Kai-Uwe Bux

Groups and Spaces

Volume I: Groups

Preface

This set of notes is based on my lectures “Important Groups” and “My Favorite Groups” which I taught at Cornell University in Spring 2002 and Spring 2003. The goal was to discuss the most important examples of infinite discrete groups in considerable depth and detail.

Being “important” is rather a sociological concept than a mathematical one. A group is important if Mathmaticians are interested in it: Important groups are those, we talk about on dinner parties; and if you fail to know them, you will become a social outcast. Of course, any list of those groups is open to debate, and the one to follow reflects many of my personal idiosyncrasies.

What is in and what is not

The focus of this lecture is on particular groups. So you will not find a lot of general theorems like the following miracle.

> A finitely generated group has a solvable word problem if and only if it embeds into a simple subgroup of a finitely presented group.

Those gadgets are sometimes mentioned, sometimes even used, but we will not bother to prove many of these results.

Instead you will find results like these:

- The infinite cyclic group is amenable, whereas the other finitely generated free groups are not.
- Arithmetic groups are residually finite.
- The group $\mathrm{Out}(F_n)$ has only finitely many finite subgroups up to conjugacy.

The groups dicussed in this lecture can be aranged in two main blocks: On the one hand we discus free groups, Thompson’s groups, and Grigorchuk’s groups. Here the focus of the discussion is about growth and amenability. On the other hand, we will meet arithmetic groups, mapping class groups of surfaces and the groups $\mathrm{Out}(F_n)$. Here the discussion will be centered around the construction of nice spaces, upon which these groups act.

Thanks ...

... to David Revelle for proofreading my crappy notes.

... to Jim Belk for helping me considerably in understanding amenability and Thompson's group F.

... to David Benbennick for debugging the exercises and problems.

... to Ference Gerlitz for teaching me about the subgroup conjugacy problem in free groups.

... to Rotislav Grigorchuk for explaining to me the solution of the conjugacy problem in his first group.

Contents

Part I

Free Groups

Chapter 1

The Infinite Cyclic Group

The infinite cyclic group $C_\infty = \mathbb{Z}$ is the most simple infinite discrete group. It can be given in various ways:

- by a presentation $C_\infty = \langle x \rangle$.
- as the fundamental group of the circle $C_\infty = \pi_1(\mathbb{S}^1)$.
- as a group of matrices $C_\infty = \left\{ \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} \,\middle|\, s \in \mathbb{Z} \right\}$.

So we can read off that C_∞ is finitely presented. It acts freely and cocompactly on the real line, the universal cover of $\mathbb{S}^1$. The real line is contractible and has dimension 1. Thus, we infer:

- C_∞ is of type F.
- C_∞ has geometric dimension 1.
- C_∞ is linear.
- C_∞ is torsion free.

Moreover, we will see:

[1.1.2] C_∞ is residually finite.

[1.2.30] C_∞ is amenable.

[1.3.3] C_∞ does not have Kazhdan's property (T).

[1.4.13] C_∞ has two ends and, up to commensurability, it is characterized by this property.

[1.4.26] C_∞ has linear growth and, up to commensurability, it is characterized by this property.

1.1 Residual Finiteness

Definition 1.1.1. A group G is residually blah if for every non-trivial element g, there is a blah quotient of G wherein g does not become trivial.

Theorem 1.1.2. *C_∞ is residually finite.*

Proof. The infinite cyclic group has factors of any order. An element will not be trivial in any factor whose order is relatively prime to the order of the given element. **q.e.d.**

1.2 Amenability

Definition 1.2.1. Let G be a group acting on a set X. The set X is called G-amenable if there is a G-invariant finitely additive probability measure on the system of all subsets of X.

A group G is called amenable if the left-action of G on itself by multiplication turns G into a G-amenable set.

Example 1.2.2. The counting measure shows that finite groups are amenable.

