mathscr{A})$ that fix the parameters pointwise.

The claims below follow from careful inspection of the definitions.

Lemma 31.6. (I) *The composition of interpretations is an interpretation.*

(II) *An interpretation of a structure* $\mathscr{B}$ *in a structure* $\mathscr{A}$ *induces a well defined homomorphism from* $\mathrm{Aut}(\mathscr{A})$ *to* $\mathrm{Aut}(\mathscr{B})$.

(III) *In* (II), *if the interpretation uses parameters, then* $\mathrm{Aut}(\mathscr{A})$ *must be replaced by the subgroup that fixes the set of parameters pointwise.*

(IV) *Bi-interpretations induce isomorphisms of automorphism groups and if parameters are involved, the automorphism groups must be replaced by the subgroups that fix the parameters pointwise.*

Bi-interpretations involving F and their consequences are discussed in Section 35.

32. Undecidability of $\mathrm{Th}(F)$

We will give two proofs that the elementary theory of F is hereditarily undecidable. This section contains the first proof and uses an interpretation with parameters. The second proof uses an interpretation without parameters and is contained in Section 34. Before that, Section 33 uses first order arguments to recover the action of F on $[0,1]$ from the algebra of F. Some of the tools developed there will be used in Section 34. For the purpose of showing undecidability, there is no need for a proof using an interpretation without parameters. However, it is an interesting display of techniques.

The fact that the elementary theory of F is hereditarily undecidable is interesting, but this does not make F particularly unique. A general fact is that any virtually solvable infinite group that is not virtually abelian interprets arithmetic and thus has a hereditarily undecidable elementary theory. See the comments and references in the first paragraphs of Section 5 of [5].

The proof in this section is short, and is based on the next lemma which follows from the general fact just mentioned. However a self contained proof is given in Lemma 5.1 of [5].

Lemma 32.1. *The group* $\mathbf{Z}\wr\mathbf{Z}$ *interprets* $(N,+,\times)$ *with parameters.*

Proposition 32.2. *The Thompson group* F *interprets* $\mathbf{Z}\wr\mathbf{Z}$ *with parameters.*

Proof. Let $a = x_0$ and let $b = x_1^2 x_2^{-1} x_1^{-1}$. As a tree pair, we have

$b =$ 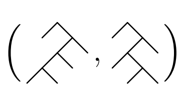

making $(\frac{1}{2}, \frac{3}{4})$ the support of b. Since $\frac{1}{2}a = \frac{3}{4}$, Lemma 4.7 gives that $\langle a, b\rangle$ equals $\langle b\rangle \wr \langle a\rangle$ and is isomorphic to $\mathbf{Z} \wr \mathbf{Z}$. We next show that $\langle a, b\rangle$ is a definable subset in F.

We use $C(f)$ and $C(A)$ to denote, respectively, the centralizers of $f \in F$ and $A \subseteq F$. From Proposition 5.17 we have $C(a) = \langle a\rangle$ which gives that $\phi(f, a) := (af = fa)$ defines $\langle a\rangle$ using the parameter a. The set $B = \{b^{a^n} \mid n \in \mathbf{Z}\}$ is defined by $\theta(f, a, b) := (\exists g \in C(a), f = b^g)$. The supports of the elements of B cover $(0, 1)$ except for a countable discrete set of points in $(0, 1)$. From Proposition 5.17 again, we have $C(B) = \langle B\rangle$ which is then defined by $\sigma(f, a, b) := (\forall g \in B, fg = gf)$. Now $\langle a, b\rangle = C(B)C(a)$ which is defined by $\rho(f, a, b) := (\exists g \in C(B), \exists h \in C(a), f = gh)$.

The group operations on $\langle a, b\rangle$ are just the restrictions of the group operations of F to $\langle a, b\rangle$ and are definable. ☐

THEOREM 32.3. *The elementary theory of Thompson's group F is hereditarily undecidable.*

PROOF. The claim now follows from Lemma 31.6(I), Theorem 31.2, and Lemmas 31.1, 31.4, and 32.1, and Proposition 32.2 ☐

33. Recovering the action

It is sometimes the case that the action of a group G on a space X is sufficiently flexible so that two conditions are met. First the centralizer of G in the full group of homeomorphisms of X is trivial, and second the space X and the action of G can be reconstructed solely from the algebraic structure of G. When this happens, every automorphism of G is realized as conjugation by a unique homeomorphism of X. That is, the automorphism group of G is isomorphic to the normalizer of G in the homeomorphism group of X.

Further, if for $i \in \{1, 2\}$ we have a group G_i acting on a space X_i that has the properties just referred to, then every isomorphism from G_1 to G_2 is induced by a unique homeomorphism from X_1 to X_2. In particular, isomorphism of the groups implies topological conjugacy of the actions.

These results have tremendous power in the study of automorphisms, isomorphisms and homomorphic embeddings of the groups to which they apply, and these results apply to a very large number of groups in the Thompson family.

In Section 33.1 below we discuss one of the most general theorems of this type, and certainly one of the most quoted theorems that apply to Thompson's groups. We only give a reference for the details of the proof, but in Section 33.2, we give the details of a more specialized result that applies to groups acting on the line, and in particular to Thompson's group F. We do so to illustrate the techniques and because many of the mechanics will be used later in Section 34.

33.1. The Rubin theorem. The following theorem is a rewording of one of the main results in Rubin 1989 [171]. The result has been modified by Rubin many times to take into account extra structure on the spaces involved. Rubin's original proof left space for such modifications and a more streamlined proof is in Belk-Elliott-Matucci 2024 [8].

THEOREM 33.1. *If for $i \in \{1, 2\}$ the group G_i acts faithfully on a locally compact, Hausdorff spaces X_i with no isolated points, and for each open $U \subseteq X_i$ and each $p \in U$, the closure of the oribit of p under $\{g \in G \mid \mathrm{supp}(g) \subseteq U\}$ contains an open set, then for each isomorphism $\theta : G_1 \to G_2$ there is a unique homeomorphism $h : X_1 \to X_2$ so that for all $(g, x) \in G_1 \times X_1$ we have $h(xg) = (h(x))\theta(g)$.*

We will not go into the details of the proof. The core of the argument is that if a group G and space X satisfy the hypotheses satisfied by each pair (G_i, X_i) in Theorem 33.1, then the elementary theory of the group G is strong enough to define the poset of the interiors of the closures of the supports of elements of G. The space X is then recovered from the set of ultrafilters on the poset.

An action that satisfies the hypotheses of Theorem 33.1 is called a *Rubin action* in [8]. Many groups in the Thompson family are given as Rubin actions, and Theorem 33.1 is applied to answer questions about their automorphisms and monomorphisms. For automorphisms, we have the following corollary.

COROLLARY 33.1.1. *If G has a Rubin action on X, then the automorphism group of G is a quotient of the normalizer of G in $\mathrm{Homeo}(X)$. If the centrlizer of G in $\mathrm{Homeo}(X)$ is trivial, then the quotient homomorphism is an isomorphism.*

33.2. The Bieri-Strebel argument. What follows is a minor reworking of an argument from Bieri-Strebel 2016 [15]. We regard F as the structure $(F, \times, ^{-1}, 1)$. When $\leq$ is a total order on a set, we use $\leq'$ to denote the relation where $a \leq' b$ if and only if $b \leq a$.

THEOREM 33.2. *In the structure F, there are definable without parameters a set E, an equivalence relation R on E, and relations $\leq$ and $\leq'$ on E so that if $\mathfrak{D} = \mathbf{Z}[\frac{1}{2}] \cap (0, 1)$, then these interpret $(\mathfrak{D}, =, \leq, \leq')$ with possibly a needed switch of $\leq$ with $\leq'$ on E. Further there is an* *action of F on the set $\widetilde{E}$ of equivalence classes in E so that the interpretation map from E to $\mathfrak{D}$ induces a conjugacy from the action of F on $\widetilde{E}$ to the action of F on $\mathfrak{D}$.*

The proof will follow a key lemma and a sequential development of definable sets.

Conjugating elements of F by the order reversing homeomorphism $t \mapsto 1-t$ on $[0,1]$ induces an automorphism of F which interchanges the relations $\leq$ and $\leq'$ in the above statement. Thus there is no way that the elementary theory of F can express a preference for one of $\leq$ and $\leq'$ over the other, and the specifics of the interpretation map require information outside of the elementary theory of F.

Technically, we could do the interpretation with paramenters and make every element of F a parameter. Then the action of F on $\widetilde{E}$ would be first order definable taking the form of an infinite set of unary operations indexed over F. It is not clear that this is of any interest.

The main ingredient in proving Theorem 33.2 is a first order formula in the algebra of F that is equivalent to a geometric statement about the group action. The binary operation $\times$ in the structure $(F, \times, ^{-1}, 1)$ will be written as adjacency. The following formula is the basis for all that follows.

$$\alpha(f,g,d) := (f \neq 1) \wedge (g \neq 1) \wedge (d \neq 1) \wedge \forall k,\ [f, g^{d^k}] = 1. \tag{33.1}$$

To discuss the geometric equivalent, we introduce some terminology. We say that an order preserving permutation d on a linearly ordered set is *directional* if d is not the identity, and so that if there is a t with $td > t$, then for all t we have $td \geq t$. Note that d directional implies that if there is a t with $td < t$, then for all t we have $td \leq t$. If d is directional and there is a t with $td > t$, then we say that the d has *positive direction* and has *negative direction* otherwise.

If X and Y are non-empty subsets of I, then we say that $t \in (0,1)$ *separates* X from Y if X lies in one component of $I \setminus \{t\}$ and Y lies in the other. In such case we write $X < Y$ if some (eqivalently all) $(p,q) \in X \times Y$ satisfy $p < q$, and write $X > Y$ otherwise.

Lemma 33.3. *Given f, g, d in F we have $\alpha(f,g,d)$ holds if and only if none is the identity, d is directional, and there is a $t \in (0,1)$ that separates X the support of f from Y the support of g. If $\alpha(f,g,d)$ holds, then $X < Y$ if and only if d has positive direction.*

Proof. We consider the "if" direction first. Under its assumptions, every conjugate d^k of d is directional with the same direction as

d, so every d^k moves each point of $\operatorname{supp}(g)$ no closer to the support of f. Thus f and the conjugate of g by d^k will commute.

For the other direction we prove the inverse. This breaks into two cases. For the first there are s and t in $(0,1)$ with $sd < s$ and $td > t$. For the second there are $p < q < r$ in $(0,1)$ with p and r in the support of (say) f, and q in the support of g. In both cases, there is a p in the support of f, a q in the support of g, and a t in the suppport of d so that $q - p$ and $td - t$ have opposite signs. It suffices to consider $q > p$ and $td < t$. The case $p > q$ has an argument that is identical except for necessary changes of inequalities.

The equality $[f, g^{d^k}] = 1$ holds if either or both of f and g are replaced by their inverses. So we can assume that $pf > p$ and $qg > q$. In particular we have

$$pf^{-3} < pf^{-2} < pf^{-1} < p < q < qg < qg^2.$$

The following calculations verify and totally obscure the fact that a conjugate d^k of d can distribute the orbit (q, qg, qg^2) among the intervals (pf^{-3}, pf^{-2}), (pf^{-2}, pf^{-1}) and (pf^{-1}, p) so that f and g^{d^k} cannot possibly commute. The idea is to place q and qg in the first interval and qg^2 in the third. We can assume that q is a dyadic, so we have that its images under various elements of F are dyadics as well.

There is an open interval J about t so that Jd is disjoint from J, and there are open intervals J_0 and J_1 in J so that we have

$$J_0d < J_1d < J_0 < J_1.$$

Choose dyadics $a_1 < a_2 < a_3$ with a_1 and a_2 in (pf^{-3}, pf^{-2}) and $a_3 \in (pf^{-1}, p)$, and $b_1 < b_2 < b_3$ with b_1 and b_2 in J_0 and $b_3 \in J_1$. Note that

$$a_1 < a_2 < a_3 < q < qg < qg^2, \text{ and}$$
$$b_1d < b_2d < b_3d < b_1 < b_2 < b_3.$$

By the transitivity properties of F from Section 5, there is a $k \in F$ so that

$$b_1dk = a_1, \ b_2dk = a_2, \ b_3dk = a_3,$$
$$b_1k = q, \ b_2k = qg, \ b_3k = qg^2.$$

We have

$$(qk^{-1}dk) = a_1, \ (qk^{-1}dk)g^{d^k} = (qgk^{-1}dk) = a_2, \text{ so}$$
$$pf^{-3} < (qk^{-1}dk) < (qk^{-1}dk)g^{d^k} < pf^{-2}, \text{ and}$$
$$pf^{-2} < (qk^{-1}dk)g^{d^k}f < pf^{-1}.$$

But

$$(qk^{-1}dk)g^{d^k} < pf^{-2} < (qk^{-1}dk)f, \text{ so}$$
$$(qk^{-1}dk)(g^2)^{d^k} < (qk^{-1}dk)fg^{d^k},$$

and

$$(qk^{-1}dk)(g^2)^{d^k} = qg^2k^{-1}dk = a_3 \in (pf^{-1}, p), \text{ giving}$$
$$(qk^{-1}dk)g^{d^k}f < pf^{-1} < (qk^{-1}dk)(g^2)^{d^k} < (qk^{-1}dk)fg^{d^k}.$$

So $g^{d^k}f$ and fg^{d^k} disagree on $(qk^{-1}dk)$ and we have $[f, g^{d^k}] \neq 1$. □

We can now define formulas based on $\alpha(f, g, d)$ and give their geometric equivalents. Justifications of the verbal claims in what follows are left as exercises to the reader. We will also abbreviate mercilessly.

$$\beta(d) := \exists f, \exists g, \alpha(f, g, d) : \tag{33.2}$$

d is directional.

$$\gamma(d_1, d_2) := \exists f, \exists g, \alpha(f, g, d_1) \wedge \alpha(f, g, d_2) : \tag{33.3}$$

d_1 and d_2 are directional in the same direction.

Let K be the set defined by β. Now γ defines an equivalence relation $\tilde{\rightarrow}$ on K with exactly two equivalence classes. We will use $\bar{d}$ to denote the equivalence class of $d \in K$ under $\tilde{\rightarrow}$ and will use $-\bar{d}$ to denote the opposite class.

We now have $\bar{d}_1 = \bar{d}_2$ implies $\alpha(f, g, d_1) \Leftrightarrow \alpha(f, g, d_2)$ and $\bar{d}_1 = -\bar{d}_2$ implies $\alpha(f, g, d_1) \Leftrightarrow \alpha(g, f, d_2)$. To abbreviate further, we will write $f \leq_{\bar{d}} g$ for $\alpha(f, g, d)$.

$$\delta(f, g, h, d) := \alpha(f, g, d) \wedge \alpha(g, h, d) : \tag{33.4}$$

$d \in K$ and $f \leq_{\bar{d}} g \leq_{\bar{d}} h$.

$$\epsilon(f, g, d) := \alpha(f, g, d) \wedge \forall h, \neg\delta(f, h, g, d) : \tag{33.5}$$

$d \in K$, $f \leq_{\bar{d}} g$, and the closures of the supports of f and g have a single point in common that we will denote $\dot{\epsilon}(f, g, d)$.

In parallel to an observation above, we have that $\bar{d}_1 = \bar{d}_2$ implies $\epsilon(f, g, d_1) \Leftrightarrow \epsilon(f, g, d_2)$ and $\bar{d}_1 = -\bar{d}_2$ implies $\epsilon(f, g, d_1) \Leftrightarrow \epsilon(g, f, d_2)$

The formula $\epsilon(f, g, d)$ identifies points that the group acts on. The point common to the closures of the supports of f and g must necessarily be dyadic. We let E be the set defined by ϵ. We need to identify when two elements of E should be seen as the same, and when one element is less than another.

$$\begin{aligned}\zeta(f_1,g_1,d_1,f_2,g_2,d_2) :=& \epsilon(f_1,g_1,d_1)\wedge\epsilon(f_2,g_2,d_2)\wedge \\ &\Big(((\bar{d}_2=\bar{d}_1)\wedge(f_1\leq_{\bar{d}_1} g_2))\vee \\ &((\bar{d}_2=-\bar{d}_1)\wedge(f_1\leq_{\bar{d}_1} f_2))\Big):\end{aligned} \tag{33.6}$$

$d_1\in K$, $d_2\in K$, $(f_1,g_1,d_1)\in E$, $(f_2,g_2,d_2)\in E$, and whether or not $\bar{d}_1=\bar{d}_2$, the common point of the closure of the supports of f_1 and g_1 is no greater, in the direction of d_1, than the corresponding point for f_2 and g_2.

In this situation we write $\dot\epsilon(f_1,g_1,d_1)\leq_{\bar{d}_1}\dot\epsilon(f_2,g_2,d_2)$. For each $d\in K$, this gives a partial order on E, all elements of K in the same class under $\tilde{\rightarrow}$ give the same order and elements in different classes under $\tilde{\rightarrow}$ give opposite orders.

$$\begin{aligned}\eta(f_1,g_1,d_1,f_2,g_2,d_2) :=& (\dot\epsilon(f_1,g_1,d_1)\leq_{\bar{d}_1}\dot\epsilon(f_2,g_2,d_2))\wedge \\ &(\dot\epsilon(f_2,g_2,d_2)\leq_{\bar{d}_1}\dot\epsilon(f_1,g_1,d_1)):\end{aligned} \tag{33.7}$$

the points $\dot\epsilon(f_1,g_1,d_1)$ and $\dot\epsilon(f_2,g_2,d_2)$ are the same point.

The formula η defines an equivalence relation on E and we let $\widetilde{E}$ be the set of equivalence classes. The order $\leq_{\bar{d}}$ on E induces a total order on $\widetilde{E}$. If each element in $\widetilde{E}$ is identified with the dyadic rational in $(0,1)$ that is the common point of closure of support of a representative pair of functions, then the order $\leq_{\bar{d}}$ is identical to the usual order on $\mathbf{Z}[\frac{1}{2}]\cap(0,1)$ if d has positive direction, and the orders are the opposite otherwise.

Proof of Theorem 33.2. We use the formulas given above. The set E is defined by the formula ϵ, and hopefully the reader has verified that η defines an equivalence relation on E with set $\widetilde{E}$ of eqivalence classes, that each equivalence class corresponds to a single point in $\mathfrak{D}=\mathbf{Z}[\frac{1}{2}]\cap(0,1)$, and that the two equivalence classes of $\tilde{\rightarrow}$ on K give inverse total orders $\leq$ and $\leq'$ on $\widetilde{E}$. Lastly, sending each class in $\widetilde{E}$ to its corresponding point in $\mathfrak{D}$ is a bijection from $\widetilde{E}$ to $\mathfrak{D}$ that carries the pair $\{\leq,\leq'\}$ on $\widetilde{E}$ to the pair $\{\leq,\leq'\}$ on $\mathfrak{D}$.

The action of F on $\widetilde{E}$ will be by conjugation in that for $h\in F$, if $\epsilon(f,g,d)$ holds then so does $\epsilon(f^h,g^h,d)$ and we will let the image of the triple (f,g,d) under the action of h be (f^h,g^h,d). In $\mathfrak{D}$, we have that $(\dot\epsilon(f,g,d))h=\dot\epsilon(f^h,g^h,d)$ which verifies the claimed conjugacy. ☐

Theorem 33.2 can be extended to give an isomorphism implies topological conjugacy result similar to that of Theorem 33.1. The following is Theorem E16.3 of Bieri-Strebel 2016 [15] and first appeared as Theorem E3 in Bieri-Strebel 1985 [13]. Item (4) in the statement uses a term given in Definition 5.1.

For the proof see [15]. However, the reader can also derive the proof by checking that the proofs of Lemma 33.3 and Theorem 33.2 use only the properties of F given as (1–4) below, and supplying the necessary arguments to go from the conclusion of Theorem 33.2 to the conclusion below.

THEOREM 33.4. *Let G act as a group of order preserving permutations on a dense, countable, linearly ordered set A having no extremal elements. Assume the following properties.*

(1) *Some $d \in G$ is directional.*
(2) *There is a non-identity $g \in G$ whose support is bounded both above and below.*
(3) *There is a triple (f, g, d) satisfying $\epsilon(f, g, d)$.*
(4) *The action of G on A is o-6-transitive.*

Then if Φ is an isomorphism from F to G, there is a unique bijection $\Theta : \mathbf{Z}[\frac{1}{2}] \cap (0,1) \to A$ that conjugates the action of F on $\mathbf{Z}[\frac{1}{2}] \cap (0,1)$ to that of G on A in that for each $f \in F$ and $t \in \mathbf{Z}[\frac{1}{2}] \cap (0,1)$, we have $\Theta(tf) = (\Theta(t))(\Phi(f))$. Further Θ is either order preserving or order reversing.

34. A parameter free interpretation of arithmetic

We give the result of Altinel-Muranov 2009 [5] that Thompson's group F interprets $(\mathbf{N}, +, \times)$ without parameters. As mentioned before, an interpretation without parameters is not necessary to show that the elementary theorey of F is undecidable, but the techniques are interesting. Our argument for the following varies from that in [5] in that we use material from Section 33.2.

THEOREM 34.1. *The elementary theory of Thompson's group F interprets $(\mathbf{N}, +, \times)$ without parameters.*

PROOF. What will be shown directly is that F interprets $(\mathbf{N}, +, |)$ where $|$ is the relation "divides." However, the constant 1 in $\mathbf{N}$ is easily definable in $(\mathbf{N}, +, |)$, and from R. M. Robinson 1951 [170] we have the formula

$$p_1(n, k) := \forall m, (n|m \Leftrightarrow (k|m \wedge (k+1)|m))$$

which holds if and only if $n = k(k+1)$. This is then used with

$$p_2(n,k,l) := (k+l)(k+l+1) = k(k+1) + l(l+1) + n + n$$

which holds if and only if $n = kl$.

To interpret $(\mathbf{N}, +, |)$ in F, we will need a definable set in F with a map to $\mathbf{N}$. The set will be

$$B_{(\geq 1)} = \{f \in F \mid 0f'_+ = 1f'_- \geq 1\}, \tag{34.1}$$

and the map from $B_{(\geq 1)}$ to $\mathbf{N}$ will be $f \mapsto \log_2(0f'_+)$. This will require a sequence of definitions which, as might be expected, will be focused on what happens near 0 and 1.

We will use the formula α (33.7)) in the structure $(F, \times, ^{-1}, 1)$ as given in Section 33.2, as well as the defined set K based on β (33.2) consisting of directional elements of F. We also make use of the fact that we can effectively quantify over defined sets, and that for definable sets in the universe of a structure, the relations of equality, containment and proper containment of sets are first order expressible.

We cannot differentiate the end of $[0,1]$ near 0 and the end near 1, but given $d \in K$, we can constently refer to the two ends of $[0,1]$ as being "to the left" or "to the right" with respect to d. This gives us the ability to inspect properties an element of F at the two ends. For $d \in K$, we can define

$$F^d_{(L1)} = \{f \mid \exists g, \alpha(g,f,d)\}, \text{ and}$$
$$F^d_{(R1)} = \{f \mid \exists g, \alpha(f,g,d)\}.$$

Assuming that d has positive direction, a characterization of the elements in $F^d_{(L1)}$ and $F^d_{(R1)}$ is

$$F^d_{(L1)} = \{f \mid f \neq 1 \wedge 0f'_+ = 1\}, \text{ and}$$
$$F^d_{(R1)} = \{f \mid f \neq 1 \wedge 1f'_- = 1\},$$

where f'_+ and f'_- are the right and left derivatives, respectively.

However $F^d_{(L1)} \cap F^d_{(R1)}$ is independent of which $d \in K$ is used. So we are justified in writing

$$F_{(1)} = F^d_{(L1)} \cap F^d_{(R1)} = \{f \mid f \neq 1 \wedge 0f'_+ = 1f'_- = 1\}.$$

This symmetry will apply in several instances below and in such cases we will not define $F^d_{(Rx)}$ for various symbols x, and let the reader fill in the details for that and $F_{(x)}$. As above, we will give descriptions of $F^d_{(Lx)}$ assuming that d has positive direction.

Note that $F' = [F,F] = F_{(1)} \cup \{1\}$, but given that every element of F' is a product of two or fewer commutators (Proposition 5.13), we also have $F' = \{f \mid \exists w, \exists x, \exists y, \exists z, f = [w,x][y,z]\}$. Since $\{1\}$ is definable, F' is definable in two ways.

The next definitions are of subgroups near the endpoints of $(0,1)$ and not of behavior at the endpoints. Given $d \in K$, $f \in F^d_{(L1)}$ and $g \in F^d_{(R1)}$, we define

$$\begin{aligned} F^d_{[fL]} &= \{h \mid \alpha(h,f,d)\}, \text{ and} \\ F^d_{[gR])} &= \{h \mid \alpha(g,h,d)\}, \end{aligned}$$

Assuming that d has positive direction, if s is the infemum of the support of f and t is the supremum of the support of g, then $F^d_{[fL]}$ is the subgroup of F, isomorphic to F, of those elements whose support is $[0,s]$ and $F^d_{[gL]}$ is similar with support $[t,1]$. Not only are these groups defined by s and t, these groups also define s and t. This is used next where we show that we can define

$$F^d_{(L<1)} = \{f \mid 0f'_+ < 1\}, \quad \text{and} \quad F^d_{(R<1)} = \{f \mid 1f'_- < 1\}.$$

Note that $0f'_+ < 1$ if and only if f carries every sufficiently small interval $(0,s)$ properly into itself, and for $1f'_- < 1$, every sufficiently small interval $(t,1)$ is carried properly into itself. This discussion was for the case where d has positive direction. We set

$$F^d_{(L<1)} = \{f \mid \exists g \in F^d_{(L1)}, \forall h \in F^d_{(L1)}, (\alpha(h,g,d) \Rightarrow f^{-1}F^d_{[hL]}f \subsetneq F^d_{[hL]})\},$$

with parallel definition of $F^d_{(R<1)}$, and we set $F_{(<1)} = F^d_{(L<1)} \cap F^d_{(R<1)}$.

We set $F^d_{(L>1)} = \{f \mid f^{-1} \in F^d_{(L<1)}\}$ and $F^d_{(R>1)} = \{f \mid f^{-1} \in F^d_{(R<1)}\}$, and set $F_{(>1)} = F^d_{(L>1)} \cap F^d_{(R>1)}$.

The six sets $F^d_{(L1)}$, $F^d_{(L>1)}$, $F^d_{(L<1)}$, $F^d_{(R1)}$, $F^d_{(R>1)}$, $F^d_{(R<1)}$ give independent control over what is happening at the two ends of $[0,1]$.

Since we can pick out slopes above 1 from other slopes, we can restrict what we can increase slopes by. So we can define

$$F^d_{(L2)} = \{f \mid 0f'_+ = 2\},$$

with the use of $F^d_{(L\geq 1)} = F^d_{(L1)} \cup F^d_{(L>1)}$ by setting

$$F^d_{(L2)} = \{f \in F^d_{(L>1)} \mid \forall g \in F^d_{(L>1)}, \exists h \in F^d_{(L\geq 1)}, g(fh)^{-1} \in F^d_{(L1)}\},$$

with parallel definition for $F^d_{(R2)}$, and we set $F_{(2)} = F^d_{(L2)} \cap F^d_{(R2)}$. In words, $0f'_+ = 2$ if and only if multiplying $0f'_+$ by powers of 2 with non-negative exponents yields all powers of 2 with positive exponents.

Recall from Definition 5.15 that a dyadic orbital of $f \in F$ is a component of the support of f to which all non-dyadic isolated fixed points have been added. We can pick out those elements of F that have exactly one dyadic orbital by

$$F^1 = \left\{ f \mid \neg \Big(\exists d \in K, \exists g, \exists h, [\alpha(g,h,d) \wedge f = gh] \Big) \right\}.$$

Now $F^1_{(2)} = F^1 \cap F_{(2)}$ is the set of $f \in F$ with one dyadic orbital and with $0f'_+ = 1f'_- = 2$. Because of the slopes at 0 and 1, such an f must have at least one fixed point in $(0,1)$ which must necessarily not be dyadic.

From Proposition 5.17 we have $C_F(f) = \langle f \rangle$ for every $f \in F^1_{(2)}$.

We set $F_{(\geq 1)} = F_{(>1)} \cup F_{(1)}$, and now the set $B_{\geq 1}$ from (34.1) is definable by the formula

$$\phi(f) := \exists g \in F_{(2)}, \exists h \in C_F(g) \cap F_{(\geq 1)}, fh^{-1} \in F'.$$

Note that $\exists g$ can be replaced by $\forall g$ with the same result.

As mentioned before, the map from $B_{\geq 1}$ to $\mathbf{N}$ is $f \mapsto \log_2(0f'_+)$. The preimage of $=$ on $\mathbf{N}$ is the restriction of $f \sim g \Leftrightarrow fg^{-1} \in F'$ to $B_{\geq 1}$. The preimage of addition is the restriction of multiplication in F to $B_{\geq 1}$.

To address divisibility, we first get rid of the annoying cases involving 0. We have for a and b in $\mathbf{N}$,

$$a|b \leftrightarrow [a \neq 0 \wedge (b = 0 \vee \exists k \neq 0, b = ak)].$$

So we will work with $B_{>1} = B_{\geq 1} \setminus F'$ and model $\exists k \neq 0, b = ak$.

We have to deal with the fact that $C_F(f) = \langle f \rangle$ fails even for a function whose only dyadic orbital is all of $(0,1)$ if f has proper roots. However by Lemma 34.2 below, for every $f \in B_{>1}$, there is $h \in F$ with $hf^{-1} \in F'$ so that h has only one dyadic orbital that is all of $(0,1)$ and which has no proper roots.

Now we note for f and g in $B_{>1}$ that $0g'_+ = (0f'_+)^k$ for some $k \in \mathbf{Z}$ if and only if

$$\forall h \in F', \exists j \in C_F(hf), gj^{-1} \in F'.$$

Thus the premimage of $|$ is definable which completes the proof. ☐

LEMMA 34.2. *For positive integers m and n, there is $f \in F$ with one dyadic orbital so that $0f'_+ = 2^m$ and $1f'_- = 2^n$, and so that f has no proper root.*

PROOF. Any f with $0f'_+ = 2^m$ has its leftmost break at a point $p_0 = (x_0, y_0)$ with $x_0 < \frac{1}{2}$ and $y_0 = x_0 2^m < 1$. Now for $k \geq 0$, a straightforward calculation gives that the leftmost break of f^k is at $(x_0/2^{mk-m}, x_0/2^{mk})$. But $x_0/2^{mk} = y_0/2^{mk-m} < 1/2^{m(k-1)}$ which even for $k = 2$ is less than $1/2$. But by the transitivity properties of F, for any $m > 0$ there is an $f \in F$ with $0f'_+ = 2^m$ and whose leftmost break has y-coordinate in the interval $(1/2, 1)$.

To get only one dyadic orbital, f can be built to have exactly one isolated fixed point, and that point be non-dyadic. The leftmost break can have y-coordinates in $(1/2, 2/3)$, and the isolated fixed point can be at $2/3$ where the slope of f is $1/4$ on a brief interval. $\square$

35. Bi-interpretations and their consequences

In Lasserre 2014 [128] it is shown that F is bi-interpretable with $(\mathbf{Z}, +, \times)$. We briefly discuss this here and touch on some of the implications.

Whether F is bi-interpretable with $(\mathbf{Z}, +, \times)$ was raised as Question 2 in [5]. In the paragraph before that question, it is pointed out that recursive function theory and the solvability of the word problem in F together guarantee an interpretation of F in $(\mathbf{Z}, +, \times)$. No other structural details about F are needed. See [5] for the relevant references. Because of the extra requirements on a bi-interpretation in Definition 31.5, the existence of interpretations of F in $(\mathbf{Z}, +, \times)$ and the reverse is not enough to give a bi-interpretation.

The argument in [128] uses the guaranteed interpretation of F in $(\mathbf{Z}, +, \times)$ together with a very carefully constructed interpretation of $(\mathbf{Z}, +, \times)$ in F to give the bi-interpretation. Because the interpretation of $(\mathbf{Z}, +, \times)$ in F in [128] requires a much longer argument than not only the proof of Proposition 32.2 and but also of Theorem 34.1, and because the reverse interpretation uses little from the structure of F, we only make a few comments about the proof and refer the reader to [128] for details.

The proofs of Proposition 32.2 and Theorem 34.1 make use of the properties of centralizers in F that follow from Proposition 5.17. These properties are also used in [128] as well as the finer control that double centralizers give which follow from Corollary 5.17.1. The arguments in [128] factor through $\mathbf{Z} \wr \mathbf{Z}$ as they do in the arguments of Altinel-Muranov 2009 [5]. The bi-interpretation in [128] is with parameters. That it is not possible to have a bi-interpretation without parameters follows from Lemma 31.6(IV).

Also shown in [128] is that Thompson's group T is bi-interpretable with $(\mathbf{Z}, +, \times)$. As one step, it is shown that F is definable in T, a fact stated, but not shown in [5].

35.1. Consequences. We give some consequences of Lasserre's results.

If $\mathscr{A} = (A, \Sigma)$ is a structure, then we refer to the cardinality of A as the cardinality of $\mathscr{A}$. For a cardinality κ, a theory T is *κ-categorical* if there is only one isomorphism class of models for T of cardinality κ. An L-structure $\mathscr{A}$ is κ-categorical if has cardinality κ and its elementary theory is κ-categorical. That is, any L-structure of cardinality κ with the same elementary theory as $\mathscr{A}$ is isomorphic to $\mathscr{A}$.

We will focus on the countably infinite cardinal $\aleph_0$. The following is found as part of Theorem 7.3.1 in [112].

Theorem 35.1. *A complete theory T over a countable signature is $\aleph_0$-categorical if and only if for some (and thus every) countable model $\mathscr{A} = (A, \Sigma)$ for T the action of* $\mathrm{Aut}(\mathscr{A})$ *on A is oligomorphic.*

Definition 35.2. An action of a group G on a set X is *oligomorphic* if for each $n \geq 1$, the coordinatewise action of G on X^n has finitely many orbits.

The action of $\mathrm{Aut}(G)$ on a group G that has an infinite, finitely generated subgroup cannot be oligomorphic. For if $(g_1, \ldots, g_n)$ generate an infinite subgroup $S \leq G$, then $\{(g_1, \ldots, g_n, s) \mid s \in S\}$ represent infinitely many orbits in G^{n+1}. In particular any group with a non-torsion element does not have an $\aleph_0$-categorical theory.

Thus we know that there is a countable group not isomorphic to F that has the same elementary theory as F. What such a group looks like is not clear. However from the discussion that follows, such an example cannot be finitely generated.

Weaker than $\aleph_0$-categorical on a finitely generated, countably infinite group G is *quasi-finitely axiomatizable* (QFA). A group G is QFA if there is a set S of finitely many sentences (equivalently, one sentence) in the elementary theory of G so that any finitely generated group that satisfies S is isomorphic to G.

Stronger than two groups having the same elementary theory, is for there to be an elementary embedding from one to the other. A function $f : G \to H$ is an elementary embedding of for every formula in the signature of groups, a tuple in G satisfies the formula if and only if the image of the tuple satisfies the formula in H. Note that applying this to formulas with no free variables shows that an elementary embedding induces an isomorphism on the elementary theories. Also, preserving the formula $x \neq y$ shows that an elementary embedding is truly an embedding. A group G is prime if it has an elementary embedding into every group that is a model for $\mathrm{Th}(G)$.

The relevance of the terms just defined is the following.

THEOREM 35.3. *A finitely generated group that is bi-interpretable with* $(\mathbf{Z}, +, \times)$ *is both QFA and prime.*

This is announced in Khelif 2007 [123] as Lemma 1. Proofs of QFA are given in Proposition 2.28 of Aschenbrenner-Khélif-Naziazeno-Scanlon 2020 [6] and Theorem 10 of Kharlampovich-Myasnikov-Sohrabi 2022 [122]. Proof of prime is in Corollary 6 of [122].

From Lassere's result, we have the following.

THEOREM 35.4. *The Thompson groups* F *and* T *are QFA and prime.*

The group T was the first simple group known to have these properties.

For each $G \in \{F, T\}$ if we set $\mathscr{M}(G)$ to be the set of countable models of $\mathrm{Th}(G)$, then there are groups in $\mathscr{M}(G)$ not isomorphic to G, but for every $H \in \mathscr{M}(G)$, there is an elementary embeddeing of G into H.

The paper [122] introduces the concept *rich*, an invariant property of bi-interpretability. A group that is rich has a first order theory that has the strength of a theory (called weak second order theory) that is strictly between first order theory and second order theory. We refer the reader to [122] for details. The ring Z is rich, and so F and T are both rich.

