Conjugacy and centralisers in Thompson's group T

David Matthew Robertson

April 2019

Thesis submitted for the degree of
Doctor of Philosophy

to the

School of Mathematics, Statistics and Physics
Newcastle University
Newcastle upon Tyne
United Kingdom

This thesis was typeset using LaTeX, equipped with a modified version of the tufte-latex[1] classes. The aim was to mimic the style of Aaron Turon's thesis.[2] Illustrations were created with Ti*k*Z.[3] Finally, the bibliography was managed using biblatex and biber.[4]

Conjugacy and centralisers in Thompson's group T
David Matthew Robertson, April 2019

[1] https://tufte-latex.github.io/tufte-latex/
[2] https://people.mpi-sws.org/~turon/
[3] http://www.ctan.org/pkg/pgf
[4] https://www.ctan.org/pkg/biblatex and https://www.ctan.org/pkg/biber

Acknowledgements

I SHOULD START by thanking those who are most important, namely my sister Helen and parents Jackie and Graeme. Your support over the years means so much. I will always be grateful.

Newcastle's maths, stats and physics department is a different, larger and much quirkier kind of family. Thanks to the school for funding my research and for giving me the chance to teach. Thank you to the administrative staff for keeping the whole place running with warmth, friendliness and kindness. In particular, many thanks to those I know best: Georgina Kay-Black, Jackie Martin, and Lauren Daley.

Thank you to the IT and computing staff, and in particular to Michael Beaty. I've learnt much more about Linux/Unix than I ever expected.[1] Thank you for sharing your time, energy and knowledge. Speaking of TEXnical subjects, I (and probably most mathematicians) would like to thank the community at `tex.stackexchange.com` for demystifying the mysterious and answering the unanswerable.

Thank you to my supervisor Andrew Duncan, for your patience, guidance, insight[2] and humour. Thank you also to my second supervisor Sarah Rees, for your advice and support. I'd also like to thank the other boffins who've encouraged me:[3] at home, Alina Vdovina, James Waldron, Martina Balagović, Peter Jørgensen, Stefan Kolb; from further afield: Ana Khukhro, Brita Nucinkis, Collin Bleak, Ewa Bieniecka, José Burillo, Laura Ciobanu, Montse Casals-Ruiz, and Yago Antolín.

The commuters of the Go North East X9/X10 service made the A19 much more enjoyable than it has any right to be. Thanks in particular to June Young for supplying **i** puzzle pages and crosswords.[4]

The troublemakers David Stewart, Richard Boys, Stuart Douglas, and Stuart Hall deserve credit for being part of *the greater education*. A small army[5] of other students, postdocs and visitors made contributions on that front too. Muchas gracias to Alex Svalova, Amit Seta, Anne-Kathleen Malchow,[6] Clarissa Barratt, David Cushing, David Pescod,[7] Gavin Whitaker, George Stagg,[8] Hayley Moore, Jack Aiston, Jack Walton,[9] Keith Newman,[10] Liam Dobson[11] Matthew Robinson, Nathan Barker,[12] Nick Loughlin, Philipp Edelmann, Sarah Jowett, Siddiqua Mazhar, and Stephen Johnson. To those before me: thank you for leading the way. To those who started with me: thanks for the memories. To those only just beginning: enjoy yourselves!

Joe T. Matthews and Tom 'anything but' Bland aren't listed above, because they deserve extra-special thanks. Shame I ran out of space.

David Matthew Robertson, April 2019

[1] . . . which has been really handy these last few months!

[2] I never truly appreciated the infinite number of ways to say *'hmmm. . . '* until I observed the master at work.

[3] Your support meant (and still means) a lot to a nervous newbie.

[4] I will never forget the word *legerdemain*.

[5] There isn't room to do you all justice; let me mention those I know best.

[6] Aspall aficionado.

[7] *Woof!*

[8] Computing craftsman and Nethack nurturer.

[9] Source of scripting shenanigans.

[10] Partner in pedantry.

[11] Master of the quantum realm!

[12] This thesis is all your fault—I would never have got into Thompson's groups without your introduction!

Abstract

The Thompson family of groups $F < T < V$ are well-known as interesting (counter-) examples in group theory. Working algebraically with these objects is difficult, and yet the groups are computationally tractable. For instance, T and V are infinite simple groups; despite this, both are finitely presented.

This thesis studies the middle group T, using the piecewise-linear function point of view. We present a solution to the conjugacy problem in this group, adapting the approach of Kassabov and Matucci[1] to the same problem in F. Conjugacy of elements in T was shown[2] to be decidable by Belk and Matucci; however our approach constructs explicit conjugators (when they exist). Later, we refine the description given by Matucci[3] of nontorsion elements' centralisers in T.

[1] Kassabov and Matucci 2012.

[2] Belk and Matucci 2014.

[3] Matucci 2008, Chapter 7.

* * *

The first chapter introduces the world of Thompson's groups. The sections on *cyclic order*, the generalised groups $\mathrm{PL}_{S,G}$ and groupoid $\mathcal{PL}_2$, and on the *Cantor space* are particularly important for readers interested in the rest of the thesis.

The second chapter discusses Thompson's groups from a dynamical point of view. We summarise how F, T and V rerrange the interval, noting the distinction between dyadic and nondyadic points. Focussing on T, we introduce the *rotation number* and explain what we can learn from it. Amidst all this we present a number of intermediate results, forming a toolkit for use in later chapters.

The third chapter studies conjugacy in T. We narrow the search space by finding constraints that a conjugator must satisfy. Next, we break the conjugacy problem into a search for a *coarse* and *fine* conjugator, the product of which—if they exist—is a bona fide conjugator. We solve these search problems,[4] and thus solve the conjugacy problem in T.

[4] Lemma 3.3.2 and Algorithms 3.3.5 and 3.4.8.

In the fourth chapter, we study element centralisers in T via a particular group extension. We focus on nontorsion elements, providing small details missing from Matucci's proof which identifies the extension's kernel.[5] We explain how to find the size of the extension's quotient, by reducing the problem to a search for coarse conjugators.[6]

[5] Theorem 4.2.1 and Remark 4.2.2.

[6] Algorithm 4.2.3.

The final chapter describes the extension structure of $C_T(\alpha)$ in more detail.[7] We do so by classifying α into one of four cases. In all but one case, this extension splits (as a wreath or direct product); in the remaining case, we identify[8] exactly when the extension splits (again as a wreath product). In each case, we describe the centraliser's structure[9] in terms of integer parameters. We then show how to construct[10] an element of T whose centraliser has a given list of parameters.

[7] Section 5.2.

[8] See Proposition 5.2.15 and the summary in in Theorem 5.3.2.

[9] Proposition 5.2.17.

[10] Corollaries 5.2.6, 5.2.16, 5.2.18, 5.2.23 and 5.2.31.

Contents

Introduction xiii
Notation . . . xvii
Summary of contributions . . . xviii

1 The Thompson family of groups 1
1.1 The main trio of groups . . . 1
1.1.1 Thompson's group F . . . 1
1.1.2 Tree pair diagrams . . . 4
1.1.3 Thompson's group T . . . 7
1.1.4 Thompson's group V . . . 10
1.2 Other points of view . . . 11
1.2.1 Generalised gradients and breakpoints . . . 11
1.2.2 More piecewise-linear maps . . . 13
1.2.3 Cantor space homeomorphisms . . . 16

2 The dynamics of Thompson's groups 21
2.1 Actions of Thompson's groups . . . 21
2.1.1 Dyadics versus nondyadics . . . 21
2.1.2 The action on the rationals . . . 23
2.1.3 Rotations . . . 24
2.1.4 The rotation number . . . 26
2.2 Revealing pairs . . . 29
2.3 Behaviour near periodic points . . . 35
2.3.1 Important points for V . . . 36
2.3.2 Important points for T . . . 38
2.3.3 Partitioning at important points . . . 42
2.3.4 Gradients . . . 44
2.4 Aligning finite lists of points . . . 47
2.4.1 Alignment involving a dyadic point . . . 48
2.4.2 Alignment involving no dyadics . . . 50

3 Conjugacy in T 53
3.1 Constraints on conjugators . . . 54
3.2 The miracle of F . . . 56
3.3 Finding coarse conjugators . . . 58
3.3.1 No important points . . . 59
3.3.2 At least one important point . . . 60
3.3.3 Reducing to a search for fine conjugators . . . 60
3.4 Finding fine conjugators . . . 62
3.4.1 Restriction to a representative arc . . . 63
3.4.2 Reassembly . . . 67

4 Centralisers in T 71
4.1 Centralisers in F 72
4.2 Centralisers of nontorsion elements in T 76
4.2.1 Identifying the kernel 77
4.2.2 Identifying the quotient 78

5 Nontorsion centraliser structure 81
5.1 Examples 81
5.2 Extension structure in detail 88
5.2.1 Case 1: dyadic self-similar elements 88
5.2.2 Case 2: dyadic offset-similar elements 94
5.2.3 Case 3: nondyadic self-similar elements 104
5.2.4 Case 4: nondyadic, offset-similar elements 109
5.3 Summary 114

A Deferred calculations 117
A.1 Script source code 119

B To split, or not to split? 121

Bibliography 125

List of Figures

1.1 Three example elements of F . . . 2
1.2 A rectangle diagram for x_0 . . . 3
1.3 Stacking rectangle diagrams . . . 3
1.4 Examples of small binary trees . . . 4
1.5 The infinite binary tree $\mathcal{T}_2$ . . . 4
1.6 Tree structure of standard dyadic intervals . . . 5
1.7 Tree pair diagrams for x_0 and x_1 . . . 5
1.8 Multiple tree pairs for x_0 . . . 5
1.9 Tree pair diagram for x_0^{-1} . . . 6
1.10 Why an alternative transitivity axiom is invalid . . . 8
1.11 An example element of T . . . 8
1.12 Annulus diagram for an element of T . . . 9
1.13 x_0 as an element of T . . . 9
1.14 An element of T as a right-continuous function on I . . . 9
1.15 Cannon-Floyd-Parry's generator c for T . . . 10
1.16 An example element of V . . . 11
1.17 Cannon-Floyd-Parry's generator π_0 for V . . . 11
1.18 An element in $\mathrm{PL}^{\mathrm{rest}}$ but not $\mathrm{PL}^{\mathrm{flat}}$ . . . 12
1.19 PL_2 maps $[0,1] \to [0,1/2]$ . . . 13
1.20 A schematic illustrating the groupoid $\mathcal{PL}_2$ . . . 14
1.21 More PL_2 maps $[0,1] \to [0,1/2]$ . . . 15
1.22 A basic neighbourhood in the boundary of $\mathcal{T}_2$ . . . 17
1.23 Cantor set construction steps are levels of $\mathcal{T}_2$ . . . 18
1.24 Maps involving the Cantor set and Cantor space . . . 19
1.25 The element $x_0 \in F$ as a transducer . . . 19
1.26 A transducer which does not model an element of V . . . 19

2.1 An element of F as a dynamical system . . . 21
2.2 The dyadic rotation by $3/8$ . . . 24
2.3 The effect of conjugating by a dyadic rotation . . . 25
2.4 An element of $T \setminus F$ with zero rotation number . . . 26
2.5 The rotation number is not a homomorphism . . . 29
2.6 A different tree pair for Figure 1.16's β . . . 29
2.7 Examples of operations on binary trees . . . 30
2.8 A schematic overview of a tree pair . . . 30
2.9 An example of attractors and repellers . . . 31
2.10 A revealing pair . . . 31
2.11 A revealing pair exhibiting all types of IACs . . . 32
2.12 An orbit of two repelling periodic points . . . 32
2.13 Schematic of a source-sink chain . . . 33

2.14 Periods are not preserved over flow lines 34
2.15 An element of V with differing dynamics about $1/2$ 37
2.16 Important points are isolated points 37
2.17 Important points are boundary points 39
2.18 An element with $1/3$ as a fixed point 41
2.19 Gradients at fixed points create attractors and repellers . . 42
2.20 The circle partitioned by important points 42
2.21 More ways to partition the circle 43
2.22 Elements of T with only nondyadic important points . . . 45
2.23 Conjugate elements of T with nonpreserved gradients . . . 45
2.24 Plots of functions from Figure 2.23 45

3.1 Schematic illustrating the conjugation of torsion elements . 59
3.2 Pointwise periodic intervals 61
3.3 Sketch illustrating the representative arc 65
3.4 Two cases for partitioning Q 65

4.1 Cyclic versus F blocks in an F-like centraliser 76
4.2 Examples illustrating the inequality on h 79

5.1 Notation for discussing centraliser examples 81
5.2 A nontorsion element α whose centraliser can be identified 81
5.3 An element $\beta = \alpha\rho_{1/2}$ similar to α 82
5.4 The square of β from Figure 5.3 83
5.5 A more complicated element $\gamma \in F$ 83
5.6 The centralising element ω 84
5.7 The product $\delta = \gamma\rho_{1/2}$ 84
5.8 The square of δ from Figure 5.7 85
5.9 An element ϵ similar to α. 85
5.10 Decomposition of $C_F(\epsilon)$ 85
5.11 An element $\zeta = \epsilon\rho_{1/2}$ similar to β 86
5.12 The square of ζ 86
5.13 Summary of centraliser examples 87
5.14 Names for the four centraliser cases 88
5.15 Schematic of an element α whose centraliser splits 89
5.16 The splitting element δ 89
5.17 Schematic of a wreath product centraliser 91
5.18 Step in constructing a wreath centraliser 92
5.19 The building block x for a wreath centraliser 93
5.20 An element with prescribed wreath centraliser 93
5.21 Schematic of a dyadic offset-similar element 94
5.22 Constructing a centralising element 96
5.23 Building a splitting element from a root of X 101
5.24 A splitting element which yields a wreath product 102
5.25 Schematic of dyadic offset similar centraliser 102
5.26 A copy of α from Figure 2.22 108
5.27 The element β from Figure 2.22 114

5.28 Nontorsion centraliser summary 116

A.1 The element δ^2 from Figure 5.7 117

B.1 Some ingredients introduced in Example 5.1.4 124
B.2 δ and δ^2 from Example 5.1.6 124
B.3 Elements of F related to δ^2 124

Introduction

In this thesis we present two main results. The third chapter describes a solution[1] to the conjugacy problem in T which recovers an explicit conjugator whenever a conjugator exists.

[1] Theorem 3.4.9.

Theorem A. *The conjugacy search problem in T has an effective solution.*

The final chapter exhibits the centraliser of a nontorsion element $\alpha \in T$ as a group extension and details[2] the structure of this extension.

[2] Theorem 5.3.2, but see also Section 5.3.

Theorem B. *Let $\beta \in T$ be nontorsion with rotation number $\rho = p/q$ in lowest terms. Matucci expresses $C_T(\beta)$ as a particular group extension. This extension is nonsplit if and only if*

- $\rho \neq 0$,
- *β has a dyadic important point, and*
- *β^q does* not *have a qth root with a fixed point.*

Otherwise the extension splits as a wreath or direct product.

These results build upon the work of Kassabov-Matucci[3] and Bleak-Kassabov-Matucci,[4] respectively.

[3] Kassabov and Matucci 2012.

[4] Matucci 2008, Chapter 7.

* * *

Chapter 1 sets the scene, by defining the main trio $F < T < V$ of Thompson's groups. F was originally introduced in the 1960s in connection with Richard J. Thompson's work on logic and associativity. Today, F and its many generalisations form an ever-expanding family of groups, and provide a rich source of interesting problems, examples and counter-examples in group theory.

Thompson's groups are sometimes called *chameleons*, because they have many different definitions and appear in many different contexts. We can see F, T and V via generators and relators using a finite presentation; as automorphism groups[5] of a certain algebra studied by Jónsson and Tarski; as homeomorphism groups of the Cantor set or other fractals[6]; as groups of *tree-pair diagrams*; as the 'geometry group of the associative law[7]'; and as groups of certain piecewise-linear (PL) functions either on the unit interval or on the whole real line.[8] While we use other technologies, we focus on the PL function point of view in particular, so that we can use the approach of and results from Kassabov-Matucci's article.

[5] K. S. Brown 1987; Higman 1974.

[6] Belk and Forrest 2018.

[7] Dehornoy 2005.

[8] Brin and Squier 2001.

After defining F as a group of PL functions, we carefully explain how these correspond to tree-pair diagrams. Doing so also helps to establish our notation, in which functions are placed to the right of

We define F, T and V in Section 1.1.

their arguments—see page xvii. To generalise to T, we need to define PL maps on the circle which preserve orientation. To pin down the meaning of 'preserve orientation' we introduce the idea of a *cyclic order*, which is modelled by a ternary (rather than binary) relation. For completeness, we give a definition of V, though we don't present any results regarding V in this thesis.

Next, we introduce the generalised groups $\mathrm{PL}_{S,G}(J)$, which allow us to vary the permitted breakpoints S, permitted gradients G, and the domain J of our functions. Particularly important are the subsets of *one-bump* and *almost one-bump* functions; these help us to decompose elements of F into simple building blocks. To facilitate this decomposition it's helpful to work with 'Thompson-like' maps whose domains are not equal to their codomains. We define these formally, noting that the collection of such maps form a groupoid we call $\mathcal{PL}_2$. To close our introduction, we define the Cantor set C and identify it with the Cantor space $\partial\mathcal{T}_2$ of infinite binary strings. This interacts nicely with tree pair diagrams (which can be described using finite binary strings), and will later allow us to identify the *important points* of an element.

The objects in this paragraph are introduced in Section 1.2.

Chapter 2 views elements of T as discrete dynamical systems on the circle. We emphasise that our group elements are not just functions, but objects that push points through an orbit over time. Here we introduce some of the key ideas behind our later arguments, for instance: the distinction between dyadic and nondyadic points on the circle; reducing problems in T to problems in F using conjugation by a rotation; and the notion of the *rotation number*.

This paragraph's material features in Section 2.1.

In order to classify the orbits of single points under an element of V, we give an overview of *revealing tree-pairs*, introduced by Brin[9] and later studied by Salazar-Díaz.[10] From these authors' work, we learn that every element of T has finitely many isolated periodic points, which can be located directly from a revealing pair. These points are called *important*,[11] usually when thought of as points in the Cantor space. We propose the name important points for their images in $[0, 1]$ or the circle, to highlight the fact that a dyadic important point can have *two* important preimages. We can learn much from about an element α from these points x, in particular by studying one-sided gradients at x.

[9] Brin 2004, Section 10.

[10] O. P. Salazar-Díaz 2010.

Revealing pairs are introduced in Section 2.2; the discussion regarding important points takes place in Section 2.3.

[11] This term was introduced in the article Bleak and O. Salazar-Díaz 2013, after Corollary 2.3.

Suppose we have two finite lists of points on the circle, both of the same size. We say that these can be *aligned* if there is an element $\delta \in T$ which sends the points of the first list to points of the second. Put differently, the lists are alignable if some element of T restricts to a bijection between the two lists. We explain how to determine if two lists are alignable, and how to construct an aligning map δ if so. We do so by reducing to the analogous problem for F. This is our first instance of a formal *search problem*.

This alignment problem is discussed in Section 2.4.

Chapter 3 studies conjugacy in T. In its introduction, we summarise the literature on conjugacy in Thompson's groups. There are two

important pieces to highlight for our purposes. Firstly, an article of Belk and Matucci[12] gives a unified approach to solving the conjugacy decision problem in F, T and V by using their *strand diagram* technology. They do not construct explicit conjugators (which would solve the corresponding *search* problem), though we suggest that this may be possible by studying their notion of a 'cutting path'. Secondly, the article of Kassabov-Matucci mentioned earlier solves the conjugacy search problem (and more) in F. Let us also mention the full conjugacy invariant for F provided by Gill–Short,[13] itself based on work of Brin–Squier.[14]

[12] Belk and Matucci 2014.

[13] Gill and Short 2013.

[14] Brin and Squier 2001.

In passing, we also mention an article by Barker, Duncan and the author[15] which discusses the power conjugacy problem in V. The current author contributed to this article by correcting a lemma of Higman; without this patch, Higman's solution to the conjugacy problem in V was incomplete. (The first complete solution was given by Salazar-Díaz.[16]) Alongside this article, the author has also written a collection of Python scripts[17] to perform computations in F, T and V. These scripts also produced Ti*k*Z code for the plots and tree pair diagrams used throughout this document. While this article and software are not part of this thesis, this work was part of the author's PhD project, and so we include it here for completeness. Note that other software is available: in particular we note Roman Kogan's *nvTrees*,[18] later modified by Collin Bleak to form *vTrees*.

[15] Barker, Duncan and Robertson 2016, For the focus of the author's contribution, see the article's Remark 4.15.

[16] O. P. Salazar-Díaz 2010.

[17] Robertson 2015.

[18] Kogan 2008.

After its introduction, the third chapter proceeds by establishing constraints on conjugators. The key idea is if $\alpha^\gamma = \beta$, then the behaviour of γ at a point x determines the behaviour of γ on the α-orbit of x. We can attempt to use this idea to build a conjugator γ, though we have issues like well-definedness and breakpoint locations to deal with. We go on to establish Theorem A, by breaking the search for conjugators into a search for *coarse* and *fine* conjugators. The former boils down to the alignment problem solved at the end of Chapter 2; the latter is essentially a conjugacy problem in F, which can solve thanks to the work of Kassabov-Matucci. We prove[19] that the algorithms we describe are correct and terminate.

Constraints on conjugators are discussed in Section 3.1; the search for coarse and fine conjugators is covered in Sections 3.3 and 3.4, respectively.

[19] Algorithms 3.3.5 and 3.4.8.

Chapter 4 begins again with a summary of the literature. The main source for centralisers in T is a chapter of Matucci's thesis, representing joint work with Bleak-Kassabov-Matucci. Before diving into that, we review the structure of element centralisers in F, based on the article of Kassabov-Matucci. As hinted at in earlier chapters, we are able to use their results to learn about the centralisers of elements of T which have fixed points. This is achieved using conjugation by a rotation, which usually makes an element of T look like an element of F.

This review takes place in Section 4.1.

The chapter of Matucci's thesis mentioned above studies centralisers in T, and related groups of circle homeomorphisms. With his collaborators, he expresses the centraliser of an element as one of two group extensions, depending on whether the element in question has

finite or infinite order. We focus on the infinite order case, reproducing Matucci's proof which identifies the extension's kernel and quotient. He shows that the kernel is contained in an F-type element centraliser; we provide[20] a missing detail which shows that this containment is an equality. Matucci et al also observe that its quotient is finite cyclic. We point out[21] that we can compute the order of this quotient, by using our solution to the conjugacy search problem.

This material features in Section 4.2.

[20] Remark 4.2.2.

[21] Algorithm 4.2.3.

Finally, Chapter 5 identifies the last piece needed to understand these extensions: knowledge of how the quotient conjugates the kernel. We consider a series of examples[22] which give us a flavour of what to expect. In particular, we exhibit an example[23] whose centraliser extension does not split. Next, we begin a general analysis of $C_T(\alpha)$ (for α nontorsion) by breaking the problem down into four separate cases. Our cases are distinguished based upon the existence of dyadic important points and whether or not the rotation number is zero. Having a dyadic important point makes things easier; without it, we need to constantly check that we are not constructing conjugators possessing nondyadic breakpoints. Life is also easier with rotation number zero, since we can appeal to results in F; without it, orbits of points wrap multiple times around the circle, making it harder to see how centralisers work.

[22] Section 5.1

[23] Example 5.1.6, with some calculations deferred to Appendices A and B.

Each case follows a similar routine:[24] we define a *block form* for α, and show that this lets us easily write down centralising elements with small rotation number $\rho > 0$. Then we explain how to find a *maximal block form*, which produces centralising elements with $\rho > 0$ minimal. From maximal block forms, we can directly observe the extension structure of $C_T(\alpha)$. All in all, we establish Theorem B after handling each case in turn. To complete the story, we explain in each case how to build a centraliser with given parameters.

[24] The four cases make up Section 5.2, and are summarised in Section 5.3.

* * *

Future work. Kassabov and Matucci execute a four-step plan[25] to solve the *simultaneous* conjugacy problem in F (and F-like groups). Together, their first two steps solve the conjugacy problem, whereas their third step involves understanding the intersection of element centralisers. The results of this thesis establish the first two steps, and hopefully make a contribution towards the third step. In the future, perhaps it will be possible to continue generalising this plan to T.

[25] See the introduction to Chapter 3.

While we do not discuss it in this thesis, Matucci's thesis also gives the centraliser of a torsion element as a group extension. Martínez-Pérez and Nucinkis generalise[26] this to the groups $T_r(\Sigma)$ parameterised by a 'Cantor algebra' Σ. It would be interesting to see how much of the nontorsion extension structure found here carries through to this more general setting.

[26] Martínez-Pérez and Nucinkis 2013.

Notation

We'll constantly be working with the action of a group G on a space X. To distinguish these two types of object, we'll use Greek letters $\alpha, \beta, \gamma, \ldots \in G$ to denote group elements[27] and Roman letters $a, b, c, \ldots \in X$ to denote points in the space.

[27] We make an exception for the usual generators x_0 and x_1 of F.

The cyclic groups of finite order n and infinite order are denoted by $\mathbb{Z}_n$ and $\mathbb{Z}$, respectively. We usually think of $\mathbb{Z}_n$ as the integers $\{0, 1, \ldots, n-1\}$ under addition modulo n.

Let u be a finite word over the binary alphabet $\{\mathtt{0}, \mathtt{1}\}$. We use $[\![u]\!]$ to denote the standard dyadic interval in $I = [0,1]$ with address[28] u. Similarly, we use $[u]$ to denote the set of infinite binary strings with u as a prefix. The middle-thirds Cantor set is denoted C, and is homeomorphic to the Cantor space $\partial\mathcal{T}_2$ of infinite binary strings. We typically work with the latter.

[28] For example, $[\![\mathtt{0}]\!] = [0, 1/2]$, $[\![\mathtt{00}]\!] = [0, 1/4]$ and $[\![\mathtt{001}]\!] = [1/8, 1/4]$. See Definition 1.1.4 and the discussion following it for full details.

The interior, closure and boundary of a subspace A are denoted A°, $\overline{A}$ and ∂A, respectively.

We choose to always act on the right-hand side of an expression, using juxtaposition to denote the result. This means a point x is sent by a group element α to an element $x\alpha$. Consequently, we conjugate group elements according to the rule $\alpha^\beta = \beta^{-1}\alpha\beta$. Our shorthand for 'X is acted on by α' is '$X \circlearrowleft \alpha$.'

To be consistent, we will use the same convention when dealing with ordinary functions. An input x is fed in to the left of a function α, with output $x\alpha$. This can be extended into a 'pipeline' of composition by placing additional functions $\alpha, \beta, \ldots$ to the right. Thus our law of composition reads $x(\alpha\beta) = (x\alpha)\beta$. This matches the notation Thompson and McKenzie used[29] for right-actions of permutation groups. Thompson later adopted[30] a hybrid approach: see the table below.

[29] McKenzie and Thompson 1973, above Definition 0.1.
[30] Thompson 1980, also above Definition 0.1.

Under this scheme, the left, right and two-sided derivatives of α at x are denoted $x^-\alpha'$, $x^+\alpha'$ and $x\alpha'$, respectively. To clarify messier expressions involving derivatives, we place points in square brackets and functions in round brackets, using a centred dot to denote multiplication in $\mathbb{R}$. For example, the chain rule reads $x(\alpha\beta)' = x\alpha' \cdot [x\alpha]\beta'$.

Let $\phi\colon X \to Y$ be a function. The restriction of ϕ to a subset $W \subseteq X$ is denoted $\phi|_W$. If in addition $Y = X$, i.e. if ϕ maps X to itself, then the fixed point set $\{x \in X \mid x\phi = x\}$ is denoted $\mathrm{Fix}(\phi)$.

	postcompose on	
input from	left	right
left	$(x)\alpha\beta = ([x]\beta)\alpha$	$(x)\alpha\beta = ([x]\alpha)\beta$
right	$\alpha\beta(x) = \alpha(\beta[x])$	$\alpha\beta(x) = \beta(\alpha[x])$

Four possible ways to denote the evaluation of a composition. Thompson used the lower-right option; the UK typically uses the lower-left option. We use the upper-right option, which preserves the order of function symbols, agrees with the domain → codomain notation, and matches the left-to-right order English is read in.

Summary of contributions

Let us briefly describe the material in each chapter, highlighting (to the best of our knowledge) the author's contributions.

The material in Chapter 1 is largely standard.

Chapter 2 discusses more technical details, but again is largely standard. We translate Proposition 2.1.4 to the language of the Cantor space in Lemma 2.1.5; the result is not new, but this particular phrasing is. In Section 2.2, we say nothing new whatsoever about revealing pairs. In Section 2.3 we discuss a version of important points suitable for the action of T on the circle. Our discussion in Section 2.4 of aligning points with T is new. With that said, we do make use of Kassabov–Matucci's ideas in this section; we prove our result by reducing to one of theirs.

Chapter 3 is mostly new material. We are not aware of anyone using a framework similar to ours in Section 3.1. In Sections 3.3 and 3.4, our solution to the conjugacy search problem in T is new; but we note that it is heavily inspired by Kassabov–Matucci's solution to the same problem in F.

Chapter 4 presents further background information. The material in Section 4.1 is a re-expression of Kassabov–Matucci's results on centralisers in F, or otherwise an argument implicit in Matucci's thesis. Remark 4.2.2 is new, and addresses a small gap in an argument of Matucci. Algorithm 4.2.3, and the material following it, is also new.

Chapter 5 is entirely original.

1 *The Thompson family of groups*

To get started, we introduce the extended family of Thompson's groups. We start by defining the *main trio*—namely F, T and V—and explain the 'tree pair diagram' calculus for computing in these groups. There are many groups which generalise beyond this trio, but we only introduce the family most relevant to our purposes, which relaxes the restrictions on gradients and breakpoints. Finally we define Cantor space, and explain how it can be useful when working with Thompson's groups.

1.1 The main trio of groups

The Thompson in question is Richard J. Thompson, who introduced[1] his now eponymous groups in connection with his work on associativity. It did not take long for them to become recognised as a source of unusual or interesting examples in group theory. Thompson showed that T and V are finitely-presented, infinite simple groups. While there are other examples of groups with these three properties,[2] T and V were the first such groups to be discovered.[3] Later, F and V were used by Thompson and McKenzie[4] to construct finitely-presented groups with unsolvable word problem.

[1] Thompson 1965.

[2] Burger and Mozes 1997, Unlike T and V, the groups constructed here are additionally torsion-free!

[3] Scott 1992.

[4] McKenzie and Thompson 1973.

Today these groups—and a growing family of generalisations—are studied in their own right. The first introductory paragraph in Matucci's thesis expresses why they capture researchers' interest: in particular he notes[5] that '[e]ven though the groups have a simple definition, many questions prove to be a challenge.' Brin has referred to Thompson's groups as *chameleons*, recognising the wide array of research areas in which these groups appear. In his words,[6] 'we see that Thompson's groups have the ability to be interesting objects in many settings.'

[5] Matucci 2008, p. 1.

[6] Brin 1996.

Perhaps the most hands-on definition describes each group as a collection of *piecewise-linear* (PL) functions. The standard introduction to this point of view is Cannon, Floyd and Parry's survey article.[7] Strictly speaking, we should say 'piecewise-affine' in place of piecewise-linear. However, as Kassabov and Matucci note,[8] 'this abuse of language is now common'.

[7] Cannon, W. Floyd and Parry 1996.

[8] Kassabov and Matucci 2012, footnote on p. 3.

1.1.1 *Thompson's group F*

Definition 1.1.1. A *dyadic rational* (a *dyadic* for short) is a rational number of the form $a/2^b$, for integers $a, b \in \mathbb{Z}$ with $b \geq 0$. The set of such rationals is a ring, denoted $\mathbb{Z}[1/2]$.

Express any number x as a binary number $B.b_1b_2b_3\ldots$, and let x_n be the truncation of this string to n binary places. Then $x_n \to x$; hence the dyadics are dense in the reals.

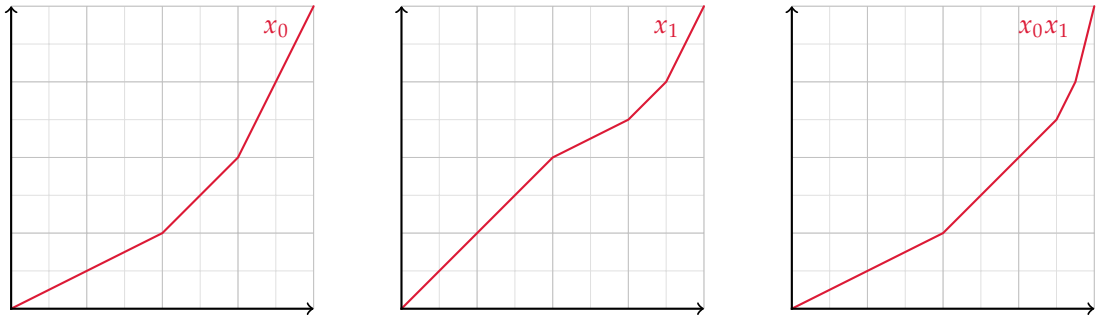

Figure 1.1: Two example elements $x_0, x_1 \in F$, together with their composition x_0x_1. We will usually use Greek letters to denote group elements, but these x_i are already named as in Cannon, Floyd and Parry.

Definition 1.1.2. Let $I = [0,1]$ be the unit interval. *Thompson's group F* is the set of PL homeomorphisms $I \to I$ with

- finitely many linear segments,
- gradients equal to integer powers of two, and with
- breakpoints (x, y) at dyadic rational coordinates.

We define the *breakpoints* of a PL function α to be the endpoints (x, y) of each linear segment of α. In an abuse of our notation, we sometimes just call x a breakpoint, since we can recover $y = x\alpha$. This convention—using x-coordinates alone—is normally used in the literature.

For example, the points $(0,0)$ and $(1,1)$ are breakpoints of every function $\alpha \in F$. In the literature, we'd just say that 0 and 1 are breakpoints of every such α.

This set forms a group under the operation of function composition, which is usually denoted by juxtaposition (c.f. the remark on page xvii).

Because each linear segment has a positive gradient, the elements of F are all increasing functions on I. In particular, every element $\alpha \in F$ fixes the minimal point 0 and fixes the maximal point 1.

We can generate F using only two elements x_0 and x_1; these are shown in Figure 1.1. The usual presentation involving this pair is[9]

$$\langle x_0, x_1 \mid [x_0x_1^{-1}, x_0^{-1}x_1x_0], [x_0x_1^{-1}, x_0^{-2}x_1x_0^2] \rangle . \tag{1.1}$$

[9] Cannon, W. Floyd and Parry 1996, Section 3.

The relators can be made simpler at the cost of adding more generators, resulting in the alternative presentation

$$\langle x_0, x_1, x_2, \ldots \mid x_k^{-1}x_nx_k = x_{n+1} \text{ for } k < n \rangle .$$

We never make use of a presentation for F: these are just included for completeness.

There are other useful presentations for F too: see Dehornoy[10] and Lodha-Moore.[11]

[10] Dehornoy 2005.

[11] Lodha and Moore 2016, Section 3.

At first glance, it seems a little awkward to perform calculations in this group. Imagine working out the composition of

$$tx_0 = \begin{cases} 0 + 1/2(t-0) & \text{if} \quad 0 \le t < 1/2 \\ 1/4 + \quad (t-1/2) & \text{if} \quad 1/2 \le t < 3/4 \\ 1/2 + 2(t-3/4) & \text{if} \quad 3/4 \le t < 1 \end{cases} \tag{1.2}$$

with

$$tx_1 = \begin{cases} 0 + 1/2(t-0) & \text{if} \quad 0 \le t < 3/4 \\ 3/8 + \quad (t-3/4) & \text{if} \quad 3/4 \le t < 7/8 \\ 1/2 + 4(t-7/8) & \text{if} \quad 7/8 \le t < 1. \end{cases} \tag{1.3}$$

The computation would be routine, but also tedious and unenlightening. To avoid this, it's worth thinking a little more carefully about how we can construct and represent an element of F.

Suppose we have a generic element $\alpha \in F$. Name the x-coordinates of its breakpoints $0 = x_0 < x_1 < \cdots < x_n = 1$. Each of these $n+1$ numbers is a dyadic rational, and the list partitions the interval into n *cells* $\{[x_i, x_{i+1}]\}_{0 \le i < n}$. Call such a list a *dyadic partition* of the unit interval. The breakpoints' y-coordinates give us a second partition $0 = y_0 < y_1 < \cdots < y_n = 1$. Because both lists have the same size, we can completely describe α by specifying that it is the unique function mapping $[x_i, x_{i+1}]$ linearly and increasingly to $[y_i, y_{i+1}]$.

Do all such pairs of partitions correspond to an element of F? The answer is *no*. The lists $x_i\colon 0 < 1/2 < 1$ and $y_i\colon 0 < 1/4 < 1$ are perfectly valid dyadic partitions with the same length, but the PL map induced by them would send $[1/2, 1] \mapsto [1/4, 1]$, with gradient $3/4 \div 1/2 = 3/2$. This is not a power of two, and so the corresponding function does not belong to F. Thus we cannot allow ourselves to use any old pair of dyadic partitions.

The solution is the notion of a *dyadic subdivision*, which is a stricter form of dyadic partition. We define this recursively by two rules:

1. The list $0 < 1$ is a dyadic subdivision of I.
2. If $x_0 < \cdots < x_i < x_{i+1} < \cdots < x_n$ is a dyadic subdivision of I, so too is $x_0 < \cdots < x_i < m < x_{i+1} < \cdots < x_n$, where $m = (x_i + x_{i+1})/2$. The resulting partition is called an *expansion* of the former, having subdivided the interval $[x_i, x_{i+1}]$.

In words: we start with the whole interval and cut it exactly in half. We are then allowed to repeatedly cut any subinterval into exactly half. These subintervals always take the form $[a/2^n, (a+1)/2^n]$ for integers $a, n \ge 0$ with $a + 1 \le 2^n$. Any interval of this form is called a *standard dyadic interval*.

We may now take two dyadic subdivisions with the same number of cells, and use the same recipe as before (ith cell mapped linearly to the ith cell) to build a PL function. Call the result of this a *dyadic rearrangement*. Because a standard dyadic interval's width is always a power of two, we will find that a dyadic rearrangement is always an element $\alpha \in F$. Conversely, given any element $\alpha \in F$, we may choose N sufficiently large such that α is linear on every standard dyadic interval of width 2^{-N}. Each of these intervals is mapped to some other standard dyadic interval, because each linear segment has gradient equal to a power of two. Hence α is a dyadic rearrangement.

Proposition 1.1.3. *F is exactly the group of dyadic rearrangements of the unit interval.*

Figure 1.2 shows an example of a *rectangle diagram*,[12] a kind of schematic for reasoning about elements $\alpha, \beta \in F$.We can compute with

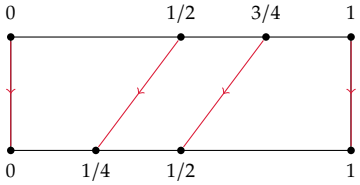

Figure 1.2: The element x_0, viewed as a dyadic rearrangement. The top and bottom edges of this rectangle are a pair of dyadic subdivisions which define x_0. We call this kind of visualisation a *rectangle diagram*.

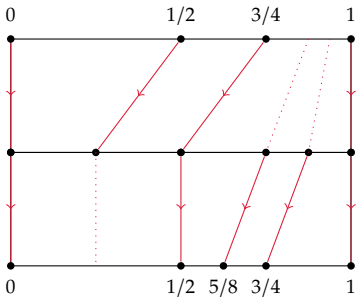

Figure 1.3: We can compose elements by stacking their rectangle diagrams. This is x_0 on top of x_1, corresponding to the product x_0x_1. The result need not immediately be a rectangle diagram. Here we had to add extra dotted lines because the interval $[3/4, 1]$ is not mapped linearly by x_0x_1; nor is $[0, 1/2]$ mapped linearly by $(x_0x_1)^{-1}$.

[12] The name 'rectangle diagrams' was used in Cannon, W. Floyd and Parry 1996, above Example 1.2; they attribute the idea to Thurston.

these diagrams: to build a rectangle diagram for $\alpha\beta$, stack a diagram for α on top of a diagram for β. To collapse the stack down to a single rectangle, we may need to cut the dyadic partitions into smaller pieces (as in Figure 1.3). This is because we need a finer dyadic subdivision to ensure that each subinterval $[u, v]$ of α's domain is sent to an interval $[u\alpha, v\alpha]$ which is mapped linearly by β. To keep things balanced, any refinements we make to α's domain partition need to be reflected by refinements to β's range partition.

1.1.2 *Tree pair diagrams*

As nice as these visualisations are, we can't hope to use them for extended calculations. If we have very fine subdivisions, we'd need to be able to draw rectangle diagrams very accurately! Is there a compromise—a graphical notation which describes dyadic partitions without literally drawing them in the plane?

Definition 1.1.4. For our purposes, a *binary tree* is a directed graph T with at least one vertex, satisfying the following properties.

Tree. T is connected and contains no (undirected) cycles.

Rooted. T has exactly one vertex with indegree 0. This vertex is called the *root* of T.

Full binary. Every vertex has outdegree 2 or 0. Such vertices are called *internal vertices* and *leaves*, respectively.

Ordered. A node's outward edges point to its *children*, which are ordered into a *left child* and a *right child*.

A small number of examples—including a *caret*—are shown in Figure 1.4.

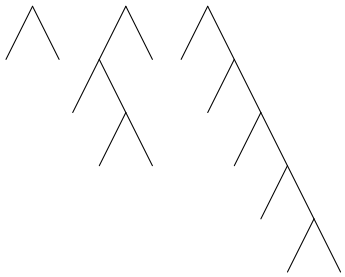

Figure 1.4: A forest of three binary trees. We always draw parent vertices above their children, and left children to the left of right children; in other words, our trees grow downwards.
Left: the second smallest possible binary tree, called a *caret*. **Right:** an unbalanced tree. This one in particular is called a *right vine*.

The infinite binary tree $\mathcal{T}_2$ (see Figure 1.5) is the union of all finite binary trees.[13] It can be defined explicitly as the graph with vertex set $\{\mathtt{0},\mathtt{1}\}^*$, that is the set of all finite strings over $\{\mathtt{0},\mathtt{1}\}$. The empty string ϵ is the root of $\mathcal{T}_2$; the left child of a string w is $w\mathtt{0}$, and the right child is $w\mathtt{1}$. The string w corresponding to a given vertex is called the *address* of the vertex. Its *length* $|w|$ is the number of binary digits used to form the string; note that $|\epsilon| = 0$.

Apart from $\mathcal{T}_2$, we will only be interested in finite binary trees (since these describe elements of F, see below). Thus when we say "binary tree", we usually mean "finite binary tree".

[13] That is, the union taken after identifying roots with roots, left children with left children and right children with right children.

We write letters of an alphabet in a `a teletype font family`. For instance, we write $|\mathtt{1}| = 1$. Variables denoting letters or strings of letters are set in serif italics as with other variables.

Figure 1.5: (An approximation of) the infinite binary tree $\mathcal{T}_2$. For each $n \geq 0$, there are exactly 2^n vertices whose address has length n.

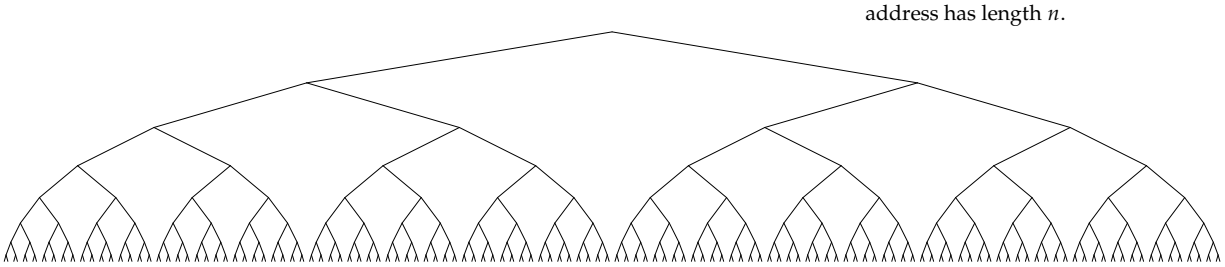

The graph structure of $\mathcal{T}_2$ is exactly the inclusion structure on the set of standard dyadic intervals. This is because we have a largest such interval, $I = [0,1]$, and any such interval $[a,b]$ splits into two smaller ones: the left child $[a,(a+b)/2]$ and the right child $[(a+b)/2,b]$. Applying this rule repeatedly, we see for example that vertices with addresses `00000` and `1001` correspond to the intervals $[0,1/32]$ and $[9/16,10/16]$, respectively. We use $[\![w]\!]$ to denote the standard dyadic interval with address w.

If T is a binary tree, we can identify the root of T with the root of $\mathcal{T}_2$. Identifying left (right) children in T with left (right) children in $\mathcal{T}_2$ allows us to see T as a subtree of $\mathcal{T}_2$, and so give addresses to T's vertices. We will usually identify a vertex with its address.

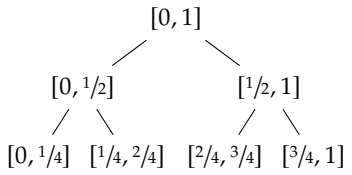

Figure 1.6: A small part of $\mathcal{T}_2$, labelled with standard dyadic intervals.

If we collect together the standard dyadic intervals associated to the leaves of a *finite* binary tree T, we recover a dyadic subdivision[14] of I. Thus we can use a pair of finite binary trees with the same number of leaves to represent an element of F.

[14] The root corresponds to the trivial subdivision $0 < 1$, and each internal vertex of T with address a describes an expansion which subdivides $[\![a]\!]$.

Definition 1.1.5. A *tree pair diagram* is a triple (D, ℓ, R). Here D and R are finite binary trees with the same number of leaves, called the *domain* and *range trees* respectively. The entry ℓ is a bijection from the leaves of D to the leaves of R.

These parameters determine an element $\alpha \in F$ (later T or V) as follows: if $d\ell = r$ then α linearly maps the standard dyadic interval $[\![d]\!]$ to the interval $[\![r]\!]$. We say that the triple (D, ℓ, R) is a *tree pair diagram for* α.

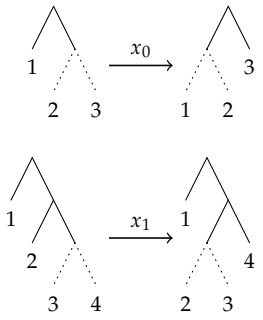

Figure 1.7: Tree pair diagrams for x_0 and x_1. To aid the eye, we use dashed lines to denote edges which belong to one of D and R, but not both.

The leaves of a binary tree are ordered by the lexicographic ordering of their addresses. For instance, the domain tree of x_0 in Figure 1.7 orders its leaves $1 < 2 < 3$ because `0` $<_{\text{lex}}$ `10` $<_{\text{lex}}$ `11`. When working in F, our group elements always preserve the order of the interval, so our leaf bijection must in turn preserve the order of the leaves. This means that ℓ is redundant as far as F is concerned—though we'll see shortly that it's needed for the supergroups T and V.

Why should we bother with tree pair diagrams? After all, they're just a graphical means to denote two dyadic subdivisions. Their usefulness becomes apparent when we need to perform computations in F. All the computations needed to compose and invert PL maps can be replaced by simpler (and non-numeric) operations on tree pairs.

TREE PAIR DIAGRAMS FOR A GIVEN ELEMENT ARE NOT UNIQUE. Suppose (D, ℓ, R) is a tree pair for α, and pick any leaf d of D. Let D' be the result of attaching a caret to D at d; similarly let R' be the result of attaching a caret to $d'\ell$. Then (D', ℓ', R') is also a tree pair for α. Here ℓ' sends the left (right) child of d to the left (right) child of $d\ell$, and otherwise behaves as ℓ does. This process of attaching two carets is called *expanding* the tree pair; we say that (D', ℓ', R') is an *expansion* of (D, ℓ, R). The reversal of an expansion is called a *contraction*. If a tree pair cannot be contracted, it is said to be *reduced*.

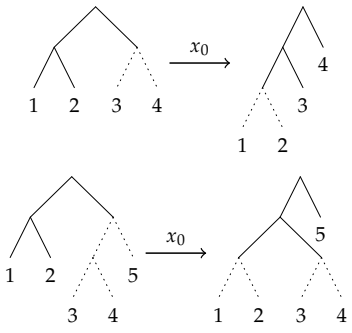

Figure 1.8: Two alternative tree pairs for x_0; compare to Figure 1.7.

Proposition 1.1.6. *Every element of F has a unique* reduced *tree pair diagram, formed by repeatedly contracting until contraction is no longer possible.*[15] *Moreover, different reduced tree pair diagrams describe different elements of F, so F is in bijective correspondence with the set of reduced tree pair diagrams.*

[15] To prove the first sentence, use Newman's Diamond Lemma, as in e.g. Baader and Nipkow 2008, Lemma 2.7.2.

While the uniqueness of reduced tree pairs is useful, being reduced is often too restrictive for our needs. Often a larger tree pair gives a more informative description of a given element in T. This is made precise using Brin's notion of a *revealing pair*, which we discuss briefly in Section 2.2. For now, we'll see that being able to expand and contract lets us perform computations in F.

Suppose we are given two tree pairs (A, ℓ, B) representing $\alpha \in F$ and (C, m, D) representing $\beta \in F$. How can we find a tree pair diagram for the product $\alpha\beta$? We can always expand to new tree pairs (A', ℓ', B') and (C', m', D') such that $B' = C'$. Then their product is given by removing the intermediate trees and composing leaf bijections. In symbols, multiplication is given by

$$(A', \ell', B') \cdot (B', m', D') = (A', \ell' m', D') .$$

This expansion corresponds to subdividing dyadic partitions. Expanding so that $B' = C'$ ensures that $\alpha\beta$ acts linearly on the intervals corresponding to the leaves of A'.

The smallest choice for $B' = C'$ is $B \cup C$; we formally define the union of binary trees in Definition 2.2.1.

We illustrate this with an example, using the tree pairs in Figure 1.7, to compute a tree pair for the product $x_0 x_1$. To do so, we need to expand our input tree pairs to larger pairs with five leaves.

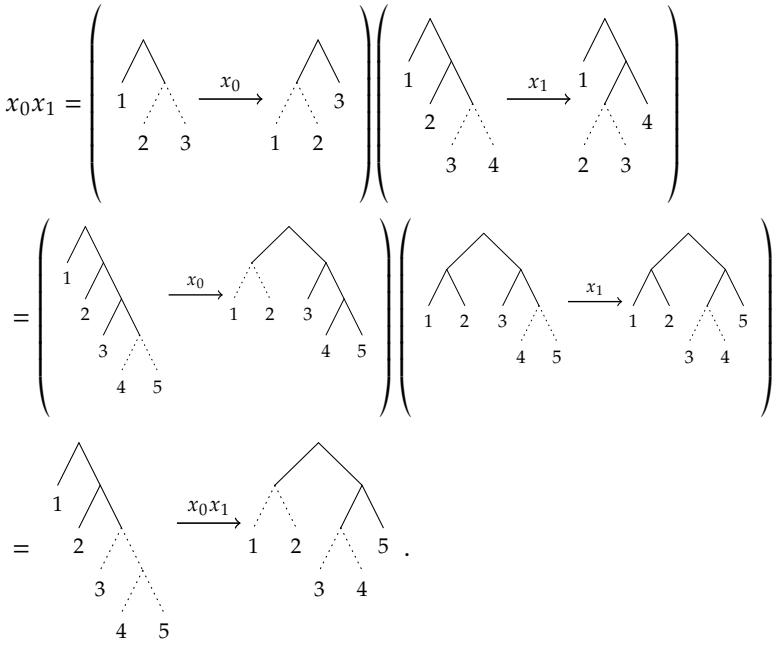

Because the tree pairs for the identity element all take the form (X, id, X), it follows that inverses are given by

$$(A, \ell, B)^{-1} = (B, \ell^{-1}, A) .$$

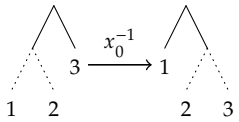

Figure 1.9: A tree pair diagram for x_0^{-1}. Effectively, we just took the diagram in Figure 1.7 and reversed the arrow. Strictly speaking, we also inverted the leaf bijection—though this isn't immediately obvious since $\mathrm{id}^{-1} = \mathrm{id}$.

Note that we can invert a tree pair directly, whereas multiplication is only defined up to equivalence tree pairs.

1.1.3 *Thompson's group T*

Our next group T is defined as a group of maps of the circle S^1, rather than the interval I. We think of the circle as the quotient space $[0,1]/\{0 \sim 1\}$, or alternatively as the quotient $\mathbb{R}/\mathbb{Z}$. This allows us to use 0 as a fixed origin for the circle.

We will need an orientation of the circle, to generalise the orientation of I defined by the binary relation $<$. To define a cyclic order systematically, we'll need to use a *ternary* relation.[16]

[16] See e.g. Tararin 2001 for more details.

Definition 1.1.7 (Cyclic order)**.** Let X be a set. A *ternary relation R on X* is a subset $R \subseteq X^3$. We write $x \to y \to z$ to mean that $(x,y,z) \in R$. Such a relation is called a *cyclic order* on X if it satisfies the following axioms, for any elements $u, x, y, z \in X$.

Cyclicity. If $x \to y \to z$, then $y \to z \to x$.

Antisymmetry. If $x \to y \to z$, then it is *not* true that $z \to y \to x$.

Transitivity. If $x \to y \to z$ and $x \to z \to u$, then $x \to y \to u$.

Irreflexivity. If $x \to y \to z$, then x, y, z are pairwise distinct.

Totality. If x, y, z are pairwise distinct, then either $x \to y \to z$ or $z \to y \to x$.

If these axioms are satisfied, whenever $x \to y \to z$ we say that the list of points x, y, z is *in cyclic order*.

[17] We suggest reading this relation as 'y is between x and z'.

Notation There doesn't seem to be a universally agreed-upon notation for expressing that three points are in cyclic order. Some authors write $(x,y,z) \in R$ explicitly; others write $R(x,y,z)$ or use juxtaposition to say xyz; yet another alternative is $[x,y,z]$ or (x,y,z). As mentioned above, we will use the notation[17] $x \to y \to z$. This extends to the shorthand $x_0 \to \cdots \to x_{n-1}$, which means that that $x_i \to x_{i+1} \to x_{i+2}$ for every $0 \le i < n-2$. If $n \in \{1,2\}$ then the shorthand has no meaning: in these circumstances $x_0 \to \cdots \to x_{n-1}$ is a vacuously true statement.

We will often use make use of a list of points $x_0, \ldots, x_{n-1}$ for which $x_0 \to \cdots \to x_{n-1} \to x_0$ (even in the cases[18] $n = 1$ or 2). wrapping back around to x_1 ensures that our list of points winds around the circle exactly once (c.f. Figure 1.10).

[18] Read literally, $x_0 \to \cdots \to x_{n-1} \to x_0$ means

$x_0 \to x_0$ when $n = 1$, and
$x_0 \to x_1 \to x_0$ when $n = 2$.

The former is meaningless; the latter is always false, by the irreflexivity axiom.

To get around this, when $n = 1$ or 2 we define the shorthand $x_0 \to \cdots \to x_{n-1} \to x_0$ to be vacuously true. To be explicit, this applies when there is at least one explicit arrow after the $\to \cdots \to$ shorthand, and when the first and last points (x_0 here) are equal.

We define open, closed and half-open arcs from x to $z \neq x$ by the formulae

$$(x,z) = \{\, y \in X \mid x \to y \to z \,\} \qquad [x,z] = (x,z) \cup \{x, z\}$$
$$[x,z) = (x,z) \cup \{x\} \qquad (x,z] = (x,z) \cup \{z\}.$$

Note that we don't define arcs (x,x) or $[x,x]$ from and to the same point.

Notice for instance that $X = (x,z) \cup [z,x]$, due to the totality axiom.

At first glance, the transitivity axiom doesn't look like the usual notion of transitivity for a binary relation. However, if we omit the initial "$x \to$" from each relation, the axiom reads

$$\text{if } y \to z \text{ and } z \to u, \quad \text{then } y \to u,$$

which looks reassuringly like the transitivity we are familiar with for binary relations. In this way, we can 'cut' our cyclic order at x to create a new linear order $<_x$ on $S^1 \setminus \{x\}$, defined by

$$y <_x z \quad \text{if and only if} \quad x \to y \to z\,. \tag{1.4}$$

Indeed, one way to view a cyclic order is as a family of linear orders.[19]

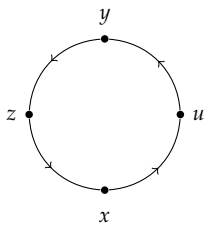

Figure 1.10: Certainly $x \to y \to z$ and $y \to z \to u$, but it is *not* true that $x \to y \to u$. In terms of our shorthand notation, we have $x \to y \to z \to u$. Problems start occurring when we travel more than once round the circumference of our circle. This is characterised by the fact that $x \to y \to z \to u \to x$ is false.

[19] Calegari 2004, Definition 2.2.1

Remark 1.1.8. We might guess that the transitivity axiom takes the form

$$\text{if } x \to y \to z \text{ and } y \to z \to u, \quad \text{then } x \to y \to u.$$

This is *not* equivalent to the transitivity axiom in Definition 1.1.7, and it is not a correct model of circular order—see Figure 1.10.

Definition 1.1.9. Let $S^1 = [0,1]/\{0 \sim 1\}$ be the circle. Also let $R_0 = \{(x,y,z) \in [0,1) \mid x < y < z\}$, where $<$ is the usual order on $[0,1)$. We define a cyclic order R on the circle by the formula

$$R = R_0 \cup \{(y,z,x) \mid (x,y,z) \in R_0\} \cup \{(z,x,y) \mid (x,y,z) \in R_0\}\,.$$

R is called the *cyclic closure* of R_0.

In our visualisations, $x \to y \to z$ will mean that the anticlockwise arc from x to z passes through y.

The circle is equipped with a metric defined on points $0 \le x, y < 1$ by $d(x,y) = \min\{|x-y|, 1-|x-y|\}$. For instance, $d(1/3, 2/3) = \min\{1/3, 2/3\} = 1/3$ and $d(8/9, 1/9) = \min\{7/9, 29\} = 2/9$. Addition and subtraction on the circle are always evaluated modulo 1. This means that $x - \epsilon \to x \to x + \epsilon$ always holds true, for any point x and distance $\epsilon < 1/2$.

With that out of the way, we can make a concrete definition of T.

Definition 1.1.10. Let $S^1 = [0,1]/\{0 \sim 1\}$ be the circle. *Thompson's group* T is the set of orientation-preserving[20] PL homeomorphisms $S^1 \to S^1$ with

- finitely many linear segments,
- gradients equal to integer powers of two, and with
- breakpoints at dyadic rational coordinates.

This set forms a group under function composition.

Just like in F, a given element $\alpha \in T$ will map each standard dyadic interval of width 2^{-N} linearly, provided N is sufficiently large. In

[20] On a linearly ordered set like I, 'orientation-preserving' meant that $x < y \implies x\alpha < y\alpha$. With a cyclic order, 'orientation-preserving' means that $x \to y \to z \implies x\alpha \to y\alpha \to z\alpha$.

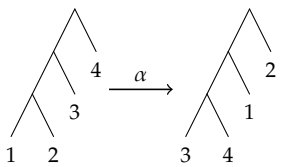

Figure 1.11: An example element $\alpha \in T$. Because the domain and range trees are identical, α must have finite order dividing the number of leaves.

other words, the linear segments can be still be described by a pair of dyadic subdivisions of the same size. What's new is that elements of T no longer have to fix the point 0, which is always fixed by every element of F. Indeed, elements of T need not fix any part of the circle whatsoever—see Figure 1.14 for instance. This allows us to have nontrivial torsion elements[21] in T; in contrast, there are no such elements in F (c.f. Equation (1.5) and Proposition 2.3.19, respectively).

Say the domain and range of $\alpha \in T$ is partitioned into standard dyadic intervals $D_0, \dots, D_{n-1}$ and $R_0, \dots, R_{n-1}$. One extra piece of information is required: we need to specify how the $\{D_i\}$ are mapped to the $\{R_i\}$. As elements of T are continuous functions preserving cyclic order, we need only choose a single domain interval D_i and specify its image $D_i\alpha = R_j$. Then continuity leads us to conclude that $D_{i+k}\alpha = R_{j+k}$ for each k, with subscripts modulo n.

To represent this on a tree pair diagram, we make use of the leaf bijection ℓ in Definition 1.1.5. The range tree's leaves no longer have to be labelled $1, \dots, n$ from left to right; instead, any cyclic permutation of these numbers is permitted. Thus we have to take a little more care when multiplying, to make sure we correctly compose the leaf bijections. For instance, if α is the element shown in Figure 1.11, then

$$\alpha^2 = \left(\begin{array}{c} 4 \\ 3 \\ 1\ 2 \end{array} \xrightarrow{\alpha} \begin{array}{c} 2 \\ 1 \\ 3\ 4 \end{array} \right) \left(\begin{array}{c} 4 \\ 3 \\ 1\ 2 \end{array} \xrightarrow{\alpha} \begin{array}{c} 2 \\ 1 \\ 3\ 4 \end{array} \right) \tag{1.5}$$

$$= \begin{array}{c} 4 \\ 3 \\ 1\ 2 \end{array} \xrightarrow{\alpha^2} \begin{array}{c} 4 \\ 3 \\ 1\ 2 \end{array} = \mathrm{id}\ .$$

Remark 1.1.11. As currently defined, the elements of F are not literally elements of T (since functions $I \to I$ are not literally functions $S^1 \to S^1$). How do we settle this against our earlier claims that $F < T$? We have a number of different options.

1. Let π be the quotient map $I \to S^1$ which identifies 0 with 1. Any element $\alpha\colon I \to I$ in F must fix the point 0 and must fix the point 1. This ensures that there is a continuous map $\alpha_0\colon S^1 \to S^1$ satisfying $\pi\alpha_0 = \alpha\pi$. Now $\alpha_0 \in T$ because α_0 inherits the breakpoint and gradient properties of α.

   Because the process of replacing $\alpha \mapsto \alpha_0$ is reversible[22] and respects composition, it describes an embedding $F \hookrightarrow T$. Informally, we make F look like T by bending the interval into a circle—see Figure 1.13. In doing so, F becomes the subgroup of T whose elements

[21] We can construct an element $\alpha \in T$ with any given order $n \in \mathbb{N}$ as follows. Pick a tree S with n leaves, and use this as the domain and range tree of α. Then define α to map the ith leaf of S to the next leaf in the circular ordering.

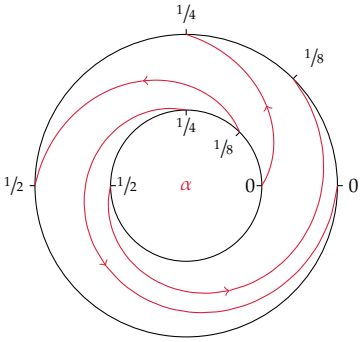

Figure 1.12: To adapt the rectangle diagrams introduced in Figure 1.2 to T, we need replace the the top and bottom intervals by circles. We've done so here using an annulus. The element visualised here is α from Figure 1.11.

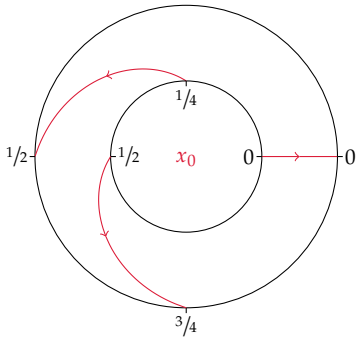

Figure 1.13: The element $x_0 \in F$ displayed as a circle homeomorphism fixing 0.

