The University of Melbourne
Department of Mathematics and Statistics

Honours Thesis

Introduction to Thompson's Group F

Author:
Daniel Yeow
largestprime@gmail.com

Supervisor:
Dr. Lawrence Reeves
L.Reeves@ms.unimelb.edu.au

November 3, 2006

Acknowledgements

First and foremost, I would like to thank my supervisor Lawrence Reeves whose guidance, aid and understanding was invaluable and indispensable throughout the course of what has been a very long year. It has been an honour and a pleasure to work under such a gifted and personable mentor on such interesting and, more importantly, beautiful mathematics.

I would also like to thank my parents whose (often thankless) counsel and generous financial and moral support has enabled me to pursue my course of study with minimal outside worry.

I would also like to thank the Melbourne University Mathematics and Statistics Society, the Melbourne University Student Counselling service and the Melbourne University Student Union.

Contents

1 Introduction

Thompson's groups F, T and V were defined by Richard Thompson in 1965 in connection with his work in logic. He and McKenzie used them to construct finitely-presented groups with unsolvable word problems [6]. In this paper we will be concerning ourselves primarily with Thompson's group F.

Thompson's group F is a group that seems to appear in many different and diverse areas of mathematics. Obviously, group theory is one of them seeing as Thompson's group (as the name suggests) is a group. It also pops up in areas as diverse as cryptography [12] and, less surprisingly, combinatorics [7].

Thompson's group is fairly widely researched and is still a very active area as indicated by the recent publication dates for many of the references of this paper. There are still some important questions which are yet to be answered. Questions like: 'Is Thompson's group amenable?'.

In this paper, we will give a brief exposition on what Thompson's group is exactly and elaborate more on some of the finer details of the bridges which connect different ways of looking at the group. This will lead on to a discussion of a canonical normal form for the group which we will prove is also unique. It follows on from this that Thompson's group contains a free abelian subgroup of infinite rank and that it is torsion-free. It is also known that, (for reasons which will not be covered fully in this paper) the group is of type FP_∞, [3]. Since groups that contain an infinite rank free abelian subgroup can not have finite cohomological dimension, F is an example of a torsion-free FP_∞ group which is not of finite cohomological dimension, [3].

We will then explore the concept of dead end elements and look at an example of one in Thompson's group as an illustration, not only of what a peculiar thing a dead end element is, but also as a brief look at some of the properties of the Cayley graph of Thompson's group. After this, there will be a very brief discussion on amenability and an overview of where the current research is at with regard to determining the amenability (or not) of Thompson's group. Finally, we will look at a potential practical application of Thompson's group in the area of public key cryptography.

2 Different Representations

Thompson's group F has many different representations which allow us flexibility when determining behavioural properties of the group. There is the standard group presentation of course, but in addition, there are also piecewise linear homeomorphisms, rectangle diagrams and pairs of rooted binary trees. Each of these allow us to get a different perspective on this fascinating group. In this section, we will focus on showing that these different representations are logically equivalent.

2.1 Group presentation

Thompson's group is encountered combinatorially as given by its standard finite and infinite presentations. First we have the finite:

$$F = \langle x_0, x_1 \mid x_1^{-1}x_2x_1 = x_3,\ x_1^{-1}x_3x_1 = x_4 \rangle \tag{1}$$

where $x_2 = x_0^{-1}x_1x_0$, $x_3 = x_0^{-2}x_1x_0^2$ and $x_4 = x_0^3x_1x_0^{-3}$. Then there is the infinite presentation:

$$\mathcal{P} = \langle x_k, k \geq 0 \mid x_i^{-1}x_jx_i = x_{j+1} \text{ if } i < j \rangle \tag{2}$$

We begin by showing that these two presentations give the same group.

We take the finite presentation and add abbreviations to make it infinite. We take the $x_n = x_0^{-1}x_{n-1}x_0$ for $n > 1$ rule and modify it to get x_n by itself on the left hand side.

$$\begin{aligned} x_n &= x_0^{-1}x_{n-1}x_0 \\ \text{since } x_{n-1} &= x_0^{-1}x_{n-2}x_0 \\ x_n &= x_0^{-1}x_0^{-1}x_{n-2}x_0x_0 \\ \Rightarrow x_n &= x_0^{-(n-1)}x_1x_0^{(n-1)} \end{aligned}$$

We add this to the finite presentation F to make it infinite:

$$F' = \langle x_0, x_1, x_2, \ldots | x_1^{-1}x_2x_1 = x_3, x_1^{-1}x_3x_1 = x_4, x_i = x_0^{-(i-1)}x_1x_0^{(i-1)} \ \forall i \geq 2 \rangle$$

Since we've only added an abbreviation, $F \cong F'$.

It now remains to show that F' is equivalent to the inifinte presentation denoted by $\mathcal{P}$ above. If we let $i = 0$ in $\mathcal{P}$, we get $x_0^{-1}x_j x_0 = x_{j+1}$. Doing the same thing in F with $x_i = x_0^{-(i-1)}x_1 x_0^{(i-1)}$ we get:

$$\begin{aligned} x_{j+1} &= x_0^{-j}x_1 x_0^{j} \\ x_j &= x_0^{-(j-1)}x_1 x_0^{(j-1)} \\ x_0^{-1}x_j x_0 &= x_0^{-1}x_0^{-(j-1)}x_1 x_0^{(j-1)}x_0 \\ &= x_0^{j}x_1 x_0^{j} \\ &= x_{j+1} \end{aligned}$$

Which we will build on shortly.

We can also show that given any (i, j) pair (where (i, j) is a more convenient way of writing $x_i^{-1}x_j x_i$), the $(i + 1, j + 1)$ pair will also work. i.e. $(i, j) \Rightarrow (i + 1, j + 1)$ where $0 < i < j$.

$$\begin{aligned} x_{i+1}^{-1}x_{j+1}x_{i+1} &= x_0^{-1}x_i^{-1}x_j x_i x_0 \\ &= x_0^{-1}x_{j+1}x_0 \quad \text{by } (i,j) \\ &= x_{j+1} \end{aligned}$$

Now, taking stock of what we have shown, we may describe our results using a lattice of points on the (i, j) plane. Taking into account our base cases of $x_1^{-1}x_2x_1 = x_3$, $x_1^{-1}x_3x_1 = x_4$ and all the cases of the form $x_0^{-1}x_i x_0 = x_{i+1}$, and in view of the fact that $(i, j) \Rightarrow (i+1, j+1)$ we have all the points which lie on the j-axis, the line $j = i + 1$ and $j = i + 2$. The lattice, as it stands, looks like the diagram Figure 1(a).

It remains to show that the points which we already have (filled in in Figure 1(a)), imply all of the points which satisfy the conditions $i < j$ and $i \geq 0$.

This can be done by showing that, if we assume the points $(1, j - 1)$ and $((j - 1), j)$ are satisfied, then the point (i, j) is satisfied (subject to our restrictions on i and j). Once we have shown this, then it follows from the cases which we already have that, in fact, all cases are satisfied.

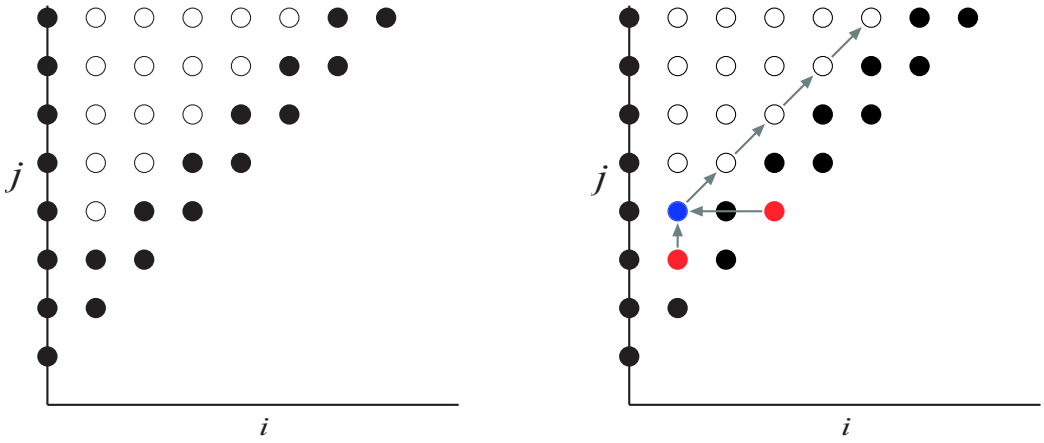

(a) Our initial information.

(b) The red dots imply the blue dot.

Figure 1: The lattice and how we use the points which we have already established (in black and red) in order to establish the ones which he haven't yet established (in blue). Once we have our new blue dot, it automatically establishes all points lying on the diagonal indicated by the arrows. We repeat this process to fill in all the dots in the lattice.

$$\begin{aligned} x_{j+1} &= x_{j-1}^{-1} x_j x_{j-1} \\ &= x_{j-1}^{-1} x_1^{-1} x_{j-1} x_1 x_{j-1} && \text{by } (j-1, j) \\ &= x_1^{-1} x_{j-2}^{-1} x_{j-1} x_{j-2} x_1 && \text{by } (1, i-2) \\ &= x_1^{-1} x_j x_1 && \text{by } (j-2, j-1) \end{aligned}$$

In Figure 1(b) the two points highlighted in red imply the point highlighted in blue. It is also interesting to note that points which lie on the same horizontal line are, in fact, the same element of the group. (this is, in fact, exactly what the infinite presentation says in $x_i^{-1} x_j x_i = x_{j+1}$).

This shows that the two presentations are, in fact, the same.

$$\begin{aligned} & F' \cong \mathcal{P} \\ \therefore \quad & F \cong \mathcal{P} \text{ as required.} \end{aligned}$$

Since there are two generators, x_0 and x_1, one can easily see that every element x_k in the group can be expressed in the form:

$$\overbrace{x_{i_1}^{-1}x_{i_2}^{-1}\ldots x_{i_{k-2}}^{-1}x_0^{-1}}^{k-1 \text{ of these}} x_1 \overbrace{x_0 x_{i_{k-2}} x_{i_{k-3}} \ldots x_{i_1}}^{k-1 \text{ of these}}$$

Where each $i_1, i_2, \ldots i_{k-2}$ is either a 0 or a 1.

We also have a convenient expression for the normal form which we will introduce now and elaborate on towards the end of this section.

$$x_0^{b_0}x_1^{b_1}\ldots x_n^{b_n}x_n^{-a_n}\ldots x_1^{-a_1}x_0^{-a_0} \tag{3}$$

Where the a_i's and b_i's can take whole number values (in practice, most are zero) and n is finite. A unique normal form for each word can be made from this if we add the condition that whenever both x_i and x_i^{-1} occur, then so does x_{i+1} or x_{i+1}^{-1}. This will become more intuitively apparent towards the end of this section as well.

2.2 Piecewise linear homeomorphisms

Another way of looking at Thompson's group is by way of piecewise linear homeomorphisms. These are homeomorphisms from the interval $[0, 1]$ to $[0, 1]$. The functions satisfy these four conditions:

1. The function is piecewise linear
2. The function is differentiable except at finitely many points
3. Each of these points is a dyadic rational number, i.e. a rational number whose denominator is a power of 2.
4. On the intervals of differentiability, the derivatives are powers of 2.

In this case, the elements of the group are simply the homeomorphisms themselves and the operation is the composition of functions. It is fairly easy to see that, since the $f'(x)$ is always a power of 2, $f(x_i)$ where x_i is a point of non-differentiability, is also a dyadic rational number. Therefore any f^{-1} is well-defined and so every element of this group has an inverse (geometrically, a reflection about the line $y = x$). The identity is simply the line $f(x) = x$. This group is a subgroup of the group of all homeomorphisms from $[0, 1]$ to $[0, 1]$. An example of two such functions are given below.

$$A(x) = \begin{cases} x/2, & 0 \leq x \leq \frac{1}{2} \\ x - \frac{1}{4}, & \frac{1}{2} < x \leq \frac{3}{4} \\ 2x - 1 & \frac{3}{4} < x \leq 1 \end{cases} \qquad B(x) = \begin{cases} x, & 0 \leq x \leq \frac{1}{2} \\ \frac{x}{2} + \frac{1}{4} & \frac{1}{2} < x \leq \frac{3}{4} \\ x - \frac{1}{8} & \frac{3}{4} < x \leq \frac{7}{8} \\ 2x - 1 & \frac{7}{8} < x \leq 1 \end{cases}$$

Figure 2: The graphs of the piecewise linear homeomorphisms $A(x)$ and $B(x)$ given above.

It turns out, as we will see later, that these functions are in fact the generators of Thompson's group.

2.3 Rectangle diagrams

Thompson's group may also be interpreted geometrically by way of a *rectangle diagram* representing $f(x)$. These diagrams simply have the *top* representing the preimage of the above functions and the *bottom* representing the image of the function.

We shall define our group of rectangle diagrams G to be the group of all rectangle diagrams (corresponding to piecewise linear homeomorphisms) which are generated by the generators g_0 and g_1 given by the rectangle diagrams in Figure 3. Right multiplication of w by an element s is given by placing the rectangle representing s on top of the rectangle representing w. The various lines in the rectangles get *joined up* in a manner which will become more obvious in later examples. Inverses of elements are given by reflecting the rectangles about the horizontal axis (effectively switching image and preimage).

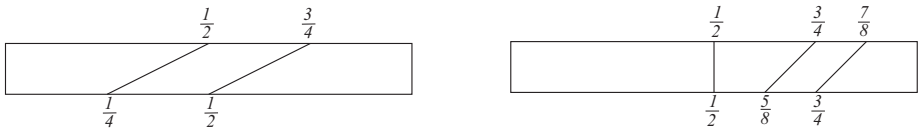

(a) g_0 (b) g_1

Figure 3: The rectangle diagrams of g_0 and g_1.

(Note that g_0 and g_1 correspond to the piecewise linear homeomorphisms $A(x)$ and $B(x)$ given in the previous subsection).