Some relations are known between the properties just discussed. Neither QFA nor prime is equivalent to being bi-interpretable with $(\mathbf{Z}, +, \times)$. In Nies 2003 [164], Theorem 5.1 says that the Heisenberg group $UT_3^3(\mathbf{Z})$ (the group of 3×3 upper triangular matrices over $\mathbf{Z}$ with ones on the diagonal) is QFA, and Corollary 5.3 (of a "well known" criterion) says that $UT_3^3(\mathbf{Z})$ is prime. In Nies 2007 [165], Theorem 7.16 shows that $UT_3^3(\mathbf{Z})$ (called $UT_3(\mathbf{Z})$ in the later paper) is not bi-interpretable with $(\mathbf{Z}, +, \times)$. Corollary 3.2 of [165] is that prime does not imply QFA, and it is commented that as of that paper whether QFA implies prime is an open question. Private communication from Myasnikov says that $UT_3(\mathbf{Z})$ is not rich and that the proof is not obvious.

36. End notes

An early example of interpretation rather than definition of one structure in another is in Mal′cev 1960 [146]. The paper can also be found as Chapter 15 of Wells 1971 [145].

A very different class of groups is shown in Giraudet-Glass-Truss [83] to have undecidable theories by interpreting the integers.

Reconstruction theorems (deducing an action and the space acted upon from the algebra of the group) have a history with specialized results applied to spaces with specialized structures. Some history is found in Rubin 1989 [171] at the end of the introductory section, and in Section 1.8 of Rubin-Yomdin 2005 [172]. Hasson-Bennet 2018 [101] is a Rubin memoriam and gives some history of Rubin's part of the subject. A paper by Mann and Wolff 2019 [147] has reconstruction theorems that cover techniques of Rubin and other techniques. The recent paper of Koberda-de la Nuez González 2025 [124] is in a very recent paper in the area. Reconstruction in $PL_+(I)$ is done from a different approach in Hyde-Moore 2023 [115].

The power of the first order language of certain groups such as the Thompson groups shows up in several ways. The rich groups of Kharlampovich-Myasnikov-Sohrabi 2022 [122] (which has an informative introduction) are a class of groups where the first order theory has power strictly between that of first order theory and second order theory. Investigations for groups of homeomorphisms of manifolds are in Koberda-de la Nuez González 2023–4 [126, 127].

The expressiveness of first order formulas is exploited in Brin 2005 [29], Taylor 2017 [188], and Bleak-Brin-Moore 2021 [21], although the results there cannot be called first order results.

CHAPTER 6

Variants of Thompson's groups

Chapter contents

37. Introduction

[1]Lack of familiarity with Chapter 2 will make this chapter difficult to read.

We can modify the definitions of F, T and V to produce variants that share properties with the original F, T and V and are thus members of the Thompson family. We can combine groups that are already in the Thompson family with other groups to produce marriages that also share properties with the members of the Thompson family. We only consider countable groups. Among desired properties are strong control over normal subgroups (simplicity would be nice) and some level of finiteness. Finitely presented is preferred over finitely generated.

We will give a sampling of variants and marriages. We will omit many details, in part because some details are similar to those in Chapter 2, and in part because as the family grows, so does the amount of material to be covered.

Thompson groups can be derived from actions on spaces, from combinatorial structures, and from dynamical systems. These derivations

[1]This chapter is not complete. More material will be added in the future.

can all be varied, inviting a search for unification. Some of these have been mentioned in then end notes (Section 14) of Chapter 2.

38. The Higman and Brown variants

In Section 11.6, pairs of finite forests are introduced. It is shown that pairs of binary forests, either finite or finitary, all ultimately lead to a group isomorphic to F. We will see that the results are different when the same exercise is done for groups like T or V.

In this section we look at two changes to the definitions in Chapter 2. We change both the number of trees, and the arity of the trees. We fix integers $n \geq 2$ and $r \geq 1$, and we look at pairs of finite n-ary forests with each forest having r trees (r roots) and with a bijection between the leaves of the two forests. If the bijection is order preserving, then we get an "F-like" group denoted $F_{n,r}$. If the bijection preserves the cyclic order, we get a "T-like" group denoted $T_{n,r}$, and for arbitrary bijections, the result is a "V-like" group denoted $V_{n,r}$. For the F-like groups, we can also let $r = \infty$ to get $F_{n,\infty}$. There will be some dependence of the isomorphism types on the parameters, and a certain amount of independence. We will discuss what properties are shared with the groups F, T and V.

Note that Chapter 2 starts with homeomorphisms and works towards trees and forests. Here we start with trees and forest and work back towards homeomorphisms.

38.1. The main parameters.

Definition 38.1. Let $n \geq 2$ and $r \geq 1$ be integers, and also allow $r = \infty$. A *forest pair* of type (n, r) is a triple (Φ, σ, Ψ) where Φ and Ψ are n-ary forests (sequences of n-ary trees) of length r and σ is a bijection from the leaves of Φ to the leaves of Ψ. The leaves of each forest are ordered with the prefix order as in Definition 11.3. If σ is omitted or $r = \infty$, then the bijection is assumed to be the unique order preserving bijection.

At this point n-ary splittings on the forests, and matched n-ary splittings on the pairs can be defined as in Sections 6 through 8, together with matched n-ary refinements and the equivalence classes that they generate. Putting opportunistic multiplication on the classes results in the groups $F_{n,r}$, $T_{n,r}$ and $V_{n,r}$, as mentioned above, depending on the nature of the bijections σ.

We limit the dependence of $F_{n,r}$, $T_{n,r}$ and $V_{n,r}$ on r for finite r. In Proposition 38.3 below, we will completely will eliminate the dependence of $F_{n,r}$ on r for $1 \leq r \leq \infty$. For finite r, we explode trees.

If T is a finite n-ary tree over the alphabet $\{0,\ldots,n-1\}$, then $E(T)$ is the sequence of length n consisting of the subtrees $(T_0, T_1, \ldots, T_{n-1})$ with the subtrees denoted as in Definition 8.3. If $\Phi = (\Phi_0, \ldots, \Phi_{r-1})$ is a finite forest of n-ary trees of length r, then we explode the last tree of Φ to get $E(\Phi) = (\Phi_0, \ldots, \Phi_{r-2}, E(\Phi_{r-1}))$ of length $r + (n - 1)$. And if (Φ, σ, Ψ) is a forest pair of type (n, r) then $E(\Phi, \sigma, \Psi) = (E(\Phi), \sigma', E(\Psi))$ is a forest pair of type $(n, r + (n-1))$ where σ' agress with σ on the leaves of the trees in Φ except for the last tree and is defined on the leaves of $E(\Phi_{r-1})$ to commute with the natural inclusions of the trees in $E(\Phi_{r-1})$ and $E(\Psi_{r-1})$ into Φ_{r-1} and Ψ_{r-1}, respectively.

Lemma 38.2. *The function E from forests pairs of type (n, r) to forest pairs of type $(n, r + (n - 1))$ commutes with matched binary splittings and induces isomorphisms $G_{n,r} \to G_{n,r+(n-1)}$ for G any one of F, T, or V.*

The full story of the dependence of the groups $T_{n,r}$ and $V_{n,r}$ groups on the parameters is more complicated and only partly known for the $T_{n,r}$. This will be discussed later. We first discuss the groups $F_{n,r}$.

38.2. Combinatorics and the $F_{n,r}$.

38.2.1. *Isomorphism types.* For $1 \le r \le \infty$, we show that the isomorphism types of the groups $F_{n,r}$ depend only on n and not r. We get presentations for the groups and review consequences of the presentation. The ability to move back and forth between $F_{n,r}$ and $F_{n,s}$ for various r and s will prove very useful.

We can successively apply the operation E to a finite forest. Since each application changes only the last tree, there is a stable notion of applying E infinitely many times. Since our trees are finite (have only finitely many vertices), then for any finite forest Φ, the result is a finitary forest $E^\infty(\Phi)$ (all but finitely many trees are trivial) of n-ary trees. Now for a forest pair (Φ, Ψ) (the omitted bijection is assumed order preserving) of type (n, r) with $r < \infty$, the pair $E^\infty(\Phi, \Psi) = (E^\infty(\Phi), E^\infty(\Psi))$ is a finitary forest pair of n-ary trees. With the usual definitions of matched n-ary splittings, equivalence classes and multiplication on the classes, we get a group $F_{n,\infty}$ based on pairs of finitary forests of n-ary trees and order preserving bijections between the sets of leaves. We have the following.

Proposition 38.3. *Fix an integer $n \ge 2$. For each $r \ge 1$, the operation E^∞ induces an isomorphism from $F_{n,r} \to F_{n,\infty}$. Thus with n fixed, all the $F_{n,r}$ are isomorphic.*

The notations $F_{n,r}$ and $F_{n,\infty}$ still have use, and F_n will ambiguously refer to any of these isomorphism types.

38.2.2. *Fractions, a presentation and consequences.* In a fintary n-ary forest, let us call a vertex *internal* if it has children, and let us call a leaf trivial if it is also a root (and thus is in a trivial tree). Induction says that if $k > 0$ is the number of internal vertices, then the number of non-trivial leaves is $k(n-1)+1$. The case $k = 0$ is an exception which can only be cured by declaring that the forest of trivial trees has one non-trivial leaf.

Building F_n from the pairs of finitary n-ary forests will give a presentation for F_n. The finitary n-ary forests form a monoid $\mathcal{F}_n$ as done in Sections 11.2.3 and 11.2.4 for $\mathcal{F}$. The method for describing generators of $\mathcal{F}_n$ is useful when dealing with both finite and finitary forests.

Let Φ be an n-ary forest, finite or finitary. With $k \leq \infty$ the number of leaves, with $i < k$, and with leaves numbered in order starting at 0, let $\Phi_{\downarrow i}$ be the result of an n-ary splitting of Φ at leaf i. With X_n the trivial, finitary, n-ary forest, the generators of $\mathcal{F}_n$ are the $X_{n\downarrow i} = (X_n)_{\downarrow i}$. Specifically, $X_{n\downarrow i}$ is a finitary, n-ary forest with a single non-trivial tree having one internal vertex and where that tree is the i-th tree of the forest.

Proposition 38.4. *The group $F_{n,\infty}$ is a group of fractions of $\mathcal{F}_n$. The presentation*

$$\langle X_{n\downarrow 0}, X_{n\downarrow 1}, \dots \mid X_{n\downarrow j} X_{n\downarrow i} = X_{n\downarrow i} X_{n\downarrow j+n-1} \text{ whenever } i < j \rangle \tag{38.1}$$

is a monoid presentation for $\mathcal{F}_n$. Using $g_{n,i}$ as a formal symbol, a corresponding group presentation for $F_{n,\infty}$ and for F_n is

$$\langle g_{n,0}, g_{n,1}, \dots \mid g_{n,j} g_{n,i} = g_{n,i} g_{n,j+n-1} \text{ whenever } i < j \rangle. \tag{38.2}$$

Proof. That $F_{n,\infty}$ is a group of fractions is established as done for F in Section 11.2.4 through Corollary 11.6.3. The argument that (38.1) is a presentation for $\mathcal{F}_n$ can be copied from the proof of Proposition 11.7. The last claim is standard from the first two points and because F_n represents the isomorphism class of $F_{n,\infty}$. □

Corollary 38.4.1. *The group F_n is generated by the n elements in $A_n = \{g_{n,i} \mid 0 \leq i < n\}$. The abelianization of F_n is $\mathbf{Z}^n$, and F_n is isomorphic to F_m if and only if $n = m$. The abelianization homomorphism can be described by letting $\mathbf{Z}^n$ be the free abelian group on the integers modulo n and sending $g_{n,0}$ to 0, and each $g_{n,i}$, $i \in \mathbf{Z}_{\geq 1}$, to $j \in \{1, \dots, n-1\}$ with $j \equiv i \pmod{n-1}$.*

Proof. Each generator in (38.1) is conjugate by a power of $g_{n,0}$ to an element of A_n. All the relations in (38.1) are conjugacy relations and add no information if the generators are declared to commute. □

Arguing as in Lemma 9.5, we see that a seminormal form for words in the $g_{n,i}^{\pm 1}$ is

$$g_{n,i_0} \cdots g_{n,i_k} g_{n,j_l}^{-1} \cdots g_{n,j_0}^{-1}, \quad i_0 \le \cdots \le i_k \ne j_l \ge \cdots \ge j_0. \tag{38.3}$$

LEMMA 38.5. *Every non-trivial normal subgroup of F_n contains the commutator subgroup.*

PROOF. The proof can be copied from that of Proposition 9.7 up to the last argument that shows that $[x_0, x_1] = 1$ modulo any non-trivial normal subgroup in F. That argument can be adapted to show that $[g_{n,i}, g_{n,j}] = 1$ for all $0 \le i < j < n$ modulo any non-trivial normal subgroup in F_n. □

More information about the commutator subgroup will have to wait until Section 38.3 where we represent the F_n as groups of homeomorphisms.

Let $\sigma_n : F_n \to F_n$ be defined by $\sigma_n(g_{n,i}) = g_{n,i+1}$. This is an endomorphism of F_n, and $\sigma_n^n = C_0\sigma_n$ composing right to left and with C_0 conjugation by $g_{n,0}$. The following is a parallel to the material in Section 10.2. Details are left to the reader.

PROPOSITION 38.6. *The triple $(F_n, \sigma_n, g_{n,0})$ is an initial object in the category of groups with an endomorphism whose first and n-th powers differ by an inner automorphism. Further if η is a morphism in that category from (F, σ, x_0) to some (G, ϕ, g), then either η is an injection, or the image of η is abelian.*

We also have a parallel to Section 10.3. For an integer $k \ge 0$, if we let $F_{n,\ge k}$ be generated by the $g_{n,i}$ with $i \ge k$, then $F_{n,\ge k}$ is a subgroup of $F_n = F_{n,\ge 0}$ isomorphic to F_n which is the image of F_n under the endomorphism σ_n^k. If we use $k = n-1$, then σ_n^{n-1} carries $F_{n,\ge 1}$ isomorphically to $F_{n,\ge n}$ which is exactly what conjugation by $g_{n,0}$ does. Thus F_n is the HNN extension of $F_{n,\ge 1}$ using σ_n^{n-1} and stable letter $g_{n,0}$. We have that F_n is an ascending HNN extension of itself.

38.3. Actions. Most of the information in Chapter 2 about F came from its existence as a group of homeomorphisms. Here we resurrect this point of view for the F_n. Parallel to the treatment of the groups using pairs of n-ary forests, we also have the treatment with pairs of n-ary partitions and thus also piecewise linear homeomorhpisms. We have the following.

LEMMA 38.7. *The group $F_{n,r}$ is naturally isomorphic the set of self homeomorphisms of $[0, r]$, under composition, that are piecewise linear,* *have slopes restricted to integral powers of* n*, and whose breakpoints are confined to the elements of* $\mathbf{Z}[\frac{1}{n}]$*.*

PROOF. The proof can be copied with trivial modifications from the proof of Proposition 6.14 after associating the i-th tree of a finite n-ary forest of length r to an n-ary subdivision of the interval $[i, i+1]$ in $[0, r]$. □

If we now associate a finitary n-ary forest with a sequence of n-ary subdivisions of the intervals $[i, i+1]$ in $[0, \infty)$, we see that a finitary n-ary forest pair leads to a self homeomorphism of $[0, \infty)$ that is a shift by a multiple of $n-1$ near infinity. We have the following.

LEMMA 38.8. *The group* $F_{n,\infty}$ *is naturally isomorphic to the set of self homeomorphisms* f *of* $[0, \infty)$ *under composition that are piecewise linear, have slopes restricted to integral powers of* n*, whose breakpoints are confined to* $\mathbf{Z}[\frac{1}{n}]$*, and for which there are integers* j *and* k *so that* $t > j$ *implies* $tf = t + k(n-1)$*.*

Turning the $g_{n,i}$ into piecewise linear homeomorphisms on $[0, \infty)$ gives generators $y_{n,i}$ for the PL version $F_{n,\geq 0}$ of $F_{n,\infty}$. The $y_{n,i}$ are defined for non-negative integers i as follows.

$$ty_{n,i} = \begin{cases} t, & t \leq i, \\ n(t-i)+i, & i \leq t \leq i+1, \\ t+n-1, & t \geq i+1. \end{cases} \tag{38.4}$$

If we temporarily define $z_{n,i} = y_{n,i-1}$ for all integers $i \geq 1$, and define $z_{n,0}$ by setting $tz_{n,0} = t+(n-1)$, then we see that $z_{n,j}z_{n,i} = z_{n,i}z_{n,i+(n-1)}$ holds for all $i < j$. Thus the non-abelian group $F_{n,\mathbf{R}}$ generated by the $z_{n,i}$ is a quotient of F_n and thus isomorphic to F_n.

LEMMA 38.9. *The group* $F_{n,\mathbf{R}}$ *consists of those self homeomorphisms* f *of* $\mathbf{R}$ *under composition that are piecewise linear, have slopes restricted to integral powers of* n*, whose breakpoints are confined to* $\mathbf{Z}[\frac{1}{n}]$*, and for which there are integers* j *and* k*,* p *and* q *so that* $t > j$ *implies* $tf = t + k(n-1)$ *and* $t < p$ *implies* $tf = t + q(n-1)$*.*

38.4. Generators for the $F_{n,r}$. From Lemmas 38.7 and 38.8 we see that the elements of $F_{n,r}$ viewed as acting on $[0, r]$ are the elements of $F_{n,\infty}$ viewed as acting on $[0, \infty)$ that fix all points outside $[0, r]$. This embedding of $F_{n,r}$ into $F_{n,\infty}$ and the isomorphism using E^∞ of Proposition 38.3 between $F_{n,r}$ and $F_{n,\infty}$ will help us understand generators of $F_{n,r}$. This will also help to understand the $T_{n,r}$.

It is useful to know for which finite forests Φ is $E^\infty(\Phi)$ the trivial forest. Generalizing the vines of Definition 8.13, we say that an n-ary tree is a vine if the number of internal vertices all of whose children are leaves is one. The right vine is the vine in which every internal vertex is the rightmost child of its parent. To avoid notational conflict with the groups $V_{n,r}$, we will use $R_{n,k}$ to denote the right n-ary vine with k internal vertices. It is now clear that a finite n-ary forest Φ has $E^\infty(\Phi)$ the trivial infinite forest if and only if every tree in Φ is trivial except possibly the last which must then be a right vine.

From the presentation (38.2) we know that $\{g_{n,i} \mid 0 \le i < n\}$ is a generating set for $F_{n,\infty}$. Using the notation of Proposition 38.4, we can express $g_{n,i}$ as the finitary forest pair $(X_{n\downarrow i}, X_n)$. If Φ_r is the finite, trivial n-ary forest of length r, and $i < r-1$, then $X_{n\downarrow i} = E^\infty(\Phi_{r\downarrow i})$, and $X_n = E^\infty(\Phi_{r\downarrow(r-1)})$. Further $\Phi_{r\downarrow i}$ and $\Phi_{r\downarrow(r-1)}$ have the same number of leaves. To generate $F_{n,r} \simeq F_{n,\infty}$, we need $g_{n,i}$ with $0 \le i < n$ up to $n-1$, so it will be convenient to continue under the assumption $r > n$. The reader can do the alterations needed for smaller values of r.

38.4.1. *Glides.* We now see that for $r > n$, the finite, n-ary forest pairs $(\Phi_{r\downarrow i}, \Phi_{r\downarrow(r-1)})$ generate a group isomorphic to $F_{n,r}$ as defined in Section 38.1, and which is isomorphic using E^∞ to the subgroup of $F_{n,\infty}$ generated by the pairs $(X_{n\downarrow i}, X_{n\downarrow(r-1)})$ for $0 \le i < n$.

For $i < j$, pairs such as $(X_{n\downarrow i}, X_{n\downarrow j})$ and $(\Phi_{r\downarrow i}, \Phi_{r\downarrow j})$ are called *glides* in Section 4.10 of [34], and given the notation γ_{ij} or $\gamma_{i,j}$. We will use this notation ambiguously for elements of $F_{n,\infty}$, and for elements of $F_{n,r}$ when $j < r$. In terms of the generators of $F_{n,\infty}$, we have $\gamma_{ij} = g_{n,i}g_{n,j}^{-1}$. For $i < j < k$, we have $\gamma_{ij}\gamma_{jk} = \gamma_{ik}$. As in [34], we will use γ_i for $\gamma_{i,i+1}$ giving

$$\gamma_{ij} = \gamma_i\gamma_{i+1}\cdots\gamma_{j-1}. \tag{38.5}$$

For $r > n$, we see from (38.5) that $F_{n,r}$ is generated by $\{\gamma_i \mid 0 \le i < r-1\}$.

The support of γ_i in $F_{n,\infty}$ acting on $[0,\infty)$ or any $F_{n,r}$ acting on $[0,r]$ with $r \ge i+2$ is $[i, i+2]$. From Corollary 38.4.1, the product $p = \gamma_1\gamma_2\cdots\gamma_{n-1} = \gamma_{1,n} = g_{n,1}g_{n,n}^{-1}$ is in the commutator subgroup of F_n.

38.5. Transitivity. The various actions given above of F_n are not transitive on $\mathbf{Z}[\frac{1}{n}]$ for $n > 2$. We show that the orbits on $\mathbf{Z}[\frac{1}{n}]$ are separated by residues modulo $n-1$.

An element t of $\mathbf{Z}[\frac{1}{n}]$ is of the form $t = p/n^q$ for integers p and q. Since $n \equiv 1 \pmod{n-1}$, setting $\phi_n(t) \equiv p \pmod{n-1}$ has ϕ_n well defined in that it is independent of the chosen p and q, and makes ϕ_n a ring homomorphism from $\mathbf{Z}[\frac{1}{n}]$ to $\mathbf{Z}/(n-1)\mathbf{Z}$. We let Δ_n denote the kernel of ϕ_n. The reader can show the following.

Lemma 38.10. *In the actions of F_n on $(0,1)$, of $F_{n,r}$ on $(0,r)$, of $F_{n,\geq 0}$ on $(0,\infty)$, and of $F_{n,\mathbf{R}}$ on $\mathbf{R}$, two elements s and t of $\mathbf{Z}[\frac{1}{n}]$ in the given domain are in the same orbit of the action if and only if $\phi_n(s) = \phi_n(t)$. For all of these actions and for each $k \geq 1$ and $t \in \mathbf{Z}[\frac{1}{n}]$, the action on the intersection of the coset $(t + \Delta_n)$ with the domain of the action is o-k-transitive.*

We can give more information than Lemma 38.10 with the following parallel to Lemma 5.3. We state it for the action of $F_{n,\mathbf{R}}$ on $\mathbf{R}$, but it applies with appropriate modifications to all the actions in Lemma 38.10.

Lemma 38.11. *Let A and B be finite subsets of $\mathbf{Z}[\frac{1}{n}]$ with $|A| = |B|$ and let X be a collection of closed intervals with pairwise disjoint interiors in I so that each $J \in X$ has its endpoints in A and no points of A in its interior. Let σ be an order preserving bijection from A to B so that $\phi_n(a\sigma) = \phi_n(a)$ holds for all $a \in A$, and for each $J \in X$ let g_J be the restriction to J of an element of $F_{n,\mathbf{R}}$ which agrees with σ on the endpoints of J. Then there is an element $f \in F_{n,\mathbf{R}}$ with bounded support so that f agrees with σ on A and for each $J \in X$, f agrees with g_J on J.*

38.5.1. *A more transitive group.* There is a Thompson-like group that is more transitive. Consider the group $\widetilde{F}_n$ generated by the $y_{n,i}$ and the function $t \mapsto t+1$. The elements are piecewise linear, have slopes integral powers of n, and have breakpoints confined to $\mathbf{Z}[\frac{1}{n}]$. But near near $\pm\infty$, the action is translation by integers and not just multiples of $n-1$. The abelianization of this group is $\mathbf{Z}^2$, and so the group cannot be isomorphic to any F_n for $n \geq 3$. The group is fully transitive on all of $\mathbf{Z}[\frac{1}{n}]$.

To understand the relation $F_n \leq \widetilde{F}_n$, note that the for every generator of $\widetilde{F}_n$, the translations near ∞ and $-\infty$ differ by an integral multiple of $n-1$. So if for $f \in \widetilde{F}_n$, we let $\tau_+(f)$ be the translation by f near ∞ and $\tau_-(f)$ be the translation by f near $-\infty$, then $\tau_+(f) \equiv \tau_-(f) \pmod{n-1}$. So $f \mapsto \tau_+(f)$ is a homomorphism from $\widetilde{F}_n$ to $\mathbf{Z}/(n-1)\mathbf{Z}$ whose kernal is F_n, and $|\widetilde{F}_n : F_n| = n-1$.

The action of $\widetilde{F}_n$ on $\mathbf{R}$ invites the use of forests indexed over $\mathbf{Z}$ rather than $\mathbf{N}$. If this is done, then it is no longer sufficient to rely on the total order on the leaves to determine the bijection between the leaves if the bijection is omitted. So an origin is needed and instructions on where it is mapped. Binary forests indexed over $\mathbf{Z}$ are developed for F in Belk-Brown 2005 [9]. If the construction in [9] is altered to use n-ary forests, then the result is a combinatorial model for $\widetilde{F}_n$.

38.6. About the F_n. Many properties of F apply to the F_n. To a large degree this is because, like F, the F_n represent faithfully into $PL_+(I)$. We leave the verification to the reader.

We have the following about the local behavior of the action of F_n on I.

LEMMA 38.12. *For an integer $n > 2$, the following hold regarding the action of F_n on $I = [0, 1]$.*

(1) *For every rational r in $(0, 1)$, there is an $f \in F_n$ for which r is an isolated fixed point.*
(2) *If $f \in F_n$ has an isolated rational fixed point r which in reduced terms has a prime factor p of the denomator with p prime to n, and the slope of f at r is n^k, then then $p|(n^k - 1)$.*
(3) *If $f \in F_n$ has an irrational fixed point t, then f is fixed on an open interval about t.*

More globally, we have the following. Arguments can be found in Sections 4.4 through 4.7, 5.2, 5.4 and 9.4.

LEMMA 38.13. *Let $n \geq 2$ be an integer. The group F_n is torsion free, is bi-orderable and locally indicable, has trivial center, has its subgroups closed under wreath product with $\mathbf{Z}$, and has that each non-abelian subgroup contains a copy of $\mathbf{Z} \wr \mathbf{Z}$ and is neither free nor nilpotent. There are infinitely differentiable actions of F_n on $\mathbf{R}$, $\mathbf{R}_{\geq 0}$ and $(0, 1)$.*

About the commutator subgroup, we have the following. Arguments can be found in Sections 5.5 and 10.

LEMMA 38.14. *For each positive integer $n \geq 2$, the first F_n' and second F_n'' commutator subgroups of F_n are equal, simple, not finitely generated, contained in every non-trivial normal subgroup of F_n, and act o-k-transitively on each coset $t + \Delta_n$ with $t \in \mathbf{Z}[\frac{1}{n}]$ and k a positive integer. Further every element of F_n' is a product of two or fewer commutators.*

The discussion in Section 5.6 about the centralizers of elements of F applies with no real change to the F_n. We will not bother to repeat the statements here.

Similarly, the discussion running through Propositions 6.17, 7.5, 8.9, and Section 11.4 about reduced representatives and normal forms also applies with no real change to the F_n. The curious reader can write out the details.

38.6.1. *Mutual embedability.*

LEMMA 38.15. *If $2 \le m < n$, then F_m and F_n homomorphically embed in each other.*

PROOF. All symbols in subscripts below are non-negative integers.

It is easy to show F_m embeds in F_n for all $m < n$. The group F_n is generated by the $g_{n,i}$, $i \ge 0$, of (38.1). We will use some but not all of the $g_{n,i}$ as generators of a copy of F_m in F_n.

Each $j \ge 1$ is uniquely expressed as $j = 1 + q(m-1) + r$ with $0 \le r < m-1$. With j expressed this way, let $j' = 1 + q(n-1) + r$. Let $0' = 0$. Note that $j < k$ implies $j' < k'$.

For each $j \ge 1$, let $f_{m,j} = g_{n,j'}$. Let $f_{m,0} = g_{n,0}$. For brevity, write f_j for $f_{m,j}$ and g_i for $g_{n,i}$. For $j < k = 1 + q(m-1) + r$, we have

$$\begin{aligned} f_k^{f_j} =& g_{k'}^{g_{j'}} = g_{k'+n-1} = g_{1+q(n-1)+r+n-1} \\ =& g_{1+(q+1)(n-1)+r} = f_{1+(q+1)(m-1)+r} \\ =& f_{k+m-1}. \end{aligned}$$

So the $f_j = f_{m,j}$ satisfy the relations of the presentation (38.2) for F_m and we have homomorphism $F_m \to F_n$. Since the f_j do not commute, Lemma 38.14 says this must be a monomorphism.

Since F_2 embeds in every F_n, to show that F_n embeds in every F_m with $m < n$ it suffices to show that F_n embeds in F_2. The group F_2 is generated by the x_i, $i \ge 0$, of (9.1). Let $f_i = x_i^{n-1}$. Now for $i < j$, we have

$$f_j^{f_i} = (x_j^{n-1})^{x_i^{n-1}} = x_{j+n-1}^{n-1} = f_{j+n-1}$$

as required. The comments made at the end of the embeding of F_m into F_n now apply. □

From Lemma 5.25, we have the following.

COROLLARY 38.15.1. *For $n \ge 2$, each F_n has exponential growth.*

38.7. On the $T_{n,r}$. Elements of the group $T_{n,r}$ are represented by pairs (Φ, σ, Ψ) where Φ and Ψ are forests of type (n, r) and σ is a bijection from the leaves of Φ to the leaves of Ψ that preserves the cyclic order of the leaves. As usual, matched n-ary refinements and the resulting equivalence relations and classes are defined and opportunistic multiplication is applied to arrive at a group. If for $0 \le i < r$, we regard the i-th tree in each of the forests Φ and Ψ as defining an n-ary partition of $[i, i+1]$, then we obtain the following.

PROPOSITION 38.16. *The group $T_{n,r}$ acts faithfully on the circle $\mathbf{R}/r\mathbf{Z}$ of length r by piecewise linear, orientation preserving homemorphisms where the slopes are integral powers of n, where the breakpoints are confined to $\mathbf{Z}[\frac{1}{n}]$ and where $\mathbf{Z}[\frac{1}{n}]$ is preserved as a set.*

In the discussion that follows, it will be useful to view the of elements of $T_{n,r}$ both as homeomorphisms and as forest pairs.

38.7.1. *Transitivity and an invariant.* We lift $T_{n,r}$ to $\mathbf{R}$. Let $PL_n(\mathbf{R})$ be those piecewise linear self homeomorphisms of $\mathbf{R}$ whose slopes are integral powers of n, whose breakpoints form a discrete subset of $\mathbf{Z}[\frac{1}{n}]$ and which preserve $\mathbf{Z}[\frac{1}{n}]$ as a set. If $\overline{r}$ is the translation $t\overline{r} = t + r$, and if $C(\overline{r})$ is the centralizer of $\overline{r}$ in $PL_n(\mathbf{R})$, then $T_{n,r}$ is isomorphic to $C(\overline{r})/\langle\overline{r}\rangle$.

Consider the homomorphism $\phi_n : \mathbf{Z}[\frac{1}{n}] \to \mathbf{Z}/(n-1)\mathbf{Z}$ of Section 38.5 and its kernel Δ_n. We refer to $\phi_n(a)$ as the "residue" of $a \in \mathbf{Z}[\frac{1}{n}]$ modulo $n-1$. If Φ is an n-ary forest of length r, then a vertex v in Φ is associated to an interval $[s,t]$ in $[0,r]$. We declare $\phi_n(v)$ to be $\phi_n(s)$. We note that, modulo $n-1$, we have $\phi_n(t) = \phi_n(s) + 1$, and if v has children identified as $v0, v1, \ldots, v(n-1)$ reading from left to right, then $\phi_n(vi) = \phi_n(v) + i$ making $\phi_n(v0) = \phi_n(v(n-1)) = \phi_n(v)$.

In parallel to Lemma 38.10, each $f \in PL_n(\mathbf{R})$ induces a rotation of the cosets of Δ_n in $\mathbf{Z}[\frac{1}{n}]$ and thus a rotation on $\mathbf{Z}/(n-1)\mathbf{Z}$. For $T_{n,r}$, we must take into account the identification of cosets of Δ_n that takes place in $\mathbf{R}/r\mathbf{Z}$, or equivalently under the action of $\overline{r}$. The representatives of Δ_n in $\mathbf{Z}$ are the multiples of $n-1$, and the images of these multiples in $\mathbf{Z}/r\mathbf{Z}$ are the residues $\{j(n-1) \pmod r \mid j \in \mathbf{N}\}$. These are all $j(n-1) + kr$ between 0 and $r-1$ and are thus the multiples in that interval of $d = \gcd(n-1, r)$. There are thus d cosets of Δ_n in $\mathbf{Z}[\frac{1}{n}]/r\mathbf{Z}$.

For $f \in T_{n,r}$ given by $(\Phi, \sigma\Psi)$ we let

$$\theta(f) = (\phi_n(v\sigma) - \phi_n(v)) \pmod d$$

for any leaf v of Φ. This is well defined since σ rotates the order of the leaves. It gives us a normal subgroup of $T_{n,r}$.

LEMMA 38.17. *With $d = \gcd(n-1, r)$, taking $f \in T_{n,r}$ to $\theta(f)$ in $\mathbf{Z}/d\mathbf{Z}$ is a homomorphism whose kernel $T^0_{n,r}$ of θ is a normal subgroup of $T_{n,r}$ of index d in $T_{n,r}$.*

For the rest of our discussion of $T_{n,r}$, we will consistently use d for $\gcd(n-1, r)$.

38.7.2. *Generators and some relations.* It will pay to have enough room to discuss generating sets without having a lot of special cases. From Lemma 38.2, we can alter r modulo $n-1$ so that $r > n$. Note that this makes no change to our invariant d. We will make $r > n$ an occasional assumption below.

The glides γ_i from Section 38.4 act on $\mathbf{R}$ with support in the interval $[i, i+2]$, Since $r > 2$, we can view γ_i as acting on $\mathbf{R}/r\mathbf{Z}$ and thus as an element of $T_{n,r}$. We have $\gamma_i = \gamma_j$ if and only if $i \equiv j \pmod r$, giving us r different glides of the form γ_i.

With $T_{n,r}$ acting on $[0, r]$ with the endpoints identified, the subgroup of $T_{n,r}$ fixing 0 can be viewed as $F_{n,r}$ which, from Section 38.4, is generated by the $(r-1)$ glides in $\{\gamma_i \mid 0 \le i < r-1\}$.

We introduce some rotations. They start with the rotation ϱ of order r acting as $t\varrho = t + 1 \pmod r$ on the circle $\mathbf{R}/r\mathbf{Z}$, but we need more. Let $\Phi_{r,0} = \Phi_r$, the finite, trivial n-ary forest of length r from Section 38.4. Let $\Phi_{r,k+1}$ be the result of doing an n-ary splitting on the rightmost leaf of $\Phi_{r,k}$. Specifically $\Phi_{r,k}$ starts with $r-1$ trivial trees, and the last tree is the right vine $R_{n,k}$. Let ϱ_k be given by the pair $(\Phi_{r,k}, \sigma_k, \Phi_{r,k})$ where σ_k takes leaf i to leaf $i+1$ in the usual numbering. We have $\varrho_0 = \varrho$.

Lemma 38.18. *(I) The following relations hold where subscripts of the γ's are treated modulo r.*

(i) $\varrho^{-1}\gamma_i\varrho = \gamma_{i+1}$.
(ii) $\gamma_{r-2}^{-1}\varrho = \varrho_1$.
(iii) $\gamma_{r-1}\gamma_0\gamma_1 \cdots \gamma_{r-2} = (\varrho_1)^{n-1}$.
(iv) $\varrho^{n-1} = \gamma_{r-n}\gamma_{r-n+1} \cdots \gamma_{r-1}\gamma_0 \cdots \gamma_{r-2}$.

(II) For $r > n \ge 2$, the group $T_{n,r}$ is generated by ϱ and γ_0. The kernel $T^0_{n,r}$ of θ from Lemma 38.17 contains ϱ^d with $d = \gcd(n-1, r)$ and equals the subgroup Γ of $T_{n,r}$ generated by the glides in $\{\gamma_i \mid 0 \le i < r\}$.

Proof. Items (i), (ii), and (iii) are verified by direct checking. Item (iv) can be derived from (iii) using (ii) to first eliminate ϱ_1, and then (i) to move appearances of ϱ to the right and gather all appearances of ϱ together.