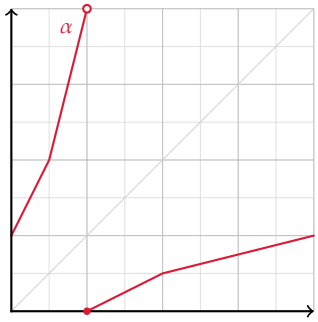

Figure 1.14: The element $\alpha \in T$ from Figures 1.11 and 1.12, displayed as a PL map on $[0, 1)$ with a discontinuity. As the graph never crosses the diagonal $y = x$, α has no fixed points.

[22] Given α_0, define α by setting $t\alpha = t\alpha_0$ for $0 \le t < 1$ and $1\alpha = 1$.

fix 0.

2. Alternatively, we could make T look more like F. Because we view the circle S^1 as $[0,1]/\{0 \sim 1\}$, we can see a function $\alpha_0 \in T$ as a function $\alpha\colon [0,1) \to [0,1)$. If $0\alpha_0 \in 0$, then this function α is continuous (and belongs to the restriction of F to $[0,1)$. Otherwise $0\alpha_0 \neq 0$, which means α will contain a discontinuity at the point[23] where α_0 wraps around from 1 to 0.

   All in all, we can see T as a group of PL bijections $[0,1) \to [0,1)$ which are right-continuous (see Figure 1.14). Since any element of F can be recovered from its restriction to $[0,1)$, we obtain an embedding $F \hookrightarrow T$.

3. A third option is to regard F and T as being groups of (equivalence classes of) tree pair diagrams. Then F is literally a subgroup of T.

4. In Section 1.2.3 we introduce *Cantor space*, and in particular the Cantor space $\partial\mathcal{T}_2$ consisting of infinite binary strings. We can view F, T and V (the latter defined below) as groups of homeomorphisms of Cantor space; this is arguably the most natural way to think about Thompson's groups. From this point of view, the subgroup relationships $F < T < V$ are immediate. Even better, we don't need to spend effort worrying about the difference between $[0,1]$, $[0,1)$ and S^1: we have just one ambient space permuted by the main trio of Thompson's groups.

[23] Specifically, as we approach $(0\alpha_0^{-1}, 1)$ from the left, we jump down to $(0\alpha_0^{-1}, 0)$.

Remark 1.1.12. We can augment Presentation (1.1) to form a presentation for T. In addition to x_0 and x_1, we require only one new generator c which allows us to cyclically permute leaves. Cannon-Floyd-Parry show[24] that four extra relations (in addition to the two between the the x_i) are sufficient to build a presentation of T.

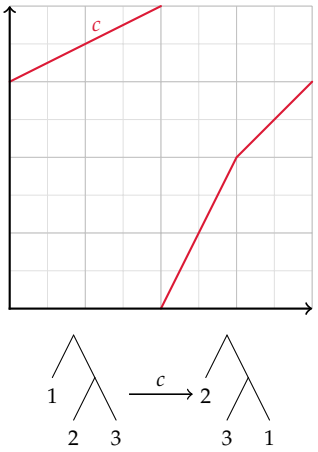

Figure 1.15: Cannon-Floyd-Parry's generator c for T. Despite being only of order three, c can be used in conjunction with x_0 and x_1 to produce any element of T (which can have arbitrarily large order).

[24] Cannon, W. Floyd and Parry 1996, Section 5.

1.1.4 *Thompson's group V*

The largest of Thompson's main trio can be again defined as a group of PL functions on the interval. We take the viewpoint of Item 2 of Remark 1.1.11, thinking of our elements as right-continuous functions α to and from $[0,1)$. This allows discontinuities to exist for α (and not just its derivative like in F).

Definition 1.1.13. *Thompson's group V* is the set of right-continuous PL bijections $[0,1) \to [0,1)$ with

- finitely many linear segments,[25]
- gradients equal to integer powers of two, and with
- breakpoints at dyadic rational coordinates.

This set forms a group under function composition.

[25] Strictly speaking, the linear segments are now closed on the left and open on the right.

Unlike in T, elements of V can have arbitrarily many discontinuities—compare Figures 1.14 and 1.16. Thus if we have two dyadic subdivisions D and R with the same number of cells n, there are no restrictions on how V must send D's cells to R's cells—any of the $n!$ matchings describe a valid element of V. The tree pair approach to V is perfectly valid too—to accommodate V, we just allow an arbitrary bijection ℓ between the leaves in a diagram (D, ℓ, R).

Because V allows us total freedom to rearrange tree pair leaves (or to rearrange subdivisions, if you prefer), its elements exhibit a vast variety of behaviours. To see this, let G be any finite group, viewed as a permutation group on n elements $\{1, 2, \ldots, n\}$. Choose a binary tree D with n leaves, and number the leaves $1, \ldots, n$. Then the set of tree pair diagrams $\{(D, \ell, D) \mid \ell \in G\}$ is a subgroup of V isomorphic to G. Thus every finite group embeds into V! In contrast, the finite subgroups of T are exactly the cyclic groups.[26]

To identify and control the new behaviour possible in V, it is useful to work with carefully chosen tree pair diagrams called *revealing pairs*—see Section 2.2.

We can further augment Cannon-Floyd-Parry's presentation for T (Remark 1.1.12) to form a presentation for V, again using only one new generator π. In addition to CFP's six relations, we need to add a further eight relations[27] to account for the introduction of π. More recently, Bleak and Quick have produced[28] a 'human-interpretable' presentation with 3 generators and 8 relations; Tietze transformations reduce this to 2 generators and 7 relations.

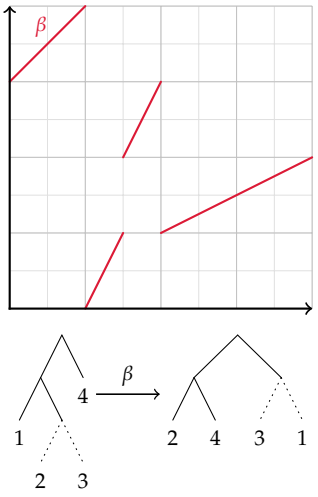

Figure 1.16: An example element $\beta \in V$, which turns out to have order 5. We have not drawn the discontinuity markers for the sake of clarity.

[26] Geoghegan and Varisco 2017, Theorem 3.1.

[27] Cannon, W. Floyd and Parry 1996, Section 6.
[28] Bleak and Quick 2017.

1.2 Other points of view

There are many different tools and languages we can use to study Thompson's family of groups. (Indeed, this is partly why these groups interest so many researchers!) In this section we introduce other ways to think about the main trio, and some generalised Thompson groups which are will be relevant later.

1.2.1 *Generalised gradients and breakpoints*

We can tweak the definition of F to define a similar group with different choice of gradients and breakpoints.

Definition 1.2.1. Let S be an additive subgroup of $\mathbb{R}$ and let G be a subgroup of the multiplicative group $\{g \in \mathbb{R}_{>0} \mid Sg = S\}$. We define the group $\mathrm{PL}_{S,G}(I)$ to be the group of PL homeomorphisms $I \to I$ with

- finitely many linear segments,
- gradients in G, and with
- breakpoints' coordinates in $S \times S$.

As usual, I denotes the unit interval $[0, 1]$.

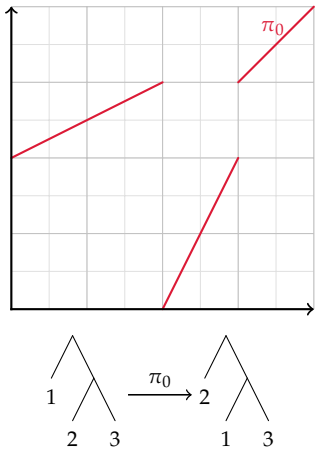

Figure 1.17: Cannon-Floyd-Parry's generator π_0 for V. This provides a non-cyclic leaf permutation, which can be used together with x_0, x_1 and c to produce any element of V.

We use the notation of Kassabov-Matucci.[29] For completeness, we note that $\mathrm{PL}_{S,G}(I)$ is called $G(I;A;p)$ by Bieri-Strebel[30] and $F(\ell,A,P)$ by Stein.[31] Our S (resp. G and I) corresponds to Bieri-Strebel's A (resp. P and I) and Stein's A (resp. P and ℓ).

There are three important choices of gradient and breakpoints to highlight. Taking $S = \mathbb{Z}[1/2]$ and $G = 2^{\mathbb{Z}} = \{2^n \mid n \in \mathbb{Z}\}$ yields Thompson's group F. We write PL_2 as shorthand for $\mathrm{PL}_{\mathbb{Z}[1/2],2^{\mathbb{Z}}}$. On the other hand, taking $S = \mathbb{R}$ and $G = \mathbb{R}_{>0}$ yields the largest group of this type, with no restrictions on gradients or breakpoints. We write PL_+ as shorthand for $\mathrm{PL}_{\mathbb{R},\mathbb{R}_{>0}}$. If is often convenient to work in this group when we want to reason about the PL nature of the functions being studied, and not the distinction F makes between dyadics and nondyadics.

In the middle lies the choice $S = \mathbb{Q}$ and $G = \mathbb{Q}_{>0}$, which we denote as $\mathrm{PL}_{\mathbb{Q}}$ for short. We can think of $\mathrm{PL}_{\mathbb{Q}}$ as an approximation to PL_+ in which computation is feasible, since computers can work with rationals, but not reals. All of these sit inside the full group $\mathrm{Homeo}_+(I)$ of increasing homeomorphisms of the interval.

By changing these functions to be orientation-preserving maps of the circle, we get a family $\mathrm{PL}_{S,G}(S^1)$ of PL homeomorphisms of S^1. Important examples are $T = \mathrm{PL}_2(S^1)$ and $\mathrm{PL}_+(S^1)$; again these sit inside the full group $\mathrm{Homeo}_+(S^1)$ of orientation-preserving homeomorphisms. Similarly, using right-continuous functions on $[0,1)$ yields a V-like family of groups.

[29] Kassabov and Matucci 2012, Section 2.1.
[30] Bieri and Strebel 2016.
[31] Stein 1992.

If we have a compact interval $J \subseteq \mathbb{R}$ with endpoints in S, we define $\mathrm{PL}_{S,G}(J)$ to be the group of PL homeomorphisms $J \to J$ subject to the same restrictions as in Definition 1.2.1. Let J_1 and J_2 be two such intervals. If the symbol $\mathrm{PL}_{S,G}$ is one of the three cases PL_2, $\mathrm{PL}_{\mathbb{Q}}$ or PL_+ above, then the groups $\mathrm{PL}_{S,G}(J_1)$ and $\mathrm{PL}_{S,G}(J_2)$ are isomorphic. This is not true[32] for general parameters S and G.

There are two natural definitions[33] of $\mathrm{PL}_{S,G}(J)$ when J has endpoints not in S.

[32] Kassabov and Matucci 2012, Remark 9.5
[33] Kassabov and Matucci 2012, Remark 2.2. Our group $\mathrm{PL}^{\mathrm{flat}}$ is their group $\mathrm{PL}^{\mathrm{Fix}=I\setminus J}$.

Definition 1.2.2. Let S and G be as in Definition 1.2.1. Also let $J = [j_1, j_2]$ and K be compact real intervals with $J \subseteq K$ and $\partial K \in S$. The groups $\mathrm{PL}^{\mathrm{rest}}_{S,G}(J)$ and $\mathrm{PL}^{\mathrm{flat}}_{S,G}(J)$ are defined to be

$$\mathrm{PL}^{\mathrm{rest}}_{S,G}(J) = \{\alpha\big|_J \mid \alpha \in \mathrm{PL}_{S,G}(K) \text{ fixes } j_1 \text{ and } j_2\}$$
$$\text{and} \quad \mathrm{PL}^{\mathrm{flat}}_{S,G}(J) = \{\alpha \in \mathrm{PL}_{S,G}(K) \mid t\alpha = t \text{ for all } t \in K \setminus J\}.$$

Any element of $\mathrm{PL}^{\mathrm{flat}}_{S,G}(J)$ can be uniquely recovered from its restriction to J, so we can think of this group as containing functions $J \to J$. If we do so, then $\mathrm{PL}^{\mathrm{flat}}_{S,G}(J)$ is a subgroup of $\mathrm{PL}^{\mathrm{rest}}_{S,G}(J)$.

What's the difference between these groups? If $\partial J \subseteq S$ then these groups are identical to $\mathrm{PL}_{S,G}(J)$. For each α in $\mathrm{PL}^{\mathrm{flat}}_{S,G}(J)$, the left gradient $j_1^-\alpha'$ and right gradient $j_2^+\alpha'$ at J's endpoints must be equal to 1, since

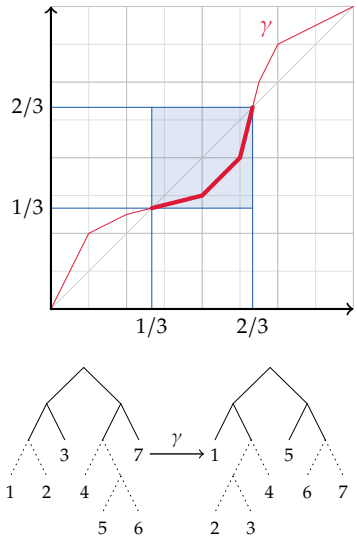

Figure 1.18: Let $J = [1/3, 2/3]$ and $K = I$. The element $\gamma \in F$ is such that the **restriction of γ to J** is in $\mathrm{PL}^{\mathrm{rest}}_2(J)$, but γ is not in $\mathrm{PL}^{\mathrm{flat}}_2(J)$. ('Rest' stands for 'restriction'.)

beyond these points α is the identity. If one of J's endpoints, say j_1, is not in S, then we must have an internal gradient $j_1^+\alpha'$ also equal to 1. In this case, the initial gradient of $\alpha\big|_J$ cannot be steep (> 1) or shallow (< 1); hence the term 'flat'. In contrast, there is no such restriction on $j_1^+\alpha'$ (or $j_2^-\alpha'$) if $\alpha \in \mathrm{PL}^{\mathrm{rest}}_{S,G}(J)$, allowing the latter group to contain more elements—see Figure 1.18.

In short, we have $\mathrm{PL}^{\mathrm{flat}}_{S,G}(J) \leq \mathrm{PL}^{\mathrm{rest}}_{S,G}(J)$. Equality holds if $\partial J \subseteq S$; otherwise the equality might not hold.

We will want to work with the fullest possible set of PL functions later, so we make the convention that $\mathrm{PL}_{S,G}$ always refers to $\mathrm{PL}^{\mathrm{rest}}_{S,G}$.

Finally, Kassabov and Matucci define[34] three subsets of $\mathrm{PL}_{S,G}(J)$. Their dynamics are relatively straightforward, so functions in these sets will be important objects of study in later chapters.

[34] Kassabov and Matucci 2012, Definitions 2.3–2.4.

Definition 1.2.3. Let $J = [j_1, j_2]$ be a real interval. We define two sets (in fact semigroups) of functions

$$\mathrm{PL}^{<}_{S,G}(J) = \{\, \alpha \in \mathrm{PL}_{S,G}(J) \mid t\alpha < t \text{ for all } j_1 < t < j_2 \,\}$$
$$\text{and} \quad \mathrm{PL}^{>}_{S,G}(J) = \{\, \alpha \in \mathrm{PL}_{S,G}(J) \mid t\alpha > t \text{ for all } j_1 < t < j_2 \,\}$$

whose graphs are strictly below (above) the diagonal. Elements in either set are called *one-bump functions*. One-bump functions also belong to a larger set

$$\mathrm{PL}^{0}_{S,G}(J) = \{\, \alpha \in \mathrm{PL}_{S,G}(J) \mid \text{if } t\alpha = t \text{ for some } j_1 < t < j_2, \text{ then } t \notin S \,\}$$

whose elements are known as *almost one-bump functions*. In this larger set, elements' graphs may cross the diagonal—but if so, their gradients cannot change while crossing. Notice that $\mathrm{id} \notin \mathrm{PL}^{0}_{S,G}(J)$.

For instance, x_0 in Figure 1.1 is a one-bump function below the diagonal; its inverse x_0^{-1} is the same, but with graph above the diagonal. In the same figure, x_1 is *not* a one-bump function.

See γ in Figure 1.18 for an example of an almost one-bump function on $[0, 1]$ which is not one-bump.

Almost one-bump functions are a useful collection of 'building blocks', because any function in $\mathrm{PL}_{S,G}(J)$ can be formed by gluing together almost one-bump functions and copies of the identity. This allows Kassabov and Matucci to reduce much of their analysis in $\mathrm{PL}_{S,G}(J)$ to only consider almost one-bump functions—a strategy we will use too.

1.2.2 *More piecewise-linear maps*

So far, we have only spoken about piecewise linear maps to and from the same interval, say X. Later we will want to consider PL maps on another interval Y, and it will be useful to have a way of producing maps on Y from maps on X.

Definition 1.2.4. Let X and Y be compact real intervals. We define $\mathrm{PL}_2(X, Y)$ to be the set of PL homeomorphisms $X \to Y$ with

- finitely many linear segments,
- gradients equal to integer powers of two; and with

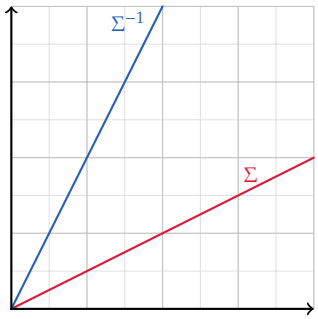

Figure 1.19: The simplest PL_2 map $\Sigma\colon [0,1] \to [0,1/2]$ has only one linear segment. The same is true of its inverse Σ^{-1}.

Notice that we do not require the endpoints (or 'exterior breakpoints') ∂X or ∂Y to be dyadic.

- interior breakpoints at dyadic rational coordinates.

Maps of this form are called PL_2 maps, or sometimes said to be *Thompson-like*. The set of such maps is written $\mathrm{PL}_2(X, Y)$. If we take $Y = X$, we obtain the group $\mathrm{PL}_2(X, X) = \mathrm{PL}_2(X)$. If ∂X is not dyadic, this is the larger group $\mathrm{PL}_2^{\mathrm{rest}}(X)$ rather than the smaller $\mathrm{PL}_2^{\mathrm{flat}}(X)$.

Claim 1.2.5. *The category $\mathcal{PL}_2$ defined below is a groupoid.*

- *The category's objects are compact intervals $X, Y \subseteq \mathbb{R}$.*
- *The morphism set between X and Y is $\mathrm{PL}_2(X, Y)$.*
- *Morphisms $f\colon X \to Y$ and $g\colon Y \to Z$ are composed using function composition.*

To establish this we'll need to consider the gradient of a composition of functions; we'll do so by making use of the *chain rule* from single variable calculus. Except for the identity, the PL functions f we're considering have at least one breakpoint x. At such points, the left- and right-derivatives x^+f' and x^-f' exist but are not equal, so the two-sided derivative xf' will not exist. Thus we need to use a one-sided version of the chain rule.

Proposition 1.2.6 (One-sided chain rules)**.** *Let $f\colon A \to B$ and $g\colon B \to \mathbb{R}$ be functions defined on open intervals $A, B \subseteq \mathbb{R}$, and let $x \in A$ be some point.*

- *Suppose that f is increasing on some interval $[x, x+\epsilon]$. If the right-derivatives x^+f' and $[xf]^+g'$ exist, then the right-derivative $x^+(fg)'$ exists and is equal to the product $x^+f' \cdot [xf]^+g'$ of the first two derivatives.*
- *Suppose that f is increasing on some interval $[x-\epsilon, x]$. If the left-derivatives x^-f' and $[xf]^-g'$ exist, then the left-derivative $x^-(fg)'$ exists and is equal to the product $x^-f' \cdot [xf]^-g'$ of the first two derivatives.*

Note that every PL_2 map is left-differentiable everywhere except for the left endpoint of its domain. The symmetric statement obtained by swapping the roles of 'left' and 'right' is also true.

Proof of Claim 1.2.5. First we check that $\mathcal{PL}_2$ is a category. The identity morphisms id_X are just the identity functions $\mathrm{id}\colon X \to X$ in $\mathrm{PL}_2(X)$. Function composition is associative, but we should check that $f \in \mathrm{PL}_2(X, Y)$ and $g \in \mathrm{PL}_2(Y, Z)$ implies that $fg \in \mathrm{PL}_2(X, Z)$. Let $X = [x_0, x_1]$ and $Z = [z_0, z_1]$ be the domain and codomain of fg.

Gradients. Let L be any linear segment of the product fg. The domain of L contains a point $x \neq x_1$ at which the right-derivative x^+f' exists. We know then that $xf \neq x_1f$, and so the right-derivative $[xf]^+g'$ exists. These gradients are integer powers of two; in particular they are positive, so f is increasing at x. The one-sided chain rule applies, telling us that $x^+(fg)' = x^+f' \cdot [xf]^+g'$. This is a product of integer powers of two, so is itself an integer power of two. Thus all linear segments L have a permissible gradient.

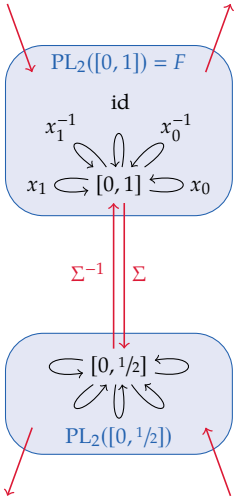

Figure 1.20: A schematic illustrating a very small part of the groupoid $\mathcal{PL}_2$. Functions with the same domain and range X, form a group $\mathrm{PL}_2(X)$. We can think of maps between different sets, say $\Sigma\colon X \to Y$, as 'converting' between $\mathrm{PL}_2(X)$ and $\mathrm{PL}_2(Y)$.

This will also help establish our non-standard notation for the derivative: a superscript + or − indicates a direction, and $\cdot$ denotes multiplication in $\mathbb{R}$.

We really are making use of the fact that f is locally increasing here. E.g. if f were is decreasing to the right of x instead, we'd have to use the rule $x^+(fg)' = x^+f' \cdot [xf]^-g'$.

In particular, this proof explains why F is a group—something we asserted without proof earlier.

Finitely many linear segments. We certainly have two 'extreme' breakpoints (x_0, z_0) and (x_1, z_1) at the endpoints of our intervals. Since f and g are continuous, all other 'internal' breakpoints of fg occur when $x^+(fg)' \neq x^-(fg)'$. Using one-sided chain rules, we conclude that $x^+f' \cdot [xf]^+g' \neq x^-f' \cdot [xf]^-g'$. This implies that $x^+f' \neq x^-f'$ or $[xf]^+g' \neq [xf]^-g'$ (possibly both). In other words, if x is a breakpoint of fg then x is a breakpoint of f or xf is a breakpoint of g.[35] This shows that there are finitely many breakpoints (hence finitely many linear segments) of fg.

[35] The converse implication is false. Take an internal breakpoint x of f and set $g = f^{-1}$. Then xf is an internal breakpoint of g, but x is not a breakpoint of $fg = \mathrm{id}_X$, since the latter has no internal breakpoints!

Internal breakpoints. As noted above, for each breakpoint (x, xfg) we know that x is a breakpoint of f or xf is a breakpoint of g. If both options hold, then the coordinates are dyadic by definition of f and g.

If only the former option holds, then x is dyadic and we know that (xf, xfg) lies within a linear segment L of g. The points of L—say (y', z')—satisfy $z' = sg + 2^i(y' - s)$, where i is an integer and (s, sg) is an endpoint of L with dyadic coordinates. Thus $xfg = sg + 2^i(xf - s)$. The right-hand side is a combination of dyadic rationals, so xfg is also a dyadic rational.

If only the latter option holds, then xf and xfg are dyadic and we know that (xf, xfg) lies within a linear segment L of f. Its points—say (x', y')—satisfy $y' = sf + 2^i(x' - s)$, where i is an integer and (s, sf) is an endpoint of L with dyadic coordinates. Thus $xf = sf + 2^i(x - s)$, which rearranges to $x = 2^{-i}(xf - sf) + s$. Again, x is a combination of dyadic rationals, so is a dyadic rational itself.

Thus fg is a PL_2 map.

To make $\mathcal{PL}_2$ a groupoid, we need an inversion map. Any PL_2 map $f\colon X \to Y$ is a homeomorphism, so is invertible. Thus our inversion map is just the ordinary inversion of functions. However, we need to check that $f^{-1}\colon Y \to X$ is PL_2. The reasoning is similar to the above. Briefly: there is a bijection between linear segments of f and f^{-1}: take the segment S of f to the segment S^{-1} of f^{-1} whose domain is the image of S. Thus f^{-1} has a finite number of segments. The gradient of S^{-1} is the reciprocal of the gradient of S, so is an integer power of two. There is also a bijection between breakpoints: (y, x) is a breakpoint of f^{-1} if and only if (x, y) is a breakpoint of f. The latter has dyadic coordinates, so the former does too. □

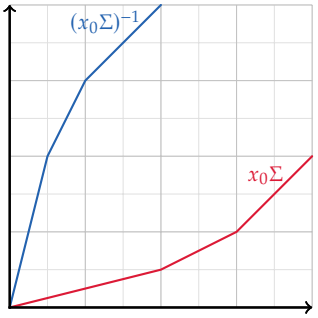

Figure 1.21: The product $x_0\Sigma$ is a more complicated PL_2 map $[0,1] \to [0,1/2]$. Its inverse map $[0,1/2] \to [0,1]$ is also shown.

In passing, we note that Matucci has spoken of a different object called *Thompson's groupoid* $\mathcal{F}$, in relation to strand diagrams.[36]

[36] Matucci 2008, p. 25.

Claim 1.2.7. *Let G be a groupoid. For any object $X \in G$, the set $G(X)$ of maps $X \to X$ is a group. Let $Y \in G$ be another object. If there is a groupoid map $\Sigma\colon X \to Y$, then $G(X)$ is isomorphic to $G(Y)$. An isomorphism $G(X) \to G(Y)$ is given by $\cdot^\Sigma\colon \alpha \mapsto \alpha^\Sigma = \Sigma^{-1}\alpha\Sigma$.*

In particular, if Σ is a PL_2 map $X \to Y$ then $\mathrm{PL}_2(X)^\Sigma = \mathrm{PL}_2(Y)$.

We call this isomorphism a *local conjugation by* Σ. We say 'local' because this conjugation only makes sense for the PL_2 maps $X \to X$; it ignores all the other maps in $\mathcal{PL}_2$.

Another way to express this claim is that vertex groups in a connected component of a groupoid are all isomorphic.[37] The *change of basepoint isomorphism*[38] is a famous instance of this result in algebraic topology.

[37] Cohen 1989, Chapter 3.

[38] Munkres 2000, Sections 51–52.

Proof. Pick group elements $\alpha, \beta \in G(X)$. The conjugation described above is a homomorphism, because $(\Sigma^{-1}\alpha\Sigma)(\Sigma^{-1}\beta\Sigma) = \Sigma^{-1}(\alpha\beta)\Sigma$. The kernel is trivial, because $\Sigma^{-1}\alpha\Sigma = \mathrm{id}_Y$ implies $\alpha = \Sigma\,\mathrm{id}_Y\,\Sigma^{-1} = \Sigma\Sigma^{-1} = \mathrm{id}_X$. The map is surjective, because $\gamma \in G(Y)$ is the image of $\Sigma\gamma\Sigma^{-1} \in G(X)$. □

Claim 1.2.8. *Let X and Y be compact real intervals with dyadic endpoints. Then $\mathrm{PL}_2(X) \cong \mathrm{PL}_2(Y)$.*

Proof. Partition X into standard dyadic intervals, then do the same for Y. If the partitions have a different number of cells, subdivide a cell in the smaller partition. This adds one more cell to P; repeat until the partitions have the same number of cells. Say our partitions are $X\colon x_0 < \cdots < x_n$ and $Y\colon y_0 < \cdots < y_n$. We define a map $\Sigma\colon X \to Y$ by stipulating that Σ linearly maps $[x_i, x_{i+1}]$ to $[y_i, y_{i+1}]$ (as we did in Section 1.1.1). Now Σ is a PL_2 map: in particular, it has the correct gradients, because the width of each segment's domain and range is an integer power of 2. By the claim above, we see that $\mathrm{PL}_2(X)$ and $\mathrm{PL}_2(Y)$ are isomorphic, via a conjugation by Σ. □

Remark 1.2.9. If the ratio of the width of Y to the width of X is a power of two, then we can choose $\Sigma\colon X \to Y$ to consist of a single line segment, as in Figure 1.19 for example. Otherwise we will need to partition X and Y, then subdivide. This could potentially introduce arbitrary choices and produce different maps Σ, hence different isomorphisms $\mathrm{PL}_2(X) \to \mathrm{PL}_2(Y)$. Thus the recipe we have given does not produce a *canonical* isomorphism.

Say for instance that we have three cells, and need to expand the partition to four cells. Which of the three should we subdivide?

In later chapters, we won't worry about this too much—we just need to get our hands on any PL_2 map $\Sigma\colon X \to Y$, in order to explicitly work with an isomorphism $\cdot^{\Sigma}\colon \mathrm{PL}_2(X) \to \mathrm{PL}_2(Y)$.

Brown's article[39] gives more details on groupoids, and a thorough motivation for their study.

[39] R. Brown 1987.

1.2.3 *Cantor space homeomorphisms*

Finite binary trees describe dyadic subdivisions well, but they don't give us a natural way to talk about the points $x \in I$. At best, we can only approximate x with smaller and smaller dyadic intervals. We can formalise this by using infinite rather than finite addresses.

Notation 1.2.10 (Working with infinite strings)**.** Let $x = x_1x_2\ldots$ be an infinite string over an alphabet A. Any substring of the form $x_1 \ldots x_n$

for $n \geq 0$ is called a *prefix* of x; this includes the empty string ϵ when $n = 0$. An infinite substring of the form $x_n x_{n+1} \dots$ for $n \geq 1$ is called a *suffix* or *tail* of x.

If $a = a_1 \dots a_n$ is a finite nonempty string, we denote by $\bar{a}$ the infinite string $aaaa \cdots = (a_1 \dots a_n)(a_1 \dots a_n) \cdots$. The smallest string g such that $\bar{a} = \bar{g}$ is called the *minimal generator* of $\bar{a}$. The *left cyclic shift* of a is the string $a^{\triangleleft} = a_2 \dots a_n a_1$. Repeatedly cycling left i times yields the string $a^{\triangleleft i} = a_{i+1} \dots a_n a_1 \dots a_i$, with subscripts modulo n.

If x takes the form $x = x_1 \dots x_n \bar{a}$, we say that x is *eventually repeating* and call $\bar{a}$ a *repeating tail* of x. Note that x can also be written as $x_1 \dots x_n a_1 \overline{a_2 \dots a_n a_1} = x_1 \dots x_n a_1 \overline{a^{\triangleleft}}$, so $\overline{a^{\triangleleft}}$ is also a repeated tail of x.

Definition 1.2.11. The *boundary* of $\mathcal{T}_2$ is the set $\partial\mathcal{T}_2 = \{\mathtt{0},\mathtt{1}\}^{\mathbb{N}}$ of sequences of binary digits. To be explicit, its elements are infinite strings $s = s_1 s_2 s_3 \dots$ with no restrictions on the pattern of digits. If u is a finite binary string, we define the *cone below* u to be the set

$$[u] = \{\, us \mid s \in \partial\mathcal{T}_2 \,\}.$$

It has the same cardinality as $\partial\mathcal{T}_2$, namely $|2^{\mathbb{N}}| = |\mathbb{R}|$. We equip $\partial\mathcal{T}_2$ with the topology[40] generated by the cones below each finite string $u \in \{\mathtt{0},\mathtt{1}\}^*$. We call the resulting topological space *the Cantor space* (see Remark 1.2.13).

This is a specific instance of the Gromov boundary of a hyperbolic space X. Formally this consists of equivalence classes of geodesic rays in X, but the tree structure of $\mathcal{T}_2$ means that we can describe everything using binary strings only.[41]

When equipped with the topology above, the Cantor space $\partial\mathcal{T}_2$ is homeomorphic to the Cantor set C. In the construction of the latter set, we divide an interval into two parts. This leaves us with a left and right subinterval—hence the connection to $\mathcal{T}_2$ in Figure 1.23. The full construction is as follows.

Definition 1.2.12. The *Cantor set* C is a subset of the interval I constructed by an iterated function system. We delete the open middle third of $[0,1]$, leaving us with the subset $C_1 = [0, 1/3] \cup [2/3, 1]$. A scaled down version of this procedure is applied to the connected components of C_1: we delete the open middle third of the two remaining intervals. This results in a set C_2 consisting of four disjoint intervals, each of length $1/9$. The process is repeated, yielding a sequence $\{C_n\}_{n\in\mathbb{N}}$ of subsets in which the nth term is a disjoint union of 2^n intervals of length 3^{-n}.

The Cantor set C is defined as the intersection of all C_n. Its topology is the subspace topology inherited from I.

The natural homeomorphism $\phi\colon \partial\mathcal{T}_2 \to C$ is given by the formula $(s_1 s_2 \dots)\phi = \sum_{n\in\mathbb{N}} (2s_n)\, 3^{-n}$. The idea is to turn a string $s = s_1 s_2 \dots$ over $\{\mathtt{0},\mathtt{1}\}$ into a string $s' = s'_1 s'_2 \dots$ over $\{\mathtt{0},\mathtt{2}\}$, by replacing every $\mathtt{1}$ with a $\mathtt{2}$. The result is then interpreted as the infinite ternary expansion of a real number $s\phi = 0.s'_1 s'_2 \dots$. For instance, take the string $s =$

[40] If we give $\{\mathtt{0},\mathtt{1}\}$ the discrete topology, the topology we are describing is exactly the product topology on $\prod_{n\in\mathbb{N}} \{\mathtt{0},\mathtt{1}\}$.

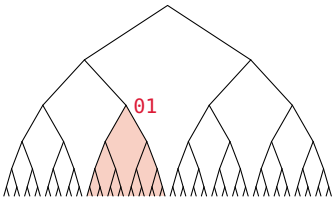

Figure 1.22: The cone [01] below $u = \mathtt{01}$.

[41] See Kapovich and Benakli 2002 for more details.

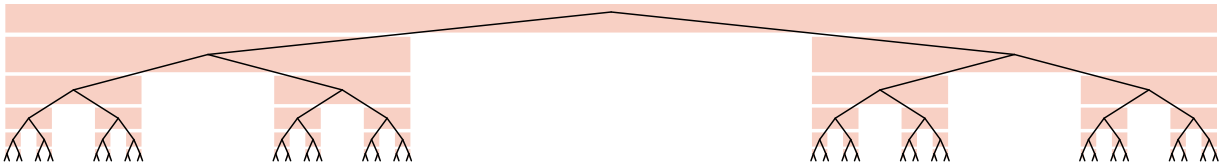

Figure 1.23: The correspondence between Cantor set construction steps C_n and the vertices of $\mathcal{T}_2$.

$\mathtt{0}\bar{\mathtt{1}} \in \partial\mathcal{T}_2$. This becomes a ternary string $s' = \mathtt{0}\bar{\mathtt{2}} \in \{\mathtt{0},\mathtt{2}\}^{\mathbb{N}}$, which we interpret as an infinite ternary expansion $s\phi = 0.0222\cdots \in C$. This is a geometric series with sum $s\phi = 1/3 \in C$.

Remark 1.2.13. Some authors reserve the term 'Cantor set' exclusively for the subset $C \subseteq I$ constructed by removing middle thirds. The Cantor set can be characterised purely topologically: every totally disconnected, perfect, compact, metric space is homeomorphic to the Cantor set.[42] Such spaces are called *Cantor spaces*, and $\partial\mathcal{T}_2$ in particular is called *the* Cantor space.

[42] Willard 2004, Corollary 30.4.

For our purposes the infinite string notation is most useful, so we will work with the Cantor space $\partial\mathcal{T}_2$ throughout this thesis.

There is a natural surjection π from the Cantor space to the interval I. An infinite string $s = s_1 s_2 \ldots$ is interpreted by π as the infinite *binary* expansion $0.s_1 s_2 \ldots$ of a real number $s\pi = \sum_{n\in\mathbb{N}} s_n 2^{-n}$ (c.f. the map ϕ above). Even better, π is continuous. To see this, let $U \subseteq I$ be an open subset.[43] For each $x \in U$, there is a standard dyadic interval $[\![w_x]\!]$ such that $x \in [\![w_x]\!] \subseteq U$. The inverse image $[\![w_x]\!]\pi^{-1}$ is a cone $[w_x]$, which is open. Then $U\pi^{-1} = (\bigcup_{x\in U}[\![w_x]\!])\,\pi^{-1}$ is the union $\bigcup_{x\in U}[w_x]$ of open sets, and hence open. In short, π allows us to use the Cantor space to reason about the interval I.

[43] We give I the topology it inherits as a subspace of $\mathbb{R}$ under the usual topology. This means e.g. that the intervals $[0,\epsilon)$ and $(1-\epsilon,1]$ are open in I, for all $\epsilon > 0$.

There is a price to pay for this: π is not injective. To see this, consider the infinite strings $x = \mathtt{1}\bar{\mathtt{0}}$ and $y = \mathtt{0}\bar{\mathtt{1}}$. The former has image $x\pi = 1/2$, which is equal to the image $y\pi = 1/4 + 1/8 + 1/16 + \cdots = 1/2$ of the latter. More generally, let $s_1 \ldots s_{n-1}\mathtt{1}$ be a finite string ending in $\mathtt{1}$. Then the distinct infinite strings $s_1 \ldots s_{n-1}\mathtt{1}\bar{\mathtt{0}}$ and $s_1 \ldots s_{n-1}\mathtt{0}\bar{\mathtt{1}}$ are both mapped by π to the real number z with finite binary expansion $z = 0.s_1 \ldots s_{n-1}\mathtt{1}$. This phenomenon[44] occurs exactly when z is a dyadic rational, because the dyadics are exactly the numbers with finite binary expansions. So nondyadics have exactly one preimage under π.

[44] This is the same mechanism which causes the decimal expansions $0.\bar{9}$ and $1.\bar{0}$ to be equal. See also Figure 2.17 for a visualisation.

The rationals $\mathbb{Q}$ have binary expansions which are eventually repeating. In $(0,1)$, these expansions take the form $x = 0.a\bar{b}$. If x is not dyadic, then the repeating tail is unique up to cyclic shifts (see Notation 1.2.10). Otherwise there are two possible repeating tails $\bar{\mathtt{0}}$ and $\bar{\mathtt{1}}$, as noted above.

There are two natural choices of metric to use on the Cantor space. Let $x = x_1 x_2 \ldots$ and $y = y_1 y_2 \ldots$ be two points in $\partial\mathcal{T}_2$. First there is the metric $d_{\text{thirds}}(x,y) = |x\phi - y\phi|$ obtained by restricting the usual metric on $[0,1]$ to the middle-thirds Cantor set C. The second metric is easier to define in terms of strings. If $x = y$ then set $d(x,y) = 0$; otherwise there is a maximal index $k \geq 0$ for which $x_k = y_k$. We define

$$C \underset{\phi^{-1}}{\overset{\phi}{\leftrightarrows}} \partial\mathcal{T}_2 \xrightarrow{\ \pi\ } [0,1]$$

$$\begin{aligned} [u]\phi &\longleftarrow [u] \longmapsto [\![u]\!] \\ [0,1/3]\cap C &\longleftarrow [\mathtt{0}] \longmapsto [\![\mathtt{0}]\!] = [0,1/2] \\ [2/3,1]\cap C &\longleftarrow [\mathtt{1}] \longmapsto [\![\mathtt{1}]\!] = [1/2,1] \end{aligned}$$

Figure 1.24: The Cantor set C is connected to the Cantor space $\partial\mathcal{T}_2$ via the homeomorphism ϕ. The map π continuously projects onto the unit interval $I = [0,1]$. The examples demonstrate how π 'reattaches' the intervals which remain after deleting a middle third during the construction of the Cantor set.

$d(x,y) = 2^{-k}$ in these circumstances; this is the width of the smallest standard dyadic interval $[\![x_1 \dots x_k]\!]$ containing both x and y.

We will typically work with the latter metric since we prefer to work with infinite strings and the Cantor space $\partial\mathcal{T}_2$. Fortunately, the topologies generated by d_{thirds} and d are the same as the topology generated by the cones (as in Definition 1.2.11).

The Cantor space is a kind of cover of the interval (via π), so we can lift elements of V to maps on the Cantor set. The recipe is the following: take a tree pair diagram (D, ℓ, R) for $\alpha \in V$, and let $d_1, \dots, d_n$ be the addresses of the leaves of D. The Cantor space version of α, $\hat{\alpha}$ say, is defined on the cone $[d_i]$ below d_i by the formula

$$(d_i b_1 b_2 b_3 \dots)\,\hat{\alpha} = (d_i \ell) b_1 b_2 b_3 \dots , \tag{1.6}$$

for any infinite binary string $b_1 b_2 \dots$. This is compatible with π in the sense that $\hat{\alpha}\pi = \pi\alpha$. Working over the Cantor space $\partial\mathcal{T}_2$ instead of $[0,1]$ has its advantages, because the space $\partial\mathcal{T}_2$ itself distinguishes between approaching a dyadic s from the left and the right. For example, the left-hand limit $x \to (1/2)^-$ corresponds to the point $\mathtt{0\bar{1}} \in \partial\mathcal{T}_2$, whereas the right-hand limit $x \to (1/2)^+$ corresponds to $\mathtt{1\bar{0}}$. This is useful when elements of V exhibit different behaviour either side of breakpoints—see the discussion surrounding Figure 2.15.

In the light of Equation (1.6), one way to generalise Thompson's groups is to a group of functions which rearrange infinite strings. We might insist that the rearrangement is performed by a *finite-state transducer*.[45] Informally, this is a simple computer which switches between a finite set of states Q. The transducer is in exactly one state $q \in Q$ at any given time. A computation step is performed by reading a binary digit b as input. Two things happen: the transducer outputs a sequence b' of binary digits, and changes to a new state $q' \in Q$ (possibly $q' = q$). The output word b' and new state q' are entirely determined by the input digit b and previous state q.

An element $\alpha \in V$ will always be represented by a transducer with an 'echo state' e. In this state, the transducer reads in a digit b, immediately outputs the same digit and remains in state e. (See Figure 1.25, and compare with Equation (1.6)). This preserves the tail of an infinite string, once we have read as input the address of an interval on which α acts linearly. Transducers without an echo state

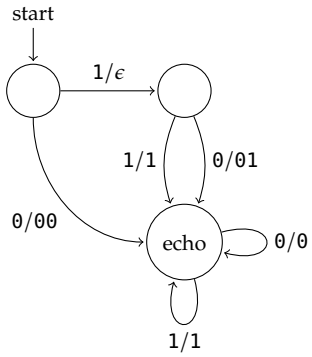

Figure 1.25: A transducer on the input and output alphabet $\{\mathtt{0},\mathtt{1}\}$ corresponding to $x_0 \in F$. The 'x/y' notation is short for 'read x, then write y'. The symbol ϵ stands for the empty word.

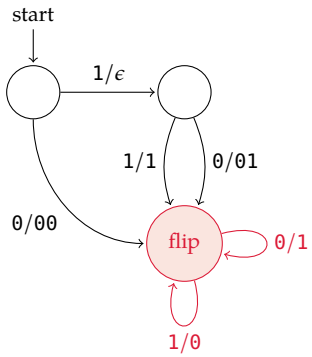

Figure 1.26: If we modify the echo state to flip digits $\mathtt{0} \leftrightarrow \mathtt{1}$, the resulting transducer does not model an element of V. Instead, it first applies $x_0 \in F$, then reflects each cone $[\mathtt{0}]$, $[\mathtt{10}]$ and $[\mathtt{11}]$ about their midpoint.

[45] For a rigorous definition, see Grigorchuk, Nekrashevych and Sushchanskiĭ 2000.

(e.g. Figure 1.26) describe new homeomorphisms of the Cantor space; see for instance the articles of Grigorchuk et al, Bleak et al and Aroca.[46]

[46] Grigorchuk, Nekrashevych and Sushchanskiĭ 2000; Bleak, Donoven and Jonušas 2017; Aroca 2018.

2 *The dynamics of Thompson's groups*

A fruitful way to think of Thompson's groups is to consider their elements as being *discrete dynamical systems*.

Definition. A discrete dynamical system[1] on a space X is a homomorphism $\Gamma\colon \mathbb{Z} \to \mathrm{Homeo}(X)$. The space X is called the *phase space*, and the integers $\mathbb{Z}$ model *time*.

Such a system is determined entirely by the homeomorphism at time point 1, say $\gamma = 1\Gamma \in \mathrm{Homeo}(X)$. Points $x \in X$ are mapped through their *orbit* $O_x = \{x(t\Gamma) = x\gamma^t \mid t \in \mathbb{Z}\}$.

In our context, the phase space X is either the interval $[0,1]$, the circle S^1 or the Cantor space $\partial\mathcal{T}_2$; our homeomorphisms γ are elements of F, T or V. Rather than study a Thompson group element γ directly, we'll study its orbits, to see how γ rearranges points in X. This is particularly advantageous, since an element γ need not act nicely on the set of standard dyadic intervals (i.e. on $\mathcal{T}_2$); at best, we have only a partial action. For example, the image of $[0, 1/2]$ under γ in Figure 2.1 is $[0, 3/8]$—which is not a standard dyadic interval.

[1] Strictly speaking this is a *reversible* system, since we are rearranging X with homeomorphisms. More generally we could modify X with endomorphisms, which need not have an inverse—let alone a continuous inverse. Then we'd need to use $\mathbb{N}$ to model time.

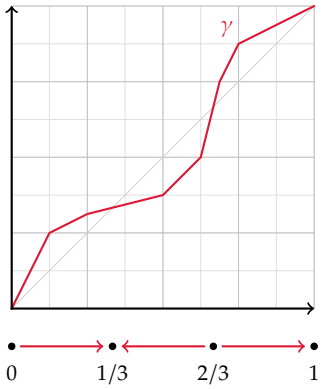

Figure 2.1: As $t \to \infty$, the powers $\gamma, \gamma^2, \ldots, \gamma^t, \ldots$ of γ from Figure 1.18 move points away from the sources 0 and 2/3, limiting on sinks 1/3 and 1. Each of these four sources and sinks is fixed by γ. This is summarised in the sketch below the graph, which is known as a *phase portrait* in the theory of dynamical systems.

2.1 Actions of Thompson's groups

This chapter serves as a 'toolkit': a collection of minor results for use later. To get things started, we study how the Thompson groups rearrange the points of the interval (or if you prefer, the points of the Cantor space). Later we specialise to the action of T on the circle.

2.1.1 *Dyadics versus nondyadics*

The first thing to emphasise is the distinction between dyadic and nondyadic points. Let G be one of the groups F, T or V, and take a linear segment of a generic element $\alpha \in G$. This segment is described by an equation of the form $y = y_0 + M(x - x_0)$, where (x_0, y_0) is an endpoint with dyadic coordinates and $M = 2^m$ is a power of two. If x is dyadic, then y is formed by adding, multiplying and subtracting other dyadics—so $y = x\alpha$ is dyadic.[2] Conversely, if y is dyadic then $x = y\alpha^{-1}$ is dyadic by the same kind of argument. This shows that each of the three actions $I \circlearrowleft G$ permute the dyadics (and hence their complement, the nondyadics) in $[0,1]$. The upshot is that the dyadics are first-class citizens—perhaps not a surprise given their prominence in Definitions 1.1.2, 1.1.10 and 1.1.13 of F, T and V.

How are the dyadics in $[0,1]$ divided into orbits under Thompson's groups? For starters, F must preserve the minimum and maximum

[2] We have already made use of this argument in the proof of Claim 1.2.5.

elements of the interval, being a group of order-preserving transformations. Thus $\{0\}$ and $\{1\}$ are orbits of size one under F. Any remaining dyadics $x = a/2^n$ and $x' = a'/2^{n'}$ must be minimal elements of standard dyadic intervals X and X'. Because $x, x' \notin \{0, 1\}$, we may assume that both X and X' contain neither 0 or 1.[3] Choose dyadic subdivisions D and D' containing X and X', respectively. We may expand these subdivisions so that they contain the same number of cells, and so that X and X' have the same index in their respective partitions. The element of F corresponding to the modified subdivisions D and D' must then map x to x'.

[3] If $1 \in X$, replace X with its left child.

Now we complete the story for T and its supergroup V. As noted above, any element $\alpha \in T$ must send 0 (now identified with 1) to a dyadic. It is now possible that $0\alpha \neq 0$: for instance, α in Figure 1.14 sends $0 \mapsto 1/4$. Then an element $\beta \in F$ can be used to send 0α to any other dyadic. Thus T (and hence V) act transitively on the dyadics in $S^1 \cong [0, 1)$.

We summarise with a lemma.

Lemma 2.1.1 (Behaviour on dyadics)**.** *Every element of F fixes 0 and fixes 1. The actions of*

- *F on the dyadics in $(0, 1)$, and*
- *T and V on the dyadics in S^1*

are transitive.

The observation above—that we can move dyadics to other dyadics—is worth formalising. As Cannon-Floyd-Parry explain,[4] we can actually map sorted lists of dyadics to sorted lists of dyadics using F.

[4] Cannon, W. Floyd and Parry 1996, Lemma 4.2.

Lemma 2.1.2 (F aligns dyadics)**.** *Let $x_0 < \cdots < x_n$ and $y_0 < \cdots < y_n$ be dyadic partitions.[5] We can construct an element $\alpha \in F$ such that $x_i\alpha = y_i$ for each i simultaneously. Thus F acts transitively on the set of ordered dyadic partitions with k components, for each k.*

[5] As defined above Proposition 1.1.3.

Sometimes it's not enough to map points to points. The next lemma[6] allows us to join up (or 'glue') pre-existing maps from subintervals to other subintervals.

[6] Matucci 2008, Lemma 4.1.5. Our quoted result is not exactly that written by Matucci: we allow our intervals to share endpoints, and more carefully handle the case when 0 or 1 belongs to the intervals in question.

Lemma 2.1.3 (Extension of partial maps)**.** *Let $I_1, \ldots, I_k \subseteq [0, 1]$ be a family of intervals $I_i = [a_i, b_i]$ whose endpoints $a_1 < b_1 \leq a_2 < b_2 \leq \cdots < b_{k-1} \leq a_k < b_k$ are all dyadic. Let $J_1, \ldots, J_k$ be another family of intervals $J_i = [c_i, d_i]$ whose endpoints satisfy the same conditions. Assume that*

- *$b_i = a_{i+1}$ if and only if $d_i = c_{i+1}$;*
- *$a_1 = 0$ if and only if $c_1 = 0$; and*
- *$b_k = 1$ if and only if $d_k = 1$.*

Suppose that $\gamma_i \colon I_i \to J_i$ is a PL_2 map for each i. Then there exists a map $\Gamma \in F$ such that $\Gamma\big|_{I_i} = \gamma_i$ for each i simultaneously.

2.1.2 *The action on the rationals*

More generally, if x is a member of some ring $R \leq \mathbb{R}$ which contains the dyadics $\mathbb{Z}[1/2]$, then $y = x\alpha$ belongs to R if and only if $x \in R$, for all $\alpha \in G$. Thus G permutes $R \cap I$, and also the difference $(R \setminus \mathbb{Z}[1/2]) \cap I$. The action on the rationals $R = \mathbb{Q}$ will be particularly important (see Lemma 2.3.10). Kassabov and Matucci[7] specify when two rationals in $(0,1)$ belong to the same orbit under F.

[7] Kassabov and Matucci 2012, Proposition 6.18.

Proposition 2.1.4. *Let $x = 2^t m/n$ and $y = 2^k p/q$ be rational numbers in $(0,1)$, where $t, k \in \mathbb{Z}$ and m, n, p, q are odd integers with $\gcd(m,n) = \gcd(p,q) = 1$. Then there is a $\gamma \in F$ such that $x\gamma = y$ if and only if $n = q$ and*

$$p = 2^R m \pmod{n}$$

for some $R \in \mathbb{Z}$. Moreover there is an algorithm which constructs such an element γ if the above condition is satisfied.

Their result is easier to state if we use the infinite address notation for points in the Cantor space.

Lemma 2.1.5. *Let G be one of the groups F, T or V. Let $x, y \in (0,1)$ be rational numbers, with binary expansions $x = 0.a\bar{b}$ and $y = 0.c\bar{d}$ such that $\bar{b} \neq \bar{1} \neq \bar{d}$.*

Assume that b and d are minimal generators[8] for $\bar{b}$ and $\bar{d}$. Then x and y belong to the same orbit under G if and only if b is a cyclic permutation of d.

[8] See Notation 1.2.10

Proof. The second assumption is necessary to handle the fact that dyadics have two binary expansions: see the discussion after Definition 1.2.12.

Let $\alpha \in G$ belong to one of Thompson's groups. There is a finite prefix w of $x = a\bar{b}$ such that α maps $[\![w]\!]$ linearly onto its image. Any prefix longer than w also has this property, because $[\![w\mathtt{0}]\!]$ and $[\![w\mathtt{1}]\!]$ are subintervals of $[\![w]\!]$. Extending w if necessary, we may assume w takes the form $w = ab^n c$, where c is a prefix of b—say $b = cd$. Then $x = wd\overline{cd} = w\overline{dc} = wb^{\overline{\lhd|c|}}$. This is sent by α to $x\alpha = (w\alpha)\overline{b^{\lhd|c|}}$, so the cyclic permutation condition on b and d is necessary for x and y to share an orbit.

In the other direction, let $d = b^{\lhd i}$ so that $d = uv$ and $b = vu$, where $|v| = i$. Rewrite x as $x = a\overline{vu} = av\overline{uv} = av\bar{d}$. We may assume that av and c both contain at least one $\mathtt{0}$ and one $\mathtt{1}$. If not, replace a and c with the strictly longer strings ab and cd. This works because ab^n and cd^m are prefixes of x and y respectively, and $x, y \notin \{0,1\}$ have addresses with at least one of each binary digit. Let γ be the PL_2 map $[\![av]\!]$ to $[\![c]\!]$ consisting of one linear segment. Invoke Lemma 2.1.3 to construct an element $\Gamma \in F$ which restricts to γ. Then $x\Gamma = av\gamma\bar{d} = c\bar{d} = y$. Thus the cyclic permutation condition is sufficient for x and y to share an orbit. □

This lets us partition $\mathbb{Q} \cap [0,1]$ into orbits under F. Firstly, both 0 and 1 belong to orbits of size one; all other dyadics belong to the same

orbit. A nondyadic is represented—up to cyclic shifts—by the minimal generator for its binary expansion's repeating tail (i.e. b above). To select a unique representative, choose the cyclic shift of b which is minimal under the lexicographic order.

In T and V, the dyadics including 0 and 1 form a single orbit. The nondyadics are broken into the same orbits as in F.

2.1.3 *Rotations*

In this and the following section, we consider dynamical systems on the circle S^1, rather than the on interval I or Cantor space $\partial\mathcal{T}_2$. This is because we want to discuss the group T later in the thesis, rather than F or V. With this in mind, we need to work in the group $\mathrm{Homeo}_+(S^1)$ of orientation-preserving homeomorphisms, rather than the larger group $\mathrm{Homeo}(S^1)$ of all homeomorphisms. Working in the smaller group excludes the maps which *reverse* orientation, including e.g. reflection maps.

One important class we do have at our disposal is the collection of *rotations*. These are parameterised by an angle s, which may be taken to be any real number $0 \le s < 1$. (Our angles are in units such that a rotation by 1 is a rotation through 2π radians.) The rotation by s is the map ρ_s which sends $x \mapsto x + s \bmod 1$. Its inverse element is the rotation ρ_{-s} through the same angle in the opposite direction.

When working in Thompson's groups, we don't have access to the full group of rotations. If $s \notin \mathbb{Z}[1/2]$ is not dyadic, then ρ_s would map the point 0 to s. But this is impossible, because we know from Lemma 2.1.1 that T permutes the dyadics in S^1. So suppose otherwise that $s = a/2^n$. To construct $\rho_s \in T$, let D be the tree pair with 2^n leaves, each of depth n. Label these $1, \dots, 2^n$ from left to right. Let ℓ be the permutation of the leaves of D which sends leaf i to leaf $i + a \bmod 2^n$. Then (D, ℓ, D) is a tree pair for ρ_s.

Thus the rotations inside T are precisely the dyadic rotations. The set of such maps forms a subgroup of T, isomorphic to the *Prüfer 2-group*

$$\mathbb{Z}(2^\infty) = \langle \{ g_n \}_{n\in\mathbb{N}} \mid g_1^2 = \mathrm{id},\ g_{n+1}^2 = g_n \rangle \cong \mathbb{Z}[1/2]/\mathbb{Z}\,,$$

where the generator g_n corresponds to the rotation $\rho_{2^{-n}}$. See e.g. Rotman[9] for more details.

This next result isn't deep at all, but it does highlight an important connection between F and T.

Lemma 2.1.6. *Let $x \in S^1$ be dyadic. The stabiliser $T_x = \{ \alpha \in T \mid x\alpha = x \}$ of x is exactly the conjugate F^{ρ_x}.*

Proof. First suppose that $\alpha \in T_x$. Because x is dyadic, $\rho_{-x} \in T$ and so $\alpha^{\rho_{-x}} \in T$ also. The calculation $0\rho_x\alpha\rho_{-x} = x\alpha\rho_{-x} = x\rho_{-x} = 0$ shows that this conjugate fixes 0, and hence belongs to F. Thus $\alpha = (\alpha^{\rho_{-x}})^{\rho_x} \in F^{\rho_x}$.

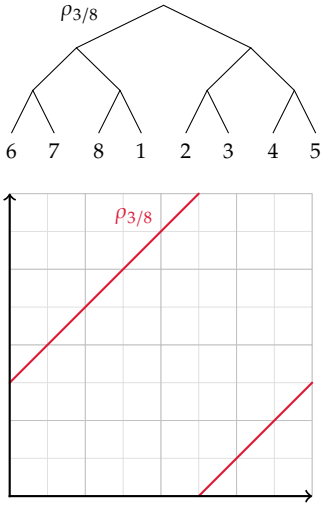

Figure 2.2: The dyadic rotation by 3/8. The tree pair diagram (D, l, R) uses the same tree twice, so we've abbreviated this by only drawing R. The domain tree is labelled $1, \dots, 8$ as usual.

[9] Rotman 1995, Section 10, in particular Theorem 10.13 and Exercise 10.5.

Conversely, given $\beta \in F$, the conjugate $\beta^{\rho_x} \in F^{\rho_x}$ is in T. This conjugate maps x to $x\rho_{-x}\beta\rho_x = 0\beta\rho_x = 0\rho_x = x$, so $\beta^{\rho_x} \in T_x$. □

We'll reuse the idea on a number of occasions: if $\beta \in T$ has a dyadic point x of interest, we can 'pretend' that $\beta \in F$ and that x is 0. Then hopefully we can use our knowledge of F to our advantage. We achieve this via conjugation by a dyadic rotation. When viewed graphically, this conjugation translates a function's graph along the diagonal— see Figure 2.3.

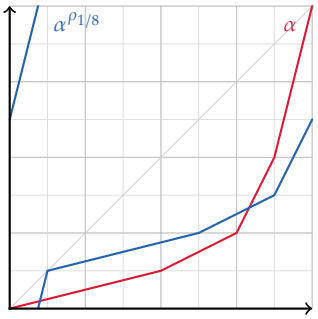

Figure 2.3: An element $\alpha \in F$ and its conjugate $\alpha^{\rho_{1/8}} \in T$. The blue graph is obtained from the red graph by translating along the vector $(1/8, 1/8)$.

Here's an example of this pretending in action, in which we generalise our earlier gluing result (Lemma 2.1.3) to T.

Lemma 2.1.7 (Extension of partial maps). *Let $I_1, \ldots, I_k \subseteq S^1$ be a family of closed arcs $I_i = [a_i, b_i]$ with dyadic endpoints. Suppose that any two arcs intersect either trivially or at a common endpoint. Let $J_1, \ldots, J_k \subseteq S^1$ be another family of closed arcs $J_i = [c_i, d_i]$ with the same properties. For each i, assume that $b_i = a_j$ for some j if and only if $d_i = c_j$.*

Suppose that $\gamma_i \colon I_i \to J_i$ are PL_2 maps. Then there exists a map $\Gamma \in T$ such that $\Gamma\big|_{I_i} = \gamma_i$ for each i simultaneously.

Proof. Let $X \colon x_1 \to \ldots \to x_n \to x_1$ be the points of the set $\{a_i, b_i \mid 1 \leq i \leq k\}$ with duplicates removed. Similarly let $Y \colon y_1 \to \ldots \to y_m \to y_1$ be the points of the set $\{c_i, d_i \mid 1 \leq i \leq k\}$ with duplicates removed. The 'assume that…' condition tells us that $n = m$, and that $[x_i, x_{i+1}] = I_j$ for some j if and only if $[y_i, y_{i+1}] = J_j$ for the same i and j.

Now we rotate our partitions' points so that they include zero. Introduce new linear partitions

$$\begin{aligned} X\rho_{-x_1} = X' \colon 0,\ x_2\rho_{-x_1},\ \ldots,\ x_n\rho_{-x_1} \\ \text{and} \quad Y\rho_{-y_1} = Y' \colon 0,\ y_2\rho_{-y_1},\ \ldots,\ y_n\rho_{-y_1}\,, \end{aligned}$$

of the interval $[0, 1]$, and write $I'_j = I_j\rho_{-x_1}$ and $J'_j = J_j\rho_{-y_1}$. Let γ'_i denote the PL_2 map $\gamma'_i = \rho_{+x_1}\gamma_i\rho_{-y_1}$. The chain of domains and codomains is

$$I'_j \xmapsto{\rho_{+x_1}} I_j \xmapsto{\gamma_i} J_j \xmapsto{\rho_{-y_1}} J'_j\,,$$

so γ'_j maps $I'_j \to J'_j$. As all the points involved are dyadic, we can invoke Lemma 2.1.3 to yield an element $\Gamma' \in F$ such that $\Gamma'\big|_{I'_j} = \gamma'_j$ for each j. Then define $\Gamma = \rho_{-x_1}\Gamma'\rho_{+y_1}$, noting that $\Gamma \in TFT \subseteq T$. The restriction of Γ to I_j is

$$\begin{aligned} \Gamma\big|_{I_j} &= \left(\rho_{-x_1}\Gamma'\rho_{+y_1}\right)\big|_{I_j} \\ &= \rho_{-x_1}\ \Gamma'\big|_{I'_j}\ \rho_{+y_1} \\ &= \rho_{-x_1}\rho_{+x_1}\gamma_j\rho_{-y_1}\rho_{+y_1} \\ &= \gamma_j\,, \end{aligned}$$

as required. □

2.1.4 *The rotation number*

Unfortunately, the rotation maps account for only a small proportion of orientation-preserving homeomorphisms $\alpha \in \mathrm{Homeo}_+(S^1)$. However, every such α has a 'nearest' rotation that—in some sense—best approximates α. This idea is made precise by the *Poincaré rotation number* of α, which we define and discuss below. For a thorough treatment of the rotation number, see Katok and Hasselblatt[10] or Bleak et al.[11]

[10] Katok and Hasselblatt 1995, Chapter 11.
[11] Bleak, Kassabov and Matucci 2011.

Definition 2.1.8 (Rotation number)**.** Let α be an orientation-preserving homeomorphism of the circle. Choose a lift $\bar{\alpha}\colon \mathbb{R} \to \mathbb{R}$ of α, via the covering map $\pi\colon \mathbb{R} \to S^1$ which sends $x \mapsto x \bmod 1$. We define the *rotation number* $\mathrm{rot}(\alpha)$ to be the limit

$$\mathrm{rot}(\alpha) = \lim_{n\to\infty} \frac{x\bar{\alpha}^n - x}{n} \bmod 1 \in \mathbb{R}/\mathbb{Z}\,, \tag{2.1}$$

which always exists and is independent of the choices of point x and lift $\bar{\alpha}$ used.

For the sake of convenience, we always take our rotation number in the interval $[0, 1)$ so that we don't have to identify numbers that differ by an integer. When specifying a rational rotation number p/q, we assume $0 \le p < q$ and $q > 0$ are coprime. In the special case $p = 0$, we take $q = 1$.