Remark 1.2.3. Given a finitely additive G-invariant probability measure μ on the set X, we can define a left invariant mean on the set of bounded real-valued functions $\mathrm{L}(X, \mathbb{R})$. A mean on $\mathrm{L}(X, \mathbb{R})$ is a linear map $\mathrm{L}(X, \mathbb{R}) \to \mathbb{R}$ taking constant functions to their unique value. Such a mean is left-invariant, if it is invariant with respect to the canonical left-action of G on $\mathrm{L}(X, \mathbb{R})$ given by

$$(gf)(x) := f\big(g^{-1}x\big)$$

for $g \in G$ and $f \in \mathrm{L}(X, \mathbb{R})$.

The left invariant mean associated to a finite probability measure is constructed as follows:

1. For characteristic functions χ_Y, define

$$\int_{x \in X} \chi_Y \,\mathrm{d}_\mu\, x := \mu(Y)\,.$$

2. Extend this definition to finite linear combinations of characteristic functions.

3. For an arbitrary bounded function f, put

$$\int_{x\in X} f \,\mathrm{d}_\mu x := \lim_{\varepsilon\to+0} \int_{x\in X} f_\varepsilon \,\mathrm{d}_\mu x$$

where f_ε is any linear combination of characteristic functions that is everywhere ε-close to f. Using $\mu(X) = 1$, it is easy to see that the limit on the right hand side does not depend on the chosen ε-approximations f_ε.

Invariance of the mean follows from the invariance of μ. Beware that this mean, although we denote it by the $\int$-sign behaves not always as an analyst might expect.

Conversely, every left-invariant mean on $\mathrm{L}(X, \mathbb{R})$ induces a left-invariant finitely additive probability measure on G: just assign to each subset the mean of its characteristic function. We have thus proved one of the many characterizations of amenability:

> A G-set X is amenable if and only if there is a G-invariant mean on $\mathrm{L}(X, \mathbb{R})$.

A group acts on itself by multiplication from the left and from the right. We can use invariant means to improve upon the invariance of a measure:

Proposition 1.2.4. *An amenable group G has a bi-invariant finitely additive probability measure $\tilde{\mu}$.*

Proof. Let μ be a left-invariant finitely additive probability measure on G. Observe that

$$\mu^{-}(S) := \mu\big(S^{-1}\big)$$

defines a right-invariant measure on G. It is easy to check that

$$\tilde{\mu}(S) := \int_{g\in G} \mu\big(Sg^{-1}\big) \,\mathrm{d}_{\mu^-} g$$

is bi-invariant: We have

$$\begin{aligned}
\tilde{\mu}(hS) &= \int_{g\in G} \mu\big(hSg^{-1}\big) \,\mathrm{d}_{\mu^-} g \\
&= \int_{g\in G} \mu\big(Sg^{-1}\big) \,\mathrm{d}_{\mu^-} g \\
&= \tilde{\mu}(S)
\end{aligned}$$

and

$$\begin{aligned}\tilde{\mu}(Sh) &= \int_{g\in G} \mu\big(Shg^{-1}\big)\,\mathrm{d}_{\mu^-}\, g \\ &= \int_{g\in G} \mu\big(Shg^{-1}\big)\,\mathrm{d}_{\mu^-}\, gh^{-1} \\ &= \int_{g\in G} \mu\Big(S\left(gh^{-1}\right)^{-1}\Big)\,\mathrm{d}_{\mu^-}\, gh^{-1} \\ &= \tilde{\mu}(S)\,.\end{aligned}$$

q.e.d.

The main goal of this section is to show

Theorem 1.2.5. *The infinite cyclic group is amenable.*

which will be proved as Corollary (1.2.30).

1.2.1 Følner Sequences

Definition 1.2.6. A Følner sequence for a G-set X is a sequence (F_i) of finite subsets $F_i \subseteq X$ such that

$$\lim_{i\to\infty} \frac{|gF_i \triangle F_i|}{|F_i|} = 0 \qquad \text{for all } g \in G.$$

Here $M \triangle N$ denotes the symmetric difference of two sets:

$$M \triangle N := (M \cup N) - (M \cap N)$$

Remark 1.2.7. It is very easy to see that, for any sequence (F_i) of finite subset sets in the G-set X and any group element $g \in G$,

$$\lim_{i\to\infty} \frac{|gF_i \triangle F_i|}{|F_i|} = 0 \qquad \text{if and only if} \qquad \lim_{i\to\infty} \frac{|gF_i \cap F_i|}{|F_i|} = 1.$$

So either condition can be used to define Følner sequences.