From Item (i), we know that $F_{n,r}$ is in $\langle \varrho, \gamma_0 \rangle$. Modulo $F_{n,r}$, every element of $T_{n,r}$ equals some rotation (P, σ, P). And each rotation of the form (P, σ, P) is conjugate by an element of $F_{n,r}$ to a power of some ϱ_k. From Item (ii) we have ϱ_1 in $\langle \varrho, \gamma_0 \rangle$. With all elements of $F_{n,r}$ available, parallels to Item (ii) can be built to inductively get all the ϱ_k in $\langle \varrho, \gamma_0 \rangle$. This shows the first sentence in (II).

From (i), every element of $T_{n,r}$ can be written as a product of glides followed by a power of ϱ. Since we have $\theta(\varrho) = 1$ and θ of any glide is 0, we have that ϱ^d and glides generate $T^0_{n,r}$. In particular $\Gamma \subseteq T^0_{n,r}$.

But ϱ has order r, so $\varrho^r = 1 \in \Gamma$ and (iv) has $\varrho^{n-1} \in \Gamma$, so $\varrho^d \in \Gamma$ and $T^0_{n,r} \subseteq \Gamma$. This proves the second sentence in (II). $\square$

38.7.3. *A second invariant.* We work with an element f of $T_{n,r}$ as given by a forest pair (Φ, σ, Ψ). We also note that f can be given as a product of glides followed by a power of ϱ. Our invariant will live in $\mathbf{Z}^d \rtimes \mathbf{Z}/d\mathbf{Z}$ where $\mathbf{Z}^d$ is regarded as the free abelian group generated by the symbols e_i, $0 \le i < d$, and our preferred generator τ of $\mathbf{Z}/d\mathbf{Z}$ acts on the d-tuple α in $\mathbf{Z}^d$ so that $(\tau(\alpha))_i = \alpha_{i+1}$ with subscripts using arithmetic modulo d. Note that $\tau(e_i) = e_{i-1}$.

Let Φ be an n-ary forest of length r and let v be a vertex of Φ. Let $\phi_n(v) \in \mathbf{Z}/(n-1)\mathbf{Z}$ be as defined in Section 38.7.1. Since $d|(n-1)$, we can let $\rho_d(v) \in \{0, \dots, d-1\}$ be the residue of $\phi_n(v) \pmod d$. Let $I(\Phi)$ be the set of internal vertices of Φ. Define

$$\rho_d(\Phi) = \sum_{v \in I(\Phi)} e_{\rho_d(v)} \in \mathbf{Z}^d. \tag{38.6}$$

Let f in $T_{n,r}$ be given by (Φ, σ, Ψ), and let $\theta(f) \in \mathbf{Z}/d\mathbf{Z}$ be as given in Lemma 38.17. Set

$$\eta(f) = \left(\rho_d(\Phi) - \tau^{\theta(f)}(\rho_d(\Psi)), \tau^{\theta(f)}\right) \in \mathbf{Z}^d \rtimes \mathbf{Z}/d\mathbf{Z}. \tag{38.7}$$

LEMMA 38.19. *For $n \ge 2$, $r \ge 1$ and $d = \gcd(n-1, r)$, the map $\eta : T_{n,r} \to \mathbf{Z} \rtimes \mathbf{Z}/d\mathbf{Z}$ is a well defined homomorphism. We have $\eta(\gamma_{ij}) = (e_{i'} - e_{j'}, 0)$ with $i' \equiv i \pmod d$ and $j' \equiv j \pmod d$. The image of η is all $(\alpha, \tau^k) \in \mathbf{Z}^d \rtimes \mathbf{Z}/d\mathbf{Z}$, where the sum of the entries in α is 0 in $\mathbf{Z}$.*

PROOF. None of the conclusions specifically mention r. As mentioned at the beginning of Section 38.7.2, we can increase r modulo $n-1$ so that Lemma 38.18 applies.

All hinges on the well definedness of η. For well definedness, it suffices to show that for $f = (\Phi, \sigma, \Psi)$, the value of η does not change when a matched n-ary splitting is applied to (Φ, σ, Ψ).

A matched binary splitting of (Φ, σ, Ψ) applies a binary splitting to some a v of Φ and a leaf $v\sigma$ of Ψ, turning v and $v\sigma$ into internal vertices of their respective trees, and making no other changes to the internal vertices. We have $\rho_d(v\sigma) = \rho_d(v) + \theta(f)$, and the modified bijection takes $v0$ to $(v\sigma)0$ where $\rho_d(v0) = \rho_d(v)$ and $\rho_d((v\sigma)0) = \rho_d(v\sigma)$. So we have added $e_{\rho_d(v)}$ to both $\rho_d(\Phi)$ and $\tau^{\theta(f)}\rho_d(\Psi)$, proving well definedness.

Now that η is well defined, it can be noted that it was built to be a homomorphism. The value of $\eta(\gamma_{ij})$ is immediate from the definition. Since ϱ can be represented by a pair of trivial forests, we have $\eta(\varrho) = (0, \tau)$. Since η is a homomorphism and any $f \in T_{n,r}$ is of the form $w\varrho^k$

where w is a product of glides, the claim about the image of η follows from our knowledge of $\eta(\gamma_{ij})$ and $\eta(\varrho)$. $\square$

38.7.4. *Normal subgroups.*

PROPOSITION 38.20. *For $2 \leq n$, $r \geq 1$ and $d = \gcd(n-1, r)$, the subgroup $T^0_{n,r}$ of Lemma 38.18 consists of those $f \in T_{n,r}$ where the second coordinate of $\eta(f)$ is 0. Let $T^c_{n,r}$ be the commutator subgroup of $T_{n,r}$ and let $T^s_{n,r}$ be the kernel of η. Then $T^s_{n,r} \triangleleft T^c_{n,r} \triangleleft T^0_{n,r} \triangleleft T_{n,r}$. The group $T^s_{n,r}$ is the second commutator subgroup of $T_{n,r}$, and is also the commutator subgroup of $T^0_{n,r}$. We have the following quotients.*

1. $T_{n,r}/T^0_{n,r} \simeq \mathbf{Z}/d\mathbf{Z}$.
2. $T_{n,r}/T^c_{n,r} \simeq \mathbf{Z}/d\mathbf{Z} \times \mathbf{Z}/d\mathbf{Z}$.
3. $T^0_{n,r}/T^s_{n,r} \simeq \mathbf{Z}^{d-1}$.

Lastly, $T^s_{n,r}$ is simple and contained in any non-trivial subgroup of $T_{n,r}$ normalized by the elements of $T^s_{n,r}$. Unless $d = 1$, $T^s_{n,r} \neq T^c_{n,r}$.

PROOF. None of the conclusions (e.g.., the right sides of (1)–(3)) specifically mention r, and as in the proof of Lemma 38.19, we can increase r modulo $n-1$ to make it as large as necessary for arguments.

From Lemma 38.17, the group $T^0_{n,r}$ is the kernel of $\theta : T_{n,r} \to \mathbf{Z}/d\mathbf{Z}$. The homomorphism is onto since $\theta(\varrho) = 1$. This verifies (1). Further if π is the projection from $\mathbf{Z}^d \rtimes \mathbf{Z}/d\mathbf{Z}$ to the abelian second factor $\mathbf{Z}/d\mathbf{Z}$, then $T^0_{n,r}$ is the kernel of $\pi\eta : T_{n,r} \to Z/d\mathbf{Z}$, and the image $\eta(T^0_{n,r})$ lies in the abelian $\mathbf{Z}^d \leq \mathbf{Z}^d \rtimes \mathbf{Z}/d\mathbf{Z}$. Denoting by $(T^0_{n,r})'$ the commutator subgroup of $T^0_{n,r}$, we have from all of this that

$$(T^0_{n,r})' \triangleleft T^s_{n,r} \triangleleft T^0_{n,r} \triangleleft T_{n,r}, \text{ and } T^c_{n,r} \triangleleft T^0_{n,r}.$$

This makes $T^0_{n,r}/T^s_{n,r}$ a quotient of $T^0_{n,r}/(T^0_{n,r})'$.

We wish to show $T^s_{n,r} \subseteq (T^0_{n,r})'$. From Lemma 38.18, $T^0_{n,r}$ is generated by the glides in $\{\gamma_i \mid 0 \leq i < r\}$ and so from Lemma 38.19, $\eta(T^0_{n,r})$ is all $(\alpha, 0) \in \mathbf{Z}^d \rtimes \mathbf{Z}/d\mathbf{Z}$ where the sum of the entries in α is 0 in $\mathbf{Z}$. So $\eta(T^0_{n,r}) = T^0_{n,r}/T^s_{n,r}$ is isomorphic to $\mathbf{Z}^{d-1}$. We compare this to $T^0_{n,r}/(T^0_{n,r})'$.

All equalities that follow in this paragraph and the next are modulo $(T^0_{n,r})'$. From Lemma 38.18, $T^0_{n,r}$ is generated by glides and contains ϱ^d, so we have that $\gamma_i = \gamma_j$ if $i \equiv j \pmod{d}$. Thus $T^0_{n,r}/(T^0_{n,r})'$ is no bigger than a quotient of $\mathbf{Z}^d$. Lemma 38.18(iv) says that ϱ^{n-1} equals a product of $r+n-1$ consecutive γ_i, and so we have ϱ^{n-1} equal to the product of d consecutive γ_i raised to the power $(r+n-1)/d$. Letting t be the product of d consecutive γ_i, we have $\varrho^{n-1} = t^{(r+n-1)/d}$.

From the remarks at the end of Section 38.4.1, we have $t^{(n-1)/d} = 1$. This gives $\varrho^{n-1} = t^{r/d}$. We will twice exploit the fact (†) that $(n-1)/d$ and r/d are relatively prime. Now $t^{r/d}$ has order dividing $(n-1)/d$ and $\varrho^r = 1$ gives that ϱ^{n-1} has order dividing r/d. So $t^{r/d} = 1$ using (†). Now $t = 1$, again using (†). So $T^0_{n,r}/(T^0_{n,r})'$ is no bigger than a quotient of $\mathbf{Z}^{d-1}$. So $T^s_{n,r} \subseteq (T^0_{n,r})'$. Also (3) has been verified.

Since $T^s_{n,r} = (T^0_{n,r})' \subseteq T^c_{n,r}$, we have all the containments claimed in the statement and all claims of normalities. We look at (2).

Since $T^s_{n,r} \subseteq T^c_{n,r}$, the quotient $T_{n,r}/T^c_{n,r}$ is isomorphic to the abelianization of $T_{n,r}/T^s_{n,r}$ or the abelianization of the image of η. The image of η is in $\mathbf{Z}^d \rtimes \mathbf{Z}/d\mathbf{Z}$ with Z^d generated by e_i, and the image generated by $\eta(\varrho) = (0,1)$ and the $\eta(\gamma_i) = (e_i - e_{i+1}, 0)$ with all subscripts treated modulo d. The action of ϱ is to conjugate γ_i to γ_{i+1} so after abelianization, all the γ_i become a single element γ. Now γ^d is equivalent to $\gamma_0\gamma_1 \cdots \gamma_{d-1}$ which is in the kernel of η. So $\gamma^d = 1$ in the abelianization. For $1 \le k < d$, the image of γ^k is $e_0^k - e_1^k$ and we have that $T_{n,r}/T^c_{n,r}$ is isomorphic to $\mathbf{Z}/d\mathbf{Z} \times \mathbf{Z}/d\mathbf{Z}$.

We prove the simplicity of $T^s_{n,r} = (T^0_{n,r})'$. We use Higman's arguments from [106] as used in the proof of Proposition 5.12, and prove that $T^0_{n,r}$ is abelian modulo the normalizer in $T^s_{n,r}$ of any non-trivial element in $T^s_{n,r}$. The group $T^0_{n,r}$ is generated by the γ_i, and γ_i commutes with γ_j unless $|i-j| = 1$. But the support of $\Gamma_i = \langle \gamma_i, \gamma_{i+1} \rangle$ is an interval of length 3. By raising r, we can assume we have $r \ge 5$.

Let us temporarily denote by F_t the subgroup of $T_{n,r}$ that fixes $t \in \mathbf{Z}[\frac{1}{n}]$. This is isomorphic to F_n and is a subgroup of $T^0_{n,r}$ since it fixes t. The image of F_t under η is abelian, and so its commutator subgroup F_t' is a subgroup of $T^s_{n,r}$. We choose t not in the closure of the support of Γ_i, and the transitivity of F_t' and Higman's argument from the proof of Proposition 5.12 completes the proof. Details are left to the reader.

Finally, the difference between (2) and (3) shows that $T^s_{n,r} \ne T^c_{n,r}$ when $d \ne 1$. □

COROLLARY 38.20.1. *For $n \ge 2$ and $r \ge 1$, the group $T_{n,r}$ is simple if and only if* $\gcd(n-1, r) = 1$.

38.7.5. *Isomorphisms and mutual embedability.* As of this writing, there is not complete classification of the $T_{n,r}$ up to isomorphism, nor is the embedability relation on the $T_{n,r}$ completely understood. We give some partial information.

The following is immediate from Proposition 38.20 and Lemma 38.2.

Proposition 38.21. *Let* $n \geq 2$, $m \geq 2$, $r \geq 1$ *and* $s \geq 1$ *be integers.*

(I) If $T_{n,r}$ *is isomorphic to* $T_{m,s}$, *then* $gcd(r, n-1) = \gcd(s, m-1)$.

(II) If $r \equiv s \pmod{n-1}$, *then* $T_{n,r}$ *is isomorphic to* $T_{n,s}$.

(III) The groups $T_{n,n-1}$ *and* $T_{m,m-1}$ *are isomorphic if and only if* $n = m$.

We can say something about the embeddings between the $T_{n,r}$ by looking at torsion. It is easy to show that if $f \in T_{n,r}$ has finite order, then f can be represented by some (P, σ, Q) with $P = Q$. So if the order of f is j, then the number of leaves of P must be a multiple of j. The number of leaves of an n-ary forest of length r is $r + k(n-1)$ for some integer k. Thus $\mathrm{Tor}(T_{n,r})$, the set of finite orders of elements of $T_{n,r}$ must be the set of divisors of $\{r + k(n-1) \mid k \in \mathbf{N}\}$. This gives necessary conditions for embedding.

Lemma 38.22. *If* $T_{m,s}$ *embeds homomorphically in* $T_{n,r}$, *then the divisors of* $\{s + j(m-1) \mid j \in \mathbf{N}\}$ *must be a subset of the divisors of* $\{r + k(n-1) \mid k \in \mathbf{N}\}$.

We get sufficient conditions by imitating the proof of part of Lemma 38.15. There, F_n is embedded in F_2 by using $g_i = x_i^{n-1}$ as a generator of F_n. This works, because in the monoid of finitary, binary forests, x_i^{n-1} represents a single non-trivial tree with n leaves and root at position i. Many other words (such as $x_i x_{i+1} x_{i+2} \cdots x_{i+(n-1)}$) could have done as well if used consistently.

In the monoid of finitary n-ary forests, it is possible to build a tree at position i with $1 + k(n-1)$ leaves for any $k \in \mathbf{N}$. For example, $g_{n,i}^k$ will work. Further, the $r + j(n-1)$ leaves for some $j \in \mathbf{N}$ of a finite n-ary forest of length r can be used as attaching vertices for the roots of a finite n-ary forest of length $r + j(n-1)$.

Lemma 38.23. *Let* $n \geq 2$, $m \geq 2$, $r \geq 1$ *and* $s \geq 1$ *be integers with* $n \leq m$ *and* $r \leq s$. *If* $(n-1)|(s-r)$ *and* $(n-1)|(m-1)$, *then* $T_{m,s}$ *embeds homomorphically into* $T_{n,r}$.

Proof. This is mostly left to the reader. Given (P, σ, Q) representing an element in $T_{m,s}$, finding (P', σ', Q') that represents the corresponding element in $T_{n,r}$ is straightforward. To show that this is a homomorphism, one notes that multiplication in $T_{m,s}$ hinges on sequences of matched m-ary splittings and their effect on the rotations. If the shape of the n-ary tree having m leaves is used consistently throughout, then a single matched m-ary splitting in $T_{m,s}$ corresponds to a fixed sequence of matched n-ary splittings in $T_{n,r}$. It then follows that the assignment is a homomorphism. $\square$

38.8. On the $V_{n,r}$. Much can be adapted from the discussion for V in Section 13. We will work with the view of V as a group of actions on the Cantor set $\mathfrak{C}$. To move to $V_{n,r}$, we will start with the n-ary Cantor set $\mathfrak{C}_n$ regarded as elements of X_y^ω, and then look at $\tilde{r} \times \mathfrak{C}_n$ which we denote by $r\mathfrak{C}_n$ where $\tilde{r} = \{1, \ldots, r\}$. Now elements of $f\mathfrak{C}_n$ are elements of $\tilde{r} \times X_y^\omega$ consisting of words w that start with $w_0 \in \tilde{r}$ and that are followed by elements of X_y^ω.

A prefix set U for $r\mathfrak{C}_n$ has the predictable definition and prefix sets are naturally in one-to-one correspondence with leaf sets of n-ary forests of length r. We have the following parallel to Lemma 13.12 and its corollary.

LEMMA 38.24. *For integers $n \geq 2$ and $r \geq 1$, the following is true. Let $A = \{u_1\mathfrak{C}_n, \ldots, u_k\mathfrak{C}_n\}$ and $B = \{v_1\mathfrak{C}_n, \ldots, v_k\mathfrak{C}_n\}$ each be a set of pairwise disjoint cones that do not cover $r\mathfrak{C}_n$. Then there is an element $f \in V_{n,r}$ so that for each $i \in \{1, \ldots, k\}$, the restriction of f to $u_i\mathfrak{C}_n$ is the rigid cone map to $v_i\mathfrak{C}_n$.*

COROLLARY 38.24.1. *The group $V_{n,r}$ is closed under deferment.*

We did not prove a parallel to the next Lemma in Section 13.

LEMMA 38.25. *If H is a finite subgroup of $V_{n,r}$, then there is an n-ary forest Θ of length r so that every $h \in H$ can be represented as (Θ, σ, Θ) for some bijection σ on the leaf set $\Lambda(\Theta)$ of Φ.*

PROOF. Each $h \in H$ is represented by a forest pair whose domain forest is some Φ_h. Let Ψ be the union of all the Φ_h. By this we mean that the i-th tree of Ψ is the union of all the i-th trees of the Φ_h.

Let Θ be the union of all the forests in $\{(\Psi)h \mid h \in H\}$. By $(\Psi)h$ we mean the forest whose leaf set is $(\Lambda(\Psi))h$. See Lemma 8.8. Every leaf of Θ is a leaf of $(\Psi)h$ for some h. We claim that Θ has the properties that we seek.

Let u be in $\Lambda(\Theta)$ and g in H. If ug is not a leaf of Θ, then either (i) a proper descendant v of ug is a leaf of Θ or (ii) ug is a proper descendant of a leaf v of Θ. In case (i), $v = wh$ for some leaf w of Ψ and some $h \in H$, and $z = vg^{-1} = w(hg^{-1})$ is a proper descendant of u. But our construction of Θ guarantees that no leaf of any $(\Psi)h$ is a proper descendant of a leaf of Θ. In the case (ii), $u = w'h'$ for some $w' \in \Lambda(\Psi)$ and $h' \in H$, making $z' = ug = w'(h'g)$ a proper descendant of $v \in \Lambda(\Theta)$ which is again impossible. ☐

The discussion of normal subgroups is less complex than for the $T_{n,r}$, but not quite as trivial as for V. We start with the following where the definitions are in the same spirit as Section 13.2.

PROPOSITION 38.26. *Fix integers $n \geq 2$ and $r \geq 1$. Any two proper transpositions are conjugate in $V_{n,r}$, and any two proper 3-cycles are conjugate in $V_{n,r}$. The proper transpositions generate $V_{n,r}$ and the normal closure of any proper transposition is all of V. If n is even, the normal closure of any non-trivial element contains a proper transposition, and $V_{n,r}$ is simple. If n is odd, the normal closure of any non-trivial element contains a proper 3-cycle and $V_{n,r}$ has a simple subgroup of index 2.*

PROOF. We will give the details for the last sentence that covers the case when n is odd. The rest is left to the reader guided by the material in Section 13.2. Only minor modifications will be needed, and Lemma 38.25 also helps.

We write permutations as products of cycles such as $(a\ b\ c)(d\ e)$ since more than transpositions will be involved.

As in the proof of Theorem 13.17, a non-trivial f moves some cone $u\mathfrak{C}_n$ rigidly off itself. Let $v\mathfrak{C}_n = (u\mathfrak{C}_n)f$. There are at least $u0$, $u1$ and $u2$ as children of u, and $v0$, $v1$ and $v2$ as children of v. The commutator of f and $(u0\ u1)$ is $g = (u0\ u1)(v0\ v1)$. The commutator of g with $(u0\ v0)$ is $h = (u0\ v0)(u1\ v1)$. If $j = (u0\ u1)(v0\ v1\ v2)$, then $hjh^{-1}j^{-1}$ is $(u1\ v1\ v2)$ and is in the normal closure of f.

We turn our attention to the simple subgroup of index 2. Given a permutation on a finite totally ordered set J, we define

$$\chi(\sigma) = \#\{(i,j) \in J \times J \mid i < j,\ i\sigma > j\sigma\} \pmod 2$$

where $\#S$ is the number of elements of a finite set S. The value $\chi(\sigma)$ counts modulo 2 the number of pairs that σ puts out of order, and determines whether σ is odd or even.

Given $f = (P, \sigma, Q)$ representing an element of $V_{n,r}$, we declare $\chi(f) = \chi(\sigma)$ and need to argue well definedness. If a matched n-ary splitting is applied to the leaf v of P and the leaf $v\sigma$ of Q giving (P', σ', Q'), then the set S' of disturbed orders for $\chi(\sigma')$ differs from the set S for $\chi(\sigma)$ in that every time v appears in a pair in S, that pair is replaced by n pairs in S'. No other changes take place. If n is odd, the replacement increases the count by $n-1$ for each appearance of v in a pair in S and the parity does not change. That χ is a surjective homomorphism to $\mathbf{Z}/2\mathbf{Z}$ now follows from the details of opportunistic multiplication.

The kernel of χ are those (P, σ, Q) where σ is an even permutation. Every even permutation is a product of 3-cycles, and every non-trivial normal subgroup N of $V_{n,r}$ has been shown to contain all proper 3-cycles. It is straightfoward to argue that all f in the kernel of χ will be in N. □

We will use $V_{n,r}^+$ to denote the simple index two subgroup of $V_{n,r}$ from Proposition 38.26 when n is odd, and $V_{n,r}$ when n is even. From the results in the next section, we have that the $V_{n,r}^+$ form an infinite family of countable, simple groups.

38.8.1. *Isomorphisms and mutual embeddability.* The groups $V_{n,r}$ and $V_{n,r}^+$ are classified up to isomorphism. That is, there are conditions on the pairs (n, r) that are both necessary and sufficient for the corresponding groups to be isomorphic. The conditions are that the groups given by (n, r) and (m, s) are isomorphic if and only if $n = m$ and $\gcd(r, n-1) = \gcd(s, m-1)$. That the conditions are necessary is shown in Higman 1974 [107]. That the conditions are sufficient is shown in Pardo 2011 [168]. The almost four decade gap between the two is significant. Below we will give Higman's proof that the conditions are necessary.

That the conditions are sufficient needs more machinery and will be shown in Chapter 7. There it is shown that the $V_{n,r}$ embed as characteristic subgroups of endomorphism rings of certain algebras. The extra structure carried by the algebras gives tools sufficient to build isomorphisms between the rings, which then induce isomorphisms between the $V_{n,r}$. As of yet, no one has built the corresponding isomorphisms between the $V_{n,r}$ directly.

The necessity of the conditions for isomorphism is shown by counting conjugacy classes of torsion elements in the $V_{n,r}$. We can start with a more general discussion of conjugacy classes of finite subgroups of the $V_{n,r}$.

Let H be a finite subgroup of some $V_{n,r}$ regarded as the image of some embedding $\alpha : X \to V_{n,r}$. Let Θ be a forest guaranteed by Lemma 38.25 for which every $h \in H$ can be represented as (Θ, σ, Θ) for some permutation σ on $\Lambda(\Theta)$. The orbits of σ are representations of H as transitive permutation groups. There are only finitely many possible isomorphism classes of such transitive representations (including the trivial permutation on a set of size 1). For the isomorphism class of the group H, fix a k-tuple $\mathbf{X} = (X_1, \ldots, X_k)$ of the full list of the isomorphism classes of transitive permutation representations of H. We will use $|X_i|$ to denote the size of the underlying set of the representation X_i. We will insist that $|X_1| = 1$ and $|X_k| = |H|$. Then there is a k-tuple of non-negative integers $\mathbf{c} = (c_1, \ldots, c_k)$ satisfying $|\Lambda(\Theta)| = \sum_1^k c_i |X_i|$. Since Θ is an n-ary forest with r trees, we must have

$$\sum_1^k c_i |X_i| \equiv r \pmod{n-1}. \tag{38.8}$$

The tuple $\mathbf{c}$ is a function of α, Θ and $\mathbf{X}$ which we can denote $\mathbf{c}(\alpha,\Theta,\mathbf{X})$.

If β is another inclusion of H into $V_{n,r}$ keeping a forest Ψ invariant, with $\mathbf{c}(\beta,\Psi,\mathbf{X}) = \mathbf{c}(\alpha,\Theta,\mathbf{X})$, then the transitivity properties of $V_{n,r}$ guarantee that β and α are conjugate. As a partial converse, if $|\Psi| = |\Theta|$ and an element represented by (Θ,σ,Ψ) conjugates α to β, then $\mathbf{c}(\alpha,\Theta,\mathbf{X}) = \mathbf{c}(\beta,\Psi,\mathbf{X})$.

If $\mathbf{c}$ is any k-tuple of non-negative integers satisfying the equivalence in (38.8), then there is an inclusion γ of H into $V_{n,r}$ keeping invariant a forest Θ with $\mathbf{c}(\gamma,\Theta,\mathbf{X}) = \mathbf{c}$. We discuss when different k-tuples of non-negative integers give conjugate inclusions.

Let α and β be two inclusions of H into $V_{n,r}$, and let Θ_α and Θ_β, respectively, be the forests kept invariant by the two inclusions. We can symbolize the situation by $(\Theta_\alpha,\alpha,\Theta_\alpha)$ and $(\Theta_\beta,\beta,\Theta_\beta)$. If (Ψ,σ,Ψ) conjugates α to β, then carrying out the computaton of

$$(\Psi,\sigma^{-1},\Psi)(\Theta_\alpha,\alpha,\Theta_\alpha)(\Psi,\sigma,\Psi) = (\Theta_\beta,\beta,\Theta_\beta)$$

goes through relevent matched n-ary refinements to arrive at

$$(\Psi',(\sigma^{-1})',\Psi')(\Theta'_\alpha,\alpha',\Theta'_\alpha)(\Psi',\sigma',\Psi') = (\Theta'_\beta,\beta',\Theta'_\beta)$$

with appropriate equalities of certain of the forests. Most important will be that $\Theta'_\alpha = \Theta'_\beta$ which we can denote Φ, and that α' and β' are conjugate by σ'. From this we get $\mathbf{c}(\alpha',\Phi,\mathbf{X}) = \mathbf{c}(\beta',\Phi,\mathbf{X})$. We look at the relation between $\mathbf{c}(\alpha,\Theta_\alpha,\mathbf{X})$ and $\mathbf{c}(\alpha',\Phi,\mathbf{X})$.

If a single matched n-ary splitting is done to $(\Theta_\alpha,\alpha,\Theta_\alpha)$ at a leaf v, and v is fixed by α, then the domain and range forests will remain equal. If v not fixed by α, then the domain and range forests will become unequal and can only be made equal by doing matched n-ary splittings at every leaf in the orbit of v under αH. In both cases, one orbit becomes n orbits, each isomorphic to the orbit of v as permutation representations of H. If X_i is the permutation type of the action of H on the orbit of v and c_i is the i-th entry in $\mathbf{c}(\alpha,\Theta_\alpha,\mathbf{X})$, then c_i is replaced by $c_i + (n-1)$ after the splittings are done. Note that c_i cannot be 0 in this situation, both before and after the splitting. We define an equivalence to mirror this.

We write $c \overset{\circ}{\equiv} d \pmod{n-1}$ to mean that either $c = d = 0$ or that both c and d are not 0 and $c \equiv d \pmod{n-1}$. For k-tuples $\mathbf{c}$ and $\mathbf{d}$, we write $\mathbf{c} \overset{\circ}{\equiv} \mathbf{d} \pmod{n-1}$ if $\overset{\circ}{\equiv} \pmod{n-1}$ holds coordinate by coordinate. In the situation being discussed, we get $\mathbf{c}(\alpha,\Theta_\alpha,\mathbf{X}) \overset{\circ}{\equiv} \mathbf{c}(\alpha',\Phi,\mathbf{X}) \pmod{n-1}$, and it then follows that $\mathbf{c}(\alpha,\Theta_\alpha,\mathbf{X}) \overset{\circ}{\equiv} \mathbf{c}(\beta,\Theta_\beta,\mathbf{X}) \pmod{n-1}$. The converse is clear and we have the following.

LEMMA 38.27. *Two k-tuples $\mathbf{c}$ and $\mathbf{d}$ satisfying the equivalence in (38.8) give rise to conjugate inclusions of H into $V_{n,r}$ if and only if* $\mathbf{c} \overset{\circ}{\equiv} \mathbf{d} \pmod{n-1}$.

We will use Lemma 38.27 together with (38.8) to count conjugacy classes of cyclic subgroups of prime power order. If H is cyclic of prime power order p^a, then the full list of the X_i for H is $(X_1, \dots X_{a+1})$ with $|X_i| = p^{i-1}$. In particular, if H is cyclic of prime order p, then the full list of the X_i for H is (X_1, X_2) with $|X_1| = 1$ and $|X_2| = p$.

We will need the Chinese remainder theorem (Proposition 50.2 in the appendix), and the following elementary facts about solutions to equivalences.

LEMMA 38.28. *(I) If $a|m$ and for some x, we have $ax \equiv r \pmod{m}$, then $a|r$.*

(II) The equation $kx \equiv r \pmod{m}$ has exactly $g = \gcd(k, m)$ solutions modulo m if $g|r$.

The lemma is stated and proved as Lemma 50.3 in the Appendix.

THEOREM 38.29. *(I) If the prime p does not divide $n-1$, then there are n conjugacy classes of elements of order n in $V_{n,r}$ and in $V^+_{n,r}$.*

(II) If $V_{n,r}$ and $V_{n',r'}$ are isomorphic or if $V^+_{n,r}$ and $V^+_{n',r'}$ are isomorphic, then $n = n'$.

PROOF. (I) The number of conjugacy classes of inclusions of $\mathbf{Z}/p\mathbf{Z}$ in $V_{n,r}$ is the number of solutions of (†) $c_1 + c_2 p \overset{\circ}{\equiv} r \pmod{n-1}$ with $c_2 \neq 0$. Note that $c_1 = 0$ and $c_1 = n-1$ do not give conjugate inclusions. With p prime to $n-1$, there is a unique solution with $1 \leq c_2 \leq n-1$ to (†) for each value of c_1 with $0 \leq c_1 \leq n-1$, and there are n conjugacy classes in $V_{n,r}$. If n is even, then $V^+_{n,r} = V_{n,r}$, and if n is odd, then $n-1$ is even and p is odd. An element of order p is an even permutation and lives in $V^+_{n,r}$. Since p is prime to $n-1$, from the Chinese remainder theorem there are infinitely many m satisfying both $m \equiv r \pmod{n-1}$ and $m \equiv 2 \pmod{p}$. An element h of order p acts faithfully on the leaves of some tree whose number of leaves is such an m. The action must fix at least two leaves u and v. The element transposing u and v is in $V_{n,r} \setminus V^+_{n,r}$ and commutes with h. Thus a conjugation of h can be accomplished by an element of $V^+_{n,r}$.

(II) If $n \neq n'$, then there is some prime p that is prime to both $n-1$ and $n'-1$ and the relevant groups have different numbers of conjugacy classes of elements of order p. □

THEOREM 38.30. *(I) If p is prime and $a > 0$ is the largest integer for which $p^a|(n-1)$, then there is an $(a+1)$-tuple $\mathbf{d} = (d_0, \dots, d_a)$ so* *that for each r modulo $(n-1)$, if $b \geq 0$ is the largest integer for which $p^b | \gcd(r, n-1)$, then the number of conjugacy classes of elements of order p^a in $V_{n,r}$ or $V^+_{n,r}$ is $\sum_0^b d_i$.*

(II) If $V_{n,r}$ and $V_{n,r'}$ are isomorphic or if $V^+_{n,r}$ and $V^+_{n,r'}$ are isomorphic, then $\gcd(r, n-1) = \gcd(r', n-1)$.

Proof. (I) The transitive actions of $\mathbf{Z}/p^a\mathbf{Z}$ on finite sets are p^k-cycles for $k \leq a$. So the conjugacy classes of embeddings of $\mathbf{Z}/p^a\mathbf{Z}$ are in one-to-one correspondence with solutions to

$$\sum_{i=0}^{a} c_i p^i \overset{\circ}{\equiv} r \pmod{n-1} \quad \text{subject to } c_a \neq 0. \tag{38.9}$$

The union of the classes under $\overset{\circ}{\equiv}$ (mod $n-1$) of solutions to (38.9) over all values of r is simply the set of classes of $(a+1)$-tuples $\mathbf{c}$ under $\overset{\circ}{\equiv}$ (mod $n-1$) with $c_a \neq 0$. Let $\mathcal{C}$ denote this union. For $0 \leq k \leq a$, let $\mathcal{C}_k$ be those elements of $\mathcal{C}$ where $c_0 = c_1 = \cdots = c_{k-1} = 0$ and $c_k \neq 0$.

Fix r modulo $(n-1)$. We will show that the number of solutions to (38.9) depends only on b for that r.

If $k > b$, then p^k does not divide $\gcd(r, n-1)$ and thus does not divide r. By (I) of Lemma 38.28, no element of $\mathcal{C}_k$ is a solution to (38.9) for r.

If $k \leq b$, then a solution to (38.9) for r in $\mathcal{C}_k$ is a solution to

$$c_k p^k \overset{\circ}{\equiv} r - \left(\sum_{i=k+1}^{a} c_i p^i \right) \pmod{n-1}. \tag{38.10}$$

Note that the expressionon the right is simply r if $k = b = a$.

With $p^k | (n-1)$, we have $\gcd(p^k, n-1) = p^k$ which divides p^b which in turn divides r and thus also divides the right hand side of (38.10). So from (II) of Lemma 38.28, we know that there are exactly p^k solutions to (38.10) for c_k modulo $(n-1)$ (which by definition of $\mathcal{C}_k$ excludes $c_k = 0$) for every combination of values of c_i, $k < i \leq a$ and up to $\overset{\circ}{\equiv}$ (mod $n-1$), with $c_a \neq 0$. The number of different expressions on the right is $|\mathcal{C}_k|/(n-1)$ and the total number of solutions to (38.10) is $p^k |\mathcal{C}_k|/(n-1)$ which we will denote by d_k. This gives that the number of solutions to (38.9) is $\sum_{k=0}^{b} d_k$. This completes the claim in (I) for $V_{n,r}$.

Some of the argument to transfer the result for (I) to the $V^+_{n,r}$ can be done as in Theorem 38.29. In all cases, we will show that elements of order p^a in $V^+_{n,r}$ that are conjugate in $V_{n,r}$ are also conjugate in

$V^+_{n,r}$. What will remain to discuss is the relation between the conjugacy classes in the two groups when n is odd and $V^+_{n,r} \neq V_{n,r}$.

As in Theorem 38.29, we will show that each element of order p^a in $V^+_{n,r}$ commutes with an element in $V_{n,r} \setminus V^+_{n,r}$. If p is odd, the argument can be copied from Theorem 38.29. If $p = 2$, then the orbits of an element h of order p^a, $a \neq 0$, all have sizes an integral power of 2, with at least one orbit having more than one element. The action on this orbit is an odd cycle, and there is an element of $V_{n,r} \setminus V^+_{n,r}$ whose action is this cycle and the identity everywhere else. This element commutes with h.

The discussion of conjugacy classes breaks into cases. We assume n is odd. If p is odd, then all the conjugacy classes of elements of order p^a lie in $V^+_{n,r}$.