Proposition 2.1.9. *The rotation number operator* $\mathrm{rot}(\cdot)$ *has the following properties.*

1. *The rotation number of* ρ_s *is* s.
2. *If* $m \in \mathbb{Z}$ *then* $\mathrm{rot}(\alpha^m) = m\,\mathrm{rot}(\alpha)$.
3. rot *is invariant under conjugation in* $\mathrm{Homeo}_+(S^1)$.
4. $\mathrm{rot}(\alpha) = 0$ *if and only if* α *has a fixed point.*
5. $\mathrm{rot}(\alpha) = p/q \in \mathbb{Q}/\mathbb{Z}$ *in lowest terms if and only if* α *has a periodic orbit of length* q*. If this is the case, then* all *periodic points are in size* q *orbits under* α.

Let $\alpha \in T$. The fourth property tells us that if $\alpha \in F$ then $\mathrm{rot}(\alpha) = 0$, since any element of F fixes $\partial I = \{0, 1\}$ pointwise. The converse is false: see Figure 2.4.

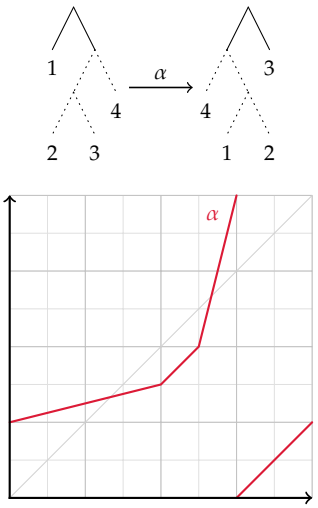

Figure 2.4: This element $\alpha \in T \setminus F$ fixes $1/3$ and $2/3$. Thus $\mathrm{rot}(\alpha) = 0$.

Note that fourth property is the special case of the fifth property where $q = 1$.

The fifth property does not tell us what the numerator of $\mathrm{rot}(\alpha)$ is. Suppose that $x \in S^1$ is fixed by α^q and choose a lift $\bar{\alpha}$ of α. Then $x\bar{\alpha}^q = x + p$ for some integer p. In prose, this means that the point x is wound p times around the circle by α^q. It follows that $x\bar{\alpha}^{mq} = x + mp$, for any integer m. Let r_n be the nth term of the sequence in Equation (2.1). Then

$$r_{mq} = \frac{x\bar{\alpha}^{mq} - x}{mq} = \frac{mp}{mq} = \frac{p}{q} \pmod 1\,,$$

which shows that $(r_{mq})_{m\geq 1}$ is a convergent subsequence of $(r_n)_{n\geq 1}$. Since we know that the larger sequence has a limit $\mathrm{rot}(\alpha)$, it must be equal to the subsequence's limit, namely $\mathrm{rot}(\alpha) = p/q$.

In short: the numerator of a rational rotation number counts how many times a periodic orbit wraps around the circle before returns to its starting position.

Lemma 2.1.10 (Finding the rotation number)**.** *Suppose $\alpha \in \mathrm{Homeo}_+(S^1)$ has a finite orbit $X = \{x_0 \to \cdots \to x_{q-1} \to x_0\}$ of points in the circle, enumerated in circular order. Let $x_0\alpha = x_p$. Then $\mathrm{rot}(\alpha) = p/q$ in lowest terms.*

Proof. The orbit of x_0 is traced out by α in the order

$$x_0 \mapsto x_p \mapsto x_{2p} \mapsto \cdots \mapsto x_{(q-1)p} \mapsto x_{qp} = x_0\,,$$

with subscripts modulo q. Because this list is the entire orbit X, we know that p has order q in the additive group $\mathbb{Z}_q$ of integers modulo q. This means that $\gcd(p, q) = 1$ and hence $\mathrm{lcm}(p, q) = pq$. We see that x_0 returns to its original position for the first time after q applications of α. In doing so, x_0 winds around the circle p times. So $\mathrm{rot}(\alpha) = p/q$. □

In short, we can read off a rotation number given only a finite orbit. In the other direction, once we know α has a rational rotation number, we can immediately break down a finite set permuted by α into orbits.

Lemma 2.1.11 (Rotation number determines actions)**.** *Let $X = \{x_0 \to \cdots \to x_{n-1} \to x_0\}$ be a nonempty finite set of points in the circle permuted by $\alpha \in \mathrm{Homeo}_+(S^1)$. Write $\mathrm{rot}(\alpha) = p/q$ in lowest terms. Then:*

1. *X consists of $t = n/q$ orbits of size q. The orbit containing x_j is*

$$\{x_j \to x_{j+t} \to x_{j+2t} \to \cdots \to x_{j+(q-1)t} \to x_j\}$$

   (expressed in cyclic order rather than orbit order).

2. *The action $X \circlearrowleft \alpha$ is given by $x_j\alpha = x_{j+pt}$, for each $0 \leq j < n$. In particular $x_j\alpha^k = x_{j+t}$, where $k = p^{-1} \pmod q$.*

That is, k is the unique integer $0 \leq k < q$ such that $pk \equiv 1 \pmod q$. This number k is the *multiplicative inverse* of p modulo q.

Proof. As all finite orbits under α must have size q (Proposition 2.1.9.5), we know that $n = qt$ for some integer t. Write $x_0\alpha = x_i$. Since α preserves the cyclic order on X, we know $x_j\alpha = x_{j+i}$ for each j, with subscripts modulo n. We know that $x_j\alpha^q = x_{j+qi} = x_j$. The rotation number tells us we have wound p times around the circle, so we have $qi = pn$. Rearranging yields $i = pn/q = pt$. Finally, we know $k = p^{-1} \pmod q$ exists because p and q are coprime. Then $x_j\alpha^k = x_{j+pkt} = x_{j+t}$, since subscripts are modulo n. □

Since the rotation number is an invariant of conjugacy (Proposition 2.1.9.3), we immediately obtain the following corollary.

Corollary 2.1.12. *Let $X \subseteq S^1$ be a finite set permuted by conjugate elements α and β. Then the actions $X \circlearrowleft \alpha$ and $X \circlearrowleft \beta$ are identical.*

ROTATION NUMBERS MAY VERY WELL BE IRRATIONAL, e.g. take the rotation by $1/\sqrt{2}$. If so, α is a 'wild' automorphism, and its orbits are qualitatively different[12] to those of an element with rational rotation number. Fortunately, the next theorem tells us that the elements of T have rational rotation number (so are not 'wild'). It is originally due to Ghys-Sergiescu,[13] though other proofs later followed.[14] Our proof makes use of the 'revealing pair' technology, which we'll see in Section 2.2.

[12] Katok and Hasselblatt 1995, Table above Theorem 11.2.9.

[13] Ghys and Sergiescu 1987.

[14] Liousse 2005; Calegari 2007; Burillo et al. 2009; Bleak, Kassabov and Matucci 2011; Diller and Lin 2017.

Theorem 2.1.13. *Every element of T has rational rotation number. Moreover every rational $x \in \mathbb{Q}/\mathbb{Z}$ is the rotation number of some element $\alpha \in T$.*

Proof. To compute the rotation number of $\alpha \in T$, choose a revealing pair (D, ℓ, R) for α. We will see later (in Remark 2.2.9) that we can always find a point x in a finite orbit. Then Proposition 2.1.9.5 tells us that $\mathrm{rot}(\alpha)$ is rational.

To construct[15] $\alpha \in T$ with rotation number $1/q$, let D be any tree with q leaves, labelled $1 < \cdots < q$ in left-to-right order. Define β by the tree pair diagram (D, ℓ, D), where ℓ sends the ith leaf to the $(i+1)$th leaf modulo q. Call the leaves $u_0, \ldots, u_{q-1}$ in left-to-right order, and let x_i be the left endpoint of $[\![u_i]\!]$. Then $X = \{x_0 \to \cdots \to x_{n-1} \to x_0\}$ is a finite orbit of points under α, with $x_0\alpha = x_1$. We conclude from Lemma 2.1.10 that $\mathrm{rot}(\alpha) = 1/q$. Then Proposition 2.1.9.2 informs us that $\mathrm{rot}(\alpha^p) = p/q$. □

[15] Bleak et algo beyond this construction and provide an embedding $\mathbb{Q}/\mathbb{Z} \hookrightarrow T$; see Bleak, Kassabov and Matucci 2011.

The rotation number is a powerful tool to have at our disposal. For starters, we use it to show that the only periodic points under F are fixed points.

Lemma 2.1.14. *Suppose that $\alpha \in F$ has a periodic point $x\alpha^n = x$ with $n \neq 0$. Then $x\alpha = x$.*

Proof. As $\alpha \in F$, we must have $0\alpha = 0$, so $\mathrm{rot}(\alpha) = 0/1$ by Proposition 2.1.9.4. Then all finite orbits under α have size 1, by Proposition 2.1.9.5.

For a more direct proof, let $O = \{x\alpha^i \mid i \in \mathbb{Z}\}$ be the orbit in question Seeking a contradiction, assume $|O| \geq 2$ and write $m = \min O$. Then $m < m\alpha$. As F preserves orientation (i.e. linear order) we have $m < m\alpha^i$ for each positive integer i. But then we'd conclude that $m < m\alpha^n = m$, which is nonsense. □

With that said, the rotation number is not all-powerful. It is not a homomorphism $\mathrm{Homeo}_+(S^1) \to \mathbb{R}/\mathbb{Z}$, nor a homomorphism $T \to \mathbb{Q}/\mathbb{Z}$; see for example Figure 2.5. However, it is a homomorphism when restricted to certain subgroups of homeomorphisms.[16]

[16] Matucci 2008, Theorem 6.2.1, published in Bleak, Kassabov and Matucci 2011, Lemma 1.8.

Theorem. *Let $G \leq \mathrm{Homeo}_+(S^1)$ have no nonabelian free subgroups. Then the rotation number is a homeomorphism $G \to \mathbb{R}/\mathbb{Z}$.*

In the next lemma, we use Proposition 2.1.9 to explain why the following result[17] about circle homeomorphisms fails for T.

[17] Matucci 2008, Lemma 7.1.2.

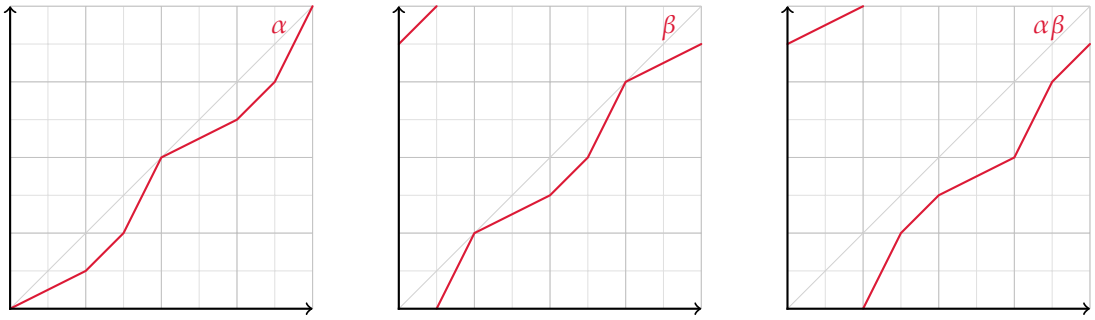

Figure 2.5: Two elements $\alpha, \beta \in T$ with rotation number 0. Their product $\alpha\beta$ has rotation number 5/6, so the rotation number is *not* a homomorphism $T \to \mathbb{R}/\mathbb{Z}$. Compare with Figure 6.4 of Matucci's thesis.

Lemma 2.1.15. *Let G be one of the groups $\mathrm{Homeo}_+(S^1)$ or $\mathrm{PL}_+(S^1)$. Every torsion element of G is conjugate to a rotation.*

Proof that this is false in T. In short: T doesn't contain enough rotations.

In more detail, let $\alpha \in T$ have finite order $q > 2$ and suppose that q is odd. Then α and any of its conjugates have rotation number $\mathrm{rot}(\alpha) = p/q$, for some integer p coprime to q. So the only conjugate of α which could be a rotation is $\rho_{p/q}$. But p/q is not dyadic, because $q > 2$ is odd. Hence $\rho_{p/q} \notin T$. □

2.2 Revealing pairs

The previous section tried to discuss the big picture, explaining how Thompson's groups act on spaces as a whole. Now we want a much smaller picture: we'll examine the orbit of single points x under a single element α.

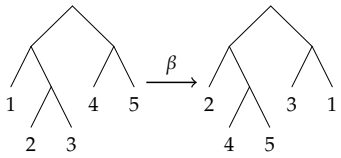

Figure 2.6: Another tree pair diagram for β from Figure 1.16.

Not all tree pair diagrams are created equally. A given element $\alpha \in V$ has a unique minimal tree pair diagram, but this doesn't always give us a clear picture of how α rearranges the interval. Take β from Figure 1.16 for instance. Without tracing through the orbit of points or intervals, it's difficult to see how β behaves. If we add a caret to the domain tree at address 1, we obtain the slightly larger tree pair as in Figure 2.6. Crucially, the domain and range trees are now equal. Hence β can be seen as a permutation of the leaf labels $\{1, \ldots, 5\}$. Its order must therefore divide 5!, and certainly be finite.

What is it about this larger tree pair that allows it to better describe the action of β on the interval? We turn to Brin[18] and Salazar-Díaz[19] for the answer. Alternatively see Bleak et al.[20]

[18] Brin 2004, Section 10.
[19] O. P. Salazar-Díaz 2010, Section 3.
[20] Bleak, Bowman et al. 2013, Section 4.1.

Definition 2.2.1 (Operations on binary trees)**.** Let D and R be two binary trees, viewed as subtrees of the infinite binary tree $\mathcal{T}_2$. The union $D \cup R$ is the subtree whose vertices belong to at least one of D and R. The intersection $D \cap R$ is the subtree whose vertices belong to both D and R. The difference $D \setminus R$ is the subforest of D whose vertices belong to D and are not internal vertices of R. Figure 2.7 illustrates these operations with examples.

The leaves of $D \cap R$ are partitioned into three types.

- Leaves of both D and R. These are called *neutral leaves*.

- Leaves of D which are internal vertices of R. These are *roots of* $R \setminus D$.
- Leaves of R which are internal vertices of D. These are *roots of* $D \setminus R$.

When (D, ℓ, R) is a tree pair, we denote the intersection $D \cap R$ with solid edges and the differences $D \setminus R$ and $R \setminus D$ with dotted edges.

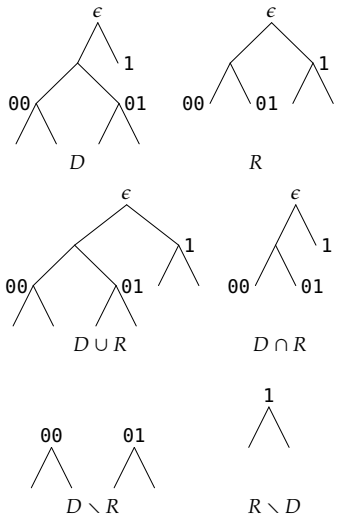

Figure 2.7: Four operations displayed for two binary trees D and R of different sizes.

Definition 2.2.2 (Attractors and repellers)**.** Let d_1 be a leaf of $D \setminus R$ in a tree pair (D, ℓ, R). Its *iterated augmentation chain* (IAC) is the maximal sequence $d_1, \ldots d_n$ which

- contains no element more than once, and
- has the property that $d_i \ell = d_{i+1}$ for $1 \leq i < n$.

Then $d_2, \ldots, d_n$ are neutral leaves, and $r = d_n \ell$ is a leaf of $R \setminus D$. We call $d_n \ell$ the *target* of the IAC, and n the *length* of the IAC.

We call d_1 a *repeller* if the address of r is a prefix of the address of d_1. In this situation r is the root of the component of $D \setminus R$ containing d_1. The binary string s such that $d_1 = rs$ is called the *spine* of the repeller. Since $d_1 \alpha^n = r$, we see that $d_1 s \alpha^n = d_1$ for a repeller d_1.

Dually, we call r an *attractor* if the address of d_1 is a prefix of the address of r. In this situation d_1 is the root of the component of $R \setminus D$ containing r. The binary string s such that $r = d_1 s$ is called the *spine* of the attractor. This time we see that $d_1 \alpha^n = d_1 s$ for an attractor d.

Remark 2.2.3. It would be nice to maintain the distinction between α acting on infinite strings and ℓ acting on finite strings (the tree pair's leaves' addresses). In practice, we blur the lines a little: when convenient, we define $dw\alpha = d\ell w$, for some finite—possibly empty—string w over $\{\mathtt{0}, \mathtt{1}\}$. (Compare to in Equation (1.6)). For instance, when we say above that $d_1 s \alpha^n = d_1$ for a repeller, we really mean that

$$d_1 s \alpha^n = d \ell^n s = rs = d_1 \,.$$

We don't define the image of an interval vertex u of D under α, because there is no guarantee that α maps $[\![u]\!]$ linearly to its image: perhaps $[\![u]\!]$ contains a breakpoint. Thus there is no need for α to act nicely on the vertices of $\mathcal{T}_2$; we only have a partial action.

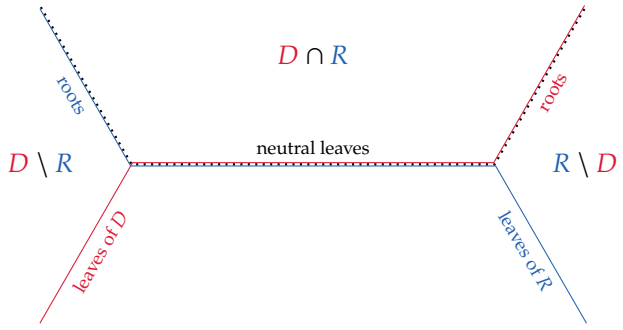

Figure 2.8: A schematic overview of a tree pair (D, ℓ, R). The leaves of the intersection are denoted by a dotted black line.

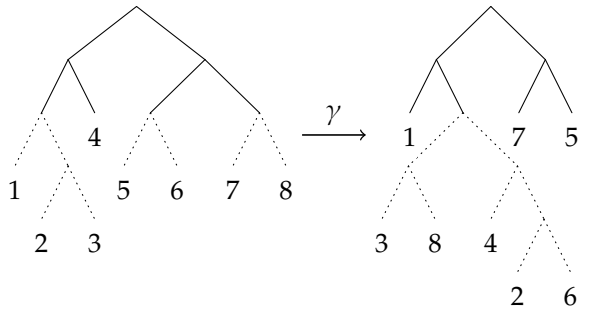

Figure 2.9: In this tree pair diagram, the leaf with address `000` is a repeller and the leaf with address `0110` is an attractor. Note that the components of $D \setminus R$ below addresses `10` and `11` do not contain repellers.

Definition 2.2.4. A tree pair (D, ℓ, R) is called *revealing* if every connected component of $D \setminus R$ contains a repeller and every connected component of $R \setminus D$ contains an attractor. The leaves of $D \setminus R$ which are not attractors are called *sources*; dually the leaves of $R \setminus D$ which are not repellers are called *sinks*.

The tree pair for γ in Figure 2.9 is *not* revealing: the components of $D \setminus R$ with leaves labelled 5, 6 and 7, 8 do not contain repellers. Fortunately, Brin explains[21] that we will always be able to find a revealing pair for a given element.

[21] Brin 2004, Lemmas 10.2–10.6.

Proposition 2.2.5. *There is an algorithm which can construct a revealing pair for a given element $\alpha \in V$.*

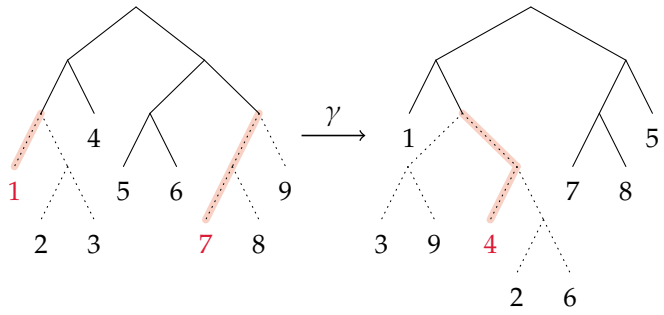

Figure 2.10: Having added a caret to D at `001` and to R at `001`α `= 01`, we obtain a new pair (D', ℓ', R') for α which *is* revealing. We have highlighted the spines in red to emphasise the attractor and repeller leaves.

Any given element has infinitely many revealing pairs, and there need not be a unique minimal revealing pair.[22]

[22] O. P. Salazar-Díaz 2010, Remark after Claim 16.

The killer feature of a revealing pair is that it gives us a complete description[23] of the orbit of any point x in the Cantor space $\partial\mathcal{T}_2$, and hence any point in the interval I. We give a brief overview of the details here.

[23] Brin 2004, Proposition 10.1; O. P. Salazar-Díaz 2010, Claim 5, Lemma 4

Say we have a revealing pair (D, ℓ, R) for $\alpha \in V$, and let $x = x_1 x_2 \ldots$ be a point in the Cantor space. There is exactly one leaf d of D which is a prefix of the infinite string x, so that $x = d y_1 y_2 \ldots$ for some digits y_i. We can partition the leaves of a revealing pair's domain tree into three kinds of (partial) orbits:

1. Finite cycles of neutral leaves under α;

2. IACs coming from a repeller, or leading to an attractor; which leaves
3. IACs coming from a source and leading to a sink. These are known as *source-sink chains*.

The behaviour of d under α and the tail $Y = y_1 y_2 \dots$ completely determine the orbit of x. We will demonstrate this by explicitly computing the translate $x\alpha^M$. To do so, write M in the unique form $M = kn + m$, where $k \in \mathbb{Z}$ and $0 \le m < n$.

Finite cycles The first option is the most straightforward. Suppose $d = d_i$ belongs to a finite cycle of leaves, say $d_1, \dots, d_n$ with $d_n\alpha = d_n\ell d_1$. It follows that $d_i\alpha^n = d_i\ell^n = d_i$, so we may directly compute

$$x\alpha^{kn+m} = (d_iY)\alpha^{kn}\alpha^m = d_i\alpha^{kn}\alpha^m Y = d_i\alpha^m Y = d_{i+m}Y$$

with subscripts modulo n. We see that x belongs to an orbit of size n under α.

Repellers Suppose d belongs to a repeller IAC $d_1, \dots, d_n$ with spine s; say $d = d_i$. First let the tail Y of x be the specific string $\bar{s}$. Observe that

$$\begin{aligned} x\alpha^n &= d_i\bar{s}\alpha^n & &= d_1\alpha^{i-1}\bar{s}\alpha^n \\ &= d_1\bar{s}\alpha^{i-1}\alpha^n & &= d_1(s\bar{s})\alpha^n\alpha^{i-1} \\ &= (d_1s)\alpha^n\bar{s}\alpha^{i-1} & &= d_1\bar{s}\alpha^{i-1} \\ &= d_1\alpha^{i-1}\bar{s} & &= d_i\bar{s} \qquad = x\,. \end{aligned}$$

It follows that $x\alpha^{kn+m} = x\alpha^m = x\alpha^m = d_{i+m}\bar{s}$, for integers $k \in \mathbb{Z}$, $0 \le m < n$ and with subscripts modulo n. Once again, x belongs to an orbit of size n under α. The points in this orbit ($d_i\bar{s}$ for $1 \le i \le n$) are called *repelling periodic points*, since they lie under the IAC of a repeller.

Now assume we have any other tail $Y \ne \bar{s}$, so that x is near[24] to—but distinct from—our previous point. Rather than give a full description of the orbit of x here, we show in the next paragraph that a forward translate $x\alpha^i$ lies below a leaf in a source-sink chain. Then we can deduce the orbit of x once we have described the orbit of points below a source-sink chain below.

Let $p \ge 0$ be maximal with the property that $Y = s^pY'$, and let r be the target of the IAC containing d. We can uniquely write Y' as a concatenation $Y' = zY''$, where z is the address of a leaf of $D \setminus R$ relative to r. Then $z \ne \epsilon$, and we note that $z \ne s$ (by the maximality of p). The orbit of x is the same as the orbit of

$$\begin{aligned} x\alpha^{-(i-1)}\alpha^{(n+1)p} &= d_iY\alpha^{-(i-1)}\alpha^{(n+1)p} \\ &= d_1Y\alpha^{(n+1)p} \\ &= d_1s^pzY''\alpha^{(n+1)p} \\ &= d_1s^p\alpha^{(n+1)p}zY'' \\ &= rzY''\,. \end{aligned}$$

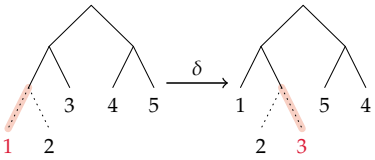

Figure 2.11: This revealing pair for δ exhibits all types of IACs.

- `000` is a repeller, with target `00`;
- `001` is a source, with target the sink `010`;
- `01` is a root of $R \setminus D$, with target the attractor `011`;
- `10` and `11` form a cycle of neutral leaves.

As in Figure 2.10, the spines are highlighted in red.

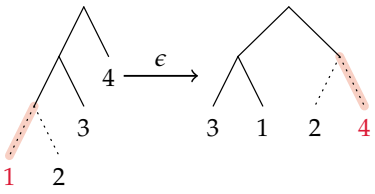

Figure 2.12: This revealing pair for ϵ has an repeller at `000`. Its IAC has length two, moving to `01`, before reaching the target of `00`. The repelling periodic point $\bar{0}$ is sent by ϵ to $0\bar{1}0$ and then to $\bar{0}$ again, forming an orbit of size 2. In contrast, the attracting periodic point $\bar{1}$ is fixed by ϵ.

[24] Let N be the length of the address of d_i. Then $0 < d(x, d_i\bar{s}) \le 2^{-N}$. (See the discussion of metrics after Definition 1.2.12.)

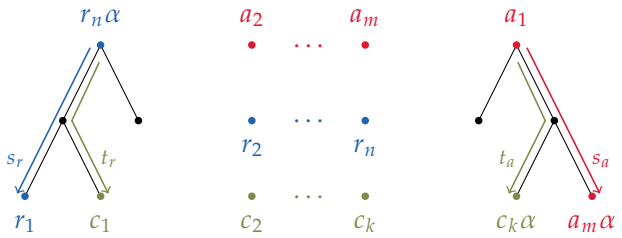

Figure 2.13: A schematic of the nodes involved in a source-sink chain. There are three IACs involved: that of a repeller, an attractor and a source. The left (right) tree stands for a generic component of $D \setminus R$ ($R \setminus D$). The middle vertices stand for neutral leaves. Note that the IACs may have different lengths, despite the picture suggesting otherwise.

This last string sits in the cone below rz. The leaf at this address is a source for the pair (D, l, R), due to the properties of z above.

Attractors The attractors in the revealing pair (D, ℓ, R) for α are precisely the repellers of the revealing pair (R, ℓ^{-1}, D) for α^{-1}. Thus a similar analysis of orbit structure can be undertaken for attractors.

Suppose that a leaf a of R is an attractor with spine s, and let a be the target of the IAC $d_1, \ldots, d_n$. The upshot is that the n points of the form $x = d_i\bar{s}$ belong to a single orbit, and are called *attracting periodic points*. Any other point x' below any d_i has a translate $x'\alpha^M$ which lies below a source in the revealing pair.

Source-sink chains Finally, $d = c_i$ could be part of an IAC $c_1, \ldots, c_k$ starting at a source. It must be the case that $c_k\ell$ is a sink. If not, then $c_k\ell$ is an attractor, meaning c_1 would be a root of $R \setminus D$. But sources are leaves of $D \setminus R$, so c_1 would be both strictly below and strictly above a leaf of R—which is absurd. This explains why sources are joined by their IACs to sinks.

Let's set up our notation (see Figure 2.13).

- Let $c_1, \ldots, c_k$ be the IAC of a source being studied. The target of this, $c_k\ell$, is a sink.
- Let r_1 be the repeller in the same component of $D \setminus R$ as the source c_1. It has an IAC $r_1, \ldots r_n$ and spine s_r. Then $r_n\alpha s_r = r_1$.
- Let $a_1, \ldots, a_m$ be the IAC whose target is the attractor $a_m\alpha$ in the same component of $R \setminus D$ as the sink $c_k\alpha$. Its spine s_a satisfies $a_1 s_a = a_m\alpha$.

We now compute the orbit of the leaf c_i under α. By definition $c_i\alpha^j = c_{i+j}$ for $1 - i \leq j \leq k - i$. Continuing forward, the next image leaf is $c_k\alpha = a_1 t_a$. Let $M \geq 0$ be an integer written uniquely in the form $M = pm + j$, for $p \in \mathbb{Z}$ and $0 \leq j < m$. We compute

$$c_i\alpha^{k-i+1+M} = c_k\alpha\alpha^M = a_1 t_a \alpha^M = a_1\alpha^j\alpha^{pm} t_a = a_j s_a^p t_a \,.$$

For the backward orbit of c_1, set $N = pn + j > 0$ with $p \in \mathbb{Z}$ and

$0 \le j < n$. Then

$$c_i\alpha^{1-i-n-N} = c_1\alpha^{-n-N} = r_n\alpha t_r\alpha^{-n}\alpha^{-N} = r_1\alpha^{-j}t_r\alpha^{-pn}$$
$$= r_{1-j}\alpha^{-pn}t_r = r_{1-j}s_r^p t_r\,,$$

with subscripts modulo n. We see that each image $c_i\alpha^M$ is a distinct leaf of $\mathcal{T}_2$, so c_i and any point $x = c_iY$ below it lie in an infinite orbit.

We can say more about the orbit of x. In the forward direction (again with $M = pm + j$), we move to

$$x\alpha^{k-i-+1+M} = c_i\alpha^{k-i-+1+M}Y = a_j s_a^p t_a Y\,.$$

This limits on the periodic sequence whose $(pm + j)$th term is $a_j\overline{s_a}$. Specifically, the base two logarithm of the term-to-term distance is at most $-|a_j| - p|s_a|$, so this distance tends to zero as $p, M \to \infty$. But this periodic sequence is an orbit of repelling periodic points, as we saw earlier! Similarly, the backwards orbit moves through the points

$$x\alpha^{1-i-n-N} = c_i\alpha^{1-i-n-N}Y = r_{1-j}s_r^p t_r Y$$

(with $N = pn + j$ as above). The sequence with $(pn + j)$th term $r_{1-j \bmod n}\overline{s_r}$ is an orbit of attracting periodic points. The base two logarithm of the term-to-term distance is $-|r_{1-j \bmod n}| - p|s_r|$, so again the distance tends to zero.

We see that the orbit of x converges to (and emerges from) a limit cycle. Informally, points x are pushed out from an orbit of repelling fixed points; pushed forward below a source-sink chain for a finite time; then they are drawn forever towards attracting fixed points.

We summarise with a proposition.

Proposition 2.2.6. *Let $\alpha \in V$, and think of α as a homeomorphism of the Cantor space $\partial\mathcal{T}_2$ or interval I. We can fully describe the orbit of $x \in \partial\mathcal{T}_2$ (hence $x\pi \in I$) given a revealing pair for α. In particular:*

- *There are finitely many maximal cones (maximal standard dyadic intervals) which have a finite orbit under α. Each of these is fixed pointwise by a power of α.*
- *Excluding these, there are finitely many isolated points with a finite orbit under α. These are classified as either attracting or repelling.*
- *This finite amount of data can be determined algorithmically by constructing a revealing pair.*
- *All other points belong to an infinite orbit under α. Their orbits move away from repelling periodic orbits and towards attracting periodic orbits.*

Remark 2.2.7. Suppose we have a source-sink chain leading from a repeller r to an attractor a. Let x and y be the associated attracting and repelling fixed points. Since these points are 'linked' by a source-sink chain, we might guess that they must have the same period. This is

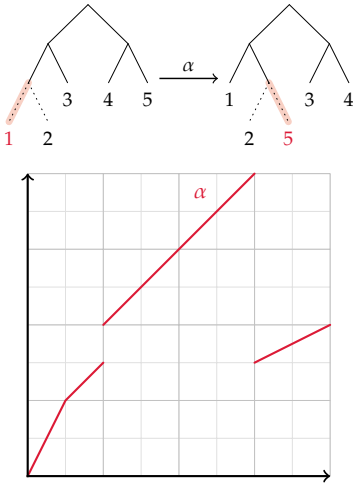

Figure 2.14: The repelling fixed point $\mathtt{00\bar{0}}$ is fixed by α, and the attracting fixed point $\mathtt{01\bar{1}}$ has order 3 under α. This is despite the fact that they are linked by the source-sink chain $\mathtt{001} \mapsto \mathtt{010}$.

false, as Figure 2.14 shows. With that said, under an element $\alpha \in T$ all periodic points must have the same order—see Proposition 2.1.9.5.

Revealing pairs enable us to give a concise proof of the next lemma. It was proven using revealing pairs by Bleak et al.[25], who attribute the result to Burillo et al.[26]

[25] Bleak, Bowman et al. 2013, Lemma 4.4.
[26] Burillo et al. 2009, Proposition 6.1.

Lemma 2.2.8. *An element $\alpha \in V$ has finite order if and only if it can be described by a tree pair whose domain and range tree are equal.*

Proof. Let (D, ℓ, R) be a revealing pair for a finite order element α. This pair cannot have any attractors or repellers, or else some point x would have an infinite orbit under α. Thus both $D \setminus R$ and $R \setminus D$ are empty, so $D = R$.

On the other hand if (D, ℓ, D) is a tree pair for α, then ℓ is a permutation of the leaves of D. Thus α has finite order equal to the order of ℓ. □

Remark 2.2.9. A similar argument shows that every element $\alpha \in V$ must have a periodic point. Choose a revealing pair (D, ℓ, R) for α. If this pair has a cycle of neutral leaves, then any point below any of these leaves lies in a finite orbit. Otherwise there are no such cycles. Then $D \neq R$, so $D \setminus R$ and $R \setminus D$ are nonempty. Then the pair has at least one attractor and at least one repeller, so α has an (attracting or repelling) periodic point.

2.3 Behaviour near periodic points

The simplest possible behaviour in a dynamical system is a point x which the system moves back to itself after a finite time (possibly immediately). These points have much to tell us about the system in question. From a group-theoretical viewpoint in particular, suppose that $\alpha^\gamma = \beta$ and that $x\alpha^k = x$. Then $x\gamma = x\alpha^k\gamma = x\gamma(\gamma^{-1}\alpha^k\gamma) = x\gamma\beta^k$. In words, points in a finite orbit under α get sent by conjugators into a finite orbit of the same size under β. There's nothing special about Thompson's group here—this is just how conjugation interacts with a group action. The point is that a group action can reveal information about the group which is hard to discern from a purely algebraic point of view.

In this section, we will use the knowledge we gain from revealing pairs to describe what happens near to periodic points. We saw there were two types of periodic points—those under cycles of leaves, and others below attractors and repellers. We distinguish these two types topologically, then go on to study the second kind—called 'important points'—in particular. Our aim is to describe in more detail how important points influence the dynamics of elements of T.

We will denote by $\mathrm{Fix}(\beta)$ the points of the interval fixed by $\beta \in V$. When working in T, the same notation will refer to fixed points in the circle.

2.3.1 *Important points for V*

Let d be an attractor or repeller with spine s for $\alpha \in V$. Then $x = d\bar{s}$ belongs to a finite orbit under α of attracting or repelling periodic points. We saw that all infinite orbits lie below source-sink chains. These diverge from repelling periodic points and converge to attracting periodic points. Assembling this information into a graph yields the *train tracks* and *flow graphs* introduced by Bleak et al.[27] They are useful for summarising the global dynamics of an element of V. Up to the right notion of equivalence, flow graphs in particular serve as a conjugacy invariant.

[27] Bleak, Bowman et al. 2013, Sections 4.2–4.3.

In short, the attracting and repelling periodic points are very significant in determining the orbits of an element α. Let us give them a name.[28]

[28] Bleak and O. Salazar-Díaz 2013, after Corollary 2.3

Definition 2.3.1. Let (D, ℓ, R) be a revealing pair for $\alpha \in V$. The *important points* I_α of α are the points of the form $u\bar{s}\alpha^i \in \partial\mathcal{T}_2$, where i is an integer, and u is an attractor or repeller with spine s. All other points $y \in \partial\mathcal{T}_2 \setminus I_\alpha$ are called *unimportant*.

This definition depends on a choice of revealing pair for α. Salazar-Díaz's thesis[29] describes how all revealing pairs for α can be obtained from one another, via moves called 'rollings'. If R is a revealing pair for α, it can be checked that any rolling S of R yields the same set of important points as R. Thus we are justified in referring to these points as being important points *of* α, without reference to a revealing pair.

[29] O. P. Salazar-Díaz 2010, Section 3.5.

Why do we care about important points in Cantor space? We could use the map $\pi\colon \partial\mathcal{T}_2 \twoheadrightarrow I$ (from Section 1.2.3) to project from the Cantor space to the interval, and then look for analogous points in I. However, doing so gives us a blurrier picture with less information. This is part of why Cantor space is the natural space for studying V

Let's try an example: consider α from Figure 2.15. The point $x = \mathtt{01}\bar{\mathtt{1}}$ is important for α, and in particular x is fixed by α. On the other hand take $y = \mathtt{100}\bar{\mathtt{0}}$. Its forward orbit is

$$\mathtt{100}\bar{\mathtt{0}} \mapsto \mathtt{110}\bar{\mathtt{0}} \mapsto \mathtt{111}\bar{\mathtt{0}} \mapsto \mathtt{1111}\bar{\mathtt{0}} \mapsto \cdots \mapsto \mathtt{1}^n\bar{\mathtt{0}} \mapsto \ldots;$$

this contains infinitely many elements, and so y is not an important point for α We see that x and y have different dynamical behaviours. In particular, x is important whereas y is not. This distinction is lost if we move to the interval, because the images $x\pi = y\pi = 1/2$ are the same. To recover the two different behaviours, we need to spend extra effor taking one-sided limits. In this example we have $(1/2)^+\alpha = 3/4$ and $(1/2)^+\alpha' = 1$ from the right, whereas $(1/2)^-\alpha = 1/2$ and $(1/2)^-\alpha' = 1/2$ from the left.

THE NEXT CLAIM follows immediately from the study of revealing pairs earlier: in particular, from Proposition 2.2.6 and Lemma 2.2.8.

Claim 2.3.2. *There are always finitely many important points I_α of an element $\alpha \in V$. Moreover, $|I_\alpha| = 0$ if and only if α has finite order.*

In the next lemma, we give a topological characterisation of an element's important points.

Lemma 2.3.3. *Let $\alpha \in V$, and let P be the set of points in $\partial\mathcal{T}_2$ in finite orbits under α. Then I_α is the set of isolated points of P.*

Proof. First we show that important points x are isolated in P; then we show that unimportant points $y \in P$ are not isolated in P.

Take x to be an important point. Then we can write $x = d\bar{s}$, for some leaf address d such that either $ds\alpha^n = d$ or $d\alpha^n = ds$. To show x is an isolated point of P, we must produce a neighbourhood of x which contains no other points of P. We demonstrated in Proposition 2.2.6 that $[d]$ is such a neighbourhood, because all points in $[d] \setminus \{x\}$ belong to infinite orbits (which either attract to or repel from x). Thus important points are isolated in P.

The unimportant points in finite orbits $y \in P$ are those below a leaf e in a neutral cycle of leaves. To show y is not an isolated point of P, we must show that any neighbourhood U of y contains a point $z \in P$ distinct from y. Such a neighbourhood U must contain an open set of the form $[u]$, where u is a prefix of y. Let v be the longer of the two words e and u, so that $y \in [v] \subseteq [e] \cap U$. Then any point $z \in [v] \setminus \{y\}$ is in $[e]$, so must have a finite orbit. Hence $z \in P$, which shows that y is not isolated in P. □

The following lemma uses the characterisation to identify important points of powers α^i—without constructing a revealing pair for α^i.

Lemma 2.3.4. *Let $\alpha \in V$ and $0 \neq i \in \mathbb{Z}$. Then the important points $I(\alpha)$ and $I(\alpha^i)$ are equal.*

Proof. For a fixed value of i, let P_i denote the points $x \in \partial\mathcal{T}_2$ in finite orbits under α^i. The following statements are equivalent.

- $x \in P_1$.
- $x\alpha^n = x$ for some $n \neq 0$.

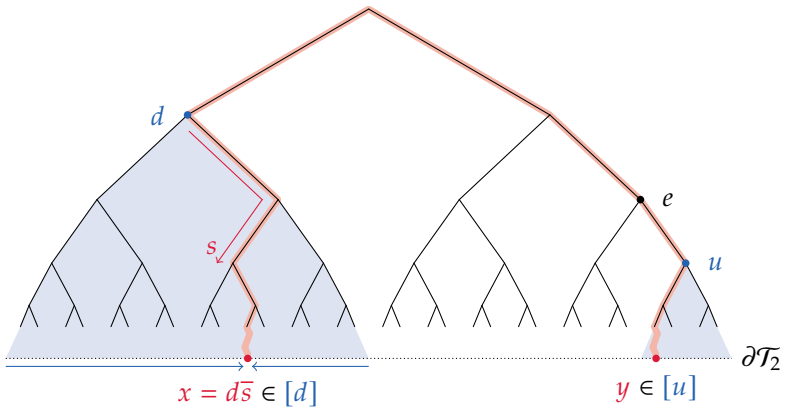

Figure 2.16: Schematic illustrating the proof of Lemma 2.3.3. In this example we have $u = e1$, so that e is a proper prefix of $u' = u$. The horizontal arrows are meant to suggest that $[d]$ is attracted towards the attractor x.

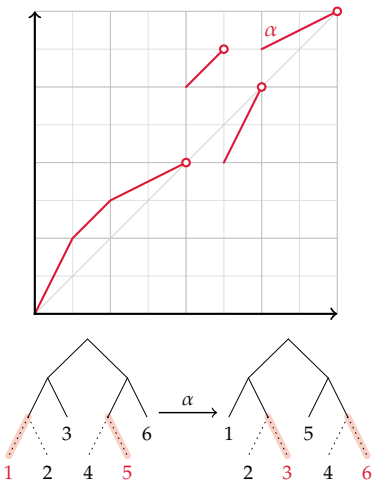

Figure 2.15: **Above:** a revealing pair for an element $\alpha \in V$. **Below:** the element α viewed as a right-continuous map $I \to I$.
The two preimages $(1/2)\pi^{-1} \subseteq \partial\mathcal{T}_2$ have different dynamical behaviour: one is fixed, whereas the other is attracted to $11\bar{1}$.

- $x\alpha^{mi} = x$ for some $m \neq 0$.
- $x \in P_i$.

This shows that $P_i = P_1$. Take the subset of isolated points on both sides to see that $I(\alpha^i) = I(\alpha)$. □

We conclude with a short—but important—observation about important points.

Lemma 2.3.5. *An element $\alpha \in V$ permutes its important point set I_α.*

Proof. This is immediate from Definition 2.3.1: if $u\bar{s}\alpha^i$ is important, so too is $u\bar{s}\alpha^{i+1}$.

Alternatively, we could make use of the topological characterisation above. Our element α must map its set of periodic points P to P, or else the points in P wouldn't be periodic! The restriction $\alpha\big|_P : P \to P$ of α to P is a homeomorphism, and so must preserve all topological properties of P. In particular, $\alpha\big|_P$ preserves the isolated points of P. □

2.3.2 *Important points for T*

In the rest of this thesis, we're going to work with the action of T on the circle S^1. We'll still want to use important points, along with all the information we have at our disposal from the study of V. To keep things in order, it's worth establishing definitions specific to T and S^1.

Definition 2.3.6. Let $\pi\colon \partial\mathcal{T}_2 \twoheadrightarrow I$ be the projection map from Section 1.2.3, and let $\phi\colon I \twoheadrightarrow S^1$ be the projection map which identifies 0 and 1. Define $\pi_0 = \pi\phi\colon \partial\mathcal{T}_2 \to S^1$ to be their composition. The *important points* of an element $\alpha \in T$ are the points in the image $I_\alpha\pi_0 \subset S^1$. All other points $y \in S^1$ are called *unimportant*.

When it is clear from context that we are working with the action $S^1 \circlearrowleft T$ of T on the circle, we shall simply call these the *important points* of α. In this circumstance, we'll drop the map π_0 and refer to the important point set as just I_α, or occasionally $I(\alpha)$ if α is replaced by a complicated expression.

Much of what we know about important points from Proposition 2.2.6 applies to important points too.[30] We saw that points $x \in \partial\mathcal{T}_2$ in infinite orbits move away from and converge to the orbits of important points (for V). As the projection map $\pi_0\colon \partial\mathcal{T}_2 \twoheadrightarrow S^1$ is continuous, it must preserve limits. Thus we see that points $x\pi_0 \in S^1$ in infinite orbits *for T* move away from and converge to the orbits of important points *for T*.

[30] Though as noted after page 36, the picture becomes a little blurrier after projecting to S^1.

The following is another example where important points exhibit the same behaviour in T and V.

Lemma 2.3.7. *Let $\alpha \in T$.*

1. *I_α is finite; it is empty if and only if α has finite order.*

2. *Let $0 \neq i \in \mathbb{Z}$. The important points of $I(\alpha)$ and $I(\alpha^i)$ are equal.*

3. *The element $\alpha \in T$ permutes its important point set I_α.*

Proof. These follow from Claim 2.3.2, and Lemmas 2.3.4 and 2.3.5, respectively. □

Just like in V, we have a topological characterisation of important points for T.

Lemma 2.3.8. *Let $I(\alpha')$ be the important points (for T) of an element $\alpha' \in T$. Then $I(\alpha') = \partial P'$, where $P' \subseteq S^1$ is the set of points in finite orbits under $\alpha' \in T$.*

Proof. In this proof, primed symbols refer to the circle S^1, and unprimed symbols refer to the Cantor space $\partial\mathcal{T}_2$. We can view $\alpha' \in T$ as an element $\alpha \in V$, where the latter acts on Cantor space $\partial\mathcal{T}_2$. These actions $S^1 \circlearrowleft \alpha'$ and $\partial\mathcal{T}_2 \circlearrowleft \alpha$ are compatible because[31] $\alpha\pi_0 = \pi_0\alpha'$.

First we show that $I(\alpha') \subseteq \partial P'$. Take an important point $x \in \partial\mathcal{T}_2$ for $\alpha \in V$, so that $x\pi_0 \in S^1$ is important for $\alpha' \in T$. Now x must be of the form $x = d\bar{s}$, for some finite strings s and d (c.f. Lemma 2.3.3). To show that $x\pi_0 \in \partial P'$, we must show that any neighbourhood U' of $x\pi_0$ contains a point not in P'. As π_0 is continuous, the inverse image $U = U'\pi_0^{-1}$ is a neighbourhood of x. This means there is some basic open set $[u]$ with $x \in [u] \subseteq U$. Let v be the longer of the two words d and u, so that $x \in [v] \subseteq [d] \cap U$. Our study of revealing pairs shows that every point $z \in [v] \setminus \{x\} \subseteq U$ belongs to an infinite orbit (repelled from or attracted to x). Thus our neighbourhood U' contains a point $z\pi_0 \in U\pi = U'$ in an infinite orbit,[32] which establishes $x \in \partial P'$.

Let $\mathrm{rot}(\alpha') = p/q$ in lowest terms. Then all finite orbits under α' (and α) have size q, by Proposition 2.1.9.5. This means that $P' = \mathrm{Fix}(\alpha'^q)$. Now P' is closed, because fixed point sets of continuous maps $\beta\colon X \to X$ in a Hausdorff space X are closed. Hence $\partial P' \subseteq P'$.

To finish the proof, let $y' \in S^1$ be unimportant for α'. If $y' \notin P'$ then $y' \notin \partial P'$ (by the paragraph above), so assume otherwise that $y' \in P'$ has a finite orbit. We will explain how to find a neighbourhood U' of y' contained within P'. Our argument breaks into two cases.

Suppose y' is nondyadic. Then y' has exactly one preimage $y \in y'\pi_0^{-1} \subseteq P$. We know y is unimportant and has a finite orbit,[33] so $y \in [e] \subseteq P$ lies below a leaf e in a cycle of neutral leaves. Applying π_0, we see that $y' \in [\![e]\!] \subseteq P'$. Because y' is nondyadic, y' belongs to the interior of $[\![e]\!]$. Thus we may take $U' = [\![e]\!]^\circ$.

Suppose y' is dyadic. Then y' has two preimages: $y_{\mathtt{0}}$ with tail $\bar{\mathtt{0}}$, and $y_{\mathtt{1}}$ with tail $\bar{\mathtt{1}}$. Again, each y_i is unimportant and belongs to a finite orbit so, each $y_i \in [e_i] \subseteq P$ lies below a leaf e_i in a neutral cycle of leaves. Applying π_0, we see that $y' \in [\![e_i]\!]$ for both i. Set $U' = ([\![e_{\mathtt{0}}]\!] \cup [\![e_{\mathtt{1}}]\!])^\circ$. If $y' \in [\![e_i]\!]^\circ$ for some i, then $y \in U'$. Otherwise y' is the right endpoint of $[\![e_{\mathtt{1}}]\!]$ and the left endpoint of $[\![e_{\mathtt{0}}]\!]$, so again $y \in U'$.

[31] That is, α and α' are topologically semiconjugate (via π_0).

[32] The infinite orbit O_z cannot project onto a finite set under π_0, because each point in the circle has at most two preimages under π_0. (Dyadics have two preimages; nondyadics have one.)

[33] Argue as in sidenote 32: the orbit of y cannot be infinite, as could not project onto the finite orbit of y' under π_0.

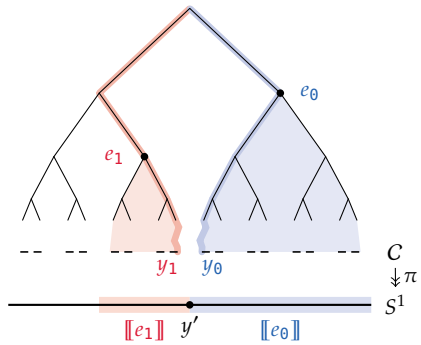

Figure 2.17: Schematic illustrating the case of a dyadic unimportant point y' with two preimages $y_i \in \partial\mathcal{T}_2$. We've drawn $\partial\mathcal{T}_2$ as the middle-thirds Cantor set C here, to emphasise the distance between $y_{\mathtt{0}}$ and $y_{\mathtt{1}}$.

In both cases, there is an open set U' contained in P' and containing y'. This shows that $y' \in (P')^\circ$, and that $y' \notin \partial P'$ in particular. □

Remark 2.3.9. Let $\alpha' \in T$ have reduced rotation number p/q. We argued in the proof above that $P' = \mathrm{Fix}(\alpha^q)$, and so we have $I_\alpha = \partial \mathrm{Fix}(\alpha^q)$.

With this remark in mind, whenever Kassabov and Matucci work with $\partial \mathrm{Fix}(\alpha)$ in F, we will work with $\partial \mathrm{Fix}(\alpha^q)$ in T.

With that out of the way, which points can be important points for an element $\beta \in T$? As noted above, $I_\beta = \partial \mathrm{Fix}(\beta^q)$ where q is the reduced rotation number denominator. To simplify things, we put $\alpha = \beta^q$ and look for important points $z \in \partial \mathrm{Fix}(\alpha)$.

Because our functions α are piecewise linear, there are two ways they can have a fixed point. Firstly, a linear segment of α could have gradient $m = 1$ and run along along the diagonal $y = x$. Alternatively, a segment's gradient could be $m \neq 1$ and cross the diagonal at exactly one point. Thus our fixed point set's connected components are either fixed intervals or fixed singletons. What are the boundaries of these components? At the endpoints of a fixed segment we have a breakpoint, so the endpoints must be dyadics, and in particular they must be rational. On the other hand, the isolated fixed points are their own boundary. Where can they occur?

Take a linear segment of α with gradient $M = 2^m \neq 1$ and initial breakpoint (x_0, y_0). Its points (x, y) satisfy the equation $y - y_0 = M(x - x_0)$, where x_0 and y_0 are dyadics. We locate a fixed point z by putting $x = y = z$; rearranging shows that[34]

$$z = (y_0 - 2^m x_0)/(2^m - 1) = s/(2^m - 1)$$

is the dyadic integer s divided by an integer. Again we see that z is neccessarily rational.

[34] In the more general setting of $\mathrm{PL}_{S,G}$, isolated fixed points take the form $s/(g-1)$, where $S \in S$ and $1 \neq g \in G$. Kassabov and Matucci call the set of such points Q_S. See Kassabov and Matucci 2012, Theorem 6.12.

Can *every* rational in the interval be a important fixed point? The answer is 'yes', as the next lemma shows.

Lemma 2.3.10. *A point $z \in S^1 = [0,1]/\{0 \sim 1\}$ is a important fixed point of some element of T if and only if z is rational.*

Proof. We have just seen that any such fixed point z must be rational; now we explain how to build an element $\alpha \in T$ with z as a important fixed point. To keep things simple, we'll actually construct an element $\alpha \in F$. First note that $\partial \mathrm{Fix}_{S^1}(x_0) = \partial\{0\} = \{0\}$, where x_0 is the element defined in Equation (1.2) and Figure 1.1. So 0 can be a important fixed point. For the rest of the proof, we take z to be a rational number in $(0,1)$.

Write z as a repeating base two number $z = 0.a\bar{b}$, where $a = a_1 \ldots a_n$ and $b = b_1 \ldots b_m$ are finite binary strings with $n \geq 0$ and $m \geq 1$. Then $z = a\bar{b}\pi$. Invoke Lemma 2.1.3 to build an element $\alpha \in F$ which linearly maps $[\![ab]\!]$ to $[\![ab^2]\!]$. It follows that α fixes z, because

$$z\alpha = ab\bar{b}\alpha\pi = ab\alpha\bar{b}\pi = ab^2\bar{b}\pi = a\bar{b}\pi = z\,.$$

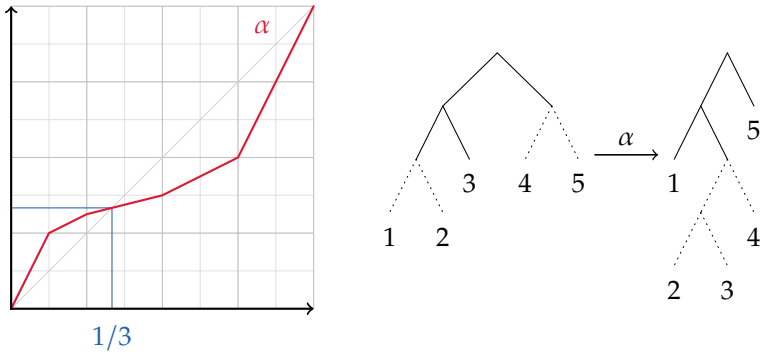

Figure 2.18: We illustrate the proof of Lemma 2.3.10. Since $1/3 = \sum_{n \in \mathbb{N}} 4^{-n}$ we can write $1/3 = 0.\overline{01} = 0.\overline{0101}$. We construct a tree pair for F where `01` and `0101` are leaves of the domain and range trees, respectively. The resulting element α fixes the point $1/3$.

Now we explain why $z \in \partial \operatorname{Fix}(\alpha)$ is important. To do so we need to explain why every neighbourhood of z contains points not fixed by α. At least one of the one-sided derivatives (say in the positive direction) of α at z is $2^{-m} < 1$. If z is not dyadic, then α must be continuous at z. The other one-sided derivative (say in the negative direction) must also be 2^{-m}, or else z would be a breakpoint. Then no point $z \pm \epsilon$ in the immediate vicinity of z is fixed by α. Otherwise z is dyadic. We can modify[35] α outside of $[\![ab]\!]$ so that it is continuous and has a gradient not equal to 1 in the other (say negative) direction. Then no point $z - \epsilon$ in the immediate vicinity of z is fixed by α. □

[35] E.g. use Lemma 2.1.3.

Corollary 2.3.11. *A point $z \in S^1 = [0,1]/\{0 \sim 1\}$ is a important point of some element $\beta \in T$ if and only if z is rational. The period of z under β can be any positive integer.*

Sketch proof. As discussed above Lemma 2.3.10, to study important points it suffices to study important points which are fixed. So our first assertion follows from the same lemma.

We can build an element $\beta \in T$ with z as a important point of period n as follows. Choose a tree A with leaves $a_1, \ldots, a_n$. Then write $z = a_i t \bar{s} \pi$ as an infinite binary string, for some i and some strings t, s with $|s| \geq 1$. Use Lemma 2.3.10 to find an element $\alpha \in F$ for which $w = t\bar{s}\pi$ is a important fixed point of α. Choose a tree pair (D, R) for α, say with m leaves each.

Form new a tree D' (resp. R') by attaching a copy of D (R) to each a_i. The new trees both have mn leaves in total. Then define β to have the tree pair (D', ℓ, R'), where ℓ sends the ith leaf of D' to the $(i+m)$th leaf of R'. Then z belongs to an orbit of size n under β. Moreover, points near z are attracted to or repelled from z by β^n; this occurs because points near w are attracted to or repelled from w by α. □

The key thing to take away from this result is that any rational can be important—even a nondyadic. This is a potential source of difficulty, since (broadly speaking) it is more difficult to work with and around nondyadic points z. We have to spend extra effort ensuring that z (or any other point depending on z) is *not* a breakpoint of any elements we're considering.

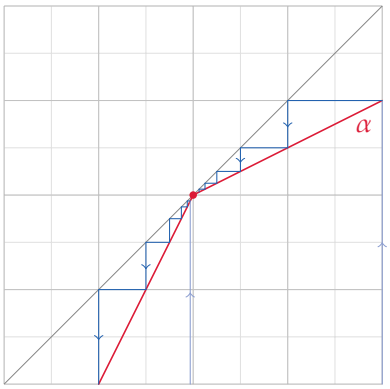

Figure 2.19: An illustration to complement the proof of Lemma 2.3.12. In this sketch, a function α has two different gradients $x^+\alpha' = 1/2$ and $x^-\alpha' = 2$ either side of a fixed point x. The grid represents the square $[0,1]^2$.

Points x_0 are fed into the function from the x-axis. These 'staircase diagrams' show the iterations $x_0\alpha^n$ starting with $x_0 = 63/128$ and $x_0 = 1$, for $1 \leq n \leq 5$. Gradients less than 1 cause attraction to fixed points, whereas those greater than 1 repel points away from x.

2.3.3 *Partitioning at important points*

Revealing pairs showed us that infinite orbits in $\partial\mathcal{T}_2$ jump all over the place, as they diverge from and converge to limit cycles—namely, the orbits of important points for V. If we consider a suitable power α^q instead of α, this 'jumping' vanishes. Instead, we find that intervals are stretched away from repelling important points and squashed towards attracting important points. Limit cycles of size n become separated into n separate, stationary limit points—namely, the important points themselves.

We're going to mirror this approach, using important points for T instead of for V, and using the circle S^1 instead of Cantor space $\partial\mathcal{T}_2$. Let $\alpha \in T$ be a generic element. We know α has finitely many important points. Enumerate them and sort them into circular order, say as $r_0 \to \cdots \to r_{n-1} \to r_0$ and let $P_i = [r_i, r_{i+1}]$ for each $0 \leq i < n$. Let $\mathrm{rot}(\alpha) = p/q$ in lowest terms, so that $I(\alpha^q) = I_\alpha$. If we have $n = 0$ important points, then $\alpha^q = \mathrm{id}$, so the entire circle remains unchanged by α^q.

Otherwise $n \geq 1$, so assume now that we have at least one important point and cell. Each cell P_i is fixed setwise by α^q, and may be equipped with a linear order $<_{r_i}$ formed by cutting the cyclic order on S^1 at r_i.[36] For each cell, we will show that one of its endpoints either attracts, repels, or leaves unchanged the points within P_i. The other endpoint must then repel, attract or leave unchanged the points of P_i. Note that if $n = 1$ then these two endpoints are the same point on the circle!

Lemma 2.3.12. *Let α have rotation number p/q in lowest terms. Partition S^1 into cells $P_i = [r_i, r_{i+1}]$ as above, and write $m_i = r_i^+(\alpha^q)'$. Points in the cell interior (r_i, r_{i+1}) are*

- *fixed pointwise by α^q if $m_i = 1$,*
- *repelled from r_i and attracted to r_{i+1} by α^q if $m_i > 1$, and otherwise*
- *repelled from r_{i+1} and attracted to r_i by α^q if $m_i < 1$.*

Expressed in other words, this result says that the cell behaviour under α^q is determined by the right-gradient $r_i^+(\alpha^q)'$ at the start of each cell. See Figures 2.20 and 2.21 for a visualisation.

[36] See Definition 1.1.7, and Equation (1.4) in particular.

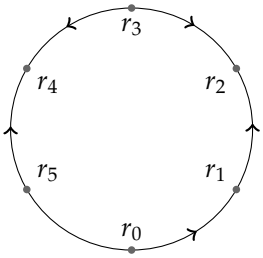

Figure 2.20: The forward orbits of α^q will either leave points fixed—as on the arc $[r_5, r_0]$—or move them away from one and towards another important point. Each important point may attract, repel, or have no influence on its left and right side—see Remark 2.3.13 for more on this.

In this example, r_0 and r_5 repel on one side and attract on neither; r_1 attracts on one side and repels on the other; r_2 and r_4 attract on both sides; and r_3 repels on both sides.

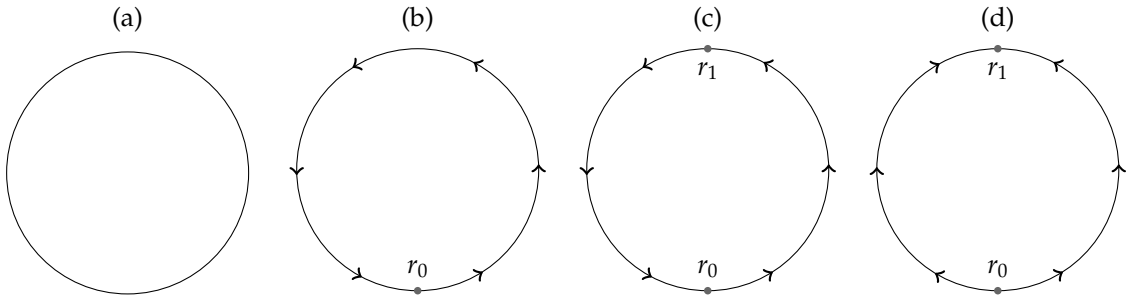

Figure 2.21: The circle can be divided up into cells P_i by the n important points r_i of an element $\alpha \in T$.

(a) Here $n = 0$ and $\alpha^q = \mathrm{id}$. The entire circle is fixed pointwise by α^q.

(b) Now $n = 1$ so $\alpha^q \neq \mathrm{id}$. Points are repelled from r_0 to the right, only to be attracted to r_0 from the left.

(c) Now $n = 2$. Both important points attract material from their left and repel it to their right. Over time, the powers of α^q will send all unimportant points clockwise until they approach a important point.

(d) Again $n = 2$, but this time points are repelled on both sides by r_0, and attracted to r_1 from both sides.

Proof. Suppose for starters that $m = 1$. Then α^q has to fix the entire cell P_i. If not, let $x = \inf\{\, y \in (r_i, r_{i+1}) \mid y\alpha^q \neq y \,\}$ be the infimum of points not fixed by α^q. Then x must be a breakpoint, with $x^-(\alpha^q)' = 1$ and $x^+(\alpha^q)' \neq 1$. But then x would belong to $\partial \operatorname{Fix}(\alpha^q)$, so $x \in (r_i, r_{i+1})$ would be a important point. As we have already sorted I_α into cyclic order, no point in the cell interior is a important point: a contradiction.

Next is the case when $m > 1$. View the cell P_i as a linearly ordered set[37] Let s be the first breakpoint of α^q to the right of r_i. Then α^q is locally defined by $(r_i+\epsilon)\alpha^q = r_i + m\epsilon$, for all ϵ in the range $r_i \leq r_i+\epsilon \leq s$. Choose such an ϵ and fix a point $x = r_i + \epsilon$. As $m > 1$ it follows that $x\alpha^{qN} = r_i + m^N\epsilon$ for any $N \leq 1$. This converges to r_i as $N \to -\infty$, showing us that x is repelled from r_i.