Exercise 1.2.8. For any sequence (F_i) of finite subsets in X, the set

$$\left\{ g \in G \;\middle|\; \lim_{i\to\infty} \frac{|gF_i \triangle F_i|}{|F_i|} = 0 \right\}$$

is a subgroup of G.

Lemma 1.2.9. *The "balls of radius i" form a Følner sequence for the infinite cyclic group.*

Proof. The ball of radius i is the subset $\{-i, -i+1, \ldots, i-1, i\}$. It is obvious that for a fixed group element g, very large balls will have large overlaps with their g-translate. **q.e.d.**

We will show below that a group is amenable if it has a Følner sequence. This will complete the proof that the infinite cyclic group is amenable.

1.2.2 Ultralimits

Definition 1.2.10. Let M be a set. A filter on M is a set $\mathcal{F}$ of subsets of M satisfying:

1. $\emptyset \notin \mathcal{F}$.
2. If $F \in \mathcal{F}$ and $M \supseteq H \supseteq F$, then $H \in \mathcal{F}$.
3. If $F \in \mathcal{F}$ and $H \in \mathcal{F}$, then $(F \cap H) \in \mathcal{F}$.

Observation 1.2.11. *Since a filter does not contain the empty set but is closed with respect to forming intersections, no two sets in a filter are disjoint.* **q.e.d.**

Example 1.2.12. For any non-empty subset $S \subseteq M$, the system

$$\{F \subseteq M \mid F \supseteq S\}$$

is a filter. It is called the principal filter induced by S.

Example 1.2.13. If M is infinite, the set of cofinite sets (complements of finite sets) is a filter $\mathcal{CF}_M$.

Example 1.2.14. A directed set is a partially ordered set D such that any two elements have a common upper bound. A coinitial segment is a subset $F \subseteq D$ satisfying:

if $\alpha \in F$ and $\alpha < \beta$, then $\beta \in F$.

The set of supersets of non-empty coinitial segments is a filter $\mathcal{CI}_T$.

Exercise 1.2.15. Show that $\mathcal{CF}_{\mathbb{N}} = \mathcal{CI}_{\mathbb{N}}$.

All filters on M are comparable with respect to inclusion: a filter $\mathcal{F}$ is called finer that a filter $\mathcal{H}$ if every subset of M that belongs to $\mathcal{H}$ also belongs to $\mathcal{F}$. Obviously, this just means $\mathcal{H} \subseteq \mathcal{F}$. A filter that cannot be refined (a finest filter) is called an ultrafilter.

Example 1.2.16. Fix a subset $S \subset M$ and a filter $\mathcal{F}$ that contains neither S nor $M - S$. Then

$$\mathcal{F}_S := \{T \mid M \supseteq T \supseteq (S \cap F) \text{ for some } F \in \mathcal{F}\}$$

is a filter finer than $\mathcal{F}$.

From this example, we immediately infer:

Corollary 1.2.17. *Let $\mathcal{U}$ be an ultrafilter on M. For every subset $S \subseteq M$ either $S \in \mathcal{U}$ or $(M - S) \in \mathcal{U}$.* **q.e.d.**

Since ascending unions of filters are filters, Zorn's Lemma immediately implies:

Lemma 1.2.18. *Every filter is contained in an ultrafilter.* **q.e.d.**

Definition 1.2.19. Let M be a set, $\mathcal{F}$ a filter on M, and X a topological space. A family $(x_m)_{m \in M}$ of points in X $\mathcal{F}$-converges to a point $x \in X$ if, for every open neighbourhood U of x,

$$\{m \in M \mid x_m \in U\} \in \mathcal{F}.$$

In this case, we say $x = \mathcal{F}\text{-}\lim_{m \in M} x_m$ is an $\mathcal{F}$-limit of $(x_m)_{m \in M}$.

Example 1.2.20. Ordinary convergence of sequences is the same as $\mathcal{CF}_{\mathbb{N}}$-convergence.

Example 1.2.21. A net in X over D is a family of points in X indexed by a directed set D. Convergence for nets is defined as $\mathcal{CI}_D$-convergence.