A useful thing to establish would be a geometrical interpretation of how the elements of this group behave. For example, in the section about the group presentation, we showed that the element x_3 could be represented either by $x_1^{-1}x_2x_1$ (which, in terms of the two generators is $x_1^{-1}x_0^{-1}x_1x_0x_1$) or by $x_0^{-2}x_1x_0^2$. We can show, quite easily, that if we consider the rectangle diagram counterparts to the elements in the group presentation (x_0's counterpart would be g_0 and x_1's counterpart would be g_1), that these two presentations of the same element are in fact the same. Later in the section, we will make this more precise by proving that these two groups are isomorphic. For now:

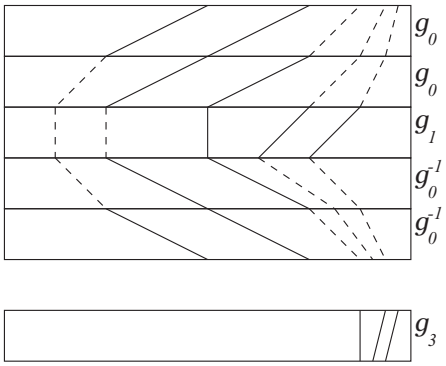

(a) $g_0^{-2}g_1g_0^2$

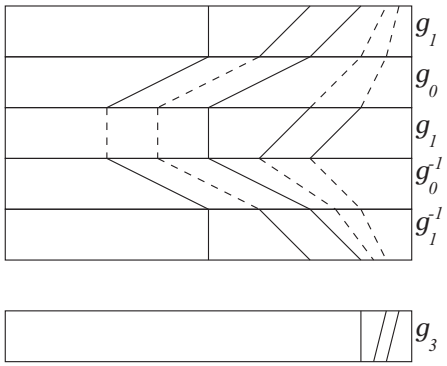

(b) $g_1^{-1}g_0^{-1}g_1g_0g_1$

Figure 4: The rectangle diagrams showing the equivalence of $g_0^{-2}g_1g_0^2$ and $g_1^{-1}g_0^{-1}g_1g_0g_1$ which are both g_3.

In fact, it follows on that every successive g_i moves the first vertical line in the rectangle to the right by $\frac{1}{2^i}$. The point is illustrated in Figure 5 with the rectangle diagrams for g_1, g_2, g_3 and g_4.

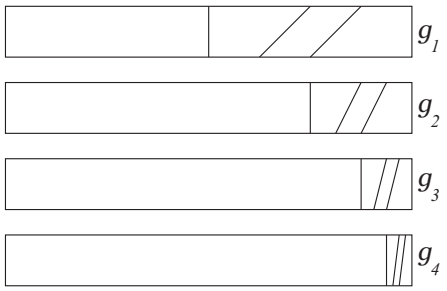

Figure 5: The rectangle diagrams of g_1, g_2, g_3 and g_4.

2.4 F and G isomorphic

Let F be Thompson's group as given by either the finite or infinite group presentations (proved equivalent at the beginning of the section) and let G be the group of all Rectangle diagrams corresponding to the group of piecewise linear homeomorphisms from [0,1] to [0,1] as discussed in sections 2.2 and 2.3. We proceed to tie all the preceding sections together by supplying the connection between them and showing that the normal form for F is unique.

We want to show that the groups F and G are isomorphic.

2.4.1 Homomorphism

First show that there is a homomorphism from F to G. Recall g_0 and g_1 from Figure 3 in the previous subsection.

We want to define a homomorphism $\phi : F \to G$ by setting $\phi(x_0) = g_0$ and $\phi(x_1) = g_1$.

The first test is to see whether this gives a homomorphism. We use the fact that any homomorphism from F to G will always map the identity in F to the identity in G. The identity in G is just a blank rectangle (representing the piecewise linear homeomorphism $f(x) = x$). So we take our relators in F.

$$x_1^{-1}x_2x_1 = x_3 \text{ and } x_1^{-1}x_3x_1 = x_4$$

rearrange them to get the identity on one side like so:

$$x_1^{-1}x_2x_1x_3^{-1} = 1 \text{ and } x_1^{-1}x_3x_1x_4^{-1} = 1$$

We apply ϕ to both sides of each equation. $\phi(1) = 1$ and, if we are correct in saying that ϕ is a homomorphism, then $\phi(x_1^{-1}x_2x_1x_3^{-1})$ and $\phi(x_1^{-1}x_3x_1x_4^{-1})$ should both equal 1 as well.

First we split our expression up by applying $\phi(ab) = \phi(a)\phi(b)$ a few times to get $\phi(x_1^{-1})\phi(x_2)\phi(x_1)\phi(x_3^{-1})$ and $\phi(x_1^{-1})\phi(x_3)\phi(x_1)\phi(x_4^{-1})$. Now we take a look at the rectangle diagrams to see if we indeed get the identity.

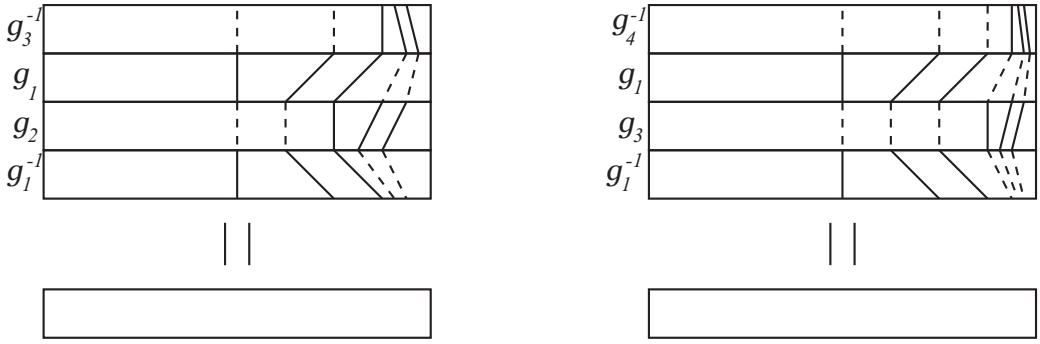

(a) $g_1^{-1}g_2g_1g_3^{-1} = 1$ (b) $g_1^{-1}g_3g_1g_4^{-1} = 1$

Figure 6: Rectangle diagrams of the relators from the finite presentation F.

And indeed we do in both cases.

Since all the elements in the group F, ($x_0, x_1, x_2 \ldots$ etc.), are defined on the generators and the relators, it follows that the homomorphism $\phi : x_0 \to g_0, x_1 \to g_1$ extends to a homomorphism $\phi : F \to G$ and we can now assert that $\phi(x_n) = g_n$.

In order to show that the homomorphism is an isomorphism, we now need to show that it is both injective and surjective.

2.4.2 Injectivity

To show that ϕ is injective, we need to show that $\phi(a) = \phi(b)$ implies $a = b$. For this, we use a normal form for F which we show is unique.

Given any word in F we can use the following rules (which we obtained from rearranging the relators in the infinite presentation) to rearrange our word into our desired normal form.

$$\begin{aligned} x_j x_i &= x_i x_{j+1} \\ x_i^{-1} x_j &= x_{j+1} x_i^{-1} \end{aligned} \tag{4}$$

The first one ensures that we can always get the x_i's in increasing or decreasing order (depending on which half of the normal form we find ourselves in).

The second one ensures that we can swap a positive for a negative in order to get all the positive exponents on the left and the negative ones on the right.

We already have a candidate for the normal form:

$$x_0^{b_0} x_1^{b_1} \dots x_n^{b_n} x_n^{-a_n} \dots x_1^{-a_1} x_0^{-a_0} \tag{5}$$

Moreover, we can rearrange any given word to give a normal form such that whenever both x_i and x_i^{-1} occur then so must x_{i+1} or x_{i+1}^{-1}. This is because, if both x_i and x_i^{-1} occur and neither x_{i+1} nor x_{i+1}^{-1} do, we can swap them both *inwards* with higher subscripted elements (which are necessarily towards the centre, due to the structure that our normal form takes) and keep repeating this process until both x_i and x_i^{-1} are next to each other and cancel each other out. We claim that this is gives a unique normal form.

Given that we have established that $\phi : F \to G$ is a homomorphism we simply map all the x_i's to g_i's which gives a normal form (not surprisingly) which looks like equation (5) except with g's instead of x's – (6).

$$g_0^{b_0} g_1^{b_1} \dots g_n^{b_n} g_n^{-a_n} \dots g_1^{-a_1} g_0^{-a_0} \tag{6}$$

All we now need to show is that this normal form, now in terms of g_i's is unique.

Notice that every element $g \in G$ (the rectangles) has a *leftmost vertical line* which we will define as being the vertical line at the point $1 - 2^{-i}$ of any given g_i-rectangle (remembering that the top line of the rectangle represents the preimage in the piecewise linear homeomorphism and the bottom line represents the image). The lines that divide a rectangle diagram are simply indications of when the gradient of the piecewise linear homeomorphisms change. The lines themselves correspond to the points of non-differentiability (which is why they all occur on dyadic rational numbers). Now we turn our attention to the next line in the rectangle, right of the vertical line which we just discussed. Taking the full length of the rectangle to be 1, the ratio of the distance from the start of the vertical line to the start of the next line to the distance from the end of the vertical line to the end of the next line corresponds to the right-hand derivative of the piecewise linear homeo-

morphism at the point $x = 1-2^{-i}$. To illustrate this we observe the case of x_1:

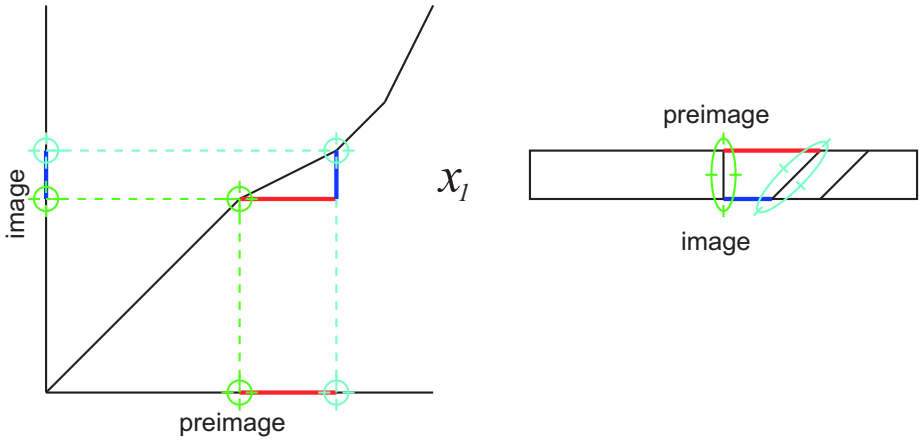

Figure 7: Diagram illustrating the relation between lines in the rectangle diagram and points on the piecewise linear homeomorphism using x_1.

Taking our normal form in terms of g_n's $\in G$ as described above. We find i in this normal form such that it is the smallest subscript and note that the right-hand derivative of the piecewise linear homeomorphism which corresponds to the rectangle represented by the normal form is 2^{-j} where j is the g_i-exponent sum in the normal form. Moreover, any other normal form for g with subscripts $\geq i$ will have the same g_i-exponent sum.

The reason for this, is because we know that for any given g_n, the rectangle diagram's leftmost vertical line occurs at the point $1 - 2^{-n}$. Therefore, any g_m where $m > n$ would, when composed with its smaller g_n's only affect the final rectangle diagram at points past $1 - 2^{-m}$ leaving the crucial section between $1 - 2^{-n}$ and $1 - 2^{-(n+1)}$ unchanged. It is easy to see then, that the only way to change this is with another g_n or g_n^{-1} (or a g_l where $l < n$, but since we are picking a smallest subscript, we needn't worry about this).

Figure 8 clearly shows that, as long as we pick n to be the smallest subscript, the right-hand derivative at $x = n$ (represented by the ratio between the thick blue and green horizontal lines on the rectangle diagram as illustrated further above) depends **only** on the g_n-exponent sum. We can now proceed with the proof of the uniqueness of the normal form.

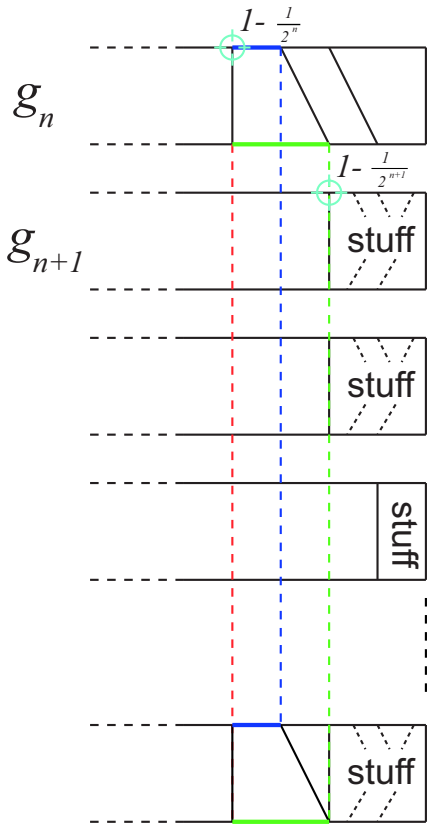

Figure 8: Diagram showing the effect of the *leftmost vertical line* in the rectangle diagrams representing elements of the group G.

Our proof is by contradiction. First we assume that there are two different normal forms for the same element of the group, say G. We choose such a pair so that they have minimal total length. That is, out of all the different group elements, we pick one which, out of all the different normal forms which represent it, we find two such that their sum is minimal. We specify that that total length is necessarily non-zero and positive. So we have:

$$\text{normalform1} = \text{normalform2} \tag{7}$$

Firstly, both can neither begin nor end with g_i or g_i^{-1} respectively, otherwise we could just cancel them by either left-multiplying by g_i^{-1} or right-multiplying by g_i contradicting minimality. Since they both represent the same group element, both these normal forms must necessarily correspond to the same rectangle (piecewise linear homeomorphism). This implies that

they must then have the same g_i-exponent sum. We now turn our attention to g_i (or g_i^{-1}) – the lowest exponent and the g_i-exponent sum. From earlier in the paragraph, we know that if one of the normal forms begins with x_i then the other cannot, and similarly with x_i^{-1}. In the case where one starts with g_i and the other finishes with g_i^{-1} we note that, because g_i is the smallest exponent, this would make it impossible for both normal forms to have the same g_i-exponent sum. What we are left with then, is that one of the normal forms must contain g_i or g_i^{-1} and the other cannot. However, in order for them both to have the same g_i-exponent sum, the normal form which contains one of g_i or g_i^{-1} must therefore contain both.

$$\begin{aligned} \text{normalform1} &= g_i \ldots \text{stuff} \ldots g_i^{-1} \\ \text{normalform2} &= g_j \ldots \text{otherstuff} \ldots g_k^{-1} \\ i &< j, k \text{ (by minimality of } i) \end{aligned} \tag{8}$$

Without loss of generality, lets say that normalform1 is the one which contains both. The equality between rectangles represented by normal forms is thus $g_i p g_i^{-1} = \text{normalform2}$, where p is just normalform1 with the g_i and g_i^{-1} removed from either side. It is important to note that, normalform1 has subscripts which are all $\geq i$ and, since normalform2 does not contain g_i or g_i^{-1}, all of its subscripts must be $\geq i+1$. Then, with a little rearranging, we get $p = g_i^{-1}(\text{normalform2})g_i$ (or $p = (\text{normalform2})^{g_i}$).