What happens as $N \to \infty$? We certainly have $r_i < x < x\alpha^q < r_{r+1}$. As α^q preserves the linear order on the cell P_i, we know $r_i < x\alpha^q < x\alpha^{2q} < r_{i+1}$. Continue by induction to see that $r_i < x < \cdots < x\alpha^{qN} < r_{i+1}$ for all $N > 0$. Now $x\alpha^{qN}$ must converge to a important point as $N \to \infty$. Since the (linear) distance $|x\alpha^{qN} - r_i|$ increases with N, the forward orbit of x cannot converge to r_i; instead, the orbit must be attracted to r_{i+1}.

We now know that any point $y \in (r_i, r_{i+1})$ is contained in the interval $(x\alpha^{-Mq}, x\alpha^{Mq})$, for some $M > 0$. Since α^q is orientation preserving, the translate $y\alpha^N$ belongs to $(x\alpha^{-Mq+N}, x\alpha^{Mq+N})$. Taking limits as $N \to \pm\infty$, we see that the entire cell interior (r_i, r_{i+1}) is attracted to r_{i+1} and repelled from r_i.

The case $m < 1$ is similar; we could even replace α with α^{-1} and argue as above. Then every point in (r_i, r_{i+1}) is attracted to r_i and repelled from r_{i+1}. □

Remark 2.3.13 (Dyadics versus nondyadics). There exists an element $\alpha \in T$ with only nondyadic important points—see Figure 2.22 on page 45.

If a important point r is attracting (repelling) on one side and not on the other, then the gradients $r^-\alpha^{q\prime}$ and $r^+\alpha^{q\prime}$ must be different. So r is a breakpoint, and hence must be dyadic. Taking the contrapositive, we see that nondyadic important points r must have the same gradients either side. It's impossible for r to be neither attracting or repelling on both sides—then it wouldn't be important.[38] So nondyadic important points must either attract on both sides or repel on both sides.

[37] See Equation (1.4).

[38] See Remark 2.3.9.

2.3.4 *Gradients*

One of the first comments in the part of Matucci's thesis which deals with conjugacy[39] is that taking gradients at fixed points results in a homomorphism. Unfortunately, we don't necessarily have any fixed points in T, so it's not clear that we can use this approach. What can we do instead? We start with the calculation behind all of this, which boils down to careful use of the chain rule.

[39] Matucci 2008, Section 4.1.2.

Claim 2.3.14. *Let α and γ be right-differentiable (left-differentiable) functions under composition with positive gradients. Suppose that $x\alpha = x$ is a fixed point and that γ is invertible. Then the right (left) derivatives $[x\gamma]^{\pm}(\alpha^{\gamma})'$ and $x^{\pm}\alpha'$ are equal.*

Proof. The chain rule tells us that

$$(\alpha^{\gamma})' = (\gamma^{-1}\alpha\gamma)' = (\gamma^{-1})' \cdot (\gamma^{-1}\alpha') \cdot (\gamma^{-1}\alpha\gamma'),$$

so evaluating[40] the right-hand derivatives at $x\gamma$ yields

$$[x\gamma]^{+}(\alpha^{\gamma})' = [x\gamma]^{+}(\gamma^{-1})' \cdot ([x\gamma\gamma^{-1}]^{+}\alpha') \cdot ([x\alpha]^{+}\gamma').$$

The right-hand side simplifies to

$$[x\gamma]^{+}(\alpha^{\gamma})' = [x\gamma]^{+}(\gamma^{-1})' \cdot (x^{+}\gamma') \cdot (x^{+}\alpha'). \tag{2.2}$$

after applying $x\alpha = x$, cancelling $\gamma\gamma^{-1}$ and reordering the factors. Finally, since $\gamma^{-1}\gamma = \mathrm{id}$, the chain rule informs us that

$$1 = (\gamma^{-1}\gamma)' = (\gamma^{-1})' \cdot (\gamma^{-1}\gamma'),$$

so

$$1 = [x\gamma]^{+}(\gamma^{-1})' \cdot ([x\gamma\gamma^{-1}]^{+}\gamma') = [x\gamma]^{+}(\gamma^{-1})' \cdot (x^{+}\gamma').$$

Applying this to Equation (2.2) yields $[x\gamma]^{+}(\alpha^{\gamma})' = x^{+}\alpha'$.

The calculation for the left-derivative is identical: replace every superscript plus with a superscript minus. □

[40] We're using the fact that our functions are locally increasing here: see Proposition 1.2.6.

Corollary 2.3.15. *Let $\alpha^{\gamma} = \beta$ in a group of right-differentiable (left-differentiable) functions under composition. Suppose that x is fixed by α, β and γ. Then the right (left) derivatives of α and β at x are equal.*

Proof. This is the special case where $x\gamma = x$. □

In F, Matucci applies this observation to elements α and β whose important points I_α and I_β are equal.[41] If these elements are conjugate in F, say $\alpha^{\gamma} = \beta$, their gradients at points $x \in I_\alpha = I_\beta$ must agree. Conversely, mismatching gradients will help us filter out nonconjugate pairs of elements.

[41] Matucci 2008, Section 4.1.1.

Can we use a similar approach to study conjugate elements $\alpha, \beta \in T$? Suppose again that their important point sets $I_\alpha = I_\beta$ are identical. Unlike in F, these sets need not be fixed pointwise by α and β (or even

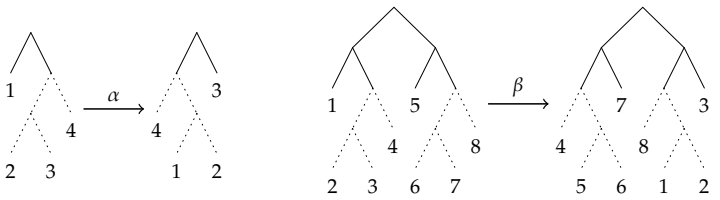

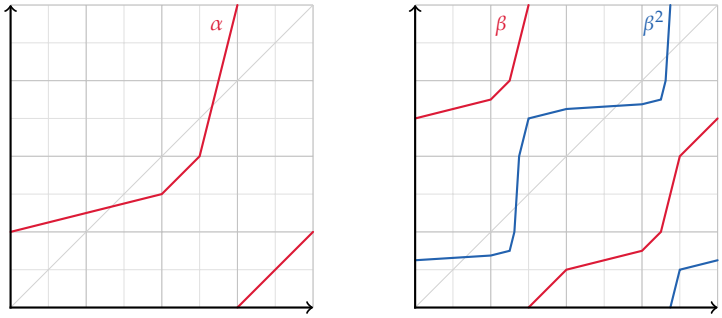

Figure 2.22: Two elements $\alpha, \beta \in T$ with only nondyadic important points $I_\alpha = \{1/3, 2/3\}$ and $I_\beta = \{1/6, 1/3, 2/3, 5/6\}$. The rotation numbers are $\mathrm{rot}(\alpha) = 0$ and $\mathrm{rot}(\beta) = 1/2$. Also shown is the plot of the square β^2.

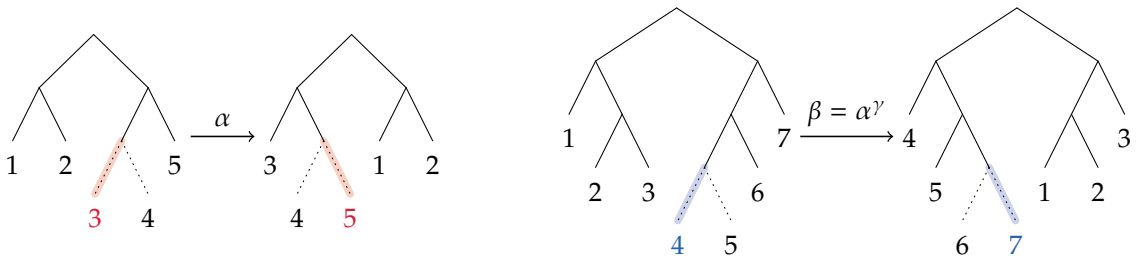

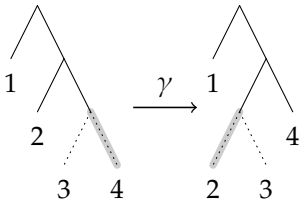

Figure 2.23: The elements $\alpha, \beta \in T$ are conjugate via $\gamma \in F$.

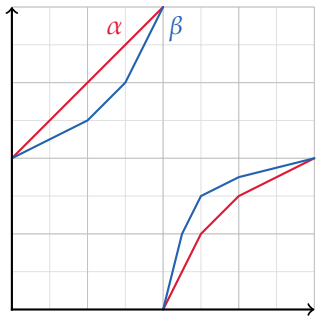

Figure 2.24: Plots of $\alpha, \beta \in T$ from Figure 2.23. Notice the different gradients $(1/2)^{\pm}\alpha'$ and $(1/2)^{\pm}\beta'$.

a conjugator γ). Fortunately, we know that the actions of α and β on this set are identical by Corollary 2.1.12. This weakens our hypotheses to $x\alpha^i = x\beta^i$ for $1 \leq i < q$. This doesn't give us much to work with when repeating the previous calculation; indeed Figures 2.23 and 2.24 show that gradients at important points need not be conserved by conjugation in T.

We need a bigger picture, with global rather than local information. The answer is the 'speed homomorphism' of Bleak et al.[42] In short, we don't get enough information about dynamics by looking at a single gradient at a point—we need to investigate gradients along this point's orbit.

[42] See Bleak, Bowman et al. 2013, Lemma 7.5. Their speed is the base 3 logarithm of our speed, because they are working with the middle thirds Cantor set C.

Definition 2.3.16. Let $O \subseteq S^1$ be a finite subset of the circle. For any element $\alpha \in \mathrm{PL}_+(S^1)$, define the *right-speed of α at O* to be the quantity[43]

$$\alpha S^+_O = \prod_{x \in O} (x^+ \alpha') \,,$$

where the superscript + indicates a right-hand derivative. The *left-speed* S^-_O is the defined by the same formula, but using left-hand derivatives (x^- instead of x^+).

[43] *Speed* isn't quite the right word for this quantity. Since we're taking the product of derivatives at multiple points, we can construct examples where taking O to be larger results in a larger speed. Perhaps we should be using the geometric mean of $\{ x^+\alpha' \mid x \in O \}$ instead, as a kind of average speed rather than total speed?

Claim 2.3.17. *Let O be a finite subset of the circle permuted by a group $G \leq \mathrm{PL}_+(S^1)$. Then S^+_O and S^-_O are homomorphisms $G \to \mathbb{R}^{>0}$ which are constant on conjugacy classes.*

Proof. Suppose that α and γ permute O. Then

$$\begin{aligned}(\alpha\gamma)\, S^+_O &= \prod_{x \in O} x^+ (\alpha\gamma)' \\ &= \prod_{x \in O} (x^+\alpha') \cdot ([x\alpha]^+ \gamma') \\ &= \prod_{x \in O} (x^+\alpha') \prod_{x \in O} ([x\alpha]^+ \gamma') \,.\end{aligned}$$

Now as α permutes O, we can substitute $y = x\alpha$ and iterate over $y \in O\alpha = O$. We obtain

$$(\alpha\gamma)\, S^+_O = \alpha S^+_O \prod_{y \in O} (y^+ \gamma') = \alpha S^+_O \cdot \gamma S^+_O \,.$$

Now the speed $(\alpha^\gamma) S^+_O$ of a conjugate is the product $(\gamma S^+_O)^{-1} \cdot \alpha S^+_O \cdot \gamma S^+_O$ of speeds. Since the codomain group $\mathbb{R}^{>0}$ is abelian, we see that the latter speed is just αS^+_O.

The calculation for S^-_O is identical: replace + with − as before. □

Note that Corollary 2.3.15 is the special case of the previous result when $O = \{ x \}$ is a singleton.

Remark 2.3.18 (Orbit speeds are gradients of powers). Suppose we take O

to be a finite orbit $\{ x\alpha^i \}_{i=0}^{q-1}$. Applying the chain rule twice, we see that

$$\begin{aligned} x^+(\alpha^q)' &= x^+(\alpha^{q-1})' \cdot [x\alpha^{q-1}]^+\alpha' \\ &= x^+(\alpha^{q-2})' \cdot [x\alpha^{q-2}]^+\alpha' \cdot [x\alpha^{q-1}]^+\alpha' . \end{aligned}$$

Continue this until the leftmost factor is $x^+(\alpha^{q-(q-1)})' = x^+\alpha'$. At this point we obtain

$$x^+(\alpha^q)' = \prod_{i=0}^{q-1} [x\alpha^i]^+\alpha' = \alpha S_O^+ ,$$

which shows that the speed over O is exactly the gradient of α^q at x.

If we were to take another representative $y = x\alpha^j \in O$ instead of x, exactly the same calculation would apply. It follows that $y^+(\alpha^q)' = x'^+(\alpha^q)'$. To put this in prose, gradients of qth powers are the same everywhere on a size-q orbit.

To close this subsection we prove a classical result about F (and F-like groups). We include it here to emphasise the fact that this is an argument about gradients, so can be seen as a consequence of the dynamics of F acting on the interval. Alternatively we could prove this using the rotation number, as in Lemma 2.1.14.

Proposition 2.3.19. *Let J be a compact interval. Then $\mathrm{PL}_{S,G}(J)$ is torsion-free.*

Proof. Let $\mathrm{id} \neq \alpha \in \mathrm{PL}_{S,G}(J)$, and let

$$x = \sup\{ t \in J \mid s^+\alpha' = 1 \text{ for all points } s < t \} .$$

Then the interval (possibly singleton) $[0, x]$ is fixed pointwise by α, and $m = x^+\alpha' \neq 1$. We saw above that the map $\langle\alpha\rangle \to G$ sending α^i to $x^+(\alpha^i)'$ is a homomorphism, so $x^+(\alpha^i)' = m^i$. As $m \neq 1$, we have $x^+(\alpha^i)' \neq 1$ for each integer $i \neq 0$. Thus α^i is not the identity function, for each $i \neq 0$. □

2.4 Aligning finite lists of points

Let $X\colon x_1 < \cdots < x_n$ be a sorted list of points in the interval I. We call such a list a *linear partition*, since it expresses I as the union[44] of intervals $[0, x_1] \cup \cdots \cup [x_i, x_{i+1}] \cup \cdots \cup [x_n, 1]$. In Section 1.1.1 we saw special instances of this definition, namely dyadic partitions and dyadic subdivisions. Let $Y\colon y_1 < \cdots < y_n$ be another linear partition with the same number of points. We say that X and Y can be *aligned* by F if there is an element $\alpha \in F$ with $x_i\alpha = y_i$ for every index i.

Lemma 2.1.2 shows that any two dyadic partitions can be aligned by F. Kassabov and Matucci go one step further,[45] explaining that we can determine if two linear partitions of *rational* numbers can be aligned with F.

[44] In the usual definition of a partition $S = \bigcup_{\lambda\in\Lambda} T_\lambda$, the cells T_λ are pairwise disjoint. In contrast, the intersection of any two of our cells is either empty or a singleton $\{ x_i \}$.

[45] Kassabov and Matucci 2012, Corollary 6.14.

Theorem 2.4.1 (Searching for alignments of with F)**.** *Suppose we have two linear partitions $0 < y_1 < \cdots < y_m < 1$ and $0 < z_1 < \cdots < z_m < 1$ where each y_i and z_i is rational. There is an algorithm which constructs an element $\alpha \in F$ such that $y_i = z_i\alpha$ for each i, or else determines that no such α exists.*

Now our task is to extend this result to T.[46] A *circular partition* is a finite list $X\colon x_0, \ldots, x_{n-1}$ of points in the circle with the property that $x_0 \to x_1 \to \cdots \to x_{n-1} \to x_0$. This partitions the circle into *cells* $[x_i, x_{i+1}]$ for $0 \le i < n$, with subscripts modulo n. Let $Y\colon y_0, \ldots, y_{n-1}$ be another circular partition of the same length. We say that X and Y can be *aligned* by T if there is an element $\alpha \in T$ and index $0 \le j < n$ such that $x_i\alpha = y_{i+j}$ for every index i, with subscripts modulo n. Such an element α is said to align the two partitions with *index difference* j. Note that the parameter j allows for up to n ways in which the circular partitions can be aligned, while there's only one way to align linear partitions. An alignment with $j = 0$ is said to align the circular partitions *directly*.

[46] It's probably more accurate to say that we reduce the analogous results for T to those of Kassabov and Matucci in F.

We can generalise the problem solved by Theorem 2.4.1 to the following problem in T.

Problem 2.4.2. Let $X\colon x_0 \ldots, x_{n-1}$ and $Y\colon y_0, \ldots, y_{n-1}$ be two circular partitions in which every element x_i and y_i is rational. Construct an element $\gamma \in T$ such that $x_i\gamma = y_i$ for every i; or otherwise demonstrate that no such γ exists.

As stated, the problem asks us to align the circular partitions directly. This is not a loss of generality: we may replace partition Y with its j-fold left cyclic shift

$$Y^{\lhd j}\colon y_j, \ldots, y_{n-1}, y_0, \ldots, y_{j-1}\,.$$

Then any solution to Problem 2.4.2 can be used to determine if X and $Y^{\lhd j}$ can be aligned directly. Such an alignment is exactly an alignment of X and Y with index difference j. Repeating this for each $0 \le j < n$, we will determine which of the n potential alignments are possible in T.

2.4.1 *Alignment involving a dyadic point*

One strategy to find alignments in T would be to modify the partitions with elements of T, sending x_1 and y_1 (say) to zero. Then we could apply Theorem 2.4.1 to the modified partitions, and use its result to solve the problem for the original problems. Unfortunately, in T we can only map rationals to zero if they are dyadic.[47]

[47] See Section 2.1.1.

Lemma 2.4.3. *We can solve Problem 2.4.2 when one of the x_i or y_i is dyadic.*

Proof. If there is no index i for which x_i and y_i are both dyadic, it is not possible to align X and Y directly. So assume such an i exists, and

relabel the partitions' indices so that x_0 and y_0 are dyadic. Rotate the points of X by $-x_0$ and those of Y by $-y_0$ to obtain new *linear* partitions

Note the similarity to the proof of Lemma 2.1.7.

$$X' = X\rho_{-x_1} : 0,\ x_1\rho_{-x_1},\ \ldots,\ x_{n-1}\rho_{-x_1}$$
$$\text{and} \quad Y' = Y\rho_{-y_1} : 0,\ y_1\rho_{-y_1},\ \ldots,\ y_{n-1}\rho_{-y_1}\,.$$

Note that the maps ρ_{-x_0} and ρ_{-y_0} are elements of T, because x_0 and y_0 are dyadic. Use Theorem 2.4.1 to search for an element $\gamma' \in F$ aligning X' and Y' directly. If such an element exists, it satisfies $x_i\rho_{-x_0}\gamma' = y_i\rho_{-y_0}$ for each i. Then $\gamma = \rho_{-x_0}\gamma'\rho_{+y_0} \in TFT \subseteq T$ sends each x_i to y_i.

To complete the proof, we must show that if a solution γ exists to Problem 2.4.2 aligning X and Y, then there *must* be an element γ' matching X' to Y'. This is almost immediate: if $x_i\gamma = y_i$ then

$$y_i\rho_{-y_0} = x_i\gamma\rho_{-y_0} = x_i\rho_{-x_0}(\rho_{+x_0}\gamma\rho_{-y_0})\,,$$

and so $\gamma' = \rho_{+x_0}\gamma\rho_{-y_0}$ matches X' with Y'. □

We've used a particular structure in this proof which is worth emphasising, as we'll make use of it later. First, a little formalism.

Definition 2.4.4. A *search problem* $\mathcal{P}$ may be modelled as a relation from a set I of *instances* to a set A of *answers*. If $i \in I$ is related to $a \in A$, we say that *i has answer a* or that *a is an answer for i*. For instance, we could take $I = A = \mathbb{N}$ and let $\mathcal{P}$ be the relation

For more details, see R. W. Floyd and Beigel 1994, Chapter 2.

$$a \text{ is an answer for } i \quad \Longleftrightarrow \quad a \text{ is a prime dividing } i\,.$$

Less formally, $\mathcal{P}$ is the problem 'find a prime factor of i'.

An *effective solution to $\mathcal{P}$* is an algorithm which takes an instance $i \in I$ as input and outputs an answer a for i in finite time. If no such a exists, the algorithm must terminate with no output in finite time. We say that $\mathcal{P}$ is *solvable* if there exists an effective solution to $\mathcal{P}$. An answer computed by such a solution is sometimes called a *positive certificate* for the instance i under consideration. Similarly, evidence that there does not exist an answer for i is sometimes called a *negative certificate* for i.

Remark 2.4.5 (Reducing search problems). In the previous proof, we showed that one search problem $\mathcal{P}$ *reduces* to another $\mathcal{Q}$ by arguing as follows.

- An answer a for an instance p of $\mathcal{P}$ can be algorithmically constructed from an answer a' to a particular instance q_p of $\mathcal{Q}$.

Note that the second instance q_p depends on p.

- If an answer a for p exists, then an answer a' for q_p must exist.

We can now build an effective solution for $\mathcal{P}$ given an effective solution for $\mathcal{Q}$. When an instance p is input, use the given solution to find an answer a' for q_p. If there is no such answer, then there is no answer for p by the contrapositive of the second bullet. Otherwise construct a from a'; this is possible and is a correct answer by the first bullet.

In our proof above, p is the instance "find a direct alignment of X and Y by T" of Lemma 2.4.3. On the other hand, q_p is the instance "find an alignment of X' and Y' by F" of Theorem 2.4.1. The answers to these questions are γ and γ' respectively.

Corollary 2.4.6. *Two circular partitions X and Y with the same size consisting only of dyadics can always be aligned by T.*

Proof. Rerun the proof of Lemma 2.4.3 assuming that each x_i and y_i are dyadic. The modified partitions X' and Y' would also consist only of dyadics. This time, we would be able to use Lemma 2.1.2 (instead of Theorem 2.4.1) to align the primed partitions. This lemma always returns a concrete element γ', and thus the proof would always build an aligning element γ. □

2.4.2 *Alignment involving no dyadics*

To complete our description of a solution to Problem 2.4.2, we need to handle the remaining case where no point x_i or y_i is dyadic.

Lemma 2.4.7. *We can solve Problem 2.4.2 when no partition element is dyadic.*

Proof. Let $X\colon x_1 \ldots, x_n$ and $Y\colon y_1, \ldots, y_n$ be our circular partitions. Choose a small arc $[x_0, x_{1.5}]$ with dyadic endpoints containing x_1 but no other x_i. Similarly choose $[y_0, y_{1.5}]$ with dyadic endpoints, containing y_1 but no other y_i. We obtain new partitions

Note that we're indexing our partitions from 1 to n here, instead of 0 to $n-1$.

$$\begin{aligned} X'\colon\ & x_0 \to x_1 \to x_{1.5} \to x_2 \to x_3 \to \cdots \to x_n \to x_0 \\ \text{and}\quad Y'\colon\ & y_0 \to y_1 \to y_{1.5} \to y_2 \to y_3 \to \cdots \to y_n \to y_0\,, \end{aligned}$$

both of which contain a dyadic point. Any function $\gamma' \in T$ directly aligning X' and Y'—sending $x_i \mapsto y_i$ for each $0 \le i \le n$—solves the problem for the original partitions X and Y. Use Lemma 2.4.3 to decide if there exist such a γ', and return it if so.

Now we need to explain why such a γ' must exist if the unprimed partitions can be aligned by some $\gamma \in T$. In this situation, any arc $[x_0, x_{1.5}]$ as above is mapped to some other dyadic arc $[x_0\gamma, x_{1.5}\gamma]$ which contains the point $x_1\gamma = y_1$. Because $x_n \to x_{0.5} \to x_1 \to x_{1.5} \to x_2$ is in cyclic order, so too is the list $y_n \to x_{0.5}\gamma \to y_1 \to x_{1.5}\gamma \to y_2$, because γ preserves cyclic order. Thus we can align X' and Y' if we can find a $\delta \in T$ which aligns $X'\gamma$ and Y'. Explicitly, these partitions are

$$\begin{aligned} X'\colon\ & x_{0.5} \to x_1 \to x_{1.5} \to x_2 \to x_3 \to \cdots \to x_n \to x_{0.5}\,, \\ X'\gamma\colon\ & x_{0.5}\gamma \to y_1 \to x_{1.5}\gamma \to y_2 \to y_3 \to \cdots \to y_n \to x_{0.5}\gamma \\ \text{and}\quad Y'\colon\ & y_0 \to y_1 \to y_{1.5} \to y_2 \to y_3 \to \cdots \to y_n \to y_{0.5}\,. \end{aligned}$$

Now we use a trick of Kassabov–Matucci.[48] Choose more dyadic

[48] Matucci 2008, see Figure 4.3.

points $z_{0.25}$, $z_{0.75}$, $z_{1.25}$ and $z_{1.75}$ such that the lists

$$\begin{aligned} Z\colon\ & z_{0.25} \to x_{0.5}\gamma \to z_{0.75} \to y_1 \to z_{1.25} \to x_{1.5}\gamma \to z_{1.75} \to y_2 \to \\ & \qquad \to y_3 \to \cdots \to y_n \to z_{0.25} \\ Z'\colon\ & z_{0.25} \to y_{0.5} \to z_{0.75} \to y_1 \to z_{1.25} \to y_{1.5} \to z_{1.75} \to y_2 \to \\ & \qquad \to y_3 \to \cdots \to y_n \to z_{0.25} \end{aligned}$$

are in cyclic order. Apply Lemma 2.4.3 to copies of Z and Z' which have had the $\{ y_i \mid i \in \mathbb{Z} \}$ removed. All points remaining in the copies are dyadic, so the lemma will produce a map $\delta_0 \in T$ aligning the copies. Modify δ_0 by forcing it to restrict to the identity function on the arcs $[z_{0.75}, z_{1.25}]$ and $[z_{1.75}, z_{0.25}]$; this is possible because each z_i is dyadic. The result, say $\delta \in T$, aligns the uncopied partitions Z and Z'. Thus the composition $\gamma' = \gamma\delta \in T$ aligns the partitions X' and Y' directly. □

Corollary 2.4.8 (Searching for alignments of rationals with T)**.** *Problem 2.4.2 has an effective solution.*

To close the chapter, let X and Y be the circular partitions described in Problem 2.4.2. It is possible to test for the *existence* of an aligning map $\gamma \in T$ without attempting to construct it.[49]

[49] We thank C. Bleak for highlighting this fact. The proof sketch we give is essentially that given above Kassabov and Matucci 2012, Cor. 6.14.

Lemma 2.4.9 (Deciding when rationals can be aligned with T)**.** *Let $X\colon x_0, \ldots, x_{n-1}$ and $Y\colon y_0, \ldots, y_{n-1}$ be two circular partitions in which every element x_i and y_i is rational. Each point has a binary expansion $x_i = 0.a_i\overline{b_i}$ and $y_i = 0.c_i\overline{d_i}$ with repeating tails b_i and d_i. Let*

$$B_i = \min{}_{\text{lex}}\{ b_i, b_i^{\triangleleft}, b_i^{\triangleleft 2}, \ldots \} \quad \text{and} \quad D_i = \min{}_{\text{lex}}\{ d_i, d_i^{\triangleleft}, d_i^{\triangleleft 2}, \ldots \}$$

be the lexicographically minimal cyclic permutation of these tails.

Then there exists a map $\gamma \in T$ aligning X and Y directly if and only if $B_i = D_i$ for each i.

The fact that γ aligns these partitions *directly* is not a loss of generality. To recover the general case, replace Y with a cyclic shift $Y^{\triangleleft j}$, as below the statement of Problem 2.4.2.

Sketch proof. If $x_i\gamma = y_i$ then it is neccessary that x_i and y_i have the same repeating tail up to cyclic shifts, and hence have the same lexicographically minimal repeating tail. This is established in Lemma 2.1.5 and the discussion following it.

To show that this condition is sufficient, let $\gamma_i \in T$ be such that $x_i\gamma_i = T$. (Again, Lemma 2.1.5 and the discussion following it establish that such a map exists.) Surround each point x_i with a sufficiently small closed arc I_i such that the collection of arcs $\{ I_i \}_{0 \le i < n}$ are pairwise disjoint. We may then apply Lemma 2.1.7 to glue the restriced maps $\gamma_i\big|_{I_i}$ together, forming a map $\gamma \in T$ with $\gamma\big|_{I_i} = \gamma_i\big|_{I_i}$ for each i. In particular, $x_i\gamma = x_i\gamma_i = y_i$ holds for each i. □

3 *Conjugacy in T*

The conjugacy problem (CP) in a group G is the following decision problem: given elements $\alpha, \beta \in G$, determine if there exists an element $\gamma \in G$ such that $\alpha^\gamma = \beta$. The corresponding search problem—which we consider here—requires us to construct a conjugator γ as a positive certificate of conjugacy.

We did not define a decision problem formally in Definition 2.4.4. Briefly, a decision problem is a search problem for which every instance i has exactly one answer 'yes' or 'no'.

Conjugation has been well-studied in Thompson's groups. Higman used the universal algebra viewpoint[1] to solve CP in the groups $V_{n,r}$. (Informally, these groups are versions of V whose gradients are powers of n and whose functions rearrange $[0, r)$, for integers n and r.) As part of a study of power conjugacy in $V_{n,r}$, a small error was corrected in Higman's CP solution by Barker et al.[2] The corrected algorithm was implemented by the author in Python.[3] A more concrete approach to conjugacy in V using revealing pairs to understand dynamics was given in Salazar-Díaz's thesis.[4]

[1] Higman 1974.

[2] Barker, Duncan and Robertson 2016

[3] Robertson 2015.

[4] O. P. Salazar-Díaz 2010.

In F, CP was solved by Guba and Sapir using Diagram groups.[5] Another solution to CP in F was given by Gill and Short, building 'on the geometric invariants introduced by Brin and Squier'.[6]

[5] Victor Guba and Mark Sapir 1997, Theorem 15.23.

[6] Gill and Short 2013; Brin and Squier 2001.

Matucci's thesis gives a unified approach to conjugacy in all three groups F, T and V, in joint work with Belk. The authors use *strand diagrams*, describing their strategy as 'dual' to that of Guba and Sapir. They prove[7] that an element's 'reduced annular strand diagram' (RASD) is a conjugacy invariant, and that two elements of F are conjugate if and only if they have the same RASD. Computing and comparing RASDs can be performed in linear time.[8] The approach generalises to T and V, but requires cubic time.[9] More recently, this technique was adapted by Aroca[10] for larger groups $V_n(H)$ of homeomorphisms.

[7] Belk and Matucci 2014, Theorem 2.10; see also Theorems 2.12 and 2.15.

[8] Hossain et al. 2013.

[9] Matucci 2008, Theorem 2.4.5. See also Proposition 2.1.10 for cutting paths.

[10] Aroca 2018.

The fourth chapter of Matucci's thesis—later published as an article with Kassabov—solves the *k-simultaneous conjugacy problem (SCP)* in F. This is a harder problem, taking as input $2k$ elements $\alpha_1, \ldots, \alpha_k$ and $\beta_1, \ldots, \beta_k$. The goal is to determine if there exists an element γ such that $\alpha_i^\gamma = \beta_i$ for each i simultaneously; note that the special case $k = 1$ is the ordinary conjugacy problem. Again, the corresponding search problem requires us to produce such a conjugator γ as a positive certificate.

Kassabov–Matucci's approach views $F = \mathrm{PL}_2(I)$ as a group of PL functions, and generalises naturally to the larger groups $\mathrm{PL}_{S,G}(J)$ with ∂J dyadic. They outline a four-step plan[11] to tackle SCP in these larger groups.

[11] Kassabov and Matucci 2012, Section 2.2.

Step 1. Find an element of F which aligns $\mathrm{Fix}(\alpha)$ with $\mathrm{Fix}(\beta)$.

Step 2. Solve the conjugacy problem when $\mathrm{Fix}(\alpha) = \mathrm{Fix}(\beta)$.

Step 3. Describe the intersection of element centralisers.

Step 4. Reduce the SCP to a 'restricted' conjugacy problem.

In this chapter we will generalise the first two steps to the group T. The tools required to do so have mostly been covered in Chapter 2. Since we will frequently refer to them, we abbreviate 'Kassabov and Matucci' to 'K&M'.

Our approach is reminiscent of the discussion preceding Lemma 2.1.7: we reduce a problem P in T to a problem Q in F. As long as the reduction takes finite time and calls a solution to Q a finite number of times, we can implement a solution to P.

To abuse terminology, our problem P is 'virtually Q' (in the finite-index sense of virtually cyclic, virtually hyperbolic, etc. groups).

3.1 Constraints on conjugators

In this section we present three lemmas which describe the relationship between a conjugator γ and its restriction $\gamma\big|_U$ to some subset $U \subseteq S^1$. We'll use this to see how partial maps can be extended to bona fide conjugators.

If we have a candidate conjugator γ defined on some small set U, the first lemma explains how γ must be defined on the α-orbit of U.

Lemma 3.1.1. *Let α, β and γ be permutations of a space X. Suppose $\alpha^\gamma = \beta$, and let $U \subseteq X$. Then γ is uniquely determined on the translates $U\alpha^i$ by α, β and $\gamma\big|_U$.*

Proof. Let $t \in U\alpha^i$ be a point in a given translate of U. Then $t\gamma = t\alpha^{-i}\gamma\beta^i$. The statement $t \in U\alpha^i$ is equivalent to $t\alpha^{-i} \in U$, so γ in the previous equation can be replaced by $\gamma\big|_U$. Thus γ must satisfy

$$t\gamma = t\alpha^{-i}\,\gamma\big|_U\,\beta^i \quad \text{for all } t \in U\alpha^i \text{ and all } i \in \mathbb{Z}\,, \tag{3.1}$$

so γ is uniquely determined on translates $U\alpha^i$ by α, β and $\gamma\big|_U$. □

If the translates of U cover the whole space X, then this condition is sufficient for γ to be a conjugator.

Lemma 3.1.2. *Let α, β and γ be permutations of a space X, and let $U \subseteq X$. Suppose $X = \bigcup_{i\in\mathbb{Z}} U\alpha^i$. If γ satisfies Condition (3.1), then $\alpha^\gamma = \beta$.*

A note of caution: if α and β belong to some group of interest $G \leq \mathrm{Sym}(X)$, there is no reason that a map γ satisfying Condition (3.1) also belongs to G. We have to decide if $\gamma \in G$ separately.

Proof. By hypothesis, each $t \in X$ belongs to a translate $U\alpha^i$. Then $t\alpha \in U\alpha^{i+1}$, so

$$(t\alpha)\gamma = t\alpha\alpha^{-(i+1)}\,\gamma\big|_U\,\beta^{i+1}\,;$$

whereas

$$(t\gamma)\beta = (t\alpha^{-i}\,\gamma\big|_U\,\beta^i)\beta\,.$$

As the two right-hand expressions are equal and this argument holds for all $t \in X$, we see that $\gamma\beta = \alpha\gamma$. □

This gives us a crude means to search for conjugators. First we find a set U with the property that $X = \bigcup_{i\in\mathbb{Z}} U\alpha^i$. Then we choose a

bijection $\delta_0 \colon U \to V$, for some set $V \subseteq X$. We define $\delta \colon X \to X$ by

$$t\delta = t\alpha^{-i}\delta_0\beta^i \qquad \text{for all } t \in U\alpha^i \text{ and all } i \in \mathbb{Z}\,. \tag{3.2}$$

Then we check to see if δ conjugates α to β. There are two details to verify. Firstly, because the translates $\{U\alpha^i\}$ may overlap, we need to check that δ is well-defined. We also need to verify that δ is a permutation, i.e. that $\delta^{-1} \colon X \to X$ exists. The next lemma does so, under some additional hypotheses.

Lemma 3.1.3. *Let α and β be permutations of a space X. Also let $U, V \subseteq X$ be such that $X = \bigcup_{i\in\mathbb{Z}} U\alpha^i = \bigcup_{i\in\mathbb{Z}} V\beta^i$. Suppose there exists an integer $q > 0$ with the following properties.*

1. *$U = U\alpha^q$ and $V = V\beta^q$.*
2. *$U \cap U\alpha^\ell = \emptyset$ and $V \cap V\beta^\ell = \emptyset$, for each $0 < \ell < q$.*

If $\delta_0 \colon U \to V$ is a bijection with $\delta_0^{-1}\,\alpha^q\big|_U\,\delta_0 = \beta^q\big|_V$, then the map δ defined by Equation (3.2) is a well-defined permutation, with $\alpha^\delta = \beta$.

Proof. The only means by which δ may be ill-defined occurs when $U\alpha^i \cap U\alpha^j$ is not empty, for distinct integers $i < j$. Assume that this is the case. In order for δ to be well-defined, we require that

$$t\alpha^{-i}\delta_0\beta^i = t\alpha^{-j}\delta_0\beta^j \qquad \text{for all } t \in U\alpha^i \cap U\alpha^j\,,$$

for each such pair $i < j$ as above. By multiplying on the right by β^{-i}, we see that this condition is equivalent to

$$t\alpha^{-i}\delta_0 = t\alpha^{-j}\delta_0\beta^{j-i} \qquad \text{for all } t \in U\alpha^i \cap U\alpha^j\,.$$

Let $t' = t\alpha^{-i}$. Noting that $t' \in (U\alpha^i \cap U\alpha^j)\alpha^{-i} = U \cap U\alpha^{j-i}$, this is equivalent[12] to

$$t'\delta_0 = t'\alpha^{-(j-i)}\delta_0\beta^{j-i} \qquad \text{for all } t' \in U \cap U\alpha^{j-i}\,.$$

[12] On the right-hand side, replace t with $t\alpha^{-i}\alpha^i = t'\alpha^i$.

There exist unique integers $n \geq 0$ and $0 \leq r < q$ such that $j - i = nq + r$. Since $U = U\alpha^q$, it follows that $U\alpha^{nq+r} = U\alpha^r$. If $0 < r < q$, property 2 tells us that $U \cap U\alpha^{nq+r} = U \cap U\alpha^r$ is empty. Thus the only way for $U\alpha^i \cap U\alpha^j$ to be nonempty is if $r = 0$. Now we know that $j - i = nq$, the condition we have been tracking is now

$$t'\delta_0 = t'\alpha^{-(nq)}\delta_0\beta^{nq} \qquad \text{for all } t' \in U\,. \tag{3.3}$$

Now we need to demonstrate that this statement is a consequence of our hypotheses.

Thanks to this lemma's requirements on δ_0, we know that

$$t\delta_0 = t\alpha^{-q}\delta_0\beta^q \qquad \text{for all } t \in U\,. \tag{3.4}$$

Now if $t' \in U$, then $t'\alpha^{-q} \in U$ also. Applying Equation (3.4) tells us

that $t'\alpha^{-q}\delta_0 = (t'\alpha^{-q})\alpha^{-q}\delta_0\beta^q$. Multiplying by β^q on the right yields

$$t'\alpha^{-q}\delta_0\beta^q = t'\alpha^{-2q}\delta_0\beta^{2q} \qquad \text{for all } t' \in U\,.$$

However the left-hand side of this equation is just $t'\delta_0$, again by Equation (3.4). We conclude that

$$t'\delta_0 = t'\alpha^{-2q}\delta_0\beta^{2q} \qquad \text{for all } t' \in U\,.$$

Use the same trick to see that $t'\alpha^{-2q}\delta_0 = t'\alpha^{-3q}\delta_0\beta^q$ for all $t' \in U$; then conclude that $t'\delta_0 = t'\alpha^{-2q}\delta_0\beta^{2q} = t'\alpha^{-3q}\delta_0\beta^{3q}$. Repeating this process inductively, we deduce that

$$t'\delta_0 = t'\alpha^{-mq}\delta_0\beta^{mq} \qquad \text{for all } t' \in U \text{ and all } m \geq 0,$$

and in particular this holds for $m = n$. Thus Condition (3.3) holds, meaning that δ is well-defined.

Is δ an invertible function? We know that δ restricts to q functions

$$\delta\big|_{U\alpha^i} = (\alpha^{-i}\delta_0\beta^i)\big|_{U\alpha^i} : U\alpha^i \to V\beta^i\,,$$

where $0 \leq i < q$. These restrictions are a composition of invertible functions α^{-1}, δ_0 and β, so each restriction $\delta\big|_{U\alpha^i}$ is a bijection. Now we argued above that $\{U\alpha^i\}_{0\leq i<q}$ is a partition of X; similarly, $\{V\beta^i\}_{0\leq i<q}$ is another partition of X. Thus δ is a bijection. We can conclude that $\alpha^\delta = \beta$ by applying Lemma 3.1.2 with δ in place of γ. □

3.2 The miracle of F

The aim of this short section is to justify a miraculous lemma about conjugacy in F and F-like groups. Briefly: conjugators between powers are also conjugators between roots.

Proposition 3.2.1 (The miracle of F)**.** *Let $\alpha, \beta, \gamma \in \mathrm{PL}_{S,G}(J)$, for some compact interval J. Let $n \neq 0$ be any integer. Then $\alpha^\gamma = \beta$ if and only if $(\alpha^n)^\gamma = \beta^n$.*

This miracle has two particularly nice consequences.

Corollary 3.2.2 (Roots are unique in F)**.** *et $\alpha, \beta \in \mathrm{PL}_{S,G}(J)$ and let $n \neq 0$ be an integer.*

1. *$\alpha^n = \beta^n$ if and only if $\alpha = \beta$.*
2. *$C_{\mathrm{PL}_{S,G}(J)}(\alpha) = C_{\mathrm{PL}_{S,G}(J)}(\alpha^n)$, where $C_H(\cdot)$ denotes a centraliser in H.*

Proof. These are special cases of Proposition 3.2.1. For part 1, choose $\gamma = \mathrm{id}$. For part 2, choose $\alpha = \beta$. □

How should we prove our miraculous proposition? The hard work has already been done for us by Kassabov and Matucci (K&M). We will use the following result of theirs.

Lemma 3.2.3. *Let $J = [j_0, j_1]$ be a compact interval and let $\alpha, \beta \in \mathrm{PL}^0_{S,G}(J)$ be almost one-bump functions. Fix an element $g \in G$. There is at most one element $\gamma \in \mathrm{PL}_{S,G}(J)$ such that $\alpha^\gamma = \beta$ and $j_0^+\gamma' = g$. Similarly, there is at most one element $\gamma \in \mathrm{PL}_{S,G}(J)$ such that $\alpha^\gamma = \beta$ and $j_1^-\gamma' = g$.*

In other words, this lemma says a conjugator between almost one-bump functions is entirely determined by its initial (or final) gradient.

Proof. K&M prove this result in the case[13] where $\partial J \subseteq S$. For the case where $\partial J \nsubseteq S$, we can use the larger breakpoint group $S' = \mathbb{R}$ and apply their proposition in the context of the larger group $\mathrm{PL}_{\mathbb{R},G}(J)$. This is legitimate because $\partial J \subseteq \mathbb{R}$. □

[13] Kassabov and Matucci 2012, Proposition 4.20, Remark 4.22.
While their computational requirements (see their Section 3) are not satisfied when using $\mathbb{R}$ as the breakpoint group, the uniqueness of conjugators with a given initial gradient still holds.

Proof of Proposition 3.2.1. The $\Longrightarrow$ direction is immediate; we need to establish the $\Longleftarrow$ direction. Assume then that $(\alpha^n)^\gamma = \beta^n$. We'll split our argument into various cases based on the 'bumpiness' of α^n.

No bumps. If $\alpha^n = \mathrm{id}$ then $\beta^n = (\alpha^n)^\gamma = \mathrm{id}^\gamma = \mathrm{id}$. Since F-like groups are torsion-free,[14] it follows that $\alpha = \beta = \mathrm{id}$. So it is certainly the case that $\alpha^\gamma = \beta$ holds.

[14] Proposition 2.3.19

Almost one-bump. If $\alpha^n \in \mathrm{PL}^0_{S,G}(J)$ is an almost one-bump function, then by definition $\mathrm{Fix}(\alpha^n) \cap S = \partial J$. Now

$$\begin{aligned}\mathrm{Fix}(\beta^n) \cap S &= \mathrm{Fix}([\alpha^n]^\gamma) \cap S = \mathrm{Fix}(\alpha^n)\gamma \cap S \\ &= [\mathrm{Fix}(\alpha^n) \cap S]\,\gamma = \partial J\gamma = \partial J\,,\end{aligned}$$

Argue as in Section 2.1.1 to see that $\mathrm{PL}_{S,G}(J)$ acts on $S \cap J$. This means that $S\gamma = S$, which justifies the third equality.

meaning that β^n is also an almost one-bump function.

Seeking a contradiction, suppose that α is *not* an almost one-bump function. Then there is some point $s \in S \cap J^\circ$ fixed by α. But that point s would also be fixed by α^n, contradicting our earlier assumption that α^n is almost one-bump. We conclude that $\alpha \in \mathrm{PL}^0_{S,G}(J)$ must be almost one-bump. Apply the same reasoning with β in place of α to see that $\beta \in \mathrm{PL}^0_{S,G}(J)$ is almost one-bump too.

Observe that the elements β and α^γ both conjugate β^n to itself: $(\beta^n)^\beta = \beta^n$ and

This paragraph's argument is that given in Kassabov and Matucci 2012, Lemma 4.13.

$$\begin{aligned}(\beta^n)^{(\alpha^\gamma)} &= (\alpha^\gamma)^{-1}\beta^n\alpha^\gamma = \gamma^{-1}\alpha^{-1}(\gamma\beta^n\gamma^{-1})\alpha\gamma \\ &= (\alpha^{-1}\alpha^n\alpha)^\gamma = (\alpha^n)^\gamma = \beta^n\,.\end{aligned}$$

Write $J = [j_0, j_1]$. As the elements α^n and β^n are conjugate, their initial gradients

$$j_0^+(\alpha^n)' = (j_0^+\alpha')^n \quad \text{and} \quad j_0^+(\beta^n)' = (j_0^+\beta')^n$$

must be equal, by Corollary 2.3.15. Take nth roots in the gradient group $G = 2^{\mathbb{Z}}$ to see that the initial gradients of α and β are also equal. Use Corollary 2.3.15 once more to learn that initial gradients of α and α^γ are equal. This establishes that α^γ and β share an initial

gradient and both conjugate β^n to itself. By Lemma 3.2.3 we must have $\alpha = \beta$.

Many bumps. Let $A = \partial J \cup (\partial \operatorname{Fix}(\alpha^n) \cap S)$ and $B = \partial J \cup \big(\partial \operatorname{Fix}(\beta^n) \cap S\big)$. Enumerate these sets as $A = \{a_0 < \cdots < a_m\}$ and $B = \{b_0 < \cdots < b_m\}$. Then $m \geq 1$, $J = [a_0, a_m] = [b_0, b_m]$, and $a_i\gamma = b_i$ for each i. Define $A_i = [a_i, a_{i+1}]$ and $B_i = [b_i, b_{i+1}]$ so that $A_i\alpha = A_i$, $B_i\beta = B_i$ and $A_i\gamma = B_i$. Then the restrictions $\alpha\big|_{A_i}$, $\alpha^n\big|_{A_i}$, $\beta\big|_{B_i}$ and $\beta^n\big|_{B_i}$ are all almost one-bump functions.

These sets are the important points of α^n and β^n for F.

We are in the situation where $(\alpha^n)^\gamma = \beta^n$. Restricting this equality to B_i, we see that $(\alpha^n\big|_{A_i})^{\gamma|_{A_i}} = \beta^n\big|_{B_i}$. As all three restrictions are almost one-bump functions, we can use our argument from the previous case to conclude that $(\alpha\big|_{A_i})^{\gamma|_{A_i}} = \beta\big|_{B_i}$ for each i. Thus $\alpha^\gamma = \beta$. □

We could try to extend this 'miracle' from F to T. Unfortunately, there is no miracle[15] of T, so we can only use the ideas and techniques from F up to a point.

[15] For example, let $\alpha, \beta \in T$ be the elements described in Examples 5.1.2 and 5.1.3. These are constructed such that $\alpha^2 = \beta^2$, so the powers are conjugate via $\gamma = \mathrm{id}$. But $\alpha \neq \beta$, so $\alpha^\gamma \neq \beta$ in this example. Thus there is no miracle of T.

3.3 Finding coarse conjugators

The rest of this chapter describes an algorithm to solve the conjugacy problem in T. Given $\alpha, \beta \in T$, we are asked to construct (if possible) a conjugator $\gamma \in T$ with $\alpha^\gamma = \beta$. First we should establish any conditions on α and β which are necessary for γ to exist. We know from Proposition 2.1.9.3 that the rotation numbers of conjugate elements are equal, so our first task is to compute and compare $\operatorname{rot}(\alpha)$ and $\operatorname{rot}(\beta)$ using Theorem 2.1.13. If these are not equal, we immediately conclude that the elements are not conjugate. Assume then that $\operatorname{rot}(\alpha) = \operatorname{rot}(\beta)$ is a rational number p/q in lowest terms.

As we've noted before, if $x \in S^1$ is fixed by α, then $x\gamma$ is fixed by α^γ. The converse is true too, so if $\beta = \alpha^\gamma$ then $\operatorname{Fix}(\beta) = \operatorname{Fix}(\alpha^\gamma) = \operatorname{Fix}(\alpha)\gamma$. Replacing α with α^i, we see that any conjugator γ must map $\operatorname{Fix}(\alpha^i)$ to $\operatorname{Fix}(\beta^i)$, for each $i \in \mathbb{N}$. We know by Proposition 2.1.9.5 that all finite orbits have period q. Thus any conjugator γ must satisfy $\operatorname{Fix}(\alpha^q)\gamma = \operatorname{Fix}(\beta^q)$. We give a name to maps satisfying this condition.

In F, the rotation number is always $0 = 0/1$, so Kassabov and Matucci (K&M) know their conjugators γ have to map $\operatorname{Fix}(\alpha)$ to $\operatorname{Fix}(\beta)$.

Definition 3.3.1. Let $\alpha, \beta \in T$ have the same rotation number p/q in lowest terms. An element $\delta \in T$ for which $\operatorname{Fix}(\alpha^q)\delta = \operatorname{Fix}(\beta^q)$ is called a *coarse conjugator* for α and β.

Now we explain how to search for coarse conjugators; in Section 3.3.3 we'll see how they can be used to solve the conjugacy problem.

If δ is a coarse conjugator then $\operatorname{Fix}(\alpha^q)\delta = \operatorname{Fix}(\beta^q)$. Because T is a group of homeomorphisms, it follows that $\partial \operatorname{Fix}(\alpha^q)\delta = \partial \operatorname{Fix}(\beta^q)$. In other words, the important points I_α of α are mapped by δ to the important points I_β of β. This gives us another necessary condition to establish: it must be the case that $|I_\alpha| = |I_\beta|$, or else they cannot

be aligned by any map $\delta \in T$. If this is not the case, we immediately conclude that the elements are not conjugate.

How we proceed depends on the number of important points $n = |I_\alpha| = |I_\beta|$. We handle the two cases $n = 0$ and $n > 0$ separately.

3.3.1 *No important points*

In this section we handle the case where there are no important points to align, i.e. when $\partial \operatorname{Fix}(\alpha^q) = \partial \operatorname{Fix}(\beta^q) = \emptyset$. A set has empty boundary if and only if it's closed and open. The only closed and open subsets of S^1 are the empty set and S^1 itself (because S^1 is a connected space). We can't have $\operatorname{Fix}(\alpha^q) = \emptyset$: that would mean that α has no periodic points, but this is false (see Remark 2.2.9). Thus $\operatorname{Fix}(\alpha^q) = S^1$ and $\operatorname{Fix}(\beta^q) = S^1$ are equal. Since q is the reduced rotation number denominator, we see that α and β are torsion elements of T with order q.

Notice that when we have $n = 0$ important points, the fixed point sets are automatically equal—there is no need to search for a coarse conjugator.

It turns out that torsion elements with the same rotation number are always conjugate. Put differently, the rotation number is a complete conjugacy invariant for torsion elements. To the best of the author's knowledge, this result is due to Burillo et al,[16] but see also the work of Matucci[17] and Geoghegan and Varisco.[18]

[16] Burillo et al. 2009, Corollary 6.2.

[17] Matucci 2008, Proposition 2.2.13, Lemma 7.1.2.

[18] Geoghegan and Varisco 2017, Corollary 2.6.

Lemma 3.3.2. *Torsion elements $\alpha, \beta \in T$ have the same rotation number if and only if they are conjugate in T. In this situation, we can construct a conjugator $\delta \in F$ so that $\alpha^\delta = \beta$.*

Proof. We already know—from Proposition 2.1.9.3—that the rotation number is a conjugacy invariant. We prove the other implication by constructing a conjugator $\delta \in F$.

Let the rotation number be $\operatorname{rot}(\alpha) = \operatorname{rot}(\beta) = p/q$ in lowest terms. Use A and B to denote the orbits of 0 under α and β, respectively. Arrange them into circular order, say as

$$A\colon 0 = a_0 \to a_1 \to \cdots \to a_{q-1} \to 0$$
$$\text{and} \quad B\colon 0 = b_0 \to b_1 \to \cdots \to b_{q-1} \to 0\,,$$

and let the arcs $A_i = [a_i, a_{i+1})$ and $B_i = [b_i, b_{i+1})$ be the ith cells of these partitions, closed at the left and open at the right. We know from Lemma 2.1.11 that $a_i\alpha = a_{i+p}$ and $a_i = a_0\alpha^{ik}$ where $k = p^{-1} \bmod q$. It follows that $A_i = A_0\alpha^{ik}$ and $A_0\alpha^i = A_{ip \bmod q}$. The same is true replacing the symbols a_i, A_i and α with b_i, B_i and β, respectively.

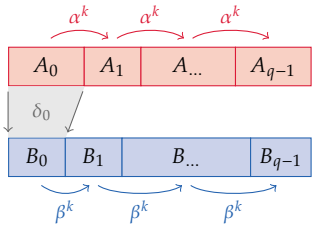

Figure 3.1: A schematic depicting the cells A_i and B_i, as well as the maps α^k, β^k and δ_0.

Use Lemma 2.1.1 to produce an element $\delta_0 \in F$ for which $a_1\delta_0 = b_1$. Then $A_0\delta_0 = B_0$, as δ_0 fixes $a_0 = b_0 = 0$. For each i, the composition $\phi_i = \alpha^{-ik}\delta_0\beta^{ik}$ is a PL_2 map $A_i \to B_i$. Use Lemma 2.1.7 to glue[19] these functions together. The result $\delta \in T$ is defined by

$$t\delta = t\alpha^{-ik}\delta_0\beta^{ik} \qquad \text{if } t \in A_i = A_0\alpha^{ik}, \text{ for each } 0 \le i < q,$$

[19] Strictly speaking, our definition of PL_2 maps and the gluing lemma require that the domain and codomain of the ϕ_i are closed. This isn't a problem: extend ϕ_i to a map $\overline{A_i} \to \overline{B_i}$ between closed arcs by defining $a_{i+1}\phi_i = b_{i+1}$. (The overbar denotes the closure operation.)

for each $0 \le i < q$; this is a particular instance of Equation (3.2).

Now apply Lemma 3.1.3, where the arcs A_0 and B_0 take on the roles of U and V, respectively. The lemma's requirements are satisfied, because

- $\{A_i\}_{i=0}^{q}$ and $\{B_i\}_{i=0}^{q}$ are both partitions of the circle;
- $A_0\alpha^i \cap A_0\alpha^j = A_{ip \bmod q} \cap A_{jp \bmod q} \neq \emptyset$ if and only if $i \equiv j \pmod q$, and similarly for B_i and B_j; and finally because
- $\alpha^q = \beta^q = \mathrm{id}$.

Thus we conclude that $\delta \in F$ is a well-defined conjugator between α and β. □

3.3.2 *At least one important point*

Now we handle the remaining case, in which there are $n > 0$ important points which must be aligned by a coarse conjugator. In this situation, α and β have infinite order, so we have to deal with points in infinite orbits converging to and from attracting and repelling points. (Contrast this to previous section, where all points belonged to size q orbits under α and β.)

Enumerate the important points in cyclic order as $I_\alpha\colon x_1, \ldots, x_n$ and $I_\beta\colon y_1, \ldots, y_n$. A coarse conjugator must align these cyclic partitions, with some index difference $0 \le j < n$. We determine if a coarse conjugator δ exists by looping over each value $0 \le j < n$, and looking specifically for a coarse conjugator δ_j with index difference j.

C.f. Section 2.4.

There is a necessary condition to worry about. Let $X_i = [x_i, x_{i+1}]$ and $Y_i = [y_i, y_{i+1}]$ be the cells corresponding to the important point partitions. Then $X_i\delta_j = Y_{i+j}$. For δ_j to exist, it is necessary that

$$\text{whenever } X_i \text{ is fixed pointwise by } \alpha^q, \text{ its image } X_i\delta_j = Y_{i+j} \text{ must be fixed pointwise by } \beta^q; \tag{3.5}$$

this is because $\mathrm{Fix}(\alpha^q)$ is sent to $\mathrm{Fix}(\beta^q)$ by any conjugator. We can determine if a cell is fixed pointwise using Lemma 2.3.12, so for each j we can check if Condition (3.5) holds true. If not, then there cannot be a coarse conjugator δ_j with index difference j. For any remaining values of j, we defer to Corollary 2.4.8 and construct—if possible—a coarse conjugator δ_j.

3.3.3 *Reducing to a search for fine conjugators*

Now that we've explained how to find coarse conjugators, how do we use them to solve the conjugacy problem?

Definition 3.3.3. Two elements $\alpha, \beta \in T$ are called *coarsely equivalent* if they have the same reduced rotation number p/q and the same periodic point sets $\mathrm{Fix}(\alpha^q) = \mathrm{Fix}(\beta^q)$. Such elements have the same important points $I = I_\alpha = I_\beta$.

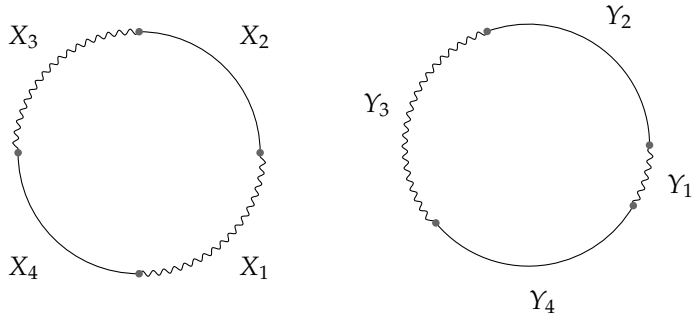

Figure 3.2: In this schematic, the even-indexed cells X_i are fixed pointwise by α^q, but odd-indexed cells are not fixed pointwise by α^q. Similarly, even-indexed (odd-indexed) cells Y_i are (are not) fixed pointwise by β^q. Condition (3.5) tells us that we should only look for coarse conjugators δ_0 and δ_2 with index differences 0 and 2, respectively, because we need to map cells fixed pointwise to cells fixed pointwise.

Suppose as above that $\alpha, \beta \in T$ have the same rotation number p/q in lowest terms. If there is a coarse conjugator δ between these two elements, then $\mathrm{Fix}([\alpha^\delta]^q) = \mathrm{Fix}([\alpha^q]^\delta) = \mathrm{Fix}(\alpha^q)\delta = \mathrm{Fix}(\beta^q)$. We also know that $\mathrm{rot}(\alpha) = \mathrm{rot}(\alpha^\delta)$, so the elements α^δ and β are coarsely equivalent.

Next, suppose that a 'full' conjugator γ with $\alpha^\gamma = \beta$ exists. Then certainly $(\alpha^\delta)^{\delta^{-1}\gamma} = \beta$. Now we need a procedure to find an element ϵ of the form $\delta^{-1}\gamma$; what are these properties?

Definition 3.3.4. Let $\alpha, \beta \in T$ be coarsely equivalent elements with important points $I = I_\alpha = I_\beta$. A map $\epsilon \in T$ with $\alpha^\epsilon = \beta$ which fixes I pointwise is called a *fine conjugator* between α and β.

The name 'fine' is unfortunate, because 'a fine' sounds and reads like 'affine'.

The key phrase in this definition is *fixes I pointwise*. The idea is that we want coarse conjugators alone to be responsible for aligning important points; once they're aligned, fine conjugators should leave them well alone.

Given coarsely equivalent elements α and β, the *aligned conjugacy problem* (ACP) is the search problem which asks us to find a fine conjugator ϵ between α and β, or determine that no such ϵ exists.

If we find a fine conjugator ϵ between α^δ and β, then we can use $\gamma = \delta\epsilon$ as a full conjugator between α and β. This suggests that we can reduce CP to ACP. We explain why this is the case with the following formal algorithm, which corresponds to step 1 of K&M's four-step outline.

Algorithm 3.3.5 (Reduction to fine conjugator search)**.** The following procedure constructs a coarse conjugator for elements α and β in T, and reduces the conjugacy problem to ACP.

1. Check that the elements have the same reduced rotation number p/q.
2. Check that they have the same number of important points $n = |I_\alpha| = |I_\beta|$.
3. If $n = 0$, use Lemma 3.3.2 to construct a conjugator γ and return it.
4. For each $0 \le j < n/q$:
   (a) For each $1 \le i \le n$, check that X_i is fixed pointwise by α^q if and only if Y_{i+j} is fixed pointwise by β^q. If not, continue to the next value of j.
   (b) Use Corollary 2.4.8 to construct a element δ_j mapping I_α to I_β with index difference j. This is our coarse conjugator. If this is not possible, continue to the next value of j.

The loop bounds are tighter than the discussion above would suggest; see the proof for a justification.

(c) Use a solution to ACP to construct an element ϵ_j which fixes I_β pointwise and satisfies $(\alpha^{\delta_j})^{\epsilon_j} = \beta$. This is our fine conjugator. If such an ϵ_j exists, return the product $\gamma_j = \delta_j\epsilon_j$ as a positive certificate of conjugation. Otherwise continue to the next value of j.

5. If any of the checks in steps 1 and 2 fail, or if we never return from the loop in step 4c, there is no conjugator $\alpha^\gamma = \beta$.

Proof of correctness. If the algorithm succeeds, then we produce a bona fide conjugator. We must explain why the algorithm will succeed if such an element γ exists, with index difference j_0 say. All of the necessary conditions ('Check that...' above) will be satisfied. The case where α and β have (the same) finite order is handled by step 3. Otherwise we reach step 4.

We know from Lemma 2.1.11 that α acts on its important points I_α by pushing them pn/q steps forward. If $p \neq 0$ then let $k = p^{-1} \bmod q$, so that α^k pushes I_α forward by $t = n/q$ steps. Write $j_0 = \ell t + j$, where ℓ and j are integers and $0 \leq j < t$. Then $\gamma' = \alpha^{-k\ell}\gamma$ is also a conjugator between α and β. It aligns the important point sets I_α and I_β with index difference $j_0 - \ell t = j$ in the range $0 \leq j < t = n/q$.

If $p = 0$ then $q = 1$, since we assume the rotation number p/q is in lowest terms. Let $\gamma' = \gamma$ and $j = j_0$ in this case. It is still true that γ' is a conjugator between α and β with index difference $0 \leq j < n/q = n$. This justifies the upper limit on the loop in step 4.

Within the body of the loop, we will pass step 4a because Condition (3.5) is another necessary condition which is satisfied by any conjugator. In step 4b, a coarse conjugator δ_j certainly exists: $\delta_j = \gamma'$ is an example. Whichever δ_j is chosen, in step 4c a fine conjugator ϵ_j completing the conjugation must exist: $\epsilon_j = \delta_j^{-1}\gamma'$ is an example. This fixes $I(\alpha^{\delta_j}) = I_\beta$ pointwise because $b\delta_j = b\gamma'$ for each $b \in I_\beta$, since the two maps δ_j and γ' have the same index difference j. Thus the algorithm produces a conjugator. □

Remark 3.3.6 (Restricting the index difference). Suppose we restrict the loop in step 4 to run for only a single value j^*, possibly outside the range $0 \leq j < n/q$. Then the resulting algorithm will construct a full conjugator γ with index difference j^*, or determine that no such γ exists.