Observation 1.2.22. *If $\mathcal{H}$ is finer than $\mathcal{F}$ then any $\mathcal{F}$-limit of a net is also an $\mathcal{H}$-limit.*

Proposition 1.2.23. *If X is Hausdorff, then $\mathcal{F}$-limits are unique.*

Proof. This is done by contradiction. Suppose there were two points a_1 and a_2 such that, for each open neighborhood U_j of a_j,

$$\{m \in M \mid x_m \in U_j\} \in \mathcal{F}.$$

Since X is Hausdorff, we can choose these two neighborhoods to be disjoint. But then, both sets

$$\{m \in M \mid x_m \in U_j\}$$

are disjoint. As $\mathcal{F}$ is a filter, this cannot happen. **q.e.d.**

Theorem 1.2.24. *Fix an ultrafilter $\mathcal{U}$ on the set M. Let C be a compact topological space. Then any family of points $(x_m)_{m\in M}$ has a $\mathcal{U}$-limit.*

Proof. Assume by contradiction that each $a \in C$ has an open neighbourhood U such that

$$\{m \in M \mid x_m \in U\} \notin \mathcal{U}.$$

Then we can, by compactness, cover C with finitely many open sets U_j such that for each j

$$\{m \in M \mid x_m \in U_j\} \notin \mathcal{U}.$$

However,

$$M = \bigcup_j \{m \in M \mid x_m \in U_j\}$$

is a cover of M by finitely many subsets. Hence one of them must be in $\mathcal{U}$ by the following Lemma 1.2.25 **q.e.d.**

Lemma 1.2.25. *Let $\mathcal{U}$ be an ultrafilter. If a finite union $S_1 \cup \cdots S_n$ belongs to $\mathcal{U}$, then, for at least one index i, we have $S_i \in \mathcal{U}$.*

Proof. Suppose $S_i \notin \mathcal{U}$ for all i. Then

$$M - S_i \in \mathcal{U} \qquad \text{for all } i$$

Hence

$$\bigcap_i M - S_i = M - \bigcup_i S_i \in \mathcal{U}$$

which contradicts the assumption $\bigcup_i S_i \in \mathcal{U}$. **q.e.d.**

Anything called a "limit" should commute with continuous functions and cross-products. The following statements just say that $\mathcal{F}$-limits behave as you would expect.

Lemma 1.2.26. *Let $f : X \to Y$ be a continuous map between compact spaces. Then for any $\mathcal{F}$-convergent family $(x_m)_{m\in M}$ in X,*

$$\mathcal{F}\text{-}\lim f_{x_m} = f_{\mathcal{F}\text{-}\lim x_m}.$$

Proof. Consider an open neighbourhood U of $f_{\mathcal{F}\text{-}\lim x_m}$. Its preimage under f is an open neighbourhood V of $\mathcal{U}\text{-}\lim x_m$. Hence

$$\{m \in M \mid f_{x_m} \in U\} = \{m \in M \mid x_m \in V\} \in \mathcal{U}.$$

Since U was arbitrary, the statement follows. **q.e.d.**

Remark 1.2.27. Let C and D be compact. For any family of pairs $((x_\alpha, y_\alpha))$ in $C \times D$, we have:

$$\mathcal{U}\text{-}\lim(x_\alpha, y_\alpha) = (\mathcal{U}\text{-}\lim x_\alpha, \mathcal{U}\text{-}\lim y_\alpha)$$

Since arithmetic operations, like $+ : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$, are continous, and since bounded sequences only use compact subsets of $\mathbb{R}$, we immediatetly conclude:

Corollary 1.2.28. *Bounded sequences of real numbers have unique ultralimits and these limits are compatible with the arithmetic operations of addition, subtraction and multiplication.*

1.2.3 From Følner Sequences to Amenability

Proposition 1.2.29. *A G-set X is G-amenable if it admits a Følner sequence.*

Proof. Suppose we have a Følner sequence (F_i). Fix an ultrafilter $\mathcal{U}$ on $\mathbb{N}$ refining the coinitial filter so that we can form ultralimits of bounded sequences of real numbers. Let Y be any subset of X. Then

$$\frac{|Y \cap F_i|}{|F_i|}$$

is a sequence in $[0, 1]$. We define the finitely additive measure μ by:

$$\mu(Y) := \mathcal{U}\text{-}\lim \frac{|Y \cap F_i|}{|F_i|}$$

Obviously $\mu(G) = 1$.