From the rule $x_i^{-1}x_j = x_{j+1}x_i^{-1}$ discussed above, (we can apply it to G since ϕ is a homomorphism) it is easy to see that the g_i^{-1} at the front of normalform2 will swap with the first element of normalform2, then the second and so on until it reaches the point in the 'middle' of the normal form where positive exponents end and negative ones begin. The g_i at the end of normalform2 will do the same until it reaches the middle and eliminates g_i^{-1}. The net result of all this is that all the subscripts of normalform2 are now one higher than they were before normalform2 was conjugated. This gives an expression which is made up entirely of elements which have subscripts $\geq i+2$. But this would mean that p only involves subscripts which are $\geq i+2$ which would contradict the assumption that normalform1 $= g_i p g_i^{-1}$ is a normal form. $\square$

$\Rightarrow$ The normal form of $g \in G$ is unique

$\Rightarrow$ If $\phi(a) = \phi(b)$ then a = b (since $g_0^{b_0} g_1^{b_1} \dots g_n^{b_n} g_n^{-a_n} \dots g_1^{-a_1} g_0^{-a_0}$ is the same as $\phi(x_0^{b_0})\phi(x_1^{b_1}) \dots \phi(x_n^{b_n})\phi(x_n^{-a_n}) \dots \phi(x_1^{-a_1})\phi(x_0^{-a_0})$)

$\Rightarrow \phi$ is injective.

2.4.3 Surjectivity

Define an inverse function of ϕ, say Δ. We know that $\phi(x_0) = g_0$ and $\phi(x_1) = g_1$ so we'll define Δ so that $\Delta(g_0) = x_0$ and $\Delta(g_1) = x_1$. Since we've defined the group G as being generated by g_0 and g_1, and shown that the normal form of $g \in G$ and $x \in F$ is unique, surjectivity follows immediately.

The groups F and G are injective and surjective

$\Rightarrow$ the groups F and G are isomorphic.

2.5 Pairs of rooted binary trees

We can also interpret elements of the group in the form of pairs of rooted binary trees. This representation is particularly useful for calculating word length. We denote these rooted binary trees by (T_-, T_+), each with the same number of exposed *leaves*. An *exposed leaf* ends in a vertex of valence (degree) 1. We number these exposed leaves from left to right, beginning with 0 . We refer to a node together with the two downward-directed edges from the node as a *caret* (see Figure 9). A caret C may have a *right child*, a caret C_R which is attached to the right edge of C. We can similarly define the *left child* C_L of the caret C. The set of all carets which stem from the right leaf of a caret C is called the right subtree of C, and we can analogously define the left subtree of C, [7].

Definition 2.1. *A* **caret** *is a rooted binary tree with exactly two edges and three vertices. Every caret has the form shown in Figure 9. Every non-leaf vertex of a rooted binary tree is the root of a caret.*

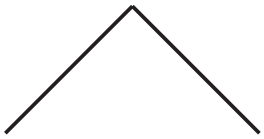

Figure 9: A Caret.

Proposition 2.2. *The generators x_0 and x_1 are represented by the tree pair diagrams in Figure 10*

(a) x_0 (b) x_1

Figure 10: The tree pair diagrams for x_0 and x_1.

It is easy to see that the leaves on these trees correspond to *standard dyadic intervals*.

Definition 2.3. *A* **standard dyadic interval** *is an interval of the form* $[\frac{a}{2^n}, \frac{a+1}{2^n}]$*, where* a *and* n *are nonnegative integers with* $a \leq 2^n - 1$*.*

With this in mind, it is perhaps easier to see why the tree pair diagrams in Figure 10 are the way that they are. In figure 11 we see the same diagrams except with the intervals which each of the leaves represent.

(a) x_0 (b) x_1

Figure 11: The tree pair diagrams for x_0 and x_1 with dyadic intervals marked in at the end of the leaves.

It is no coincidence that these intervals are the same as those which appear in the preimage and image of the corresponding piecewise linear homeomorphisms for x_0 and x_1 from section 2.2.

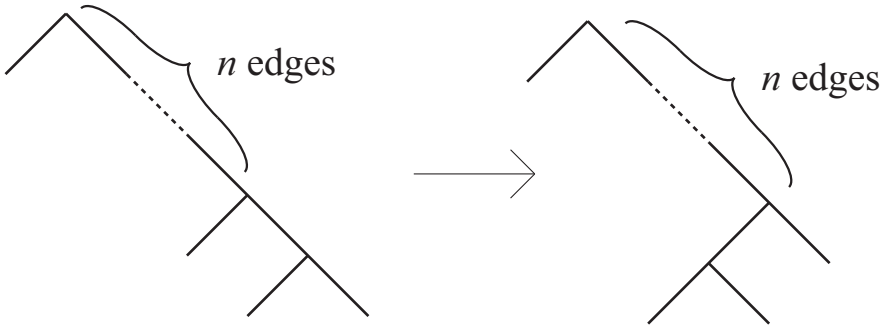

Figure 12: The reduced tree pair diagram for X_n.

Extrapolating from the pattern set in Figure 10, it is easy to see that the reduced tree diagram for x_n is the tree diagram in Figure 12.

There is just one more definition we need before we can make the connection between tree-pair diagrams and Thompson's group F given by its group presentation more precise. The composition of functions needs to be defined.

In our rectangle diagrams, when we joined them up, we often had to add new lines to make the images and preimages connect properly. In fact, it is helpful to think of these lines not as *new* lines, but of redundant lines which have, by necessity been brought into use. Perhaps an illustration might help, see Figure 13.

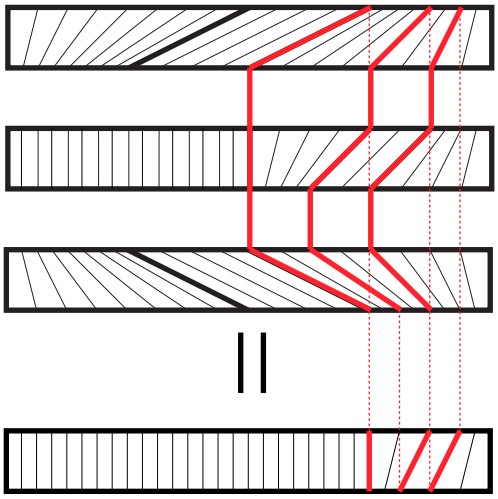

Figure 13: An illustration of 'redundant' lines (not the red ones).

The only lines which 'matter' are the ones which define the borders between regions where an interval in the preimage is taken to an interval in the image which is 2^n the size of the original interval ($n \in \mathbb{Z}$). Although it is difficult to see, these redundant lines can be equated to redundant carets – carets which prevent the tree pair diagram from being reduced. So follows the concept of a 'reduced tree'.

Definition 2.4. *A caret is* **redundant** *if both its leaves are the same in T_- and T_+.*

A reduced tree is one in which, if you number all the leaves from left to right (or right to left) in both trees in your tree pair, there will be no instances of the i^{th} and $i+1^{th}$ leaves being on the same caret (see Figure 14) in both trees.

Definition 2.5. *A pair of rooted binary trees is* **reduced** *if there are no redundant carets.*

In order to 'compose' two functions which are represented in tree pair diagram form, we need simply to ensure that all the trees have the same number of leaves. To do so, one adds so-called redundant carets in a fashion very similar to the adding of lines in the rectangle diagram.

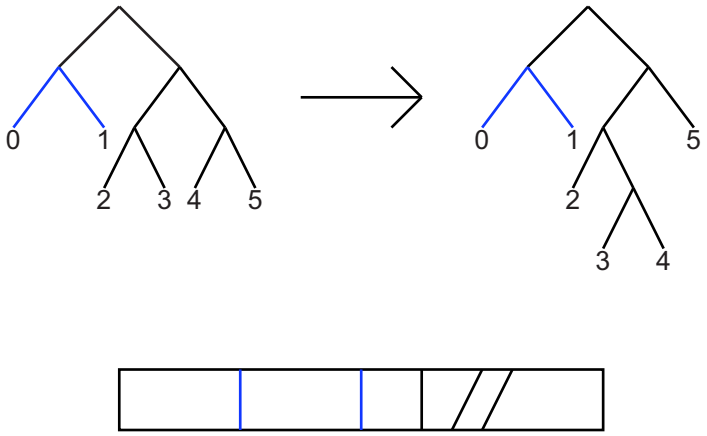

Figure 14: A non-reduced tree pair diagram (the blue caret prevents this tree pair diagram from being reduced). The corresponding rectangle diagram is also shown illustrating the fact that redundant carets essentially provide no new information on the element and can so be discarded.

The following theorem is reproduced from Fordham [11] and uses the connection with piecewise linear homeomorphisms and the standard dyadic intervals.

Theorem 2.6. *There is exactly one reduced pair of trees representing an element $f \in F$.*

Proof. Assume that (T_-, T_+) is a reduced tree diagram representing a piecewise linear homeomorphism $f \in F$. If I is a standard dyadic interval such that I is a leaf of T_- or I is not a node in T_-, then I is a leaf or one of infinitely many smaller subintervals of a leaf in T_-, so f is linear on I. If I is a standard dyadic interval such that f is linear on I and $f(I)$ is also a standard dyadic interval, then I is a leaf of T_- or is not a node of T_- since the pair of trees is reduced. Therefore, T_- is the *only* tree such that a standard dyadic interval I is a leaf of T_- or not a node of T_- if and only if f is linear on I and $f(I)$ is a standard dyadic interval. □

Our next definition allows us to make the link from tree diagram representations to our original group presentation complete. We need to define what an *exponent* is.

Definition 2.7. *Let $I_0, \ldots, I_n$ be the leaves of the tree in order from left to right. For every integer k with $0 \leq k \leq n$ let a_k be the length of the maximal arc of left edges in the tree which begins at I_k and which does not reach the right side of the tree. Then a_k is the k^{th}* **exponent** *of the tree.*

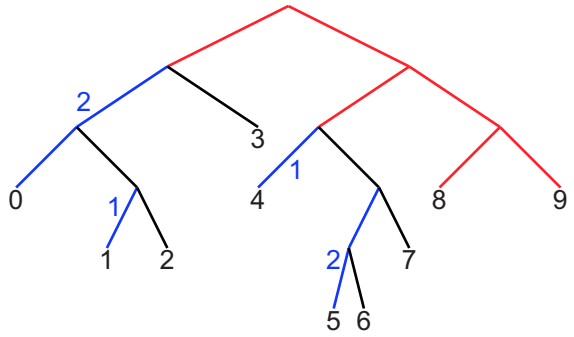

Figure 15: Diagram indicating the right side of a tree and exponents. The red 'branches' indicate the right side of the tree and the blue branches highlight the non-trivial exponents.

In Figure 15, the exponents of leaves number 0,1, 2, 3, ... 9, are 2, 1, 0, 0, 1, 2, 0, 0, 0, 0.

2.5.1 Putting it all together

While the connection between piecewise linear homeomorphisms, rectangle diagrams and the group presentation has already been established (via uniqueness of the normal form), the connection with the tree pair diagrams is not.

This connection can be given as follows. For any reduced tree pair, take the preimage tree's exponents and label them $a_0, a_1, a_2, \ldots, a_n$, then take the image tree's exponents and label them $b_0, b_1, b_2, \ldots, b_n$. The normal form is given by

$$x_0^{b_0} x_1^{b_1} \ldots x_n^{b_n} x_n^{-a_n} \ldots x_1^{-a_1} x_0^{-a_0} \tag{9}$$

This is perhaps easier to see with an example. Consider the trees in Figure 16 to demonstrate how to count exponents.

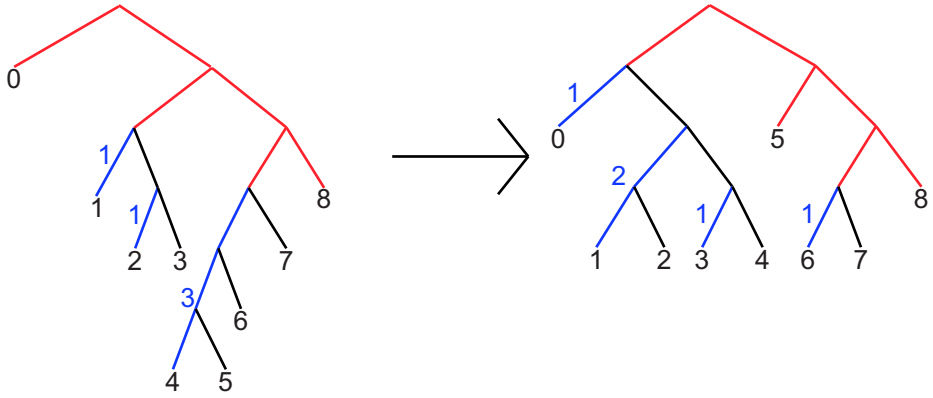

Figure 16: Counting exponents on the left and right to arrive at the normal form.

We have an exponent of 1 next to the first leaf, 1 next to the second leaf and 3 next to the fourth on the left-hand tree. On the right we have an exponent of 1 on the zero-th leaf, 2 on the first, 1 on the second, 1 on the third, and 1 on the sixth. This gives a normal form of $x_0x_1^2x_3x_6x_4^{-3}x_2^{-1}x_1^{-1}$. Notice that adding leaves to the right side of the tree (the branches in red) does nothing to the normal form but can impact on whether or not a tree-pair is reduced.

One can also construct a tree pair diagram from any given normal form by reversing this process (although it is somewhat more fiddly).