3.4 Finding fine conjugators

To complete the story, we need to explain how to solve the aligned conjugacy problem. Assume then that $\alpha, \beta \in T$ have infinite order and are coarsely equivalent. As usual, let p/q be their reduced rotation numbers and let $I \neq \emptyset$ be their important point sets. We begin with a general discussion of our strategy for finding fine conjugators $\epsilon \in T$, i.e. for finding ϵ such that $\alpha^\epsilon = \beta$ and $\epsilon\big|_I = \mathrm{id}$.

One potential avenue of attack could be to cut the circle at a important point r into an interval $[r, r+1] \subseteq \mathbb{R}$. Since α^q, β^q and a fine conjugator ϵ all fix r, we can see these elements as members of $\mathrm{PL}_2([r, r+1])$. This is an F-like group rather than a T-like group, so the results of K&M apply. We might hope that an understanding of conjugacy in this group will help us find fine conjugators. There are three potential problems to overcome with this approach.

First of all, we need to distinguish between dyadics and nondyadics. In the theorem below,[20] K&M explain how to solve the conjugacy problem in $\mathrm{PL}_2(J)$ assuming that J has has dyadic endpoints ∂J. With that said, their machinery—in particular the stair algorithm—works perfectly well in any group $\mathrm{PL}_+(J)$, with no restrictions on ∂J.

[20] Kassabov and Matucci 2012, Theorem 7.5; see also Remark 2.2.

Theorem (Kassabov–Matucci)**.** *Let $\mathrm{PL}_{S,G}(I)$ be the generalised Thompson group defined in Definition 1.2.1. Let $J \subseteq I$ be a closed subinterval with endpoints $\partial J \subseteq S$. Then the group $\mathrm{PL}_{S,G}(J)$ has solvable conjugacy problem. Moreover we can construct and enumerate all possible conjugators.*

If r is not dyadic we have a second problem to be aware of. If $\omega \in T$ is one of our elements α^q, β^q or ϵ under consideration, then we must have $r^+\omega' = r^-\omega'$. If ω corresponds to a map $\phi \in \mathrm{PL}_2([r, r+1])$ then this condition reappears as $r^+\phi' = [r+1]^-\phi'$. These gradients need not be 1, which means that we must work in $\mathrm{PL}_2^{\mathrm{rest}}([r, r+1])$ rather than in $\mathrm{PL}_2^{\mathrm{flat}}([r, r+1])$[21]; when we say PL_2 we mean $\mathrm{PL}_2^{\mathrm{rest}}$. The set of such elements ϕ forms a proper subgroup of $\mathrm{PL}_2([r, r+1])$. We need to make sure we only produce conjugators in this subgroup, so that we can recover a fine conjugator ϵ on the circle.

[21] For the definition of and distinction between 'rest' and 'flat', see Definition 1.2.2.

A third issue with this proposed method is that it searches for a fine conjugator between $\alpha, \beta \in T$ by instead searching for fine conjugators between the powers α^q and β^q. Any conjugator between roots is a conjugator between powers, but is the reverse true? In a general group the answer is 'no'. Encouragingly, the answer is 'yes' for F-like groups, as we saw in Section 3.2 above.

Hopefully the reader agrees that the search for fine conjugators does not immediately reduce to an F-like conjugacy problem. At the very least there are details to check! In the rest of this chapter, we will use some of the ideas above to describe how to find fine conjugators. The main departure from the discussion above is that we cut the circle into q intervals—not just a single interval $[r, r+1]$.

3.4.1 *Restriction to a representative arc*

Thanks to the lemmas at the start of this chapter, we only need to consider a sufficiently representative slice of the circle—a kind of fundamental domain. Let us define this slice more precisely.

Definition 3.4.1 (Representative arc)**.** Let $\alpha, \beta \in T$ be coarsely equivalent, with reduced rotation numbers p/q and equal important point sets I. Arrange I in circular order as $r_0 \to \cdots \to r_{n-1} \to r_0$. If

possible, choose r_0 to be dyadic; otherwise all r_i are nondyadic. For $0 \le j < q$, define $Q_j = [r_{jt}, r_{(j+1)t})$, where $t = n/q$. These sets are called *representative arcs* for α and β.

Note that subscripts on points r_i and arcs Q_j are reduced modulo n and q, respectively.

Because we know that α and β permute I, we have $r_i\alpha = r_i\beta = r_{i+pt}$. This means that the arcs themselves are permuted according to the rule $Q_j\alpha = Q_j\beta = Q_{j+p}$. In particular, if $k = p^{-1} \bmod q$ we have $r_i\alpha^k = r_i\beta^k = r_{i+t}$, and hence $Q_j\alpha^k = Q_j\beta^k = Q_{j+1}$. Thus an endpoint of Q_j is dyadic if and only if the other endpoint is dyadic. In fact, the endpoints of our arcs $\bigcup_{0\le j<q} \partial Q_j$ form a single orbit under α (and also a single orbit under β). Because of this—and because of the convention on r_0 in Definition 3.4.1—we have two cases.

1. We could have $\partial Q_j \subseteq \mathbb{Z}[1/2]$ for each j. This occurs when I contains at least one dyadic point; I may or may not contain nondyadics.
2. The other option is that $\partial Q_j \cap \mathbb{Z}[1/2] = \emptyset$, for each j. This occurs when all points of I are nondyadic.

We cannot avoid the latter case,[22] because it is quite possible that a non-torsion element of T has no dyadic important points—see Figure 2.22 for an example.

[22] We cannot avoid this case, but our definition means we don't have to handle this case unless we have no other choice!

Despite naming them 'arcs', we think of each Q_j as subinterval of $\mathbb{R}$ rather than a circular arc. In particular, if $q = 1$ there is a single representative arc Q_0 equal to the whole circle S^1. We will think of this arc as the real interval $[r, r+1)$, in order to handle this case in the same way as the case $q > 1$.

The following claim will show that we only need to search for fine conjugators over a single representative arc.

Claim 3.4.2. *Let $\alpha, \beta \in T$ be coarsely equivalent, with rotation number p/q in lowest terms and representative arcs $\{Q_j\}_{0\le j<q}$.*

1. *A fine conjugator ϵ between α and β is uniquely determined by its restriction ϵ_0 to Q_0.*
2. *Any orientation-preserving homeomorphism $\epsilon_0 \colon Q_0 \to Q_0$ satisfying $(\alpha^q\big|_Q)^{\epsilon_0} = \beta^q\big|_Q$ uniquely extends to a fine conjugator $\epsilon \in \mathrm{Homeo}_+(S^1)$ given by*
$$t\epsilon = t\alpha^{-j}\epsilon_0\beta^j \qquad \text{if } t \in Q_0\alpha^j. \tag{3.6}$$

There's nothing special about Q_0 over the other Q_j: we just need to choose one representative arc in particular.

There is no guarantee that $\epsilon \in T$—we'll have to check this separately.

Proof. Let $U = V = Q_0 = [r_0, r_t)$ be our chosen representative arc. For any integer i we have

$$U\alpha^i = V\beta^i = Q_{ip \bmod q}. \tag{3.7}$$

Firstly, this means that the translates $\{U\alpha^i\}_{i\in\mathbb{Z}}$ cover the circle. Apply Lemma 3.1.1 to see that item 1 above holds.

Secondly, because the $\{Q_j\}$ are pairwise disjoint, Equation (3.7) means that properties 1 and 2 from Lemma 3.1.3 hold in our scenario. Apply that very lemma—with our map ϵ_0 in place of its δ_0—to construct

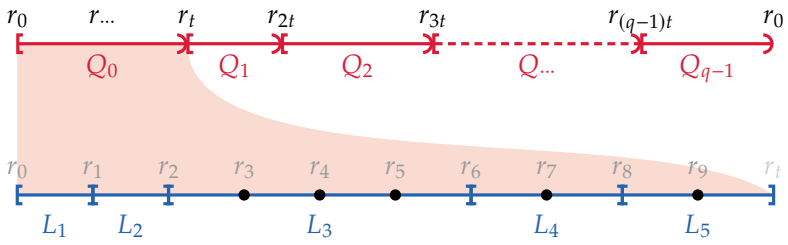

Figure 3.3: We split $Q = Q_0$ into subintervals L_i by partitioning Q at dyadic important points (I). On the L_i containing nondyadic important points (●), α^q and β^q are almost one-bump functions. On the remaining L_i, these powers are either both the identity or both one-bump functions.

N.B. In this sketch, ∂Q_0 consists of dyadic points; but this need not always be the case.

the fine conjugator ϵ. (Specifically, Equation (3.6) comes from Equation (3.2)). The uniqueness of ϵ is guaranteed by item 1. Finally, note that ϵ is an orientation-preserving homeomorphism because the restrictions $\epsilon\big|_{Q_0\alpha^j} = \left(\alpha^{-j}\epsilon_0\beta^j\right)\big|_{Q_0\alpha^j}$ are themselves orientation-preserving homeomorphisms defined over pairwise disjoint sets $\{Q_0\alpha^j\}$. This establishes item 2 above. □

From now on, we drop the subscript 0 and refer to our chosen representative arc as Q, rather than Q_0.

To proceed, we further divide Q into smaller 'blocks' on which the behaviour of α^q and β^q are manageable. Let $I' = I \cap \overline{Q} \cap \mathbb{Z}[1/2]$ be the set of dyadic important points in $\overline{Q}$. Add in the (potentially nondyadic) endpoints ∂Q to form the set $I'' = I' \cup \partial Q$. This contains at least two points, since we treat Q as a real interval and not a circular arc. Arrange I'' in linear order as $q_1 < \cdots < q_{m+1}$, and define the 'blocks' to be the intervals $L_i = [q_i, q_{i+1}]$, for $1 \le i \le m$. As the endpoints are important, each block is fixed setwise by α^q and β^q. Our choice of cut points I'' ensures that either $\alpha^q\big|_{L_i}$ and $\beta^q\big|_{L_i}$ are almost one-bump functions, or $\alpha^q\big|_{L_i} = \beta^q\big|_{L_i} = \mathrm{id}_{L_i}$.

We will construct the representative part $\epsilon_0 = \epsilon\big|_Q$ of a fine conjugator ϵ by searching for the potential restrictions $\epsilon_i = \epsilon_0\big|_{L_i}$ for $1 \le i < m$ and gluing them together. So now we need to search for a PL_2 map $\epsilon_i \colon L_i \to L_i$ which conjugates $\alpha^q\big|_{L_i}$ to $\beta^q\big|_{L_i}$. For brevity, drop the subscript i and restriction symbols; thus we're trying to find a conjugator ϵ between α^q and β^q in $\mathrm{PL}_2(L)$. Also write $L = [r, s]$, so that we have a name for each endpoint.

Suppose first that L is fixed pointwise by α^q. As $\mathrm{Fix}(\alpha^q) = \mathrm{Fix}(\beta^q)$ we conclude that L is also fixed pointwise by β^q. Then $\alpha^q = \beta^q = \mathrm{id}_L$ are conjugate in $\mathrm{PL}_2(L)$, and we may choose any element of $\mathrm{PL}_2(L)$ as our conjugator ϵ. To keep things simple, we can choose $\epsilon = \mathrm{id}_L$.

Otherwise L is not fixed pointwise by α^q, nor by β^q. Then these maps restrict to non-identity, almost one-bump functions on L. We turn to the following proposition[23] which greatly reduces our search space. In short, the result says that there is at most one conjugator $\epsilon \in \mathrm{PL}_2(L)$ with a given initial or final gradient. Its proof explicitly constructs a map $\zeta \in \mathrm{PL}_2(L)$ using the 'stair algorithm', and then checks to see if ζ conjugates α^q to β^q. If so, we can use $\epsilon = \zeta$ and we're done; if not, there is no conjugator ϵ on L.

Proposition 3.4.3. *Let $L = [r, s]$ be an interval with $\partial L \subseteq S$, and assume*

dyadic important points	
do exist	do not exist
r_0 dyadic	no r_i dyadic
$I'' = I'$	$I'' = \partial Q$, $I' = \emptyset$
$m \ge 1$ block(s)	$m = 1$ block

Figure 3.4: A summary of how we divide Q into blocks, depending on whether ∂Q is dyadic. If we have a dyadic important point then $\partial Q \subseteq I'$, so that $I'' = I'$. If there are no dyadic important points then then I' is empty, meaning that $m + 1 = |I''| = |\partial Q| = 2$.

[23] Kassabov and Matucci 2012, Proposition 4.20 and Remark 4.22. The computation requirements are discussed in their Section 3.

that S and G satisfy certain computational requirements. Let $\alpha, \beta \in \mathrm{PL}^0_{S,G}(L)$ be almost one-bump functions. Fix a gradient $g \in G$.

We can algorithmically construct a conjugator $\epsilon \in \mathrm{PL}_{S,G}(L)$ between α and β with initial gradient $r^+\epsilon' = g$. If such an element ϵ exists, it is uniquely determined[24] *by the initial gradient $r^+\epsilon'$. Similarly, we can construct a conjugator $\bar{\epsilon}$ with a given final gradient $s^-\bar{\epsilon}' = g$, or determine that no such $\bar{\epsilon}$ exists. Again, such an $\bar{\epsilon}$ is uniquely determined by $s^-\bar{\epsilon}'$.*

[24] See also our discussion in and around Lemma 3.2.3.

We normally use the breakpoint group $S = \mathbb{Z}[1/2]$ and gradient group $G = 2^{\mathbb{Z}}$. Unfortunately, doing so here would prevent us from using the proposition, because the endpoints $\partial L = \{r, s\}$ are both nondyadic. Instead, we note that our endpoints are always rational, as they're important points (see Remark 2.3.13). The choice of breakpoint and gradient groups $S = \mathbb{Q}$ and $G = \mathbb{Q}_{>0}$ satisfies K&M's computational requirements,[25] so we can apply Proposition 3.4.3 to the group $\mathrm{PL}_{\mathbb{Q}}(L)$. Since we are seeking conjugators in $\mathrm{PL}_2(L)$, we'll only use initial gradients $g \in 2^{\mathbb{Z}}$ as inputs (even though any rational gradient is permitted by K&M's algorithm). If this results in a conjugator $\epsilon \in \mathrm{PL}_{\mathbb{Q}}(L)$, we can inspect gradients and breakpoints to determine if $\epsilon \in \mathrm{PL}_2(L)$.

[25] Kassabov and Matucci 2012, Example 9.1.

Next we follow K&M in establishing bounds[26] on the initial gradient of a conjugator ϵ (if one exists at all).

[26] Kassabov and Matucci 2012, Section 7.

Lemma 3.4.4. *Suppose $L = [r, s]$ is not fixed pointwise by $\alpha^q \in \mathrm{PL}_{S,G}(L)$, nor by $\beta^q \in \mathrm{PL}_{S,G}(L)$. Let $A = r^+(\alpha^q)'$. If α^q and β^q are conjugate in $\mathrm{PL}_{S,G}(L)$, then they are conjugate via a map $\epsilon \in \mathrm{PL}_{S,G}(L)$ whose initial gradient $r^+\epsilon'$ belongs to the real interval with endpoints 1 (inclusive) and A (exclusive).*

Proof. If the elements are conjugate via $\delta \in \mathrm{PL}_2(L)$, they are also conjugate via $\alpha^{qj}\delta$ for each $j \in \mathbb{Z}$. Let $g = r^+\delta'$. The initial gradient of $\alpha^{\pm q}\delta$ is

$$r^+(\alpha^{\pm q}\delta)' = r^+(\alpha^{\pm q})' \cdot [r\alpha^{\pm q}]^+\delta' = A^{\pm 1} \cdot r^+\delta' = A^{\pm 1}g\,.$$

It follows that $\alpha^{jq}\delta$ is a conjugator between α^q and β^q with initial gradient $A^j g$, for each $j \in \mathbb{Z}$.

Let $m = \lfloor \log_A g \rfloor$. We must have $A \neq 1$, or else α^q would be the identity on some small interval containing r—which would mean that α^q is not an almost one-bump function. If $A > 1$ we have $A^m \leq g < A^{m+1}$, so $1 \leq A^{-m}g < A$. Hence the initial gradient of $\epsilon = \alpha^{-mq}\delta$ belongs to $[1, A)$. Otherwise we have $A < 1$. In this case $A^{m+1} < g \leq A^m$, so $A < gA^{-m} \leq 1$. This time, the initial gradient of ϵ belongs to $(A, 1]$. □

$\lfloor x \rfloor$ denotes the floor of x, the greatest integer $n \leq x$.

Since the intervals $(A, 1]$ and $[1, A)$ contain finitely many powers of two, we can enumerate all of these powers 2^k. For each power, we use Proposition 3.4.3 to construct a conjugator ϵ_k over L with $r^+\epsilon_k' = 2^k$. If there is no such element ϵ_k, then there cannot exist a fine conjugator δ between α^q and β^q.

Because we can also use Proposition 3.4.3 to test for the existence of conjugators with a given *final* gradient, we obtain the following.

Lemma 3.4.5. *Let α^q, β^q and L and δ be as in Lemma 3.4.4 and let $B = r^+(\alpha^q)'$. If the restrictions $\alpha^q\big|_L$ and $\beta^q\big|_L$ are conjugate in $\mathrm{PL}_{S,G}(L)$, then they are conjugate via a map $\epsilon \in \mathrm{PL}_2(L)$ whose* final *gradient $s^-\epsilon'$ belongs to the real interval with endpoints 1 (inclusive) and B (exclusive).*

Sketch proof. The proof is very similar: replace A with B; r^+ with s^-; and the word 'initial' with 'final'. □

We summarise our discussion with a claim.

Claim 3.4.6. *Let $\alpha, \beta \in T$ be coarsely equivalent and let L be a block as discussed above. If L is fixed pointwise by α^q, then any function $\epsilon \in \mathrm{PL}_2(L)$ conjugates $\alpha^q\big|_L$ to $\beta^q\big|_L$; in particular we may take $\epsilon = \mathrm{id}_L$. Otherwise there exists a conjugator between these two restrictions if and only if there exists a conjugator $\epsilon_g \in \mathrm{PL}_2(L)$ with initial (final) gradient g belonging to a constructable finite set Γ. For each initial (final) gradient $g \in \Gamma$, we can effectively construct such an element ϵ_g, or prove that no such ϵ_g exists.*

Thus we can decide if $\alpha^q\big|_L$ and $\beta^q\big|_L$ are conjugate in $\mathrm{PL}_2(L)$, and produce a conjugator if so.

3.4.2 *Reassembly*

To solve the aligned conjugacy problem we now turn to brute force. We use Claim 3.4.6 to enumerate conjugators on each L_i. Gluing these together, we obtain a conjugator ϵ_0 on the representative arc Q. This extends to a conjugator ϵ defined on the whole circle, according to Claim 3.4.2. Now certainly ϵ must be an element of $\mathrm{PL}_{\mathbb{Q}}(S^1)$. By inspecting gradients and breakpoints, we can then check if $\epsilon \in T$.

The next lemma enables a small optimisation to this process: it allows us to determine if $\epsilon \in T$ given only ϵ_0.

Lemma 3.4.7. *Let $\alpha, \beta \in T$ be coarsely equivalent, and let $Q = [r, s]$ be a representative arc. Suppose $\epsilon_0 \in \mathrm{PL}_2(Q)$ is extended to a candidate conjugator $\epsilon \in \mathrm{PL}_{\mathbb{Q}}(S^1)$. If ∂Q consists of dyadics, then $\epsilon \in T$. If ∂Q is nondyadic and $q = 1$, then $\epsilon \in T$ if and only if*

$$s^-\epsilon_0' = r^+\epsilon_0'. \tag{3.8}$$

If ∂Q is nondyadic and $q > 1$, then $\epsilon \in T$ if and only if

$$s^-\epsilon_0' = s(\alpha^{-k})' \cdot r^+\epsilon_0' \cdot r(\beta^k)', \tag{3.9}$$

where $k = p^{-1} \bmod q$.

Proof. The map ϵ is formed by gluing the functions $\alpha^{-j}\epsilon_0\beta^j$ together. Call these functions $\epsilon_j \colon Q\alpha^j \to Q\alpha^j$, for each $0 \le j < q$. We know ϵ must have finitely many linear segments (each with gradient equal to an integer power of two) because each ϵ_i is a PL_2 map. To conclude

We can't use Lemma 2.1.7 to conclude that $\epsilon \in T$, because ∂Q need not be dyadic.

that $\epsilon \in T$, we need to locate the breakpoints of ϵ, and verify that each is dyadic. The breakpoints in the interior of $Q\alpha^j$ are dyadic, because each ϵ_j is a PL_2 map. Any remaining breakpoints are in $\partial(Q\alpha^j) = (\partial Q)\alpha^j$ for some j. Because T acts on the dyadics, if ∂Q consists of dyadics then all breakpoints of every ϵ are dyadic, so $\epsilon \in T$.

In the remaining case, ∂Q consists of nondyadics. First we handle the situation where $q = 1$ and hence $Q = S^1$. In these circumstances, r and $s = r+1$ represent the same point r on the circle. It's possible that r is a nondyadic breakpoint of ϵ. This *does not* occur if and only if the initial and final gradients of ϵ_0 are equal, which justifies Condition (3.8).

Otherwise ∂Q is nondyadic and $q > 1$. We know from the definition of a representative arc that $s = r\alpha^k = r\beta^k$, where $k = p^{-1} \bmod q$. For each $0 \le j < q$, the point $x_j = r\alpha^{(j+1)k}$ is the right endpoint of Q_j and the left endpoint of Q_{j+1}. Now $\epsilon \in T$ if and only each of these gluing points x_j is *not* a breakpoint. The left-hand derivative is

$$\begin{aligned}[r\alpha^{(j+1)k}]^-\epsilon' &= [r\alpha^{(j+1)k}]^-(\alpha^{-jk}\epsilon_0\beta^{jk})' \\ &= [r\alpha^{(j+1)k}]^-(\alpha^{-jk})' \cdot [r\alpha^k]^-\epsilon_0' \cdot [r\alpha^k\epsilon_0]^-(\beta^{jk})' \\ &= [r\alpha^{(j+1)k}]^-(\alpha^{-jk})' \cdot s^-\epsilon_0' \cdot s^-(\beta^{jk})', \qquad (3.10)\end{aligned}$$

The third equality holds because $s = r\alpha^k$ and $s\epsilon_0 = s$.

according to the chain rule. Similarly the right-hand derivative is

$$\begin{aligned}[r\alpha^{(j+1)k}]^+\epsilon' &= [r\alpha^{(j+1)k}]^+(\alpha^{-(j+1)k}\epsilon_0\beta^{(j+1)k})' \\ &= [r\alpha^{(j+1)k}]^+(\alpha^{-(j+1)k})' \cdot r^+\epsilon_0' \cdot r^+(\beta^{(j+1)k})'. \qquad (3.11)\end{aligned}$$

We are using the fact that $r\epsilon_0 = r$ here.

Thus $\epsilon \in T$ if and only if Expressions (3.10) and (3.11) are equal, for every value of j. After some rearrangement, this condition takes the form

$$\frac{s^-\epsilon_0'}{r^+\epsilon_0'} = \frac{[r\alpha^{(j+1)k}](\alpha^{-(j+1)k})'}{[r\alpha^{(j+1)k}](\alpha^{-jk})'} \cdot \frac{r(\beta^{(j+1)k})'}{s(\beta^{jk})'}. \qquad (3.12)$$

We are always differentiating α and β at a point t of the form $t = r\alpha^i$, for some i. As r is nondyadic, so too are the points in its orbit. Thus the left and right derivatives of α and β at t are the same. This allows us to drop some of the superscripts ± on derivatives.

Let A_j (resp. B_j) be the left (resp. right) fraction on the right-hand side of Equation (3.12). Apply the chain rule to the numerator of A_j to see that

$$A_j = \frac{[r\alpha^{(j+1)k}](\alpha^{-jk})' \cdot [r\alpha^k](\alpha^{-k})'}{[r\alpha^{(j+1)k}](\alpha^{-jk})'} = [r\alpha^k](\alpha^{-k})' = s(\alpha^{-k})'$$

is independent of j. Using the fact that $s = r\alpha^k = r\beta^k$, we see in a similar fashion that

$$B_j = \frac{r(\beta^k)' \cdot [r\beta^k](\beta^{jk})'}{[r\beta^k](\beta^{jk})'} = r(\beta^k)'$$

is also independent of j. Thus we see that $\epsilon \in T$ if and only if

$$\frac{s^-\epsilon'}{r^+\epsilon'} = A_jB_j = s(\alpha^{-k})' \cdot r(\beta^k)',$$

which is equivalent to Condition (3.9). Note that this is the equality between Expressions (3.10) and (3.11) in the special case $j = 0$. □

Now we assemble our work in this section into an algorithm to search for fine conjugators.

Algorithm 3.4.8 (Finding fine conjugators)**.** The following procedure solves the aligned conjugacy problem by building[27] a fine conjugator ϵ for coarsely equivalent elements $\alpha, \beta \in T$. Let their common rotation number be p/q in lowest terms, and let I denote their shared important point set.

[27] If no such ϵ exists, the algorithm returns the message 'no fine conjugator'.

1. Split the circle at the orbit of a important point r, forming representative arcs $Q = Q_0, \ldots, Q_{q-1}$. If possible, choose r to be dyadic.
2. Partition $Q = [r, s]$ at its interior dyadic points, breaking it into blocks $L_1, \ldots, L_m$.
3. If ∂Q is dyadic:
   (a) For each $1 \le i \le m$:
      i. Use Claim 3.4.6 to construct a conjugator $\epsilon_i \in \mathrm{PL}_2(L_i)$.
      ii. If this is not possible, there is no fine conjugator between α and β. We terminate the algorithm, returning the message 'no fine conjugator'.
   (b) Assemble $\epsilon_0 \colon Q \to Q$ by gluing together $\epsilon_1, \ldots, \epsilon_m$.
4. If ∂Q is nondyadic:[28]
   (a) Let $A = r^+(\alpha^q)'$. If $A > 1$, let M be the interval $[1, A)$; otherwise let $M = (A, 1]$.
   (b) For each integer power of two $g = 2^j$ in M:
      i. Use Proposition 3.4.3 to build a conjugator $\epsilon_{0,g} \in \mathrm{PL}_2(Q)$ with initial gradient g. If this is not possible, continue to the next value of g.
      ii. If $q = 1$, let $k = 0$; otherwise let $k = p^{-1} \bmod q$. Check that $s^- \epsilon'_{0,g} = s(\alpha^{-k})' \cdot g \cdot r(\beta^k)'$. If the check passes, store $\epsilon_{0,g}$ as ϵ_0 for use later, and break the loop over g.
   (c) If we exhaust all values of g without storing an element ϵ_0, then there is no fine conjugator between α and β. We terminate the algorithm, returning the message 'no fine conjugator'.
5. Extend ϵ_0 to $\epsilon \colon S^1 \to S^1$ using Claim 3.4.2. Return ϵ.

[28] In this situation, $m = 1$ and there is only one block $L_1 = Q$. See Figure 3.4.

Proof of correctness. First we show that any value ϵ returned by the algorithm is a fine conjugator in T. The only point at which this occurs is step 5. The claim referred to in this step guarantees that ϵ is a conjugator between α and β, provided that ϵ_0 is a conjugator between $\alpha^q\big|_Q$ and $\beta^q\big|_Q$. If ∂Q is nondyadic then this is true thanks to the success of step 4(b)i. Otherwise ∂Q is dyadic, and ϵ_0 is a conjugator over Q because it is formed from the ϵ_i in step 3b, each of which is a conjugator between $\alpha^q\big|_{L_i}$ and $\beta^q\big|_{L_i}$. In either case, ϵ_0 fixes the important points

in Q, and hence ϵ fixes all important points I. Thus $\epsilon \in \mathrm{PL}_{\mathbb{Q}}(S^1)$ is a fine conjugator. Finally we check that $\epsilon \in T$. Lemma 3.4.7 tells us this is automatically the case when ∂Q is dyadic. Otherwise we know from the success of step 4(b)ii that Condition (3.8) or (3.9) (as appropriate) is satisfied, so $\epsilon \in T$ by the same lemma.

Now we need to explain why the algorithm must produce an output ϵ whenever a fine conjugator γ exists. If ∂Q is dyadic, then step 3(a)i will always succeed, because $\gamma\big|_{L_i}$ is a conjugator between $\alpha^q\big|_{L_i}$ and $\beta^q\big|_{L_i}$. We then proceed to step 5 without any problems.

The alternative is that ∂Q is nondyadic, in which case $m = 1$ and $Q = L_1$ (see Figure 3.4). In step 4(b)i, there must exist a conjugator ϵ_0 with initial gradient in the given range M. An example of the form $\epsilon_{0,g} = (\alpha^{xq}\gamma)\big|_{L_1}$ is guaranteed to exist with initial gradient $g = [r(\alpha^q)']^x \cdot r\gamma'$, according to Lemma 3.4.4. We now show that step 4 is successfully completed when (if) the loop considers this particular example $\epsilon_{0,g}$; this means passing the gradient check in step 4(b)ii.

Let $x \in \{r, s\}$ and $\omega \in \{\alpha, \beta, \gamma\}$. When considering step 4, we have $x \notin \mathbb{Z}[1/2]$. Since $\omega \in T$, we know that $x^-\omega' = x^+\omega'$ in this situation. Thus we may drop some of the superscripts $\pm$ for brevity, as in Lemma 3.4.7.

If $q = 1$ then $\epsilon_{0,g} = \alpha^{xq}\gamma$ is an element of T, thought of as a PL_2 map on $[r, s] = [r, r + 1]$. So it will certainly be the case that $s^-\epsilon_0 = r^+\epsilon_0$, because $r = s \bmod 1$ is a nondyadic, which cannot be a breakpoint of $\alpha^{xq}\gamma \in T$. We pass the gradient check and successfully exit the loop.

Otherwise $q > 1$, and we need to check that

$$s^-\epsilon'_{0,g} = s(\alpha^k)' \cdot g \cdot r(\beta^k)'. \tag{3.13}$$

By applying the chain rule, we see that $s^-\epsilon'_{0,g} = s(\alpha^{xq})' \cdot s\gamma' = [s(\alpha^q)']^x \cdot s\gamma'$. With this and the value of g above in hand, we see that Equation (3.13) is equivalent to

$$[s(\alpha^q)']^x \cdot s\gamma' = s(\alpha^k)' \cdot [r(\alpha^q)']^x \cdot r\gamma' \cdot r(\beta^k)'.$$

Because $s = r\alpha^k$ belong to the same α-orbit and are fixed by α^q, their derivatives under α^q are equal (by Remark 2.3.18). Thus we can simplify to obtain the equivalent equation

$$s\gamma' = s(\alpha^k)' \cdot r\gamma' \cdot r(\beta^k)'.$$

This is true by Lemma 3.4.7,[29] and so Equation (3.13) is satisfied. This means we will pass the gradient check and complete step 4. We move on to step 5, which will return a conjugator $\epsilon \in T$, as explained in the first paragraph of this proof. □

[29] Apply the lemma with γ in place of ϵ. As $\gamma \in T$ is a fine conjugator, Condition (3.9) must hold—which is exactly what we wanted to show.

With this proof and Algorithm 3.3.5 above in hand, we have the tools to construct conjugators in T.

Theorem 3.4.9. *The conjugacy search problem in T has an effective solution.*

4 *Centralisers in T*

What should we do if one conjugator is not enough? Suppose $\alpha^\gamma = \beta$ in some group G. If δ is any other conjugator with $\alpha^\delta = \beta$, then the difference $\delta\gamma^{-1}$ belongs to the centraliser $C_G(\alpha)$. Conversely if $\epsilon \in C_G(\alpha)$ then $\alpha^{\epsilon\gamma} = \beta$. Thus the set of all conjugators $C_G(\alpha, \beta) = \{\delta \in G \mid \alpha^\delta = \beta\}$ is exactly the right coset $C_G(\alpha)\gamma$ represented by a (any) conjugator γ. On the other hand, the difference $\delta^{-1}\gamma$ belongs to $C_G(\beta)$. Again if $\epsilon \in C_G(\beta)$, we see that $\beta^{\epsilon\gamma^{-1}} = \alpha$, meaning that $\alpha^{\gamma\epsilon^{-1}} = \beta$. Thus $C_G(\alpha, \beta)$ is also the left coset $\gamma C_G(\beta)$. In short: element centralisers parameterise the set of available conjugators. For this reason, the study of centralisers is naturally linked with the study of conjugators.

Let us briefly summarise what is known about centralisers for each of the main trio of Thompson groups. Guba and Sapir's work on diagram groups includes the result[1] that element centralisers in F are direct products of $\mathbb{Z}$ and F; in a follow-up article they show[2] that these are undistorted in F. From the piecewise-linear point of view, Kassabov and Matucci[3] (K&M) obtain the same direct product result, and show that a similar result holds for the centraliser of any finitely generated subgroup of F. As a corollary, they note that the intersection of any finite number element centralisers $\{C_F(\alpha_i)\}_{1\le i\le n}$ is equal to an intersection of two element centralisers $C_F(\beta_1) \cap C_F(\beta_2)$.

[1] Victor Guba and Mark Sapir 1997, Corollary 15.36.

[2] V.S. Guba and M.V. Sapir 1999, Theorem 34.

[3] Kassabov and Matucci 2012.

At the other end of the spectrum lies V. Bleak et al[4] describe an element centraliser $C_V(\alpha)$ by using revealing pairs to understand the dynamics of α. They show that $C_V(\alpha)$ is finitely generated, after expressing this centraliser as a direct product involving two kinds of factors. The first kind is a semidirect product involving the Higman-Thompson[5] group $V_{n,r}$; this is a version of V whose gradients are powers of n and whose functions permute $[0, r]$ instead of $[0, 1]$. The second is a wreath product of the base group $A \wr \mathbb{Z}$ with a finite symmetric group, for some finite group A. Recently, Bieniecka[6] has shown that every finite group A and every semidirect product $A \rtimes \mathbb{Z}$ is involved in the centraliser of some element of V.

[4] Bleak, Bowman et al. 2013.

[5] Higman 1974.

[6] Bieniecka 2018.

Martínez-Pérez, Matucci and Nucinkis consider a generalisation of V called $V_r(\Sigma)$, which consists of automorphisms of a Cantor algebra Σ. An argument they use[7] shows that the first of Bleak et al's factors is of type F_∞. The same is true of the second kind of factor, from which it follows that the whole centraliser $C_V(\alpha)$ is finitely presented (not just finitely generated). In a later article, the three authors give an explicit finite presentation[8] for centralisers of finite subgroups.

[7] Martínez-Pérez, Matucci and Nucinkis 2016, Theorem 4.9.

[8] Martínez-Pérez, Matucci and Nucinkis 2018, Section 7.

In the middle lies T. A chapter of Matucci's thesis discusses element

centralisers in T (as well as larger groups of circle homeomorphisms), handling torsion and nontorsion elements separately. He shows[9] that a torsion element α has a centraliser $C_T(\alpha)$ which is a nonsplit extension of a finite cyclic group by T. Geoghegan and Varisco separately show the same result, including the extra observation[10] that the normaliser $N_T(\alpha)$ is equal to $C_T(\alpha)$. Martínez-Pérez and Nucinkis generalise Matucci's result to the group $T_r(\Sigma)$, where again Σ is a Cantor algebra; they exhibit[11] the centraliser of a finite subgroup as a central extension.

Matucci's thesis also describes[12] nontorsion elements' centralisers as an extension. In this chapter, we fill a small gap in his proof, and study this extension in more detail. We will make use of K&M's work on centralisers in F when considering this extension's kernel.

[9] Matucci 2008, Theorem 7.1.5.

[10] Geoghegan and Varisco 2017, Theorem 1.3.

[11] Martínez-Pérez and Nucinkis 2013.

[12] Matucci 2008, Theorem 7.2.5.

4.1 Centralisers in F

We will continue to adapt results in F to results in T. This time we're concerned with results about centralisers.

Definition 4.1.1. Let G be a group containing an element α and let $0 \neq n \in \mathbb{Z}$. An *nth root* of α is an element $\omega \in G$ for which $\omega^n = \alpha$. We call ω a *root* of α if it is an nth root of α, for some $0 \neq n \in \mathbb{Z}$. If α does not have an nth root for any $|n| \geq 2$, we say that α is a *root element*. A root element ω which is a root of α is called a *minimal root* of α.

Any element commutes with any power of itself and any of its roots, so if ω is a minimal root of α, we know that $\langle\omega\rangle \leq C_G(\omega)$. With a little effort to handle nondyadic points, K&M explain[13] that this inclusion is an equality for almost one-bump functions in F. Their statement requires that $\partial J \subseteq S$. However, this requirement is not used in their proof, nor in the results they refer to (which describe centralisers in $\mathrm{PL}_+(J)$).

[13] Kassabov and Matucci 2012, Theorem 5.5.

Theorem 4.1.2. *Let $\alpha \in \mathrm{PL}^0_{S,G}(J)$ be an almost one-bump function. Then $C_{\mathrm{PL}_{S,G}(J)}(\alpha)$ is infinite cyclic, generated by a minimal root ω of α. A generator ω can be constructed algorithmically.*

Sketch proof. K&M's stair algorithm shows that one-bump functions $\beta \in \mathrm{PL}_{S,G}([a,b])$ have infinite cyclic centralisers generated by a minimal root. This is because an element which commutes with β can be uniquely reconstructed from its initial gradient $a^+\beta'$ alone.[14] It follows that $C_T(\beta)$ is isomorphic to the group $G_\beta = \{a^+\gamma' \mid \gamma \in C_T(\beta)\} \leq G$ of initial gradients. In the case of Thompson's main trio, $G = 2^{\mathbb{Z}}$ is infinite cyclic and hence $G_\beta \cong C_T(\beta)$ is also infinite cyclic. More generally, it can be shown[15] that G_β is a discrete subgroup of $\mathbb{R}_{>0}$, and hence is infinite cyclic.

Enumerate $\mathrm{Fix}(\alpha) = \{j_1 < \cdots < j_n\}$ so that Let $J = [j_1, j_n]$. Because α is almost one-bump, $\{j_2 < \cdots < j_{n-1}\} \not\subseteq S$ is a discrete set of 'nondyadic' points. Then α can be formed by gluing together the one-bump functions $\alpha\big|_{D_i}$ for $1 \leq i < n$, where $D_i = [j_i, j_{i+1}]$. If γ

[14] This reconstruction property is the root cause of the 'miracle of F': see Section 3.2.

[15] Kassabov and Matucci 2012, Theorem 5.1, Lemmas 5.2–5.3.

commutes with α, then $\gamma\big|_{D_i}$ commutes with $\alpha\big|_{D_i}$. Because the j_i are nondyadic, the final gradient of $\gamma\big|_{D_i}$ is equal to the initial gradient of $\gamma\big|_{D_i}$ (or else j_{i+1} is a breakpoint of γ). Thus the initial gradient of $\gamma\big|_{D_1}$ determines the whole function $\gamma\big|_{D_1}$; this determines the initial gradient of and thus the entire function $\gamma\big|_{D_2}$. This continues up to $\gamma\big|_{D_{n-1}}$, so that γ is entirely determined by its initial gradient. It follows that

$$C_{\mathrm{PL}_{S,G}(J)}(\alpha) \cong C_{\mathrm{PL}_{S,G}([j,k])}(\alpha\big|_{[j,k]}), \tag{4.1}$$

for any $j < k$ in $\mathrm{Fix}(\alpha)$. In particular, this group is infinite cyclic. □

The next step[16] is to use this result to consider the centraliser of an arbitrary element of F. In short, these centralisers are finite products of $\mathbb{Z}$ and F.

[16] Kassabov and Matucci 2012, Theorem 5.10.

Theorem 4.1.3 (Description of centralisers in F)**.** *Let $\alpha \in \mathrm{PL}_{S,G}(J)$. Its centraliser $C_{\mathrm{PL}_{S,G}(J)}(\alpha)$ is isomorphic to a product $\mathbb{Z}^z \times \prod_{i=1}^{f} \mathrm{PL}_{S,G}(D_i)$, for some intervals $D_i \subseteq J$ with $\partial D_i \subset S$, and some non-negative integers f and z not both zero. Moreover,[17] the centralisers $C_{\mathrm{PL}_{S,G}(J)}(\alpha)$ and $C_{\mathrm{PL}_{S,G}(J)}(\alpha^k)$ are equal, for any $k \neq 0$.*

[17] At the time of writing, the author has not been able to find the result given in this sentence in the literature.

Sketch proof. Find the 'dyadic' important points $I_\alpha \cap S = \partial\,\mathrm{Fix}(\alpha) \cap S$ and enumerate them as $r_0 < \cdots < r_n$, so that $J = [r_0, r_n]$. Define D_i to be the interval $[r_i, r_{i+1}]$ for $0 \leq i < n$. The centraliser of α is isomorphic to the direct product

$$C_{\mathrm{PL}_{S,G}(J)}(\alpha) \cong \prod_{0 \leq i < n} C_{\mathrm{PL}_{S,G}(D_i)}(\alpha\big|_{D_i}).$$

To be more specific, it is the internal direct product[18]

$$C_{\mathrm{PL}_{S,G}(J)}(\alpha) = \prod_{0 \leq i < n} C_{\mathrm{PL}_{S,G}(D_i)}(\alpha\big|_{D_i})\theta_i,$$

[18] Let I_0 be a subset of $I_\alpha \cap S$ containing $\partial J = \{r_0, r_n\}$. This internal direct product decomposition still holds if we define the D_i in terms involving only our chosen important points I_0, rather than the full collection $I_\alpha \cap S$.

where $\theta_i \colon \mathrm{PL}_{S,G}(D_i) \to \mathrm{PL}_{S,G}(J)$ is an injective homomorphism. The image $\gamma\theta_i$ of $\gamma \in \mathrm{PL}_{S,G}(D_i)$ is defined by

$$t(\gamma\theta_i) = \begin{cases} t\gamma & \text{if } t \in D_i \\ t & \text{otherwise.} \end{cases}$$

In other words, θ_i extends maps $D_i \to D_i$ to maps $J \to J$ by making them identity on the complement $J \setminus D_i$.

If $\alpha\big|_{D_i}$ is the identity then $C_{\mathrm{PL}_{S,G}(D_i)}(\alpha\big|_{D_i}) = \mathrm{PL}_{S,G}(D_i) \cong \mathrm{PL}_{S,G}(I)$; otherwise $\alpha\big|_{D_i}$ is almost one-bump, so $C_{\mathrm{PL}_{S,G}(D_i)}$ is infinite cyclic by Theorem 4.1.2. There are no other possibilities because there are no important dyadic points in the interior of D_i.

For the last part of this claim,[19] let $k \neq 0$. As $I(\alpha^k) = I_\alpha$, we obtain a very similar internal direct product decomposition

[19] This isn't explicitly written down in the literature, but can be deduced quickly given K&M's description.

$$C_{\mathrm{PL}_{S,G}(J)}(\alpha^k) = \prod_{0 \leq i < n} C_{\mathrm{PL}_{S,G}(D_i)}(\alpha^k\big|_{D_i})\theta_i$$

to that of $C_{\mathrm{PL}_{S,G}(J)}(\alpha)$. Again, each restriction is an almost one-bump

function. One possibility is that $\alpha^k\big|_{D_i} = \mathrm{id}$, which would imply that $\alpha\big|_{D_i} = \mathrm{id}$, because F-like groups are torsion-free (Proposition 2.3.19). On the other hand $\alpha^k\big|_{D_i}$ could be nontrivial; then its root $\alpha\big|_{D_i}$ would be too. In both cases, $C_{\mathrm{PL}_{S,G}(D_i)}(\alpha^k\big|_{D_i}) = C_{\mathrm{PL}_{S,G}(D_i)}(\alpha\big|_{D_i})$ (again by Theorem 4.1.2). So the two internal direct products in this proof are identical. □

Can we partly generalise this result to T? We reuse a trick from Chapters 2 and 3, where we exploit the ability to conjugate by rotations. Before that, some notation.

Definition 4.1.4. Let H be a group of homeomorphisms of a space X. We denote by H_0 the subset of H whose elements have fixed points.

It is *not* automatically true that H_0 is a subgroup of H: see for instance Figure 2.5.

Corollary 4.1.5. *Suppose $\alpha \in T \setminus \{\mathrm{id}\}$ has a dyadic fixed point r. Then $C_T(\alpha)_0$ is a subgroup of T isomorphic to a direct product $\mathbb{Z}^z \times F^f$.*

Proof. Without loss of generality, we may assume that r is important: if not, r belongs to a fixed interval which is not the entire circle as $\alpha \neq \mathrm{id}$. Then we may replace r with either endpoint of this interval.

Let $\rho = \rho_{-r}$ be the rotation by $-r$. Then $C_T(\alpha)_0$ is equal[20] to $[C_T(\alpha^\rho)_0]^{\rho^{-1}}$. Now α^ρ fixes 0, so $\alpha^\rho \in F$. We claim that $C_T(\alpha^\rho)_0$ is equal to $C_F(\alpha^\rho)$. If $\gamma \in C_T(\alpha^\rho)_0$ then γ permutes the important points $I(\alpha^\rho) = I(\alpha)\rho$; in particular this set contains $r\rho = 0$. Now $\gamma \in C_T(\alpha^\rho)_0$ has a fixed point and hence $\mathrm{rot}(\gamma) = 0$, so γ must fix 0. Thus $\gamma \in F$, so $C_T(\alpha^\rho)_0 \subseteq C_F(\alpha^\rho)$. The other containment $C_T(\alpha^\rho)_0 \supseteq C_F(\alpha^\rho)$ is immediate.

We may now argue that

$$C_T(\alpha)_0 = [C_T(\alpha^\rho)_0]^{\rho^{-1}} = C_F(\alpha^\rho)^{\rho^{-1}} \cong C_F(\alpha^\rho)\,.$$

By Theorem 4.1.3, this group is isomorphic to a product $\mathbb{Z}^z \times F^f$. In a little more detail, Section 4.1 tells us that

$$C_F(\alpha^\rho) = \prod_{0\leq i<n} C_{\mathrm{PL}_2(D_i)}(\alpha^\rho\big|_{D_i})\theta_i\,,$$

where the D_i are delimited by the dyadic important points of α^ρ. Thus

$$\begin{aligned} C_T(\alpha)_0 = C_F(\alpha^\rho)^{\rho^{-1}} &= \prod_{0\leq i<n} \left[C_{\mathrm{PL}_2(D_i)}(\alpha^\rho\big|_{D_i})\theta_i\right]^{\rho^{-1}} \\ &= \prod_{0\leq i<n} C_{\mathrm{PL}_2(D_i')}(\alpha\big|_{D_i'})\theta_i'\,, \end{aligned}$$

where $D_i' = D_i\rho^{-1}$ and θ' extends maps on D_i' to maps on J which are the identity on $J \setminus D_i'$. The intervals D_i' are delimited by the important points of α. □

[20] We're using $C_G(\alpha)^\beta = C_{G^\beta}(\alpha^\beta)$ here, and the fact that α has a fixed point iff α^β has a fixed point.

Corollary 4.1.6. *Suppose $\alpha \in T$ has $n \geq 1$ fixed points, all of which are nondyadic. Then $C_T(\alpha)_0 \cong \mathbb{Z}$, generated by a minimal root ω of α in $C_T(\alpha)_0$. Moreover $C_T(\alpha)_0 \cong C_{\mathrm{PL}_2([r,s])}(\alpha\big|_{[r,s]})$, where $r \neq s$ are distinct fixed points of α.*

Proof. Let r be a nondyadic fixed point of α. Take an element $\gamma \in C_T(\alpha)$, and view γ and α as PL_2 functions $\bar{\gamma}, \bar{\alpha}$ on $[r, r+1]$. Then certainly $\bar{\gamma} \in C_{\mathrm{PL}_2([r,r+1])}(\bar{\alpha})$; thus we have an injective homomorphism $\bar{\cdot}\colon C_T(\alpha) \hookrightarrow C_{\mathrm{PL}_2([r,r+1])}(\bar{\alpha})$. To check that this is a surjection, we need to ensure that every element β of the codomain group has the same initial and final gradient. (Then β describes a map $\hat{\beta} \in \mathrm{PL}_2(S^1)$, because r is not a breakpoint of $\hat{\beta}$.) We do so in Remark 4.2.2; thus $\bar{\cdot}$ is an isomorphism.

Now because α has no dyadic fixed points, $\bar{\alpha}$ is an almost one-bump function. Then Theorem 4.1.2 tells us that $C_{\mathrm{PL}_2([r,r+1])}(\bar{\alpha})$ is infinite cyclic, generated by a minimal root $\bar{\omega}$ of $\bar{\alpha}$. Applying the inverse of the bar isomorphism, we see that $C_T(\alpha)$ is infinite cyclic, generated by a minimal root ω of α.

The final claim is a consequence of Equation (4.1) and the discussion preceding it. □

Lemma 4.1.7. *Let $\alpha \in T$ have rotation number $\mathrm{rot}(\alpha) = 0$. Then $C_T(\alpha) = C_T(\alpha^k)$, for any integer $k \neq 0$.*

Proof. The left-hand centraliser is contained in the right-hand centraliser automatically; we just need to show that if γ commutes with α^i then γ commutes with its root α. If $\alpha = \mathrm{id}$ this is certainly true. Otherwise, enumerate the $n > 0$ important points $I_\alpha = I(\alpha^k) = \{r_0 \to \cdots \to r_{n-1} \to r_0\}$. Because $\mathrm{rot}(\alpha) = 0$ we have $\mathrm{rot}(\alpha^k) = 0$, and so $r_i\alpha = r_i\alpha^k = r_i$ for each i. If $\gamma \in C_T(\alpha^k)$, we know that γ cyclically permutes $I(\alpha^k)$: say $r_i\gamma = r_{i+j}$ for some j.

Now $\gamma \in C_T(\alpha^k)$ means exactly that $(\alpha^k)^\gamma = \alpha^k$. We have the equality[21] $\gamma = \rho_{-r_{n-j}}\rho_{r_{n-j}}\gamma$. Substitute this into the previous equation, then conjugate both sides by ρ_{-r_0} to conclude that

$$(\alpha^k)^{\rho_{-r_{n-j}}\rho_{r_{n-j}}\gamma\rho_{-r_0}} = (\alpha^k)^{\rho_{-r_0}}\,.$$

[21] The rotations ρ_{-r_0} and $\rho_{r_{n-j}}$ need not be elements of T because r_0 (and hence r_{n-j}) need not be dyadic. So this is an equality in the larger group $\mathrm{PL}^+_{\mathbb{Q}}(S^1) > T$.

Use associativity (and the fact that conjugation commutes with powers) to rebracket this equation as

$$([\alpha^{\rho_{-r_{n-j}}}]^k)^{[\rho_{r_{n-j}}\gamma\rho_{-r_0}]} = [\alpha^{\rho_{-r_0}}]^k\,. \tag{4.2}$$

The computations

$$\begin{aligned}
0 &\xmapsto{\rho_{r_{n-j}}} r_{n-j} \xmapsto{\alpha} r_{n-j} \xmapsto{\rho_{-r_{n-j}}} 0\\
0 &\xmapsto{\rho_{r_{n-j}}} r_{n-j} \xmapsto{\gamma} r_0 \xmapsto{\rho_{-r_0}} 0\\
0 &\xmapsto{\rho_{r_0}} r_0 \xmapsto{\alpha} r_0 \xmapsto{\rho_{-r_0}} 0
\end{aligned}$$

show that each of the three products in square brackets belong to F.

We never use the fact that the maps ρ_* in the proof are rotations; we only require that they are elements of T which map 0, r_0 and r_{n-j} to the right places.

Now Corollary 3.2.2.1 tells us that kth roots are unique in F, so from Equation (4.2) we conclude

$$[\alpha^{\rho_{-r_{n-j}}}]^{[\rho_{r_{n-j}}\gamma\rho_{-r_0}]} = \alpha^{\rho_{-r_0}}\,.$$

This simplifies to $\alpha^{\gamma\rho_{-r_0}} = \alpha^{\rho_{-r_0}}$; a final conjugation on both sides by ρ_{+r_0} yields $\alpha^\gamma = \alpha$. □

To close this section, we consider more carefully which direct products may occur as centralisers in F. To the author's knowledge, this was first written down in print by Guba and Sapir.[22]

[22] V.S. Guba and M.V. Sapir 1999, Above Theorem 34.

Claim 4.1.8 (Which F-centralisers are possible?)**.** *Let f and z be integers. There exists an $\alpha \in F$ with centraliser $C_F(\alpha) \cong F^f \times \mathbb{Z}^z$ if and only if f and z are not both zero and satisfy $0 \le f \le z+1$.*

Proof. "Not both zero" is necessary because otherwise $C_F(\alpha) = \mathrm{id}$; this is impossible, since $C_T(\mathrm{id}) = F \neq \{\mathrm{id}\}$, and the centraliser of any other element $\alpha \neq \mathrm{id}$ contains α, so can't be trivial.

Why is the inequality necessary? Enumerate the dyadic important points as $I_\alpha = \partial\,\mathrm{Fix}(\alpha) = \{0 = r_0, r_1, \ldots, r_{n-1}, r_n = 1\}$, say. As described above, we find the direct factors of $C_F(\alpha)$ by restricting to the 'blocks' $D_i = [r_i, r_{i+1}]$, for $0 \le i < n$. We get a factor of F if $\alpha\big|_{D_i} = \mathrm{id}_{D_i}$, and a factor of $\mathbb{Z}$ otherwise. We cannot have two consecutive F-type blocks, i.e. we cannot have $\alpha\big|_{D_i} = \mathrm{id}_{D_i}$ and $\alpha\big|_{D_i} = \mathrm{id}_{D_i}$: then r_{i+1} would not belong to $\partial\,\mathrm{Fix}(\alpha)$ so would not be important. Thus our F-type blocks (f in total) must be separated by $\mathbb{Z}$-type blocks (z in total). We need at least $z \ge f-1$ blocks to achieve this separation.

$$\begin{aligned} &F\mathbb{Z}F\mathbb{Z}F\cdots F\mathbb{Z}F \\ &F\mathbb{Z}F\mathbb{Z}F\cdots F\mathbb{Z}F\mathbb{Z}\mathbb{Z}\mathbb{Z}\cdots\mathbb{Z} \\ &\mathbb{Z}\mathbb{Z}\mathbb{Z}\cdots\mathbb{Z}F\mathbb{Z}F\mathbb{Z}F\cdots F\mathbb{Z}F \\ &F\mathbb{Z}\mathbb{Z}\mathbb{Z}\mathbb{Z}F\cdots F\mathbb{Z}\mathbb{Z}F\mathbb{Z}F\mathbb{Z}\cdots\mathbb{Z}F \end{aligned}$$

Figure 4.1: Label D_i with the symbol F (resp. $\mathbb{Z}$) if $\alpha\big|_{D_i}$ is isomorphic to F (resp. $\mathbb{Z}$). If we concatenate the labels for $D_0, D_1, \ldots, D_{n-1}$, we get a string of length $n = f + z$ over $\{F, \mathbb{Z}\}$ in which there is no substring 'FF'.

Now we let $0 \le f \le z+1$ and construct an element whose centraliser has these parameters, as follows. Set $n = f + z$ and partition the unit interval into n standard dyadic intervals D_i. On $D_0, D_2, \ldots, D_{2(f-1)}$ define α to be the identity. On all other intervals D_i, define α to be an almost one-bump function[23] on $\mathrm{PL}_2(D_i)$. Then this map has the required centraliser. □

[23] For instance, take $\alpha\big|_{D_i}$ to be a scaled copy of x_0. We can do this explicitly by declaring $\alpha\big|_{D_i} = x_0^{\Sigma_i}$, where Σ_i is a PL_2 map $I \to D_i$.

4.2 Centralisers of nontorsion elements in T

A chapter of Matucci's thesis studies[24] the centralisers of elements in orientation-preserving circle homeomorphism groups. In particular, he describes the centraliser of an element $\alpha \in T$. His analysis is split into two cases: first when α has finite order; then when α has infinite order. We present Matucci's main result for the latter case[25] below.

[24] Matucci 2008, Sections 7.1–7.2.

[25] Matucci 2008, Theorem 7.2.5. Matucci's theorem applies to $\mathrm{PL}_+(S^1)$ and $T = \mathrm{PL}_2(S^1)$, but we only consider the latter.

Theorem 4.2.1. *Let $\alpha \in T$ have rotation number $\mathrm{rot}(\alpha) = p/q$ and suppose that $\alpha^q \neq \mathrm{id}$. Pick a representative arc $Q = [r, s]$ for the action of α on the circle, as in Definition 3.4.1. Then $K = C_T(\alpha)_0$ is indeed a subgroup of $C_T(\alpha)$, and there is a short exact sequence of the form*

$$1 \to \underbrace{C_T(\alpha)_0}_{K} \hookrightarrow C_T(\alpha) \twoheadrightarrow \underbrace{C_T(\alpha)\big|_{I_\alpha}}_{H} \to 1\,. \tag{4.3}$$

The quotient H is a finite cyclic group, whereas the kernel K is isomorphic to $C_{\mathrm{PL}_2(\mathbb{Q})}(\alpha^q\big|_{\mathbb{Q}})$. In particular, if ∂Q is nondyadic then $K \cong \mathbb{Z}$.

In some sense, the fine (resp. coarse) conjugators from Chapter 3 correspond to the kernel (resp. quotient) of this extension.

Sketch proof. If γ commutes with α, then γ conjugates α to itself. The set $P = \mathrm{Fix}(\alpha^q)$ consists of the points in finite orbits under α. Now P is mapped by γ to itself, as the calculation

$$\mathrm{Fix}(\alpha^q)\gamma = \mathrm{Fix}([\alpha^q]^\gamma) = \mathrm{Fix}(\alpha^q)$$

shows. Since boundaries commute with homomorphisms, we may take boundaries on both sides to learn that $[\partial\,\mathrm{Fix}(\alpha^q)]\,\gamma = \partial\,\mathrm{Fix}(\alpha^q)$. Thus we have $I_\alpha\gamma = I_\alpha$.

Recall from Lemma 2.3.8 that $I_\alpha = \partial P$.

All in all, we have an action $I_\alpha \circlearrowleft C_T(\alpha)$. The group $K = C_T(\alpha)_0$ is the kernel of this action, consisting of the centralising elements that fix I_α pointwise. The corresponding quotient $H = C_T(\alpha)/K$ must be a finite cyclic group, because the elements of T preserve the cyclic order on the finite set I_α. □

Recall our convention regarding representative arcs Q from Definition 3.4.1: if at all possible, we choose the endpoints ∂Q to be dyadic. So ∂Q consists of nondyadics if and only if there are no dyadic important points of α.

4.2.1 *Identifying the kernel*

While we know how to characterise the elements of the kernel group $K = C_T(\alpha)_0$, it would be nice to have a way of parameterising them. Theorem 4.2.1 does this by asserting that $K \cong C_{\mathrm{PL}_2(\mathbb{Q})}(\alpha^q\big|_{\mathbb{Q}})$. We will justify this assertion in more detail.

We know that ∂Q consists of important points for α, and so is fixed pointwise by K. Thus Q is fixed setwise by K. So the set $K\big|_{\mathbb{Q}} = \{\kappa\big|_{\mathbb{Q}} \mid \kappa \in K\}$ is a group. Let $\theta\colon K \twoheadrightarrow K\big|_{\mathbb{Q}}$ be the restriction map. This is a homomorphism, because $(\gamma\delta)\big|_{\mathbb{Q}} = \gamma\big|_{\mathbb{Q}}\,\delta\big|_{\mathbb{Q}\gamma} = \gamma\big|_{\mathbb{Q}}\,\delta\big|_{\mathbb{Q}}$.

We know that α^q has rotation number 0, so must belong to K. As each element $\delta \in K$ commutes with α, we have $\delta\alpha^q = \alpha^q\delta$ for all $\delta \in K$. Apply the map θ to obtain $\delta\big|_{\mathbb{Q}}\,\alpha^q\big|_{\mathbb{Q}} = \alpha^q\big|_{\mathbb{Q}}\,\delta\big|_{\mathbb{Q}}$, which shows that $K\big|_{\mathbb{Q}} \le C_{\mathrm{PL}_2(\mathbb{Q})}(\alpha^q\big|_{\mathbb{Q}})$. Now by Claim 3.4.2 (with $\beta = \alpha$) there is a unique extension of $\delta\big|_{\mathbb{Q}}$ to $C_{\mathrm{PL}_{\mathbb{Q}}(S^1)}(\alpha)$. By uniqueness, this extension must be $\delta \in K$. This shows that θ is injective, and hence is a isomorphism.

So far we have shown that $K \cong K\big|_{\mathbb{Q}} \le C_{\mathrm{PL}_2(\mathbb{Q})}(\alpha^q\big|_{\mathbb{Q}})$. Matucci's proof concludes by classifying the group $C_{\mathrm{PL}_2(\mathbb{Q})}(\alpha^q\big|_{\mathbb{Q}})$ as an F-type centraliser (see also Section 4.1).

Remark 4.2.2 (The inclusion is an equality). Implicit in Matucci's proof is the fact that the containment $\le$ above is actually an equality. This is not explicitly shown in his thesis; we give an explicit proof of this here.

Proof. A centralising element is exactly one which commutes α to itself. Thus we may apply Claim 3.4.2 to any map $\gamma_0 \in C_{\mathrm{PL}_2(\mathbb{Q})}(\alpha^q\big|_{\mathbb{Q}})$,

resulting in a unique extension $\gamma \in C_{\mathrm{PL}_Q(S^1)}(\alpha)$. Let $Q = [r, s]$ so that we can refer to the endpoints of Q. If we can establish that $\gamma \in T$, then $\gamma \in K$ because γ_0 fixes ∂Q. Then $\gamma_0 = \gamma\big|_Q \in K\big|_Q$, which will complete the proof.

When ∂Q is dyadic, Lemma 3.4.7 guarantees that $\gamma \in T$. So suppose otherwise that ∂Q is not dyadic. This time, Lemma 3.4.7 tells us that $\gamma \in T$ only in certain circumstances. For $q > 1$, $\gamma \in T$ if and only if Condition (3.9) is satisfied with $\beta = \alpha$ and with γ_0 in place of ϵ_0. This condition reads

$$s^- \gamma_0' = s(\alpha^{-k})' \cdot r^+ \gamma_0' \cdot r(\alpha^k)' .$$

Observe that

$$s(\alpha^{-k})' \cdot r(\alpha^k)' = s(\alpha^{-k})' \cdot s\alpha^{-k}(\alpha^k)' = s(\alpha^{-k}\alpha^k)' = 1 ,$$

and so $\gamma \in T$ if and only if

$$s^- \gamma_0' = r^+ \gamma_0' . \tag{4.4}$$

This is exactly Condition (3.8), which is equivalent to $\gamma \in T$ when $q = 1$. We see that the value of q plays no role: we just need to establish Condition (4.4).