To see that this measure is additive, consider two disjoint subsets Y and Z and observe that

$$|(Y \cup Z) \cap F_i| = |Y \cap F_i| + |Z \cap F_i|$$

From this additivity follows because ultralimits commute with addition (1.2.28).

To see that the measure is left-invariant, we write

$$\left| \frac{|gY \cap F_i|}{|F_i|} - \frac{|Y \cap F_i|}{|F_i|} \right| = \left| \frac{|Y \cap g^{-1}F_i|}{|F_i|} - \frac{|Y \cap F_i|}{|F_i|} \right| \leq \frac{|Y \cap (g^{-1}F_i \triangle F_i)|}{|F_i|} \xrightarrow[i \to \infty]{} 0$$

Now the claim follows since our ultrafilter refines the cofinite filter on $\mathbb{N}$ whence ordinary limits are ultralimits for $\mathcal{U}$. **q.e.d.**

Corollary 1.2.30. *The infinite cyclic group is amenable.* **q.e.d.**

Exercise 1.2.31. Show that every abelian group is amenable:

1. Show that the direct product of two amenable groups is amenable.
2. Show that a group is amenable if all its finitely generated subgroups are amenable (i.e., locally amenable groups are amenable). Hint: The system of finitely generated subgroups inside G is a directed set. Use an ultralimit construction to obtain a measure on G from the measures on the finitely generated subgroups of G.

From (1) infer that finitely generated abelian groups are amenable. Then (2) implies that abelian groups are amenable.

Exercise 1.2.32. Add a little twist to what you did on direct products and show that a group is amenable if it has an amenable normal subgroup such that the quotient is also amenable. I.e., amenable-by-amenable groups are amenable. Infer that solvable groups are amenable.

Definition 1.2.33. A group G is foo–by–bar if there is a short exact sequence

$$1 \to H \to G \to F \to 1$$

where H is foo and F is bar. The particle "by" implies left parentheses. So "foo–by–bar–by–blah" means (foo-by–bar)–by–blah.

Exercise 1.2.34. Show that subgroups of amenable groups are amenable.

Lemma 1.2.35. *If a group G has an amenable subgroup H of finite index, it is amenable.*

Proof. Let K be the kernel of the action of G on the finite set of cosets G/H. Obviously, $K \leq H$. Hence K is amenable. Since we have the short exact sequence

$$K \hookrightarrow G \twoheadrightarrow G/K$$

wherein K is amenable and G/K is finite, G is amenable-by-amenable and therefore amenable. **q.e.d.**

Corollary 1.2.36. *Virtually solvable groups (i.e., groups that have a solvable subgroup of finite index) are amenable.* **q.e.d.**

1.3 Kazhdan's Property (T)

Definition 1.3.1. A unitary representation of a topological group G on a Hilbert space $\mathcal{H}$ is a continous homomorphism $\rho : G \to \mathrm{U}_{\mathcal{H}}$ from G to the group of bounded unitary isomorphisms of $\mathcal{H}$. If no confusion can arise, we suppress ρ in the notation and think of it as a left-action of G on $\mathcal{H}$.

Such a unitray representation is said to have almost invariant vectors, if, for any compact subset $K \subseteq G$ and any $\varepsilon > 0$, there is a unit vector $\mathbf{u} \in \mathcal{H}$ satisfying

$$|g\mathbf{u} - \mathbf{u}| < \varepsilon \qquad \text{for all } g \in K.$$

The group G has Property (T) if every unitary representation that has almost invariant vectors has an invariant vector.