From this construction, we observe that every tree pair diagram has exactly one normal form which represents it. Since we proved in the previous subsection that there is exactly one reduced pair of trees representing every element of the group, we now establish that there is a bijective correspondence between elements in the group (as given by the group presentation) and the set of all pairs of rooted binary trees.

2.6 Other properties

We begin with a curious theorem about creating elements of the group by multiplying two elements of the group together.

Theorem 2.8. *F contains a copy of $F \times F$*

Proof. We may prove this by means of an explicit construction. We need to show that $\varphi : F \times F \mapsto F$ is injective and well-defined.

Let f and $g \in F$ take values from the interval [0,1] and map them to [0,1]. We define our map $\varphi(f, g) \mapsto h$:

$$\varphi(x) = \begin{cases} \frac{1}{2}f(2x) & 0 \leq x \leq \frac{1}{2} \\ \\ \frac{1}{2}g(2x-1) + \frac{1}{2} & \frac{1}{2} \leq x \leq 1 \end{cases}$$

We need to show that $F \times F \mapsto F$ is injective. We accomplish this by seeing what the kernel maps to.

Suppose $\varphi(f, g) = id_F$

$\Rightarrow$	$\varphi(f,g)$	$=$	x		
$\Rightarrow$	$\frac{1}{2}f(2x)$	$=$	x	$\forall x \in [0, \frac{1}{2}]$	(a)
and	$\frac{1}{2}g(2x-1)+\frac{1}{2}$	$=$	x	$\forall x \in [\frac{1}{2}, 1]$	(b)
$\Rightarrow$	$f(2x)$	$=$	$2x$	by rearranging (a)	
and	$g(2x-1)$	$=$	$2x-1$	by rearranging (b)	
	$f(u)$	$=$	u	if we let	$u = 2x$
	$g(v)$	$=$	v	if we let	$v = 2x-1$
$\Rightarrow$	f	$=$	id_F	and	$g = id_F$
$\Rightarrow$	φ	is injective			

We check that the domain and codomain of f and g map to the domain and codomain of h, which indeed they do.

It is helpful to think of this as simply 'squeezing' two piecewise linear homeomorphisms to half the size and putting one in the bottom-left corner and the other in the top-right corner of the $[0, 1] \times [0, 1]$–box in the cartesian plane to construct another piecewise linear homeomorphism.

Because this new piecewise linear homeomorphism's points of non-differentiability occur on dyadic rational numbers, and because we are always multiplying our numbers by powers of 2, our new function will also have its points of non-differentiability on dyadic rational numbers.

$$\therefore h \text{ is in } F$$

□

We could, by the same process shrink elements of F indefinitely and place them on the diagonal between (0,0) and (1,1) to construct more elements of F.

Note: $\varphi(id_F, g_0) = g_1$; moreover $\varphi(id_F, g_n) = g_{n+1}$, $\varphi(\varphi(id_F, g_0)) = g_2$ and $\varphi^n(id_F, g_0) = g_n$.

Corollary 2.9. *F contains a copy of F^{2k} for all $k \in \mathbb{N}$. Repeat the process from the proof of Theorem 2.8 k times to arrive at such a subgroup.*

Theorem 2.10. *F is torsion-free. That is, every element of the group (except the identity) has infinite order.*

Proof. We know that every element of the group can be represented by a piecewise linear homeomorphism which maps the interval [0,1] to [0,1] which satisfies the conditions defined in section 2.2.

We begin at the point (0,0). We know that the point (0,0) gets mapped to (0,0) under any element of the group. Consider the point $(\varepsilon, f(\varepsilon))$ for some $\varepsilon > 0$. If $f(\varepsilon) \neq \varepsilon$ then composing f with itself once, twice, n times will take $f(f(\varepsilon))$, $f(f(f(\varepsilon)))$, etc. further and further away from ε implying that the element has infinite order.

If $f(\varepsilon) = \varepsilon$ then we move our starting point along the line $f(x) = x$ until we hit a point where the gradient isn't 1. If we never hit that point over the whole interval [0,1] then we have the identity element. If we reach such a point, then we simply proceed as above and it becomes clear that all elements of the group except the identity have infinite order. Therefore the group is torsion-free. □

Proposition 2.11. *F contains a free abelian subgroup of infinite rank.*

Proof. The elements $x_0x_1^{-1}$, $x_2x_3^{-1}$, $x_4x_5^{-1}$, ..., commute and are linearly independent. □

This is noteworthy because it seems counterintuitive that a group which has a finite presentation with only two generators can contain an isomorphic copy of $\mathbb{Z}^\infty$.

F is also torsion-free and is of type FP_∞. This is explored more fully in [3]. This is due in part to the fact that it is torsion-free and in part to properties relating to it cohomological dimension and contractability. Things which are beyond the scope of this paper. It is noteworthy that this group was the first example of a torsion-free infinite-dimensional FP_∞ group, see [4].

3 Word Length

We have a group and several different but equivalent ways of viewing this group. A natural question to ask would be how to calculate the word length of elements of the group. This leads us to the Cayley graph of the group and, eventually, to the rather curious phenomenon of *dead end elements* – elements from which, no matter which way you travel (on the Cayley graph) the distance from your starting point does not increase. Another way of thinking about it, for those of us who are less visually inclined in our thinking, is to say that given any word w, adding either $x_0^{\pm 1}$ or $x_1^{\pm 1}$ to the right hand side of the word does not increase the word length.

3.1 Calculating word length

To calculate the length of a word in our group, we take a word in the normal form, reduce it to whichever generating set we choose (usually x_0 and x_1 from our finite presentation), and basically *count* the number of letters in the word (this process is non-trivial, hence the usefulness of Thompson's group F in public key cryptography). We also have to keep in mind that exponents also increase the length of words.

To understand the form that a dead end element takes, it is helpful to look at the tree-pair diagrams which correspond to these dead end elements and understand the form which they take. Obviously, to understand how the dead end elements work in the tree-pair diagrams, we need to know how to calculate word length from a tree-pair diagram.

We will examine a method of calculating word length which is known as Fordham's method. It involves an algorithm in which one labels the carets a certain way and pairs up corresponding carets in T_- to T_+. One then reads values off a table to see how each caret-pair is *weighted* and the sum of all the caret-pair weights is the word length, [7]. A very in-depth treatment of the algorithm is also given in Fordham's paper, [11].

First we examine how to label the carets. We start by numbering them in some ordered way. For simplicity, we begin with 1 and go from left to right

(see diagram). An intuitive way to think about it is to place the number "n" in some sense *between* the leaves "$n-1$" and "n". Once we have done this, we divide the carets into the following disjoint types:

1. L_0. The first caret on the left side of the tree. In other words, caret number 1.
2. L_L. Any left caret which isn't L_0. By convention, the root caret of the tree (the one in the middle) is considered to be a left caret.
3. I_0. An interior caret which has no right child.
4. I_R. An interior caret which has a right child.
5. R_I. Any right caret (on the right side of the tree, as discussed above) numbered k with the property that caret $k+1$ is an interior caret.
6. R_{NI}. Any right caret numbered k with the property that there exists a caret numbered $k+2$ or greater which is an interior caret.
7. R_0. A right caret with no higher-numbered interior carets.

We then look at each individual caret and check which type they are in T_- and T_+. We look up the weights in table 1.

Table 1: Table of weights of pairs of carets

	R_0	R_{NI}	R_I	L_L	I_0	I_R
R_0	0	2	2	1	1	3
R_{NI}	2	2	2	1	1	3
R_I	2	2	2	1	3	3
L_L	1	1	1	2	2	2
I_0	1	1	3	2	2	4
I_R	3	3	3	2	4	4

Let us look at a very simple example: x_2.

The table for this particular example is quite simple (see Table 2).

One of the first things that one notices is that every tree has exactly one L_0 which will necessarily map to L_0 and always has a weight of 0. Every tree

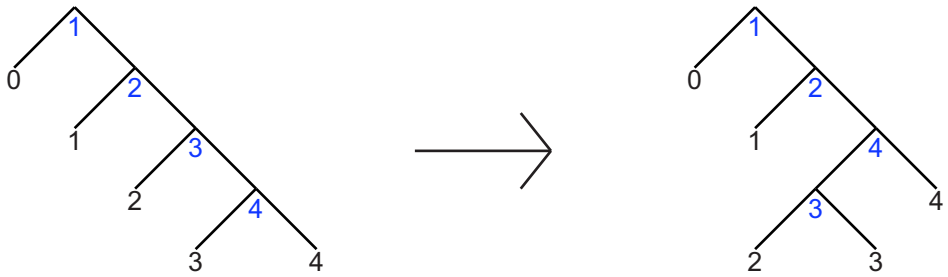

Figure 17: The tree pair diagram for x_2 with carets labelled.

Table 2: Table of caret weights for x_2

Caret number	T_-	T_+	weight
1	L_0	L_0	0
2	R_0	R_I	2
3	R_0	I_0	1
4	R_0	R_0	0
		Total	3

will also have at least one R_0 which will also necessarily map to an R_0 in the other tree in the pair.

Notice also that the total isn't 1 as we expect. This is because it gives the word length in terms of the generators x_0 and x_1. In this case, x_2 is, when expressed in terms of the generators, $x_0^{-1}x_1x_0$ and the length of that word is clearly 3.

Let us now try a *slightly* more complicated example, $x_2^2x_1^{-1}$.

The table (Table 3) for this example is slightly more involved.

We now get a total of 5. This is because, when we expand it out $x_2^2x_1^{-1}$ becomes $x_0^{-1}x_1x_0x_0^{-1}x_1x_0x_1^{-1}$. Notice that a x_0 and a x_0^{-1} appear next to each other and can thus be canceled out giving $x_0^{-1}x_1x_1x_0x_1^{-1}$ whose word length is clearly 5.

Let us now look at a more complicated example. We take the tree pair diagram for $x_0^2x_1x_6x_3^{-1}x_0^{-2}$. We proceed to label the diagram according to the

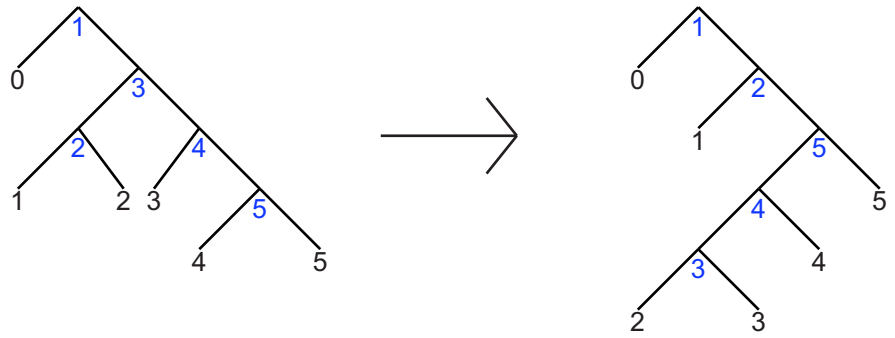

Figure 18: The tree pair diagram for $x_2x_1^{-1}$ with carets labelled.

Table 3: Table of caret weights for $x_2x_1^{-1}$

Caret number	T_-	T_+	weight
1	L_0	L_0	0
2	I_0	R_I	3
3	R_0	I_0	1
4	R_0	I_0	1
5	R_0	R_0	0
		Total	5

algorithm above.

We then draw the table (Table 4) to determine which caret types go to which.

This is a little bit more difficult to verify, but we can easily show that when we expand $x_0^2x_1x_6x_3^{-1}x_0^{-2}$ into generators, we get 11.

$$\begin{aligned} & x_0^2x_1x_6x_3^{-1}x_0^{-2} \\ = \; & x_0^2x_1(x_1^{-1}x_0^{-2}x_{0\text{or}1}^{-1}x_0^{-1}x_1x_0x_{0\text{or}1}x_0^2x_1)(x_1^{-1}x_0^{-1}x_1^{-1}x_0x_1)x_0^{-2} \\ = \; & x_{0\text{or}1}^{-1}x_0^{-1}x_1x_0x_{0\text{or}1}x_0^1x_1^{-1}x_0x_1x_0^{-2} \end{aligned} \tag{10}$$

(remembering $\overbrace{x_{i_1}^{-1}x_{i_2}^{-1}\ldots x_{i_n}^{-1}x_0^{-1}}^{k-1 \text{ of these}} x_1 \overbrace{x_0x_{i_n}x_{i_{n-1}}\ldots x_{i_1}}^{k-1 \text{ of these}}$ from the earlier sections)

Although it may seem fairly arbitrary which of x_0 or x_1 we decide to "un conjugate" our x_n's with, there is a method to it. This is outlined more explicitly in the last section of this paper, where the word problem in Thompson's

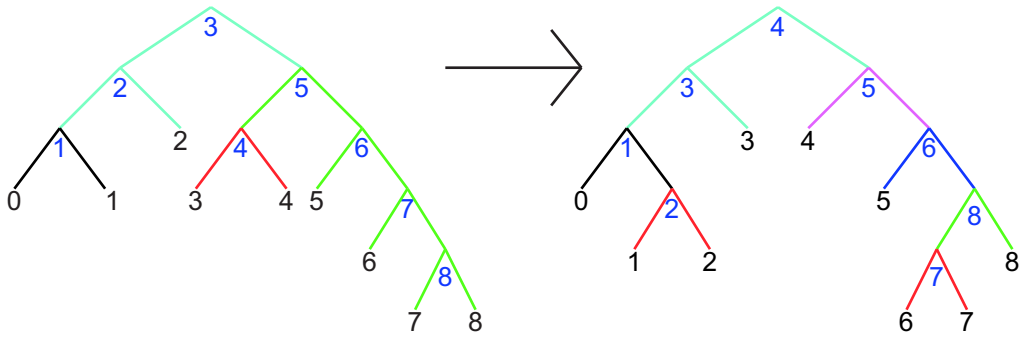

Figure 19: We introduce a colour scheme to aid us in distinguishing caret types. The black caret represents L_0, light blue L_L, dark blue R_I, green R_0, red I_0 and purple R_{NI} (type I_R is not represented in this example).