Because ∂Q is nondyadic, our convention in Definition 3.4.1 means that the interior Q° contains no dyadic fixed points. Thus $\alpha^q\big|_Q$ is an almost one-bump function $Q \to Q$. According to Theorem 4.1.2, the centraliser $C_{\mathrm{PL}_2(Q)}(\alpha^q\big|_Q)$ is infinite cyclic, generated by a minimal root ω of $\alpha^q\big|_Q$. Say $\omega^u = \alpha^q\big|_Q$ with $u > 0$, so that $\gamma_0 = \omega^i$ for some integer i. Because r and $s = r\alpha^k$ belong to the same α-orbit and are fixed by α^q, their derivatives under α^q are equal (by Remark 2.3.18). Because $\omega \in \mathrm{PL}_2(Q)$ it must fix $\partial Q = \{ r, s \}$ pointwise. These two facts justify the equalities

$$\begin{aligned} [r^+\omega']^u &= r^+(\omega^u)' = r^+(\alpha^q\big|_Q)' = s^-(\alpha^q\big|_Q)' \\ &= s^-(\omega^u)' = [s^-\omega']^u . \end{aligned}$$

We conclude that $r^+\omega' = s^-\omega'$, because uth roots are unique in the multiplicative group $G = 2^{\mathbb{Z}}$. Hence

$$\begin{aligned} s^-\gamma_0' &= s^-(\omega^i)' = [s^-\omega']^i = [r^+\omega']^i \\ &= r^+(\omega^i)' = r^+\gamma_0' . \end{aligned}$$

This establishes Condition (4.4), which completes the proof. □

4.2.2 *Identifying the quotient*

The quotient of Sequence (4.3) is the group of permutations $H = C_T(\alpha)\big|_{I_\alpha}$ induced on I_α by the centraliser. Since the action preserves the cyclic order on I_α, the rightmost factor is isomorphic to a cyclic

group, of order h say. If $\mathrm{rot}(\alpha) = p/q$ then $h \geq q$, since the important points have size q orbits under α. To determine the actual value of h, we turn to brute force.

Algorithm 4.2.3. Let $\alpha \in T$ with rotation number $\mathrm{rot}(\alpha) = p/q$. The order h of the quotient in Sequence (4.3) can be determined algorithmically as follows.

1. Enumerate the important points of α as $I_\alpha = \{x_1 \to \ldots \to x_n \to x_1\}$ in cyclic order.
2. For each proper divisor j of n/q, in ascending order:
   (a) Use Remark 3.3.6 to decide if there exists a conjugator γ_j between α and itself with index difference j; that is, with $x_i\gamma_j = x_{i+j}$ for each i.
   (b) If so, break the loop and retain the current value of j.
3. If we did not break the previous loop, return $h = q$.
4. Otherwise return $h = n/j$ (using the value of j from the break).

Proof of correctness. The algorithm always terminates and returns a value—we need to show it returns the correct value. Let G be the group of *all* cyclic permutations of I_α, i.e. the group of functions $\{\gamma_j : I_\alpha \to I_\alpha\}_{j=0}^{n-1}$ where $x_i\gamma_j = x_{i+j}$. The group operation is $\gamma_j\gamma_k = \gamma_{j+k}$ with subscripts modulo n, so G is cyclic of order n. The subgroups of cyclic groups $H \leq G$ are determined[26] by their order $h = |H|$, which must divide $|G|$. Explicitly, if $d = n/h$ then H is the subgroup $H = \langle\gamma_d\rangle$.

[26] Rotman 1995, Lemma 2.15.

Now let H be the right-hand side of Sequence (4.3). Since we have the inclusion $\langle\alpha\big|_{I_\alpha}\rangle \leq H \leq G$, it follows that $q \mid h \mid n$ by taking group orders. Then $d = (n/h) \mid (n/q)$. One possibility is that $n/h = n/q$, in which case $h = q$. Then the loop in step 2 never breaks, and we correctly return $h = q$ in step 3. Otherwise the loop breaks at $j < n$, in which case $\langle\gamma_j\rangle \leq H$. Now j is minimal with this property, because the loop in step 2 is evaluated in ascending order. This means that no element $\gamma_{j'}$ with $1 \leq j' < j$ belongs to H; hence $H = \langle\gamma_j\rangle$ has order $h = n/j$, which we return as our output. This justifies step 4. □

As noted in the proof above, the group orders q, h and n divide each other (from left to right), so certainly $q \leq h \leq n$. Each inequality may be strict or not; there are then four possible cases, each of which does indeed occur—see Figure 4.2 for references to examples. The first inequality is strict if and only if then $H = C_T(\alpha)$ has a strictly finer action[27] than $\langle\alpha\rangle$ on I_α. The second inequality is strict if and only if H acts intransitively on I_α.

Example	$h < n$	$h = n$
$q < h$	5.1.7	5.1.2
$q = h$	5.1.8	5.1.3

Figure 4.2: These examples exhibit all four cases of the inequality $q \leq h \leq n$. This shows that the bounds on h cannot be strengthened.

[27] i.e. each α-orbit is strictly contained in an H-orbit.

Remark 4.2.4. Let $G_2 \leq G$ be the group of cyclic permutations of I_α which map dyadics to dyadics and nondyadics to nondyadics. Because T preserves the dyadics, it must be the case that $H \leq G_2$. Thus we can refine the inequality $q \leq h \leq n$ to $q \leq h|G_2| \leq n$.

Algorithm 4.2.3 does more than identify the order $h = |H|$: it also produces an element $\gamma_j \in C_T(\alpha)$ whose restriction to I_α, η say, generates the quotient H. (If we don't break the loop in step 2 then $h = n/q$, so $j = n/q$ is the minimal index difference. We may take $\gamma_j = \alpha$ as an element with this difference.) We may then take γ_j^i to be a preimage of $\eta^i \in H$, for each power $0 \leq i < h$. This collection of preimages is a *transversal* for K.

If we understand the kernel K, quotient H, and how a chosen transversal conjugates K, then we gain an understanding of the extension (in this case $C_T(\alpha)$). For instance, if we have presentations for K and H and understand the transversal, we can construct[28] a presentation for $C_T(\alpha)$. So far, we've concentrated on understanding the first two ingredients K and H. In the rest of this thesis, we'll seek to understand the final piece of the puzzle: how exactly does our transervsal for H conjugate K?

[28] Johnson 1997, Section 10.2.

5 *Nontorsion centraliser structure*

Fix a nontorsion element $\alpha \in T$. In this chapter, we will explain how to determine the structure of Sequence (4.3). We do so by classifying such elements α into one of four categories, and from these we classify $C_T(\alpha)$ into one of five types.

5.1 Examples

To begin, we present a series of examples to illustrate Theorem 4.2.1 and explore the questions pose at the end of Section 4.2.2. Each example is constructed by gluing together shrunk copies of elements of F. In the interests of brevity, we introduce some notation to facilitate this gluing.

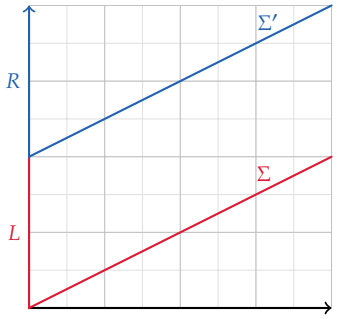

Figure 5.1: We introduce two PL_2 maps Σ and Σ' for discussing centraliser examples.

Notation 5.1.1. Let I, L and R be the real intervals $[0,1]$, $[0,1/2]$ and $[1/2,1]$ respectively. Define $\Sigma\colon I \to L$ to be the PL_2 map $t \mapsto t/2$. Similarly let $\Sigma'\colon I \to R$ be the map $t \mapsto t/2 + 1/2$. 'Conjugating' by these maps results in isomorphisms $\cdot^{\Sigma}\colon F = \mathrm{PL}_2(I) \to \mathrm{PL}_2(L)$ and $\cdot^{\Sigma'}\colon F \to \mathrm{PL}_2(R)$ (c.f. Claim 1.2.5).

When thought of as a permutation of I, the dyadic rotation $\rho_{1/2}$ can be defined in terms of this notation as

$$t\rho_{1/2} = \begin{cases} t + 1/2 & \text{if } t \in L \\ t - 1/2 & \text{if } t \in R. \end{cases}$$

Notice then that $\Sigma' = \Sigma\,\rho_{1/2}\big|_L$ and $\Sigma = \Sigma'\,\rho_{1/2}\big|_R$.

Example 5.1.2. The map α illustrated in Figure 5.2 has been constructed by gluing two copies of x_0 together. Specifically, α is the result of the recipe

$$t\alpha = \begin{cases} tx_0^{\Sigma} & \text{if } t \in L \\ tx_0^{\Sigma'} & \text{if } t \in R. \end{cases}$$

The important points $I_\alpha = \{0, 1/2\}$ are fixed pointwise by α, so the rotation number is $\mathrm{rot}(\alpha) = 0$ with denominator $q = 1$. Thus our representative arc for α is $Q = I$. If $H \cong \mathbb{Z}_h$ denotes the quotient of Sequence (4.3) as above, then α projects onto the identity of H.

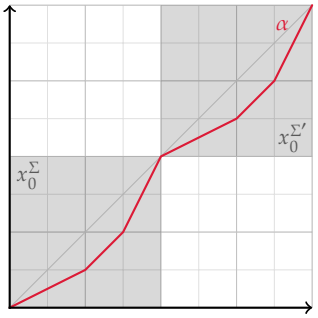

Figure 5.2: The element α from Figure 2.5 consists of two copies of the generator $x_0 \in F$, rescaled to permute the standard dyadic intervals $[0, 1/2]$ and $[1/2, 1]$.

By design, α commutes with the dyadic rotation $\rho_{1/2} \in T$. To be explicit, the two products in question are given by the casewise formulas

$$t\alpha\rho_{1/2} = \begin{cases} tx_0^{\Sigma}\rho_{1/2} & \text{if } t \in L \\ tx_0^{\Sigma'}\rho_{1/2} & \text{if } t \in R \end{cases} \tag{5.1}$$

and

$$t\rho_{1/2}\alpha = \begin{cases} t\rho_{1/2}x_0^{\Sigma'} & \text{if } t \in L \\ t\rho_{1/2}x_0^{\Sigma} & \text{if } t \in R. \end{cases} \tag{5.2}$$

If $t \in L$, then

$$\begin{aligned} tx_0^{\Sigma}\rho_{1/2} &= t\Sigma^{-1}x_0\Sigma\rho_{1/2} \\ &= t(\Sigma'\rho_{1/2})^{-1}x_0\Sigma' \\ &= t\rho_{1/2}\Sigma'^{-1}x_0\Sigma' \\ &= t\rho_{1/2}x_0^{\Sigma'}, \end{aligned} \tag{5.3}$$

showing that $t\alpha\rho_{1/2} = t\rho_{1/2}\alpha$ in this case. A similar calculation holds for $t \in R$, and so $\rho_{1/2}$ commutes with α.

Our centralising element $\rho_{1/2}$ pushes the important points d steps forward, where $d = 1$. So $C_T(\alpha)$ transitively permutes I_α, and hence the quotient group H must be of size $h = n/d = 2$. We see that the inequalities $q \le h \le n$ are realised as $1 < 2 = 2$ in this scenario.

As noted above, $\rho_{1/2}$ gives us a transversal[1] σ of the kernel K in Sequence (4.3). If γ_i is the map which cyclically pushes I_α forward by i steps, then our transversal sends $\gamma_0 \mapsto^\sigma \mathrm{id}$ and sends $\gamma_1 \mapsto^\sigma \rho_{1/2}$. This is not just a function between sets: the relation $\gamma_1^2 = \gamma_0$ is preserved by σ, because $\rho_{1/2}^2 = \mathrm{id}$. Since σ is a homomorphism (a splitting), the centraliser splits[2] as a semidirect product $C_T(\alpha) \cong K \rtimes \mathbb{Z}_2$.

To understand the group structure of this semidirect product, we need to identify K and describe how conjugating by $\rho_{1/2}$ affects K. Matucci's analysis[3] of centralisers in F shows that the kernel K is isomorphic to

$$\begin{aligned} K \cong C_{\mathrm{PL}_2(Q)}(\alpha^1\big|_Q) &= C_F(\alpha) \\ &\cong C_{\mathrm{PL}_2(L)}(\alpha\big|_L) \times C_{\mathrm{PL}_2(R)}(\alpha\big|_R) \\ &= \langle \alpha\big|_L \rangle \times \langle \alpha\big|_R \rangle \\ &\cong \mathbb{Z} \times \mathbb{Z}. \end{aligned}$$

Since $\alpha\big|_L^{\rho_{1/2}} = \alpha\big|_R$, these two direct factors are swapped by $\rho_{1/2}$. This shows that $C_T(\alpha)$ is isomorphic to the wreath product $\mathbb{Z} \wr \mathbb{Z}_2$ arising from the action $\{L, R\} \circlearrowleft \mathbb{Z}_2$.

For our next example, we'll consider a related element whose rotation number is nonzero. The idea is to use the same building blocks (two shrunk copies of x_0), but this time we'll build an element which exchanges the two copies.

Example 5.1.3. With α as in Example 5.1.2, let $\beta = \alpha\rho_{1/2}$ be the element shown in Figure 5.3. The new element β is constructed much like α: we use two copies of x_0, according to the recipes given in Equations (5.1) and (5.2). The important point set is $I_\beta = I_\alpha = \{0, 1/2\}$; since $0\beta = 1/2$ we see that the rotation number is $\mathrm{rot}(\beta) = 1/2$, with denominator $q = 2$. As β transitively permutes I_β, we have $d = 1$ and $h = n/d = 2$. In this example, our inequality $q \le h \le n$ is realised as $2 = 2 = 2$.

[1] By 'a transversal σ' we really mean a transversal function $\sigma\colon H \to C_T(\alpha)$, as described in Robinson 1996, Section 11, p. 311 in particular.

[2] Rotman 1995, Lemma 7.20.

[3] Kassabov and Matucci 2012, Section 5. See also our Section 4.1.

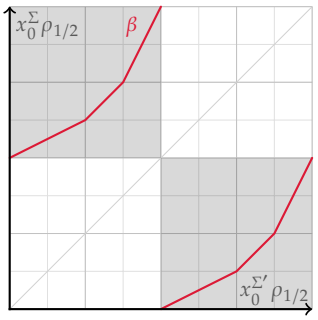

Figure 5.3: The product $\beta = \alpha\rho_{1/2}$ $\in T \setminus F$.

We take $Q = L$ as our representative arc for β. We note that β is the product of two elements which commute with $\rho_{1/2}$, so β itself must commute with this rotation. Thus we can use the transversal σ to split the extension $C_T(\beta)$, exactly as in Example 5.1.2.

The kernel of the extension is $C_T(\beta)_0 \cong C_{\mathrm{PL}_2(L)}(\beta^2|_L)$. We can directly compute

$$t\beta^2 = \begin{cases} t(x_0^\Sigma\, \rho_{1/2})\,(\rho_{1/2} x_0^\Sigma) \\ t(x_0^{\Sigma'} \rho_{1/2})\,(\rho_{1/2} x_0^{\Sigma'}) \end{cases} = \begin{cases} t(x_0^2)^\Sigma & \text{if } t \in L \\ t(x_0^2)^{\Sigma'} & \text{if } t \in R, \end{cases}$$

which shows in particular that $\beta^2|_L$ is a shrunken copy of x_0^2. Thus $C_{\mathrm{PL}_2(L)}(\beta^2|_L) \cong C_F(x_0^2) = \langle x_0 \rangle$, again by Matucci's study of centralisers in F-like groups. The pseudo-conjugate $x_0^\Sigma = \alpha|_L$ is a generator for $C_{\mathrm{PL}_2(L)}(\beta^2|_L)$. Thus the extension's kernel is generated by the unique extension of this element to $C_T(\beta)$, namely α.

In summary, the centraliser $C_T(\beta)$ splits as a semidirect product $C_T(\beta)_0 \rtimes \mathbb{Z}_2 = \langle \alpha \rangle \rtimes \langle \rho_{1/2} \rangle$. Since α commutes with $\rho_{1/2}$, this must in fact be a direct product $C_T(\beta) = \langle \alpha \rangle \times \langle \rho_{1/2} \rangle \cong \mathbb{Z} \times \mathbb{Z}_2$.

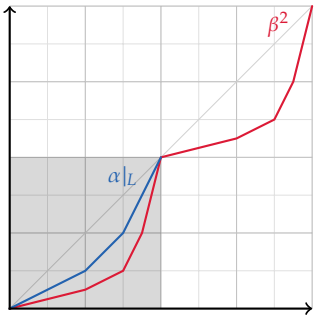

Figure 5.4: The square of β from Figure 5.3. The shaded region is the set L^2 containing the graph of $\beta^2|_L$. Also shown is the graph of $\alpha|_L$, which generates $C_{\mathrm{PL}_2(L)}(\beta^2|_L)$.

The next two examples are similar in spirit to those above. We slightly complicate things by ensuring that our elements have different (but conjugate) graphs over L and R. Much of our discussion goes through unchanged for our new version of α, but the story is not so simple for the analogue of β.

Example 5.1.4. Let $\gamma, \phi \in F$ be the maps shown in Figure 5.5. In particular, γ is given by the recipe

$$t\gamma = \begin{cases} t x_0^\Sigma & \text{if } t \in L \\ t(x_0^\phi)^{\Sigma'} & \text{if } t \in R. \end{cases}$$

This element has similar dynamical behaviour to that of α in Example 5.1.2: both consist of one-bump functions below the diagonal on L^2 and R^2. Consequently, the rotation number is again $\mathrm{rot}(\gamma) = 0$, with reduced denominator $q = 1$. Our representative arc for γ is $Q = I$; we only have one block. The important point set is $\{0, 1/2\}$, fixed pointwise by γ. What is the value of h? By Algorithm 4.2.3, we must determine if there exists an element $\omega \in T$ conjugating γ to itself with $0\omega = 1/2$.

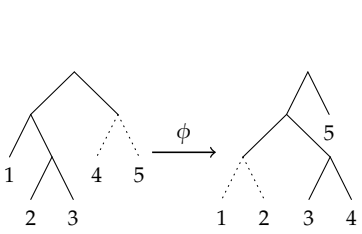

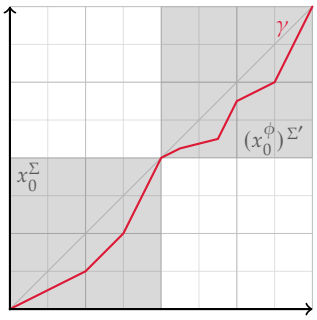

Figure 5.5: We construct an element $\gamma \in F$. In the lower-left quadrant L^2, γ is a scaled copy of x_0. In contrast, the graph of γ in the upper-right quadrant R^2 is a scaled copy of x_0^ϕ, where $\phi \in F$ is the map with the given tree pair diagram.

From thin air, we conjure an element ω defined by

$$t\omega = \begin{cases} t\phi^{\Sigma}\rho_{1/2} & \text{if } t \in L \\ t(\phi^{-1})^{\Sigma'}\rho_{1/2} & \text{if } t \in R \end{cases}$$

(see Figure 5.6). We see that $0\omega = 1/2$ as required, but is $\gamma^{\omega} = \gamma$? The relevant computations are

$$t\omega\gamma = \begin{cases} t(\phi^{\Sigma}\rho_{1/2})(x_0^{\phi\Sigma'}) & \text{if } t \in L \\ t([\phi^{-1}]^{\Sigma'}\rho_{1/2})(x_0^{\Sigma}) & \text{if } t \in R \end{cases}$$

and

$$t\gamma\omega = \begin{cases} t(x_0^{\Sigma})(\phi^{\Sigma}\rho_{1/2}) & \text{if } t \in L \\ t(x_0^{\phi\Sigma'})([\phi^{-1}]^{\Sigma'}\rho_{1/2}) & \text{if } t \in R. \end{cases}$$

The expressions for $t \in L$ are equal, because

$$t\omega\gamma = t\Sigma^{-1}\left(\phi[\Sigma\rho_{1/2}(\Sigma')^{-1}]\phi^{-1}\right)x_0\phi\Sigma' = t\Sigma^{-1}x_0\phi\Sigma\rho_{1/2} = t\gamma\omega\,,$$

where the bracketed factor is equal to the identity. A similar (but increasingly messy) calculation holds over $t \in R$. Thus $h = 2$, and the inequality $q \leq h \leq n$ is realised as $1 < 2 \leq 2$, just as in Example 5.1.2.

Does ω have order 2? Its square is given by

$$t\omega^2 = \begin{cases} t(\phi^{\Sigma}\rho_{1/2})([\phi^{-1}]^{\Sigma'}\rho_{1/2}) & \text{if } t \in L \\ t([\phi^{-1}]^{\Sigma'}\rho_{1/2})(\phi^{\Sigma}\rho_{1/2}) & \text{if } t \in R. \end{cases}$$

If $t \in L$, then

$$t\omega^2 = t\Sigma^{-1}\left(\phi\Sigma\rho_{1/2}(\Sigma')^{-1}\phi^{-1}\right)\Sigma'\rho_{1/2} = t\Sigma^{-1}\Sigma = t\,,$$

and once again a similar calculation holds for $t \in R$. We see that $\omega^2 = \mathrm{id}$, which means that $C_T(\gamma) \cong C_F(\gamma) \rtimes \langle\omega\rangle \cong \mathbb{Z}^2 \rtimes \mathbb{Z}_2 \cong \mathbb{Z} \wr \mathbb{Z}_2$ as in Example 5.1.2.

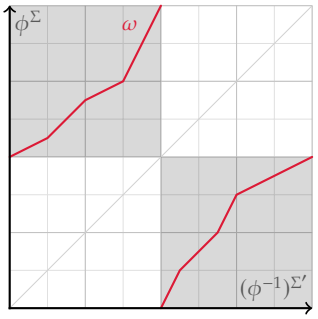

Figure 5.6: The centralising element ω.

The wreath corresponds to the action $\{L, R\} \circlearrowleft \langle\omega\rangle$.

Remark 5.1.5. Let $\chi \in T$ be the map defined by $\chi\big|_L = \mathrm{id}$ and $\chi\big|_R = \phi^{\Sigma'}$. A direct computation verifies that $\alpha^{\chi} = \gamma$. So with the benefit of hindsight, it's no coincidence that γ and α have isomorphic centralisers.

In our next example, we define $\delta = \gamma\rho_{1/2}$ by analogy with $\beta = \alpha\rho_{1/2}$ in Example 5.1.3. These two elements' centralisers are not as similar as we might expect.

Example 5.1.6. Let $\delta = \gamma\rho_{1/2}$ be the map illustrated in Figure 5.7. To be explicit, δ is defined casewise by

$$t\delta = \begin{cases} tx_0^{\Sigma}\rho_{1/2} & \text{if } t \in L \\ t(x_0^{\phi})^{\Sigma'}\rho_{1/2} & \text{if } t \in R. \end{cases}$$

Then $\mathrm{rot}(\delta) = 1/2$ with denominator $q = 2$; the important points $I_\delta = \{0, 1/2\}$ are permuted transitively by δ, so $d = 1$ and $h = n/d = 2$;

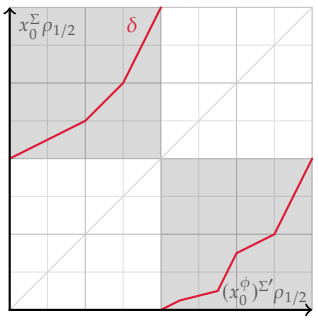

Figure 5.7: The product $\delta = \gamma\rho_{1/2}$.

and we may again take $Q = L$ as a representative arc for δ. The inequality $q \leq h \leq n$ is realised as $2 = 2 = 2$.

We'll start by identifying the kernel $K = C_T(\delta)_0$ of the extension, which we know to be isomorphic to $C_{\mathrm{PL}_2(L)}(\delta^2|_L)$. The square is given by

$$t\delta^2 = \begin{cases} t(x_0^\Sigma \rho_{1/2})\,([x_0^\phi]^{\Sigma'} \rho_{1/2}) \\ t([x_0^\phi]^{\Sigma'} \rho_{1/2})\,(x_0^\Sigma \rho_{1/2}) \end{cases} = \begin{cases} t(x_0 x_0^\phi)^\Sigma & \text{if } t \in L \\ t(x_0^\phi x_0)^{\Sigma'} & \text{if } t \in R, \end{cases}$$

so the kernel is isomorphic to $C_{\mathrm{PL}_2(L)}([x_0 x_0^\phi]^\Sigma) \cong C_F(x_0 x_0^\phi)$. As $X = x_0 x_0^\phi$ is a one-bump function, its centraliser in F is generated by a minimal root of X. The initial gradient of X is $1/4$, and it can be checked[4] that no element $w \in C_F(X)$ exists with initial gradient $1/2$. (Again we are using Kassabov and Matucci's work[5] on F-like centralisers here.) Thus $C_F(X) = \langle X \rangle$ is infinite cyclic, and so $C_{\mathrm{PL}_2(L)}(\delta^2|_L)$ is generated by $X^\Sigma = \delta^2|_L$. The unique extension of $\delta^2|_L$ to $C_T(\delta)$ is δ^2, so $C_T(\delta)_0 = \langle \delta^2 \rangle \cong \mathbb{Z}$.

Because $h = 2$, we know that $K = \langle \delta^2 \rangle$ has index 2 in $C_T(\delta)$. The missing coset must be $\langle \delta^2 \rangle \delta$, meaning that $C_T(\delta) = \langle \delta \rangle \cong \mathbb{Z}$. Notice that the centraliser extension does not split in this case,[6] as there is no element of order 2 in $\langle \delta \rangle$.

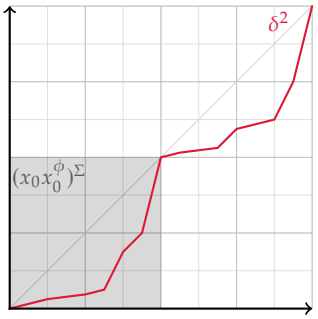

Figure 5.8: The square of δ. The shaded area is L^2, containing the graph of $\delta^2|_L$. This restriction is equal to a scaled copy of $x_0 x_0^\phi$. Within R^2, $\delta^2|_R$ is a scaled copy of a different map $x_0^\phi x_0$.

[4] See Appendix A.

[5] Kassabov and Matucci 2012, Theorem 5.6.ii.

[6] See also Appendix B for another argument to this end.

The examples so far have all had $q = h$, because we constructed elements α (say) whose important points were transitively permuted by α. Another pair of examples shows that this is not necessary. The key is to replace the use of x_0 in previous examples with x_1.

Example 5.1.7. In this example, we construct an element $\epsilon \in T$ with parameters $q < h < n$. Let $\epsilon \in T$ be the map defined by

$$t\epsilon = \begin{cases} t x_1^\Sigma & \text{if } t \in L \\ t x_1^{\Sigma'} & \text{if } t \in R. \end{cases}$$

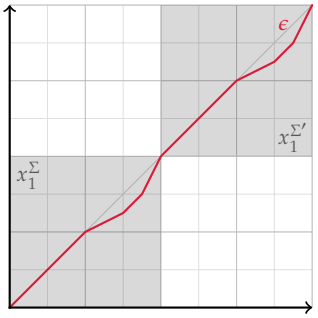

Figure 5.9: The element ϵ is constructed in the same manner as α from Example 5.1.2, using x_1 in place of x_0.

Since $\epsilon \in F$, its rotation number is $\mathrm{rot}(\epsilon) = 0$, with denominator $q = 1$. Its important points are $I_\epsilon = \{0, 1/4, 1/2, 3/4\}$, each of which is fixed by ϵ. Our representative arc is $Q = I$.

To identify the quotient of the extension, we first note that $\rho_{1/2} \in C_T(\epsilon)$. The calculation to justify this is the same as that given in Equations (5.1) to (5.3), after the symbol x_0 is replaced with x_1 throughout. Thus we know that $q \geq 2$. In fact $q = 2$, because no element of the centraliser can map 0 to $1/4$ or $3/4$. This is because[7] the gradients of $\epsilon^q = \epsilon$ differ: $0^+\epsilon' = 1$, whereas $1/4^+\epsilon' = 3/4^+\epsilon' = 1/2$. Thus the inequality $q \leq h \leq n$ is realised as $1 < 2 < 4$. Once again, we see that the extension splits.

[7] Or: centralisers send fixed intervals to fixed intervals—see Condition (3.5).

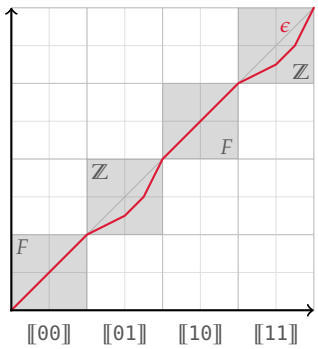

Figure 5.10: The F-centraliser of ϵ splits as a direct product over four factors. The first, second, third and fourth factors' elements rearrange the intervals [[00]], [[01]], [[10]] and [[11]], respectively.

Arguing as in Examples 5.1.2 and 5.1.4, we see that the kernel of the extension is $C_F(\epsilon)$. This in turn is the product

$$\begin{aligned} C_F(\epsilon) &\cong C_{\mathrm{PL}_2(L)}(\epsilon|_L) \times C_{\mathrm{PL}_2(R)}(\epsilon|_R) \\ &= C_{\mathrm{PL}_2(L)}(x_1^\Sigma) \times C_{\mathrm{PL}_2(R)}(x_1^{\Sigma'}) \end{aligned}$$

$$\begin{aligned}
&\cong \left(C_{\mathrm{PL}_2(\llbracket \mathtt{00} \rrbracket)}(\mathrm{id}) \times C_{\mathrm{PL}_2(\llbracket \mathtt{01} \rrbracket)}(x_0^{\Sigma_{\mathtt{01}}})\right) \\
&\qquad \times \left(C_{\mathrm{PL}_2(\llbracket \mathtt{10} \rrbracket)}(\mathrm{id}) \times C_{\mathrm{PL}_2(\llbracket \mathtt{11} \rrbracket)}(x_0^{\Sigma_{\mathtt{11}}})\right) \\
&= \left(F^{\Sigma_{\mathtt{00}}} \times \langle x_0^{\Sigma_{\mathtt{01}}} \rangle\right) \times \left(F^{\Sigma_{\mathtt{10}}} \times \langle x_0^{\Sigma_{\mathtt{10}}} \rangle\right), \qquad (5.4)
\end{aligned}$$

where Σ_s is the PL_2 map $I \to [\![s]\!]$ with a single linear segment, for any string s over $\{\mathtt{0}, \mathtt{1}\}$. (For instance, $\Sigma_{\mathtt{0}} = \Sigma$ and $\Sigma_{\mathtt{1}} = \Sigma'$.) To reconstruct the extension, we note that our chosen preimage $\rho_{1/2}$ for the generator of $H \cong \mathbb{Z}_2$ swaps the bracketed factors in Expression (5.4). Thus the entire centraliser is the semidirect product $C_T(\epsilon) \cong (F \times \mathbb{Z})^2 \rtimes \mathbb{Z}_2 \cong (F \times \mathbb{Z}) \wr \mathbb{Z}_2$.

By analogy with previous examples, we again study the product of our previous element with $\rho_{1/2}$. Much is similar to the case of β in Example 5.1.3; however we have to take a little more care this time, because $q < h$.

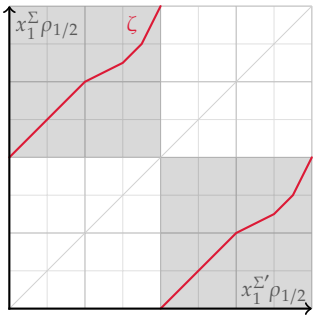

Figure 5.11: The product $\zeta = \epsilon\rho_{1/2} \in T \setminus F$.

Example 5.1.8. Let $\zeta = \epsilon\rho_{1/2}$, as shown in Figure 5.11. The rotation number is $\mathrm{rot}(\zeta) = 1/2$ with denominator $q = 2$. The important point set is $I_\zeta = \{0, 1/4, 1/2, 3/4\}$, of size $n = 4$. Each important point is shifted two steps forward by ζ.

The rotation $\rho_{1/2}$ commutes with ζ, because ζ is a product of two elements which also commute with $\rho_{1/2}$. This shows $h \geq 2$. Once again, no element of the centraliser can map 0 to $1/4$ or $3/4$ (just as in Example 5.1.7). This time the relevant gradients are $0^+(\zeta^2)' = 1$ versus $1/4^+(\zeta^2)' = 3/4^+(\zeta^2)' = 1/4$. This shows that $h = 2$, and so we can use $\rho_{1/2}$ to split the extension. We note that the inequality $q \leq h \leq n$ is realised as $2 = 2 < 4$.

The extension's kernel is isomorphic to $C_{\mathrm{PL}_2(L)}(\zeta^2\big|_L)$, which is

$$\begin{aligned}
C_{\mathrm{PL}_2(L)}(\zeta^2\big|_L) &= C_{\mathrm{PL}_2(L)}([x_1^2]^{\Sigma}) \\
&\cong C_{\mathrm{PL}_2(\llbracket \mathtt{00} \rrbracket)}(\mathrm{id}) \times C_{\mathrm{PL}_2(\llbracket \mathtt{01} \rrbracket)}([x_0^2]^{\Sigma_{\mathtt{01}}}) \\
&= F^{\Sigma_{\mathtt{00}}} \times \langle x_0^{\Sigma_{\mathtt{01}}} \rangle \\
&\cong F \times \mathbb{Z}.
\end{aligned}$$

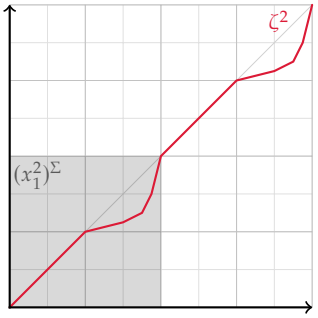

Figure 5.12: The square of ζ. The shaded region is the set L^2, in which ζ^2 is a scaled copy of x_1^2.

Note that this is the left-hand bracketed factor in Expression (5.4). To finish, we need to know how conjugation by $\rho_{1/2}$ affects the kernel K, so we need to work with elements of $K \leq C_T(\zeta)$ instead of an isomorphic copy of K.

Any element $\omega \in C_{\mathrm{PL}_2(L)}(\zeta^2\big|_L)$ lifts to a unique element $\bar{\omega} \in C_T(\zeta)_0$ given by

$$t\bar{\omega} = \begin{cases} t\omega & \text{if } t \in L \\ t\,\zeta^{-1}\big|_R\,\omega\,\zeta\big|_L & \text{if } t \in R. \end{cases}$$

The two products involving $\bar{\omega}$ and $\rho_{1/2}$ are

$$t\bar{\omega}\rho_{1/2} = \begin{cases} t\omega\rho_{1/2} & \text{if } t \in L \\ t\,\zeta^{-1}\big|_R\,\omega\,\zeta\big|_L\,\rho_{1/2} & \text{if } t \in R \end{cases}$$

Example	Name	Centraliser	$q \le h \le n$
5.1.2	α	$\mathbb{Z} \wr \mathbb{Z}_2$	$1 < 2 = 2$
5.1.3	β	$\mathbb{Z} \times \mathbb{Z}_2$	$2 = 2 = 2$
5.1.4	γ	$\mathbb{Z} \wr \mathbb{Z}_2$	$1 < 2 = 2$
5.1.6	δ	$\mathbb{Z}$	$2 = 2 = 2$
5.1.7	ϵ	$(F \times \mathbb{Z}) \wr \mathbb{Z}_2$	$1 < 2 < 4$
5.1.8	ζ	$(F \times \mathbb{Z}) \times \mathbb{Z}_2$	$2 = 2 < 4$

Figure 5.13: A summary of the centralisers we've been discussing. This isn't a comprehensive survey of nontorsion elements of T: for instance, we have not considered an element with nondyadic fixed points.

and

$$t\rho_{1/2}\bar{\omega} = \begin{cases} t\rho_{1/2}\,\zeta^{-1}\big|_R\,\omega\,\zeta\big|_L & \text{if } t \in L \\ t\rho_{1/2}\omega & \text{if } t \in R. \end{cases}$$

For the moment, we suppose that $t \in L$. Then $\bar{\omega}\rho_{1/2} = \rho_{1/2}\bar{\omega}$ if and only if

$$t\omega\rho_{1/2} = t\rho_{1/2}\,\zeta^{-1}\big|_R\,\omega\,\zeta\big|_L\ .$$

With some rearrangement, we see that

$$t\omega = t\rho_{1/2}\,\zeta^{-1}\big|_R\,\omega\,\zeta\big|_L\,\rho_{1/2} = t\omega^{(\zeta|_L\rho_{1/2})} = t\omega^{(x_1^\Sigma)}\,.$$

Thus $\rho_{1/2}$ commutes with $\bar{\omega}$ over L if and only if ω commutes with x_1^Σ. Put differently, we need $\omega \in C_{\mathrm{PL}_2(L)}(\zeta^2\big|_L) = C_{\mathrm{PL}_2(L)}([x_1^2]^\Sigma)$ to be contained in $C_{\mathrm{PL}_2(L)}(x_1^\Sigma)$. This is true if and only if $C_F(x_1^2) \le C_F(x_1)$. Matucci's analysis of centralisers in F shows that these two groups are in fact equal to $\langle x_1 \rangle$ (see also Theorem 4.1.3).

A similar argument applies when we suppose that $t \in R$, so we conclude that $\bar{\omega}$ commutes with $\rho_{1/2}$. Then the centraliser extension is a direct product $C_T(\zeta) = C_T(\zeta)_0 \times \langle \rho_{1/2} \rangle \cong (F \times \mathbb{Z}) \times \mathbb{Z}_2$.

This concludes our series of examples, which are summarised in Figure 5.13. Five of our examples split: two of these are direct products and the other three as wreath products. The element δ was the odd one out, since for some reason its centraliser extension is nonsplit. In the remainder of this chapter, we generalise this approach to study a generic nontorsion element of T. We'll see that the wreath and direct products above are part of a general pattern; why δ was different to the other examples; and we'll see what happens when there are no dyadic important points to work with.

Remark 5.1.9. Up to equivalence, there are three extensions of $\mathbb{Z}_2$ by $\mathbb{Z}$, namely

$$\begin{aligned} \mathbb{Z} &\hookrightarrow \mathbb{Z} \times \mathbb{Z}_2 \twoheadrightarrow \mathbb{Z}_2 && \text{(direct product)} \\ \mathbb{Z} &\hookrightarrow \mathbb{Z} \rtimes \mathbb{Z}_2 \twoheadrightarrow \mathbb{Z}_2 && \text{(infinite dihedral)} \\ 2\mathbb{Z} &\hookrightarrow \mathbb{Z} \twoheadrightarrow \mathbb{Z}_2 && \text{(division by 2)} \end{aligned}$$

We have seen that $C_T(\beta)$ and $C_T(\delta)$ are isomorphic to the first and last extensions, respectively. Does there a nontorsion element $\eta \in T$ whose

rem. points	rotation number: zero	rotation number: nonzero
some dyadic	dyadic self-similar DSS Section 5.2.1	dyadic offset-similar DOS Section 5.2.2
all nondyadic	nondyadic self-similar NSS Section 5.2.3	nondyadic offset-similar NOS Section 5.2.4

Figure 5.14: Names and abbreviations for our four types of nontorsion element, along with the section each is discussed in.

centraliser is infinite dihedral? The answer is 'no', because the centre of $C_T(\eta)$ contains[8] $\langle \eta \rangle$, which is nontrivial; however the infinite dihedral group is centreless.

[8] Notice that we don't need the hypothesis that η is nontorsion here; we just need $\eta \neq \mathrm{id}$.

5.2 Extension structure in detail

In this section we aim to describe the structure of Sequence (4.3) in more detail, generalising from the examples we have already seen. We break our analysis into four cases using two independent criteria: firstly, does the element have a dyadic important point; and secondly, does the element have zero rotation number? We give each case a name (see Figure 5.14) and handle each in its own subsection. Each type of element $\alpha \in T$ is studied according to a four-step plan.

1. Define a *block form* for α: some data which give a recipe for constructing α.
2. Use the block form to find formulas for centralising elements $\gamma \in C_T(\alpha)$.
3. Check that every element $\alpha \in T$ of the current case has a *maximal block form*.
4. Use the maximality to find *all* elements $\gamma \in C_T(\alpha)$, then use their formulas from step 2 to deduce the group structure of $C_T(\alpha)$.

5.2.1 *Case 1: dyadic self-similar elements*

We now demonstrate that the recipe used to construct α, γ and η in Examples 5.1.2, 5.1.4 and 5.1.7 can be generalised, always yielding split centralisers.

Normally x denotes a point rather than a group element. We use x here to remind the reader of x_0 and x_1 from our examples.

Definition 5.2.1. Let x be any element in $F \setminus \{\mathrm{id}\}$. Let $d_0 \to d_1 \to \cdots \to d_{h-1} \to d_0$ be a circular partition in which every point is dyadic, for some $h \geq 1$. Define circular arcs $D_i = [d_i, d_{i+1}]$ with subscripts mod h, for each $0 \leq i < h$. Each D_i is called a *block*. Finally, let $\Sigma_i : I \to D_i$ be a collection of PL_2 maps, with one map for each i.

In terms of these ingredients, we define an element $\alpha \in T$ by the recipe

$$t\alpha = \begin{cases} tx^{\Sigma_i} & \text{if } t \in D_i. \end{cases}$$

Such an element α is called *dyadic self-similar*. The data $(x; \Sigma_0, \dots, \Sigma_{h-1})$ is called a *block form* for α with h blocks. If it is not possible to express α in block form with $h' > h$ blocks, we say that a block form with h blocks is *maximal*.

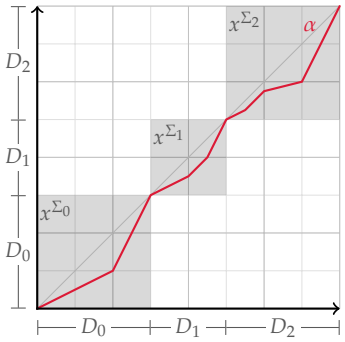

Figure 5.15: An example element an element α of the form described in Definition 5.2.1. This example was constructed using $x = x_0$ and three blocks D_i.

In an attempt to mimic Example 5.1.4, we could instead have defined

$$t\alpha = \begin{cases} tx^{\phi_i \Sigma_i} & \text{if } t \in D_i, \end{cases} \tag{5.5}$$

for some maps $\phi \in F$. However, this is unnecessary: we can define $\Sigma_i' \colon I \to D_i$ to be the PL_2 product $\phi_i \Sigma_i$. Then the map given in Equation (5.5) has block form $(x; \Sigma_i')$ as above. In this way, the maps Σ_i can 'absorb' a conjugation of x in F.

For $n > 0$ we have $\alpha^n\big|_{D_i} = (\alpha\big|_{D_i})^n = (x^{\Sigma_i})^n = (x^n)^{\Sigma_i}$. As $x \neq \mathrm{id}$, we see that $\alpha \neq \mathrm{id}$.

Because a self-similar element α is divided up into blocks and behaves similarly on each block, we can immediately write down a certain centralising element.

Lemma 5.2.2. *Let α be a dyadic self-similar element, with a block form $(x; \Sigma_i)$ containing h blocks. Define an element $\delta \in T$ by the formula*

$$t\delta = \begin{cases} t\Sigma_i^{-1}\Sigma_{i+1} & \text{if } t \in D_i, \end{cases}$$

with subscripts on Σ modulo h. Then δ commutes with α and has order h.

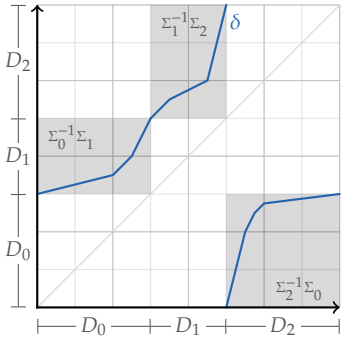

Figure 5.16: The splitting element $\delta \in C_T(\alpha)$, which cyclically shifts the blocks $\{D_i\}$ one step forward.

Proof. First we check that δ is a well-defined element of T. If $h = 1$ then δ is already well-defined. Otherwise, the only nonempty intersections among the blocks $D_i = [d_i, d_{i+1}]$ are those of the form $D_i \cap D_{i+1} = \{d_{i+1}\}$. On the one hand we have

$$d_{i+1}\,\delta\big|_{D_i} = d_{i+1}\Sigma_i^{-1}\Sigma_{i+1} = 1\Sigma_{i+1} = d_{i+2}\,;$$

on the other we have

$$d_{i+1}\,\delta\big|_{D_{i+1}} = d_{i+1}\Sigma_{i+1}^{-1}\Sigma_{i+2} = 0\Sigma_{i+2} = d_{i+2}\,,$$

so δ is well-defined. Since each Σ_i is a PL_2 map, so too is each restriction $\delta\big|_{D_i}$. Hence $\delta \in T$. We note that $D_i\delta = D_{i+1}$.

Let $t \in D_i$. Then the expressions

$$t\alpha\delta = tx^{\Sigma_i}(\Sigma_i^{-1}\Sigma_{i+1}) \quad = t\Sigma_i^{-1}x(\Sigma_i\Sigma_i^{-1})\Sigma_{i+1} \quad = t\Sigma_i^{-1}x\Sigma_{i+1}$$

and

$$t\delta\alpha = t(\Sigma_i^{-1}\Sigma_{i+1})x^{\Sigma_{i+1}} = t\Sigma_i^{-1}(\Sigma_{i+1}\Sigma_{i+1}^{-1})x\Sigma_{i+1} = t\Sigma_i^{-1}x\Sigma_{i+1}$$

are equal, so we see that $\delta \in C_T(\alpha)$.

Next, we see that

$$\begin{aligned} t\delta^h &= t\,\delta\big|_{D_i}\,\delta^{h-1} \\ &= t(\Sigma_i^{-1}\Sigma_{i+1})\,\delta\big|_{D_{i+1}}\,\delta^{h-2} \\ &= t\Sigma_i^{-1}(\Sigma_{i+1}\Sigma_{i+1}^{-1})\Sigma_{i+2}\delta^{h-2} \\ &= t\Sigma_i^{-1}\Sigma_{i+2}\delta^{h-2} \\ &\;\;\vdots \\ &= t\Sigma_i^{-1}\Sigma_{i+j}\delta^{h-j}\,, \end{aligned}$$

for each index $0 < j \leq h$. In particular, the choice $j = h$ yields $t\delta^h = t\Sigma_i^{-1}\Sigma_i = t$, so δ has order dividing h. This order must be exactly h, because δ transitively permutes the arcs $\{D_i\}_{i=0}^{h-1}$. □

Our plan is to use δ to split Sequence (4.3). In order to do that, we need to know that δ corresponds to a generator of the quotient group H. In other words, we need to make the order of δ—which is h—as large as possible. Thus we want to work with maximal block forms, where the block count h is as large as possible.

Now we characterise the set of dyadic self-similar elements, and in doing so we will explain how to obtain a maximal block form for such an element. We make use of the element δ to prove that this block form really is maximal.

Lemma 5.2.3. *An element $\alpha \in T$ is dyadic self-similar if and only if has a dyadic important point and rotation number 0. We can algorithmically determine a maximal block form for any such α.*

Proof. First suppose that α is dyadic self-similar, with block form $(x; \Sigma_i)$. Then α fixes each ∂D_i pointwise, so $\mathrm{rot}(\alpha) = 0$. As $x \neq \mathrm{id}$, there is some point $z \in I$ which is not fixed by x. Then $z\Sigma_i$ is not fixed by x^{Σ_i}, so α is not the identity. Hence $\mathrm{Fix}(\alpha)$ is neither empty nor the entire circle. The same is true of any connected component A of $\mathrm{Fix}(\alpha)$. We must have $\partial A \subseteq \mathbb{Z}[1/2]$, and so α has a dyadic important point.

Now suppose that α has a dyadic important point r_0 and rotation number 0. We must have $\alpha \neq \mathrm{id}$ because the identity has no important points. We will explain how to construct a minimal block form for α. Enumerate all the important points $I_\alpha = \{r_0 \to \cdots \to r_{n-1} \to r_0\}$ in cyclic order. Now any element of $C_T(\alpha)$ permutes I_α with some index difference. Use Algorithm 4.2.3 to find an element $\gamma \in C_T(\alpha)$ with minimal index difference $0 < d \leq n$. Then $n = dh$, where h is the order of $\gamma\big|_{I(\alpha)}$. Define our blocks to be $D_i = [r_{id}, r_{(i+1)d}]$ for each $0 \leq i < h$; that is, take $d_i = r_{id}$. We note that $D_i\alpha = D_i$ and $D_i\gamma = D_{i+1 \bmod h}$.

Choose[9] a PL_2 map $\Omega\colon I \to D_0$, and use it to define $x = \alpha\big|_{D_0}^{\Omega^{-1}}$. Then x is a PL_2 map $I \to I$, i.e. $x \in F$. We must have $x \neq \mathrm{id}$, because otherwise $\alpha = \mathrm{id}$, which we noted above is impossible. To finish, define $\Sigma_i = \Omega\,\gamma^i\big|_{D_0}$ for each i. These are PL_2 maps $I \to D_i$ as required. The

[9] See Remark 1.2.9.

calculation

$$x^{\Sigma_i} = \left(\alpha\big|_{D_0}{}^{\Omega^{-1}}\right)^{\Omega\gamma^i|_{D_0}} = \alpha\big|_{D_0}{}^{\gamma^i|_{D_0}} = \alpha^{(\gamma^i)}\Big|_{D_i} = \alpha\big|_{D_i}$$

shows that α is dyadic self-similar.

Why is $(x;\Sigma_i)$ a maximal block form? If there were a block form with $h' > h$ blocks, we could construct an element $\delta' \in C_T(\alpha)$ with order h', using Lemma 5.2.2. Then δ' would cyclically permute I_α with index difference $n/h' < n/h = d$. This is a contradiction, because γ has minimal index difference d. □

Proposition 5.2.4. *Let $\alpha \in T$ be dyadic self-similar with a maximal block form $(x;\Sigma_i)$ containing h blocks. Then $C_T(\alpha)$ splits as the wreath product $C_F(x) \wr \mathbb{Z}_h$ arising from the action $\{D_i\}_{i=1}^h \circlearrowleft \mathbb{Z}_h$. The right-hand factor is generated by the image of δ from Lemma 5.2.2.*

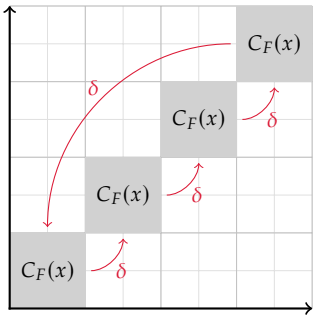

Figure 5.17: A schematic of the centraliser $C_T(\alpha)$. The h different copies of $C_F(x)$ (one for each D_i) are permuted cyclically by the splitting element δ.

Proof. The quotient in Sequence (4.3) is the restriction of $C_T(\alpha)$ to the important point set I_α. As explained in the proof of Lemma 5.2.3, because the block count h is maximal, the quotient must have order h. Since δ has order h, the extension splits.

The extension's kernel is $K = C_T(\alpha)_0$, which is isomorphic to the direct product $\prod_{0\le i<h} C_{\mathrm{PL}_2(D_i)}(\alpha|_{D_i})$ by Corollary 4.1.5. Write $A_i = C_{\mathrm{PL}_2(D_i)}(\alpha|_{D_i})$ for the ith factor of this product, so that

$$C_T(\alpha) \cong \left(\prod_{0\le i<h} A_i\right) \rtimes \langle\delta\rangle\,.$$

Now the endpoints of each D_i are all dyadic, because $r_{di} = r_0\delta^i$ and r_0 is dyadic. Thus

$$A_i = C_{\mathrm{PL}_2(D_i)}(\alpha|_{D_i}) = C_{F^{\Sigma_i}}(x^{\Sigma_i}) \cong C_F(x)\,;$$

Note that the equality $\mathrm{PL}_2(D_i) = F^{\Sigma_i}$ is a special instance of Claim 1.2.7.

in other words, the factor A_i is parameterised by elements $y \in C_F(x)$. Explicitly, the parameter y corresponds to the element $\omega_{y;i} \in A_i$ of the form

$$t\omega_{y;i} = \begin{cases} ty^{\Sigma_i} & \text{if } t \in D_i \\ t & \text{if } t \notin D_i \end{cases}$$

To conclude, we need to determine the action of the quotient on the kernel. What is the result of conjugating $\omega_{y;i}$ by δ? If $t \in D_{i+1}$ then

$$\begin{aligned} t\delta^{-1}\,\omega_{y;i}\,\delta &= t\Sigma_{i+1}^{-1}\Sigma_i\omega_{y;i}\,\delta = t\Sigma_{i+1}^{-1}\Sigma_i(y^{\Sigma_i})\Sigma_i^{-1}\Sigma_{i+1} \\ &= ty^{\Sigma_{i+1}} = t\omega_{y;i+1}\,, \end{aligned}$$

and if $t \in D_j \neq D_{i+1}$ then

$$\begin{aligned} t\delta^{-1}\,\omega_{y;i}\,\delta &= t\Sigma_j^{-1}\Sigma_{j-1}\omega_{y;i}\,\delta = t\Sigma_j^{-1}\Sigma_{j-1}(\mathrm{id}_{D_{j-1}})\Sigma_i^{-1}\Sigma_{i+1} \\ &= t\,\mathrm{id}_{D_j} = t\omega_{y;i+1}\,. \end{aligned}$$

We see that $\omega_{y;i}^{\delta} = \omega_{y;i+1}$, so conjugation by δ maps the ith factor A_i isomorphically to $A_{i+1 \bmod h}$. In doing so, the choice of representative $y \in C_F(x)$ is preserved.[10] This shows that $C_T(\alpha)$ is isomorphic to the wreath product $C_F(x) \wr \mathbb{Z}_h$ arising from the action $\{D_i\} \circlearrowleft \langle\delta\rangle$. □

[10] Informally, conjugating by δ induces the 'correct' isomorphism between A_i and A_{i+1}.

Remark 5.2.5. In the light of this proof, it is worth highlighting the structure of a generic element $\omega \in C_T(\alpha)_0$. It must be a product of the form

$$\omega = \omega_{y_1;1}\omega_{y_2;2}\cdots\omega_{y_h;h},$$

parameterised by h elements $y_i \in C_F(x)$. Thus ω is given by the casewise formula

$$t\omega = t y_i^{\Sigma_i}, \quad \text{for } t \in D_i \text{ and } 0 \le i < h.$$

There are no constraints that hold between the $\{y_i\}$: we may freely choose h elements that commute with x.

To complete the story, we ask 'which groups can be centralisers of dyadic self-similar elements?'

Corollary 5.2.6. *Let α be dyadic self-similar. Then $C_T(\alpha)$ is isomorphic to $(F^f \times \mathbb{Z}^z) \wr \mathbb{Z}_h$, for some integer parameters $z \ge 0$, $0 \le f \le z+1$ and $h > 0$ such that $(0,0) \ne (f,z) \ne (1,0)$.*

Moreover, for each such triple (f,z,h) there exists a dyadic self-similar element α whose centraliser has these parameters.

Proof. The isomorphism follows from Matucci's classification of centralisers in F: see our Theorem 4.1.3 and Claim 4.1.8 above, together with Proposition 5.2.4. The latter claim explained it was possible to construct an element x whose centraliser was isomorphic to $F^f \times \mathbb{Z}^z$, with f and z in the given range. Now $(f,z) = (1,0)$ if and only if the graph of x consists of a single fixed interval, which in turn occurs if and only if $x = \mathrm{id}$. Since the definition of block form requires $x \ne \mathrm{id}$, we exclude the possibility $(f,z) = (1,0)$. This establishes that our conditions are necessary.

We show the same conditions are sufficient by building a specific element α whose centraliser has parameters (f,z,h). Given f and z, partition the unit interval into $f+z$ standard dyadic intervals $E_j = [e_j, e_{j+1}]$, for $0 \le j < f+z$. Over $E_0, E_2, \ldots, E_{2(f-1)}$, define $x\big|_{E_j}$ to be the identity. On the remaining z intervals $\{E_j\}$, define $x\big|_{E_j} = (x_0^{\,j})^{\Xi_j}$, where Ξ_j is some PL_2 map $I \to E_j$ and $x_0 \in F$ is the standard generator shown in Figure 1.1. To build α, choose a partition of the unit circle into h standard dyadic intervals $D_0, \ldots, D_{h-1}$, and select PL_2 maps

Figure 5.18: Schematic illustrating the break down of $I = [0,1]$ into standard dyadic intervals E_j. On f of these, x is just the identity, leading to a centraliser isomorphic to F over each such interval. No two F-type intervals are adjacent. On the remaining z intervals, x is a specific one-bump function, and hence has a cyclic centraliser.

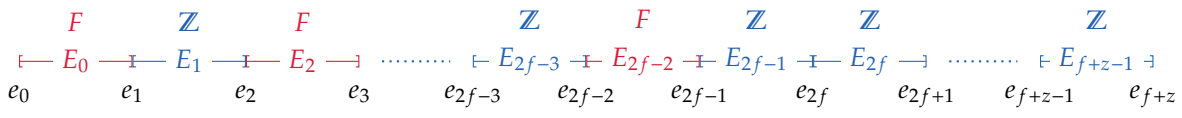

$\Sigma_i : I \to D_i$. As usual, we define $\alpha\big|_{D_i}$ to be x^{Σ_i}, so that α has block form $(x;\Sigma_i)$ with h blocks. As noted above, $x \neq \mathrm{id}$ because $(f,z) \neq (1,0)$.

The fixed point set of α is

$$\mathrm{Fix}(\alpha) = \bigcup_i \mathrm{Fix}(\alpha\big|_{D_i}) = \bigcup_i \mathrm{Fix}(x^{\Sigma_i}) = \bigcup_i \mathrm{Fix}(x)\Sigma_i \,,$$

so now we need the fixed point set of x. By construction, this is the union

$$\mathrm{Fix}(x) = E_0 \cup E_2 \cup \cdots \cup E_{2(f-1)} \cup \{e_{2f}, \ldots, e_{f+z}\} \,.$$

of intervals and isolated points. Noting that $e_{f+z}\Sigma_i = e_0\Sigma_{i+1}$, we can express $\mathrm{Fix}(\alpha)$ as

$$\mathrm{Fix}(\alpha) = \bigcup_i \Big(E_0 \cup E_2 \cup \cdots \cup E_{2(f-1)} \cup \{e_{2f}, \ldots, e_{f+z-1}\}\Big)\Sigma_i \,.$$

Each of the finitely many sets within the union is closed and pairwise disjoint, meaning that boundary of union is the union of boundaries. In symbols,

$$\begin{aligned} I_\alpha = \partial\,\mathrm{Fix}(\alpha) &= \bigcup_i \partial\Big(E_0 \cup E_2 \cup \cdots \cup E_{2(f-1)} \cup \{e_{2f}, \ldots, e_{f+z}\}\Big)\Sigma_i \\ &= \bigcup_i \{e_0, e_1, e_2, \ldots, e_{f+z-1}\}\Sigma_i \\ &= \{e_j\Sigma_i \mid 0 \le j < f+z \text{ and } 0 \le i < h\} \,. \end{aligned}$$

Setting $d = f + z$, we see that α has $n = hd$ important points $I_\alpha = \{r_{j+id} = e_j\Sigma_i\}$.

To conclude our proof, we need to show the the block form $(x;\Sigma_i)$ is maximal. We do so by showing that that there is no element $\gamma \in C_T(\alpha)$ which permutes the important points $r_k \mapsto r_{k+d'}$, for some $0 < d' < d$. Then we will know that the centraliser quotient has order exactly $h = n/d$.

Seeking a contradiction, suppose we have such a γ. Then the functions α and α^γ are equal, and hence so too are their right-derivatives at any point. Differentiate α at $e_1\Sigma_i \in E_1\Sigma_i \subseteq D_i$ to see that

$$\begin{aligned} [e_1\Sigma_i]^+(\alpha\big|_{D_i})' = [e_1\Sigma_i]^+(x^{\Sigma_i}\big|_{E_1\Sigma_i})' &= [e_1\Sigma_i]^+(x\big|_{E_1}{}^{\Sigma_i})' = {e_1}^+(x\big|_{E_1})' \\ &= {e_1}^+(x_0^{\Xi_1})' = [e_1\Xi_1^{-1}]^+ x_0' = 0x_0' = 2^{-1} \,. \end{aligned}$$

After some simplification, differentiating α^γ at the same point $e_1\Sigma_i$ yields

$$[e_1\Sigma_i]^+(\alpha^\gamma)' = [e_1\Sigma_i\gamma^{-1}]^+\alpha' \,. \tag{5.6}$$

We note that

$$e_1\Sigma_i\gamma^{-1} = r_{id+1}\gamma^{-1} = r_{id+1-d'} = r_{(i-1)d+d-d'+1} = e_{d-d'+1}\Sigma_{i-1} \,,$$

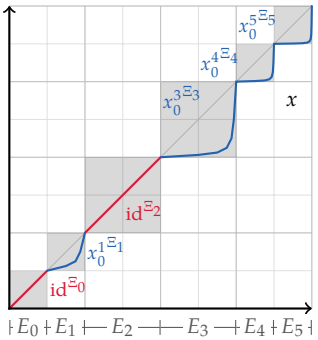

Figure 5.19: We choose a particular element $x \in F$ whose centraliser is isomorphic to $F^f \times \mathbb{Z}^z$. In this example, $f = 2$ and $z = 4$.

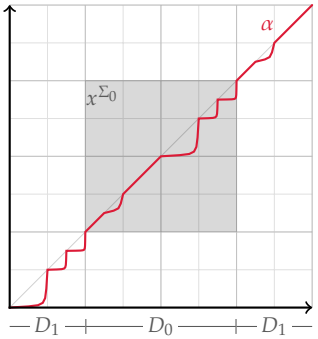

Figure 5.20: Continuing Figure 5.19, we select $h = 2$ dyadic arcs to cover the circle, namely $D_0 = [1/4, 3/4]$ and $D_1 = [3/4, 1/4]$. Selecting PL_2-maps $\Sigma_i : I \to D_i$, we define α on D_i to be x^{Σ_i}. We've only shaded $D_0^2 \subset (S^1)^2$ for clarity.

We're making use of Claim 2.3.14 to handle derivatives of conjugates here.

so we may continue the calculation in Equation (5.6) to see that

$$\begin{aligned}[e_1\Sigma_i]^+(\alpha^\gamma)' &= [e_{d-d'+1}\Sigma_{i-1}]^+(\alpha|_{E_{d-d'+1}\Sigma_{i-1}})' \\ &= [e_{d-d'+1}\Sigma_{i-1}]^+(x|_{E_{d-d'+1}}{}^{\Sigma_{i-1}})' \\ &= {e_{d-d'+1}}^+(x|_{E_{d-d'+1}})' \,.\end{aligned}$$

Depending on the value of $d-d'+1$, the restriction of x is either the identity or a scaled version of ${x_0}^{d-d'+1}$. In the former case, the derivative is 1; in the latter it is

$$\begin{aligned}{e_{d-d'+1}}^+([{x_0}^{d-d'+1}]^{\Xi_{d-d'+1}})' &= [e_{d-d'+1}\Xi_{d-d'+1}^{-1}]^+({x_0}^{d-d'+1})' \\ &= 0({x_0}^{d-d'+1})' \\ &= 2^{-(d-d'+1)} \,.\end{aligned}$$

Hence either $2^{-1}=1$ or $2^{-1}=2^{-(d-d'+1)}$. The first possibility is false; the latter holds if and only if $d=d'$, which contradicts $0<d'<d$. This contradiction completes the proof. □

5.2.2 *Case 2: dyadic offset-similar elements*

Examples 5.1.3, 5.1.6 and 5.1.8 describe the centraliser of elements $\beta, \delta, \zeta \in T$ with nonzero rotation number and a dyadic important point. Can we formulate a description of a generic such element?

Definition 5.2.7. Let $d_0 \to d_1 \to \ldots \to d_{h-1} \to d_0$ be a circular partition in which every point is dyadic, for some $h \geq 2$. Define blocks $D_i = [d_i, d_{i+1}]$ for each $0 \leq i < h$, and let $\Sigma_i : I \to D_i$ be a PL_2 map, for each i in the same range. Next, let q be a divisor of h, say $h = sq$. Also let $0 < p < q$ be an integer coprime to q, so that $0 < ps < h$. Finally, select a map $x_i \in F$ for each $0 \leq i < h$, subject to the constraint that the product $X = \prod_{0\leq j<q} x_{i+jps}$ is independent of i and not equal to the identity. (Subscripts on the x_i are modulo h.)