If $p = 2$, the classes will be different, but in a predictable way. Conjugacy classes in $V_{n,r}$ are in one-to-one correspondence with solutions to (38.9) up to $\overset{\circ}{\equiv} \pmod{n-1}$. With $p = 2$ each non-trivial orbit of an element in such a class is an odd permutation and the element will be an even permutation if and only if $\sum_{i=1}^{a} c_i$ is even. Using the fact that $(n-1)/2^a$ is odd and not equal to $(n-1)$, we see that mapping of $(a+1)$-tuples given by

$$(c_0, c_1, \ldots, c_{a-1}, c_a) \mapsto (c_0, c_1, \ldots, c_{a-1}, c_a + (n-1)/2^a)$$

up to $\overset{\circ}{\equiv} \pmod{n-1}$ is a fixed point free permutation of the solutions of (38.9) taking even permutations to odd permutations and vice versa. Hence each conjugacy class in $V_{n,r}$ of an element of order p^a has exactly half its elements in $V^+_{n,r}$. Thus dividing each d_k by two gives a tuple that works as claimed for $V^+_{n,r}$.

(II) If $\gcd(r, n-1) \neq \gcd(r', n-1)$, then for some prime p, different powers of p divide $\gcd(r, n-1)$ and $\gcd(r', n-1)$. From (I), the relevant groups will have different numbers of conjugacy classes of elements of order p^a with a as given in (I). □

Lemma 6.5 of [107] shows that no further necessary conditions for isomorphisms can be found by counting conjugacy classes of finite subgroups. But in fact the conditions in Theorems 38.29 and 38.30 are sufficient. Showing this uses representations of the $V_{n,r}$ into other structures, and will be covered in Section 40.

Theorem 7.5 of [107] gives a partial converse to Theorems 38.29 and 38.30. We refer the reader to [107] for the proof. The statement of the result is that if c is a divisor of n, then $V_{n,r}$ and $V_{n,cr}$ are isomorphic. In Section 41, this is compared to the full converse that is proven there.

39. End notes

The material on the F_n is from Brown 1987 [34]. Material on the $T_{n,r}$ is from [34] and Liousse-Hmili Ben Ammar 2020 [111]. Material on the $V_{n,r}$ is from Higman 1974 [107].

This chapter is far from complete. There are more variations of the definition to cover, and we have not even touched on marriages. The history of both goes back to the beginning of the subject with the first variations appearing in Higman 1974 [107], and Brown 1987 [34], and the first marriages in Thompson 1980 [189], and Scott 1982 [174, 176, 175]. Variations and marriages are a big part of the Thompson culture and more will be covered in later editions.

CHAPTER 7

Representations

Chapter contents

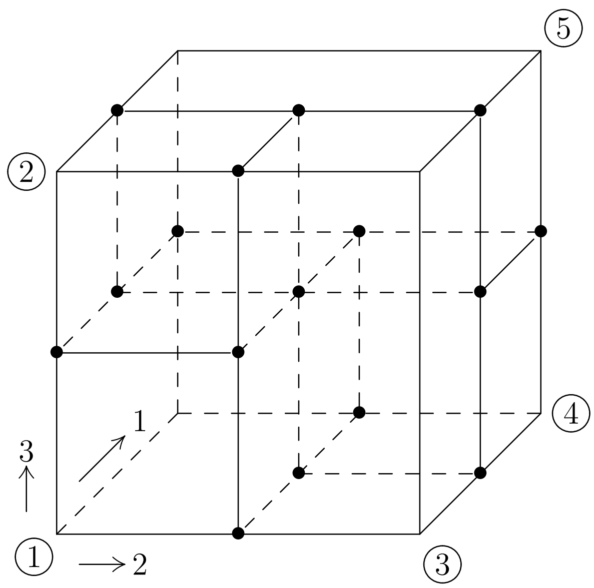

① $:= y_{2,1}y_{3,1}$ ② $:= y_{1,1}y_{2,1}y_{3,2}$ ③ $:= y_{1,1}y_{2,2}$
④ $:= y_{1,2}y_{2,2}y_{3,1}$ ⑤ $:= y_{1,2}y_{3,2}$

[1] Sections 40 and 41 of this chapter deliver material promised in the end notes (Section 21) of Chapter 3. The goal of these sections is to introduce the Leavitt rings and algebras and then use them to classify the groups $V_{n,r}$ up to isomorphism. We start with the Leavit algebras and how they are built, but as mentioned in Section 21 one could start from simpler principles.

40. Leavitt algebras

The Thompson groups $V_{n,r}$ of Section 38.8 have faithful representations of as characteristic subgroups of certain rings with extra structure. By this we mean that the images of the representations are invariant under automorphisms of the rings that preserve the extra structure. In particular, structure preserving isomorphisms between the rings induce isomorphisms of the $V_{n,r}$. The representations prove to be useful since they ultimately allow us to classify the $V_{n,r}$ up to isomorphism. Universalities arise that are parallel to those brought up in Section 18. We give a bit more detail and background.

In Section 18, the algebra A of Jónsson-Tarski 1961 [119] is developed together with its connection to Thompson's group V. The algebra A is the free algebra on one variable in a variety $\mathcal{V}$, where $\mathcal{V}$ is universal for the property that the free algebra in $\mathcal{V}$ on one variable is also free on two variables. The group V is isomorphic to $\mathrm{Aut}(A)$.

In a series of papers 1956–62 [136, 137, 138], Leavitt develops rings with analogous properties to those of the Jónsson-Tarski algebras. For an integer $n > 1$, the Leavitt ring L_n is universal for the property that the free right module over L_n of rank 1 is also free of rank n. That is, L_n^1 and L_n^n are isomorphic as L_n-modules.

Pardo 2011 [168] shows that if the endomorphism ring $M_r(L_n)$ of the free right module over L_n of rank r is endowed with a certain involution and partial order, then the Thompson group $V_{n,r}$ is isomorphic to a characteristic subgroup of $M_r(L_n)$. The notation $M_r(L_n)$ reflects that it is the ring of $r \times r$ matrices over L_n. It follows that an order and involution preserving isomorphism from $M_r(L_n)$ to $M_{r'}(L_{n'})$ induces an isomorphism from $V_{n,r}$ to $V_{n',r'}$.

Abrams-Ánh-Pardo 2008 [1] finds sufficient conditions for the existence of isomorphisms between some pairs of the $M_r(L_n)$, and in [168], this result is enhanced to give sufficient conditions for more pairs of the $M_r(L_n)$. When applied to the characteristic subgroups $V_{n,r}$, this shows that the criteria from Higman 1974 [107] that are necessary for isomorphisms between the $V_{n,r}$ to exist are also sufficient. This gives a

[1]This chapter is not complete. More material will be added in the future.

complete classifications of the isomorphisms types of the $M_r(L_n)$ and of the $V_{n,r}$.

Dicks-Martínez-Pérez 2014 [58] reproves and extends the results of [1] and [168]. The extensions cover groups in the Thompson family that we have not yet discussed in Chapter 6. We will not cover the extensions here, but will use techniques from [58] to keep track of certain details.

In this section, we describe the rings L_n and their generalizations to algebras $L_k(1,n)$ over an arbitrary field k. The ring L_n can be thought of as $L_{\mathbf{Z}}(1,n)$. In Section 40.4 we derive a characterization of the subgroup of $M_r(L_n)$ that is isomorphic to $V_{n,r}$. What we show is is minutely more general than what is shown in [168] and [58] in that we only need to assume the preservation of the partial order on $M_r(L_n)$ and not the involution. This change is mostly cosmetic since we continue to work with the involution, and the amount of checking that is eliminated is negligible.

In Section 41, we derive the results of [1] and [168] following [58] and refer the reader to [58] for the generalization to other Thompson groups. Unfortunately, the letters we choose for our subscripts agree with the notation of Section 38.8 and other papers, and not with the notation in [58].

40.1. The definition. Here we build the rings $L_{\mathbf{Z}}(1,n)$ and the corresponding algebras $L_k(1,n)$ for a field k. We will use L_n as a shorthand for $L_{\mathbf{Z}}(1,n)$. Starting with Section 40.4, we will focus entirely on the rings $L_n = L_{\mathbf{Z}}(1,n)$.

40.1.1. *Conventions.* All rings and algebras will have 1. For a ring or algebra L, the elements of the free L-module L^n over L will be row vectors with entries from L, and transformations from L^n to L^m will be $n \times m$ matrices acting on the right. The endomorphism ring of L^n will be $M_n(L)$, the $n \times n$ matrices over L. We use $\mathbf{I}_n$ for the $n \times n$ identity matrix. For matrices A and B, we write $A \oplus B$ for the block matrix $\begin{pmatrix} A & 0 \\ 0 & B \end{pmatrix}$.

40.1.2. *The algebras.* Fix an integer $n > 1$. We will derive what must be true for an algebra L to have L^1 as a free right L-module of rank 1 isomorphic to L^n as a free right L-module of rank n. There must be a transformation A thought of as a row matrix $(y_1, \dots, y_n)$ from L^1 to L^n, and a transformation B thought of as a column matrix $(x_1, \dots, x_n)^{tr}$ (with ^{tr} indicating transpose) from L^n to L^1 so that $AB = \mathbf{I}_1$ and

$BA = \mathbf{I}_n$. Expanding, we get the following n^2+1 relations.

$$\sum_1^n y_i x_i = 1, \tag{40.1}$$

$$x_i y_j = \begin{cases} 1, & i = j, \\ 0, & i \neq j. \end{cases} \tag{40.2}$$

Given a field k, the Leavitt algebra $L_k(1,n)$ is the k-algebra generated by $\mathbf{X} = \{x_1, \ldots, x_n, y_1, \ldots, y_n\}$ subject to the relations in (40.1) and (40.2). One obtains the Leavitt ring $L_n = L_{\mathbf{Z}}(1,n)$ similarly.

We will have need to refer often to the set $\mathcal{M}_n$ of monomials in $L_k(1,n)$. Endowed with the multiplication of $L_k(1,n)$, the set $\mathcal{M}_n$ becomes a monoid. From Corollary 40.1.1 below, $\mathcal{M}_n$ is the set of all finite (possibly empty) words in the alphabet $\mathbf{X}$ modulo the relations in (40.2). Elements of $L_k(1,n)$ are equivalence classes of k-linear combinations of elements of $\mathcal{M}_n$.

Note that if we define $y_i^* = x_i$ and $x_i^* = y_i$, and if we extend * to all monomials by $(uv)^* = v^*u^*$, then we have $(y_ix_i)^* = y_ix_i$ and $(x_iy_j)^* = x_jy_i$, and the operation * preserves (40.1) and (40.2). So extending * to all of of $L_k(1,n)$ linearly gives an algebra involution on $L_k(1,n)$. As an algebra with involution *, the sets $X_y = \{y_1, \ldots, y_n\}$ and $X_x = \{x_1, \ldots, x_n\}$ each generate $L_k(1,n)$.

If M is a matrix over a ring L, then M^* will denote the $*$-transpose of M. That is, $(M^*)_{i,j} = (M_{i,j})^*$.

We say that an element ν of $L_k(1,n)$ is *unitary* if $\nu\nu^* = 1 = \nu^*\nu$. As a preview to the important Lemma 40.6 below, the reader can take as an exercise to show that if σ is a permutation on $\{1, \ldots, n\}$, then $\sum_1^n y_i x_{i\sigma} = \sum_1^n y_i y_{i\sigma}^*$ is unitary. Unitaries are units and the set of unitaries forms a group under the ring multiplication. But over $\mathbf{R}$ with $n = 2$ and $0 \neq a \in \mathbf{R}$, if we set $\nu_i = a - y_ix_i$, $i \in \{1,2\}$, then $\nu_i^* = \nu_i$, and we get $\nu_i\nu_i^* = a^2 + (1-2a)y_ix_i$, and $\nu_1\nu_2 = \nu_2\nu_1 = a^2 - a$. So ν_i is a unit when $a^2 - a = 1$. The normal forms from Section 40.2 below will show that this gives examples of units that are not unitary.

40.2. Normal forms and positive elements. A normal form for elements of a Leavitt algebra was given in Lemma 1 of 1956 [136]. We will introduce two normal forms, one general and one restricted, and give some consequences. We will use Newman's lemma from Appendix A (Section 47) to establish their uniqueness. We introduce terminology that will apply to both forms.

Every element of $L_k(1,n)$, with $n > 1$ an integer and k a field or the integers, is a finite k-linear sum of monomials from $\mathcal{M}_n$. We will say that a monomial uv^* in $\mathcal{M}_n$ is *standard* (or is a standard monomial) to mean that u and v are words (possibly empty) over $X_y = \{y_1, \ldots, y_n\}$. If uv^* is standard, then $(uv^*)^* = vu^*$ is standard as well.

Every monomial in $L_k(1, n)$ can be represented as a standard monomial since by (40.2) we can remove any appearance of $x_i y_j$ as a subword from a monomial without changing the element represented. If $x_i y_j$ with $i \neq j$ appears, then that summand can be replaced by 0. If $x_i y_i$ appears, then that appearance of $x_i y_i$ can be replaced by 1. The reader can use Newman's lemma (Lemma 47.2) to verify that every element of $\mathcal{M}_n$ is represented by a unique standard monomial.

The distributive law lets us combine summands having identical monomials. If no such combinations are possible, then the sum will be referred to as *linearly reduced*.

40.2.1. *The general form.* This form applies to $L_k(1, n)$ for k a field or the integers and will let us distinguish between elements. It is essentially the form of [136]. Note that we can rewrite the relation (40.1) as

$$y_n x_n = 1 - \sum_{1}^{n-1} y_i x_i. \tag{40.3}$$

This makes it reasonable to ask that there be no subword of the form $y_n x_n$ in any summand of our normal form. We can now claim the following.

LEMMA 40.1. *For k a field or the integers, each element of $L_k(1, n)$ is represented by a unique linearly reduced sum of standard monomials in which there is no appearance of a subword of the form $y_n x_n$.*

PROOF. We will introduce two types of reductions and argue that the set of reductions is terminating and locally confluent. We work with k-linear sums of monomials.

A reduction of Type I is an elimination, using (40.2), of an appearance of $x_i y_j$ in a summand followed by a complete linear reduction (combining of like terms). That this is well defined in spite of different ways to linearly reduce follows from the associative law of addition.

A reduction of Type II is an elimination of an appearance of $y_n x_n$ using (40.3) in a summand followed by a complete linear reduction. Again this is well defined.

To argue that random applications of the two types of reduction must terminate in an irreducible, we use a complexity. Given a k-linear sum of monomials, let S be the set of monomials (without coefficients) used in the sum. That is, S eliminates duplication. We let c_1 be the number of appearances of all symbols in S. We let c_2 be the number of appearances of the symbol y_n in S. For example if $S = \{y_2y_2, y_1y_2x_1\}$ and $n = 2$, then $c_1 = 5$ and $c_2 = 3$.

Since the starting sum might not be linearly reduced, we discuss what changes occur to the complexity under reductions other than the first reduction. A reduction of Type I lowers c_1 and either leaves c_2 the same or reduces it. A reduction of Type II might raise c_1 since new monomials are introduced, but it will lower c_2. Thus pairs (c_1, c_2) ordered lexicographically with c_2 the most significant form a complexity that is always lowered by either type of reduction. This shows that reductions are terminating.

If two reductions of Type I are available to a sum, then they must be at disjoint sites. Checking that the two reductions can be combined is complicated by the fact that each must end with a linear reduction. A number of cases must be checked (including cases where the two sites are in either the same monomial or in different monomials) and are left to the reader. The same comment applies to two reductions of Type II and reductions of different types at disjoint sites. We must consider a sum in which a reduction of the two different types is available and they affect overlapping sets of letters. This happens when either $y_nx_ny_n$, $x_ny_nx_n$, $y_nx_ny_i$, or $x_iy_nx_n$, $i \neq n$, appears in a summand.

If $y_nx_ny_n$ appears in a monomial, then a reduction of Type I replaces the triple with y_n. A reduction of Type II creates n monomials in which the triple is replaced by $1y_n = y_n$ in one monomial and replaced by $y_ix_iy_n = 0$, $i \in \{1, \ldots, n-1\}$, in the remaining $n-1$ monomials, and the results are seen to be the same. A similar argument covers the appearance of $x_ny_nx_n$ in a monomial.

If $y_nx_ny_i$, $i \neq n$, appears in a monomial, then a reduction of Type I gives 0. A reduction of Type II produces n monomials with the triple replaced by $1y_i = y_i$ in one monomial, replaced by $y_ix_iy_i$ in one monomial with its coefficient negated, and replaced by $y_jx_jy_i = 0$ in the remaining $n-2$ monomials. But by the previous paragraph the triple $y_ix_iy_i$ in the negated monomial can be reduced to a single negated monomial with $y_ix_iy_i$ replaced y_i. Thus the ultimate sum is 0 whichever reduction is applied first. A similar argument covers the appearance of $x_iy_nx_n$, $i \neq n$ in a monomial.

Thus the process is locally confluent and uniqueness follows from Newman's lemma (Lemma 47.2). □

The following corollaries apply to $L_k(1, n)$ for $n > 1$ and k a field or the integers, and to the submonoid $\mathcal{M}_n$ where k is not relevant.

COROLLARY 40.1.1. *The monoid $\mathcal{M}_n$ is presented with generators $\mathbf{X} = \{x_1, \ldots, x_n, y_1, \ldots, y_n\}$ and relations (40.2). The standard monomials give a set of normal forms for the elements of $\mathcal{M}_n$. Different standard monomials represent different elements of $\mathcal{M}_n$ and $L_k(1,n)$.*

Corollary 40.1.1 makes $\mathcal{M}_n$ isomorphic to the polycyclic monoid generated by X_y since we can define the polycylic monoid generated by X_y to be the monoid generated by $X_y \cup X_x$ subject to the relations (40.2). For example, see Section 7 of Birget 2004 [16]. The polycyclic monoids were mentioned in the end notes (Section 21) of Chapter 3. In Section 40.4.2, we will use the actions of the polycyclic monoids on Cantor sets.

COROLLARY 40.1.2. *If u and v are words over X_y, then v^*u is 0 if and only if neither u nor v is a prefix of the other. If in addition, u and v have the same length, then $v^*u = 1$ if $u = v$ and $v^*u = 0$ otherwise.*

40.2.2. *The restricted form.* An alternate normal form in a restricted setting can be based on (40.1) and (40.2). The restricted setting is specified in the statement of the next lemma.

LEMMA 40.2. *In $L_n = L_{\mathbf{Z}}(1,n)$ any element that can be represented by a sum of monomials from $\mathcal{M}_n$ with no negative coefficients can be represented by a unique linearly reduced sum of standard monomials with no negative coefficients so that no set S of summands is of the form $S = \{c_i u(y_i x_i)v^* \mid 1 \leq i \leq n\}$ with u and v words over $X_y = \{y_1, \ldots, y_n\}$ and $0 < c_i \in \mathbf{Z}$.*

Note that appearances of $y_n x_n$ are allowed in this form.

PROOF. Again we appeal to Newman's lemma. The outline of the argument is similar to that of Lemma 40.1 and we leave many details to the reader.

As in Lemma 40.1, a reduction of Type I is an elimination using (40.2) of an appearance of $x_i y_j$ in a summand followed by a complete linear reduction (combining of like terms).

A reduction of Type III starts with a set $\{c_i u y_i x_i v^* \mid 1 \leq i \leq n\}$, $c_i > 0$, of summands as in the statement of the lemma. With d the minimum value of the c_i, the reduction uses (40.1) to replace the sum $\sum_i c_i u y_i x_i v^*$ by

$$duv + \sum_i (c_i - d) u y_i x_i v^*$$

followed by a complete linear reduction.

Showing that random applications of the two types of reductions must terminate is easier than in Lemma 40.1. The argument that the reductions are locally terminating is similar to the argument in Lemma 40.1 and is left to the reader. The claim follows from Newman's lemma (Lemma 47.2). □

This has an important corollary.

Corollary 40.2.1. *The subset of $L_n = L_{\mathbf{Z}}(1,n)$ consisting of those elements representable by a sum of standard monomials with no negative coefficients is closed under addition and multiplication and is thus a subsemiring of L_n. This subset is also closed under $*$.*

Proof. Multiplication, addition and the reductions used to get the form of Lemma 40.2 introduce no negative coefficients. The last claim is immediate. □

40.2.3. *Positive elements.* The partial orders on the rings, modules and their endomorphisms will be defined by knowing the positive elements. All notions of positivity will have Corollary 40.2.1 as a basis.

In the ring $L_n = L_{\mathbf{Z}}(1,n)$, let PL_n be the set of all elements other than 0 that can be represented by a sum of standard monomials with no negative coefficients. We refer to PL_n as the *positive* elements of L_n and define $<$ on Ł$_n$ by $\alpha < \beta$ if $\beta - \alpha \in PL_n$. The corresponding $\leq$ is a partial order on L_n by Corollary 40.2.1.

A vector or matrix over L_n will be positive if it is not zero and all its non-zero entries are positive. We denote the positives in $M_r(L_n)$ by $PM_r(L_n)$. From Corollary 40.2.1, we know that products and sums of positive matrices over L_n are positive, and that if M is a $p \times q$ positive matrix, then $(PL_n^p)M \subseteq PL_n^q$. Also if A and B are positive, so is $A \oplus B$.

We extend the notion of unitaries from L_n to matrices over L_n and say that a $p \times q$ matrix M over L_n is unitary if $MM^* = \mathbf{I}_p$ and $M^*M = \mathbf{I}_q$. If M is a positive matrix, then M^* is also positive. We use $PUM_r(L_n)$ to denote the set of positive unitaries in $M_r(L_n)$.

Unitaries are units, and positive unitaries have positive inverses, but

$$\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}^{-1} = \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}$$

shows that not all positive units have positive inverses. We will show in Section 40.4.3 that $PUM_r(L_n)$ equals the set of positive units of $M_r(L_n)$ with positive inverses.

To set the tone for Theorem 40.10, the reader can verify that those $A \in GL(n, \mathbf{Z})$ where both A and A^{-1} have no negative entries are exactly the permutation matrices (defined below before Lemma 40.6).

40.3. Ideals. We show that there are very few 2-sided ideals.

PROPOSITION 40.3. *If k is a field, then $L_k(1,n)$ is simple. The only proper quotients of $L_{\mathbf{Z}}(1,n)$ are $\mathbf{Z}/r\mathbf{Z}$.*

The bulk of the work in proving the proposition is in the following lemma. The wording is motivated by Theorem 1.13 of [51].

LEMMA 40.4. *Let k be a field or the integers, and let Σ be a k-linear sum of standard monomials in $L_k(1,n)$ in the normal form of Lemma 40.1. Let $j \in \mathbf{Z}$ be greater than the length of any of the monomials (ignoring the coefficient) used in Σ. If cuv^* with $c \neq 0$ and uv^* standard is a term in Σ, then $(uy_n^{2j}y_1)^*\Sigma(vy_n^{2j}y_1) = c$.*

PROOF OF PROPOSITION 40.3. For a non-zero ideal I, all the coefficients of summands of elements of I expressed in the form of Lemma 40.1 are elements of I. If the scalars are a field, then these coefficients are invertible, so $1 \in I$. If the scalars are the integers, then the set of scalars in I is an ideal $r\mathbf{Z}$ of $\mathbf{Z} \subseteq L_{\mathbf{Z}}(1,n)$ that generates I. □

PROOF OF LEMMA 40.4. Let Σ, cuv^*, and j be as in the statement of the lemma.

A summand of $u^*\Sigma v$ is $u^*(cuv^*)v = c(u^*uv^*v) = c$. If dst^* with st^* standard is a summand in Σ different from cuv^*, we consider $u^*(dst^*)v$. If $u^*(dst^*)v = d(u^*st^*v) = d$, then checking cases and using Corollary 40.1.2 we see that $dst^* = duv^*$ which is not possible by our assumptions on Σ. So c is the constant term of $u^*\Sigma v$. Returning to $u^*(dst^*)v$, and again checking cases, we see that $u^*(dst^*)v = def^* \neq 0$ with e and f non-empty words over X_y can only occur if $|s| > |u|$ and $|t| > |v|$. Since u and v cannot both end in y_n, we see that e and f cannot both end in y_n.

Let $\Sigma' = u^*\Sigma v$. We know that c is the constant term of Σ', and we know that $2j$ is greater than the length of any of the monomials used in Σ'. One summand of $x_n^{2j}\Sigma' y_n^{2j}$ is $x_n^{2j}cy_n^{2j} = c$. We argue that all other summands of $x_n^{2j}\Sigma' y_n^{2j}$ are zero or are of a very restrictive form.

Consider dst^* with st^* standard that is a summand of Σ' different from c. Now $dx_n^{2j}st^*y_n^{2j}$ is zero unless both s and t are proper (by the choice of j) prefixes of y_n^{2j}. From a remark above, we cannot have both s and t non-empty and proper prefixes of y_n^{2j}. Thus $dx_n^{2j}st^*y_n^{2j} = 0$ unless the dst^* is of the form dy_n^m or dx_n^m for some $m \neq 0$, in which case $dx_n^{2j}st^*y_n^{2j}$ is also of the form dy_n^m or dx_n^m for some $m \neq 0$.

Now $\Sigma'' = x_n^{2j}\Sigma' y_n^{2j}$ has the form c plus summands of the form dy_n^m or dx_n^m for various values of $m \neq 0$. This gives $x_1\Sigma'' y_1 = c$. □

Corollary 40.4.1. *If k is a field, then a homomorphism with domain $L_k(1,n)$ is trivial or a monomorphism. A homomorphism with domain $L_{\mathbf{Z}}(1,n)$ is a monomorphism if its image is infinite.*

40.4. Two representations and a characterization. In this section we work with the Leavitt ring $L_n = L_{\mathbf{Z}}(1,n)$. We wish to prove that $V_{n,r}$ has a faithful representation into $M_r(L_n)$ whose image is exactly $PUM_r(L_n)$. Building a representation whose image is contained in $PUM_r(L_n)$ is reasonably straightforward, and will be done in two steps. The first step will be to build a representation of $V_{n,1}$ into $L_n = M_1(L_n)$.

To help characterize the image of the full representaion of $V_{n,r}$ into $M_r(L_n)$, we will also build a representation of $M_r(L_n)$ into the ring of endomorphisms of $C(r\mathfrak{C}_n, \mathbf{Z})$, the abelian group of continuous functions from r copies $r\mathfrak{C}_n$ of the n-ary Cantor set $\mathfrak{C}_n$ to $\mathbf{Z}$. Here and in the rest of this discussion, elements of $\mathfrak{C}_n$ will be viewed as infnite words over X_y. To save typing, we use K_n to denote $C(\mathfrak{C}_n, \mathbf{Z})$, and we note that $C(r\mathfrak{C}_n, \mathbf{Z})$ is isomorphic to K_n^r. Since $V_{n,r}$ already acts on $r\mathfrak{C}_n$, there is an induced action of $V_{n,r}$ on K_n^r, and this will guide the representation of $M_r(L_n)$.

To keep the actions parallel to each other we will be careful about which sides the actions are on. Since we have chosen to have $V_{n,r}$ to act on r copies of the n-ary Cantor set $r\mathfrak{C}_n$ on the right, the induced action of $V_{n,r}$ on K_n^r will be on the left. To be consistent, the elements in $M_r(L_n)$ will also act on the left.

40.4.1. *Representing $V_{n,r}$.* Before tackling $V_{n,r}$, we look at $V_{n,1}$.

We say that U is a prefix set for X_y if every infinite word in the alphabet X_y has a unique prefix in U. We order X_y using the order of the subscripts of the y_i, and prefix sets inherit the prefix order that is derived from the order on X_y. Using this order, we regard U as a row vector and U^* as a column vector so that if $|U| = d$, then UU^* is a 1×1 matrix and U^*U is an $d \times d$ matrix. We have the following.

Lemma 40.5. *If U is a prefix set for X_y with d elements, then $UU^* = \mathbf{I}_1$ and $U^*U = \mathbf{I}_d$.*

Proof. For the first equality, we note that $U = X_y$ is a prefix set for X_y and the equality holds in this case by (40.1). Every prefix set is obtained from X_y by a sequence of n-ary splittings where U' is obtained from U by an n-ary splitting at a given $u \in U$ if $u \in U$ is replaced by $\{uy_i \mid y_i \in X_y\}$. The equality follows easily by induction.

For the second equality, we note that since U is a prefix set, if $u \neq v$ in U, then neither u nor v is a prefix of the other. The equality follows from Corollary 40.1.2. ∎

We now look at $V_{n,1}$. Let (U, σ, V) be given with U and V prefix sets for X_y and $\sigma : U \to V$ a bijection. Using the prefix orders on U and V, the bijecttion σ can be also be regarded as a bijection on $\{1, \ldots, d\}$ with $d = |U| = |V|$ so that $v_{i\sigma} = (u_i)\sigma$.

Let M_σ be the $d \times d$ permutation matrix with entries $m_{i,j}$ that corresponds to σ in that $m_{i,j} = 1$ if $j = i\sigma$ and is otherwise 0. The inverse of M_σ is the (conjugate) transpose M_σ^*. Then $UM_\sigma V^* = \sum u_i v_{i\sigma}^*$. We will refer to $UM_\sigma V^*$ as the element of L_n corresponding to (U, σ, V). The element corresponding to (V, σ^{-1}, U) is $VM_\sigma^*, U^* = (UM_\sigma V^*)^*$.

LEMMA 40.6. *(I) Let (U, σ, V) be given with U and V prefix sets for X_y and $\sigma : U \to V$ a bijection. Then $UM_\sigma V^* \in L_n$ is unitary.*

(II) If we view the underlying n-ary Cantor set acted on by $V_{n,1}$ as words over X_y, and we take $f \in V_{n,1}$ represented by (U, σ, V) as in (I) to $\theta(f) = UM_\sigma V^*$, then θ is a well defined homomorphic embedding from $V_{n,1}$ into the positive unitaries of L_n.

PROOF. (I) Let $d = |U| = |V|$. We have

$$UM_\sigma V^*(UM_\sigma V^*)^* = UM_\sigma V^* V M_\sigma^* U^* = U\mathbf{I}_d U^* = UU^* = \mathbf{I}_1,$$

and similarly for $(UM_\sigma V^*)^* UM_\sigma V$. This argument is slick, but it is also worth writing out that since V is a prefix set, we have

$$\sum_i u_i v_{i\sigma}^* \sum_j v_{j\sigma} u_j^* = \sum_i \sum_j u_i v_{i\sigma}^* v_{j\sigma} u_j^* = \sum_i u_i u_i^* = 1 \tag{40.4}$$

because $v_{i\sigma}^* v_{j\sigma} = 0$ unless $i = j$.

(II) Invariance of $\theta(f)$ under n-ary splitting is straightforward by noting that after a splitting that replaces some y_i in U with the n-tuple $(y_i y_1, y_i y_2, \ldots, y_i y_n)$ in U and a corresponding replacement of $(y_i)\sigma$ in V, the modified σ' will take $y_i y_j$ to $(y_i)\sigma y_j$, and near the end of the calculation of (40.4), a sum of the form

$$\sum_{j=1}^{n} (y_i y_j)(y_i y_j)^* = y_i y_i^*$$

will appear. That θ is a homomorphism is immediate. It is an embedding because $V_{n,1}$ has no normal subgroups other than possibly a subgroup of index 2. □

To get a representation of $V_{n,r}$, we modify the represention of $V_{n,1}$. Let $f \in V_{n,r}$ be represented by (U, σ, V) where U and V are sequences of words in X_y of the same length d. The sequence U is the concatenation of a sequence of r sequences $(U_1, \ldots, U_r)$ where each U_i is a prefix set for the n-ary Cantor set $\mathfrak{C}_n$ under the prefix order. Similarly V is the concatenation of $(V_1, \ldots, V_r)$ of similar description. We do not require that for each i, the length of U_i is the same as that of V_i. The bijection σ is then between sets of size d and is to be regarded as a bijection on $\{1, \ldots, d\}$ so that we can overuse the letter σ and write $(u_i)\sigma = u_{i\sigma}$.

We let M_σ be the $d \times d$ permutation matrix for σ. With each U_i and V_i a row vector, we let $M_U = U_1 \oplus \cdots \oplus U_r$ and $M_V = V_1 \oplus \cdots \oplus V_r$ (see Section 40.1.1), and let $M_f = M_U M_\sigma (M_V)^*$. Each of M_U and M_V is an $r \times d$ matrix, and M_f is an $r \times r$ matrix. We note that $M_U(M_U)^* = \mathbf{I}_r = M_V(M_V)^*$ and $(M_U)^* M_U = \mathbf{I}_d = (M_V)^* M_V$. As in (40.4) we get

$$\begin{aligned} M_f(M_f)^* &= (M_U M_\sigma (M_V)^*)(M_U M_\sigma (M_V)^*)^* \\ &= M_U M_\sigma (M_V)^* M_V (M_\sigma)^* (M_U)^* \\ &= M_U M_\sigma \mathbf{I}_d (M_\sigma)^* (M_U)^* \\ &= M_U \mathbf{I}_d (M_U)^* = \mathbf{I}_r, \end{aligned}$$

and similarly $(M_f)^* M_f = \mathbf{I}_r$. Invariance under matched splittings is as in Lemma 40.6 and we get the following. The embedding comes from the fact that the image has more than two elements.

Proposition 40.7. *Sending $f \in V_{n,r}$ to $\theta(f) = M_f$ in $M_r(L_n)$ makes θ a homomorphic embedding into $PUM_r(L_n)$.*

40.4.2. *Representing $M_r(L_n)$.* We build the representation in steps, and we start with monomials. Monomials will first act by partial bijections on $\mathfrak{C}_n$ on the right, and then the action will be linearized to act on K_n on the left.

Because of our choices about the sides of our actions, we will have to twist unnaturally so as to make the actions cooperate naturally with the algebra. Doing so will make clear that the actions do in fact commute with the multiplication.

Elements of $\mathcal{M}_n$ are words of the form uv^* where u and v are finite words over X_y. Recall that elements of $\mathfrak{C}_n$ are viewed as infinite words over X_y. Given a standard monomial $uv^* \in \mathcal{M}_n$ in L_n and a word w in X_y^ω, we set

$$w(uv^*) = \begin{cases} vw', & w = uw', \\ \text{undefined}, & \text{otherwise}. \end{cases} \tag{40.5}$$

That is, uv^* takes $u\mathfrak{C}_n$ rigidly to $v\mathfrak{C}_n$ and is undefined on the complement of $u\mathfrak{C}_n$.

We can relate (40.5) to the algebra of L_n by contorting things. In parallel to the view of $\mathfrak{C}_n$ as words in X_y^ω, we can also regard $\mathfrak{C}_n$ as the set of infinite words over $X_x = \{x_1, \ldots, x_n\} = \{y_1^*, \ldots, y_n^*\}$ written from right to left (i.e., backwards). We "equate" $w \in X_y^\omega$ written left to right with $w^* \in X_x^\omega$ written right to left. Now $(w')^*u^*uv^* = (w')^*v^*$.

Note that the set of partially defined transformations $u\mathfrak{C}_n \to v\mathfrak{C}_n$ that are rigid ($u\alpha \mapsto v\alpha$ for all $\alpha \in \mathfrak{C}_n$) is closed under composition if we add an element that is undefined everywhere since two cones are either disjoint or nested. The set of such transformations is thus a symmetric inverse semigroup (see Lawson [129] or [133]).

If we specialize the definition to $y_i = y_i1$ and $x_i = 1x_i$, we get

$$\begin{aligned} w(x_i) &= (y_i w), \\ w(y_i) &= \begin{cases} w', & w = y_i w', \\ \text{undefined}, & \text{otherwise.} \end{cases} \end{aligned} \tag{40.6}$$

If we view $\mathfrak{C}_n$ as the ends of the complete n-ary tree drawn with the root at the top, then x_i "lowers" $\mathfrak{C}_n$ to the cone $y_i\mathfrak{C}_n$, and y_i "raises" the cone $y_i\mathfrak{C}_n$ to the entire Cantor set. These are not elements of of the Thompson monoid $\mathscr{M}$ of Section 16, but they are elements of the monoid $invM_{n,1}$ of Birget 2009 [17].

Using $(w')^*u^*uv^* = (w')^*v^*$ from above, or simply checking cases, it is seen that the action commutes with composition in $\mathcal{M}_n$. From (40.6) the action preserves the relations (40.2). From Corollary 40.1.1, we have a homomorphism from $\mathcal{M}_n$ to the partial symmetries of $\mathfrak{C}_n$ and it is clear from the definition that this is injective.

To fulfill the status of $\mathcal{M}_n$ as an abstract inverse semigroup (every a has a unique a^{-1} for which $aa^{-1}a = a$ and $a^{-1}aa^{-1} = a^{-1}$) we put $(uv^*)^* = vu^*$ in the role of $(uv^*)^{-1}$. Its status as an inverse semigroup gives certain early facts which are also obvious from the representation as a symmetric inverse semigroup. Idempotents are all of the form uu^*, and idempotents commute. The idempotents are in one-to-one correspondence with the cones in $\mathfrak{C}_n$. That $\mathcal{M}_n$ is an (abstract) inverse semigroup and the facts just given can also be derived as an interesting exercise directly from the presentation of Corollary 40.1.1.