Table 4: Table of caret weights for $x_0^2 x_1 x_6 x_3^{-1} x_0^{-2}$

Caret number	T_-	T_+	weight
1	L_0	L_0	0
2	L_L	I_0	2
3	L_L	L_L	2
4	I_0	L_L	2
5	R_0	R_{NI}	2
6	R_0	R_I	2
7	R_0	I_0	1
8	R_0	R_0	0
		Total	11

group is discussed. Here we used the flexibility in the x_6 to eliminate the $x_0^2 x_1$ sequence to illustrate the point that there is no simple algorithm for determining word length just from taking a word in its normal form.

Finally, we know that each reduced tree pair diagram in a unique representative of an element in F, but what assurances have we that Fordham's method gives the *minimum length* of the element?

We begin with a lemma introducing criteria for minimal length which we will denote by l. We will also denote φ as any function which gives the length of a word (though not necessarily the minimal length).

Lemma 3.1. *Given a presentation of a group G and a function $\varphi : G \to \{0, 1, 2, \ldots\}$, if φ has the properties:*

1. $\varphi(\mathrm{id}_G) = 0$ *(id_G is the identity element in G);*
2. $\varphi(g) = 0 \Rightarrow g = \mathrm{id}_G$*;*
3. *if $g \in G$ and x is a generator of G then $\varphi(g) - 1 \leq \varphi(gx)$;*
4. *for any non-identity element $g \in G$, there is at least one generator x of G such that $\varphi(gx) = \varphi(g) - 1$;*

then $\varphi(g) = l(g)$ for all $g \in G$.

Theorem 3.2. *For any element $f \in F$, the word length given by Fordham's method is the minimum length word.*

The proof of this theorem relies on verifying Lemma 3.1, a long and arduous process which is covered in detail in Fordham [11], as is the proof of the lemma itself.

3.2 Dead end elements

Why do we need such a long and convoluted way of measuring word length when we can simply expand our elements to be in terms of the generators and count off the word length? As we saw above, the algorithm for expanding any given word to arrive at a minimal word length is a fairly nontrivial process whereas Fordham's method relies only on the geometric properties of the tree pair diagrams which whose form is dictated by the normal form.

3.2.1 An example of a dead end element

Let us begin by examining an example of such a dead end element: $x_0^2x_1x_6x_3^{-1}x_0^{-2}$. This is an example of a dead end element of length 11. We can demonstrate this as follows by right multiplying the word by all the generators.

$$\begin{aligned} & x_0^2x_1x_6x_3^{-1}x_0^{-2}x_0 \\ = \; & x_0^2x_1x_6x_3^{-1}x_0^{-1} \end{aligned}$$

Whose length is clearly less than the original word

$$\begin{aligned} & x_0^2x_1x_6x_3^{-1}x_0^{-2}x_0^{-1} \\ = \; & x_0^2x_1(x_0^{-5}x_1x_0^2x_1^{-1}x_0^3)x_0^{-2}x_0^{-1} \\ = \; & x_0^2x_1x_0^{-5}x_1x_0^2x_1^{-1} \\ = \; & x_0^2x_1x_0^{-3}x_3x_1^{-1} \\ = \; & x_0^2x_1x_0^{-3}x_1^{-1}x_0^{-1}x_1x_0x_1x_1^{-1} \\ = \; & x_0^2x_1x_0^{-3}x_1^{-1}x_0^{-1}x_1x_0 \end{aligned}$$

Conveniently, in terms of the generators. The word length is 10, which is less than 11.

$$\begin{aligned} & x_0^2x_1x_6x_3^{-1}x_0^{-2}x_1 \\ = \; & x_0^2x_1x_6x_3^{-1}x_3x_0^{-2} \\ = \; & x_0^2x_1x_6x_0^{-2} \end{aligned}$$

Which is, again, clearly shorter than the original word.

$$\begin{aligned} & x_0^2x_1x_6x_3^{-1}x_0^{-2}x_1^{-1} \\ = \; & x_0^2x_1x_6x_3^{-1}x_3^{-1}x_0^{-2} \\ = \; & x_0^2x_1x_6x_0^{-2}x_1^{-2}x_0^2x_0^{-2} \\ = \; & x_0^2x_1x_0^{-2}x_1^{-2}x_0^{-1}x_1x_0x_1^2x_0^2x_0^{-2}x_1^{-2} \\ = \; & x_0^2x_1x_0^{-2}x_1^{-2}x_0^{-1}x_1x_0 \end{aligned}$$

We can read the length off as 10, which is, once again, less than 11. Refer to Table 5 for a list of normal forms for each of the above elements.

Table 5: Table of normal forms

	Normal form
wx_0	$x_0^2x_1x_6x_3^{-1}x_0^{-1}$
wx_0^{-1}	$x_0^2x_1x_6x_3^{-1}x_0^{-3}$
wx_1	$x_0^2x_1x_6x_0^{-2}$
wx_1^{-1}	$x_0^2x_1x_6x_3^{-2}x_0^{-2}$

The explicit calculation by Fordham's method is presented at the end of this subsection, after some discussion on how the generators act on the tree diagrams of the elements of F.

To understand how this works, we must first understand how, graphically, the generators x_0 and x_1 act on trees. We already know what their tree pair diagrams look like. Recall:

(a) x_0 (b) x_1

Figure 20: Tree pair diagrams for x_0 and x_1.

Rather curiously, when they are applied to larger trees, they do the same things but on a larger scale. The Figure 21 illustrates this.

(a) x_0 acting on a bigger tree. (b) x_1 acting on a bigger tree.

Figure 21: Diagrams showing how x_0 and x_1 act on larger trees.

In Figure 21, the little triangles with the letters a, b, c or d in them represent (possibly empty) subtrees.

The reason this is important is that when we consider elements w and wx_0 for example, we need to make note of what happens to the tree pair for w and how the types of carets that w carries is affected by right multiplying with a generator.

Table 6 shows us what happens given different conditions on the starting tree (T_-) and observing what happens to a particular caret. First we consider elements $w = (T_-, T_+)$ and wx_0. Caret C is the root caret of T_-.

Table 6: The effect of right multiplying w by x_0

Condition on T_-	Initial type of caret C	New type of caret C	Increase if C paired with	Decrease if C paired with
$S_{RL} \neq \emptyset$	L_L	R_I	R_*,I_*	L_L
$S_{RL} = \emptyset$,$S_{RR} \neq \emptyset$	L_L	R_{NI}	R_*,I_R	L_L,I_0
$S_{RL} = \emptyset$,$S_{RR} = \emptyset$	L_L	R_0	R_{NI},R_I,I_R	R_0,L_L,I_0

We see a similar table for $w = (T_-, T_+)$ and wx_0^{-1}. Caret C is the caret C_R of T_-.

Table 7: The effect of right multiplying w by x_0^{-1}

Condition on T_-	Initial type of caret C	New type of caret C	Increase if C paired with	Decrease if C paired with
$S_{RL} \neq \emptyset$	R_I	L_L	L_L	R_*,I_*
$S_{RL} = \emptyset$,$S_{RR} \neq \emptyset$	R_{NI}	L_L	L_L,I_0	R_*,I_R
$S_{RL} = \emptyset$,$S_{RR} = \emptyset$	R_0	L_L	R_0,L_L,I_0	R_{NI},R_I,I_R

It is perhaps not surprising that the tables are essentially the same up to a swap of columns 2 and 3, and columns 4 and 5. The following are the tables for wx_1 and wx_1^{-1}. Caret C is the caret C_{RL} in the first instance and C_R in the second.

Table 8: The effect of right multiplying w by x_1

Condition on T_-	Initial type of caret C	New type of caret C	Increase if C paired with	Decrease if C paired with
$S_{RLR} \neq \emptyset$	I_R	R_I	none	any
$S_{RLR} = \emptyset, S_{RR} \neq \emptyset$	I_0	R_{NI}	R_0, R_{NI}	L_L, I_*, R_I
$S_{RLR} = \emptyset, S_{RR} = \emptyset$	I_0	R_0	R_{NI}	L_L, I_*, R_I, R_0

Again, we expect the column pairs 2,3 and 4,5 to swap, as indeed they do.

Table 9: The effect of right multiplying w by x_1^{-1}

Condition on T_-	Initial type of caret C	New type of caret C	Increase if C paired with	Decrease if C paired with
$S_{RRL} \neq \emptyset$	R_I	I_R	any	none
$S_{RRL} = \emptyset, S_{RRR} \neq \emptyset$	R_{NI}	I_0	L_L, I_*, R_I	R_0, R_{NI}
$S_{RRL} = \emptyset, S_{RRR} = \emptyset$	R_0	I_0	L_L, I_*, R_I, R_0	R_{NI}

Reasoning carefully through the above tables, we find that we can write down a general form for dead end elements in terms of tree pair diagrams. Which the reader will notice is the same form as the tree pair diagram which $w = x_0^2 x_1 x_6 x_3^{-1} x_0^{-2}$ takes.

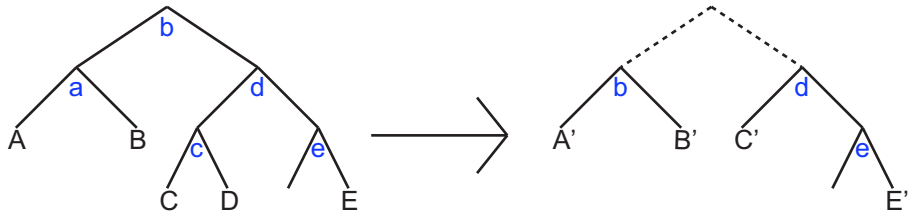

Figure 22: The general form of dead end elements. Capital letters represent possibly empty subtrees with the exception of E and E' which are nonempty. Consequently, caret d is of type R_{NI} in at least one tree, [7].

Now that we have a better idea of how the generators act on elements of the groups in terms of tree pair diagrams, let us return to our concrete example and examine the tree pair diagrams for wx_0, wx_0^{-1}, wx_1 and wx_1^{-1} and see

how they compare to w where w is $x_0^2 x_1 x_6 x_3^{-1} x_0^{-2}$.

Lets see how the tables from above (which told us which pairs C had to be to increase/decrease) relate to our example. Also included with each tree pair diagram is the table of caret pairs and weights to get a sense of how each multiplication by a generator affects the word length. These tables are also useful to get a quantitative feel for how the word length is directly affected.

Note: the word length is always 10 because in order for w to be a dead end element, the length of the word must decrease when right-multiplied by any of $x_0^{\pm 1}$ or $x_1^{\pm 1}$. Since the space is discrete $|w|$ must decrease by exactly 1 in order to satisfy the triangle inequality (because it is also a metric space).

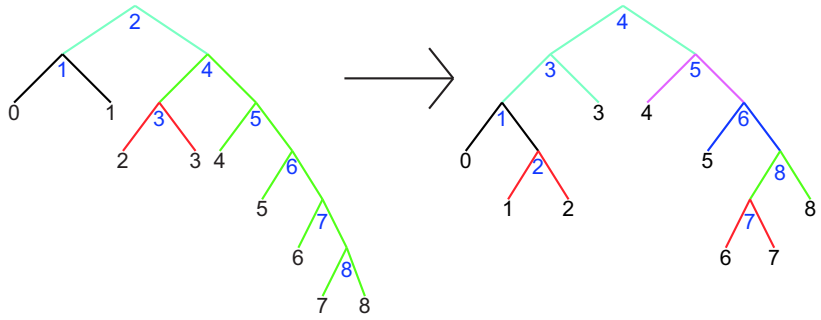

Figure 23: The tree pair diagram for wx_0.

Table 10: Table for wx_0

Caret number	T_-	T_+	weight
1	L_0	L_0	0
2	L_L	I_0	2
3	I_0	L_L	2
4	R_0	L_L	1
5	R_0	R_{NI}	2
6	R_0	R_I	2
7	R_0	I_0	1
8	R_0	R_0	0
		Total	10

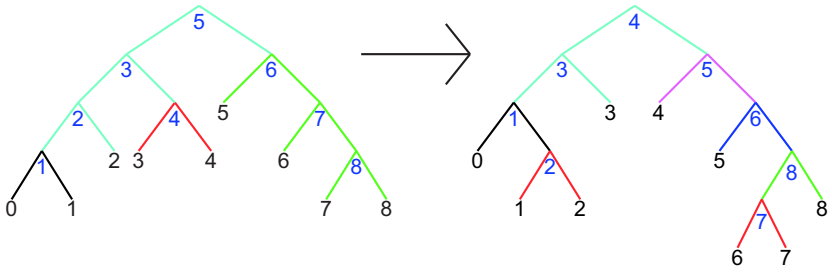

Figure 24: The tree pair diagram for wx_0^{-1}.

Table 11: Table for wx_0^{-1}

Caret number	T_-	T_+	weight
1	L_0	L_0	0
2	L_L	I_0	2
3	L_L	L_L	2
4	I_0	L_L	2
5	L_L	R_{NI}	1
6	R_0	R_I	2
7	R_0	I_0	1
8	R_0	R_0	0
		Total	10

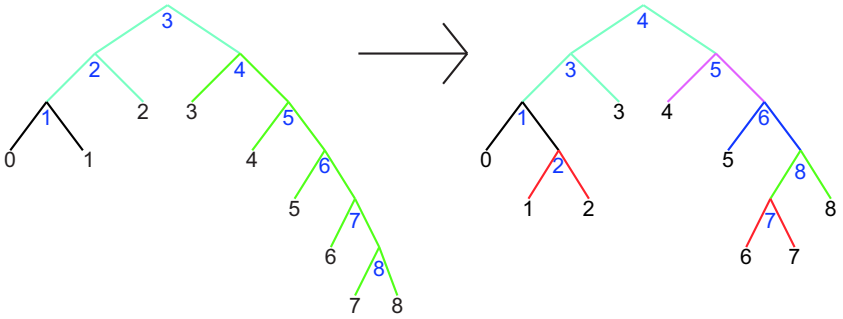

Figure 25: The tree pair diagram for wx_1.

Table 12: Table for wx_1

Caret number	T_-	T_+	weight
1	L_0	L_0	0
2	L_L	I_0	2
3	L_L	L_L	2
4	R_0	L_L	1
5	R_0	R_{NI}	2
6	R_0	R_I	2
7	R_0	I_0	1
8	R_0	R_0	0
		Total	10

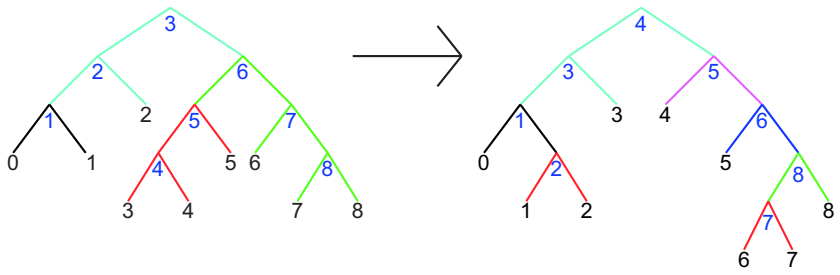

Figure 26: The tree pair diagram for wx_1^{-1}.