From these ingredients, we define an element $\beta \in T$ by the recipe

$$t\beta = \left\{ t\Sigma_i^{-1} x_i \Sigma_{i+ps} \quad \text{if } t \in D_i. \right.$$

Any element of this form is said to be *dyadic offset-similar* (DOS). The data tuple $(\{x_i\}; \{\Sigma_i\}; p; q)$ is called a *block form* for β, with h blocks. If it is not possible to express β in block form with $h' > h$ blocks, we say that a block form with h blocks is *maximal*.

There are a few changes to note from Definition 5.2.1. This time, we require at least $h \geq 2$ blocks, which are no longer fixed by the result of our recipe: instead, $D_i\beta = D_{i+ps} \neq D_i$. Each of these is rearranged by (a suitable conjugate of) its own element $x_i \in F$; we no longer use only a single $x \in F \setminus \{1\}$. We note that

$$\begin{aligned}X^{x_i} &= x_i^{-1}(x_i x_{i+ps} x_{i+2ps} \cdots x_{i+(q-1)ps})[x_i] \\ &= x_{i+ps} x_{i+ps+ps} \cdots x_{i+ps+(q-2)ps}[x_{i+ps+(q-1)ps}] = X \,,\end{aligned} \tag{5.7}$$

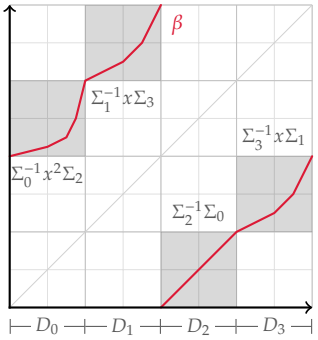

Figure 5.21: Schematic of a dyadic offset-similar element β. Here we have $h = 4$ blocks; the other numerical parameters are $p = 1$, $q = 2$ and $s = 2$. Our elements of F are $(x_0, x_1, x_2, x_3) = (x^2, x, \mathrm{id}, x)$, where x is one of the standard generators of F (usually denoted by x_0).

The product X is independent of i: each of its $h = 4$ possible values

$$\begin{aligned}x_0x_2 &= x^2\,\mathrm{id} = x^2 \\ x_1x_3 &= x^2 \\ x_2x_0 &= \mathrm{id}\,x^2 = x^2 \\ x_3x_1 &= x^2\end{aligned}$$

are all equal to x^2.

which shows that each $x_i \in C_F(X)$.

Claim 5.2.8. *Let β have block form $(x_i; \Sigma_i; p; q)$. Then* $\mathrm{rot}(\beta) = p/q \neq 0$ *in lowest terms.*

Proof. We know that β cyclically permutes the intervals $\{D_i\}_{0 \leq i < h}$, sending $D_i \mapsto D_{i+ps}$. Since there are $h = qs$ blocks in total, the rotation number is $\mathrm{rot}(\beta) = ps/qs = p/q$. This is in lowest terms because p and q are coprime. As $0 < p < q$, $q \geq 2$ and $\mathrm{rot}(\beta) \neq 0$. □

The original goal of introducing block forms for elements α with rotation number 0 was to make it easy to write down centralising elements with small (hopefully minimal) index difference. We can do the same for β with nonzero rotation number—with some extra effort. First we establish some helpful notation; then we derive conditions for such a centralising elements θ to exist. As a corollary we point out that such a θ can always be constructed.

Notation 5.2.9. Let β have block form $(x_i; \Sigma_i; p; q)$ and let X be the product $X = \prod_{0 \leq j < q} x_{i+jps}$. We define the symbol $\pi_{i,m} = \prod_{0 \leq j < m} x_{i+jps}$, for any integer $0 \leq i < h$ and any natural number $m \geq 1$. This has the following properties, where we treat the first subscript i on π as an integer modulo h.

1. $\pi_{i,m+1} = \pi_{i,m} x_{i+mps}$, and hence $\pi_{i,m+m'} = \pi_{i,m} \pi_{i+mps,m'}$.
2. $\Sigma_i^{-1} \pi_{i,m} \Sigma_{i+mps} = \beta^m\big|_{D_i}$.
3. $\pi_{i,q} = X$, for any value of i; hence $\pi_{i,qk} = X^k$.
4. $\pi_{i,m} \in C_F(X)$, because it's a product of elements $x_i \in C_F(X)$.

Lemma 5.2.10. *Let β have block form $(x_i; \Sigma_i; p; q)$ with h blocks, and let X be the product $X = \prod_{0 \leq j < q} x_{i+jps}$. Then there exists an element $\theta \in C_T(\beta)$ with* $\mathrm{rot}(\theta) = 1/h$ *if and only if there exist maps $\theta_0, \ldots, \theta_{h-1} \in C_F(X)$ satisfying the simultaneous equations*

$$x_i \theta_{i+ps} = \theta_i x_{i+1}, \quad \text{for all } 0 \leq i < h. \tag{5.8}$$

Proof. If $\theta \in C_T(\beta)$ then θ maps D_i to D_{i+1}. Thus we have $\theta\big|_{D_i} = \Sigma_i^{-1}(\Sigma_i\, \theta\big|_{D_i}\, \Sigma_{i+1})^{-1} \Sigma_{i+1}$. As shorthand, define $\theta_i = \Sigma_i\, \theta\big|_{D_i}\, \Sigma_{i+1}^{-1}$ for each i. Then $\theta\big|_{D_i} = \Sigma_i^{-1} \theta_i \Sigma_{i+1}$ holds for every $0 \leq i < h$. Now the products $\theta\beta$ and $\beta\theta$ are equal if their restrictions to D_i are equal, for all i. The former restriction is

$$\theta\big|_{D_i}\, \beta\big|_{D_{i+1}} = \Sigma_i^{-1} \theta_i \Sigma_{i+1} \Sigma_{i+1}^{-1} x_{i+1} \Sigma_{i+1+ps} = \Sigma_i^{-1} \theta_i x_{i+1} \Sigma_{i+1+ps}\,,$$

whereas the latter restriction is

$$\beta\big|_{D_i}\, \theta\big|_{D_{i+1}} = \Sigma_i^{-1} x_i \Sigma_{i+ps} \Sigma_{i+ps}^{-1} \theta_{i+ps} \Sigma_{i+ps+1} = \Sigma_i^{-1} x_i \theta_{i+ps} \Sigma_{i+ps+1}\,.$$

As these expressions are equal, we learn that

$$\Sigma_i^{-1} \theta_i x_{i+1} \Sigma_{i+1+ps} = \Sigma_i^{-1} x_i \theta_{i+ps} \Sigma_{i+ps+1} \quad \text{for each } 0 \leq i < h,$$

which is equivalent to Condition (5.8). Repeatedly apply this condition using different values of i to see that

$$\begin{aligned}\theta_i &= x_i \theta_{i+ps} x_{i+1}^{-1} \\ &= x_i (x_{i+ps} \theta_{i+2ps} x_{i+ps+1}^{-1}) x_{i+1}^{-1} \\ &= x_i x_{i+ps} x_{i+2ps} \theta_{i+3ps} x_{i+2ps+1}^{-1} x_{i+ps+1}^{-1} x_{i+1}^{-1} \qquad (5.9) \\ &\vdots \\ &= \pi_{i,m} \theta_{i+mps} \pi_{i+1,m}^{-1}, \qquad \text{for any integer } m \geq 0.\end{aligned}$$

In particular, take $m = q$ so that $\theta_{i+mps} = \theta_i$ and that both π symbols become X. Then we see that $\theta_i \in C_F(X)$, as claimed.

Conversely, given $\{\theta_i\}_{0 \leq i < h} \subset C_F(X)$ satisfying Condition (5.8), define[11] an element $\theta \in T$ by the formula $\theta\big|_{D_i} = \Sigma_i^{-1} \theta_i \Sigma_{i+1}$ for $0 \leq i < h$. Then θ maps D_i to $D_{i+1 \bmod h}$ for all $0 \leq i < q$, so $\mathrm{rot}(\theta) = 1/h$. Repeat the above calculations for $\theta\big|_{D_i} \beta\big|_{D_{i+1}}$ and $\beta\big|_{D_i} \theta\big|_{D_{i+1}}$; the two functions will be equal precisely because our $\{\theta_i\}$ satisfy Condition (5.8). This holds for all $0 \leq i < q$, so $\theta\beta = \beta\theta$. □

[11] Formally, use Lemma 2.1.7 to see that this is a well-defined map $\theta \in T$ instead of $\theta \in \mathrm{Homeo}_+(S^1) \setminus T$.

Corollary 5.2.11. *Let β and X be as in Lemma 5.2.10. For any elements $\theta_0, \ldots, \theta_{s-1} \in C_F(X)$, there is a unique element $\theta \in C_T(\beta)$ with $\theta\big|_{D_j} = \Sigma_j^{-1} \theta_j \Sigma_{j+1}$ for each $0 \leq j < s$. Moreover* $\mathrm{rot}(\theta) = 1/h$.

Proof. Adopt the notation used in Lemma 5.2.10 to describe the given block from for β. Suppose we have elements $\theta_0, \ldots, \theta_{s-1} \in C_F(X)$. To prove this result, we will establish that there are *unique* elements $\theta_s, \ldots, \theta_{h-1}$ such that the whole collection $\{\theta_0, \ldots, \theta_{h-1}\}$ satisfies Condition (5.8). For each $0 \leq j < s$, we must define $\theta_{ps}, \theta_{1+ps} \ldots, \theta_{s-1+ps}$ by $\theta_{j+ps} = x_j^{-1} \theta_j x_{j+1}$, in order to satisfy to satisfy Condition (5.8) when $i = j$. For the same values of j, we must similarly define

$$\begin{aligned}\theta_{j+2ps} &= x_{j+ps}^{-1} \theta_{j+ps} x_{j+ps+1} \\ &= x_{j+ps}^{-1} x_j^{-1} \theta_j x_{j+1} x_{j+ps+1} \\ &= \pi_{j,2}^{-1} \theta_j \pi_{j+1,2},\end{aligned}$$

to satisfy Condition (5.8) for $i = j + ps$. In general, we define $\theta_{j+mps} = \pi_{j,m}^{-1} \theta_j \pi_{j+1,m}$; this is necessary to satisfy the required condition when $i = j + (m-1)ps$, for each $0 < m < q$ and each $0 \leq j < s$.

With all our θ_i now defined, we need to verify that the last remaining case ($i = j + (q-1)ps$) of Condition (5.8) holds, i.e. to verify that

$$x_{j+(q-1)ps} \theta_{j+qps} = \theta_{j+(q-1)ps} x_{j+(q-1)ps+1} \qquad (5.10)$$

for each $0 \leq j < s$. The left-hand side of Condition (5.10) is the product $x_{j+(q-1)ps} \theta_{j+qps}$, since $qs = 0 \pmod h$; this is equal to the right-hand side because

$$\theta_{j+(q-1)ps} x_{j+(q-1)ps+1} = \pi_{j,q-1}^{-1} \theta_j \pi_{j+1,q-1} \, x_{j+(q-1)ps+1}$$

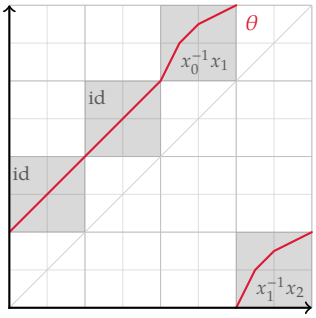

Figure 5.22: The result of our recipe for constructing a centralising element $\theta \in C_F(\beta)$, where β is the element illustrated in Figure 5.21. We choose θ to be the identity on D_0 and D_1; this determines θ on the remaining D_i.

$$
\begin{aligned}
&= \pi_{j,q-1}^{-1}\theta_j\pi_{j+1,q} \\
&= (Xx_{j+(q-1)ps}^{-1})^{-1}\theta_j X \\
&= x_{j+(q-1)ps}X^{-1}\theta_j X \\
&= x_{j+(q-1)ps}\theta_j\,.
\end{aligned}
$$

Thus Condition (5.8) is satisfied for every $0 \le i < h$. □

We now give a dynamical characterisation of DOS elements. For each DOS element β, we explain how to construct a maximal block form for β.

Lemma 5.2.12. *An element $\beta \in T$ is dyadic offset-similar if and only if it has a dyadic important point and nonzero rotation number. We can algorithmically determine a maximal block form for any such β.*

Proof. First assume that β has a block form $(x_i; \Sigma_i; p; q)$ with h blocks. We saw that $\mathrm{rot}(\beta) \neq 0$ in Claim 5.2.8, so now we need to find a dyadic important point. The whole important point set is $I_\beta = \partial\,\mathrm{Fix}(\beta^q)$. On a given block, the qth power is

$$\beta^q\big|_{D_i} = \Sigma_i^{-1}\pi_{i,q}\Sigma_{i+qps} = \Sigma_i^{-1}X\Sigma_i\,,$$

by Notation 5.2.9. Now since $X \neq \mathrm{id}$ there must be a dyadic important point[12] $r \in I_X = \partial\,\mathrm{Fix}(X)$. It follows that $r\Sigma_i$ is a important point of β; this is also dyadic because r is dyadic and Σ_i is a PL_2 map.

[12] Let $[0, s]$ be the first linear segment of X, for some $s < 1$. If this segment has gradient 1, take $r = s$; otherwise take $r = 0$. Then $r \in \partial\,\mathrm{Fix}(X)$.

In the other direction, we take an element β with rotation number $p/q \neq 0$ in lowest terms which possesses a dyadic important point r_0. We will describe an algorithm to produce a maximal block form for β. Enumerate all important points in cyclic order as $I_\beta = \{r_0 \to r_1 \to \cdots \to r_{n-1} \to r_0\}$. Use Algorithm 4.2.3 to find an element $\gamma \in C_T(\beta)$ with minimal index difference $0 < d \le n$. Then $n = dh$ for some integer h; use d to divide the important points into h blocks $D_i = [r_{id}, r_{(i+1)d}]$. As each r_i lies in a size q orbit under β, we must have $h = qs$ for some $s \ge 1$; thus $h \ge q \ge 2$.

As γ commutes with β, it must also commute with β^q. This means that the quotient of $C_T(\beta^q)$ (from the structure of Sequence (4.3)) has order at least h (possibly larger). But the power β^q is dyadic self-similar, so it has a (not necessarily maximal) block form $(X; \Sigma_i)$ with h blocks. Since the important points $I(\beta^q) = I_\beta$ are the same, we may assume without loss of generality that the codomain of Σ_i is D_i. Thus we obtain $\beta^q\big|_{D_i} = \Sigma_i^{-1}X\Sigma_i$.

Define $\delta \in T$ by $\delta\big|_{D_i} = \Sigma_i^{-1}\Sigma_{i+1}$ as in Lemma 5.2.2, so that $\delta \in C_T(\beta^q)$. Now certainly $\beta \in C_T(\beta^q)$, so the product $\epsilon = \beta\delta^j \in C_T(\beta^q)$ also. Choose $j = -ps$ so that $\epsilon \in C_T(\beta^q)_0$ has rotation number 0. From Remark 5.2.5, we know that $\epsilon\big|_{D_i} = x_i^{\Sigma_i}$, for some maps $x_i \in C_F(X)$. Then

$$\beta\big|_{D_i} = \epsilon\big|_{D_i}\,\delta^{ps}\big|_{D_i} = (\Sigma_i^{-1}x_i\Sigma_i)(\Sigma_i^{-1}\Sigma_{i+ps}) = \Sigma_i^{-1}x_i\Sigma_{i+ps}\,.$$

Now that we have formulas for $\beta\big|_{D_i}$ and $\beta^q\big|_{D_i}$, we can restrict the equality $\beta^q = \beta\beta\cdots\beta$ to D_i to learn

$$\begin{aligned}\Sigma_i^{-1} X \Sigma_i &= (\Sigma_i^{-1} x_i \Sigma_{i+ps})(\Sigma_{i+ps}^{-1} x_{i+ps} \Sigma_{i+2ps}) \cdots (\Sigma_{i+(q-1)ps}^{-1} x_{i+(q-1)ps} \Sigma_{i+qps}) \\ &= \Sigma_i^{-1} \Big(\prod_{0 \le j < q} x_{i+jps} \Big) \Sigma_i \,.\end{aligned}$$

We see that the product $X = \prod_{0\le j<q} x_{i+jps}$ is independent of i, and $X \neq 1$. Thus β has block form $(x_i; \Sigma_i; p; q)$.

Finally we consider the maximality of h. We deliberately chose γ to have minimal index difference d, and so the image of γ in the last term H of Sequence (4.3) has maximal order $h = n/d$. If our block form for β were not maximal, there would be a block form with $h' > h$ blocks. But then we could construct an element $\theta' \in C_T(\beta)$ with rotation number $1/h'$ using Corollary 5.2.11. This would yield an element of order $h' > h$ in H, which is a contradiction. We conclude that our block form is indeed maximal. □

With maximal block forms in hand, we can begin to describe the structure of Sequence (4.3) for $C_T(\beta)$.

Lemma 5.2.13. *Let $\beta \in T$ be dyadic offset-similar with maximal block form $(x_i; \Sigma_i; p; q)$ consisting of h blocks. Set $X = \prod_{0\le m<q} x_{i+mps}$. Then there is a short exact sequence*

$$C_F(X)^s \hookrightarrow C_T(\beta) \twoheadrightarrow \mathbb{Z}_h \,, \tag{5.11}$$

This sequence is a specific instance of Sequence (4.3).

where $h = qs$. Given an element $\theta \in C_T(\beta)$ which is a preimage of $1 \in \mathbb{Z}_h$, we can explicitly describe how θ conjugates the kernel of this extension.[13]

[13] See Equation (5.14).

Proof. Given a choice of representative arc Q, we know that the kernel of Sequence (4.3) for $C_T(\beta)$ is isomorphic to $C_{\mathrm{PL}_2(Q)}(\beta^q\big|_Q)$. Since this arc needs to cover $1/q$th of the circle, we take $Q = D_0 \cup \cdots \cup D_{s-1}$, where D_i is the domain of Σ_i. Because of the condition on the $\{x_i\}$ in the definition of offset-similar block form, we know that $\beta^q\big|_{D_i} = \Sigma_i^{-1} X \Sigma_i$. Thus, by the argument of Theorem 4.1.3, our kernel is a direct product of other centralisers, namely

$$\begin{aligned} C_{\mathrm{PL}_2(Q)}(\beta^q\big|_Q) &\cong \prod_{0\le j<s} C_{\mathrm{PL}_2(D_j)}(\beta^q\big|_{D_j}) && \text{(5.12a)} \\ &= \prod_{0\le j<s} C_{F^{\Sigma_j}}(X^{\Sigma_j}) && \text{(5.12b)} \\ &\cong \prod_{0\le j<s} C_F(X) \,. && \text{(5.12c)} \end{aligned}$$

Reminder. Isomorphism (5.12a) decomposes the centraliser over Q into the product of centralisers over the $\{D_j\}$. This is valid because the endpoints ∂D_j are dyadic important points. Isomorphism (5.12c) is a conjugation by the PL_2 map Σ_j^{-1}.

Any element $\theta \in C_T(\beta)$ with rotation number $1/h$ projects onto $1 \in \mathbb{Z}_h$. We know that such an element exists thanks to Corollary 5.2.11. We also know that $1/h$ is the smallest nonzero rotation number possible, because our block form (and hence h) is maximal. Thus the extension's quotient is $\mathbb{Z}_h$, as claimed.

We saw in Equation (5.12) that our kernel's elements were parameterised by tuples $(\omega_0, \ldots, \omega_{s-1}) \in C_F(X)^s$. After applying the inverse of Isomorphism (5.12c) followed by the inverse of Isomorphism (5.12a), our parameters yield an element $\omega \in \mathrm{PL}_2(Q)$ defined by $\omega\big|_{D_j} = \omega_j^{\Sigma_j}$, for each $0 \leq j < s$. To obtain a fully-fledged element of T, we need to apply the inverse of the isomorphism $C_T(\beta)_0 \cong C_{\mathrm{PL}_2(Q)}(\beta^q\big|_Q)$ discussed in Theorem 4.2.1 and Remark 4.2.2. Take ω and form its unique extension $\bar{\omega} \in C_T(\beta)$; this is given by

$$\bar{\omega}\big|_{D_{j+mps}} = (\omega\big|_{D_j})^{\beta^m\big|_{D_j}} = (\omega_j^{\Sigma_j})^{\Sigma_j^{-1}\pi_{j,m}\Sigma_{j+mps}} = \omega_j^{\pi_{j,m}\Sigma_{j+mps}}, \tag{5.13}$$

where $0 \leq j < s$ and $0 \leq m < h$. Thus the conjugate $\bar{\omega}^\theta$ is given by

$$\bar{\omega}^\theta\big|_{D_{j+mps}} = (\bar{\omega}\big|_{D_{j+mps-1}})^{\theta|_{D_{j+mps-1}}}.$$

Now $\bar{\omega}^\theta$ is also an element of the extension's kernel, so it must be described by s parameters in $C_F(X)$. What is the jth such parameter? First we handle the cases $0 < j < s$. Write $\theta\big|_{D_i} = \Sigma_i^{-1}\theta_i\Sigma_{i+1}$ as in the proof of Lemma 5.2.10. Then

$$\begin{aligned}
\bar{\omega}^\theta\big|_{D_{j+mps}} &= (\bar{\omega}\big|_{D_{j-1+mps}})^{\theta|_{D_{j-1+mps}}} \\
&= (\omega_{j-1}^{\pi_{j-1,m}\Sigma_{j-1+mps}})^{\Sigma_{j-1+mps}^{-1}\theta_{j-1+mps}\Sigma_{j+mps}} \\
&= \omega_{j-1}^{\pi_{j-1,m}\theta_{j-1+mps}\Sigma_{j+mps}} \\
&= \omega_{j-1}^{\theta_{j-1}\pi_{j,m}\Sigma_{j+mps}},
\end{aligned}$$

where the last equality is due to Equation (5.9). Comparing with Equation (5.13), we see that the jth parameter for $\bar{\omega}^\theta$ is $\omega_{j-1}^{\theta_{j-1}}$.

The case $j = 0$ is a little different: here $j + mps - 1 = mps - 1$ needs to be expressed in the form $(s-1) - s + mps = (s-1) + (m-k)ps$, where $k = p^{-1} \bmod q$. Then

$$\begin{aligned}
\bar{\omega}^\theta\big|_{D_{mps}} &= (\bar{\omega}\big|_{D_{(s-1)+(m-k)ps}})^{\theta|_{D_{(s-1)+(m-k)ps}}} \\
&= (\omega_{s-1}^{\pi_{s-1,m-k}\Sigma_{mps-1}})^{\Sigma_{mps-1}^{-1}\theta_{(s-1)+(m-k)ps}\Sigma_{mps}} \\
&= \omega_{s-1}^{\pi_{s-1,m-k}\theta_{(s-1)+(m-k)ps}\Sigma_{mps}} \\
&= \omega_{s-1}^{\theta_{s-1}\pi_{s,m-k}\Sigma_{mps}},
\end{aligned}$$

again thanks to Equation (5.9). Now since $\pi_{0,m} = \pi_{0,k}\pi_{s,m-k}$, we can rewrite this last line as

$$\bar{\omega}^\theta\big|_{D_{mps}} = \omega_{s-1}^{(\theta_{s-1}\pi_{0,k}^{-1})\pi_{0,m}\Sigma_{mps}}.$$

Comparing with Equation (5.13) shows that the parameter defining $\bar{\omega}^\theta$ over D_0 is $\omega_{s-1}^{\theta_{s-1}\pi_{0,k}^{-1}}$.

In summary, the action of θ on $C_F(X)^s$ is given by

$$(\omega_0, \ldots, \omega_{s-1})^\theta = (\omega_{s-1}^{\theta_{s-1}\pi_{0,k}^{-1}}, \omega_0^{\theta_0}, \omega_1^{\theta_1}, \ldots, \omega_{s-2}^{\theta_{s-2}}), \tag{5.14}$$

where the $\omega_j \in C_F(X)$ are arbitrary parameters for elements $\bar{\omega}$ of the extension's kernel. This completes our proof. □

Now we will study the conditions under which Sequence (5.11) for $C_T(\beta)$ splits. This occurs if and only if we can find an element θ of order $h = \left|C_T(\beta)\big|_{I_\beta}\right|$ in the centraliser. The next lemma tells us how to determine if such a θ exists, given a maximal block form. Even better, if such an element exists then we can construct an example.

Proposition 5.2.14. *Let $\beta \in T$ have maximal block form $(x_i; \Sigma_i; p; q)$ with h blocks. Set $X = \prod_{0\le j<q} x_{i+jps}$. Then Sequence (5.11) splits if and only if X has a qth root in F.*

An equivalent statement of this result is that $C_T(\beta)$ splits if and only if β^q has a qth root in F.

Proof. To split the extension, we need to find an element $\theta \in C_T(\beta)$ whose image in the right-hand term $\mathbb{Z}_h$ of Sequence (5.11) has order h. Without loss of generality, we may assume that $\mathrm{rot}(\theta) = 1/h$. If not, then say that $\mathrm{rot}(\theta) = p'/h$ in lowest terms. Set $k' = p'^{-1} \bmod h$, so that $\theta^{k'}$ has rotation number $p'k'/h = 1/h \pmod 1$.

To start, we assume that such a θ exists and construct a qth root of X. Arguing as in the proof of Lemma 5.2.10, we may write $\theta\big|_{D_i} = \Sigma_i^{-1}\theta_i\Sigma_{i+1}$ for some maps $\theta_i \in C_F(X)$. First we note that

$$\theta^h\big|_{D_0} = (\Sigma_0^{-1}\theta_0\Sigma_1)(\Sigma_1^{-1}\theta_1\Sigma_2)\cdots(\Sigma_{h-1}^{-1}\theta_{h-1}\Sigma_0) = \Sigma_0^{-1}\theta_0\theta_1\cdots\theta_{h-1}\Sigma_0\,.$$

For each $0 \le m < q$ define $\Theta_m = \prod_{0\le j<s}\theta_{ms+j}$. Then $\Theta_m \in C_F(X)$, and the previous equation now reads $\theta^h\big|_{D_0} = \Sigma_0^{-1}\Theta_0\Theta_1\cdots\Theta_{q-1}\Sigma_0$. Set $k = p^{-1} \bmod q$, so that $kps \equiv s \pmod h$. Use this together with Equation (5.9) to conclude that

$$\theta_{j+ms} = \theta_{j+mkps} = \pi_{j,mk}^{-1}\theta_j\pi_{j+1,mk}$$

for each m as above. This means that

$$\begin{aligned}\Theta_m &= \theta_{mkps}\theta_{mkps+1}\cdots\theta_{mkps+(s-1)}\\ &= \left(\pi_{0,mk}^{-1}\theta_0\pi_{1,mk}\right)\left(\pi_{1,mk}^{-1}\theta_1\pi_{2,mk}\right)\cdots\left(\pi_{s-1,mk}^{-1}\theta_{s-1}\pi_{s,mk}\right)\\ &= \pi_{0,mk}^{-1}\theta_0\theta_1\cdots\theta_{s-1}\pi_{s,mk}\\ &= \pi_{0,mk}^{-1}\Theta_0\pi_{s,mk}\,.\end{aligned}$$

Expressing each Θ_m in terms of Θ_0 only, we learn that

$$\theta^h\big|_{D_0} = \Sigma_0^{-1}\Theta_0\left(\pi_{0,k}^{-1}\Theta_0\pi_{s,k}\right)\left(\pi_{0,2k}^{-1}\Theta_0\pi_{s,2k}\right)\cdots \cdots\left(\pi_{0,(q-1)k}^{-1}\Theta_0\pi_{s,(q-1)k}\right)\Sigma_0\,.$$

Now because $kps \equiv s \pmod h$ we have $\pi_{0,k}\pi_{s,mk} = \pi_{0,(m+1)k}$. It follows that the $q-2$ factors of the form $\pi_{s,mk}\pi_{0,(m+1)k}^{-1}$ in the previous equation are all equal to $\pi_{0,k}^{-1}$. The previous equation now reads

$$\theta^h\big|_{D_0} = \Sigma_0^{-1}\left(\Theta_0\pi_{0,k}^{-1}\right)^{q-1}\Theta_0\pi_{s,(q-1)k}\Sigma_0\,;$$

rewrite the right-hand side to obtain the expression

$$\begin{aligned}\theta^h\big|_{D_0} &= \Sigma_0^{-1}\left(\Theta_0\pi_{0,k}^{-1}\right)^q \pi_{0,k}\pi_{s,(q-1)k}\Sigma_0 \\ &= \Sigma_0^{-1}\left(\Theta_0\pi_{0,k}^{-1}\right)^q \pi_{0,qk}\Sigma_0 \\ &= \Sigma_0^{-1}\left(\Theta_0\pi_{0,k}^{-1}\right)^q X^k\Sigma_0\end{aligned} \tag{5.15}$$

for $\theta^h\big|_{D_0}$.

The above holds for any $\theta \in C_T(\beta)$ with $\mathrm{rot}(\theta) = 1/h$. Now we also that $\theta^h = \mathrm{id}$. Applying this to Equation (5.15) yields

$$\Sigma_0\,\mathrm{id}_{D_0}\,\Sigma_0^{-1} = \left(\Theta_0\pi_{0,k}^{-1}\right)^q X^k\,.$$

Since $\Sigma_0\,\mathrm{id}_{D_0}\,\Sigma_0^{-1}$ is equal to $\mathrm{id}_I = \mathrm{id}_F$, the above is equivalent to the equation

$$X^{-k} = \left(\Theta_0\pi_{0,k}^{-1}\right)^q$$

in which all terms belong to $C_F(X)$. Next, set $\phi = \Theta_0\pi_{0,k}^{-1}$ so that $\phi^q = X^{-k}$. Each of the elements $\{\theta_i, x_i\}_{0\le i<h}$ belongs to $C_F(X)$, so we must also have $\phi \in C_F(X)$. There is an integer ℓ such that $1 = \gcd(p,q) = kp + q\ell$. Then $\phi^{-qp} = X^{kp} = X^{1-q\ell}$, which tells us that $\phi^{-qp}X^{q\ell} = X$. But then $X = (\phi^{-p}X^{\ell})^q$, so X has a qth root as claimed.

What about sufficiency? Suppose that $\psi \in F$ is a qth root of X. (We note that such a ψ must be unique, by Corollary 3.2.2.1.) As before set $k = p^{-1} \bmod q$, then define $\theta_0 = \psi^{-k}\pi_{0,k} \in C_F(X)$. For j in the range $1 \le j < s$, define $\theta_j = \mathrm{id} \in C_F(X)$. Extend this to a map $\theta \in C_T(\beta)$ using Corollary 5.2.11.

To see that θ has order h we first turn to Equation (5.15), which now reads

$$\theta^h\big|_{D_0} = \Sigma_0^{-1}\psi^{-kq}X^k\Sigma_0 = \Sigma_0^{-1}X^{-k}X^k\Sigma_0 = \mathrm{id}_{D_0}\,.$$

The restriction $\theta^h\big|_{D_i}$ is the conjugate $(\theta^h\big|_{D_0})^{\theta^i}$, which is equal to id_{D_i} by the previous equation. This holds for all i, so we conclude that $\theta^h = \mathrm{id}$ and that the order of θ divides h. Since $\mathrm{rot}(\theta) = 1/h$, the order of θ must be exactly h. □

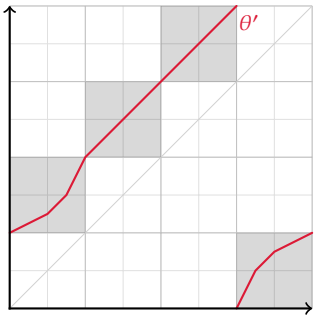

Figure 5.23: Let β be as in Figure 5.21. Then $X = x^2$. We have $p = k = 1$, $q = s = 2$ and $h = 4$, so we take ψ to be the unique qth root $\psi = x$ of X in F. Now $\pi_{0,k} = x_0 = x^2$, so we define

$$\theta_0 = \psi^{-k}\pi_{0,k} = x^{-1}x_0 = x$$
$$\theta_1 = \mathrm{id}$$

which forces

$$\theta_2 = x_0^{-1}\theta_0 x_1 = x^{-2}xx = \mathrm{id}$$
$$\theta_3 = x_1^{-1}\theta_1 x_2 = x^{-1}\,.$$

This defines $\theta \in C_F(x)$, which has rotation number $1/4$ and order 4.

Kassabov and Matucci (K&M) explain how to search for roots in F by reducing to a search for conjugators.[14] This means we can construct a splitting element θ of order h whenever it exists, or otherwise conclude that no such θ exists. If such a θ exists, we immediately conclude that $C_T(\beta) \cong C_F(X)^s \rtimes \langle\theta\rangle \cong C_F(X)^s \rtimes \mathbb{Z}_h$. To see which semidirect product this is, we need to understand how θ conjugates the kernel. This action is described by Equation (5.14). If we choose θ carefully, we will see that this action is that of a wreath product.

[14] Kassabov and Matucci 2012, Theorem 5.6.ii.

Proposition 5.2.15. *Assume the hypotheses of Lemma 5.2.13, and further assume that X has a qth root. Then $C_T(\beta)$ is a wreath product $C_F(X) \wr_s \mathbb{Z}_h$. The subscript s indicates that the wreath corresponds[15] to the action $\mathbb{Z}_s \circlearrowleft \mathbb{Z}_h$.*

[15] To be explicit, if $j \in \mathbb{Z}_s$ and $i \in \mathbb{Z}_h$, then the image of j under i is $j.i = j + i \pmod s$.

Proof. In the proof of Proposition 5.2.14, we saw that Sequence (5.11)

splits via some map θ if and only there exist maps $\theta_0, \dots, \theta_{s-1} \in C_F(X)$ satisfying

$$\theta_0 \cdots \theta_{s-1} \pi_{0,k}^{-1} = \phi\,,$$

for some fixed element $\phi \in F$ which is the unique qth root of X^{-k}. Now from Equation (5.14) we see that θ splits $C_T(\beta)$ into a wreath product if and only if each of the elements $\theta_0, \theta_1, \dots, \theta_{s-2}, \theta_{s-1}\pi_{0,k}^{-1}$ belongs to $Z(C_F(X))$, where $Z(\cdot)$ denotes the centre. Thus Sequence (5.11) splits as a wreath product if and only if

$$\begin{aligned}&\text{There exist elements } \theta_0, \dots, \theta_{s-2} \in Z(C_F(X)) \text{ and there}\\ &\text{exists } \theta_{s-1} \in C_F(X) \text{ such that } \theta_{s-1}\pi_{0,k}^{-1} \in Z(C_F(X)) \text{ and}\\ &\theta_0 \cdots \theta_{s-2}\theta_{s-1}\pi_{0,k}^{-1} = \phi.\end{aligned} \tag{5.16}$$

As ϕ is a root of a power of X, it follows from Lemma 4.1.7 that $C_F(\phi) = C_F(X)$; hence $\phi \in Z(C_F(\phi)) = Z(C_F(X))$. Additionally, $\pi_{0,k} \in C_F(X)$ because it is a product of factors $x_i \in C_F(X)$. This means we can satisfy Condition (5.16) by choosing elements

$$\theta_0 = \phi\,, \quad \theta_1 = \theta_2 = \cdots = \theta_{s-2} = \mathrm{id} \quad \text{and} \quad \theta_{s-1} = \pi_{0,k}\,.$$

Hence the extension of $\{\theta_0, \dots, \theta_{s-1}\}$ to a centraliser $\theta \in C_T(\beta)$ splits the centraliser as a wreath product. □

Corollary 5.2.16. *Let $\beta \in T$ be dyadic offset-similar. If Sequence (5.11) splits, then $C_T(\beta)$ is isomorphic to $(F^f \times \mathbb{Z}^z) \wr_s \mathbb{Z}_h$, for some integer parameters (f, z, h, s), with (f, z, h) as in Corollary 5.2.6 and with $s \geq 1$ dividing h.*

Moreover, for every such 4-tuple there is a dyadic offset-similar element β whose centraliser extension splits with these parameters.

Proof. $C_T(\beta)$ must be of the form $(F^f \times \mathbb{Z}^z) \wr_s \mathbb{Z}_h$ by Proposition 5.2.15. Here $h > 0$ is some integer divisible by $s > 0$. The left-hand factor is (isomorphic to) the centraliser $C_F(X)$, where $\mathrm{id} \neq X \in F$ can be obtained from a block form for β. The bounds on f and z are necessary by Claim 4.1.8. As in Corollary 5.2.6, we must have $(f, z) \neq (1, 0)$ because $X \neq \mathrm{id}$. This establishes the necessity of our conditions on the parameters.

To demonstrate sufficiency, we'll explain how to construct an element $\beta \in T$ whose centraliser has given parameters (f, z, h, s). Let $\alpha \in T$ be a dyadic self-similar element whose centraliser has parameters (f, z, h). (We can construct an explicit example α using Corollary 5.2.6.) Then α has a maximal block form $(x; \Sigma_i)$ with exactly h blocks. Define β to be the element with offset block form $(x_i; \Sigma_i; 1; q)$, where $h = qs$ and $x_i = x$ for all $0 \leq i \leq h$. In this case, the product $X = \prod_{0 \leq m \leq q} x_{i+mps}$ is $X = x^q$.

We claim that this is a maximal block form for β. Seeking a contradiction, suppose not. Then there is a block form for β with $h' > h$ blocks. We can use this to construct an element $\theta \in C_T(\beta)$ with

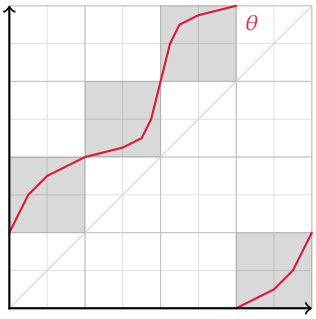

Figure 5.24: Continuing from Figure 5.23, ϕ is the unique qth root $\phi = x^{-1}$ of $X^{-k} = x^{-2}$ in F. We define a different centraliser θ' by

$$\theta'_0 = \phi = x^{-1}$$
$$\theta'_1 = \pi_{0,k} = x^2$$

which forces

$$\theta'_2 = x_0^{-1}\theta'_0 x_1 = x^{-2}x^{-1}x = x^{-2}$$
$$\theta'_3 = x_1^{-1}\theta'_1 x_2 = x^{-1}x^2\,\mathrm{id} = x\,.$$

This defines $\theta' \in C_F(x)$, which has rotation number 1/4, order 4 and splits $C_T(\beta)$ into a wreath product.

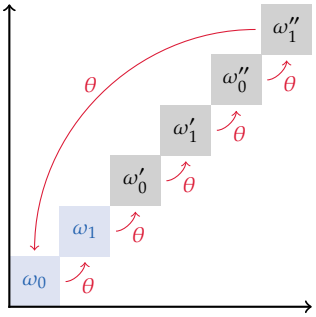

Figure 5.25: A schematic illustrating the centraliser $C_T(\beta)$, with parameters $h = 6, q = 3$ and $s = 2$. In contrast to Figure 5.17, kernel elements $\omega \in C_T(\beta)_0$ are determined by their behaviour on a representative arc Q, which consists of $s = 2$ blocks.

In this example, this behaviour is described by the pair $(\omega_0, \omega_1) \in C_F(X)^2$. The behaviour of ω on translates of Q is then given by related elements $\omega'_0, \omega''_0, \omega'_1, \omega''_1 \in C_F(X)$. Conjugating by θ permutes the blocks one step forward. If we choose θ carefully, then ω^θ is described by the pair (ω_1, ω_0), and $C_T(\beta)$ has a wreath product structure.

rotation number $1/h'$ (see Corollary 5.2.11). Observe that

$$\beta^q\big|_{D_i} = \Sigma_i^{-1} X \Sigma_i = \Sigma_i^{-1} x^q \Sigma_i = \alpha^q\big|_{D_i} \ ;$$

since this holds for all i, we see that $\beta^q = \alpha^q$. Then

$$\theta \in C_T(\beta) \leq C_T(\beta^q) = C_T(\alpha^q) = C_T(\alpha)\,,$$

where the last equality is due to Lemma 4.1.7. We see that $\theta \in C_T(\alpha)$ has rotation number $1/h'$—but this contradicts the fact α has a maximal block form with h blocks. Hence our block form for β is maximal.

To finish the proof, we note that $X = x^q$ has a qth root in F. So $C_T(\beta)$ splits as a wreath product $C_F(X) \wr_s \mathbb{Z}_h$, by Proposition 5.2.15. But $C_F(X) = C_F(x^q) = C_F(x)$ by Lemma 4.1.7, and $C_F(x) \cong F^f \times \mathbb{Z}^z$ by our choice of element α. □

To conclude this case, we briefly discuss what happens when $C_T(\beta)$ *doesn't* split—as in Example 5.1.6. We see that we can choose $\theta \in C_T(\beta)$ with rotation number $1/h$ inducing a wreath-like action on the kernel.

Proposition 5.2.17. *Assume the hypotheses of Lemma 5.2.13, and further assume that X* does not *have a qth root. Then Sequence (5.11) does not split as a wreath expansion. However, we can choose a preimage θ of $1 \in \mathbb{Z}_h$ whose conjugation action cyclically shifts the s factors of the kernel.*

Proof. The fact that Sequence (5.11) does not split was established by Proposition 5.2.14.

In order to cyclically permute the factors in Equation (5.14), we need to find elements $\theta_0, \ldots, \theta_{s-2} \in Z(C_F(X))$ and $\theta_{s-1} \in C_F(X)$ such that $\theta_{s-1}\pi_{0,k}^{-1} \in Z(C_F(X))$. There certainly is such a collection of elements: for instance, take

$$\theta_0 = \cdots = \theta_{s-2} = \mathrm{id} \qquad \text{and} \qquad \theta_{s-1} = \pi_{0,k}\,.$$

As in the second half of Proposition 5.2.14's proof, extend the maps $\theta_0, \ldots, \theta_{s-1}$ to a map $\theta \in C_T(\beta)$ with rotation number $1/h$.

For completeness, we note that θ has infinite order. Using Equation (5.15) we see that

$$\theta^h\big|_{D_0} = \Sigma_0^{-1} \left(\theta_0 \cdots \theta_{s-1}\pi_{0,k}^{-1}\right)^q X^k \Sigma_0 = \Sigma_0^{-1} X^k \Sigma_0 = \beta^q\big|_{D_0}\ ,$$

where the last equality follows from to Notation 5.2.9. Moreover

$$\theta^h\big|_{D_i} = (\theta^h\big|_{D_0})^{\theta_i} = (\beta^q\big|_{D_0})^{\theta_i} = \beta^q\big|_{D_i}$$

for each i, so $\theta^h = \beta^q$. As β is nontorsion, we see that θ is also nontorsion. □

Corollary 5.2.18. *Let $\beta \in T$ be dyadic offset-similar. If Sequence (5.11) does* not *split, then $C_T(\beta)$ is an extension $(F^f \times \mathbb{Z}^z)^s \hookrightarrow C_T(\beta) \twoheadrightarrow \mathbb{Z}_h$, for some integer parameters (f, z, s, h) as in Corollary 5.2.16 with $q = h/s \geq 2$.*

Moreover, for every such 4*-tuple there is a dyadic offset-similar element* β *whose centraliser extension does not split and has these parameters.*

Proof. The necessary conditions on (f, z, s, h) from Corollary 5.2.16 also apply here. We must have $q \geq 2$: otherwise $q = 1$, and X will have a qth root, so Sequence (5.11) will split by Proposition 5.2.14.

To build an element with such a centraliser, let $\alpha \in T$ be a dyadic self-similar element whose centraliser has parameters (f, z, h). (Again, Corollary 5.2.6 gives an explicit construction.) Then α has a maximal block form $(X; \Sigma_i)$ with h blocks. Without loss of generality we may assume that X is a root element: if not, we can use K&M's work[16] to find a minimal root X' of X. Then the element α' with block form $(X'; \Sigma_i)$ has the same centraliser parameters as α, because $C_F(X) = C_F(X')$ by Theorem 4.1.3.

[16] Kassabov and Matucci 2012, Theorem 5.6.ii.

Use this "rootless block form" to define a dyadic offset-similar element $\beta \in T$ with block form $(x_i; \Sigma_i; 1, q)$. Here the elements $x_i \in F$ are

$$x_0 = \cdots = x_{s-1} = X \qquad \text{and} \qquad x_s = \cdots = x_{h-1} = \mathrm{id}\ .$$

The products $\prod_{0 \leq m < q} x_{i+mps}$ are all equal to $X \neq 1$ by design. This means that

$$\beta^q\big|_{D_i} = \Sigma_i^{-1} X \Sigma_i = \alpha\big|_{D_i}$$

for each i, and so $\beta^q = \alpha$. We learn that $C_T(\beta) \subseteq C_T(\beta^q) = C_T(\alpha)$. This means our block form for β is maximal: if there was a block form with $h' > h$ blocks, we could create an element with rotation number $1/h'$ in $C_T(\beta)$ and hence in $C_T(\alpha)$—but this is impossible. This means that $C_T(\beta)$ is an extension $C_F(X)^s \hookrightarrow C_T(\beta) \twoheadrightarrow \mathbb{Z}_h$. As X is a root element, the extension is nonsplit by Proposition 5.2.14. □

5.2.3 *Case 3: nondyadic self-similar elements*

In this section we mimic our analysis for dyadic self-similar elements, adapting to elements α which only have nondyadic important points. Many of our calculations are identical; others are completely new, needed to ensure we are building elements of T from PL_2 maps defined between nondyadic points.

Definition 5.2.19. Let $d_0 \to d_1 \to \cdots \to d_{h-1} \to d_0$ be a circular partition with $h \geq 1$ points in which no d_i is dyadic. Define circular arcs $D_i = [d_i, d_{i+1}]$ with subscripts mod h, for each $0 \leq i < h$. Each D_i is called a *block*. For $0 \leq i < h$, let $\Sigma_i \colon D_0 \to D_i$ be a PL_2 map; similarly let $x \colon D_0 \to D_0$ be a PL_2 map. We impose the condition that

For dyadic self-similar elements there was a natural choice of interval to use as the domain of Σ_i, namely $I = [0, 1]$. This is because any two intervals with dyadic endpoints have isomorphic PL_2 groups see (Claim 1.2.8). There's no obvious choice to use in place of I in the nondyadic case, so we simply use D_0 in place of $I = [0, 1]$.

$$\text{there is a constant } c \text{ such that } d_0^+ \Sigma'_{i+1} = c \cdot d_1^- \Sigma'_i, \text{ for all } i \tag{5.17}$$

on our collection of maps $\{\Sigma_i\}$. We also require that[17]

[17] This condition is equivalent to $x \neq \mathrm{id}$.

$$1 \neq d_0^+ x' = d_1^- x' \tag{5.18}$$

and that x has no interior dyadic important points.

In terms of these ingredients, we define an element $\alpha \in T$ by the recipe

$$t\alpha = \begin{cases} t x^{\Sigma_i} & \text{if } t \in D_i. \end{cases}$$

Such an element α is called *nondyadic self-similar* (NSS). The data $(x; \{\Sigma_i\})$ is called a *block form* for α with h blocks. If it is not possible to express α in block form with $h' > h$ blocks, we say that a block form with h blocks is *maximal*.

Bookkeeping. We need to explain why our recipe produces a well-defined element α in T. The only time two blocks overlap is at $\{d_{i+1}\} = D_i \cap D_{i+1}$. This isn't a problem, because

$$\begin{aligned} d_{i+1}\Sigma_i^{-1} x \Sigma_i = d_1 x \Sigma_i = d_1 \Sigma_i = d_{i+1} \\ \text{and} \quad d_{i+1}\Sigma_{i+1}^{-1} x \Sigma_{i+1} = d_0 x \Sigma_{i+1} = d_0 \Sigma_{i+1} = d_{i+1}\,; \end{aligned}$$

we see that α is well-defined.

For each i, the restriction $\alpha\big|_{D_i}$ is a product $\Sigma_i^{-1} x \Sigma_i$ of maps in the groupoid $\mathcal{PL}_2$. So $\alpha\big|_{D_i} \in \mathrm{PL}_2(D_i)$, which means the internal breakpoints of $\alpha\big|_{D_i}$ are all dyadic. The last thing to check is that no point d_{i+1} between two blocks is a breakpoint. The left derivative is

$$d_{i+1}^-(\Sigma_i^{-1} x \Sigma_i)' = [d_{i+1}\Sigma_i^{-1}]^- x' = d_1^- x'\,,$$

where the first equality is due to Claim 2.3.14.[18] Similarly, the right derivative is

$$d_{i+1}^+(\Sigma_{i+1}^{-1} x \Sigma_{i+1})' = [d_{i+1}\Sigma_{i+1}^{-1}]^+ x' = d_0^+ x'\,.$$

These two derivatives are equal thanks to Condition (5.18). This establishes that $\alpha \in T$ (rather than $\mathrm{PL}_{\mathbb{Q}}(S^1)$).

[18] Apply the claim with (d_1, x, Σ_i) in place of the claim's (x, α, γ).

As in previous cases, the point of defining a block form is to make it easy to write down certain centralising elements. The details of how to do so are largely the same as in the dyadic self-similar case.

Lemma 5.2.20. *Let α be a nondyadic self-similar element, with block form $(x; \Sigma_i)$ consisting of h blocks. Define an element $\delta \in T$ by the formula $\delta\big|_{D_i} = \Sigma_i^{-1}\Sigma_{i+1}$, for each i with subscripts modulo h. Then δ commutes with α, has rotation number $\mathrm{rot}(\delta) = 1/h$ and has order h.*

Proof. The calculations follow exactly as in Lemma 5.2.2. However, we need to justify why $\delta \in T$, since we have introduced blocks with nondyadic endpoints. Again, we cannot have any nondyadic breakpoints within the interior of D_i, because the internal breakpoints of $\Sigma_i^{-1}\Sigma_{i+1} \in \mathrm{PL}_2(D_i, D_{i+1})$ are all dyadic (see also Claim 1.2.5). So we need to check that gradients do not change at the points d_{i+1} between blocks. The left gradient is

$$d_{i+1}^-(\Sigma_i^{-1}\Sigma_{i+1})' = d_{i+1}^-(\Sigma_i^{-1})' \cdot d_1^-\Sigma_{i+1}' = \frac{d_1^-\Sigma_{i+1}'}{d_1^-\Sigma_i'}\,,$$

whereas the right gradient is

$$d_{i+1}^{+}(\Sigma_{i+1}^{-1}\Sigma_{i+2})' = d_{i+1}^{+}(\Sigma_{i+1}^{-1})' \cdot d_0^{+}\Sigma_{i+2}' = \frac{d_0^{+}\Sigma_{i+2}'}{d_0^{+}\Sigma_{i+1}'}\,.$$

These two fractions are equal due to Condition (5.17). Hence d_{i+1} is not a breakpoint of δ for any i, meaning that $\delta \in T$. □

As before, the next step is to characterise the nondyadic self-similar elements of T. To do so, we describe a process which produces maximal block forms. Once again, the ideas are similar to the dyadic self-similar case; the new material in this proof is checking the location of breakpoints.

Lemma 5.2.21. *An element $\alpha \in T$ is nondyadic self-similar if and only if α has rotation number 0 and has $n \geq 1$ important points, all of which are nondyadic. We can algorithmically determine a maximal block form for any such α.*

Proof. First suppose that α is nondyadic self-similar, with block form $(x; \Sigma_i)$. As α fixes each $D_i = [d_i, d_{i+1}]$ setwise, we have that $d_i\alpha = d_i$ and so $\mathrm{rot}(\alpha) = 0$. The one-sided gradients

$$d_{i+1}^{+}\alpha' = d_{i+1}^{+}(\Sigma_{i+1}^{-1}x\Sigma_{i+1})' = d_0^{+}x'$$

and

$$d_{i+1}^{-}\alpha' = d_{i+1}^{-}(\Sigma_{i}^{-1}x\Sigma_{i})' = d_1^{-}x'$$

are both equal to a number distinct from 1, by Condition (5.18) on x. Hence $d_i \in \partial\,\mathrm{Fix}(\alpha) = I_\alpha$, showing that I_α is nonempty. We need to check that all points in $r \in I_\alpha$ are nondyadic. If $r \in D_i^\circ$ then $r \in \partial\,\mathrm{Fix}(x^{\Sigma_i}) = \partial\,\mathrm{Fix}(x)\Sigma_i = I_x\Sigma_i$, where I_x are the important points of x. Since these are nondyadic and Σ_i maps nondyadics to nondyadics, we see that r is nondyadic. Otherwise r is the left endpoint d_i of some block D_i, so is nondyadic by definition of our block form.

In the other direction, we now suppose that α has the properties listed in the statement of the lemma. The construction of a block form is very similar to that given in the proof of Lemma 5.2.3. Enumerate the important points I_α; find an element $\gamma \in C_T(\alpha)$ with minimal index difference $0 < d \leq n$; then define the h blocks D_i exactly as before. To define x and Σ_i we don't need to worry about using a map Ω to conjugate back to $\mathrm{PL}_2(I) = F$; instead we simply define

$$x = \alpha\big|_{D_0} \qquad \text{and} \qquad \Sigma_i = \gamma^i\big|_{D_0}\,.$$

Does this data satisfy the definition of a block form? By construction, an endpoint $d \in \partial D_i$ is a important point of α, so is nondyadic. Each Σ_i is a restriction of the PL_2 map γ, so are themselves PL_2. We also note that x cannot have any (interior) dyadic important points—as then they would be important points of α.

We have two conditions on gradients to verify. Using the chain rule

once more, we see that the Σ maps satisfy

$$\begin{aligned} d_0^+\Sigma'_{i+1} &= d_0^+(\gamma^{i+1})' \\ &= d_0^+\gamma' \cdot [d_0\gamma]^+(\gamma^i)' \\ &= d_0^+\gamma' \cdot d_1^+(\gamma^i)' \\ &= d_0^+\gamma' \cdot d_1^-(\gamma^i)' \\ &= d_0^+\gamma' \cdot d_1^-\Sigma'_i\,, \end{aligned}$$

The fourth equality holds because d_1 is not dyadic, so is not a breakpoint of γ^i.

which establishes Condition (5.17). We also have a condition on the gradients of x to establish. We calculate that

$$d_0^+x' = d_0^+\alpha' = [d_0\gamma]^+(\alpha^\gamma)' = d_1^+\alpha' = d_1^-\alpha' = d_1^-x'\,;$$

We've made use of Claim 2.3.14 here in the second equality. Again, d_1 is not a breakpoint of α, so the left and right derivatives of α at d_1 are equal.

as d_0 is nondyadic we have $d_0^-\alpha' = d_0^+\alpha'$. This gradient cannot be 1 or else d_0 would not be a important point of α, contrary to the choice of blocks D_i. Thus $d_0^+x' \neq 1$, establishing that Condition (5.18) holds. Finally, the calculation

$$x^{\Sigma_i} = (\alpha\big|_{D_0})^{\gamma^i\big|_{D_0}} = \alpha^{(\gamma^i)}\Big|_{D_i} = \alpha\big|_{D_i}$$

verifies that α can be built using x and the $\{\Sigma_i\}$.

Why is the block form $(x;\Sigma_i)$ maximal? If there were a larger block form with $h' > h$ blocks, we could again construct an element $\delta' \in C_T(\alpha)$ with order h', this time using Lemma 5.2.20. This would contradict our choice of γ and the minimality of d. □

Proposition 5.2.22. *Let $\alpha \in T$ be nondyadic self-similar with a maximal block form $(x;\Sigma_i)$ containing h blocks. Then $C_T(\alpha)$ is an extension*

$$C_{\mathrm{PL}_2(D_0)}(x) \hookrightarrow C_T(\alpha) \twoheadrightarrow \mathbb{Z}_h\,,$$

which splits as the direct product $\mathbb{Z} \times \mathbb{Z}_h$.

Proof. The kernel of Sequence (4.3) is $C_T(\alpha)_0$. This is infinite cyclic by Corollary 4.1.6, generated by a minimal root ω of α with rotation number $\mathrm{rot}(\omega) = 0$. By the same result, we can take the kernel to be

$$C_{\mathrm{PL}_2(D_0)}(\alpha\big|_{D_0}) = C_{\mathrm{PL}_2(D_0)^{\Sigma_i}}(x^{\Sigma_i}) \cong C_{\mathrm{PL}_2(D_0)}(x)$$

up to isomorphism.

Lemma 4.1.7 establishes that $C_T(\omega) = C_T(\alpha)$, so $\omega \in Z(C_T(\omega)) = Z(C_T(\alpha))$, where $Z(\cdot)$ denotes the centre. Because our block form is maximal, we know the quotient of the extension is $\mathbb{Z}_h$, so the element δ (of order h) constructed in Lemma 5.2.20 can be used to split the extension. We thus have a split central extension $\mathbb{Z} \hookrightarrow C_T(\alpha) \twoheadrightarrow \mathbb{Z}_h$, meaning that $C_T(\alpha)$ is the direct product $\mathbb{Z} \times \mathbb{Z}_h \cong \langle\omega\rangle \times \langle\delta\rangle$. □

It is worth highlighting the structure of elements $\epsilon \in C_T(\alpha)_0$ in the kernel of the extension for $C_T(\alpha)$. We saw that these are powers of a minimal root $\omega \in C_T(\alpha)_0$; say $\omega^n = \alpha$. Then ω must be given by

$\omega\big|_{D_i} = \Sigma_i^{-1} w \Sigma_i$, where w is the nth root of x in $\mathrm{PL}_2(D_0)$. Since $\epsilon = \omega^k$ for some k, we see that

$$\epsilon\big|_{D_i} = \Sigma_i^{-1} w^k \Sigma_i, \quad \text{for each } i\,. \tag{5.19}$$

We conclude our study of nondyadic self-similar elements by asking 'which centralisers are possible?'

Corollary 5.2.23. *For every integer $h \geq 1$, there is a nondyadic self-similar element α whose centraliser is isomorphic to $\mathbb{Z} \times \mathbb{Z}_h$.*

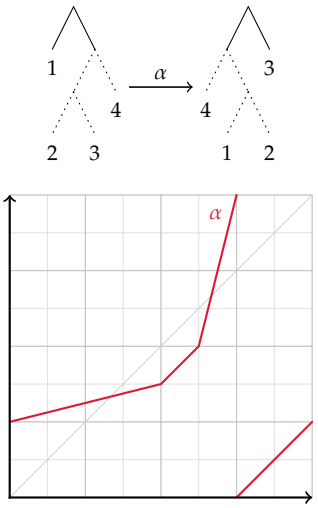

Figure 5.26: The element $\alpha \in T$ with domain and range trees (D, R), as shown in Figure 2.22.

Proof. Let T_h be a binary tree with exactly h leaves. Let their addresses be $t_0, \ldots, t_{h-1}$ in lexicographic order. Also let D and R be the domain and range trees of the element α from Figure 2.22. This element is nondyadic self-similar, because $I_\alpha = \{1/3, 2/3\}$ is fixed pointwise by α. We use α to create a related element α_h with more important points below. Note that our approach does not construct a block form directly.

Define a new tree D_h by gluing copies of D below each leaf of T_h; similarly define R_h by gluing copies of R below the leaves of T_h. To be explicit, the leaves of D_h have addresses

$$d_{4i+0} = t_i\mathtt{0}, \quad d_{4i+1} = t_i\mathtt{100}, \quad d_{4i+2} = t_i\mathtt{101} \quad \text{and} \quad d_{4i+3} = t_i\mathtt{11}$$

for $0 \leq i < h$. Similarly, R_h is the tree whose leaves have addresses

$$r_{4i+0} = t_i\mathtt{00}, \quad r_{4i+1} = t_i\mathtt{010}, \quad r_{4i+2} = t_i\mathtt{011} \quad \text{and} \quad r_{4i+3} = t_i\mathtt{1}\,,$$

again for $0 \leq i < h$. We define an element $\alpha_h \in T$ using the tree pair (D_h, σ_h, R_h), where the leaf bijection σ_h maps d_k to $r_{k+1 \bmod 4h}$. We claim that α_h is nondyadic self-similar and has centraliser isomorphic to $\mathbb{Z} \times \mathbb{Z}_h$.

To establish the first subclaim, we use the characterisation given in Lemma 5.2.21. From the observation

$$d_{4i}\alpha_h = r_{4i+1} = t_i\mathtt{010} = d_{4i}\mathtt{10}\,, \tag{5.20}$$

it follows that the infinite string $d_{4i}\overline{\mathtt{10}}$ is fixed by α_h; all other infinite strings with prefix d_{4i} are not fixed by α_h. Thus $\mathrm{rot}(\alpha_h) = 0$ and so the important points are $I(\alpha_h) = \partial\,\mathrm{Fix}(\alpha_h)$. We also observe that

$$\begin{aligned} d_{4i+1}\alpha_h &= r_{4i+2} = t_i\mathtt{011} = d_{4i}\mathtt{11} \\ r_{4i+3}\mathtt{01}\alpha_h &= d_{4i+2}\alpha_h = r_{4i+3} \\ d_{4i+3}\alpha_h &= r_{4(i+1)} = t_{i+1}\mathtt{00} = d_{4(i+1)}\mathtt{0}\,, \end{aligned} \tag{5.21}$$

which shows that the only other fixed points of α_h are those of the form $r_{4i+3}\overline{\mathtt{01}}$. Altogether, we see that

$$I(\alpha_h) = \partial\,\mathrm{Fix}(\alpha_h) = \mathrm{Fix}(\alpha_h) = \{\, d_{4i}\overline{\mathtt{10}}, r_{4i+3}\overline{\mathtt{01}} \mid 0 \leq i < h \,\}$$

is a discrete set of size $2h$. Enumerated in circular order, its points are

the infinite strings

$$I(\alpha_h)\colon\ d_0\overline{\mathtt{10}} \to r_3\overline{\mathtt{01}} \to d_4\overline{\mathtt{10}} \to r_7\overline{\mathtt{01}} \to \dots \\ \dots \to d_{4(h-1)}\overline{\mathtt{10}} \to r_{4(h-1)+3}\overline{\mathtt{01}} \to d_0\overline{\mathtt{10}}\,.$$

Because the tails of these infinite strings are neither $\bar{\mathtt{0}}$ nor $\bar{\mathtt{1}}$, each corresponds to a nondyadic point on the circle. Hence α_h is nondyadic self-similar.

We know that $C_T(\alpha_h)\big|_{I(\alpha_h)}$ is a cyclic group of size $k \geq 0$. To conclude this proof, we need to show that $k = h$. Let $\delta \in T$ be the element with tree pair (T_h, ρ, T_h), with leaf bijection given by $t_i\rho = t_{i+1 \bmod h}$. (In other words, δ cycles the leaves of T_h forward one step.) We see that δ has order h, because the same is true of ρ. From the calculations

$$d_{4i}\alpha\delta = r_{4i+1}\delta = t_i\mathtt{010}\delta = t_{i+1}\mathtt{010} = r_{4(i+1)+1}$$

and

$$d_{4i}\delta\alpha = t_i\mathtt{0}\delta\alpha = t_{i+1}\mathtt{0}\alpha = d_{4(i+1)}\alpha = r_{4(i+1)+1}\,,$$

we learn that $\alpha\delta$ and $\delta\alpha$ agree below d_{4i}. Similar calculations hold below the remaining d_{4i+j}, so we conclude that $\delta \in C_T(\alpha)$.

We note that δ cycles the important points forward with index difference 2; for instance, $d_0\overline{\mathtt{10}}\delta = d_4\overline{\mathtt{10}}$. Now there cannot be an element $\epsilon \in C_T(\alpha)$ with index difference 1. If so, ϵ would map $x = d_0\overline{\mathtt{10}}$ to $y = r_3\overline{\mathtt{01}}$. Then by Corollary 2.3.15 the gradients $x^+\alpha'_h$ and $y^+\alpha'_h$ should be equal. As $x \in [\![d_0]\!]$, we see from Equation (5.20) that the former gradient is 2^{-2}. On the other hand, $y \in [\![r_{4i+3}\mathtt{01}]\!]$; we see from Equation (5.21) that the gradient here is 2^{+2}. These gradients are different—but this is a contradiction. Thus the minimal index difference is 2, so the centraliser quotient has size $k = 2h/2 = h$. □

5.2.4 *Case 4: nondyadic, offset-similar elements*

For our last case, we adapt Section 5.2.2 to the case where our elements have only nondyadic important points.