To obtain an action of $\mathcal{M}_n$ on K_n, we replace "undefined" by 0. Given a standard monomial uv^* in $\mathcal{M}_n$, a function $f : \mathfrak{C}_n \to \mathbf{Z}$, and $w \in \mathfrak{C}_n$, we set

$$w((uv^*)f) = (w(uv^*))f = \begin{cases} (vw')f, & w = uw' \\ 0, & \text{otherwise.} \end{cases} \tag{40.7}$$

Thinking of K_n as a $\mathbf{Z}$-module, the element uu^* is a projection of K_n to the submodule of elements in K_n whose support is in the cone $u\mathfrak{C}_n$.

We extend the action of $\mathcal{M}_n$ on K_n defined by (40.7) to L_n linearly. This action cooperates with addition. The action cooperates with multiplication on $\mathcal{M}_n$, and the only further check needed is to verify that the relation (40.1) is preserved.

The Cantor set $\mathfrak{C}_n$ is the disjoint union of the cones in $\{y_i\mathfrak{C}_n \mid 1 \leq i \leq n\}$. So K_n is the direct sum

$$K_n \simeq \sum_{i=1}^{n} \oplus C(y_i\mathfrak{C}_n, \mathbf{Z}), \tag{40.8}$$

and the identity on K_n is the sum of the projections of K_n to the $C(y_i\mathfrak{C}_n, \mathbf{Z})$. Since $(y_i x_i f)w = w$ if w starts with y_i and is 0 otherwise, the $y_i x_i$ are exactly these projections. This verifies that the relation (40.1) is satisfied and we have a well defined representation of L_n in the endomorphism ring of K_n. This is also straightforward to verify formally and is left to the reader. Thus we have a homomorphism.

From Corollary 40.4.1, the homomorphism is a monomorphism if its image is infinite. The integer j takes the constant function from $\mathfrak{C}_n$ to 1 to the constant function from $\mathfrak{C}_n$ to j. We have shown the following.

Proposition 40.8. *The function based on (40.7) from the Leavitt ring $L_n = L_{\mathbf{Z}}(1, n)$ to $\mathrm{End}(K_n)$, the endomorphisms of the group of continuous functions from the Cantor set $\mathfrak{C}_n$ to the integers, is a homomorphic embedding.*

It is now straightfoward to define the representation of $M_r(L_n)$. We regard the r-tuples that form K_n^r as $r \times 1$ column matrices. The image under an $r \times r$ matrix $A \in M_r(L_n)$ with (i, j)-entries $A_{i,j} \in L_n$ of an r-tuple (f_i) in K_n^r is the r-tuple (g_i) where

$$g_i = \sum_{j}^{r} A_{i,j} f_j. \tag{40.9}$$

If some $A_{i,j} \neq 0$ with $i \neq j$ is an entry in A, then from Proposition 40.8, there is an $f \in K_n$ with $A_{i,j} f \neq 0$ and A cannot act as the identity on the tuple in K_n^r whose only non-zero coordinate is $f_j = f$. So any A acting as the identity must be diagonal. Now Proposition 40.8 says that an A acting as the identity is the identity matrix. We have shown the following.

Proposition 40.9. *The action of $M_r(L_n)$ on K_n^r just defined gives a homomorphic embedding of $M_r(L_n)$ in the endomorphisms of K_n^r.*

The composition of the representation of $V_{n,r}$ into $M_R(L_n)$ with the action of $M_r(L_n)$ on K_n^r gives the action of $V_{n,r}$ on K_n^r induced by the action of $V_{n,r}$ on $r\mathfrak{C}_n$. We have no particular use for this fact.

40.4.3. *Characterizing the image.* The following is a slight extension of Proposition 3.5 from [168].

THEOREM 40.10. *The set of positive units in $M_r(L_n)$ with positive inverses equals the set $PUM_r(L_n)$ of positive unitaries in $M_r(L_n)$. The map $\theta : V_{n,r} \to M_r(L_n)$ is a homomorphic embedding with respect to the multiplication in $M_r(L_n)$ whose image is $PUM_r(L_n)$.*

PROOF. The proof is mostly accounting. For positive $A \in M_r(L_n)$ we keep track of how often prefixes and suffixes are used. We give the idea first and precision after. For words u and v over X_y, the count $row_i(u)$ will be the number of times u is used in row i of A as the prefix of a monomial uv^*, and the count $col_j(v)$ will be the number of times v^* is used in column j of A as the suffix of a monomial uv^*. We can be more precise.

The (i,j) entry $A_{i,j}$ of A is an element $\sum c_p u_p v_p^*$ with $c_p \in \mathbf{Z}_{>0}$, and u_p and v_p words over X_y. We let $\rho_{i,j}(u) = \sum c_p$ with the sum taken over all p where $u_p = u$, and we let $\kappa_{i,j}(v) = \sum c_p$ with the sum taken over all p where $v_p^* = v^*$. Now

$$row_i(u) = \sum_j^r \rho_{i,j}(u), \text{ and}$$

$$col_j(v) = \sum_i^r \kappa_{i,j}(v).$$

CLAIM 1. *For positive, invertible A in $M_r(L_n)$ with positive inverse, and $\tilde{r} = \{1,\dots,r\}$ the following hold.*

(I) For $i \in \tilde{r}$, every word u and v over X_y has $row_i(u) \in \{0,1\}$ and $col_j(v) \in \{0,1\}$.

(II) For $i \in \tilde{r}$, the set $\{u \mid row_i(u) = 1\}$ is a prefix set for $\mathfrak{C}_n$.

(III) For $j \in \tilde{r}$, the set $\{v \mid col_j(v) = 1\}$ is a prefix set for $\mathfrak{C}_n$.

PROOF OF CLAIM 1. We eliminate sources of failure of (I)–(III). Note that the number of words u, v with $row_i(u)$ or $col_j(v)$ not zero is finite. We focus on the fact that A must carry K_n^r to itself bijectively.

If for some $j \in \tilde{r}$, the set $V_j = \{v \mid col_j(v) \geq 1\}$ is not a prefix set because some word in X_y^ω has no prefix in V_j, then A cannot be one-to-one by considering Af on an $f \neq 0$ with $f_i = 0$ for $i \neq j$ and f_j supported on some $v\mathfrak{C}_n$ where no element of $v\mathfrak{C}_n$ has a prefix in V_j. If for some $i \in \tilde{r}$, the set $U_i = \{u \mid row_i(u) \geq 1\}$ is not a prefix set because some word in X_y^ω has no prefix in U_i, then A cannot be onto by noting that $(Af)_i$ must always be zero on $u\mathfrak{C}_n$ for u where no element of $u\mathfrak{C}_n$ has a prefix in U_i.

We do not yet know that $A^{-1} = A^*$, but we can apply what we know so far to A^{-1}. So we know that for $f \in K_n^r$ regarded as a function from $\tilde{r}$ to K_n, every value of f makes a contribution somewhere to Af and also to $A^{-1}f$.

We will use the constant function $\overline{1}$ from $r\mathfrak{C}_n$ to 1 and concentrate on the fact that AA^{-1} has to take $\overline{1}$ to itself. From the remark just made about Af and $A^{-1}f$, we know that $A\overline{1}$ cannot have any value greater than 1. From this we know that every entry $A_{i,j}$ of A is a linearly reduced sum of monomials of the form $\sum 1u_pv_p^*$. Secondly, if $row_i(u) > 1$ for some $u \in U_i$, then $A\overline{1}$ is greater than 1 on $u\mathfrak{C}_n$. This proves that part of (I) about the values $row_i(u)$. Lastly, if some U_i fails to be a prefix set because there are u and u' in U_i with u' a prefix of u, then $A\overline{1}$ on $u\mathfrak{C}_n$ will also have value greater than 1. This proves (II).

If some $col_j(v)$ is greater than 1, then since all coefficients in the $A_{i,j}$ are 1, the suffix v^* must show up in at least two monomials u_1v^* and u_2v^*. Either these monomials are in the same $A_{i,j}$ and have $u_1 \neq u_2$, or the monomials are in different rows of A and the two prefixes are from different U_i. In either case it follows from the fact that (II) is known that the cones $u_1\mathfrak{C}_n$ and $u_2\mathfrak{C}_n$ are disjoint cones in $r\mathfrak{C}_n$ whether they are in the same copy of $\mathfrak{C}_n$ or not. In this situation no Af can have different values on these cones and A cannot be onto. This completes the proof of (I).

The last case to consider is if some v_1 is a proper prefix of v_2 with both v_1 and v_2 in some V_j. Then $u_1v_1^*$ and $u_2v_2^*$ are monomials in entries in A. If $v_2 = v_1v_3$, then let $u_3 = v_3^*$. Because (II) is known to hold, we have that for all f, the values of Af on $(u_1u_3)\mathfrak{C}_n$ must equal the values of Af on $u_2\mathfrak{C}_n$, and A cannot be onto. This completes the proof of (III) and the claim. □

Proof of Theorem 40.10 (cont.) From Proposition 40.7 we know that $\theta(V_{n,r})$ is contained in $PUM_r(L_n)$ which is then contained in those positive, invertible elements of $M_r(L_n)$ with positive inverses. Now Claim 1 shows that every positive, invertible element of $M_r(L_n)$ with positive inverse is in the image of θ. This completes the proof. □

41. Isomorphisms from Leavitt algebras

We continue to use L_n to denote the ring $L_{\mathbf{Z}}(1, n)$.

Theorems 38.29 and 38.30 of Section 38.8.1 show that if the groups $V_{n,r}$ and $V_{n',r'}$ are isomorphic, then $n = n'$ and $\gcd(r, n-1) = \gcd(r', n'-1)$. Our goal is to obtain a converse via the followng.

THEOREM 41.1. *If* $n' = n$ *and* $\gcd(r, n-1) = \gcd(r', n-1)$*, then there is an order preserving isomorphism of rings* $M_r(L_n)$ *and* $M_{r'}(L_n)$*.*

From this we get the following.

THEOREM 41.2. *If* $n' = n$ *and* $\gcd(r, n-1) = \gcd(r', n-1)$*, then there is an isomorphism between* $V_{n,r}$ *and* $V_{n,t'}$ *as well as an isomorphism between* $V^+_{n,r}$ *and* $V^+_{n,r'}$*.*

PROOF. This follows from Theorem 41.1 and from Theorem 40.10 which represents the various $V_{n,r}$ as characteristic subgroups of the $M_r(L_n)$. □

In Section 38.8 it is mentioned that Theorem 7.5 of [107] gives that if c is a divisor of n, then $V_{n,r}$ and $V_{n,cr}$ are isomorphic. Note that this is a consequence of Theorem 41.2 since if $c|n$, then $\gcd(c, n-1) = 1$, and $\gcd(r, n-1) = \gcd(cr, n-1)$ because they divide each other. A case not covered by combining the result from [107] with Lemma 38.2, but covered by Theorem 41.2, is the case $n = 5$, $r = 1$ and $r' = 3$. From Lemma 38.2 we know $V_{5,1}$ is isomorphic to all $V_{5,r}$ with $r \equiv 1 \pmod 4$, and $V_{5,3}$ is isomorphic to all $V_{5,r}$ with $r \equiv 3 \pmod 4$. From Theorem 7.5 of [107] we can only learn that $V_{4,r}$ is isomorphic to $V_{4,5r}$. But $5 \equiv 1 \pmod 4$, so this information cannot combine the two classes. But $\gcd(3, 5-1) = 1$, and that $V_{5,1}$ is isomorphic to $V_{5,3}$ follows from Theorem 41.2.

41.1. Reductions. Theorem 41.1 will be derived from the important special case proven in [1] that assumes $\gcd(r, n-1) = 1$. The result in this case follows if it is proven that there is an order preserving isomorphism from $M_1(L_n)$ to $M_r(L_n)$ for all r prime to $n-1$. But $M_1(L_n)$ is just the ring of 1×1 matrices over L_n and is order isomorphic to the ring L_n. So our special case is the following.

THEOREM 41.3. *When* $\gcd(r, n-1) = 1$*, there is an order preserving ring isomorphism from* L_n *to the ring of* $r \times r$ *matrices over* L_n*.*

The proof of Theorem 41.3 will be the bulk of the effort. The proof is aided by the fact that much is known about the ring L_n.

The reduction of Theorem 41.1 to Theorem 41.3 uses the following two lemmas. The first is proven as Lemma 50.4 in the Appendix.

Lemma 41.4. *Let a, b, s be integers with both a and b not zero and with $g = \gcd(a, s) = \gcd(b, s)$. Then some integer x prime to s is a solution to $ax \equiv b \pmod{s}$.*

Next is a parallel to Lemma 38.2.

Lemma 41.5. *If $r \equiv r' \pmod{n-1}$, then there is an order preserving isomorphim of L_n-modules between L_n^r and $L_n^{r'}$ and order preserving isomorphism of rings between $M_r(L_n)$ and $M_{r'}(L_n)$. For all r, the L_2 modules L_2^r are order isomorphic and all the rings $M_r(L_2)$ are order isomorphic.*

Proof. The matrix $A = (y_1, \ldots, y_n)$ from Section 40.1 gives an isomorphism from L_n^1 to L_n^n and A^* gives the reverse isomorphism. Both A and A^* are positive so the isomorphisms are order preserving.

The matrix $\mathbf{I}_{k-1} \oplus A$ and its $*$-transpose give order preserving isomorphisms between L_n^k and $L_n^{k+(n-1)}$. Conjugations using these inverse isomorphisms induce order preserving isomorphisms between $M_k(L_n)$ and $M_{k+(n-1)}(L_n)$. The first sentence of the lemma follows by composing these isomorphisms for the various k. The second sentence follows because $2 - 1 = 1$. □

Proof of Theorem 41.1 from Theorem 41.3. We assume that $\gcd(r, n-1) = \gcd(r', n-1)$. Lemma 41.4 gives a unit u in $\mathbf{Z}/(n-1)\mathbf{Z}$ with $r' = ru \pmod{n-1}$. With $\gcd(u, n-1) = 1$, we have $M_r(L_n)$ is order isomorphic to $M_r(M_u(L_n))$ where elements of the latter are regarded as block matrices with $u \times u$ matrices as the blocks. The latter is order isomorphic to $M_{ru}(L_n)$ which by Lemma 41.5 is order isomorphic to $M_{r'}(L_n)$. □

41.2. The proof of Theorem 41.3. From this point $\gcd(r, n-1) = 1$. We can make other restrictions. From the first conclusion of Lemma 41.5, we can take r as large as we like, and we assume $r > n$. From the second conclusion of Lemma 41.5, we can assume $n \geq 3$. These assumptions will see repeated use.

We must build an order preserving isomorphism. In spite of the fact that we are not forced to use the involution $*$, we will use it. It would seem crippling not to. As a ring with involution, L_n is generated by $X_y = \{y_1, \ldots, y_n\}$. We must find matrices $Y_1, \ldots, Y_n$ in $M_r(L_n)$ to be the images of the y_i. To show that $y_i \mapsto Y_i$ extends to a ring homomorphism η, we must show that the Y_i satisfy the relations (40.1) and (40.2). To show that η is injective, Proposition 40.3 tells us that it suffices to show that the image is not finite. To show that η is surjective, we must show that all the matrices $e_{i,j}$ and $y_k e_{i,j}$ are in the image of η

where $e_{i,j}$ is the $r \times r$ matrix having as its only non-zero entry a 1 in position (i, j). To show that η is an isomorphism of ordered structures, we need to show that it restricts to a bijection on the positive elements of the domain and range.

Difficulties arise from the semi-predictable relation between n and r. To set the tone for what follows, we quote [58] where it is remarked that the beautiful result from [1] "shows that two naturally defined rings are isomorphic without giving a natural reason, and there may not be one." What follows in the rest of this section is a carefully crafted mechanism that produces the desired isomorphism.

41.2.1. *Bookkeeping.* Deriving all the $e_{i,j}$ from the Y_i makes use of a semi-regular pattern set up modulo r. We define $\pi : \mathbf{Z} \to \mathbf{Z}$ by

$$i\pi = \begin{cases} i+n, & i \equiv 0 \pmod r, \\ i+(n-2), & i \equiv 1 \pmod r, \\ i+(n-1), & \text{otherwise.} \end{cases} \tag{41.1}$$

The pattern established by π is that all integer values are increased by $n-1$, except the pairs $(kr, kr+1)$ which are increased by $n-1$ and then interchanged. Thus π is a bijection.

In the following, there will be constant reference to intervals of integers. We have little reason to discuss intervals of real numbers. So for the rests of this argument, we will use $[i, j]$ to denote the set $\{k \in \mathbf{Z} \mid i \le k \le j\}$.

LEMMA 41.6. *The following are true about the mapping π.*

(I) An interval of $n-1$ consecutive integers starting with k is a set of orbit representatives of π if and only if $k \not\equiv 1 \pmod r$.

(II) The $r-1$ consecutive, disjoint sets of $n-1$ consecutive integers starting at n are all sets of orbit representatives of π.

PROOF. (I) For one direction, let $J = [k, k+(n-2)]$ with $k \not\equiv 1 \pmod r$. The only integers increased by π by less than $n-1$ are those equivalent to 1 modulo r. So J does not duplicate any orbits. No integer is increased by more than n so $J \cup \{k-1\}$, which includes $(k-1) \not\equiv 0 \pmod r$, has a representative for every orbit. But now $(k-1)\pi \le (k-1)+(n-1)$ which lies in J. So every orbit is represented in J. The other direction is clear.

(II) The least elements of the intervals described form an arithmetic sequence of constant difference $n-1$. Since $\gcd(r, n-1) = 1$, the sequence cycles through the residues modulo r. Since n is the integer immediately following 1 in the sequence, all of n and the next $r-2$ entries of the sequence are not equivalent to 1 modulo r. The claim now follows from (I). □

For $j \in [1, r-1]$ let J_j be the j-th set of consecutive $n-1$ integers described in (II) of Lemma 41.6 with $j < k$ giving $J_j < J_k$. Specifically $J_j = [n+(j-1)(n-1), n+j(n-1)-1]$, and $i < j$ gives $J_i \cap J_j = \emptyset$. By (I) of Lemma 41.6, $[2, n]$ and each J_j, $j \in [1, r-1]$ contains a unique representative for every orbit of π. For $s \in [2, r]$ and $j \in [1, r-1]$, let $s \cdot j$ be the unique element of $\mathbf{Z}$ in J_j in the π-orbit of s.

41.2.2. *The key cycle.* The sequence with constant difference $(n-1)$ used in the proof of (II) of Lemma 41.6 will show up repeatedly, in reverse, modulo r as $i \mapsto i - (n-1) \pmod r$, on the interval $[1, r]$. Our assumptions that $2 < n < r$ give that $1 < n-1 < n < r$ are four different elements of $[1, r]$ and so the transitions $n \mapsto 1$ and $(n-1) \mapsto r$ involve no common elements. If inductions are done by use of the cycle, and one or both of these transitions are not available for the induction, then one or both of "true for 1" and "true for r" will need to be assumed. The diagram below will help visualize the points made about the cycle.

$$\begin{array}{ccc} r & \longrightarrow\ \longrightarrow \cdots \longrightarrow\ \longrightarrow & n \\ \uparrow & & \downarrow \\ n-1 & \longleftarrow\ \longleftarrow \cdots \longleftarrow\ \longleftarrow & 1 \end{array} \tag{41.2}$$

The use of representatives of $\mathbf{Z}/r\mathbf{Z}$ in $[1, r]$ is so frequent, that we will use the notation $a\,[\mathrm{mod}\, r]$ to indicate the element in $[1, r]$ equivalent to a modulo r.

The key aspect of the permuation π is that π, modulo r, is essentially the inverse of the action $i \mapsto i - (n-1)\,[\mathrm{mod}\, r]$ on $[1, r]$, but it interchanges the "targets" of r and 1 in the direction travelled by π. This it interchanges the targets of $n-1$ and n under the cycle of $i \mapsto i - (n-1)\,[\mathrm{mod}\, r]$ on $[1, r]$. Thus π breaks the cycle into two orbits which are the top and bottom lines of (41.2).

41.2.3. *Building the homomorphism.* Our task is to define n matrices in $M_r(L_n)$. All of the entries will be 0 except that each column of each matrix will have a single non-zero entry. Of these rn entries $r-2$ entries will be 1, n entries will be the elements y_1 through y_n, one entry will be y_1^{r-1}, and the remaining $rn - (r-2) - 1 - n = (r-1)(n-1)$ entries will be defined with the help of the behavior of $s \cdot j$.

We let A be the row vector $(y_1, \ldots, y_n)$, $\mathbf{I}_{r-2}$ be the $(r-2) \times (r-2)$ identity matrix, and Z be the row vector

$$(z_{r+n-1}, z_{r+n}, \ldots z_{rn})$$

containing $(r-1)(n-1)+1$ entries where where

$$z_{r+n-1} = y_1^{r-1} \quad \text{and} \quad z_{r+(s\cdot j)} = y_1^{j-1} y_s.$$

The strange indexing is to place z_i in column i of the matrix Y defined below.

No two z_i have been defined to be the same element in L_n. The elements of L_n defined as some z_i include y_1^{r-1} and all $y_1^p y_q$ for all $(p, q) \in [0, r-2] \times [2, n]$. The underlyingn set of each of A and Z is a prefix set for X_y^*.

We let $y_i^* = x_i$ from Section 40.1, and for a matrix M, we let M^* be the $*$-transpose of M. From Lemma 40.5, we have all of $AA^* = \mathbf{I}_1$, $A^*A = \mathbf{I}_n$, $ZZ^* = \mathbf{I}_1$, and $Z^*Z = \mathbf{I}_k$ with $k = (r-1)(n-1)+1$. Now the $r \times rn$ matrix

$$Y = A \oplus \mathbf{I}_{r-2} \oplus Z = \begin{pmatrix} A & 0 & 0 \\ 0 & \mathbf{I}_{r-2} & 0 \\ 0 & 0 & Z \end{pmatrix}$$

satisfies

$$YY^* = \mathbf{I}_r \quad \text{and} \quad Y^*Y = \mathbf{I}_{nr}. \tag{41.3}$$

The matrix Y can be given the block structure $\begin{pmatrix} Y_1 & Y_2 & \cdots & Y_n \end{pmatrix}$ where each Y_i is an $r \times r$ matrix. We illustrate Y_1 and Y_2 below since they differ greatly from the other Y_i, and the nature of their internal structures will be used repeatedly.

$$\underbrace{\begin{bmatrix} y_1 & \cdots & y_n & 0 & 0 & \cdots & 0 & 0 \\ 0 & \cdots & 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & & \vdots & 0 & 1 & \cdots & 0 & 0 \\ \vdots & & \vdots & 0 & 0 & \ddots & 0 & 0 \\ \vdots & & \vdots & 0 & 0 & \cdots & 1 & 0 \\ 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 1 \\ 0 & \cdots & & \cdots & & \cdots & & 0 \\ \vdots & & & & & & & \vdots \\ 0 & \cdots & & \cdots & & \cdots & & 0 \end{bmatrix}}_{Y_1} \begin{array}{l} 1 \text{ row} \\ r-n \text{ rows} \\ n-1 \text{ rows} \end{array} \qquad \underbrace{\begin{bmatrix} 0 & & \cdots & & \cdots & & \cdots & 0 \\ \vdots & & & & & & & \vdots \\ 0 & & \cdots & & \cdots & & \cdots & 0 \\ 1 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 & 0 & \vdots & & \vdots \\ 0 & 0 & \ddots & 0 & 0 & \vdots & & \vdots \\ 0 & 0 & \cdots & 1 & 0 & \vdots & & \vdots \\ 0 & 0 & \cdots & 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 & 0 & z_s & \cdots & z_t \end{bmatrix}}_{Y_2} \begin{array}{l} r-n+1 \text{ rows} \\ n-2 \text{ rows} \\ 1 \text{ row} \end{array}$$

$(s = r + n - 1)$
$(t = 2r)$

The relations in (41.3) are a restatement of the relations in (40.1) and (40.2), and they verify that sending y_i in L_n to Y_i in $M_r(L_n)$ extends to a ring homomorphism η. In particular $\sum_{k=1}^{n} Y_k Y_k^* = \mathbf{I}_r$.

Each Y_i has infinite order in the additive structure and so the image of the homomorphism is infinite. As remarked, it follows from Proposition 40.3 that η is injective. The Y_i and Y_i^* are all positive (are in

$PM_r(L_n))$ and so η takes the positives in L_n (those in P) into the positives in $M_r(L_n)$. To prove that η is an isomorphism of ordered rings, we need to prove that η is surjective and carries P onto $PM_r(L_n)$. It suffices to prove the latter.

41.2.4. *On surjectivity.* We will show that all the $e_{i,j}$ and all $y_k e_{i,j}$ are in the image of P under η. This will complete the proof of Theorem 41.3.

Let $S = \eta(P)$. We know $S \subseteq PM_r(L_n)$. We show $PM_r(L_n) \subseteq S$ by building elements of S as products and sums of the Y_i and Y_i^*. As an aid we let O_k be the $k \times k$ zero matrix, and for $1 \leq i \leq r$ we define

$$\begin{aligned} E^i &= \mathbf{I}_i \oplus \mathrm{O}_{r-i}, \quad \text{and} \\ E_{r-i} &= \mathrm{O}_{r-i} \oplus \mathbf{I}_i, \quad \text{so} \\ \mathbf{I}_r &= E^r = E_0 = E^i + E_i. \end{aligned}$$

The order of acquisition in S will be first the E^i and E_i, then the $e_{i,i}$, and then the $e_{i,j}$. The $y_k e_{i,j}$ will be picked up toward the end.

41.2.5. *The E_i.* We start with one easy item and two special cases of what follows.

$$\begin{aligned} \mathbf{I}_r &= \sum_{k=1}^{n} Y_k Y_k^* \in S, \\ E^{r-(n-1)} &= Y_1 Y_1^* \in S, \\ E_{r+(n-1)} = \mathbf{I}_r - E^{r-(n-1)} &= \sum_{k=2}^{n} Y_k Y_k^* \in S. \end{aligned} \tag{41.4}$$

For part of an induction, we note that if $1 \leq i \leq (n-2)$, then

$$\begin{aligned} E^{i+r-(n-1)} &= E^{r-(n-1)} + Y_2 E^i Y_2^* = \\ &\qquad Y_1 Y_1^* + Y_2 E^i Y_2^* \in S \quad \text{if} E^i \in S, \quad \text{and} \\ E_{i+r-(n-1)} &= \mathbf{I}_r - E^{i+r-(n-1)} = \\ &Y_2(\mathbf{I}_r - E^i) Y_2^* + \sum_{k=3}^{n} Y_k Y_k^* = \\ &\qquad Y_2 E_i Y_2^* + \sum_{k=3}^{n} Y_k Y_k^* \in S \quad \text{if} E_i \in S. \end{aligned} \tag{41.5}$$

If $n \le i \le r$, then

$$E^{i-(n-1)} = Y_1 E^i Y_1^* \in S \ \text{if} E^i \in S, \quad \text{and}$$

$$(41.6) \qquad E_{i-(n-1)} = \mathbf{I}_r - E^{i+(n-1)} = Y_1 E_i Y_1^* + \sum_{k=2}^{n} Y_k Y_k^* \in S \ \text{if} E_i \in S.$$

Thus for all values of i from 1 through r except $i = (n-1)$, if $E^i \in S$, then $E^j \in S$ with $j \equiv i - (n-1) \pmod r$. Similary with the same restrictions on i, if $E_i \in S$, then $E_j \in S$ with $j \equiv i - (n-1) \pmod r$. The cycle (41.2) applies here with the transition $(n-1) \mapsto r$ not available. But E^r is in S, and while $E_r = \mathbf{O}_r$ is not relevant, the next element in the cycle $E_{r+(n-1)}$ is in S. So inductively, all E^i for $1 \le i \le r$ and all E_i for $1 \le i < r$ are in S.

41.2.6. *The* $e_{i,i}$. We start work on the $e_{i,i}$ by noting that $e_{1,1} = E^1$ and $e_{r,r} = E_{r-1}$ are both in S. We have the conjugation result $e_{j,i} e_{i,i} e_{i,j} = e_{j,j}$, and more generally we get $M^{-1} e_{i,i} M = e_{j,j}$ if the only non-zero entry in row i and in column j of M is a 1 in position (i, j).

Using Y_2^* as M, we have this type of 1 in position $(i, i+r-(n-1))$ for $i \in [1, n-2]$, and using Y_1^* as M, we have this in position $(i, i-(n-1))$ for $i \in [n+1, r]$. So the conjugations $Y_1 e_{i,i} Y_1^*$ and $Y_2 e_{i,i} Y_2^*$ both give $e_{j,j}$ with $j \equiv i - (n-1) \pmod r$ in the stated ranges.

The cycle (41.2) applies with both restrictions mentioned in Paragraph 41.2.2 effect. Since we know $e_{1,1}$ and $e_{r,r}$ are already in S, we get all $e_{i,i} \in S$ for all $1 \le i \le r$.

41.2.7. *The* $e_{i,j}$. If $m_{i,j}$ denotes the matrix that agrees with a matrix M on position (i, j) and is zero elsewhere, then $e_{i,i} M e_{j,j} = m_{i,j}$. As pointed out in Paragraphs 41.2.6, we see that in this way S obtains all $e_{i,i+r-(n-1)}$ for $i \in [1, n-2]$ from Y_2^* and all $e_{i,i-(n-1)}$ for $i \in [n+1, r]$ from Y_1^*. We will also exploit the fact that this puts some of the $y_k e_{i,j}$ in S as listed below.

$$(41.7) \qquad \begin{array}{ll} y_s e_{1,s} & \text{from } \ Y_1 \\ x_s e_{s,1} & \text{from } \ Y_1^* \\ y_1^{r-1} e_{r,n-1} & \text{from } \ Y_2 \\ x_1^{r-1} e_{n-1,r} & \text{from } \ Y_2^* \\ (y_1^{j-1} y_s) e_{r,(s\cdot j)\,[\bmod\, r]} & \text{from some } \ Y_i \\ (x_s x_1^{j-1}) e_{(s\cdot j)\,[\bmod\, r],r} & \text{from some } \ Y_i^* \end{array}$$

With $e_{i,j} e_{j,k} = e_{i,k}$, we get all $e_{i,j}$ (and $e_{j,i}$ by inversion) for pairs (i, j) in the cycle (41.2) not separated by the transitions $n \mapsto 1$ and

$n-1 \mapsto r$. These are all pairs in an orbit of π modulo $r-1$ in $[1,r]$, and thus in either the top or the bottom lines of (41.2). We only need to bridge one of the gaps between the two lines, and it will suffice to prove that $e_{1,r}$ is in S.

Expanding the statement that $ZZ^* = \mathbf{I}_1$ gives that

$$1 = y_1^{r-1}x_1^{r-1} + \sum_{j\in[1,r-1]}\sum_{s\in[2,n]} y_1^{j-1}y_sx_sx_1^{j-1}$$

holds in L_n, and

$$e_{1,r} = (y_1^{r-1}x_1^{r-1})e_{1,r} + \sum_{j\in[1,r-1]}\sum_{s\in[2,n]} (y_1^{j-1}y_sx_sx_1^{j-1})e_{1,r} \tag{41.8}$$

holds in $M_r(L_n)$.

We can factor the terms in (41.8) using $(y_pe_{i,j})(y_qe_{j,k}) = (y_py_q)e_{i,k}$, and using the elements of S listed in (41.7). We can rewrite (41.8) as

$$\begin{aligned} e_{1,r} &= (y_1e_{1,1})^{r-1}(e_{1,n-1})(x_1^{r-1}e_{n-1,r}) \\ &+ \sum_{j\in[1,r-1]}\sum_{s\in[2,n]}(y_1e_{1,1})^{j-1}(y_se_{1,s})(e_{s,(s\cdot j)\,[\bmod\ r]})(x_sx_1^{j-1}e_{(s\cdot j)\,[\bmod\ r],r}). \end{aligned}$$

This puts $e_{1,r}$ in S since each $(y_ke_{i,j})$ and $(x_ke_{i,j})$ that is used comes from (41.7) and is in S, and $e_{1,n-1}$ and $e_{s,(s\cdot j)\,[\bmod\ r]}$ have their two coordinates in the same row of (41.2) and so are in S.

41.2.8. *The* $y_ke_{i,j}$. We get $y_ke_{i,j}$ as $e_{i,1}(y_ke_{1,k})e_{k,j}$.

This completes the proof of Theorem 41.3, as well as Theorems 41.1 and 41.2.

42. End notes

The figure at the beginning of the Chapter gives a 3-D view of an example at the top of Page 207 in [58] that is related to one of the variants in the Thompson group family not yet covered in Chapter 6.

The comments made in the first few paragraphs of Section 41.2 imply that the L_n-modules L_n^j and L_n^k might not be isomorphic for $1 \le j < k < n$. In fact this is true and Theorem 41.1 gives an example of non-isomorphic free R-modules for a common ring R but which have isomorphic endomorphism rings. Theorem 41.1 does not give the first examples, and we will not prove that L_n^j and L_n^k are not isomorphic with j and k as stated. For other examples see Abrams 1997 [2]. For the non-isomorphism, one computes the monoid $\mathcal{V}(L_n)$ of the projective L_n-modules. See Chapter 3 of Abrams-Ara-Molina 2017 [3] for a discussion in a more general setting, and Example 3.2.6 in particular. The algebra $L_k(1,n)$ is there denoted $L_k(R_n)$.

Abrams-Ánh-Pardo 2008 [1] has the comment that $L_k(1, n)$ can be viewed as operators on an infinite dimensional k-vector space and given the operator norm when k is the complex numbers $\mathbf{C}$. The completion of $L_{\mathbf{C}}(1, n)$ is the Cuntz algebra $\mathcal{O}_n$. Representations of $V_{n,1}$ into $\mathcal{O}_n$ are found in Nekrashevych 2004 [159]. Simultaneous with [159] are representations of V into L_2 by Briget 2004 [16]. The paper Hughes 2012 [114] (with preprint from 2006) does representations in more general settings.

The notation $L_k(R_n)$ of [3] reflects that the algebra is based on paths obtained from the directed graph R_n (the n-leaved rose) consiting of 1 vertex and n labeled edges. If the labels are y_1 through y_n, then the paths are given by words over X_y. If other graphs are used, then other sets of words are obtained. This quickly splits into several topics. The operations of adding and deleting letters from the beginnings of the words leads to algebras called the Leavitt path algebras (see [3]) which can also be obtained directly from the graphs. The operation of deleting the first letters of infinite words defined by the graphs leads to modifications of the full shifts called shifts of finite type. These in turn lead to topological full groups which are variations of the $V_{n,r}$ [151]. If all of this is lifted to the analytic setting of the Cuntz algebras one obtains variations of the Cuntz algebras known as the Cuntz-Krieger algebras [52]. None of these topics will be addressed in this edition. No promises will be maded about future editions.

CHAPTER 8

Geometric properties

Chapter contents

43. Introduction

[1]This chapter is just a bit more than a stub. It contains one key calculation and one small example. Further expansion might come later.

This chapter looks at geometric aspects of the Cayley graph of the Thompson group F. To have geometry, we have to have a metric. Section 44 tells how to compute the word metric in F. Section 45 gives one observation about F.

Definition 43.1. If G is a group with generating set X, then the *Cayley graph* $\Gamma(G, X)$ of G with respect to X is a directed graph (digraph) with vertex set G and for each $g \in G$ and $x \in X$ there is an edge (g, gx) directed from g to gx. We think of the edge as labeled by x and if x^{-1} is in X, then there is a separate edge labeled x^{-1} directed from gx to g. In particular there is no requirement that X be closed under inverses. If $x = x^{-1}$, then we insist that there be two directed edges labeled x between g and gx, with one directed one way, and the other the other. If G and X are understood, the Cayley graph will simply be denoted Γ.

[1]This chapter is not complete. More material will be added in the future.

The definition should make it clear that multiplication of vertices on the left by an element of G is a graph isomorphism that fixes no vertex and no edge. Taken over all elements of G, this gives a free left action of G on Γ by graph isomorphisms.

Multiplying on the right is not an action but can be interpreted. A path in Γ from one vertex in Γ to another gives a word in $X \cup X^{-1}$ according to the following rules. If the path traverses an edge in a direction that agrees with the direction of the edge, then the label of the edge is a contribution to the word. Otherwise the inverse of the label of the edge is a contribution to the word. The word associated to the path is the string of contributions listed in the order of the edges of the path. If w is the word given by a path from some $f \in G$ to some $g \in G$, then $g = fw$. Thus multiplying a vertex in Γ on the right by a word in $X \cup X^{-1}$ drags the vertex along the path dictated by the word.

We get metric information about a Cayley graph from the following.