Table 13: Table for wx_1^{-1}

Caret number	T_-	T_+	weight
1	L_0	L_0	0
2	L_L	I_0	2
3	L_L	L_L	2
4	I_0	L_L	2
5	I_0	R_{NI}	1
6	R_0	R_I	2
7	R_0	I_0	1
8	R_0	R_0	0
		Total	10

3.3 Seesaw words in F

In addition to dead-end elements, there is also another class of elements which, in many ways behave very much like dead end elements. The difference is, however, that the length of the word is only reduced if you add *one* of the generators from the finite presentation of F rather than *either* of them. In addition to this, once reduced, right-multiplying by that same generator will continue to reduce the word some k number of times.

Definition 3.3. *An element w in a finitely generated group G with finite generating set X is a* seesaw word of swing k with respect to a generator g *if the following conditions hold. Let $|w|$ represent the word length of w with respect to the generating set X.*

1. *Right multiplication by both g and g^{-1} reduces the word length of w; that is, $|wg^{\pm 1}| = |w| - 1$, and for all $h \in X$ $g^{\pm 1}$, we have $|wh^{\pm 1}| \geq |w|$.*
2. *$|wg^l| = |wg^{l-1}| - 1$ for some $l \in 1, \ldots, k$ and $|wg^m h^{\pm 1}| \geq |wg^m|$ for all $h \in X$ g and integral $m \in 1, \ldots, k-1$.*
3. *$|wg^l| = |wg^{l-1}| - 1$ for some $l \in 1, \ldots, k$ and $|wg^m h^{\pm 1}| \geq |wg^m|$ for all $h \in X$ g for integral $m \in 1, \ldots, k-1$.*
4. *$|wg^{-l}| = |wg^{-l+1}| - 1$ for some $l \in 1, \ldots, k$ and $|wg^{-m} h^{\pm 1}| \geq |wg^{-m}|$ for all $h \in X$ g^{-1} for integral $m \in 1, \ldots, k-1$.*

These are called seesaw words because they behave like a balanced seesaw. When in balance, there is a two-way choice about which way to go down, but once that initial choice is made, there is only the inexorable descent downward by the same generator for a large number of steps determined by swing.

An interesting property of these seesaw words is that Thompson's group F contains seesaw words of arbitrarily large swing with respect to the generator x_0 in the standard generating set x_0, x_1, [8].

4 Amenability

One of the enduring questions regarding Thompson's group is that of whether or not it is *amenable*. Let us first discuss what it means for a group to be amenable.

4.1 Definition

Definition 4.1. *A group is G* **amenable** *if there is a* left-invariant measure *μ on G which is finitely additive and has total measure 1. In other words, if there is a function μ which sends the set of subsets of G to the interval [0,1] which satisfies the following conditions*

1. *$\mu(gA) = \mu(A)$ for all $g \in G$ and all subsets A of G.*
2. *$\mu(G) = 1$, and*
3. *$\mu(A \cup B) = \mu(A) + \mu(B)$ if A and B are disjoint subsets of G.*

To those of us who haven't studied much measure theory (and even for some of us who have) this particular definition is of limited usefulness. We want an equivalent but more intuitive formulation of the concept of amenability.

There is a more geometric way of interpreting the notion of amenability and it has to do with ratios of the *volume* of a subset of a group to the *volume* of its boundary (which we may want to think of in terms of length or surface area, although this doesn't actually make much sense in most cases). A group is thus amenable if it can be expressed as a family of subsets of finite volume which have boundaries which also have finite volume such that the volume of all the boundaries divided by the volume of the subsets themselves tends to 0, [1].

Roughly speaking, a group is amenable if the probability of a random walk of length L returning to 1 decreases more slowly than exponentially with L, [4].

There exists a stronger version of amenability – a group is said to be **elementary amenable** if it can be built up from finite groups and abelian

groups by a sequence of simple operations that result in amenable groups when applied to amenable groups. More formally:

Definition 4.2. *The class of* **elementary amenable groups** *is the smallest subclass of the class of all groups that satisfies the following conditions:*

1. *It contains all finite and all abelian groups*
2. *If G is in the subclass and H is isomorphic to G, then H is in the subclass*
3. *It is closed under the operations of taking subgroups, forming quotients, and forming extensions*
4. *It is closed under directed unions*

Remark: every elementary amenable group is amenable – however the converse is not true)

For some time it was unknown whether there existed amenable groups that are not elementary amenable. Thompson's group was thought to be a candidate for such a group. It is now known that such groups exist, see [2].

It is known that Thompson's group F is not elementary amenable, [4].

4.2 An easy example

Perhaps it would be helpful to demonstrate amenability with an example. First, let us define more precisely what a Cayley graph is.

Definition 4.3. *Let G be a group, and let $S \subseteq G$ be a set of group elements such that the identity element $\notin S$.*

The Cayley graph $\mathcal{C}$ is defined as the directed graph *having one vertex associated with each group element and directed edges (g,h) whenever $gh^{-1} \in S$.*

The Cayley graph depends on the choice of set S, and is connected if and only if S generates G.

Take the Cayley graph of$\mathbb{Z}^2$ (Fig 27). We will take a ball of radius n centred about some origin as our family of subsets, that is, if 0 is the origin, all elements g which satisfy $d(0, g) \leq n$. We define the interior as all points h in the set for which every point j satisfying $d(h, j) = 1$ means that j is also in the set. The boundary is all the points which are in the set but not in the interior.

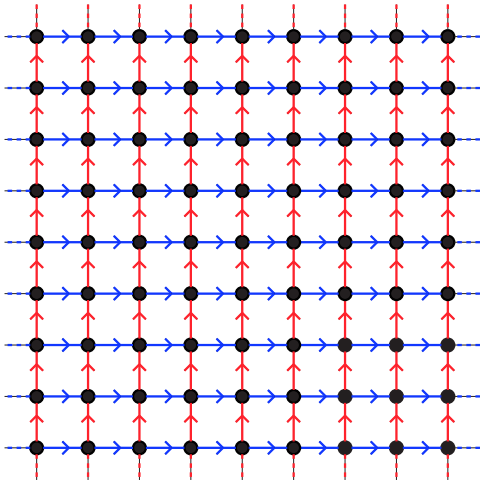

Figure 27: The cayley graph of $\mathbb{Z}^2$ which has the presentation $\langle a, b | a = b \rangle$ the red arrows represent the generator a and the blue arrows represent the generator b.

Looking at balls of radius 2 (Fig. 28(a)) and 3 (Fig. 28(b)) in $\mathbb{Z}^2$, it can be easily seen that, given any radius n, the rule for the *volume* of the boundary

is $4n$ and the *volume* of the interior is $1+4\times((n-1)+(n-2)+\ldots+2+1) = 1+2n(n-1)$. The limit:

$$\begin{aligned} \lim_{n\to\infty} \tfrac{\textit{boundary}}{\textit{interior}} &= \lim_{n\to\infty} \tfrac{4n}{1+2n(n-1)} \\ &= 0 \end{aligned} \tag{11}$$

Therefore the group $\mathbb{Z}^2$ is amenable.

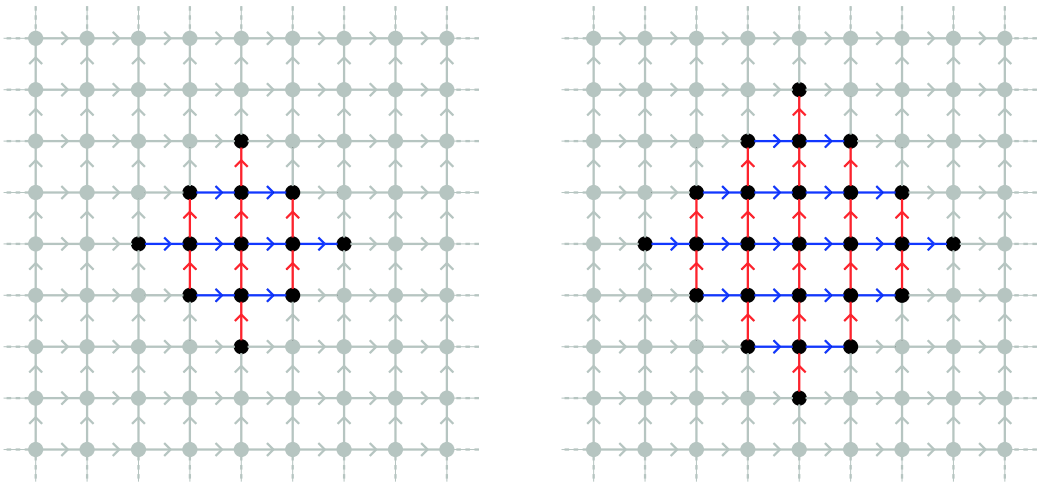

(a) A ball of radius 2 in $\mathbb{Z}^2$.

(b) A ball of radius 3 in $\mathbb{Z}^2$.

Figure 28: The balls of radius 2 and 3 highlighted on the Cayley graph of $\mathbb{Z}^2$.

4.3 Properties of the Cayley graph of F

The probability of random walks returning to the origin is closely tied with the growth and density of a group, properties which can be more easily studied by looking at Cayley graphs (the density of a graph is just the average valence of the vertices). In [10], Guba introduces a variation on the Cayley graph Γ_n which preserves most of the structure and cuts out much of the clutter. We will give a brief examination of this construction.

The Cayley graph of Thompson's group F is not an easy thing to comprehend. The author's various attempts at drawing it have all ended unhappily and with much frustration. The main reason for this is because of the way in which Thompson's group expands exponentially on its two generators (x_2 can be represented by $x_0^{-1}x_1x_0$, x_3 can be represented by $x_0^{-2}x_1x_0^2$ or $x_1^{-1}x_0^{-1}x_1x_0x_1$, x_4 can be represented by $x_0^{-3}x_1x_0^3$, $x_1^{-1}x_0^{-2}x_1x_0^2x_1$, $x_0x_1^{-1}x_0^{-1}x_1x_0x_1x_0$, or $x_1^{-2}x_0^{-1}x_1x_0x_1^2$ and so on. In general x_n can be written in 2^{n-2} different ways).

We wish to study some properties of the Cayley graph but since the Cayley graph itself is so difficult to comprehend all at once, we slowly build it up.

We draw an arrow and label it x_1. We can't conjugate an element by any other element with a subscript less than 1, so we stop and label this graph $\mathcal{C}_1$. Next, we draw an arrow and label it x_2. We can conjugate x_1 by x_0 to get x_2, so we add arrows to our graph to represent this. There are no other possibilities, so we stop and label this graph $\mathcal{C}_2$. $\mathcal{C}_3$ is slightly more complicated. We draw our x_3 arrow. x_2 can be conjugated be either x_0 or x_1 to get x_3 so we draw those respective arrows. x_2 may also be arrived at by conjugating x_1 by x_0 so we add that to our graph. Since there are no more possibilities for conjugating any of the elements in a different way, we stop and label this graph $\mathcal{C}_3$. This process can be continued indefinitely to construct $\mathcal{C}_n$. Note that arrows may *not* be added if traversing those arrows takes you *further* from the edge labelled x_n. This is our starting point. The first three $\mathcal{C}_n$'s are shown in Figure 29. Incidentally, these graphs are subgraphs of the Cayley graph of Thompson's group F and $\mathcal{C}_\infty$ is the Cayley graph of F with respect to its infinite presentation.

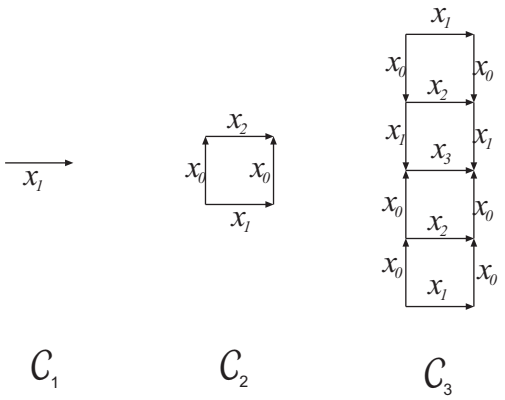

Figure 29: The first three $\mathcal{C}_n$'s.

Although not immediately obvious from Figure 29, once we get to $\mathcal{C}_4$, $\mathcal{C}_5$ and higher, these graphs become very messy and unmanageable, making it difficult to see what is happening.

We proceed to modify our graph in such a way so that it preserves the overall structure and allows us to see more clearly what is going on. We define it pictorially in Figure 30 then more precisely later.

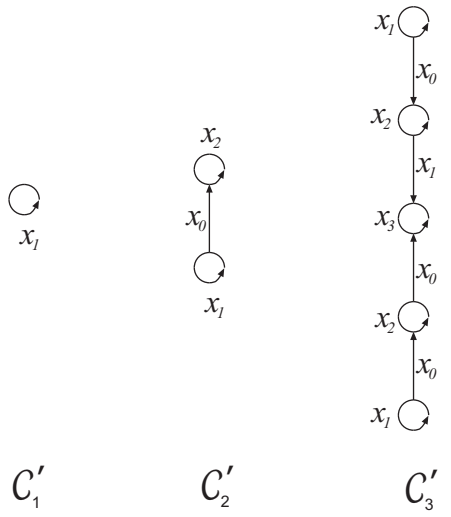

Figure 30: Diagram illustrating the intermediate step between n-subgraphs $\mathcal{C}_n$ and corresponding Γ_n's. $\mathcal{C}_1'$ corresponds to $\mathcal{C}_1$, $\mathcal{C}_2'$ to $\mathcal{C}_2$ and $\mathcal{C}_3'$ to $\mathcal{C}_3$

We simply take the horizontal edges in our subgraphs and collapse them into loops. Strictly speaking, we are *identifying* vertices at either end of these horizontal edges and forming loops, although we must keep in mind that the Γ_n we have produced is not a Cayley graph, nor does it behave much like one. Loosely speaking, we may think of the directed vertices labelled by x_i as representing the action of "conjugating by x_i".