Definition 5.2.24. Let $\{d_i\}_{i=1}^h$, $\{D_i\}_{i=1}^h$ and $\{\Sigma_i\colon D_0 \to D_i\}_{i=1}^h$ be as in Definition 5.2.19, including Condition (5.17). Next, let q be a divisor of h, say $h = sq$. Select an integer $0 < p < q$ which is coprime to q. For each $0 \leq i < h$, select maps $x_i \in \mathrm{PL}_2(D_0)$ subject to the constraint that $X = \prod_{0\leq j<q} x_{i+jps}$ is independent of i, has initial gradient $d_0X' \neq 1$ and has no internal dyadic important points. We also require that the gradients of the $\{x_i\}$ satisfy

$$d_0^+x'_{i+1} = d_1^-x'_i \qquad \text{for each } i. \tag{5.22}$$

From these ingredients, we define an element $\beta \in T$ by the recipe $\beta\big|_{D_i} = \Sigma_i^{-1}x_i\Sigma_{i+ps}$. Any element of this form is said to be *nondyadic offset-similar* (NOS). The data tuple $(\{x_i\};\{\Sigma_i\};p;q)$ is called a *block form* for β, with h blocks.

Bookkeeping. We need to explain why our recipe produces a well-defined element β in T. The only place two blocks overlap is at $\{d_{i+1}\} = D_i \cap D_{i+1}$. This isn't a problem, because the expressions

Again we are using Claim 2.3.14 here.

$$\begin{aligned} d_{i+1}\Sigma_{i+1}^{-1}x_{i+1}\Sigma_{i+ps+1} &= d_0 x_{i+1}\Sigma_{i+ps+1} = d_0\Sigma_{i+ps+1} = d_{i+ps+1} \\ \text{and} \qquad d_{i+1}\Sigma_i^{-1}x_i\Sigma_{i+ps} &= d_1 x_i\Sigma_{i+ps} = d_1\Sigma_{i+ps} = d_{i+ps+1} \end{aligned} \tag{5.23}$$

are equal—thus β is well-defined.

For each i, the restriction $\beta\big|_{D_i}$ is a product $\Sigma_i^{-1}x_i\Sigma_{i+ps}$ of maps in the groupoid $\mathcal{PL}_2$. So $\beta\big|_{D_i} \in \mathrm{PL}_2(D_i, D_{i+ps})$, which means the internal breakpoints of $\beta\big|_{D_i}$ are all dyadic. The last thing to check is that no point d_{i+1} between two blocks is a breakpoint. The left derivative is

$$\begin{aligned} d_{i+1}^-(\Sigma_i^{-1}x_i\Sigma_{i+ps})' &= d_{i+1}^-(\Sigma_i^{-1})' \cdot d_1^- x_i' \cdot d_1^-\Sigma_{i+ps}' \\ &= \frac{d_1^-\Sigma_{i+ps}'}{d_1^-\Sigma_i'} \cdot d_1^- x_i'; \end{aligned}$$

similarly the right derivative is

$$\begin{aligned} d_{i+1}^+(\Sigma_{i+1}^{-1}x_{i+1}\Sigma_{i+ps+1})' &= d_{i+1}^+(\Sigma_{i+1}^{-1})' \cdot d_0^+ x_{i+1}' \cdot d_0^+\Sigma_{i+ps+1}' \\ &= \frac{d_0^+\Sigma_{i+ps+1}'}{d_0^+\Sigma_{i+1}'} \cdot d_0^+ x_{i+1}'. \end{aligned}$$

These two quantities are equal thanks to Conditions (5.18) and (5.22). This establishes that $\beta \in T$ (rather than $\mathrm{PL}_{\mathbb{Q}}(S^1)$).

Remark 5.2.25. The calculation

$$\begin{aligned} d_0^+ X' = d_0^+\left(\prod_{0\le j<q} x_{1+jps}\right)' &= \prod_{0\le j<q} d_0^+ x_{1+jps}' \\ = \prod_{0\le j<q} d_1^- x_{jps}' = d_1^-\left(\prod_{0\le j<q} x_{jps}\right)' &= d_1^- X' \end{aligned}$$

shows that the initial and final gradients of X are equal (and distinct from 1).

Equation (5.7) still applies to the nondyadic case, so again each $x_i \in C_{\mathrm{PL}_2(D_0)}(X)$.

Does this block form allow us to write down centralising elements? The answer is a little different than in previous cases: thanks to the restrictions which follow from having nondyadic important points, it turns out for a centralising elements to transitively permute the blocks $\{D_i\}$, we must have a very specific block form.

Lemma 5.2.26. *Let $\beta \in T$ be nondyadic offset-similar with block form $(x_i; \Sigma_i; p; q)$ and let X be the product $X = \prod_{0\le j<q} x_{i+jps}$. If there exists $\theta \in C_T(\beta)$ with rotation number $1/h$, then θ must take the form $\theta\big|_{D_i} = \Sigma_i^{-1}\theta_0\Sigma_{i+1}$ for some $\theta_0 \in C_{\mathrm{PL}_2(D_0)}(X)$, and all the $\{x_i\}$ must be equal.*

Proof. Suppose that $\theta \in C_T(\beta)$ with rotation number $1/h$. As in

the proof of Lemma 5.2.10, define $\theta_i = \Sigma_i\, \theta\big|_{D_i}\, \Sigma_{i+1}^{-1} \in \mathrm{PL}_2(D_0)$ for each i. Again as before, the fact that $\theta \in C_T(\beta^q)$ establishes that $\theta_i \in C_{\mathrm{PL}_2(D_0)}(X)$ for each i. Now $X \neq 1$ has no internal dyadic important points by the definition of our block form. So X is an almost one-bump function in $\mathrm{PL}_2(D_0)$. Thus its its centraliser $C_{\mathrm{PL}_2(D_0)}(X)$ is infinite cyclic, generated by a minimal root ω of X.

Say that $\omega^n = X$ and write $D_0 = [d_0, d_1]$. From Remark 5.2.25 and the fact that $d_0^+ X' \neq 1$, we learn that

$$1 \neq (d_0^+\omega')^n = d_0^+(\omega^n)' = d_0^+ X' = d_1^- X' = d_1^-(\omega^n)' = (d_1^-\omega')^n\,,$$

and hence $1 \neq d_0^+\omega' = d_1^-\omega'$. Since each θ_i is a power of ω, we have $d_0^+\theta_i' = d_1^-\theta_i'$. Now each d_{i+1} is nondyadic, so $d_{i+1}^-\theta' = d_{i+1}^+\theta'$. The left derivative is

$$\begin{aligned} d_{i+1}^-\theta' &= d_{i+1}^-(\Sigma_i^{-1})' \cdot d_1^-\theta_i' \cdot d_1^-\Sigma_{i+1}' \\ &= \frac{d_1^-\Sigma_{i+1}'}{d_1^-\Sigma_i'} \cdot d_1^-\theta_i'\,, \end{aligned} \tag{5.24}$$

but the right derivative is

$$\begin{aligned} d_{i+1}^+\theta' &= d_{i+1}^+(\Sigma_{i+1}^{-1})' \cdot d_0^+\theta_{i+1}' \cdot d_0^+\Sigma_{i+2}' \\ &= \frac{d_0^+\Sigma_{i+2}'}{d_0^+(\Sigma_{i+1}^{-1})'} \cdot d_0^+\theta_{i+1}'\,. \end{aligned} \tag{5.25}$$

The two fractions are equal thanks to Condition (5.17), so we learn that $d_1^-\theta_i' = d_0^+\theta_{i+1}'$. Altogether, we see that each θ_i has the same initial (and final) gradient. In turn, this means that each θ_i is the same power of ω, and hence each θ_i is equal to θ_0.

Because $\theta \in C_T(\beta)$ we can again deduce that $x_i\theta_{i+ps} = \theta_i x_{i+1}$ (Condition (5.8) from before). But since each $\theta_j = \theta_0$, this reads $x_i^{\theta_0} = x_{i+1}$. All three elements in this equation belong to $C_{\mathrm{PL}_2(D_0)}(X)$, but this group is isomorphic to $\mathbb{Z}$, so is abelian. Thus $x_i = x_{i+1}$ for each i. □

This slightly surprising result demonstrates the strength of the condition 'β has only nondyadic important points.' Since we are using block forms to study and construct centralising elements, we need to specialise to this particular kind of block form.

Definition 5.2.27. A block form $(x_i; \Sigma_i; p; q)$ for a nondyadic offset-similar element β is called *constant* if $x_i = x_0$ for every i. A constant block form for β with h blocks is called *maximal* if there does not exist a constant block form for β with $h' > h$ blocks.

With this in mind, we consider the converse of Lemma 5.2.26.

Lemma 5.2.28. *Let β have constant block form $(x_0; \Sigma_i; p; q)$. Then any $\theta_0 \in C_{\mathrm{PL}_2(D_0)}(x_0)$ extends to an element $\theta \in C_T(\beta)$ given by $\theta\big|_{D_i} = \Sigma_i^{-1}\theta_0\Sigma_{i+1}$.*

Proof. Calculations similar to Equation (5.23) establish that

$$d_{i+1}\Sigma_i^{-1}x_0\Sigma_{i+1} = d_{i+2} = d_{i+1}\Sigma_{i+1}^{-1}x_0\Sigma_{i+2}\,,$$

from which we see that θ is well-defined.

To check that $\theta \in T$, we need to check that

$$d_{i+1}^-\theta' = \frac{d_1^-\Sigma'_{i+1}}{d_1^-\Sigma'_i}\cdot d_1^-\theta'_0 \qquad \text{(as in (5.24))}$$

is equal to

$$d_0^+\theta' = \frac{d_0^+\Sigma'_{i+2}}{d_0^+(\Sigma_{i+1}^{-1})'}\cdot d_0^+\theta'_0\,. \qquad \text{(as in (5.25))}$$

The two fractions are again equal due to Condition (5.17). Now $X = x_0^q$ because we have a *constant* block form, so $C_{\mathrm{PL}_2(D_0)}(X) = C_{\mathrm{PL}_2(D_0)}(x_0)$. Hence $\theta_0 \in C_{\mathrm{PL}_2(D_0)}(x_0)$ is a power of a root of X, and so must have the same initial and final gradient (because X has the same property). Thus $\theta \in T$.

Finally we see that

$$(\theta\beta)\big|_{D_i} = \Sigma_i^{-1}\theta_0\Sigma_{i+1}\Sigma_{i+1}^{-1}x_0\Sigma_{i+1+ps} = \Sigma_i^{-1}\theta_0 x_0\Sigma_{i+1+ps}$$

is equal to

$$(\beta\theta)\big|_{D_i} = \Sigma_i^{-1}x_0\Sigma_{i+ps}\Sigma_{i+ps}^{-1}\theta_0\Sigma_{i+ps+1} = \Sigma_i^{-1}x_0\theta_0\Sigma_{i+ps+1}$$

because $\theta_0 \in C_F(x_0)$. Thus $\theta \in C_T(\beta)$. □

OUR NEXT STEP (as in previous cases) is to characterise these elements β and explain how to product maximal constant block forms.

Lemma 5.2.29. *An element $\beta \in T$ is nondyadic offset-similar if and only β has nonzero rotation number* $\mathrm{rot}(\beta) = p/q$ *and has $n \geq 1$ important points, all of which are nondyadic. We can algorithmically determine a maximal* constant *block form for any such β.*

Proof. First assume that β has a (not necessarily constant) block form $(x_i; \Sigma_i; p; q)$ with h blocks. We will find that $\mathrm{rot}(\beta) = p/q \neq 0$ by arguing as in Claim 5.2.8. The whole important point set is $I_\beta = \partial\,\mathrm{Fix}(\beta^q)$. Again, $\beta^q\big|_{D_i} = \Sigma_i^{-1}X\Sigma_i$, so $\mathrm{Fix}(\beta^q) \cap D_i = \mathrm{Fix}(X^{\Sigma_i})$. This means that the important points within the interior D_i° are

$$\begin{aligned} I_\beta \cap D_i^\circ &= \partial\,\mathrm{Fix}(X^{\Sigma_i}) \cap D_i^\circ \\ &= (\partial\,\mathrm{Fix}(X)\Sigma_i) \cap (D_0^\circ\Sigma_i) \\ &= [\partial\,\mathrm{Fix}(X) \cap D_0^\circ]\Sigma_i\,. \end{aligned}$$

The set in square brackets contains no dyadics by definition of X, and PL_2 maps send nondyadics to nondyadics. So there are no dyadic important points in D_i°.

This covers all points on the circle except the $\{d_i\}$, which are nondyadic—showing that I_β is entirely nondyadic. To see that I_β is

nonempty, we will show that each d_i is a important point of β. Certainly each d_i is fixed by β^q; the right-hand gradient at this point is

$$d_i^+(\beta^q)' = d_i^+(X^{\Sigma_i})' = d_0^+ X' \neq 1\,,$$

which shows that $d_i \in \partial\,\mathrm{Fix}(\beta^q)$.

In the other direction, we take an element β with the given properties and describe an algorithm to produce a maximal constant block form for β. As in the proof of Lemma 5.2.12, find an element $\gamma \in C_T(\beta)$ with minimal index difference $0 < d \leq n$, where $dh = n$ for some h. Let the reduced rotation number of β be $p/q \neq 0$, so that $h = qs$ for some integer s. Partition the important points $I_\beta = \{\, r_0 \to \cdots \to r_{n-1} \to r_0 \,\}$ into blocks $D_i = [r_{id}, r_{(i+1)d}] = [d_i, d_{i+1}]$, for each $0 \leq d < h$. Now β^q is nondyadic *self*-similar and commutes with γ, so β^q has a block form $(X; \Sigma_i)$ with h blocks. From the definition of this block form, we know that the $\{\Sigma_i\}$ satisfy Condition (5.17) and that $1 \neq d_0^+ X' = d_1^- X'$ (from Condition (5.18)). We may impose without loss of generality that the codomain of Σ_i is D_i. Then $\beta^q\big|_{D_i} = \Sigma_i^{-1} X \Sigma_i$ for each i.

Define $\delta \in C_T(\beta^q)$ by $\delta\big|_{D_i} = \Sigma_i^{-1}\Sigma_{i+1}$, this time using Lemma 5.2.20. This time, the proof of Proposition 5.2.22 tells us that $C_T(\beta^q)$ contains the direct product $C_T(\beta^q)_0 \times \langle\delta\rangle$. Now β belongs to this group, so we can write $\beta = \epsilon\delta^{ps}$ for some $\epsilon \in C_T(\beta^q)_0$. As noted in Equation (5.19), ϵ must take the form $\epsilon\big|_{D_i} = w^k$ for some k, where $w \in \mathrm{PL}_2(D_0)$ is a minimal root of X. Then

$$\beta\big|_{D_i} = \epsilon\big|_{D_i}\, \delta^p s\big|_{D_i} = (\Sigma_i^{-1} w^k \Sigma_i)(\Sigma_i^{-1}\Sigma_{i+ps}) = \Sigma_i^{-1} w^k \Sigma_{i+ps}\,.$$

The product $\prod_{0\leq j<q} x_i$ is $(w^k)^q = X$, which meets all the requirements specified in the definition of a block form. Finally, $d_0^+ X' = d_1^- X'$ implies that $d_0^+ w' = d_1^- w'$, and hence $d_0^+(w^k)' = d_1^-(w^k)'$. This establishes Condition (5.22).

All in all, we see that β has a constant block form $(w^k; \Sigma_i; p; q)$ with h blocks. If there were a constant block form with $h' > h$ blocks, we could construct a centralising element with rotation number $0 < 1/h' < 1/h$ using Lemma 5.2.28. This would contradict the minimality of d, so our block form above must be maximal. □

Once again a maximal (constant) block form for β allow us to determine the extension structure of $C_T(\beta)$.

Proposition 5.2.30. *Let $\beta \in T$ be nondyadic offset-similar with maximal constant block form $(x_0; \Sigma_i; p; q)$. Then $C_T(\beta)$ is a central extension*

$$C_{\mathrm{PL}_2(D_0)}(x_0) \hookrightarrow C_T(\beta) \twoheadrightarrow \mathbb{Z}_h$$

which splits as the direct product $\mathbb{Z} \times \mathbb{Z}_h$.

Proof. We use $Q = D_0 \cup \cdots \cup D_{s-1}$ as our representative arc. For some integer k, we know that $C_T(\beta)$ is an extension $C_{\mathrm{PL}_2(Q)}(\beta^q\big|_Q) \hookrightarrow C_T(\beta) \twoheadrightarrow \mathbb{Z}_k$ (from Sequence (4.3)). We must have $k = h$ because our

constant block form is maximal. Because x_0 has no internal dyadic important points and ∂D_i is nondyadic, $\beta^q\big|_Q$ is an almost one-bump function and thus has infinite cyclic centraliser. By the argument preceding Equation (4.1), $C_{\mathrm{PL}_2(Q)}(\beta^q\big|_Q)$ is isomorphic to

$$C_{\mathrm{PL}_2(D_0)}(\beta^q\big|_{D_0}) = C_{\mathrm{PL}_2(D_0^{\Sigma_0})}(x_0^{\Sigma_0}) \cong C_{\mathrm{PL}_2(D_0)}(x_0)\,.$$

We saw in Lemma 5.2.28 that elements $\theta \in C_T(\beta)$ with rotation number $1/h$ take the form $\theta\big|_{D_i} = \Sigma_i^{-1}\theta_0\theta_{i+1}$, for some (any) element $\theta_0 \in C_{\mathrm{PL}_2(D_0)}(x_0)$. Its hth power is given by

$$\theta^h\big|_{D_i} = \prod_{0\le j<h} \theta\big|_{D_{i+j}} = \prod_{0\le j<h} \Sigma_{i+j}^{-1}\theta_0\Sigma_{i+j+1} = \Sigma_i^{-1}\theta_0^h\Sigma_i\,,$$

so if we take $\theta_0 = \mathrm{id}_{D_0}$ then $\theta^h = \mathrm{id}$. This means that θ splits the extension.

We could compute the conjugate of a kernel element by θ to complete our study of the extension. Instead, let $\alpha = \beta^q$. Then α is nondyadic self-similar, with block form $(x_0^q;\Sigma_i)$. Now $C_T(\beta) \le C_T(\beta^q) = C_T(\alpha)$. We proved in Proposition 5.2.22 that the latter group is abelian, so $C_T(\beta)$ must be abelian too. This means that our split extension must be a direct product $\mathbb{Z}\times\mathbb{Z}_h$. □

The last piece of the puzzle is to see which of these potential centraliser types actually exist.

Corollary 5.2.31. *For every integer $h \ge 2$, there is a nondyadic offset-similar element β whose centraliser is isomorphic to $\mathbb{Z}\times\mathbb{Z}_h$.*

Notice that the condition $h \ge 2$ is necessary, because $h \ge q \ne 1$.

Proof. Let α be a nondyadic self-similar element with a maximal block form $(x;\Sigma_i)$ containing h blocks. (Such an element exists by Corollary 5.2.23.) Define β to be the element with constant block form $(x;\Sigma_i;1;h)$. This constant block form must be maximal: if not, there is some constant block form with $h' > h$ blocks. Then we can construct a centraliser $\theta \in C_T(\beta)$ with rotation number $1/h'$, by Lemma 5.2.28. Now β^q has block form $(x^q;\Sigma_i)$, so $\beta^q = \alpha^q$. Because α has rotation number zero, we know from Lemma 4.1.7 that $C_T(\alpha) = C_T(\alpha^q)$. Thus $\theta \in C_T(\alpha)$ with rotation number $1/h' < 1/h$, which contradicts the maximality of our block form for α.

From our maximal constant block form for β, we use Proposition 5.2.30 to immediately conclude that $C_T(\beta) \cong \mathbb{Z}\times\mathbb{Z}_h$. □

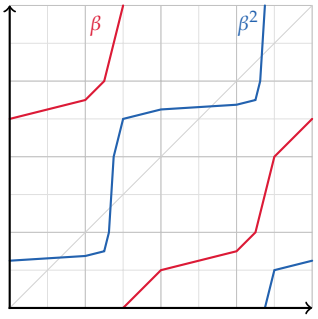

Figure 5.27: The element $\beta \in T$ from Figure 2.22 has only nondyadic important points $I_\beta = \{1/6, 1/3, 2/3, 5/6\}$. It is constructed by following the recipe in this proof, where $h = 2$ and $\alpha = \alpha_2$ is the element constructed in Corollary 5.2.23. Also shown is the graph of the square β^2.

5.3 Summary

We conclude with a summary of this chapter's findings; see also Figure 5.28.

Having analysed each of the four cases in turn, we now have a description of an arbitrary nontorsion element $\alpha \in T$. The description

allows us to determine the structure of $C_T(\alpha)$.

Theorem 5.3.1. *Every nontorsion element $\alpha \in T$ belongs to exactly one of the types DSS, DOS, NSS or NOS. The particular type of α is determined by the rotation number* $\mathrm{rot}(\alpha)$ *and the important points I_α.*

We can construct a maximal block form for any such element α. Such a block form allows us to express the centraliser $C_T(\alpha)$ as a group extension and determine its parameters.

Justification. We broke our analysis into four cases at the start of Section 5.2. Lemmas 5.2.3, 5.2.12, 5.2.21 and 5.2.29 explained how to construct maximal block forms; then Propositions 5.2.4, 5.2.15, 5.2.17, 5.2.22 and 5.2.30 showed how we can deduce the centraliser's structure from these block forms. □

We saw that one case was an odd-one out: that of dyadic offset-similar elements. Unlike the other three cases, centralisers of DOS elements are not guaranteed to split. We were able to identify exactly when this centraliser splits. We give a neater description here, which doesn't require us to find a block form.

Theorem 5.3.2. *Let $\beta \in T$ be nontorsion. Then the sequence*

$$C_T(\beta)_0 \hookrightarrow C_T(\beta) \twoheadrightarrow C_T(\beta)\big|_{I(\beta)}$$

is nonsplit if and only if

- $\mathrm{rot}(\beta) = p/q \neq 0$ *in lowest terms,*
- *β has a dyadic important point, and*
- *there does not exist a qth root ω of β^q with* $\mathrm{rot}(\omega) = 0$.

Proof. The centraliser extension always split in the DSS, NSS and NOS cases by Propositions 5.2.4, 5.2.22 and 5.2.30. This establishes that the first two bullets are necessary for the extension to be nonsplit. Proposition 5.2.14 tells us how to drill down to a necessary and sufficient condition: an element β with block form $(x_i; \Sigma_i; p; q)$ has nonsplit centraliser extension if and only if $X = \prod_{0 \leq j < q} x_{i+jps}$ does not have a qth root $w \in F$.

To see that this condition is equivalent to the third bullet, suppose such a $w \in F$ exists. Then the (self-similar) block form $(w; \Sigma_i)$ describes an element $\omega \in T$ whose qth power has block form $(X; \Sigma_i)$—so $\omega^q = \beta^q$. Conversely, a qth root ω of β^q with rotation number zero must have block form $(w; \Sigma_i)$ for some $w \in F$. Then the restriction of the equation $\omega^q = \beta^q$ to D_i reads $\Sigma_i^{-1} w^q \Sigma_i = \Sigma_i^{-1} X \Sigma_i$, and so $w^q = X$. □

Lastly, we managed to determine which parameters f, z, s and h are actually used by centralisers. Informally, any reasonable tuple of parameters are parameters for some element $\alpha \in T$.

Corollary 5.3.3. *Choose one of the four types of nontorsion elements DSS, DOS, NSS and NOS. We can construct an element $\alpha \in T$ of this type whose centraliser has the following parameters if and only if all the requirements below hold.*

- *The parameter h must be a positive integer. For DOS and NOS elements, h must be at least 2.*
- *For DOS elements, the parameter s must be a divisor of h. Otherwise there is no s parameter.*
- *For DSS and DOS elements, the parameters f and z must be integers not both zero satisfying[19] $0 \leq f \leq z+1$. Otherwise, there is no f and no z parameter.*

[19] i.e. (f, z) must the parameters for a centraliser in F: see Claim 4.1.8.

Justification. This was explained in Corollaries 5.2.6, 5.2.16, 5.2.18, 5.2.23 and 5.2.31. □

case	DSS	DOS	NSS	NOS
important points	some dyadic		all nondyadic	
rotation number	zero	nonzero	zero	nonzero
max. block forms	Lemma 5.2.3	5.2.12	5.2.21	5.2.29
centraliser structure	$(\mathbb{Z}^z \times F^f) \wr_h \mathbb{Z}_h$	either $(\mathbb{Z}^z \times F^f) \wr_s \mathbb{Z}_h$, or else nonsplit $(\mathbb{Z}^z \times F^f)^s \hookrightarrow \cdot \twoheadrightarrow \mathbb{Z}_h$	$\mathbb{Z} \times \mathbb{Z}_h$	$\mathbb{Z} \times \mathbb{Z}_h$
deduced in	Proposition 5.2.4	5.2.15 and 5.2.17	5.2.22	5.2.30
	$0 \leq f \leq z+1$, with f and z not both zero		no f, no z	no f, no z
parameters	$h \geq 1$	$h \geq 2$	$h \geq 1$	$h \geq 2$
	no s	s divides h	no s	no s
explained in	Corollary 5.2.6	5.2.16 and 5.2.18	5.2.23	5.2.31

Figure 5.28: A summary of our studies into nontorsion centralisers in T.

A *Deferred calculations*

Let L be the interval $[0, 1/2]$. In Example 5.1.6, we claimed[1] that the centraliser $C_{\mathrm{PL}_2(L)}(\delta^2|_L)$ was equal to $\langle \delta^2|_L \rangle$. We justify this now in more detail.

* * *

The map in question $\Delta = \delta^2|_L$ is explicitly given by

$$t\Delta = \begin{cases} 0 + 1/4\,(t-0) & \text{if} \quad 0 \le t < 1/8 \\ 1/32 + 1/8\,(t-1/8) & \text{if} \quad 1/8 \le t < 1/4 \\ 3/64 + 1/4\,(t-1/4) & \text{if} \quad 1/4 \le t < 5/16 \\ 1/16 + 2\,(t-5/16) & \text{if} \quad 5/16 \le t < 3/8 \\ 3/16 + 1\,(t-3/8) & \text{if} \quad 3/8 \le t < 7/16 \\ 1/4 + 4\,(t-7/16) & \text{if} \quad 7/16 \le t < 1/2. \end{cases}$$

Figure 5.7 shows that Δ is a one-bump function, so its centraliser is infinite cyclic, generated by a minimal root of Δ. The initial gradient of Δ is $1/4$, so we only need to look for a (square) root ρ with initial gradient $1/2$. We run the stair algorithm[2] with this initial gradient, and see if it produces a conjugator ρ between Δ and itself. If so, then $C_{\mathrm{PL}_2(L)}(\Delta) = \langle \rho \rangle$; otherwise $C_{\mathrm{PL}_2(L)}(\Delta) = \langle \Delta \rangle$.

The initial linearity box for this setup is $[0, 1/8]^2$, because Δ has no breakpoints internal to $I = [0, 1/8]$. The final linearity box is $[7/16, 1/2]^2$, because Δ has no breakpoints internal to $F = [7/16, 1/2]$. Let ρ be the candidate conjugator we're building. All we know for now is that

$$t\rho = \begin{cases} 1/2 t & \text{if} \quad 0 \le t < 1/8. \end{cases}$$

If there exists a conjugator ρ then $\rho = \Delta\rho\Delta^{-1}$, which implies that $\rho|_{I\Delta^{-1}} = \Delta|_{I\Delta^{-1}}\, \rho|_I\, \Delta|_{I\rho}$. Using computer-assisted calculations[3] to evaluate the right-hand side of this equation, we learn that

$$t\rho = \begin{cases} 1/16 + 1/4\,(t-1/8) & \text{if} \quad 1/8 \le t < 1/4 \\ 3/32 + 1/2\,(t-1/4) & \text{if} \quad 1/4 \le t < 5/16 \\ 1/8 + 8\,(t-5/16) & \text{if} \quad 5/16 \le t < 21/64 \\ 1/4 + 4\,(t-21/64) & \text{if} \quad 21/64 \le t < 11/32. \end{cases}$$

Now set $J = I\rho\Delta = [0, 11/32]$ and use the previous observation with J in place of I. Again with the help of a computer, we learn that[4]

$$t\rho = \begin{cases} 5/16 + 1/4\,(t-11/32) & \text{if} \quad 11/32 \le t < 3/8 \\ 41/128 + 1/8\,(t-3/8) & \text{if} \quad 3/8 \le t < 7/16. \end{cases}$$

Since $7/16$ is the left endpoint of our final linearity box we're done:

[1] Strictly speaking, in Example 5.1.6 we had an element $X = x_0 x_0^{\phi} \in F$, and claimed that $C_F(X) = \langle X \rangle$. Conjugate both sides of this equation by $\Sigma : I \to L$ to see that this equation is equivalent to

$$C_{\mathrm{PL}_2(L)}(\delta^2|_L) = \langle \delta^2|_L \rangle\,,$$

which we address in this appendix.

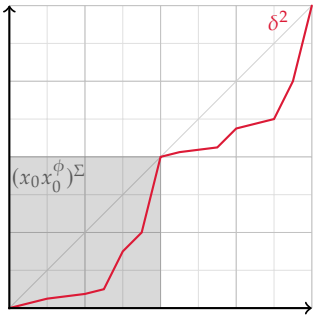

Figure A.1: The element δ^2 from Figure 5.7

[2] Kassabov and Matucci 2012, Section 4.2.

[3] Robertson 2015, See Appendix A.1 below for the script which produced these calculations.

[4] The calculation determines ρ on four more linear segments on $7/16 \le t \le 59/128$; we omit these for brevity. But notice that these segments are all in the final linearity box F^2; this is another way of seeing that no conjugator $\rho \in F$ exists.

if ρ really is a conjugator, it must have one last linear segment

$$t\rho = \left\{ {}^{21}/_{64} + M(t - {}^{7}/_{16}) \quad \text{if} \quad {}^{7}/_{16} \leq t \leq {}^{1}/_{2}, \right.$$

where $M = ({}^{43}/_{64})/({}^{1}/_{8})$. This is not an integer power of 2, so there is no conjugator $\rho \in \mathrm{PL}_2(L)$ between Δ and itself with initial gradient ${}^{1}/_{2}$. Hence $C_{\mathrm{PL}_2(L)}(\Delta) = \langle \Delta \rangle$.

A.1 Script source code

For completeness, we include the script used to run the computer calculations presented above. It uses a library written by the author, which can be obtained with

```
git clone --branch plmaps https://github.com/DMRobertson/thompsons_v.git
```

The bulk of the library is documented at `https://thompsons-v.readthedocs.io`, with experimental features briefly described at `https://thompsons-v.readthedocs.io/en/plmaps/plmaps.html`.

* * *

```
#!  /usr/bin/env python3
from thompson import *
from plmaps import *
from plmaps.util import glue

# First define our ingredients.
x0_std = PL2.from_aut(standard_generator())
phi = PL2.from_aut(Automorphism.from_dfs("110100100", "111001000"))

SigmaL = PL2([0, 1], [0, 1/2])
SigmaR = PL2([0, 1], [1/2, 1])
rho = CPL2.rotation(1/2)

# Use the ingredients to construct δ and δ^2.
def build_delta():
    gamma = glue(x0_std^SigmaL, x0_std^(phi * SigmaR), cls=CPL2)
    delta = gamma * rho
    return delta

delta = build_delta()
d2 = delta * delta
plot(delta, d2)

# Δ = δ^2|_L in this appendix is called d2_rest here.
d2_rest = d2.restriction(0, 1/2)
print(d2_rest.format_pl_segments(LaTeX=True, sfrac=True))
plot(d2_rest)

# Run Kassabov-Matucci staircase algorithm to find a generator for C_{PL_2(L)}(Δ).
gen = d2_rest.one_bump_cent_gen(verbose=True)
assert gen == d2_rest
```

B *To split, or not to split?*

These are quick notes which reconsider the element δ from Example 5.1.6. Let's set the scene. The relevant plots are reproduced in Figures B.1 and B.2.

- x_0 is the standard generator of F.
- $L = [0, 1/2]$ and $R = [1/2, 1]$ are subintervals of $I = [0, 1]$.
- $\Sigma\colon I \to L$ and $\Sigma'\colon I \to R$ are the PL_2 maps which linearly shrink I down to L and R respectively.
- $\phi \in F$ is a fixed element such that $\phi \notin C_F(x_0)$.
- $\rho_{1/2}$ is the rotation through half a turn of the circle. For brevity, we'll write $\rho = \rho_{1/2}$ in this note.

Define an element $\delta \in T$ casewise by

$$t\delta = \begin{cases} t x_0^{\Sigma} \rho & \text{if } t \in L \\ t (x_0^{\phi})^{\Sigma'} \rho & \text{if } t \in R. \end{cases}$$

The important points of δ are $\{0, 1/2\}$. These are permuted transitively by δ, and so we have a short exact sequence

$$C_T(\delta)_0 \hookrightarrow C_T(\delta) \twoheadrightarrow \mathbb{Z}_2\,,$$

as argued earlier.

Claim. *This sequence does not split.*

Proof. To prove this, we'll need to demonstrate that there does not exist an element $\mu \in C_T(\delta)$ of order 2. Seeking a contradiction, suppose that such a μ exists. What can we deduce about μ? We know that μ must permute the important points of δ. First suppose that $0\mu = 0$, so that $\mu \in F$. Now F is torsion-free, so the order of μ is either 1 or ∞; the order cannot be 2. Thus we learn that $0\mu = 1/2$, which means μ takes the form

$$t\mu = \begin{cases} t \mu_L^{\Sigma} \rho & \text{if } t \in L \\ t \mu_R^{\Sigma'} \rho & \text{if } t \in R \end{cases}$$

for some maps[1] $\mu_L, \mu_R \in F$. A direct computation shows that $\mu^2\big|_L$ is the product $\mu^2\big|_L = \mu_L^{\Sigma} \rho \mu_R^{\Sigma'} \rho$. Now since $\Sigma'\rho = \Sigma$ and $\rho = \rho^{-1}$, we can simplify this to

$$\mu^2\big|_L = \mu_L^{\Sigma} \mu_R^{\Sigma} = (\mu_L \mu_R)^{\Sigma}\,.$$

Because $\mu^2 = \mathrm{id}$ we must then have $(\mu_L \mu_R)^{\Sigma} = \mathrm{id}_L$. It follows that

$$\mu_R = \mu_L^{-1}\,. \tag{B.1}$$

[1] Explicitly, $\mu_L = (\mu\big|_L \rho)^{\Sigma^{-1}}$ and $\mu_R = (\mu\big|_R \rho)^{\Sigma'^{-1}}$.

Next, since $\mu \in C_T(\delta)$ we have $\mu\delta = \delta\mu$. Restricting this equation to L, we learn that

$$\mu_L^\Sigma \rho\, x_0^{\phi\Sigma'} \rho = x_0^\Sigma \rho\, \mu_R^{\Sigma'} \rho\,.$$

Again because $\Sigma'\rho = \Sigma$ and $\rho = \rho^{-1}$, this equation reads

$$\mu_L^\Sigma x_0^{\phi\Sigma} = x_0^\Sigma \mu_R^\Sigma\,, \qquad \text{i.e.} \qquad (\mu_L x_0^\phi)^\Sigma = (x_0\mu_R)^\Sigma\,.$$

Now conjugate by Σ^{-1} and use Equation (B.1) to conclude that

$$\mu_L x_0^\phi = x_0 \mu_L^{-1}\,. \tag{B.2}$$

This is an equation in the group F which contains just one unknown $\mu_L \in F$. Differentiate both sides and evaluate at 0 to learn that

$$0^+\mu_L' \cdot 0^+x_0' = \frac{0^+x_0'}{0^+\mu_L'}\,;$$

This implies that $(0^+\mu_L')^2 = 1$. As this is an equality in the gradient group $G = \{2^n \mid n \in \mathbb{Z}\}$, we learn that μ_L has initial gradient

$$0^+\mu_L' = 1\,. \tag{B.3}$$

Now if $\mu \in C_T(\delta)$ then certainly $\mu \in C_T(\delta^2)$. The square is given by

$$t\delta^2 = \begin{cases} t(x_0x_0^\phi)^\Sigma & \text{if } t \in L \\ t(x_0^\phi x_0)^{\Sigma'} & \text{if } t \in R. \end{cases}$$

Restricting the equality $\mu\delta^2 = \delta^2\mu$ to L yields

$$\mu_L^\Sigma \rho\, (x_0^\phi x_0)^{\Sigma'} = (x_0x_0^\phi)^\Sigma \mu_L^\Sigma \rho\,.$$

Bring the rightmost ρ over the left-hand side. Then—using the relationship between Σ, Σ' and ρ as before—we obtain

$$\mu_L^\Sigma (x_0^\phi x_0)^\Sigma = (x_0x_0^\phi)^\Sigma \mu_L^\Sigma\,,$$

and so we learn that the equality $\mu_L x_0^\phi x_0 = x_0 x_0^\phi \mu_L$ holds in F. But this reads

$$x_0^\phi x_0 = (x_0x_0^\phi)^{\mu_L}\,. \tag{B.4}$$

Certainly $\mu_L = x_0$ is a solution to Equation (B.4). However, the initial gradient of x_0 is $1/2$, so $\mu_L = x_0$ does not satisfy Equation (B.3). The full set of solutions to Equation (B.4) is the coset $C_F(x_0x_0^\phi)x_0$. We need to see if this coset contains an element with initial gradient 1. To do so, we need to determine the centraliser $C_F(x_0x_0^\phi)$.

The element $x_0x_0^\phi$ being centralised belongs to $\mathrm{PL}_2^<(I)$, because x_0 and x_0^ϕ also belong to $\mathrm{PL}_2^<(I)$. We know from Kassabov-Mattuci's work that $C_F(x_0x_0^\phi)$ generated by a minimal ('rootiest') root of $x_0x_0^\phi$. In Appendix A, we considered the one-bump function $\Delta = (x_0x_0^\phi)^\Sigma \in \mathrm{PL}_2(L)$. We used the stair algorithm to argue that $C_{\mathrm{PL}_2(L)}(\Delta) = \langle\Delta\rangle$,

because Δ was its own root element. Conjugating by Σ^{-1}, we learn that $C_F(x_0 x_0^{\phi}) = \langle x_0 x_0^{\phi} \rangle$.

This last fact means that $\mu_L \in C_F(x_0 x_0^{\phi}) x_0$ takes the form

$$\mu_L = (x_0 x_0^{\phi})^n x_0 \,,$$

for some $n \in \mathbb{Z}$. The initial gradients of $x_0 x_0^{\phi}$ and x_0 are $1/4$ and $1/2$ respectively. Hence

$$0^+\mu_L' = \frac{1}{4^n} \cdot \frac{1}{2} = \frac{1}{2^{1+2n}} \,.$$

This initial gradient must be equal to 1, by Equation (B.3). But that would mean that $1 + 2n = 0$, and there is no integer n satisfying this equation.

This contradiction leads us to conclude that there is no element $\mu_L \in F$ satisfying the neccessary conditions Equations (B.2), (B.3) and (B.4). Hence there does not exist an element $\mu \in C_T(\delta)$ of order 2. □

Remark B.0.1. Again, the argument boils down to the assertion that $X = x_0 x_0^{\phi}$ has no square root. We have no proof of this other than the computer calculations presented in Appendix A. However, the example could be adapted to work with any other element $X_{\text{alt}} \in F$ with initial gradient $0^+X_{\text{alt}}' = 1/4$ and no square root in F. Unfortunately, the author is not aware of such an element X_{alt} which *obviously* has no square root.

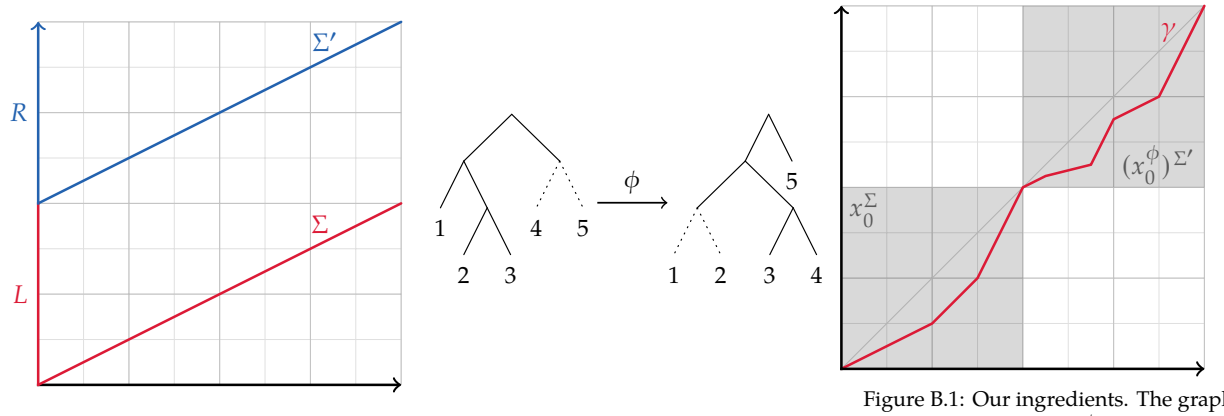

Figure B.1: Our ingredients. The graph of γ shows that $x_0 \neq x_0^\phi$, i.e. that $\phi \notin C_F(x_0)$.

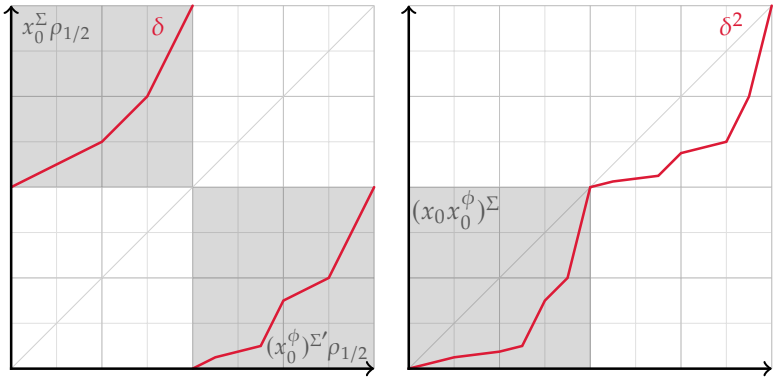

Figure B.2: The troublesome element δ and its square δ^2.

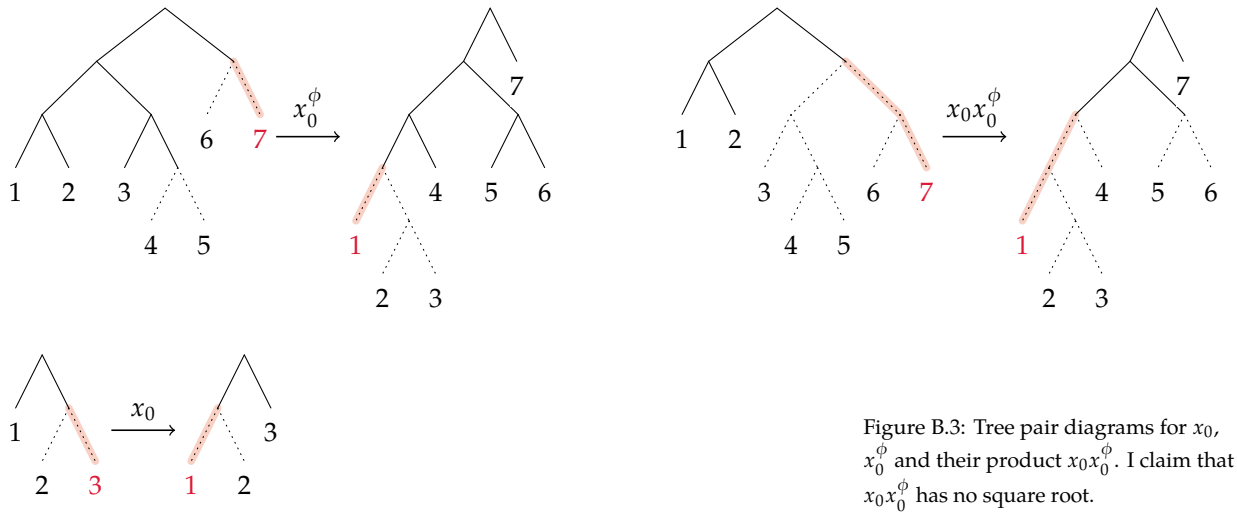

Figure B.3: Tree pair diagrams for x_0, x_0^ϕ and their product $x_0 x_0^\phi$. I claim that $x_0 x_0^\phi$ has no square root.

Bibliography

Aroca, Julio (2018). 'The conjugacy problem for Thompson-like groups'. In: arXiv: 1807.09503 [math.GR] (cit. on pp. 20, 53).

Baader, Franz and Tobias Nipkow (2008). *Term Rewriting and All That*. Cambridge: Cambridge University Press, 2008. 301 pp. ISBN: 978-1-139-17275-2. DOI: 10.1017/CBO9781139172752. MR: 1629216 (cit. on p. 6).

Barker, Nathan, Andrew J. Duncan and David M. Robertson (2016). 'The power conjugacy problem in Higman-Thompson groups'. In: *International Journal of Algebra and Computation* 26.2. ISSN: 1793-6500. DOI: 10.1142/S0218196716500144. MR: 3475063. arXiv: 1503.01032 [math.GR] (cit. on pp. xv, 53).

Belk, James and Bradley Forrest (2018). 'Rearrangement Groups of Fractals'. In: *Transactions of the American Mathematical Society*. ISSN: 0002-9947. DOI: 10.1090/tran/7386. arXiv: 1510.03133 [math.GR]. Forthcoming (cit. on p. xiii).

Belk, James and Francesco Matucci (2014). 'Conjugacy and dynamics in Thompson's groups'. In: *Geometriae Dedicata* 169 (1), pp. 239–261. ISSN: 0046-5755. DOI: 10.1007/s10711-013-9853-2. MR: 3175247. arXiv: 0708.4250 [math.GR] (cit. on pp. v, xv, 53).

Bieniecka, Ewa (2018). 'Commutativity and Free Products in Thompson's group V'. PhD thesis. University of St Andrews, 2018. URL: http://hdl.handle.net/10023/14652 (visited on 05/06/2018). Submitted (cit. on p. 71).

Bieri, Robert and Ralph Strebel (2016). *On Groups of PL-homeomorphisms of the Real Line*. Vol. 215. Mathematical Surveys and Monographs. Providence, Rhode Island: American Mathematical Society, 2016. 174 pp. ISBN: 978-1-4704-2901-0. MR: 3560537. arXiv: 1411.2868 [math.GR] (cit. on p. 12).

Bleak, Collin, Hannah Bowman et al. (2013). 'Centralizers in R. Thompson's group V_n'. In: *Groups, Geometry, and Dynamics* 7 (4), pp. 821–865. ISSN: 1661-7207. DOI: 10.4171/GGD/207. MR: 3134027. arXiv: 1107.0672 [math.GR] (cit. on pp. 29, 35, 36, 46, 71).

Bleak, Collin, Casey Donoven and Julius Jonušas (2017). 'Some isomorphism results for Thompson like groups $V_n(G)$'. In: *Israel Journal of Mathematics* 1, pp. 1–19. ISSN: 0021-2172. DOI: 10.1007/s11856-017-1580-5. MR: 3736497. arXiv: 1410.8726 [math.GR] (cit. on p. 20).

Bleak, Collin, Martin Kassabov and Francesco Matucci (2011). 'Structure Theorems for Groups of Homeomorphisms of the Circle'. In: *International Journal of Algebra and Computation* 20.6, pp. 105–121. ISSN: 1793-6500. DOI: 10.1142/S0218196711006571. MR: 2847521. arXiv: 0910.0218 [math.GR] (cit. on pp. 26, 28).

Bleak, Collin and Martyn Quick (2017). 'The infinite simple group V of Richard J. Thompson: presentations by permutations'. In: *Groups, Geometry and Dynamics* 11.4, pp. 1401–1436. ISSN: 1661-7207. DOI: 10.4171/GGD/433. MR: 3737287. arXiv: 1511.02123 [math.GR] (cit. on p. 11).

Bleak, Collin and Olga Salazar-Díaz (2013). 'Free Products in R. Thompson's Group V'. In: *Transactions of the American Mathematical Society* 365 (11), pp. 5967–5997. ISSN: 1088-6850. DOI: 10.1090/S0002-9947-2013-05823-0. MR: 3091272. arXiv: 0911.0979 [math.GR] (cit. on pp. xiv, 36).

Brin, Matthew G. (1996). 'The chameleon groups of Richard J. Thompson: automorphisms and dynamics'. In: *Publications Mathématiques de l'Institut des Hautes Études Scientifiques* 84.1, pp. 5–33. ISSN: 0073-8301. DOI: 10.1007/BF02698834. MR: 1441005. arXiv: math/9507205 [math.GR] (cit. on p. 1).

— (2004). 'Higher Dimensional Thompson Groups'. In: *Geometriae Dedicata* 108.1, pp. 163–192. ISSN: 0046-5755. DOI: 10.1007/s10711-004-8122-9. MR: 2112673. arXiv: math/0406046 [math.GR] (cit. on pp. xiv, 29, 31).

Brin, Matthew G. and Craig C. Squier (2001). 'Presentations, conjugacy, roots, and centralizers in groups of piecewise linear homeomorphisms of the real line'. In: *Communications in Algebra* 29.10, pp. 4557–4596. ISSN: 0092-7872. DOI: 10.1081/AGB-100106774. MR: 1855112 (cit. on pp. xiii, xv, 53).

Brown, Kenneth S. (1987). 'Finiteness properties of groups'. In: *Journal of Pure and Applied Algebra* 44.1–3: *Proceedings of the Northwestern conference on cohomology of groups*, pp. 45–75. ISSN: 0022-4049. DOI: 10.1016/0022-4049(87)90015-6. MR: 0885095 (cit. on p. xiii).

Brown, Ronald (1987). 'From Groups to Groupoids: a Brief Survey'. In: *Bulletin of the London Mathematical Society* 19.2, pp. 113–134. ISSN: 1469-2120. DOI: 10.1112/blms/19.2.113. MR: 0872125 (cit. on p. 16).

Burger, Marc and Shahar Mozes (1997). 'Finitely presented simple groups and products of trees'. In: *Comptes Rendus de l'Académie des Sciences, Série I* 324.7, pp. 747–752. ISSN: 0764-4442. DOI: 10.1016/S0764-4442(97)86938-8 (cit. on p. 1).

Burillo, José et al. (2009). 'Combinatorial and metric properties of Thompson's group T'. In: *Transactions of the American Mathematical Society* 361, pp. 631–652. ISSN: 1088-6850. DOI: 10.1090/S0002-9947-08-04381-X. MR: 2452818. arXiv: math/0503670 [math.GR] (cit. on pp. 28, 35, 59).

Calegari, Danny (2004). 'Circular groups, planar groups, and the Euler class'. In: *Geometry & Topology Monographs* 7: *Proceedings of the Casson Fest*. Ed. by Cameron Gordon and Yoav Rieck, pp. 431–491. ISSN: 1464-8997. DOI: 10.2140/gtm.2004.7.431. MR: 2172491. arXiv: math/0403311 [math.GT] (cit. on p. 8).

— (2007). 'Denominator bounds in Thompson-like groups and flows'. In: *Groups, Geometry, and Dynamics* 1.2, pp. 101–109. ISSN: 1661-7207. DOI: 10.4171/GGD/6. MR: 2319453. arXiv: math/0609573 [math.GR] (cit. on p. 28).

Cannon, J.W., W.J. Floyd and W.R. Parry (1996). 'Introductory notes on Richard Thompson's groups'. In: *L'Enseignement Mathématique* 42, pp. 215–256. ISSN: 0013-8584. DOI: 10.5169/seals-87877. MR: 1426438 (cit. on pp. 1–3, 10, 11, 22).

Cohen, Daniel E. (1989). *Combinatorial Group Theory. A Topological Approach*. London Mathematical Society Student Texts 14. Cambridge: Cambridge University Press, Aug. 1989. 324 pp. ISBN: 978-0-521-34936-9. DOI: 10.1017/CBO9780511565878. MR: 1020297 (cit. on p. 16).

Dehornoy, Patrick (2005). 'Geometric presentations for Thompson's groups'. In: *Journal of Pure and Applied Algebra* 203.1, pp. 1–44. ISSN: 0022-4049. DOI: 10.1016/j.jpaa.2005.02.012. MR: 21766501. arXiv: math/0407096 [math.GR] (cit. on pp. xiii, 2).

Diller, Jeffrey and Jan-Li Lin (2017). 'Rotation Numbers of Elements in Thompson's Group T'. In: *Topology Proceedings* 50, pp. 203–212. ISSN: 0146-4124. MR: 3573432. arXiv: 1609.00771 [math.DS] (cit. on p. 28).

Floyd, Robert W. and Richard Beigel (1994). *The Language of Machines. An Introduction to Computability and Formal Languages*. New York: W.H. Freeman and Company, 27th Jan. 1994. 706 pp. ISBN: 978-0-7167-8266-7 (cit. on p. 49).

Geoghegan, Ross and Marco Varisco (2017). 'On Thompson's group T and algebraic K-theory'. In: *Geometric and Cohomological Group Theory*. Ed. by Peter Kropholler et al. Vol. 444. London Mathematical Society Lecture Note Series. Cambridge University Press, Oct. 2017, pp. 34–45. ISBN: 978-1-316-77132-7. DOI: doi.org/10.1017/9781316771327.005. arXiv: 1401.0357 [math.GT] (cit. on pp. 11, 59, 72).

Ghys, Étienne and Vlad Sergiescu (1987). 'Sur un groupe remarquable de difféomorphisms du cercle'. French. In: *Commentarii Mathematici Helvetici* 62, pp. 185–239. ISSN: 0010-2571. MR: 0896095 (cit. on p. 28).

Gill, Nick and Ian Short (2013). 'Conjugacy in Thompson's group F''. In: *Proceedings of the American Mathematical Society* 141.5, pp. 1529–1538. ISSN: 0002-993. DOI: 10.1090/S0002-9939-2012-11432-7. MR: 3020840. arXiv: 0709.1987 [math.GR]. A seventh section is present on the arXiv but not in the published version. (Cit. on pp. xv, 53).

Grigorchuk, R.I., V.V. Nekrashevych and V.I. Sushchanskiĭ (2000). 'Automata, Dynamical Systems, and Groups'. Trans. from the Russian by I. Nikitin. In: *Proceedings of the Steklov Institute of Mathematics* 231, pp. 128–203. ISSN: 0081-5438. MR: 1841755 (cit. on pp. 19, 20).

Guba, Victor and Mark Sapir (1997). *Diagram groups*. Vol. 130. Memoirs of the American Mathematical Society 620. Providence, Rhode Island: American Mathematical Society, Nov. 1997. viii+117. ISBN: 978-0-8218-0639-5. DOI: 10.1090/memo/0620. MR: 1396957 (cit. on pp. 53, 71).

Guba, V.S. and M.V. Sapir (1999). 'On subgroups of R. Thompson's group F and other diagram groups'. Trans. from the Russian by J. Chadwick. In: *Sbornik: Mathematics* 190.8, pp. 1077–1130. ISSN: 1064-5616. DOI: 10.1070/sm1999v190n08ABEH000419. MR: 1725439. arXiv: math/0005114 [math.GR] (cit. on pp. 71, 76).

Higman, Graham (1974). *Finitely presented infinite simple groups*. Notes on pure mathematics 8. Canberra: Australian National University, 1974. 82 pp. ISBN: 978-0-7081-0300-5. MR: 0376874 (cit. on pp. xiii, 53, 71).

Hossain, Nabil et al. (2013). 'Deciding Conjugacy in Thompson's Group F in linear time'. In: *15th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing*. (Timişoara, Romania, Sept. 2013). IEEE, pp. 89–96. ISBN: 978-1-4799-3035-7. DOI: 10.1109/SYNASC.2013.19 (cit. on p. 53).

Johnson, D.L. (1997). *Presentations of Groups*. 2nd ed. London Mathematical Society Student Texts 15. Cambridge: Cambridge University Press, 1997. xii+216. ISBN: 978-0-521-58542-2. DOI: 10.1017/CBO9781139168410. MR: 1472735 (cit. on p. 80).

Kapovich, Ilya and Nadia Benakli (2002). 'Boundaries of hyperbolic groups'. In: *Combinatorial and Geometric Group Theory*. Contemporary Mathematics 296. Providence, Rhode Island: American Mathematical Society, 2002, pp. 39–93. ISBN: 978-0-8218-2822-9. DOI: 10.1090/conm/296. MR: 1921706. arXiv: math/0202286 [math.GR] (cit. on p. 17).

Kassabov, Martin and Francesco Matucci (2012). 'The simultaneous conjugacy problem in groups of piecewise linear functions'. In: *Groups, Geometry, and Dynamics* 6 (2), pp. 279–315. ISSN: 1661-7207. DOI: 10.4171/GGD/158. MR: 2914861. arXiv: math/0607167 [math.GR] (cit. on pp. v, xiii, 1, 12, 13, 23, 40, 47, 51, 53, 57, 63, 65, 66, 71–73, 82, 85, 101, 104, 117).

Katok, A.B. and B. Hasselblatt (1995). *Introduction to the modern theory of dynamical systems*. Encyclopedia of Mathematics and its Applications 54. Cambridge: Cambridge University Press, 1995. 802 pp. ISBN: 978-0-521-57557-7. DOI: 10.1017/CBO9780511809187. MR: 1326374. With supplementary material by Katok and Mendoza (cit. on pp. 26, 28).

Kogan, Roman (2008). *nvTrees*. Source code available at https://github.com/romwell/nvTrees. 2008. URL: http://www.math.tamu.edu/~romwell/nvTrees/ (visited on 02/03/2015) (cit. on p. xv).

Liousse, Isabelle (2005). 'Nombre de rotation, mesures invariantes et ratio set des homéomorphismes affines par morceaux du cercle'. French. In: *Annales de l'institut Fourier* 55.2, pp. 431–482. ISSN: 1777-5310. DOI: 10.5802/aif.2103. MR: 2147896 (cit. on p. 28).

Lodha, Yash and Justin Tatch Moore (2016). 'A nonamenable finitely presented group of piecewise projective homeomorphisms'. In: *Groups, Geometry and Dynamics* 10.1, pp. 177–200. ISSN: 1661-7207. DOI: 10.4171/GGD/347. MR: 3460335. arXiv: 1308.4250 [math.GR] (cit. on p. 2).

Martínez-Pérez, Conchita, Francesco Matucci and Brita E.A. Nucinkis (2016). 'Cohomological finiteness conditions and centralisers in generalisations of Thompson's group V'. In: *Forum Mathematicum* 28.5, pp. 909–921. ISSN: 1435-5337. DOI: 10.1515/forum-2014-0176. MR: 3543701. arXiv: 1309.7858 [math.GR] (cit. on p. 71).

— (2018). 'Presentations of generalisations of Thompson's group V'. In: *Pacific Journal of Mathematics*. ISSN: 0030-8730. arXiv: 1609.07058 [math.GR]. URL: https://msp.org/scripts/coming.php?jpath=pjm (visited on 28/05/2018). Forthcoming (cit. on p. 71).

Martínez-Pérez, Conchita and Brita E.A. Nucinkis (2013). 'Bredon cohomological finiteness conditions for generalisations of Thompson's groups'. In: *Groups, Geometry and Dynamics* 7.4, pp. 931–959. ISSN: 1661-7207. DOI: 10.4171/GGD/211. MR: 3134031. arXiv: 1105.0189 [math.GR] (cit. on pp. xvi, 72).

Matucci, Francesco (2008). 'Algorithms and classification in groups of piecewise-linear homeomorphisms'. PhD thesis. Cornell University, Aug. 2008. 234 pp. MR: 2712336. arXiv: 0807.2871 [math.GR]. URL: http://cornell.worldcat.org/oclc/465664820 (cit. on pp. v, xiii, 1, 15, 22, 28, 44, 50, 53, 59, 72, 76).

McKenzie, Ralph and Richard J. Thompson (1973). 'An Elementary Construction of Unsolvable Word Problems in Group Theory'. In: *Word Problems, Decision Problems and the Burnside Problem in Group Theory*. Ed. by W.W. Boone, F.B. Cannonito and R.C. Lyndon. Vol. 71. Studies in Logic and the Foundations of Mathematics. North-Holland, 1973, pp. 457–478. ISBN: 978-0-7204-2271-9. DOI: 10.1016/S0049-237X(08)71914-1. MR: 0396769 (cit. on pp. xvii, 1).

Munkres, James (2000). *Topology*. 2nd ed. Prentice Hall, 2000. xvi+537. ISBN: 978-0-13-468951-7. MR: 0464128 (cit. on p. 16).

Robertson, David M. (2015). *The thompson package. Python 3.3+ classes for computing with elements of the Higman-Thompson groups $G_{n,r}$*. 2015. URL: https://github.com/DMRobertson/thompsons_v (cit. on pp. xv, 53, 117).

Robinson, Derek (1996). *A Course in the Theory of Groups*. 2nd ed. Graduate Texts in Mathematics 80. Springer-Verlag New York, 1996. xvii+502. ISBN: 978-0-387-94461-6. DOI: 10.1007/978-1-4419-8594-1 (cit. on p. 82).

Rotman, Joseph J. (1995). *An Introduction to the Theory of Groups*. Graduate Texts in Mathematics 148. New York: Springer-Verlag, 1995. 517 pp. ISBN: 978-0-387-94285-8. DOI: 10.1007/978-1-4612-4176-8. MR: 1307623 (cit. on pp. 24, 79, 82).

Salazar-Díaz, Olga Patricia (2010). 'Thompson's group V from a dynamical viewpoint'. In: *International Journal of Algebra and Computation* 20.1, pp. 39–70. ISSN: 1793-6500. DOI: 10.1142/S0218196710005534. MR: 2655915 (cit. on pp. xiv, xv, 29, 31, 36, 53).

Scott, E.A. (1992). 'A Tour Around Finitely Presented Infinite Simple Groups'. In: *Algorithms and Classification in Combinatorial Group Theory*. Ed. by Gilbert Baumslag and Charles F. Miller III. Mathematical Sciences Research Institute Publications 23. New York: Springer-Verlag, 1992, pp. 83–119. ISBN: 978-1-4613-9732-8. DOI: 10.1007/978-1-4613-9730-4_4. MR: 1230630 (cit. on p. 1).

Stein, Melanie (1992). 'Groups of Piecewise Linear Homeomorphisms'. In: *Transactions of the American Mathematical Society* 332.2, pp. 477–514. ISSN: 1088-6850. DOI: 10.1090/S0002-9947-1992-1094555-4. MR: 1094555 (cit. on p. 12).

Tararin, V. M. (2001). 'On Automorphism Groups of Cyclically Ordered Sets'. In: *Siberian Mathematical Journal* 42.1, pp. 190–204. ISSN: 1573-9260. DOI: 10.1023/A:1004866131580. MR: 1830806 (cit. on p. 7).

Thompson, Richard J. (1965). 'Untitled handwritten notes'. Unpublished but widely circulated. 1965. URL: http://people.math.binghamton.edu/matt/thompson/ (visited on 19/04/2017) (cit. on p. 1).

— (1980). 'Embeddings into Finitely Generated Simple Groups Which Preserve the Word Problem'. In: *Word problems II*. Vol. 95. Studies in Logic and the Foundations of Mathematics. North-Holland, 1980, pp. 401–441. ISBN: 978-0-444-85343-1. DOI: 10.1016/S0049-237X(08)71348-X. MR: 0579955 (cit. on p. xvii).

Willard, Stephen (2004). *General Topology*. Mineola, New York: Dover, 2004. xii+369. ISBN: 978-0-486-43479-7 (cit. on p. 18).