DEFINITION 43.2. Given a group G and generating set X, the *norm* $\|g\|$ of an element $g \in G$ is the length of the shortest word in $X \cup X^{-1}$ that represents g. Given f and g in G, the *distance* $d(f, g)$ from f to g is $\|f^{-1}g\|$.

It is immediate that d is a metric on G, and further that if every edge in $\Gamma(G, X)$ is assigned length 1, then $d(f, g)$ is the length of the shortest path in $\Gamma(G, X)$ from f to g. Thus d extends easily to a metric on $\Gamma(G, X)$, giving the Cayley graph some geometric structure. To obtain information about the structure, it helps to be able to compute the metric. For F there is an algorithm to do so. We address this next.

44. Computing length in F

We derive an effective algorithm to compute $\|f\|$ for an $f \in F$ using the generating set $A = \{x_0^{\pm 1}, x_1^{\pm 1}\}$. We describe the setting that we will use.

44.1. Forest pairs. We will use the fact from Section 11.6 that elements of F can be represented by pairs of finitary forests, and also pairs of finite forests. Let $f \in F$ be represented as $f = \Phi\Theta^{-1}$ where Φ and Θ are finitary forests. Each of Φ and Θ ends in a tail of finite trees, and it suffices to remember an initial segment of the sequence of trees in each of Φ and Θ that includes all the non-trivial trees. We will keep enough trees in each of Φ and Θ so that the number of leaves retained in Φ equals the number of leaves retained in Θ. At this point we are dealing with structures that appear in the groupoid $P^{\pm}$ from Section

24 whose properties are laid out in Corollary 24.3.1. Multiplication and reduction to a minimal pair is discussed in Section 24.1.

We will illustrate finite forest pairs as in Section 24.1 by flippling the second forest horizontally and stacking the first forest over the second, matching the leaves. Two examples are shown in (24.3) and the process of multipliying the two figures is illustrated in (24.4) and (24.5). During the process simplifications as given below are used.

44.2. A key example. We will have more than one occasion to refer to the following example.

(44.1) $g =$

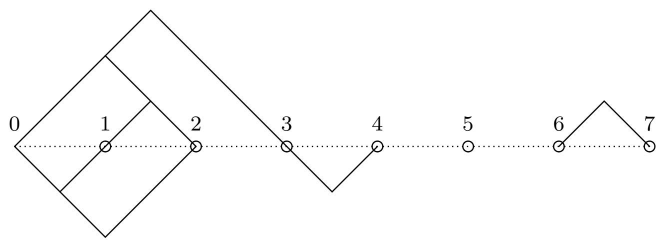

We have put in a dotted line that runs through the common leaves of the two forests and we have numbered the leaves from left to right starting from 0. The positive forest corresponds to the word $\nu_0^2\nu_1\nu_6$ in the generators of $\mathscr{F}$ and the negative forest corresponds to the word $\nu_0^2\nu_3$. We will switch to the more familiar generators x_i of F and say that the element of F given by (44.1) is $g = x_0^2x_1x_6x_3^{-1}x_0^{-2}$.

To illustrate the subtlety of the length calculation we start with a naive reduction of g to a word in the generators $A = \{x_0^{\pm 1}, x_1^{\pm 1}\}$. We have

$$\begin{aligned}
&x_0^2x_1x_0^{-5}x_1x_0^5(x_0^2x_0^{-2}x_1x_0^2)^{-1}\\
=&x_0^2x_1x_0^{-5}x_1x_0^5(x_1x_0^2)^{-1}\\
=&x_0^2x_1x_0^{-5}x_1x_0^5x_0^{-2}x_1^{-1}\\
=&x_0^2x_1x_0^{-5}x_1x_0^3x_1^{-1}
\end{aligned}$$

which has length 13. However, a slight modification of the above gives

$$\begin{aligned}
&x_0^2x_1x_0^{-5}x_1x_0^5(x_0^2x_0^{-2}x_1x_0^2)^{-1}\\
=&x_0^2x_1x_1^{-1}x_0^{-4}x_1x_0^4x_1(x_0^2x_0^{-2}x_1x_0^2)^{-1}\\
=&x_0^{-2}x_1x_0^4x_1(x_1x_0^2)^{-1}\\
=&x_0^{-2}x_1x_0^4x_1x_0^{-2}x_1^{-1}
\end{aligned} \tag{44.2}$$

which has length 11. We will soon see why 11 is the shortest possible.

44.3. Terminology. Let f be an abitrary element of F.

For calculating $\|f\|$, we adopt the following conventions and terminology. We always use an irreducible finite forest pair (Φ^+, Φ^-) to represent an $f \in F$. We picture the pair as shown in (44.1) since we want to refer to the leaves as shared by the two forests. A diagram as in (44.1) of a forest pair will be denoted by Δ.

The positive forest will be referred to as Φ^+ and the trees will be indexed from left to right from 0 and the i-th tree will be referred to as Φ^+_i. Similarly we use Φ^- and Φ^-_i for the negative forest and its trees. The leaves are usually numbered from left to right from 0 and the i-th leaf is denoted λ_i. The exception to this convention will occur when a leaf is "split" by attaching the root of a caret to that leaf. Alternate indexing will be introduced to avoid having to change the indices of the other leaves after attaching the extra caret. Details will be given when this is done.

Edges in the trees have directions, down-left and down-right in Φ^+ and up-left and up-right in Φ^-. That is, edges always are directed from node to child. A left leaf in a tree is the left child of its parent and a right leaf in a tree is the right child of its parent. Thus a leaf can be a left child in Φ^+ and a right child in Φ^-. In (44.1), λ_1 is an example of this, while λ_3 is an example of the reverse. A leaf that is a trivial tree in both Φ^+ and Φ^- is said to be doubly trivial. In (44.1), λ_5 is doubly trivial.

We need a notion of companion leaves. Two leaves in a single tree are companions if they are connected by a path in the tree with only one change of direction. In (44.1), the companion pairs are $\{\lambda_0, \lambda_2\}$, $\{\lambda_0, \lambda_3\}$ in Φ^+_0, $\{\lambda_6, \lambda_7\}$ in Φ^+_3, $\{\lambda_0, \lambda_1\}$, $\{\lambda_0, \lambda_2\}$ in Φ^-_0, and $\{\lambda_3, \lambda_4\}$ in Φ^-_1. We are actually only interested in companions of λ_0 and we use $[\lambda_0]$ to denote the set of leaves consisting of λ_0 and its companions in Δ.

We give some elementary observations. For this we need to discuss relationships in one forest, so we pick Φ^+. But the observations apply as well to Φ^-.

First if a left leaf is followed by a right leaf, then the two leaves are the leaves of an exposed caret in the forest. The extreme left leaf of a tree is a left leaf of the forest and the same is true with right replacing left twice. Second, different trees cannot have overlapping leaves in the ordering and in particular, all leaves between the extreme left leaf and the extreme right leaf of a single tree are leaves of that tree. Third, it is not possible for a leaf of a trivial tree to be immediately to the right of a left leaf or immediately to the left of a right leaf. Fourth, the extreme left leaf of a tree that is not Φ^+_0 is immediately preceeded by the extreme right leaf of another (possibly trivial) tree. A similar statement can be built about the extreme right leaf of a suitable tree.

The last observation is about $[\lambda_0]$. All contributions to $[\lambda_0]$ from considerations of Φ^+ come from Φ^+_0. If we temporarily use T to denote Φ^+_0, then we can refer to the left subtree T_0 and the right subtree T_1 of T. The only contribution to $[\lambda_0]$ from T_1 is the extreme right leaf of T. All other contributions to $[\lambda_0]$ from T come from T_0. Stated differently if the caret at the root of Φ^+_0 is removed to split Φ^+_0 into two subtrees, then the only loss to $[\lambda_0]$ is the rightmost leaf of Φ^+_0.

All comments in the previous two paragraphs apply to Φ^- after mechanical changes.

44.4. The length formula. We now can define the first of the two ingredients in the formula for an element of F. We say that a leaf λ_i is *special* if it satisfies both of the following.

(1) The leaf λ_i either doubly trivial, or is a left leaf in either Φ^+ or Φ^-.
(2) The leaf λ_i is not λ_0, and λ_{i-1} is not in $[\lambda_0]$.

We use $\#_s\Delta$ to denote the number of special leaves of Δ. The second ingredient is $\#_c\Delta$ the number of carets in Δ. We define

$$\varphi(f) = \phi(\Delta) = \#_c\Delta + 2\#_s\Delta. \tag{44.3}$$

We will prove that $\phi(f) = \|f\|$. Before we do that, let us accept the truth of the claim temporarily and look at g in (44.1). The set $[\lambda_0]$ equals $\{\lambda_0, \lambda_1, \lambda_2\lambda_3\}$. The only doubly trivial leaf is λ_5. Leaves that are left leaves in either Φ^+ or Φ^- are λ_0, λ_1, λ_3 and λ_6. However only λ_5 and λ_6 are special. There are 7 carets in the diagram in (44.1) and so $\varphi(g) = 7 + 2 \cdot 2 = 11$. Assuming our claim, the calculation in (44.2) gives as short as possible a word in $A = \{x_0^{\pm 1}, x_1^{\pm 1}\}$ that represents g.

We can now state the following.

THEOREM 44.1. *For $f \in F$, the length $\|f\|$ of f with respect to the generating set $A = \{x_0^{\pm 1}, x_1^{\pm 1}\}$ equals $\varphi(f)$ as given in (44.3).*

The following lemma dictates the outline of the proof of Theorem 44.1.

LEMMA 44.2. *Let G be a group with generating set S, and let $\ell : G \to \mathbf{N}$ be a function. Then for each $g \in G$ we will have $\ell(g) = \|g\|$ if and only if all of the following hold.*

(1) *$\ell(g) = 0$ if and only if $g = \mathbf{1}_G$.*
(2) *For all $g \in G$, we have $|\ell(gs) - \ell(g)| \leq 1$.*
(3) *For all $g \in G \setminus \{\mathbf{1}_G\}$ there is an $s \in S$ so that $\ell(gs) = \ell(g) - 1$.*

PROOF. If (1) and (2) hold, and if w is a word in the generators of length n, then a change of no more than n can occur from the value 0 of $\ell(\mathbf{1}_G)$ to the value of $\ell(w)$. If this is applied to the shortest word in the generators representing an element of G, then $\ell(g) \leq \|g\|$.

If (1) and (3) hold, take a $g \in G \setminus \{\mathbf{1}_G\}$ and with $\ell(g) = n$, find a sequence of n generators that successively multiplied on the right of g strips the ℓ value to 0. The result is $\mathbf{1}_G$, so the inverse of the sequence is a word of length n representing g. This gives $\ell(g) \geq \|g\|$.

This proves the "if" direction which is the one we actually need. The converse is straightforward and left to the reader. ☐

PROOF OF THEOREM 44.1. Let f be in F. If $f \neq \mathbf{1}$, then there is at least one caret. This verifies (1) of Lemma 44.2.

To verify (2), we will show that $|\varphi(fa) - \varphi(f)| = 1$ for all $f \in F$ and $a \in A$. We look at the four cases for fa with $a \in A$. Information gathered in this argument will be applicable in the argument for (3).

In all cases, the number of carets will change by ± 1. It will suffice to consider only those configurations where the number of carets increases by 1 and show that the number of special leaves either stays the same or decreases by 1. For if ga has one fewer caret than g, then $g = (ga)a^{-1}$ has one more caret than ga and the analysis of the arrangements where the number of carets increases will give us the desired outcome.

($a = x_0$): Since we assume that fa has one more caret than f, the tree Φ_0^- must be trivial. This "splits" λ_0 into two leaves that we will call $\lambda_{0.0}$ and $\lambda_{0.1}$ to avoid changing the indices of the other leaves in going from f to fa. The leaves $\lambda_{0.0}$ and $\lambda_{0.1}$ are not special and status of all other leaves remains the same. So the number of special leaves of fa is the same as that of f.

($a = x_1$): Now to have that fa has one more caret than f, Φ_1^- must be trivial. If u is the rightmost leaf of Φ_0^- and $v = \lambda_i$ is immediately to the right of u, then λ_i splits into $\lambda_{i.0}$ and $\lambda_{i.1}$. The leaf $\lambda_{i.0}$ is not special because u is in $[\lambda_0]$. The leaf $\lambda_{1.1}$ is not special since it is a right leaf in Φ^+ and a trivial tree in Φ^-. The status of all other leaves remains the same. Again the number of special leaves of fa is the same as that of f.

($a = x_0^{-1}$): In this cae the trees Φ_0^- and Φ_1^- are joined even if they are both trivial. Our assumption that the number of carets goes up says that the new caret does not cancel an exposed caret of Φ^+. From the observations at the end of Section 44.3, the only possible change to $[\lambda_0]$ is the addition of the rightmost leaf v of Φ_1^-. If the leaf immediately to the right of v is special in f, it is not in fa. Otherwise there are no other changes to the status of the leaves. This is the desired outcome.

Note that the addition of a caret by multiplying by x_0^{-1} results in a decrease in the number of special leaves in only one very restricted configuration. The rightmost leaf v of Φ_1^- must not be in $[\lambda_0]$ for f and the leaf immediately to the right of v must be special in f.

($a = x_1^{-1}$): In this case the trees Φ_1^- and Φ_2^- are joined even if they are both trivial, and we are assuming that the new caret does not cancel an exposed caret of Φ^+. Let p and q be the leftmost and rightmost leaves, respectively, of Φ_1^-, and let r and s be the leftmost and rightmost leaves, respectively, of Φ_2^-. The equalities $p = q$ and/or $r = s$ are possible if depending on the triviality or non-triviality of Φ_1^- and Φ_2^-. We must have $q \neq r$ because $\Phi_1^- \neq \Phi_2^-$. Whether any of p, q, r or s are in $[\lambda_0]$ depends on the structure of Φ^+ since none if these leaves are in Φ_0^-. Thus this aspect of these leaves is the same in f and fa. The leaf immediately to the left of p is definitely in $[\lambda_0]$ in f, and thus p cannot be special for f and for fa. The leaf q is special in fa if and only if it is special in f. If $r \neq s$, then each is special in fa if and only if it is special in f. If $r = s$ and is not special in f, then it is not special in fa. If $r = s$ is special in f because it is doubly trivial for f and q is not in $[\lambda_0]$, then it is not doubly trivial for fa, not a left leaf in either forest, and thus not special for fa. No other leaf can change its status from f to fa. So as in the case $a = x_0^{-1}$, the number of special leaves of fa is either the same as for f or it is fewer by 1.

As with the case $a = x_0^{-1}$ there is a very restricted configuration in which the addition of a caret by multiplying by x_1^{-1} results in a decrease in the number of special leaves. Here the leaf r immediately to the right of Φ_1^- must be doubly trivial (and thus not the last leaf of Δ) and the rightmost leaf q of Φ_1^- must not be in $[\lambda_0]$.

To verify (3), we want to lower $\varphi(f)$ by multiplying on the right by a generator. From the argument for (2) we know that if we lower the number of carets, then there are only two specific ways this can fail. We keep this in mind when we consider how Φ^- can be configured.

Case 1: The tree Φ_1^- is not trivial. In this case the only configuration in which multiplying on the right by x_1 fails to lower $\varphi(f)$ is the following. Multiplying by x_1 on the right splits the non-trivial Φ_1^- into $\Phi_{1.0}^-$ on the left and $\Phi_{1.1}^-$ on the right. The tree $\Phi_{1.1}^-$ must be trivial and its only leaf r must also be a trivial tree in Φ^+ and not be the last leaf in Δ. Also the rightmost leaf q of $\Phi_{1.0}^-$ must not be in $[\lambda_0]$. It follows that r is not in $[\lambda_0]$. If t is the leaf immediately to the right of r, then it follows that t is special for f no matter how it sits in the trees of Φ^+ and Φ^-. So in this configuration we can lower $\varphi(f)$ by multiplying by

x_0^{-1} and unite the trees Φ_0^- and Φ_1^- putting r into $[\lambda_0]$ and making t not special. Thus if Φ_1^- is not trivial, $\varphi(f)$ can always be lowered.

Case 2: The tree Φ_1^- is trivial and the tree Φ_0^- is not trivial. In this case the only configuration in which multiplying on the right by x_0 fails to lower $\varphi(f)$ is the following. Multiplying on the right by x_0 splits the non-trivial Φ_0^- into $\Phi_{0.0}^-$ on the left and $\Phi_{0.1}^-$ on the right so that the rightmost leaf v of $\Phi_{0.1}^-$ is not in $[\lambda_0]$ for fx_0^{-1}, and the leaf t immediately to the right of v is not special in f but becomes special in fx_0^{-1}. Since Φ_1^- is trivial, t is its only leaf. For fx_0^{-1} to have t be special, t is either doubly trivial or a left leaf in Φ^+. So t is not in $[\lambda_0]$ for f. There must be a leaf u immediately to the right of t since t becomes special for fx_0^{-1}. If u is a left leaf or trivial in Φ^+, then it is special for f and will become non-special if we multiply f on the right by x_0^{-1} and join ϕ_0^- and the trivial Φ_1^-. This would lower $\varphi(f)$ so we assume that u is not special. But this means u is trivial in Φ^- and a right leaf in Φ^+. This means t cannot be trivial in Φ^+ and we have already noted it cannot be a right leaf. So t is a left leaf, and t and u are the leaves of an exposed caret of Φ^+. With both t and u trivial in Φ^-, we can cancel the exposed caret by multiplying f on the right by x_1^{-1}. From Part I, this must lower $\varphi(f)$. Thus if Φ_1^- is trivial and Φ_0^- is not, we can always lower $\varphi(f)$.

Case 3: Both Φ_0^- and Φ_1^- are trivial. If λ_2 is special for f, then λ_1 is not in $[\lambda_0]$ and multiplying on the right by x_0^{-1} will put λ_1 in $[\lambda_0]$ and λ_2 will not be special for fx_0^{-1}. There are two ways λ_2 can be not special for f. If λ_1 is in $[\lambda_0]$, then λ_0 and λ_1 are the leaves of an exposed caret in Φ^+, and since Φ_0^- and Φ_1^- are trivial, fx_0^{-1} will have fewer carets than f and from Part I will have $\varphi(fx_0^{-1}) < \varphi(f)$. If λ_1 is not in $[\lambda_0]$, then for λ_2 to be not special for f, it must be trivial in Φ^- and a right leaf in Φ^+. But now λ_1 cannot be trivial in Φ^+ and since it is not in $[\lambda_0]$ it is not a right leaf in Φ^+. So it is a left leaf in Φ^+ and λ_1 and λ_2 are the leaves of an exposed caret in Φ^+. Both leaves are in trivial trees in Φ^- and so fx_1^{-1} will have fewer carets than f and from Part I will have $\varphi(fx_0^{-1}) < \varphi(f)$.

This completes the proof of Theorem 44.1. □

45. Dead ends

If G is a group with generating set X, then the Cayley graph $\Gamma = \Gamma(G, X)$ might have dead ends.

DEFINITION 45.1. With G, X and Γ as above, a vertex v of Γ with $\|v\| = n$ is a *dead end* of *depth* k if every edge path in Γ from v to a vertex u with $\|u\| > \|v\|$ includes a vertex w with $\|w\| = n - k$. If $k > 1$, then v can be referred to as a *deep pocket*.

Dead ends are discussed in Cleary-Taback 2004 [49] which also proves that there are no deep pockets in F. We will show that the example of (44.1) is a dead end. The following pictures do the trick.

element		diagram	$\#_c$	$\#_s$	φ
g	$=$	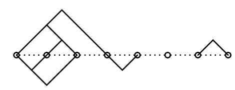	7	2	11
gx_0	$=$	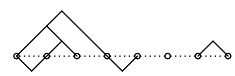	6	2	10
gx_1	$=$	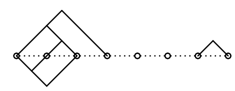	6	2	10
gx_0^{-1}	$=$	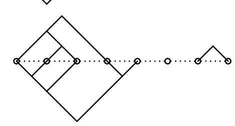	8	1	10
gx_1^{-1}	$=$	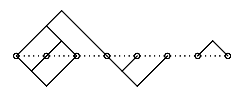	8	1	10

For gx_0 and gx_1, the number of carets goes down and the number of special leaves remains the same. For the remaining two, the number of carets goes up and one of the special leaves λ_5 becomes non-special.

45.1. Remarks and an example. Dead ends are local maxima, and the proof of Theorem 44.1 shows that the only local minimum is the identity.

The proof of Theorem 44.1 gives an algorithm for computing from an irreducible forest pair a minimal length word in $A = \{x_0^{\pm 1}, x_1^{\pm 1}\}$ representing the same element. We give an example.

Let $g = x_0^2 x_3^2 x_4^{-1} x_3^{-1} x_1^{-1} x_0^{-1}$ whose forest pair looks like the following.

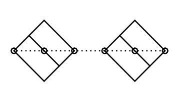

We have $\varphi(g) = 10$ because leaf 4 (second from the right) is special. No other leaf is special.

In the following we reduce the φ value by one on each line by multiplying g on the right by a generator. We start with $g_0 = g$ and multiply g_0 on the right successively by a sequence of generators to create a sequence of g_i that end at the identity.

g_i	diagram	$\#_c$	$\#_s$	φ
g_0	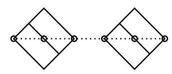	8	1	10
$g_1 = g_0 x_0$	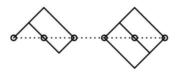	7	1	9
$g_2 = g_1 x_1$	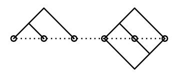	6	1	8
$g_3 = g_2 x_0^{-1}$	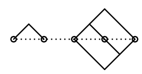	5	1	7
$g_4 = g_3 x_0^{-1}$	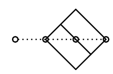	4	1	6
$g_5 = g_4 x_1$	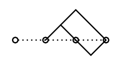	3	1	5
$g_6 = g_5 x_0^{-1}$	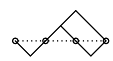	4	0	4
$g_7 = g_6 x_1$	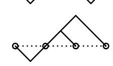	3	0	3
$g_8 = g_7 x_0$	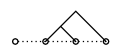	2	0	2
$g_9 = g_9 x_1^{-1}$	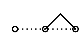	1	0	1
$\mathbf{1} = g_9 x_1^{-1}$		0	0	0

Thus

$$\begin{aligned} g =&(x_0 x_1 x_0^{-2} x_1 x_0^{-1} x_1 x_0 x_1^{-2})^{-1} \\ =&x_1^2 x_0^{-1} x_1^{-1} x_0 x_1^{-1} x_0^2 x_1^{-1} x_0^{-1} \\ =&x_1^2 x_2^{-1} x_1^{-1} x_0^2 x_1^{-1} x_0^{-1} \\ =&x_0^2 x_3^2 x_4^{-1} x_3^{-1} x_1^{-1} x_0^{-1} \end{aligned}$$

where the second line gives the expression of g as a word of length 10 in $A = \{x_0^{\pm 1}, x_1^{\pm 1}\}$ and the next two lines verify that this word is equivalent to the original expression of g as a word in the infinite generating set.

The point of the above example is to show two things. One is that locating the special leaves can take some care. The other is that the sequence of reductions can take some unexpected turns. This is illustrated by the change that brings the φ value from 5 to 4 which involves raising the number of carets to lower the number of special leaves. Note that keeping the initial leaf in the change of φ value from 7 to 6 is important to keep the subscripts of the generators accurate. It should be noted that the reductions $8 \to 7$ and $7 \to 6$, each accomplished by multiplying on the right by x_0^{-1}, are justified by noting that the expansion of x_3^{-1} in $A = \{x_0^{\pm 1}, x_1^{\pm 1}\}$ is as $x_3^{-1} = x_0^{-2} x_1^{-1} x_0^2$.

46. End notes

The existence of an algorithm to compute the word length in F was first shown in Fordham's 1995 thesis [71] and was published in [72] in 2003. Our presentation is Fordham's algorithm as modified by Guba 2004 [97], and Belk-Brown 2005 [9].

Geometric observations can be made about F without a precise calculation of the word metric. Section 2 of Burillo 1999 [38] gives an estimate for the word metric which bounded above and below from a computation from the normal form. Interestingly, the argument is based on data extracted from the graph of the element as an element of $PL_+(I)$. Applications are given in [38].

The length formula we derive in Section 44.4 is from [97]. The setting there is that of semigroup diagrams which are essentially dual to forest pairs. The basic building block is a 2-cell whose boundary is divided into three intervals, either one on top and two on the bottom, or two on top and one on the bottom. The first is dual to a caret in the positive forest, and the second is dual to a caret in the negative forest as shown below.

The definitions leading to the length formula (44.3), and the proof we give of Theorem 44.1 are bascially translations from semigroup diagrams to forest pairs of the material leading to and the proof of Theorem 4 of [97].

CHAPTER 9

Appendix

Chapter contents

47. Appendix A: Confluence, distinguished representatives, and Newman's lemma

The motivating concept here is that of an equivalence relation $\sim$ on a set X and a method of improving elements in each equivalence class so that each class contains a unique "maximally improved" element that can be taken to be the distinguished representative of that class. The model for this is a binary relation $\to$ on X that generates $\sim$ in which $x \to y$ is interpreted as "y is an improvement of x."

There is a sequence of three results each having weaker hypotheses than the previous that guarantees the existence of the distinguished element. The most general of the three is known as Newman's lemma

[161]. The fact that sometimes one of the less general results applies to a situation tends not to be of significance, and we will always just say that we are applying Newman's lemma.

An assumption common to all is that the relation $\to$ on X is *terminating*, or *Nötherian*, or *well founded*, meaning that no infinite sequence (x_i), $i \in \mathbf{N}$, exists so that for all $i \in \mathbf{N}$, we have $x_i \to x_{i+1}$. In particular, for all $x \in X$, $x \to x$ is false, and $\to$ is "maximally non-reflexive."

To discuss the remaining assumptions, let $\stackrel{+}{\to}$ be the transitive closure of $\to$, let $\stackrel{=}{\to}$ be the reflexive closure of $\to$, and let $\stackrel{*}{\to}$ be the reflexive and transitive closure of $\to$. As usual, turning a symbol around represents the inverse relation, so $x \leftarrow y$ holds if and only if $y \to x$ holds, and similiarly for the various closures.

In the following we will write $\leftarrow\to$ implies $\to\leftarrow$ to mean that if there are x, y and z with $x \leftarrow y \to z$, then there is a w with $x \to w \leftarrow z$ and similarly for the various closures. The three items in Definition 47.1 below give the hypotheses of the three results.

Definition 47.1. With $\to$ a binary relation on a set X, we define the following.

(1) We say that $\to$ is *confluent* if $x \sim y$ implies that there exists z with $x \stackrel{*}{\to} z \stackrel{*}{\leftarrow} y$.
(2) We say that $\to$ *satisfies the diamond condition* if $\leftarrow\to$ implies $\stackrel{=}{\to}\stackrel{=}{\leftarrow}$.
(3) We say that $\to$ is *locally confluent* if $\leftarrow\to$ implies $\stackrel{*}{\to}\stackrel{*}{\leftarrow}$.

To appreciate the difference between local confluence and the diamond condition, the reader can try to apply the argument in the proof of Lemma 47.2 below that assumes the diamond condition but while assuming local confluence instead. The usual picture below of the diamond condition gives rise to its name.

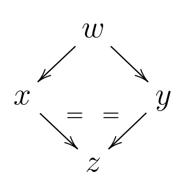

We roll the three reuslts into one and obtain the combined result as a corrolary of the following lemma. We will refer to the lemma and its corollary together as "Newman's lemma." We say that $x \in S \subseteq X$ is *irreducible* in S if there is no $y \in S$ with $x \to y$.

Lemma 47.2. *Let $\to$ be a binary relation on a set X that generates the equivalence relation $\sim$. If $\to$ is terminating, then every non-empty*

$S \subseteq X$ has an element irreducible in S. If $\to$ is confluent, then each equivalence class contains no more than one irreducible. If $\to$ satisfies the diamond condition, then it is also confluent. If $\to$ is terminating and locally confluent, then it is confluent.

Proof. If $\to$ is terminating and $S \subseteq X$ is not empty, then $\to_S$, the restriction of $\to$ to S is a terminating relation on S. So assuming $x \in S = X$ suffices. If X has no irreducible element, then every sequence $x = x_0 \to x_1 \to \cdots \to x_k$ can be extended to have one more entry and a violation of terminating is shown to exist.

If $\to$ is confluent, and x and y are two irreducibles in an equivalence class S, then there is a $z \in S$ with $x \xrightarrow{*} z \xleftarrow{*} y$. But the irreducibility of x and y implies $x = z = y$.

If $\to$ satisfies the diamond condition and $x \sim y$, then there is a sequence $(x_i \mid 0 \leq i \leq n)$ in X with $x = x_0$, $y = x_n$, and where for each $0 \leq i < n$, one of $x_i \to x_{i+1}$ or $x_i \leftarrow x_{i+1}$ holds. There are two ways to deliver the standard argument. The "graph paper" argument plots the progress from x to y on graph paper with $\leftarrow$ written vertically and $\to$ written horizontally. From our assumption,

$$\begin{array}{ccc} x_{i+1} & \longrightarrow & x_{i+2} \\ \big\downarrow & & \\ x_i & & \end{array} \qquad \text{can be replaced by} \qquad \begin{array}{ccc} & & x_{i+2} \\ & & \big\downarrow{=} \\ x_i & \xrightarrow{=} & x_{i+1} \end{array}$$

repeatedly until all arrows written horizontally come before all arrows written vertically. The "complexity" argument points out that the cardinality of the finite set

$$\{(i, j) \mid 0 \leq i < j < n,\ x_i \leftarrow x_{i+1} \text{ and } x_j \to x_{j+1}\}$$

goes down each time an instance of $\leftarrow\!\!\to$ implies $\xrightarrow{=}\xleftarrow{=}$ is used.

As preparation for the last argument, we note that the proof assuming the diamond condition can be adapted with minimal change to show that if $\xleftarrow{*}\xrightarrow{*}$ always implies $\xrightarrow{*}\xleftarrow{*}$, then $\to$ is confluent.

We now assume local confluence and terminating. The proof will be an example of Nötherian induction. Specifically if P is a predicate on elements of X with one free variable and if for all $x \in X$ we have that $P(x)$ holds whenever $P(y)$ holds for every y for which $x \xrightarrow{+} y$, then $P(x)$ holds for every $x \in X$. For if not, then $S = \{x \mid \neg P(x)\}$ is not empty and contains an element x irreducible in S. But then $P(y)$ holds for every $y \in X$ with $x \xrightarrow{+} y$, and our assumption on P then implies $P(x)$ holds.

Let $P(x)$ be the statement that for all y and z with $y \overset{*}{\leftarrow} x \overset{*}{\rightarrow} z$ there is a w with $y \overset{*}{\rightarrow} w \overset{*}{\leftarrow} z$ and assume an x for which $P(y)$ holds for every y with $x \overset{+}{\rightarrow} y$. Now if $y \overset{*}{\leftarrow} x \overset{*}{\rightarrow} z$, then such a w exists if either $y = x$ or $z = x$. So we assume there are y_0 and z_0 so that $y \overset{*}{\leftarrow} y_0 \leftarrow x \rightarrow z_0 \overset{*}{\rightarrow} z$. The rest of the argument refers to the following diagram where we have to justify the existence of all entries not yet mentioned.

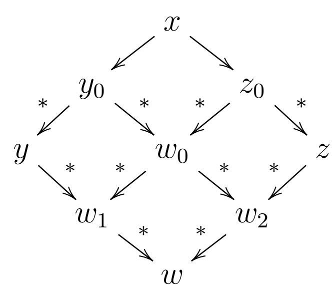

The element w_0 exists because of the assumption of local confluence. The elements w_1, w_2, and w exist because, respectively, $P(y_0)$, $P(z_0)$, and $P(w_0)$ all hold by our assumptions about x. □

Corollary 47.2.1. *If a terminating binary relation $\rightarrow$ on a set X is confluent, or satsifies the diamond condition, or is locally confluent, then every equivalence class under $\sim$ contains a unique irreducible element. Further given an element in an equivalence class, a chain of random applications of $\rightarrow$ starting with the given element must ultimately lead to the unique irreducible element of the class.*

Proof. The assumption that $\rightarrow$ is terminating implies that given any x in an equivalence class S, a random chain of applications of $\rightarrow$ starting at x must ultimately end in an element that is irreducible in S. Uniqueness of the irreducible follows from Lemma 47.2. □

48. Appendix B: Simplicial and cubical complexes

We give a quick review of those parts of simplicial and cubical complexes that we need. Almost all statements are easy exercises, and we give references for some that are not.

48.1. Simplices. For convenience, we work in $\mathbf{R}^\infty$, the vector space of infinite sequences of real numbers. In $\mathbf{R}^\infty$, the point e_i, $i \in \mathbf{N}$, has all coordinates 0, except the i-th which has the value 1.

Given $p \neq q$ in $\mathbf{R}^\infty$, the line segment $[p, q]$ from p to q is the set $\{tq + (1-t)p \mid t \in [0, 1]\}$. The points p and q are the *endpoints* of $[p, q]$ and the rest of the points in $[p, q]$ are *interior points*.

A set $S \subseteq \mathbf{R}^\infty$ is *convex* if for all p and q in S, the line segment $[p, q]$ is in S. The empty set is convex. The intersection of a collection of convex sets in $\mathbf{R}^\infty$ is convex. Given a set $A \subseteq \mathbf{R}^\infty$, the *convex hull* of A is the smallest convex set in $\mathbf{R}^\infty$ containing A and is the intersection of all convex sets in $\mathbf{R}^\infty$ that contain A.

If $V = \{v_0, v_1, \ldots, v_n\}$ is a set of $(n+1)$ points in $\mathbf{R}^\infty$, then V is *affinely independent* if $n = 0$ or if the n-tuple $(v_1 - v_0, v_2 - v_0, \ldots, v_n - v_0)$ is linearly independent. Affine independence does not depend on which element is taken as v_0. Affine independence is inherited by subsets. In $\mathbf{R}^\infty$, any set of the e_i is affinely independent.

If V is an affinely independent set of $(n+1)$ points in $\mathbf{R}^\infty$, then the *n-simplex* σ (which we also denote $[V]$) with *vertex set* V is the convex hull of V, and is referred to as the *simplex spanned by* V. The *dimension* of σ is n. If $n = 0$, the simplex is a single point. The *standard n-simplex* Δ_n is the convex hull of $\{e_0, e_1, \ldots, e_n\}$ in $\mathbf{R}^\infty$. The standard n-simplex is the set

$$\begin{aligned}\Delta_n &= \left\{ (\alpha_0, \ldots, \alpha_n) \mid \sum_{i=0}^{n} \alpha_i = 1,\ \alpha_i \geq 0,\ 0 \leq i \leq n \right\} \\ &= \left\{ \sum_{i=0}^{n} \alpha_i e_i \mid \sum_{i=0}^{n} \alpha_i = 1,\ \alpha_i \geq 0,\ 0 \leq i \leq n \right\}.\end{aligned}$$

The vertices of an n-simplex are determined by the simplex in that they are the only points in the simplex in the interior of no segment in the simplex. If σ is an n-simplex spanned by $V = \{v_0, \ldots, v_n\}$ and $f : \{e_0, \ldots, e_n\} \to V$ is such than $f(e_i) = v_i$, then an extension of f to Δ_n defined by

$$f\left(\sum_{i=0}^{n} \alpha_i e_i \right) = \sum_{i=0}^{n} \alpha_i f(e_i) = \sum_{i=0}^{n} \alpha_i v_i$$

is a homeomorphism from Δ_n onto σ that takes vertices to vertices. Thus every $x \in \sigma$ has a unique expression as $x = \sum \alpha_i v_i$ with each α_i non-negative and $\sum \alpha_i = 1$. The α_i for which this holds are the *barycentric coordinates* for that $x \in \sigma$, and each α_i is the barycentric coordinate of x at v_i.

If τ is an m-simplex spanned by $W = \{w_0, \ldots, w_m\}$, and $g : V \to W$ is any function, then g extends to the *simplicial* map $g : \sigma \to \tau$ defined by

$$g\left(\sum_{i=0}^{n} \alpha_i v_i \right) = \sum_{i=0}^{n} \alpha_i g(v_i) = \sum_{i=0}^{n} \alpha_i w_i$$

where the α_i are the barycentric coordinates of the elements of σ. The extension is surjective if and only if $g : V \to W$ is surjective, and the extension is injective if and only if $g : V \to W$ is injective.