We now construct a new visualisation, denoted by Γ_n, which has a bijective correspondence to the n-subgraphs as defined above, but which is *not* itself a Cayley graph (or subgraph of one). Guba in [10] uses this new visualisation to study properties of the growth and density of the Cayley graph. Only a rough outline is given below, for more details, refer to [10].

One can attempt to draw the graphs Γ_n $(n \geq 1)$ explicitly. If a vertex v has a loop at v labelled by x_m, it will later become clear that the valence of v is equal to m. We draw this vertex as a circle with the number m inside. If Y is a labelled graph with labels of the form x_j $(j \geq 0)$, then by $\Psi(Y)$ we denote the graph obtained from Y by increasing all subscripts of the labels by 1. We know that Γ_1 is a single loop labelled by x_1 (from $\mathcal{C}_1'$, see Figure 30). To obtain Γ_{n+1} from Γ_n $(n \geq 1)$, one has to apply Ψ to Γ_n and then attach the requisite number of x_0's and x_1's. Γ_n now begins to look a little more manageable. We can see in Figure 31 that Γ_n expands very rapidly with n.

This much simpler graph allows us to more easily see what happens as Γ_n *grows*. Although this is not the Cayley graph of Thompson's group, due to its preservation of the overall structure of the group, we can use it to determine more "structural" things such as growth and density (for more information, refer to Guba 31). We now present two lemmas useful to the understanding of growth, for the proofs the reader is referred to Guba, [10].

Lemma 4.4. *Let a_{nk} $(1 \leq k \leq n)$ be the number of verticies of Γ_n that have valence k. Then*

$$a_{nk} = \frac{k(2n-k-1)!}{(n-k)!n!} \tag{12}$$

It follows that the total number of verticies in Γ_n equals the nth catalan number, that is,

$$\frac{(2n)!}{n!(n+1)!} \tag{13}$$

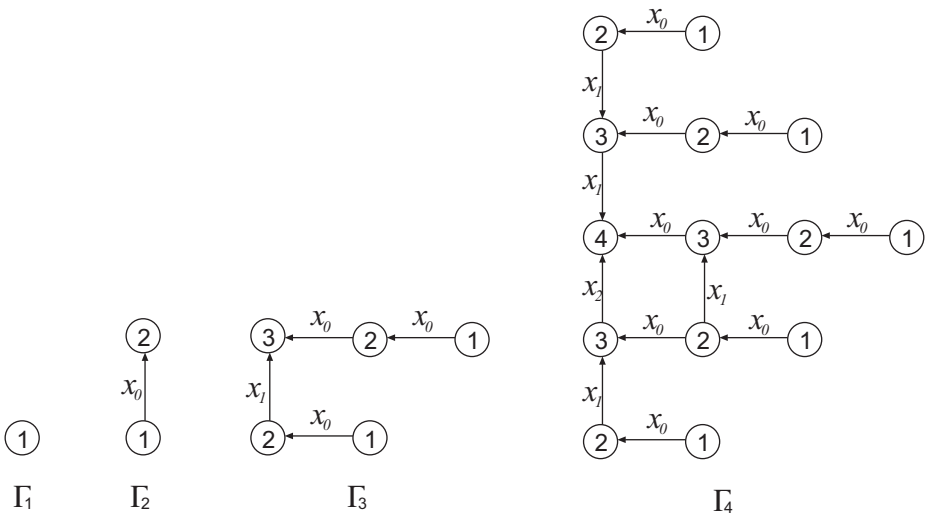

Figure 31: Diagram illustrating the expansion of the Γ_n.

We observe that when n reaches some k, the vertex representing the loop x_{k-2} has no more vertices added to it confirming that, for any circle with the number m in it, its valence will be m for sufficiently large n, as we expect.

Now if we let b_{nk} be the number of edges in Γ_k labelled by x_k, where $0 \leq k \leq n$. (Recall that if we denote a vertex by a circle with the number m inside, then this vertex has a loop labelled by x_m.)

Lemma 4.5. *For any $n \geq 2$, and $b_{10} = 0$, $b_{11} = 1$*

$$b_{n0} = b_{n1} = \frac{3(2n-2)!}{(n-2)!(n+1)!} \tag{14}$$

Studying growth rates and density are important because such things may lead to conclusions about amenability. For example:

Theorem 4.6. *The Cayley graph of a group with two generators is strongly non-amenable if and only if the density of the graph does not exceed 3.*

(Proof in 31). There are other ways of looking at the growth and density of the Cayley graph, such as through the use of computers as will be discussed in the next subsection.

4.4 Computational explorations

In order to get an intuitive feel for whether or not Thompson's group is amenable, people have studied the Cayley graph of the group using computers to try to manually *bash* out some kind of result. The Cayley graph of Thompson's group F is difficult to visualise and comprehend (indeed, most Cayley graphs tend to be difficult to visualise, at least for humans) so simply "trying out" some numbers is sometimes an instructional way of getting a feel for the problem. It is known that Thompson's group F has exponential growth, but the growth rate is unfortunately unknown at this time, [4].
We make use of the following equation:

$$p(L) = \frac{\#T(L)}{(2m)^L} \tag{15}$$

and the definition given in the previous subsection that refers to the probability of a random walk of length L returning to 1 decreasing more slowly than exponentially. $p(L)$ is simply the probability of a random walk of length L returning to its origin. In the case of F, it is the proportion of words which are equal to the identity.

This can be reformulated in a useful way in the following theorem:

Theorem 4.7. (Kesten) *A group is amenable if and only if*

$$\limsup_{L\to\infty} p(L)^{1/L} = 1$$

The direct approach at finding the numbers $p(L)$ for Thompson's group F fails even at quite small values of L due to the fact that the number of words grows exponentially, meaning the computational times get large very quickly. For example, there are 268435456 words of length 14, out of which there are 1988452 representing the identity giving a value of $p(14)^{1/14} = 0.704423677$, [4].

To summarize the methods of Burillo, Cleary and Wiest in [4], they begin by taking *samples* of words of a given length rather than *all* the words of a given length. They further improve on this by taking *balanced* words: words which have a total exponent of zero in both generators x_0 and x_1. So they considered not *all* random words, but only balanced ones. It is of note that

the ableianisation of F, (denoted F_{ab}), is $\mathbb{Z}^2$, generated by x_0 and x_1, so being balanced is equivalent to representing the trivial element in $\mathbb{Z}^2$. This gives a new expression to consider based on the samples and balanced words. Let $C(L)$ be the set of balanced words among the 4^L non-reduced words of length L in F_2, and define

$$\hat{p}(L) = \frac{\#T(L)}{\#C(L)}$$

to be the proportion of words representing the identity of F among balanced words of length L. We thus end up with (for more detailed working, refer to Burillo, Cleary and Wiest [4].)

$$\sqrt[L]{p(L)} = \sqrt[L]{\hat{p}(L)}.\sqrt[L]{\frac{\#C(L)}{4^L}}$$

Moreover, $\sqrt[L]{\frac{\#C(L)}{4^L}}$ tends to 1 as L tends to infinity, because $\mathbb{Z}^2$ is amenable (shown in the previous subsection). Thus F is amenable if and only if we have

$$\limsup_{L\to\infty} \hat{p}(L)^{1/L} = 1$$

The authors of [4] continue with what are essentially efficiency maximising steps to reduce the computational time required. They proceeded to run several testing algorithms on the "Wildebeest" 132-processor Beowulf cluster at the City University of New York. For further results, refer to [4].

It is still not known whether or not Thompson's group is amenable. However, it is known that Thompson's group is not *elementary amenable*

If it is found that Thompson's group F is amenable, then it would be an example of a finitely presented amenable but not elementary amenable group. If, however, it is found that Thompson's group F is not amenable then it would be an example of a finitely presented non-amenable group without free non-abelian subgroups, [4].

5 Applications in Cryptography

5.1 Introduction

The famous mathematician G. H. Hardy used to delight in the knowledge that the work that he did had no practical application at the time. He pursued "pure" mathematics for its own sake and not for anyone else's. Some of his work focused on number theory, in particular, factorising large numbers. It was well known at the time that multiplying two very large prime numbers together was much easier than taking a product of two very large primes and determining its prime factorisation. At the time, this had very little practical application. Now, it is the cornerstone of modern public key cryptography, which permeates through almost every area of our lives which is affected by computers, a world in which data security is becoming increasingly important, [13].

Perhaps not surprisingly, we find that Thompson's group also has applications in public key cryptography. Recently non-commutative (semi)groups have been investigated as an alternative to the standard public key exchange systems such as RSA. These use symbolic computation rather than numeric computation and take advantage of the *conjugacy search problem* in groups. This is simply a ramification of the discrete logarithm problem, well known in cryptographic protocols such as El Gamal and DSA, [12].

The conjugacy search problem works as follows: Take two elements a and b in a group G and you are given the information that $a^x = b$ for some $x \in G$. The problem is to find at least one particular x like that. In the case of Thompson's group, a^x is simply $x^{-1}ax$. The computational difficulty of this problem in some particular groups has been used in several group based cryptosystems, [12].

5.2 Summary of basic public-key cryptography

First, we should introduce some cryptographic primitives.

Definition 5.1. *A secret key cipher is a 5-tuple $(\mathcal{P}, \mathcal{C}, \mathcal{K}, \mathcal{E}, \mathcal{D})$ where the following conditions hold.*

1. *$\mathcal{P}$ is a finite set of possible **plaintexts***
2. *$\mathcal{C}$ is a finite set of possible **ciphertexts***
3. *$\mathcal{K}$ is a finite set of possible keys, also called the **keyspace***
4. *For each $K \in \mathcal{K}$, there is an encryption rule $e_K \in \mathcal{E}$ and a corresponding decryption rule $d_K \in D$. Each $e_K : \mathcal{P} \to \mathcal{C}$ and $d_K : \mathcal{C} \to \mathcal{P}$ are functions such that*

$$d_K(e_K(x)) = x, \qquad \forall x \in \mathcal{P}.$$

(Taken from the author's lecture notes from the 2006 ICE-EM/AMSI summer school, [5]).

As a general rule, we talk about "Alice" and "Bob", who are the two communicators (notice how their names begin with different letters to aid us in simplifying our system of abbreviation). Eve is an eavesdropper and/or Oscar (opponent) are adversaries. All cryptosystems also follow **Kerchoff's principle** which states that the designer of the system has to assume that the only unknown about the cryptosystem is the key. The object is usually to obtain the key.

Informally, the computational security of a cryptosystem is defined by the work factor (amount of computation) required to obtain the key.

5.3 The cryptosystem

5.3.1 Preliminaries

We already know from the previous section that the normal form of F must look like

$$x_{i_1} \dots x_{i_s} x_{j_t}^{-1} \dots x_{j_1}^{-1} \tag{16}$$

subject to the following conditions:

(NF1) $i_1 \leq \dots \leq i_s$ and $j_1 \leq \dots \leq j_t$
(NF2) If both x_i and x_i^{-1} occur, then either x_{i+1} or x_{i+1}^{-1} occurs as well.

We say that the word w is in *seminormal* form if it is of the form given above but only satisfies (NF1).

Let F be Thompson's group given by its infinite presentation and $s \in \mathbb{N}$ be a positive integer. We define the sets A_s and B_s as follows.

A_s contains elements whose normal forms have positive and negative parts which are the same length. That is, they look like:

$$x_{i_1} \dots x_{i_m} x_{j_m}^{-1} \dots x_{j_1}^{-1}$$

and

$$i_k - k < s \text{ and} j_k - k < s \text{ for every} k = 1, \dots, s \tag{17}$$

The set B_s consists of elements represented by words in the generators $x_{s+1}, x_{s+2}, \dots$. (B_s is a subgroup of F).

Proposition 5.2. *Let $a \in A_s$ and $b \in B_s$. Then $ab = ba$ in the group F.*

Proof. Let $a = x_{i_1} \dots x_{i_m} x_{j_m}^{-1} \dots x_{j_1}^{-1}$ and $b = x_{k_1}^{\varepsilon_1} \dots x_{k_l}^{\varepsilon_l}$ where $k_q > s$ for every $q \in 1, \dots, l$. By induction on l and m it is easy to show that in the group F one has

$$ab = ba = x_{i_1} \dots x_{i_m} \delta_m(b) x_{j_m}^{-1}$$

where δ_m is simply the operator that increases the indices of all generators by M. □

Proposition 5.3. *Let $s \geq 2$ be an integer. The set A_s is a subgroup of F generated by $x_0x_1^{-1}, \ldots, x_0x_s^{-1}$.*

Proof. The set A_s contains the identity and is clearly closed under taking inversions, that is $A_s = A_s^{-1}$. To show that A_s is closed under multiplication we take two arbitrary normal forms from A_s:

$$u = x_{i_1} \ldots x_{i_m} x_{j_m}^{-1} \ldots x_{j_1}^{-1}$$

and

$$v = x_{p_1} \ldots x_{p_l} x_{q_l}^{-1} \ldots x_{q_1}^{-1}$$

and show that the normal form of uv belongs to A_s. First, note that since the numbers of positive and negative letters in uv are equal, the lengths of the positive and negative letters in uv are equal, too. □

It now remains to show that the property from equation 17 of indices in the normal form of uv is satisfied.

Consider the subword in the middle of the product uv marked below:

$$uv = x_{i_1} \ldots x_{i_m} (x_{j_m}^{-1} \ldots x_{j_1}^{-1} x_{p_1} \ldots x_{p_l}) x_{q_l}^{-1} \ldots x_{q_1}^{-1}$$

and find a seminormal form for it using relations of F (by moving positive letters to the left and negative letters to the right by methods described in the earlier sections of this paper). We denote the word that we obtain by w. The word w is the product of a positive and a negative word: $w = pn$. By induction on $l+m$ one can show that both p and n satisfy the conditions from (17).

The we find normal forms for words p and n using relations of F (for p move letters with smaller indices to the left of letters with bigger indices, and for n move letters with smaller indices to the right of letters with bigger indices). By induction on the number of operations thus performed, one can show that the words that one obtains p' and n' satisfy the condition from (17). Therefore, the word $w' = p'n'$ is a seminormal form of uv satisfying the condition from (17).

Finally, we remove those pairs of generators in w' that contradict the property (NF2). Again, by induction on the number of *bad pairs*, one can show that

the result will satisfy the condition from (17). Therefore uv belongs to A_s, that is A_s is closed under multiplication, which implies that A_s is a subgroup.

Now we can show that the set of words $x_0x_1^{-1}, \ldots, x_0x_s^{-1}$ generates the subgroup A_s. Elements $x_0x_1^{-1}, \ldots, x_0x_s^{-1}$ clearly belong to A_s. To show the inclusion $A_s \leq \langle x_0x_1^{-1}, \ldots, x_0x_s^{-1} \rangle$, we construct the *Schrier graph* of $\langle x_0x_1^{-1}, \ldots, x_0x_s^{-1} \rangle$ (depicted in Figure 32) and we see that any word from A_s belongs to the subgroup on the right.

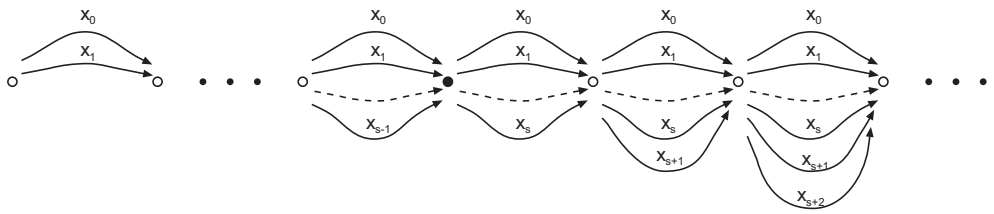

Figure 32: The Schrier graph of the subgroup $H = \langle x_0x_1^{-1}, \ldots, x_0x_s^{-1} \rangle$. The black dot denotes the right coset corresponding to H.

5.3.2 The protocol

The protocol is carried out like so:

1. Fix two positive integers s, M and a word $w = w(x_0, x_1, \ldots)$.
2. Alice randomly selects private elements $a_1 \in A_s$ and $b_1 \in B_s$. Then she reduces the element a_1wb_1 to the normal form and sends the result to Bob.
3. Bob randomly selects private elements $a_2 \in A_s$ and $b_2 \in B_s$. Then he reduces the element b_2wa_2 to the normal form and sends the result to Alice.
4. Alice computes $K_A = a_1b_2wa_2b_1$, and Bob computes $K_B = b_2a_1wb_1a_2$. Since $a_ib_i = b_ia_i$ in F, one has $K_A = K_B = K$ (an element of F), which is now Alice and Bob's common secret key.
5. Once keys are exchanged, Alice and Bob may use these to encrypt whatever message they choose using more secure symmetric key ciphers.

5.3.3 Some parameters and key generation

In a practical key exchange, it is suggested that one chooses the following parameters (given in [12]).

1. Randomly (and uniformly) select s from the interval [3,8] and M from the set 256, 258,. . . ,318,320.

2. Select the "base" word w as a product of generators

$$S_W = \{x_0, x_1, \ldots, x_{s+2}\}$$

and their inverses. This is done by first starting with an empty word which we will denote v_0. When we have a *current word* v_i, we multiply it on the right by a generator from $S_B^{\pm 1}$ and compute the normal form of the product. The word that is obtained is denoted by v_{i+1}. We continue this process until the obtained word v_{i+1} has length M.

3. Select a_1 and a_2 as products of words from

$$S_A = \{x_0 x_1^{-1}, \ldots, x_0 x_s^{-1}\}$$

and their inverses. This is done in the same way as above for w. We begin with the empty word u_0. Let u_i be the currently constructed word of length less than M. Multiply u_i on the right by a randomly chosen word from $S_A^{\pm 1}$ and compute the normal form of the product. Denote the normal form that is obtained by u_{i+1}. Continue this process until the obtained word u_{i+1} has length M.

4. Select b_1 and b_2 as products of generators from

$$S_B = \{x_{s+1}, x_{s+2}, \ldots, x_{2s}\}$$

and their inverses. To do that, as before, we begin with the empty word v_0. Multiply a current word v_i on the right by a generator from $S_B^{\pm 1}$ and compute the normal form of the product. Denote the word that is obtained by this process by v_{i+1}. Continue this process until the obtained word v_{i+1} has length M.

Note: the key space in the proposed scheme is exponential in M; moreover $|A_s(M)| \geq \sqrt{2}^M$, [12].

Define a directed labelled graph $\Gamma = (V(\Gamma), E(\Gamma))$ as follows:

- The set of vertices $V(\Gamma)$ corresponds to the set of all elements of the group F.
- The set of edges $E(\Gamma)$ contains edges $(w_1, w_2) : v_1 \to v_2$ such that $v_2 = w_1 v_1 w_2$ in the group F, with labels of two types:
  - $(w_1, 1)$, where $w_1 \in S_A^{\pm 1}$.
  - $(1, w_2)$, where $w_2 \in S_B^{\pm 1}$.

For any element $w \in F$ we denote the connected component of Γ containing w by Γ_w. From the description of the protocol it follows that w and the element $w' = a_1 w b_1$ transmitted by Alice to Bob belong to $\Gamma_w = \Gamma_{w'}$, and breaking Alice's key is equivalent to finding a label of a path from w to w' in Γ_w.

5.4 The word problem in Thompson's group

Recall that Thompson's group F has the following infinite presentation:

$$\langle x_k, k \geq 0 \mid x_i^{-1} x_j x_i = x_{j+1} \text{ if } i < j \rangle$$

The canonical normal form for an element is described in detail earlier in this paper.

Let us denote $\rho(w)$ the normal form for $w \in F$; it is unique for a given element of F. Recall that we say that a word w is in *seminormal form* if it is of the form given in equation (16) and satisfies (NF1). It is important to remember that a seminormal form is not unique. As usual, for a word w in the alphabet X, we denote the corresponding freely reduced word by $\bar{w}$.

The normal form for an element in Thompson's group can be computed in two steps:

1. Computation of a seminormal form
2. Removing *bad pairs*, that is pairs (x_i, x_i^{-1}) wherever neither x_{i+1} nor x_{i+1}^{-1} occur.

The first step is achieved by rules which the reader may recall from the first section. These are derived from the relators in the infinite presentation. They are:

$$\begin{array}{rcl} x_j x_i & \to & x_i x_{j+1} \\ x_j^{-1} x_i & \to & x_i x_{j+1}^{-1} \\ x_i^{-1} x_j & \to & x_{j+1} x_i^{-1} \\ x_i^{-1} x_j^{-1} & \to & x_{j+1}^{-1} x_i^{-1} \end{array}$$
$$(i < j)$$

Also, we should keep in mind $x_i^{-1} x_i \to 1$ for all $i \in \mathbb{N}$. We repeat these steps as many times (and in whatever order) as is necessary to get a word into seminormal form.

We denote this system by $\mathcal{R}$. A word which satisfies (NF1) by $\mathcal{R}$ (or otherwise) is called **$\mathcal{R}$-reduced**. We now have the following (very obvious) Lemma.

Lemma 5.4. *$\mathcal{R}$ terminates with a seminormal form. Moreover, a word is in seminormal form if and only if it is $\mathcal{R}$-reduced.*

Examining $\mathcal{R}$ more closely, we find that the action is similar to sorting a list of numbers but with two key differences: indices of generators may increase, and some generators may disappear altogether.

By Lemma 5.4, for any word w in the generators of F, the final result of rewrites by $\mathcal{R}$ is a seminormal form. Therefore, to compute a seminormal form we implement rewrites by $\mathcal{R}$. For convenience, we introduce a parametric function to affect a shift in index of the subscripts of elements in our seminormal form which we denote by δ_ε where $\varepsilon \in \mathbb{Z}$, defined on the set of all words in the alphabet $\{x_0^{\pm 1}, x_1^{\pm 1}, \ldots\}$ by

$$\delta_\varepsilon : x_i^{\pm 1} \mapsto x_{i+\varepsilon}^{\pm 1}$$

The function δ_ε may not be defined for some negative ε on a given word $w = w(x_{i_1}^{\pm 1}, x_{i_2}^{\pm 1}, \ldots)$, but when it is used, it is assumed that the function is defined.

We now introduce our first algorithm. This algorithm "merges" seminormal forms of two words which we'll call w_1 and w_2. Let $w_1 = p_1 n_1$ (p is for the positive part of the seminormal form and n is for the negative part, referring to the positive and negative exponents), and let $w_2 = p_2 n_2$. We *merge* these words by right multiplying w_1 by w_2. What we end up with will not satisfy (NF1) so we need to manipulate our product in order to get it into seminormal form. This process can basically be summarised below:

$$\begin{array}{c} p_1 \underbrace{n_1 p_2}\, n_2 \\ \downarrow \\ \underbrace{p_1 p_2'}\, \underbrace{n_1' n_2} \\ \downarrow \quad \downarrow \\ p \quad n \end{array}$$

Writing the steps out, he have:

1. Rewrite the subword $n_1 p_2$ of w to a seminormal form $p_2' n_1'$. Denote the result: $p_1 p_2' n_1' n_2$ by w'.

2. Rewrite the positive subword p_1p_2' of w' to a seminormal form p. Denote the result: $pn_1'n_2$ by w''.

3. Similarly, rewrite the negative subword $n_1'n_2$ of w'' to a seminormal form n. Denote the result: pn by w'''.

The word $w''' = pn$ is clearly in a seminormal form and $w = w'''$ in F.
We conclude the following important lemma about this algorithm.

Lemma 5.5. *Given two words in seminormal form w_1 and w_2, the algorithm which merges them to give another seminormal form w''' has time complexity required to compute the output which is bounded by $C(|n| + |p|)$ for some constant C.*

(For the proof of this lemma and for the merge algorithm, refer to [12]).

Now that we have algorithms for swapping parts of seminormal forms and merging them, we can essentially proceed with an algorithm for computing a seminormal form. This is because, one can take any given word and split it up into as small pieces as is necessary such that each of those pieces satisfies (NF1) and is thus a seminormal form. We can then recursively compute a seminormal form by swapping and merging until we have a word which satisfies (NF1). In the algorithm given in [12], the starting word w is split in 'half'. More precisely, w is represented by the product w_1w_2 such that $|w_1| - |w_2| \leq 1$, thus allowing for words which have an odd number of elements.

Lemma 5.6. *Let w be a word in the generators of F. For the algorithm which returns a seminormal form for w, the number of operations required for the process to terminate is $O(C|w|log|w|)$. Moreover, C is a constant independent of w.*

(Note: The base of the logarithms (as is common in computational applications) is 2).

Full details of the proof can be found in [12]. The detail that interests us the most is the asymptotic upper bound for this processes termination. Let us begin by denoting $T(n)$ to be the number of steps required for the seminormal form finding algorithm to terminate on an input of length n. Then clearly

$$T(n) = 2T(\frac{n}{2}) + C \cdot n \tag{18}$$

where the $C \cdot n$ is the complexity of merging two seminormal forms (as given in Lemma 5.5) with the sum of lengths at most $|n|$. It is a fairly trivial exercise to show that, in this case, $T(n) = O(C \cdot nlog(n))$.

Our final algorithm involves satisfying (NF2). Now that we have our word in seminormal form, we must check to see if it is in canonical *unique* normal form and, if it isn't, to make the necessary changes to it until it is.

Recall from the earlier sections that, using rules from the infinite presentation, we can move any x_i past x_j towards the centre without affecting i but affecting $j \mapsto j + 1$. The same is true for x_i^{-1} and x_j^{-1} as long as, in both cases, $i < j$.

The condition (NF2) requires that our algorithm first checks for "bad pairs", that is, the occurrence of x_i and x_i^{-1} where neither x_{i+1} nor x_{i+1}^{-1} occur. We start in the 'middle' of the word (where negative exponents meet positive ones) and work outwards until we find such a pair. Once found, all the elements of the subword which is found in between this pair are acted on by δ_1 (in other words, all the subscripts increase by 1) and the offending bad pair is removed.

Lemma 5.7. *The algorithm which takes a seminormal form (w) and returns a normal form (u) by eliminating "bad pairs" requires, at most, $D \cdot |u|$, where D is a constant independent of u.*

Putting all these lemmas together gives us the main result:

Theorem 5.8. *In Thompson's group F, the normal form of a given word w can be computed in time $O(|w|log|w|)$. (i.e. Almost linear in n).*

References

[1] G. N. Arzhantseva, J. Burillo, M. Lustig, L. Reeves, H. Short, E. Ventura, Uniform non-amenability, Advances in Mathematics, No. 197, 2005.

[2] L. Bartholdi, R. I. Grigorchuck, Z. Šuniḱ, *Branch Groups* in Handbook of Algrebra Vol 3., pp. 989–1112, 2003.

[3] K. S. Brown, R. Geoghegan, An infinite-dimensional torsion-free FP_∞ group, Inventiones mathematicae, No. 77 1984.

[4] J. Burillo, S. Cleary, B. Wiest, Computational Explorations in Thompson's Group F, Centre de Recerca Matemtica, 2005.

[5] S. Boztaş, Cryptomathematics Lecture Notes, AMSI ICE-EM Summer School 2006.

[6] J. W. Cannon, W. J. Floyd, W. R. Parry, Introductory Notes on Richard Thompson's Groups, L'Enseignment Mathematique, t. 42, 1996.

[7] S. Cleary, J. Taback, Combinatorial Properties of Thompson's group F, Transactions of the American Mathematical Society Vol 356 No. 7, 28th October 2003.

[8] S. Cleary, J. Taback, Seesaw words in Thompson's group F, arXiv:math.Gr/0310466 v2, 2004.

[9] T. M. Cover, J. A. Thomas, Elements of Information Theory, Wiley-Interscience 1991.

[10] V. S. Guba, On The Properties of the Cayley Graph of Richard Thompson's Group F, arXiv:math.GR/0211396 v1, 26th Nov 2002.

[11] S. B. Fordham, Minimal Lenth Elements of Thompson's Group F, Geometraie Dedicata 99, Kluwer Academic Publishers, 2003.

[12] V. Shpilrain, A. Ushakov, Thompson's Group and Public Key Cryptography, arXiv:math.GR/0505487 v1, 24th May 2005.

[13] S. Singh, The Code Book: The Science of Secrecy from Ancient Egypt to Quantum Cryptography, Fourth Estate, London 1999.