If σ is an n-simplex spanned by V and $W \subseteq V$ has m points, then the m-simplex τ spanned by W is an m-dimensional *face* of σ and also an $(n-m)$-codimensional face of σ. If $m < n$, then τ is a *proper face* of σ. If now $Z = V \setminus W$, and ρ is the $((m-n)-1)$-dimensional face of σ spanned by Z, then τ and ρ are *opposite faces* of σ. With Z defined this way, the points in τ are those points in σ whose barycentric coordinates at the elements of Z are zero. The *interior* of a simplex is defined as the complement of the union of its proper faces.

Two disjoint subsets A and B of $\mathbf{R}^\infty$ are *joinable* if the set of line segments $[A, B] = \{[p, q] \mid (p, q) \in A \times B\}$ has the following property: every two segments in $[A, B]$ are either disjoint, or equal, or have their only point of intersection an endpoint of both segments. If subsets A and B are joinable then the *join* $A * B$ of A and B is the union

$$A * B := \bigcup_{(a,b)\in A\times B} [a, b].$$

The join of two convex sets is convex. In the special case that A and B are joinable and B is the singleton $\{v\}$, then we write $A * v$ for $A * \{v\}$ and refer to $A * v$ as the *cone* over A with *cone point* v. The point v is a strong deformation retraction of $A * v$ and so the cone over A is always contractible. Two opposite faces τ and ρ of a simplex σ are always joinable and $\sigma = \tau * \rho$. In this case, the sum of the dimensions of τ and ρ is one less than the dimension of σ. If τ has codimension 1, and v is the vertex of σ not in τ, then $\sigma = \tau * v$.

Remark 48.1. The structure given to a simplex as a join is important. If τ and ρ are opposite faces of the simplex σ, then sending points (b, c, t) in $\tau \times \rho \times (0, 1)$ to $(1-t)b + tc$ in $\tau * \rho = \sigma$ is a homeomorphism to $\sigma \setminus (\tau \cup \rho)$. If δ is a face of σ that has non-empty intersection with both τ and ρ, then $[\delta \cap \tau]$ and $[\delta \cap \rho]$ are opposite faces of δ. If $i : \delta \to \sigma$, $j : [\delta \cap \tau] \to \tau$ and $k : [\delta \cap \rho] \to \rho$ are the natural inclusions, then the following commutes,

$$\begin{array}{ccc} [\delta \cap \tau] \times [\delta \cap \rho] \times (0,1) & \longrightarrow & \tau \times \rho \times (0,1) \\ \downarrow & & \downarrow \\ \delta \setminus ([\delta \cap \tau] \cup [\delta \cap \rho]) & \longrightarrow & \sigma \setminus (\tau \cup \rho) \end{array} \tag{48.1}$$

where the horizontal arrows are built from i, j and k and the vertical arrows are restrictions as needed of $(b, c, t) \mapsto (1 - t)b + tc$.

48.2. Simplicial complexes. A *simplicial complex* is a pair (X, Σ) where X is a topological space and Σ is a set of simplices in X so that X is the union of the simplices in Σ, the topology on X is generated by Σ in that $U \subseteq X$ is open in X if for all $\sigma \in \Sigma$, $U \cap \sigma$ is open in σ, each face of each simplex in Σ must be in Σ, and the intersection of any two simplices in Σ is empty or a face of each. Notation is frequently abused, and the complex is referred to by X if the set of simplexes Σ is understood. The space X is the *underlying space* of the simplicial complex, and the full complex (X, Σ) is often referred to as a *triangulation* of X. Every $x \in X$ is in the interior of a unique simplex in Σ. This will be commented on in Section 48.3.

If (X_1, Σ_1) and (X_2, Σ_2) are simplicial complexes then $f : X_1 \to X_2$ is simplicial if for each $\sigma_1 \in \Sigma_1$ there is a simplex $\sigma_2 \in \Sigma_2$ so that $f|_{\sigma_1}$ agrees with some simplicial map from σ_1 to σ_2. A simplicial map of complexes is continuous.

If (X, Σ) is a simplicial complex, then a *subcomplex* of (X, Σ) is a simplicial complex (Y, Γ) where Γ is a subset of Σ. If the vertex set of (X, Σ) is V, and $V' \subseteq V$, then the subcomplex (Y, Γ) of (X, Σ) *spanned by* V' has Γ consist of those simplices in Σ whose vertices are in V' and Y the union of the simplices in Γ.

48.3. Abstract simplicial complexes. Simplices are determined by their vertices, so the structure of a simplicial complex is determined the set of vertices once it is known which vertices belong to a simplex. Based on this an *abstract simplicial complex* is defined as a pair (V, S) where V is a set and S is a collection of finite, non-empty subsets of V subject to the requirement that V is the union of the elements of S and S is closed under passing to non-empty subsets. The vertices of a simplicial complex together with the subsets that are vertices of the simplices in the complex form an abstract simplicial complex.

If (V, S) is an abstract simplicial complex with V countable (the only situation we will need), then we can identify V with a set of the e_i in $\mathbf{R}^\infty$ and take Σ to be the set of simplexes that are spanned by the subsets in S. Now the union of these simplexes in Σ is a subspace X of $\mathbf{R}^\infty$ (using, say, the uniform metric) that is a candidate for the underlying space of a simplicial complex whose corresponding abstract simplicial complex is (V, S).

However if the simplices are not locally finite, then the topology on X as a subspace of $\mathbf{R}^\infty$ might not be the topology determined by the simplices in Σ. This is not a problem since we will be interested only in homotopy invariant properties of the complexes, and the main result of [61] says that the identity map on X is a homotopy equivalence between the two topologies. Therefore we accept (X, Σ) as the topological realization of (V, S). This varies from the usual definition of the topological realization, but not in an essential way. One way to look at the topological realization is to identify the points in the realization with those functions from V to $[0, 1]$ having the following restrictions. Given a function $\alpha : V \to [0, 1]$, we write α_v for its value at $v \in V$, we call $\{v \in V \mid \alpha_v > 0\}$ the support of α, we require that the support of α be finite, and we require that $\sum_{v \in V} \alpha_v = 1$. To put α in $\mathbf{R}^\infty$ with our view of V as a set of the e_i, we identify α with $x = \sum_{v \in V} \alpha_v v$. The α_v are the *barycentric coordinates* of x. The unique simplex that has x in its interior is the simplex spanned by the support of α.

If (V_1, S_1) and (V_2, S_2) are two abstract simplicial complexes and $f : V_1 \to V_2$ is a function, then we say that f is simplicial if for every $s \in S_1$ we have $f(s) \in S_2$. Such a function induces a simplicial map between the topological realizations.

If (V, S) is an abstract simplicial complex, then $|(V, S)|$ will denote its topological realization. Topological properties of $|(V, S)|$ will be referred to as properties of (V, S).

48.4. Posets and complexes. Abstract simplicial complexes can be derived from partially ordered sets. A *partially ordered set* (or poset) $(P, \leq)$ is a set P with a binary relation $\leq$ that is reflexive, transitive and anti-symmetric in that $x \leq y$ and $y \leq x$ implies $x = y$. The partial order is a *total order* if for every x and y in P either $x \leq y$ or $y \leq x$. We write $x < y$ to mean $x \leq y$ and $x \neq y$. A function f between posets is *order preserving* if for every $x \leq y$ in the domain we have $f(x) \leq f(y)$ in the range. A *chain* in a poset $(P, \leq)$ is a subset $C \subseteq P$ that is a total order under the restriction of $\leq$ to C. We will use interval notation such as $[x, z] = \{y \in P \mid x \leq y \leq z\}$, $[x, z) = \{y \in P \mid x \leq y < z\}$, and so forth.

If P is a poset, then the set of finite, non-empty chains in P forms an abstract simplicial complex that we will also refer to by P. Careful wording will avoid confusion. If $f : P \to Q$ is an order preserving function between two posets, then it is also a simplicial map of the associated simplicial complexes.

We will use $|P|$ to denote the topological realization of the complex of finite chains in P, and properties of the resulting topological space then become properties of interest. Let P and Q be two posets and f and g be two order preserving functions from P to Q. If for all $x \in P$, we have $f(x) \leq g(x)$, then f and g induce homotopic maps on the topological realizations. To discuss this, we model $|P| \times I$.

More generally if P and Q are two posets, then the product poset structure on $P \times Q$ is defined to have $(p_1, q_1) \le (p_2, q_2)$ if $p_1 \le p_2$ and $q_1 \le q_2$ both hold. Lemma 8.9 of [68, Ch. 2] says $|P \times Q|$ is homeomorphic to $|P| \times |Q|$.

Now we take $Q = \{0, 1\}$ with $0 < 1$ so $|Q| = I$. If $\sigma = \{v_0 < v_1 < \cdots < v_n\}$ is an n-simplex in P, then each

$$\sigma_i = \{(v_0, 0) < \cdots < (v_i, 0) < (v_i, 1) < (v_{i+1}, 1) < \cdots (v_n, 1)\} \quad (48.2)$$

is a simplex in $P \times \{0, 1\}$ and every simplex of $P \times Q$ is of the form in (48.2) or a face thereof. Now the assumption that $f(x) \le g(x)$ for all $x \in P$ guarantees that if we take $P \times \{0\}$ to Q by f and $P \times \{1\}$ to Q by g, then the map on $P \times \{0, 1\}$ is simplicial and extends to a continuous simplicial map.

If P contains a greatest element v, then the function $g : P \to \{v\}$ has $x \le g(x)$ for all $x \in P$, and the identity on $|P|$ is homotopic to the constant map to v, making P contractible. In fact P is the cone over $|\{w \in P \mid w < v\}|$ with cone point v.

48.5. Barycentric subdivision. If σ is a simplex with emphasis on its status as a topological space, and Σ is the collection of simplices consisting of σ and all of its faces, then (σ, Σ) is a simplicial complex with σ as underlying space. If σ is spanned by $V = \{v_0, v_1, \ldots, v_n]\}$, then

$$\bar{\sigma} = \frac{1}{n+1} \sum_{i=0}^{n} v_i$$

is the *barycenter* of σ.

Let F be the poset of faces of σ under containment. If $\tau_0 \subset \tau_1 \subset \cdots \subset \tau_k$ is a chain of faces then $\{\bar{\tau}_0, \ldots, \bar{\tau}_k\}$ is affinely independent and spans a k-simplex $[\bar{\tau}_0, \ldots, \bar{\tau}_k]$. The collection of all such simplexes is the barycentric subdivision of σ, often denoted $\operatorname{Sd} \sigma$ or σ'. The barycentric subdivision of a simplicial complex K is the set of simplices in the barycentric subdivisions of all the simplices in K. If K' is the barycentric subdivision of K, then from Lemma 6.2 of [68, Ch. 2], the underlying topological spaces of K and K' are the same.

If $(P, \le)$ is a poset and K the abstract simplicial complex of finite chains under $\le$ of elements of P, then the abstract barycentric subdivision K' of K is made of chains of chains. Specifically, K' consists of finite, non-empty chains ordered under containment, of finite, non-empty chains ordered under $\le$ of elements of P. Since the topological realizations of K and K' are homeomorphic, their topological invariants are the same.

48.6. Covering maps and covering spaces. If (V, S) is an abstract simplicial comples, and a group G acts on V (on the right, say) by simplicial maps, then we will say that the action is *simplicial*. If this action is free, then no point is fixed by any $1 \neq g \in G$. In particular no barycenter of any simplex in X is fixed by any $1 \neq g \in G$, and it follows that for any $\sigma \in S$ and any $1 \neq g \in G$, we have $\sigma g \neq \sigma$.

Assume that an action as above is free, that X is connected, that $x = \alpha : V \to [0,1]$ is a point in X, and that $\sigma \in S$ is the unique simplex that contains x in the interior of its span (i.e., σ is the support of x). Let ϵ be less than $\frac{1}{2}\min\{\alpha_v \mid v \in \sigma\}$, and let $g \neq 1$ be in G. Then the support of xg includes a $v \in V \setminus \sigma$ whose barycentric coordinate on v is greater than ϵ, and so the distance between x and xg is greater than ϵ. Thus every point $x \in X$ has an open neighborhood U so that for all $1 \neq g \in G$, $U \cap Ug = \emptyset$. From Proposition 1.40 of [104], we have that the map $p : X \to X/G$ with $p(x) = xG$ is a regular (or normal) covering map, G is the group of covering transformations (or deck transformations) of p, and G is isomorphic to $\pi_1(X/G)/p_*(\pi_1(X))$.

48.7. Stars, links and cones. If (X, Σ) is a simplicial complex and v is a vertex in Σ, then the *star* of v in the complex is the set of simplices that have v as a vertex together with their faces. Note that the set of simplices having v as a vertex is not closed under passing to proper faces if this set contains anything more than just v.

The *link* of v in the complex is the set of simplices in the star of v that do not have v as a vertex. If the star of v contains only v, then the link of v is empty. For each simplex σ having v as a vertex with τ the face of σ opposite to v, the link of v in σ is τ and all the faces of τ. It is thus legitimate to refer to the star of v as the cone over the link of v with cone point v. This holds even if the link of v is empty if we declare that $\emptyset * v = \{v\}$. Stars of vertices are contractible. If $St(v)$ is the star of v, and $Lk(v)$ is the link of v in a simplicial complex, then $St(v) \setminus \{v\}$ has the structure of $Lk(v) \times [0, 1)$, and $Lk(v)$ is a strong deformation retract of $St(v) \setminus \{v\}$.

At times a simplicial complex is built from another complex by adding cones. If L_α, $\alpha \in A$, is a family of subcomplexes of (X, σ) and v_α, $\alpha \in A$ are vertices where $\cup L_\alpha$, $\alpha \in A$, and $\{v_\alpha \mid \alpha \in A\}$ are joinable, then the result of *coning off* all the L_α is the complex

$$X \cup \bigcup_{\alpha \in A} L_\alpha * v_\alpha$$

where the simplices in the resulting complex are those of X, together with all $\sigma * \{v\}$, and $\{v\}$ itself. Note that for $\alpha \neq \beta$ in A, we have $(L_\alpha * v_\alpha) \cap (L_\beta * v_\beta) = L_\alpha \cap L_\beta$.

48.8. Simple Morse theory. The following is simplified from [11].

If X is a simplicial complex spanned by V, then a function $f : V \to \mathbf{R}$ is a *Morse function* if its image is closed and discrete, and no 1-simplex has f the same on both of its vertices. Equivalently, f is a Morse function if no two vertices in an n-simplex have the same value under f. Our Morse functions will have image in the integers.

We will filter X by subcomplexes X_i where X_i is spanned by $\{v \in V \mid f(v) \leq i\}$. We investigate how X_i changes as i increases.

We assume a Morse function f on the vertices of X that takes values only in the integers. Because f is a Morse function, a simplex σ in X_i that is not a simplex of X_{i-1} has a unique vertex v_σ with $f(v) = i$, and the addition of σ to X_{i-1} is the addition of the cone at v_σ over the face of σ opposite to v_σ. Thus for v with $f(v) = i$, the set of simplices in X_i with vertex v forms the cone over the link of v in X_i. For v with $f(v) = i$, we let $Lk_\downarrow(v, X)$ denote the link of v in X_i and refer to it as the *descending link* in X. We have arrived at our main tool.

LEMMA 48.2. *Let X be a simplicial complex spanned by V and let $f : V \to \mathbf{Z}$ be a Morse function. If X_i is the subcomplex of X spanned by $\{v \in V \mid f(v) \leq i\}$, then X_i is obtained from X_{i-1} by coning off all the $Lk_\downarrow(v, X)$ for each v with $f(v) = i$.*

48.9. Cubical complexes. Cubical complexes arise naturally, and cubical complexes can be realized as simplicial complexes.

The poset $\{0 < 1\}$ has $|\{0 < 1\}| = I$ which is also the 1-dimensional cube. Thus $\{0 < 1\}^n$ with the product poset structure from Section 48.4 will be our model for the n-cube. Elements of $\{0 < 1\}^n$ are sequences on $n = \{0, \ldots, n-1\}$ with values in $\{0, 1\}$, and two such sequences f and g have $f \leq g$ if $f_i \leq g_i$ for all $i \in n$. Any poset isomorphic to $\{0 < 1\}^n$ with this order will be an n-cube. In particular, the set of subsets of a set of n elements, ordered under inclusion, has the structure of an n-cube.

The smallest element of $\{0 < 1\}^n$ is the sequence $\overline{0}$ all of whose values are 0, and the largest element is $\overline{1}$ all of whose values are 1. A maximal chain in $\{0 < 1\}^n$ has $n + 1$ sequences and is determined by a permutation of the indices giving the locations, in order, where the value changes from 0 to 1. There are thus $n!$ maximal chains in $\{0 < 1\}^n$ and $n!$ simplices of dimension n in our model of the n-cube.

A *face* of $C = \{0 < 1\}^n$ will be a closed interval $[m, M]$ with $m \leq M$ in C. Let D be the set on which m and M agree, and let d the restriction of m or M to D. Let E be the complement of D in $n = \{0, \ldots, n-1\}$ and let $k = |E|$. Now f in C is in $[m, M]$ if and only if f agrees with d on D. So $[m, M]$ is isomorphic as a poset to the set of functions $\{0 < 1\}^E$ with E totally ordered as a subset of $n = \{0, \ldots, n-1\}$. Thus $[m, M]$ is a k-cube. We also arrive at a face of C by choosing some $D \subseteq n$ and some $d : D \to \{0 < 1\}$. With E the complement of D in n, with m equal to d on D and 0 on E, and with M equal to d on D and 1 on E, we recover the face $[m, M]$ as those elements in C agreeing with d on D. The face is proper if $D \neq \emptyset$.

A cubical complex for us will always arive as a structure that can be imposed on a simplicial complex that is derived from a poset. So we define a *cubical complex* to be a poset $(P, \leq)$ and a set K of subsets of P called cubes whose union is P, so that each $C \in K$ has the structure of a cube in the poset structure, and so that the intersection two elements of K is empty or a face (as a cube) of each. Note that every 1-cube is a 1-simplex, but in a cubical complex as just given, not every 1-simplex will be a 1-cube. For example $\{(0, 0), (1, 1)\}$ is a 1-simplex in $\{0 < 1\}^2$ but not a 1-cube. Two vertices in a cubical complex are *adjacent* if they are the endpoints of a 1-cube.

If v is a vertex in a cubical complex (P, K) as above, then the link of v with respect to the cubical structure is the abstract simplicial complex (not a cubical complex) whose vertices are the vertices adjacent to v in the cubical complex, and where a collection of these vertices forms a simplex if they lie in a single cube. The link of a vertex of a single 3-dimensional cube is a 2-simplex and all of its faces.

The link of a cubical complex and the link of the corresponding simplicial complex are topologically the same. This is because a cube and its corresponding simplicial complex have the same underlying space. We sketch the details.

With $I = [0, 1]$, let I^n model the n-cube. Let $n = \{0, \ldots, n-1\}$ index the coordinates and let $\mathbf{x} = (x_i)$ be an n-tuple in I^n. There is a permutation π on n so that $x_{\pi(0)} \leq x_{\pi(1)} \leq \cdots \leq x_{\pi(n-1)}$. Let $\alpha_0 = x_{\pi(0)}$, let $\alpha_{i+1} = x_{\pi(i+1)} - x_{\pi(i)}$, and let $\alpha_n = 1 - x_{\pi(n-1)}$. We have $\sum_0^n \alpha_i = 1$. Recursively, let $v_0 = \mathbf{1}$, the tuple with all values equal to 1, and let $v_{i+1} = v_i - e_{\pi(i)}$. That is we successively change each 1 to 0 in v_0 in the order given by π so that $v_n = \mathbf{0}$. The $\pi(j)$-th component of $\sum_0^n \alpha_i v_i$ is $\sum_0^j \alpha_i = x_{\pi(j)}$. So $\sum_0^n \alpha_i v_i = \mathbf{x}$ and $\mathbf{x}$ in the simplex spanned by the v_i. Thus I^n is contained in the union of the $n!$ different possible simplices, and the reverse containment follows because I^n is convex. That two simplices intersect in a face of each follows because the intersection of two chains is a subchain of each.

48.10. Curvature and cubical complexes. A property of the link in a cubical complex has a central role in this next topic. A simplicial complex is *flag* or a *flag complex* if a subset S of the vertices are the vertices of a simplex if and only if every pair of elements in S are the vertices of a 1-simplex. For example the link of a vertex of a 3-cube is a flag complex, but the link of a vertex in a cubical complex consisting of the proper faces of a 3-cube is not.

We need two results about CAT(0) spaces, but we omit details of the definition. See Chapter II of [26] for definitions. We can also use Theorem 48.3 below as a definition. To give some idea of the concept, Definition 1.2 of Chapter II of [26] mentions that for a real κ, a metric space has curvature no greater than κ if it is locally a CAT(κ) space.

The following is Theorem B.8 of [135].

THEOREM 48.3. *A simply connected cubical complex is* CAT(0) *if and only if its vertex links are flag complexes.*

The following is Corollary 1.5 of Chapter II in [26].

THEOREM 48.4. *For $\kappa \leq 0$, any* CAT(κ) *space is contractible.*

49. Appendix C: Ultrafilters

The bit of material we need on ultrafilters might be slightly less than standard. There are also many references that discuss ultrafilters with various differences of terminology. However, there is a commonality of concept which we do not depart from. Among many others [62, pp. X.2–X1.1], [75, pp. V.1–V.5], and [77, pp. 6.7–6.8] are possible references.

If P is a poset under $\leq$ with a minimum element 0, then a *filterbase* $\mathcal{F}$ on P is a subset of $P \setminus \{0\}$ that is directed downward in that for all p and q in $\mathcal{F}$, there is an $r \in \mathcal{F}$ with $r \leq p$ and $r \leq q$. That is, every pair of elements in $\mathcal{F}$ has a lower bound in $\mathcal{F}$.

A filterbase $\mathcal{F}$ that is upwardly closed ($p \leq q$ in P and $p \in \mathcal{F}$ implies $q \in \mathcal{F}$), is a *filter*. Clearly, the filter generated by a filterbase $\mathcal{F}$ is the set $\{p \in P \mid \exists q \in \mathcal{F}, q \leq p\}$.

If $\mathcal{F}$ and $\mathcal{G}$ are filterbases on P, then we say that $\mathcal{G}$ is *subordinate* to $\mathcal{F}$ if for every $p \in \mathcal{F}$ there is a $q \in \mathcal{G}$ so that $q \leq p$. That is, a representative of $\mathcal{G}$ can be slipped below every representative of $\mathcal{F}$. Every filter is a filterbase, so the term applies to filters as well. Note that for filters, $\mathcal{G}$ subordinate to $\mathcal{F}$ and the upwardly closed requirement puts every element of $\mathcal{F}$ in $\mathcal{G}$ and we get $\mathcal{F} \subseteq \mathcal{G}$. The converse is trivial since $p \leq p$ for all $p \in \mathcal{F}$. So for filters, $\mathcal{G}$ is subordinate to $\mathcal{F}$ if and only if $\mathcal{F} \subseteq \mathcal{G}$.

A *maximal filterbase* is a filterbase $\mathcal{M}$ so that any filterbase $\mathcal{F}$ subordinate to $\mathcal{M}$ has $\mathcal{M}$ subordinate to $\mathcal{F}$. A maximal filter is called an *ultrafilter*. Note that if $\mathcal{M}$ is an ultrafilter and $\mathcal{F}$ is a filter subordinate to $\mathcal{M}$, then $\mathcal{F} = \mathcal{M}$. A standard Zorn's lemma argument shows that every filter is contained in an ultrafilter.

An example of a filter is a *principal filter* which takes the form $\{q \in P \mid p \leq q\}$ given some specific $p \in P$. If p is not minimal in $P \setminus \{0\}$, then this filter is not an ultrafilter. However, in another setting what is usually called principal is an ultrafilter.

Let (X, T) be a topological space with basis B (regular or ordinary) for the topology T. Examples of a useful poset P ordered by inclusion could be

(1) P is the set of all subsets of X,
(2) P is the topology T,
(3) P is the basis B.

In all of these, given $p \in X$ the set $\{A \in P \mid p \in A\}$ is an ultrafilter on P called the principal ultrafilter at p. Note that in (2) and (3) an ultrafilter on P is a filter on X but in general is not an ultrafilter on X and does not generate an ultrafilter on X. It is standard that if $\mathcal{F}$ is an ultrafilter on the set of all subsets of X, then for any $\emptyset \neq A \subsetneq X$, either A or $X \setminus A$ is in $\mathcal{F}$. In (2) and (3) the filter on X generated by an ultrafilter on P will consist only of sets with non-empty interior. In general, a topological space X will have many sets A where neither A nor $X \setminus A$ have non-empty interior.

50. Appendix D: Some elementary number theory

50.1. The absolute basics. See the first chapter of any book on elementary number theory.

We assume familiarity with the consequence of the Euclidean algorithm that for integers a and b not both zero, $\gcd(a, b)$ equals the least positive linear sum with integer coefficients of a and b. We also assume the unique prime factorization of integers which we put in the following form for non-negative integers. For an integer $a \geq 0$, the prime characteristic function χ_a of a is the function from the primes to the non-negative integers giving the powers of the primes in the prime factorization of a. This is almost a logarithm.

The support of χ_a (those p with $\chi_a(p) \neq 0$) is bounded. Taking all operations coordinate by coordinate, $\chi_a + \chi_b = \chi_{ab}$, $\chi_a \leq \chi_b$ is equivalent to $a|b$, with $g = \gcd(a, b)$, $\chi_g = \inf(\chi_a, \chi_b)$, with $l = \text{lcm}(a, b)$,

$\chi_l = \sup(\chi_a, \chi_b)$, disjoint supports corresponds to relatively prime, and all this makes it clear that $gl = ab$.

50.2. The group of units modulo s.

LEMMA 50.1. *For a positive integer s, the set of integers modulo s that are prime to s, is exatly the set of units in $\mathbf{Z}/s\mathbf{Z}$ and this set forms a group under multiplication modulo s.*

PROOF. If $\gcd(a, s) = 1$, then $ax + sy = 1$ has solutions, and x is an inverse to a modulo s. The converse reverses the argument.

The set of units in a ring always forms a group, but in this case it is fun to note that if $ax + sy = 1 = bz + sw$, then $(ax + sy)(bz + sw) = 1$ and the left side has one term involving ab and the other three terms are multiples of s. □

50.3. The Chinese remainder theorem.

PROPOSITION 50.2. *Let s_i, $1 \le i \le n$, be pairwise relatively prime with $S = \prod_{i=1}^{n} s_i$. Then there is a simultaneous integer solution x to the n equations $x \equiv a_i \pmod{s_i}$ and the set of such solutions is $x + S\mathbf{Z}$.*

PROOF. We start by proving the following claim which is a rewording of the proposition for the case $n = 2$.

CLAIM 2. *For integers a, b, c, d with c and d relatively prime, there is an integer x so that $(a + c\mathbf{Z}) \cap (b + d\mathbf{Z}) = x + (cd)\mathbf{Z}$.*

We have $c\mathbf{Z} \cap d\mathbf{Z} = (cd)\mathbf{Z}$ since cd is the least common multiple of c and d. There are integer solutions to $xc - yd = 1$, and in fact for every k there are solutions to $xc - yd = k$. So there is some q in $c\mathbf{Z} \cap (b + d\mathbf{Z})$. Shifing both $c\mathbf{Z}$ and $b + d\mathbf{Z}$ down by q brings the two sequences back to $c\mathbf{Z}$ and $d\mathbf{Z}$ again, so $c\mathbf{Z} \cap (b + d\mathbf{Z}) = (q + (cd)\mathbf{Z})$. This says that every integer shift between $c\mathbf{Z}$ and $d\mathbf{Z}$ is of the form $(x + (cd)\mathbf{Z})$ and the claim follows.

Now we inductively assume that the solution set to the first $n - 1$ equations of the proposition is of the form $x' + S'\mathbf{Z}$ where $S' = \prod_{i=1}^{n-1} s_i$. By the claim, $(x' + S'\mathbf{Z}) \cap (a_n + s_n\mathbf{Z})$ is of the form $(x + S' s_n \mathbf{Z}) = (x + S\mathbf{Z})$ for some x. □

50.4. On solutions of modular equations. The following appears in Section 38.8 as Lemma 38.28. We use the usual notation (a, b) for $\gcd(a, b)$.

LEMMA 50.3. *(I) If $a|m$ and for some x, we have $ax \equiv r \pmod{m}$, then $a|r$.*

(II) The equation $kx \equiv r \pmod{m}$ has exactly (k, m) solutions modulo m if $(k, m)|r$.

Proof. (I) With $m = aq$ and $ax - r = pm = pqa$, we have $r = a(x - pq)$.

(II) We have $k = (k, m)a$, $m = (k, m)b$ and $r = (k, m)c$. If $kx \equiv r \pmod{m}$, then $kx - r = mp$ for some p or

$$\begin{aligned}(k, m)ax - (k, m)c &= (k, m)bp,\\ ax - c &= bp,\\ ax &\equiv c \pmod{b}.\end{aligned}$$

Now $(a, b) = 1$ so $ap + bq = 1$ for some p and q. There are b problems like $ax \equiv c \pmod{b}$ as c varies over the residues modulo b, and each has a solution $(ap + bq)c = c$ or $apc \equiv c \pmod{b}$ among the residues modulo b. So there is a unique solution for each problem. The x solving $ax \equiv c \pmod{b}$ solves $kx \equiv r \pmod{m}$.

If $ky \equiv r \pmod{m}$, then $ay \equiv c \pmod{b}$. So $a(x-y) \equiv 0 \pmod{b}$ and $x-y \equiv 0 \pmod{b}$ by the uniqueness just shown. So $x-y = jb$ and there are exactly (k, m) multiples of $b = m/(k, m)$ that are different residues modulo m. □

The following is used in Section 41. I learned its proof from Alexander Borisov.

Lemma 50.4. *Let a, b, s be integers with both a and b not zero and with $g = \gcd(a, s) = \gcd(b, s)$. Then some integer x prime to s is a solution to $ax \equiv b \pmod{s}$.*

Proof. If $g = 1$, then a and b are in the group of units of $\mathbf{Z}/s\mathbf{Z}$ and this group of units then contains the x required.

If $g > 1$, then $a' = a/g$, $b' = b/g$ and $s' = s/g$ are pairwise relatively prime and there is an x' prime to s' with $a'x' \equiv b' \pmod{s'}$. So $ax' \equiv b \pmod{s}$, but we don't know x' is prime to s. However any $x \in x' + s'\mathbf{Z}$ will have $a'x = b' \pmod{s'}$ and $ax = b \pmod{s}$. Let A be those primes that divide s but not s' and let B be those primes that divide s'. The set of primes that divide s is the disjoint union of A and B.

By the Chinese remainder theorem, there is an x equivalent to x' modulo s' and equivalent to 1 modulo each prime in A. No prime in B can divide x since such a prime divides s' which is prime to x'. This x is a solution as required. □

Note that Dirichlet's theorem that every arithmetic sequence contains infinitely many primes can be used in the proof of Lemma 50.4, but the lemma is much more elementary than that.

Bibliography

[1] G. Abrams, P. N. Ánh, and E. Pardo. "Isomorphisms between Leavitt algebras and their matrix rings". In: *J. Reine Angew. Math.* 624 (2008), pp. 103–132 (cit. on pp. 242, 243, 257, 259, 265).

[2] Gene Abrams. "Non-induced isomorphisms of matrix rings". In: *Israel J. Math.* 99 (1997), pp. 343–347 (cit. on p. 264).

[3] Gene Abrams, Pere Ara, and Mercedes Siles Molina. *Leavitt path algebras*. Vol. 2191. Lecture Notes in Mathematics. Springer, London, 2017, pp. xiii+287 (cit. on pp. 264, 265).

[4] Roger C. Alperin. "Notes: $PSL_2(Z) = Z_2 * Z_3$". In: *Amer. Math. Monthly* 100.4 (1993), pp. 385–386 (cit. on p. 80).

[5] Tuna Altinel and Alexey Muranov. "Interprétation de l'arithmétique dans certains groupes de permutations affines par morceaux d'un intervalle". In: *J. Inst. Math. Jussieu* 8.4 (2009), pp. 623–652 (cit. on pp. 199, 201, 202, 209, 213).

[6] Matthias Aschenbrenner, Anatole Khélif, Eudes Naziazeno, and Thomas Scanlon. "The logical complexity of finitely generated commutative rings". In: *Int. Math. Res. Not. IMRN* 1 (2020), pp. 112–166 (cit. on p. 215).

[7] G. Baumslag, E. Dyer, and A. Heller. "The topology of discrete groups". In: *J. Pure Appl. Algebra* 16.1 (1980), pp. 1–47 (cit. on p. 190).

[8] James Belk, Luna Elliott, and Francesco Matucci. "A short proof of Rubin's theorem". In: *Israel J. Math.* 267.1 (2025), pp. 157–169 (cit. on p. 204).

[9] James M. Belk and Kenneth S. Brown. "Forest diagrams for elements of Thompson's group F". In: *Internat. J. Algebra Comput.* 15.5-6 (2005), pp. 815–850 (cit. on pp. 224, 277).

[10] George M. Bergman. "Right orderable groups that are not locally indicable". In: *Pacific J. Math.* 147.2 (1991), pp. 243–248 (cit. on p. 22).

[11] Mladen Bestvina and Noel Brady. "Morse theory and finiteness properties of groups". In: *Invent. Math.* 129.3 (1997), pp. 445–470 (cit. on pp. 142, 143, 289).

[12] Robert Bieri. *Homological dimension of discrete groups*. Queen Mary College Mathematics Notes. Queen Mary College, Mathematics Department, London, 1976, pp. iii+190 (cit. on p. 142).

[13] Robert Bieri and Ralph Strebel. *On groups of PL homeomorphisms of the real line*. preprint, Math. Sem. der Univ. Frankfurt, Frankfurt am Main. 1985 (cit. on pp. 14, 105, 106, 209).

[14] Robert Bieri and Ralph Strebel. *On groups of PL homeomorphisms of the real line*. ArXiv preprint: https://arxiv.org/pdf/1411.2868.pdf. 2016 (cit. on p. 14).

[15] Robert Bieri and Ralph Strebel. *On groups of PL-homeomorphisms of the real line*. Vol. 215. Mathematical Surveys and Monographs. American Mathematical Society, Providence, RI, 2016, pp. xvii+174 (cit. on pp. 105, 204, 209).

[16] Jean-Camille Birget. "The groups of Richard Thompson and complexity". In: *Internat. J. Algebra Comput.* 14.5-6 (2004). International Conference on Semigroups and Groups in honor of the 65th birthday of Prof. John Rhodes, pp. 569–626 (cit. on pp. 138, 247, 265).

[17] Jean-Camille Birget. "Monoid generalizations of the Richard Thompson groups". In: *J. Pure Appl. Algebra* 213.2 (2009), pp. 264–278 (cit. on pp. 137, 138, 253).

[18] Collin Bleak. "A geometric classification of some solvable groups of homeomorphisms". In: *J. Lond. Math. Soc. (2)* 78.2 (2008), pp. 352–372 (cit. on p. 93).

[19] Collin Bleak. "An algebraic classification of some solvable groups of homeomorphisms". In: *J. Algebra* 319.4 (2008), pp. 1368–1397 (cit. on p. 93).

[20] Collin Bleak. "A minimal non-solvable group of homeomorphisms". In: *Groups Geom. Dyn.* 3.1 (2009), pp. 1–37 (cit. on p. 93).

[21] Collin Bleak, Matthew G. Brin, and Justin Tatch Moore. "Complexity among the finitely generated subgroups of Thompson's group". In: *J. Comb. Algebra* 5.1 (2021), pp. 1–58 (cit. on p. 216).

[22] Collin Bleak and Martyn Quick. "The infinite simple group V of Richard J. Thompson: presentations by permutations". In: *Groups Geom. Dyn.* 11.4 (2017), pp. 1401–1436 (cit. on pp. 176, 180).

[23] Marianna C. Bonanome, Margaret H. Dean, and Judith Putnam Dean. *A sampling of remarkable groups*. Compact Textbooks in Mathematics. Thompson's, self-similar, Lamplighter,

and Baumslag-Solitar. Birkhäuser/Springer, Cham, 2018, pp. xii+188 (cit. on p. 92).

[24] William W. Boone. “Between logic and group theory”. In: *Proceedings of the Second International Conference on the Theory of Groups (Australian Nat. Univ., Canberra, 1973).* 1974, 90–102. Lecture Notes in Math., Vol. 372 (cit. on p. 99).

[25] William W. Boone and Graham Higman. “An algebraic characterization of groups with soluble word problem”. In: *J. Austral. Math. Soc.* 18 (1974). Collection of articles dedicated to the memory of Hanna Neumann, IX, pp. 41–53 (cit. on pp. 99, 100).

[26] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature.* Vol. 319. Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]. Springer-Verlag, Berlin, 1999, pp. xxii+643 (cit. on p. 291).

[27] Matthew G. Brin. “The chameleon groups of Richard J. Thompson: automorphisms and dynamics”. In: *Inst. Hautes Études Sci. Publ. Math.* 84 (1996), pp. 5–33 (cit. on p. 106).

[28] Matthew G. Brin. “The ubiquity of Thompson’s group F in groups of piecewise linear homeomorphisms of the unit interval”. In: *J. London Math. Soc. (2)* 60.2 (1999), pp. 449–460 (cit. on p. 93).

[29] Matthew G. Brin. “Elementary amenable subgroups of R. Thompson’s group F”. In: *Internat. J. Algebra Comput.* 15.4 (2005), pp. 619–642 (cit. on p. 216).

[30] Matthew G. Brin. “The algebra of strand splitting. II. A presentation for the braid group on one strand”. In: *Internat. J. Algebra Comput.* 16.1 (2006), pp. 203–219 (cit. on p. 176).

[31] Matthew G. Brin and Craig C. Squier. “Groups of piecewise linear homeomorphisms of the real line”. In: *Invent. Math.* 79 (1985), pp. 485–498 (cit. on pp. 104, 105).

[32] Kenneth S. Brown. *Cohomology of groups.* Vol. 87. Graduate Texts in Mathematics. Springer-Verlag, New York-Berlin, 1982, pp. x+306 (cit. on pp. 143, 190).

[33] Kenneth S. Brown. “Presentations for groups acting on simply-connected complexes”. In: *J. Pure Appl. Algebra* 32.1 (1984), pp. 1–10 (cit. on p. 176).

[34] Kenneth S. Brown. “Finiteness properties of groups”. In: *Journal of Pure and Applied Algebra* 44 (1987), pp. 45–75 (cit. on pp. 9, 92, 105, 106, 142, 180, 181, 191, 223, 240).

[35] Kenneth S. Brown. “The geometry of finitely presented infinite simple groups”. In: *Algorithms and Classification in Combinatorial Group Theory*. Ed. by G. Baumslag and III C. F. Miller. Vol. 23. MSRI Publications. Springer-Verlag, New York, 1992, pp. 121–136 (cit. on pp. 106, 168, 180).

[36] Kenneth S. Brown and Ross Geoghegan. “An infinite-dimensional torsion-free FP_∞ group”. In: *Invent. Math.* 77 (1984), pp. 367–381 (cit. on pp. 3, 92, 103, 104, 106, 139, 181, 190).

[37] Marc Burger and Shahar Mozes. “Lattices in product of trees”. In: *Inst. Hautes Études Sci. Publ. Math.* 92 (2000), 151–194 (2001) (cit. on p. 100).

[38] José Burillo. “Quasi-isometrically embedded subgroups of Thompson’s group F”. In: *J. Algebra* 212.1 (1999), pp. 65–78 (cit. on pp. 139, 277).

[39] José Burillo. *Introduction to Thompson’s group F*. Pdf format notes: https://web.mat.upc.edu/pep.burillo/F%20book.pdf (cit. on p. 92).

[40] José Burillo, Brita Nucinkis, and Lawrence Reeves. “An irrational-slope Thompson’s group”. In: *Publ. Mat.* 65.2 (2021), pp. 809–839 (cit. on p. 139).

[41] José Burillo, Brita Nucinkis, and Lawrence Reeves. “Irrational-slope versions of Thompson’s groups T and V”. In: *Proc. Edinb. Math. Soc. (2)* 65.1 (2022), pp. 244–262 (cit. on p. 139).

[42] Stanley Burris and H. P. Sankappanavar. *A course in universal algebra*. New York: Springer-Verlag, 1981, pp. xvi+276 (cit. on p. 125).

[43] J. W. Cannon, W. J. Floyd, and W. R. Parry. “Introductory notes on Richard Thompson’s groups”. In: *Enseign. Math. (2)* 42.3-4 (1996), pp. 215–256 (cit. on pp. 9, 14, 32, 90, 92, 98, 104, 107, 176, 180).

[44] Georg Cantor. *Contributions to the founding of the theory of transfinite numbers*. Translated, and provided with an introduction and notes, by Philip E. B. Jourdain. Dover Publications, Inc., New York, N. Y., 1952, pp. ix+211 (cit. on p. 130).

[45] Pierre-Emmanuel Caprace and Bertrand Rémy. “Simplicity and superrigidity of twin building lattices”. In: *Invent. Math.* 176.1 (2009), pp. 169–221 (cit. on p. 100).

[46] Pierre-Alain Cherix, Michael Cowling, Paul Jolissaint, Pierre Julg, and Alain Valette. *Groups with the Haagerup property*. Vol. 197. Progress in Mathematics. Gromov’s a-T-menability. Birkhäuser Verlag, Basel, 2001, pp. viii+126 (cit. on p. 176).

[47] Sean Cleary. "Groups of piecewise-linear homeomorphisms with irrational slopes". In: *Rocky Mountain J. Math.* 25.3 (1995), pp. 935–955 (cit. on p. 106).

[48] Sean Cleary. "Regular subdivision in $\mathbf{Z}[\frac{1+\sqrt{5}}{2}]$". In: *Illinois J. Math.* 44.3 (2000), pp. 453–464 (cit. on p. 139).

[49] Sean Cleary and Jennifer Taback. "Combinatorial properties of Thompson's group F". In: *Trans. Amer. Math. Soc.* 356.7 (2004), pp. 2825–2849 (cit. on p. 275).

[50] A. H. Clifford and G. B. Preston. *The algebraic theory of semigroups. Vol. I.* Mathematical Surveys, No. 7. Providence, R.I.: American Mathematical Society, 1961, pp. xv+224 (cit. on pp. 62, 137, 138, 155).

[51] Joachim Cuntz. "Simple C^*-algebras generated by isometries". In: *Comm. Math. Phys.* 57.2 (1977), pp. 173–185 (cit. on pp. 138, 249).

[52] Joachim Cuntz and Wolfgang Krieger. "A class of C^*-algebras and topological Markov chains". In: *Invent. Math.* 56.3 (1980), pp. 251–268 (cit. on p. 265).

[53] Patrick Dehornoy. "Free distributive groupoids". In: *J. Pure Appl. Algebra* 61.2 (1989), pp. 123–146 (cit. on pp. 107, 134).

[54] Patrick Dehornoy. "Structural monoids associated to equational varieties". In: *Proc. Amer. Math. Soc.* 117.2 (1993), pp. 293–304 (cit. on pp. 107, 134).

[55] Patrick Dehornoy. "The structure group for the associativity identity". In: *J. Pure Appl. Algebra* 111.1-3 (1996), pp. 59–82 (cit. on pp. 107, 134).

[56] Patrick Dehornoy. *Braids and self-distributivity.* Vol. 192. Progress in Mathematics. Birkhäuser Verlag, Basel, 2000, pp. xx+623 (cit. on p. 107).

[57] R. K. Dennis and L. N. Vaserstein. "Commutators in linear groups". In: *K-Theory* 2.6 (1989), pp. 761–767 (cit. on p. 28).

[58] Warren Dicks and Conchita Martínez-Pérez. "Isomorphisms of Brin-Higman-Thompson groups". In: *Israel J. Math.* 199.1 (2014), pp. 189–218 (cit. on pp. 243, 259, 264).

[59] Leonard Eugene Dickson. *Introduction to the theory of numbers.* Dover, New York, 1957, pp. viii+183 (cit. on p. 16).

[60] Jacques Dixmier. "Traces sur les C^*-algèbres. II". In: *Bull. Sci. Math. (2)* 88 (1964), pp. 39–57 (cit. on p. 138).

[61] C. H. Dowker. "Topology of metric complexes". In: *Amer. J. Math.* 74 (1952), pp. 555–577 (cit. on p. 285).

[62] James Dugundji. *Topology.* Allyn and Bacon, Inc., Boston, Mass., 1966, pp. xvi+447 (cit. on p. 291).

[63] J. Dydak. "A simple proof that pointed, connected FANR-spaces are regular fundamental retracts of ANR's". In: *Bull. Acad. Polon. Sci. Sér. Sci. Math. Astronom. Phys.* 25 (1977), pp. 55–62 (cit. on pp. 59, 102).

[64] Jerzy Dydak and Harold M. Hastings. "Homotopy idempotents on two-dimensional complexes split". In: *Proceedings of the International Conference on Geometric Topology (Warsaw, 1978).* PWN, Warsaw, 1980, pp. 127–133 (cit. on p. 103).

[65] Jerzy Dydak and Jack Segal. *Shape theory.* Vol. 688. Lecture Notes in Mathematics. An introduction. Berlin: Springer, 1978, pp. vi+150 (cit. on p. 132).

[66] Jerzy Dydak and Jack Segal. *Shape theory.* Vol. 688. Lecture Notes in Mathematics. An introduction. Springer, Berlin, 1978, pp. vi+150 (cit. on p. 139).

[67] Samuel Eilenberg and Tudor Ganea. "On the Lusternik-Schnirelmann category of abstract groups". In: *Ann. of Math. (2)* 65 (1957), pp. 517–518 (cit. on p. 190).

[68] Samuel Eilenberg and Norman Steenrod. *Foundations of algebraic topology.* Princeton University Press, Princeton, N.J., 1952, pp. xv+328 (cit. on p. 287).

[69] David B. A. Epstein, James W. Cannon, Derek F. Holt, Silvio V. F. Levy, Michael S. Paterson, and William P. Thurston. *Word processing in groups.* Jones and Bartlett Publishers, Boston, MA, 1992, pp. xii+330 (cit. on p. 2).

[70] Daniel S. Farley. "Finiteness and CAT(0) properties of diagram groups". In: *Topology* 42.5 (2003), pp. 1065–1082 (cit. on pp. 176, 191).

[71] S. Blake Fordham. "Minimal length elements of Thompson's group F". PhD thesis. Brigham Young University, 1995 (cit. on pp. 106, 139, 277).

[72] S. Blake Fordham. "Minimal length elements of Thompson's group F". In: *Geom. Dedicata* 99 (2003), pp. 179–220 (cit. on pp. 106, 277).

[73] Peter Freyd and Alex Heller. "Splitting homotopy idempotents. II". In: *J. Pure Appl. Algebra* 89.1-2 (1993), pp. 93–106 (cit. on pp. 9, 103, 104).

[74] R. Fricke and F. Klein. *Vorlesungen über die Theorie der Automorphen Funktionen.* Vol. 1. Leipzig: Teubner, 1897 (cit. on p. 79).

[75] Steven A. Gaal. *Point set topology.* Dover Publications, Mineola, New York, 2009, pp. xi+317 (cit. on p. 291).

[76] P. Gabriel and M. Zisman. *Calculus of fractions and homotopy theory.* Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 35. Springer-Verlag New York, Inc., New York, 1967, pp. x+168 (cit. on pp. 155, 157).

[77] Michael C. Gemignani. *Elementary topology.* second. Dover Publications, Inc., New York, 1990, pp. xii+270 (cit. on p. 291).

[78] Anthony Genevois. "An introduction to diagram groups". In: *J. Comb. Algebra* 9.1-2 (2025), pp. 31–119 (cit. on pp. 93, 94).

[79] Ross Geoghegan. "The problem of pointed versus unpointed domination in shape theory". In: *Topology Proc.* 3.1 (1978), 95–107 (1979) (cit. on p. 102).

[80] Ross Geoghegan. *Topological methods in group theory.* Vol. 243. Graduate Texts in Mathematics. Springer, New York, 2008, pp. xiv+473 (cit. on pp. 142–144, 181, 189).

[81] Étienne Ghys. "Groups acting on the circle". In: *Enseign. Math. (2)* 47.3-4 (2001), pp. 329–407 (cit. on p. 21).

[82] Étienne Ghys and Vlad Sergiescu. "Sur un groupe remarquable de difféomorphismes du cercle". In: *Comment. Math. Helv.* 62.2 (1987), pp. 185–239 (cit. on pp. 77, 105).

[83] Michèle Giraudet, A. M. W. Glass, and J. K. Truss. "Undecidability of automorphism groups". In: *Math. Z.* 240.3 (2002), pp. 611–620 (cit. on p. 215).

[84] A. M. W. Glass. "The word problem for lattice ordered groups". In: *Proc. Edinburgh Math. Soc. (2)* 19 (1974), pp. 217–219 (cit. on p. 138).

[85] A. M. W. Glass. *Ordered permutation groups.* Cambridge: Cambridge University Press, 1981, pp. x+266+xlix (cit. on pp. 11, 138).

[86] A. M. W. Glass and Yuri Gurevich. "The word problem for lattice-ordered groups". In: *Trans. Amer. Math. Soc.* 280.1 (1983), pp. 127–138 (cit. on p. 138).

[87] Kurt Gödel. "Die Vollständigkeit der Axiome des logischen Funktionenkalküls". In: *Monatsh. Math. Phys.* 37.1 (1930), pp. 349–360 (cit. on p. 197).

[88] Kurt Gödel. "Über formal unentscheidbare Sätze der Principia Mathematica und verwandter Systeme I". In: *Monatsh. Math. Phys.* 38.1 (1931), pp. 173–198 (cit. on p. 198).

[89] Peter Greenberg. "Les espaces de bracelets et les complexes de Stasheff". In: *Séminaire de Théorie Spectrale et Géométrie, No.*

9, Année 1990–1991. Vol. 9. Sémin. Théor. Spectr. Géom. Univ. Grenoble I, Saint-Martin-d'Hères, 1991, pp. 111–118 (cit. on p. 105).

[90] Peter Greenberg. "Projective aspects of the Higman-Thompson group". In: *Group theory from a geometrical viewpoint*. Ed. by E. Ghys, A. Haefliger, and A. Verjovsky. ICTP conference, 1990, Triest. World Scientific, Singapore, 1991, pp. 633–644 (cit. on p. 105).

[91] Peter Greenberg. "Les espaces de bracelets, les complexes de Stasheff et le groupe de Thompson". In: *Bol. Soc. Mat. Mexicana (2)* 37.1-2 (1992). Papers in honor of José Adem, pp. 189–201 (cit. on pp. 105, 191).

[92] Peter Greenberg. "Rapport sur la géométrie $\mathrm{SL}_2\mathbf{Z}$ par morceaux". In: *Séminaire de Théorie Spectrale et Géométrie, No. 10, Année 1991–1992*. Vol. 10. Sémin. Théor. Spectr. Géom. Univ. Grenoble I, Saint-Martin-d'Hères, 1992, pp. 93–95 (cit. on p. 105).

[93] Peter Greenberg. "The Euler class for "piecewise" groups". In: *Pacific J. Math.* 155.2 (1992), pp. 283–293 (cit. on p. 105).

[94] Peter Greenberg. "Torsion and periods in some groups of homeomorphisms of the circle". In: *Math. Proc. Cambridge Philos. Soc.* 112.1 (1992), pp. 29–44 (cit. on p. 105).

[95] Peter Greenberg. "Piecewise $\mathrm{SL}_2\mathbf{Z}$ geometry". In: *Trans. Amer. Math. Soc.* 335.2 (1993), pp. 705–720 (cit. on p. 105).

[96] Peter Greenberg and Vlad Sergiescu. "An acyclic extension of the braid group". In: *Comment. Math. Helvetici* 66 (1991), pp. 109–138 (cit. on p. 105).

[97] V. S. Guba. "On the properties of the Cayley graph of Richard Thompson's group F". In: *Internat. J. Algebra Comput.* 14.5-6 (2004). International Conference on Semigroups and Groups in honor of the 65th birthday of Prof. John Rhodes, pp. 677–702 (cit. on p. 277).

[98] V. S. Guba. "The Dehn function of Richard Thompson's group F is quadratic". In: *Invent. Math.* 163.2 (2006), pp. 313–342 (cit. on p. 3).

[99] Victor Guba and Mark Sapir. "Diagram groups". In: *Mem. Amer. Math. Soc.* 130.620 (1997), pp. viii+117 (cit. on pp. 93, 94).

[100] Pierre de la Harpe. *Topics in geometric group theory*. Chicago Lectures in Mathematics. University of Chicago Press, Chicago, IL, 2000, pp. vi+310 (cit. on p. 31).

[101] Assaf Hasson and Robert Bonnet. "In memoriam: Mati Rubin 1946–2017". In: *Bull. Symb. Log.* 24.2 (2018), pp. 181–185 (cit. on p. 216).

[102] Harold M. Hastings and Alex Heller. "Splitting homotopy idempotents". In: *Shape theory and geometric topology (Dubrovnik, 1981).* Vol. 870. Lecture Notes in Math. Springer, Berlin, 1981, pp. 23–36 (cit. on pp. 102, 103, 131, 139).

[103] Harold M. Hastings and Alex Heller. "Homotopy idempotents on finite-dimensional complexes split". In: *Proc. Amer. Math. Soc.* 85.4 (1982), pp. 619–622 (cit. on pp. 103, 139).

[104] Allen Hatcher. *Algebraic topology.* Cambridge University Press, Cambridge, 2002, pp. xii+544 (cit. on pp. 190, 288).

[105] Alex Heller. "On the representability of homotopy functors". In: *J. London Math. Soc. (2)* 23.3 (1981), pp. 551–562 (cit. on p. 131).

[106] Graham Higman. "On infinite simple permutation groups". In: *Publ. Math. Debrecen* 3 (1954), pp. 221–226 (cit. on pp. 27, 231).

[107] Graham Higman. *Finitely presented infinite simple groups.* Vol. 8. Notes on Pure Mathematics. Canberra: Department of Pure Mathematics, Australian National University, 1974 (cit. on pp. 92, 93, 100, 105, 106, 137, 235, 239, 240, 242, 257).

[108] Nigel Higson and Gennadi Kasparov. "E-theory and KK-theory for groups which act properly and isometrically on Hilbert space". In: *Invent. Math.* 144.1 (2001), pp. 23–74 (cit. on pp. 175, 176).

[109] P. J. Hilton and S. Wylie. *Homology theory: An introduction to algebraic topology.* Cambridge University Press, New York, 1960, pp. xv+484 (cit. on p. 189).

[110] Morris W. Hirsch. *Differential topology.* Vol. No. 33. Graduate Texts in Mathematics. Springer-Verlag, New York-Heidelberg, 1976, pp. x+221 (cit. on p. 59).

[111] Hajer Hmili Ben Ammar and Isabelle Liousse. "Nombre de classes de conjugaison d'éléments d'ordre fini dans les groupes de Brown-Thompson". In: *Bull. Soc. Math. France* 148.3 (2020), pp. 399–409 (cit. on p. 240).

[112] Wilfrid Hodges. *Model theory.* Vol. 42. Encyclopedia of Mathematics and its Applications. Cambridge University Press, Cambridge, 1993, pp. xiv+772 (cit. on pp. 195, 196, 214).

[113] Bruce Hughes. "Local similarities and the Haagerup property". In: *Groups Geom. Dyn.* 3.2 (2009). With an appendix by Daniel S. Farley, pp. 299–315 (cit. on p. 95).

[114] Bruce Hughes. "Trees, ultrametrics, and noncommutative geometry". In: *Pure Appl. Math. Q.* 8.1 (2012), pp. 221–312 (cit. on p. 265).

[115] James Hyde and Justin Tatch Moore. “Subgroups of $\mathrm{PL}_{+}I$ which do not embed into Thompson’s group F”. In: *Groups Geom. Dyn.* 17.2 (2023), pp. 533–554 (cit. on pp. 93, 216).

[116] Antoni Janiczak. “A remark concerning decidability of complete theories”. In: *J. Symb. Log.* 15 (1950), pp. 277–279 (cit. on p. 199).

[117] Bjarni Jónsson. “Defining relations for full semigroups of finite transformations”. In: *Michigan Math. J.* 9 (1962), pp. 77–85 (cit. on p. 99).

[118] Bjarni Jónsson and Alfred Tarski. “Two general theorems concerning free algebras”. In: *Bull. Amer. Math. Soc.* 62 (1956), p. 554 (cit. on p. 100).

[119] Bjarni Jónsson and Alfred Tarski. “On two properties of free algebras”. In: *Math. Scand.* 9 (1961), pp. 95–101 (cit. on pp. 98, 100, 124, 242).

[120] Pierre Julg. “Travaux de N. Higson et G. Kasparov sur la conjecture de Baum-Connes”. In: 252. Séminaire Bourbaki. Vol. 1997/98. 1998, Exp. No. 841, 4, 151–183 (cit. on p. 175).

[121] Anatole Katok and Boris Hasselblatt. *Introduction to the modern theory of dynamical systems.* Vol. 54. Encyclopedia of Mathematics and its Applications. With a supplementary chapter by Katok and Leonardo Mendoza. Cambridge University Press, Cambridge, 1995, pp. xviii+802 (cit. on p. 78).

[122] Olga Kharlampovich, Alexei Myasnikov, and Mahmood Sohrabi. *Rich groups, weak second order logic, and applications.* 2021 (cit. on pp. 215, 216).

[123] Anatole Khelif. “Bi-interprétabilité et structures QFA: étude de groupes résolubles et des anneaux commutatifs”. In: *C. R. Math. Acad. Sci. Paris* 345.2 (2007), pp. 59–61 (cit. on p. 215).

[124] Sang-hyun Kim, Thomas Koberda, and J. de la Nuez González. “First order rigidity of homeomorphism groups of manifolds”. In: *Commun. Am. Math. Soc.* 5 (2025), pp. 144–194 (cit. on p. 216).

[125] Felix Klein. “Neue Beiträge zur Riemann’schen Functionentheorie”. In: *Math. Annalen* 21 (2 1883), pp. 141–218 (cit. on p. 79).

[126] Thomas Koberda and J. de la Nuez González. *Uniform first order interpretation of the second order theory of countable groups of homeomorphisms.* 2023. arXiv: 2312.16334 (cit. on p. 216).

[127] Thomas Koberda and J. de la Nuez González. *Locally approximating groups of homeomorphisms of manifolds.* 2024. arXiv: 2410.16108 (cit. on p. 216).

[128] Clément Lasserre. "R. J. Thompson's groups F and T are bi-interpretable with the ring of the integers". In: *J. Symb. Log.* 79.3 (2014), pp. 693–711 (cit. on p. 213).

[129] Mark V. Lawson. *Inverse semigroups*. The theory of partial symmetries. World Scientific Publishing Co., Inc., River Edge, NJ, 1998, pp. xiv+411 (cit. on pp. 137, 253).

[130] Mark V. Lawson. "Orthogonal completions of the polycyclic monoids". In: *Comm. Algebra* 35.5 (2007), pp. 1651–1660 (cit. on p. 138).

[131] Mark V. Lawson. "The polycyclic monoids P_n and the Thompson groups $V_{n,1}$". In: *Comm. Algebra* 35.12 (2007), pp. 4068–4087 (cit. on pp. 95, 138).

[132] Mark V. Lawson. "The polycyclic inverse monoids and the Thompson groups revisited". In: *Semigroups, categories, and partial algebras*. Vol. 345. Springer Proc. Math. Stat. Springer, Singapore, 2021, pp. 179–214 (cit. on p. 138).

[133] Mark V. Lawson. *Introduction to inverse semigroups*. 2023. arXiv: 2304.13580 (cit. on p. 253).

[134] Mark V. Lawson and Alina Vdovina. "Generalizations of free monoids". In: *Semigroup Forum* 109.1 (2024), pp. 167–185 (cit. on p. 95).

[135] Ian J. Leary. "A metric Kan-Thurston theorem". In: *J. Topol.* 6.1 (2013), pp. 251–284 (cit. on p. 291).

[136] W. G. Leavitt. "Modules over rings of words". In: *Proc. Amer. Math. Soc.* 7 (1956), pp. 188–193 (cit. on pp. 138, 242, 244, 245).

[137] W. G. Leavitt. "Modules without invariant basis number". In: *Proc. Amer. Math. Soc.* 8 (1957), pp. 322–328 (cit. on pp. 138, 242).

[138] W. G. Leavitt. "The module type of a ring". In: *Trans. Amer. Math. Soc.* 103 (1962), pp. 113–130 (cit. on pp. 138, 242).

[139] Yash Lodha and Justin Tatch Moore. "A nonamenable finitely presented group of piecewise projective homeomorphisms". In: *Groups Geom. Dyn.* 10.1 (2016), pp. 177–200 (cit. on p. 94).

[140] Yangxiao Luo and Shunyu Wan. *Half grid diagrams and Thompson links*. 2024. arXiv: 2408.10875 (cit. on p. 191).

[141] Roger C. Lyndon and Paul E. Schupp. *Combinatorial group theory*. Ergebnisse der Mathematik und ihrer Grenzgebiete, Band 89. Springer-Verlag, Berlin-New York, 1977, pp. xiv+339 (cit. on pp. 61, 81).

[142] A. M. Macbeath. "Groups of homeomorphisms of a simply connected space". In: *Ann. of Math. (2)* 79 (1964), pp. 473–488 (cit. on p. 176).

[143] Saunders MacLane. “Natural associativity and commutativity”. In: *Rice Univ. Studies* 49.4 (1963), pp. 28–46 (cit. on p. 135).

[144] Saunders MacLane. *Categories for the working mathematician.* Graduate Texts in Mathematics, Vol. 5. New York: Springer-Verlag, 1971, pp. ix+262 (cit. on pp. 60, 122).

[145] Anatoliĭ Ivanovič Mal′ cev. *The metamathematics of algebraic systems. Collected papers: 1936–1967.* Vol. Vol. 66. Studies in Logic and the Foundations of Mathematics. Translated, edited, and provided with supplementary notes by Benjamin Franklin Wells, III. North-Holland Publishing Co., Amsterdam-London, 1971, pp. xviii+494 (cit. on p. 215).

[146] A. I. Mal′ ceve. “Some correspondences between rings and groups”. In: *Mat. Sb. (N.S.)* 50(92) (1960). translated in Amer. Math. Soc. Transl. (2) 39 (1964), 139–206, pp. 257–266 (cit. on p. 215).

[147] Kathryn Mann and Maxime Wolff. *Reconstructing maps out of groups.* https://arxiv.org/abs/1907.03024. 2019 (cit. on p. 216).

[148] David Marker. *Model theory.* Vol. 217. Graduate Texts in Mathematics. An introduction. Springer-Verlag, New York, 2002, pp. viii+342 (cit. on pp. 195, 196).

[149] Conchita Martínez-Pérez, Francesco Matucci, and Brita E. A. Nucinkis. “Cohomological finiteness conditions and centralisers in generalisations of Thompson’s group V”. In: *Forum Math.* 28.5 (2016), pp. 909–921 (cit. on p. 95).

[150] Conchita Martínez-Pérez and Brita E. A. Nucinkis. “Bredon cohomological finiteness conditions for generalisations of Thompson groups”. In: *Groups Geom. Dyn.* 7.4 (2013), pp. 931–959 (cit. on p. 95).

[151] Kengo Matsumoto and Hiroki Matui. “Full groups of Cuntz-Krieger algebras and Higman-Thompson groups”. In: *Groups Geom. Dyn.* 11.2 (2017), pp. 499–531 (cit. on p. 265).

[152] Ralph McKenzie and Richard J. Thompson. “An elementary construction of unsolvable word problems in group theory”. In: *Word problems: decision problems and the Burnside problem in group theory (Conf., Univ. California, Irvine, Calif. 1969; dedicated to Hanna Neumann).* Ed. by F. B. Cannonito W. W. Boone and R. C. Lyndon. Studies in Logic and the Foundations of Math., Vol. 71. Amsterdam: North-Holland, 1973, pp. 457–478 (cit. on pp. 92, 98, 99, 107, 139).

[153] John Meier. *Groups, graphs and trees.* Vol. 73. London Mathematical Society Student Texts. An introduction to the geometry of infinite groups. Cambridge University Press, Cambridge, 2008, pp. xii+231 (cit. on p. 92).

[154] Ch. F. Miller III. "Decision problems in algebraic classes of groups (a survey)". In: *Word problems: decision problems and the Burnside problem in group theory (Conf., Univ. California, Irvine, Calif., 1969; dedicated to Hanna Neumann)*. Stud. Logic Found. Math., Vol. 71. North-Holland, Amsterdam-London, 1973, pp. 507–523 (cit. on p. 104).

[155] J. Donald Monk. *Mathematical logic*. Graduate Texts in Mathematics, No. 37. Springer-Verlag, New York-Heidelberg, 1976, pp. x+531 (cit. on pp. 198, 199).

[156] Terrence Napier and Mohan Ramachandran. "Thompson's group F is not Kähler". In: *Topological and asymptotic aspects of group theory*. Vol. 394. Contemp. Math. Amer. Math. Soc., Providence, RI, 2006, pp. 197–201 (cit. on p. 93).

[157] Terrence Napier and Mohan Ramachandran. "Filtered ends, proper holomorphic mappings of Kähler manifolds to Riemann surfaces, and Kähler groups". In: *Geom. Funct. Anal.* 17.5 (2007), pp. 1621–1654 (cit. on p. 93).

[158] Andrés Navas. "On the dynamics of (left) orderable groups". In: *Ann. Inst. Fourier (Grenoble)* 60.5 (2010), pp. 1685–1740 (cit. on p. 21).

[159] Volodymyr V. Nekrashevych. "Cuntz-Pimsner algebras of group actions". In: *J. Operator Theory* 52.2 (2004), pp. 223–249 (cit. on pp. 138, 265).

[160] Peter M. Neumann. "On the structure of standard wreath products of groups". In: *Math. Z.* 84 (1964), pp. 343–373 (cit. on p. 17).

[161] M. H. A. Newman. "On theories with a combinatorial definition of "equivalence."" In: *Ann. of Math. (2)* 43 (1942), pp. 223–243 (cit. on p. 280).

[162] Morris Newman. "Free subgroups and normal subgroups of the modular group". In: *Illinois J. Math.* 8 (1964), pp. 262–265 (cit. on p. 80).

[163] Graham Niblo and Lawrence Reeves. "Groups acting on CAT(0) cube complexes". In: *Geom. Topol.* 1 (1997), approx. 7 pp. (Cit. on pp. 175, 176).

[164] André Nies. "Separating classes of groups by first-order sentences". In: *Internat. J. Algebra Comput.* 13.3 (2003), pp. 287–302 (cit. on p. 215).

[165] Andre Nies. "Comparing quasi-finitely axiomatizable and prime groups". In: *J. Group Theory* 10.3 (2007), pp. 347–361 (cit. on p. 215).

[166] Maurice Nivat and Jean-François Perrot. "Une généralisation du monoïde bicyclique". In: *C. R. Acad. Sci. Paris Sér. A-B* 271 (1970), A824–A827 (cit. on p. 138).

[167] Ivan Niven and Herbert S. Zuckerman. *An introduction to the theory of numbers*. John Wiley & Sons, Inc., New York-London, 1960, pp. viii+250 (cit. on p. 16).

[168] E. Pardo. "The isomorphism problem for Higman-Thompson groups". In: *J. Algebra* 344 (2011), pp. 172–183 (cit. on pp. 235, 242, 243, 255).

[169] J. F. Plante and W. P. Thurston. "Polynomial growth in holonomy groups of foliations". In: *Comment. Math. Helv.* 51.4 (1976), pp. 567–584 (cit. on p. 20).

[170] Raphael M. Robinson. "Undecidable rings". In: *Trans. Amer. Math. Soc.* 70 (1951), pp. 137–159 (cit. on p. 209).

[171] Matatyahu Rubin. "On the reconstruction of topological spaces from their groups of homeomorphisms". In: *Trans. Amer. Math. Soc.* 312.2 (1989), pp. 487–538 (cit. on pp. 204, 216).

[172] Matatyahu Rubin and Yosef Yomdin. "Reconstruction of manifolds and subsets of normed spaces from subgroups of their homeomorphism groups". In: *Dissertationes Math.* 435 (2005), p. 246 (cit. on p. 216).

[173] E. A. Scott. "A tour around finitely presented infinite simple groups". In: *Algorithms and Classification in Combinatorial Group Theory*. Ed. by G. Baumslag and III C. F. Miller. Vol. 23. MSRI Publications. Springer-Verlag, New York, 1992, pp. 83–119 (cit. on pp. 92, 106).

[174] Elizabeth A. Scott. "A construction which can be used to produce finitely presented infinite simple groups". In: *J. Algebra* 90.2 (1984), pp. 294–322 (cit. on pp. 104, 106, 240).

[175] Elizabeth A. Scott. "A finitely presented simple group with unsolvable conjugacy problem". In: *J. Algebra* 90.2 (1984), pp. 333–353 (cit. on pp. 104, 106, 240).

[176] Elizabeth A. Scott. "The embedding of certain linear and abelian groups in finitely presented simple groups". In: *J. Algebra* 90.2 (1984), pp. 323–332 (cit. on pp. 104, 106, 240).

[177] J.-P. Serre. "Arithmetic groups". In: *Homological group theory (Proc. Sympos., Durham, 1977)*. Vol. 36. London Math. Soc. Lecture Note Ser. Cambridge Univ. Press, Cambridge-New York, 1979, pp. 105–136 (cit. on p. 190).

[178] Jean-Pierre Serre. "Cohomologie des groupes discrets". In: *Prospects in mathematics (Proc. Sympos., Princeton Univ., Princeton, N.J.,*

1970). Ann. of Math. Stud., No. 70. Princeton Univ. Press, Princeton, NJ, 1971, pp. 77–169 (cit. on p. 190).

[179] D. M. Smirnov. “Bases and automorphisms of free Cantor algebras of finite rank”. In: *Algebra i Logika* 13 (1973). English translation: Algebra and Logic 13 (1974), no. 1, 17–33 (1975.), pp. 35–62, 121 (cit. on pp. 93, 101, 130).

[180] Christophe Soulé. “Groupes opérant sur un complexe simplicial avec domaine fondamental”. In: *C. R. Acad. Sci. Paris Sér. A-B* 276 (1973), A607–A609 (cit. on p. 176).

[181] Edwin H. Spanier. *Algebraic topology*. McGraw-Hill Book Co., New York-Toronto, Ont.-London, 1966, pp. xiv+528 (cit. on p. 133).

[182] James Dillon Stasheff. “Homotopy associativity of H-spaces. I, II”. In: *Trans. Amer. Math. Soc.* 108 (1963). **1**08 (1963), 275-292; ibid, pp. 293–312 (cit. on pp. 135, 191).

[183] Melanie Stein. “Groups of piecewise linear homeomorphisms”. In: *Trans. Amer. Math. Soc.* 332 (1992), pp. 477–514 (cit. on pp. 106, 168, 191).

[184] Stephen M. Stigler. “Stigler’s law of eponymy”. In: *Trans. N.Y. Acad. Sci.* 39.1 Series II (1980), pp. 147–157 (cit. on p. 1).

[185] S. Świerczkowski. “On isomorphic free algebras”. In: *Fund. Math.* 50 (1961), pp. 35–44 (cit. on p. 101).

[186] Markus Szymik and Nathalie Wahl. “The homology of the Higman-Thompson groups”. In: *Invent. Math.* 216.2 (2019), pp. 445–518 (cit. on p. 104).

[187] Alfred Tarski. *Undecidable theories*. Studies in Logic and the Foundations of Mathematics. In collaboration with Andrzej Mostowski and Raphael M. Robinson. North-Holland Publishing Co., Amsterdam, 1953, pp. xi+98 (cit. on pp. 198, 199).

[188] Amanda Taylor. *Locally Solvable Subgroups of PLo(I)*. Thesis (Ph.D.)–State University of New York at Binghamton. ProQuest LLC, Ann Arbor, MI, 2017, p. 68 (cit. on p. 216).

[189] Richard J. Thompson. “Embeddings into finitely generated simple groups which preserve the word problem”. In: *Word problems, II (Conf. on Decision Problems in Algebra, Oxford, 1976)*. Vol. 95. Studies in Logic and the Foundations of Mathematics. North-Holland, Amsterdam-New York, 1980, pp. 401–441 (cit. on pp. 9, 100, 104, 240).

[190] Richard Joseph Thompson. *Transformational Structure of Algebraic Logics*. Thesis (Ph.D.)–University of California, Berkeley. ProQuest LLC, Ann Arbor, MI, 1979, p. 345 (cit. on p. 99).

[191] Werner Thumann. "Operad groups and their finiteness properties". In: *Adv. Math.* 307 (2017), pp. 417–487 (cit. on p. 95).

[192] C. T. C. Wall. "Finiteness conditions for CW-complexes". In: *Ann. of Math. (2)* 81 (1965), pp. 56–69 (cit. on p. 190).

[193] C. T. C. Wall, ed. *Homological group theory*. Vol. 36. London Mathematical Society Lecture Note Series. Cambridge University Press, Cambridge-New York, 1979, pp. ix+394 (cit. on pp. 103, 181, 190).

[194] Bill de Witt and Luke Elliott. *Finite Presentability of Brin-Higman-Thompson Monoids via Free Jónsson-Tarski Algebras*. 2023. arXiv: 2303.16044 (cit. on p. 131).

Index