Theorem 1.3.2. *If an infinite discrete group G has a Følner sequence, then G does not have Kazhdan's property (T).*

Proof. Let F_i form a Følner sequence. Consider the action of G on the Hilbert space $L^2(G)$ of square summable function on G. The action is given by a shift. This action has no invariant vectors. But the sequence of vectors

$$\mathbf{u}_i : g \mapsto \begin{cases} \frac{1}{\sqrt{|F_i|}} & g \in F_i \\ 0 & g \notin F_i \end{cases}$$

satisfies

$$\lim_{i \to \infty} |g\mathbf{u}_i - \mathbf{u}_i| = 0 \qquad \text{for all } g \in G$$

From this, the claim follows since compact subsets of discrete groups are finite. **q.e.d.**

Corollary 1.3.3. *The infinite cyclic group does not have Kazhdan's property (T).* **q.e.d.**

Definition 1.3.4. A group is indicable if it admits an epimorphism onto the infinite cyclic group.

Observation 1.3.5. *Every image Q of a Kazhdan (T) group G under a continous homomorphism is (T) because any representation of the quotient Q lifts to a representation of G.*

In particular, every quotient of a discrete Kazhdan (T) group has property (T). **q.e.d.**

Corollary 1.3.6. *Indicable groups do not have Kazhdan's property (T).*

Corollary 1.3.7. *The free groups and the pure braid groups do not have Kazhdan's property (T).*

Lemma 1.3.8. *A virtually finitely generated group G is finitely generated.*

Proof. Let $H = \langle h_1, \ldots, h_r \rangle$ be a finitely generated subgroup of G and let $g_1, \ldots, g_s$ be a complete set of representatives for the finitely many cosets in G/H. Then

$$G = \langle H, g_1, \ldots, g_s \rangle = \langle h_1, \ldots, h_r, g_1, \ldots, g_s \rangle .$$

q.e.d.

Proposition 1.3.9. *A (discrete) group G that has Kazhdan's property (T) is finitely generated.*

Proof. We consider the unitary representation

$$\bigoplus_{H \leq G} L^2 (G/H)$$

where H runs through all finitely generated subgroups of G. Since any compact (finite) subset of G is contained in one of these subgroups, this representation has almost invariant vectors – just consider the action of the finite subset on an appropriate summand where it fixes the coset of the identity. Since G is supposed to be Kazhdan, we conclude that the representation has an invariant vector.

Hence one of the summands has an invariant vector. Such a vector corresponds to a constant function on G/H. Hence this quotient is finite. Therefore, G is virtually finitely generated and hence finitely generated. **q.e.d.**

Corollary 1.3.10. *Locally indicable groups do not have Kazhdan's property (T).*

Proof. Being discrete and Kazhdan, the group is finitely generated. Being finitely generated and locally indicable, it is indicable. **q.e.d.**

1.4 The Geometry of the Cayley Graph

Definition 1.4.1. Let G be a group with a finite generating system Σ. The (left) Cayley graph Γ_{Σ}^{G} is a (directed and labeled) graph. Its set of vertices is G, and for each vertex $g \in G$ and each generator $x \in \Sigma$, there is an edge (labeled by x) from g to gx. Note that G acts from the left on Γ_{Σ}^{G}.

There is a corresponding notion of a right Cayley graph upon which G acts from the right.

Remark 1.4.2. We usually do not care about the direction of edges or the labeling. Thus we regard the Cayley graph as a metric space: every edge has length 1 and the distance between any two points is the length of the shortest path connecting them. This length is finite since the Cayley graph is connected – this follows from the assumption that Σ generates G: any element of G can be written as a word in the generators (and their inverses) and this determines a path connecting the group element to the identity element.

Example 1.4.3. Here is the Cayley graph $\Gamma^{C_\infty}_{\{1\}}$:

And here this is what $\Gamma^{C_\infty}_{\{2,3\}}$ looks like:

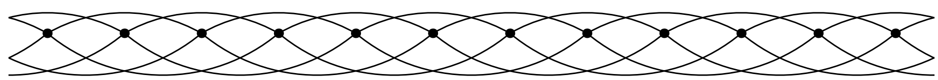

Let us provide one example of how to make use of the Cayley graph. We already know by (1.3.8) that a group is finitely generated if it has a finitely generated subgroup of finite index. As we shall see, the converse holds true, as well. This follows from applying the following lemma to actions on Cayley graphs.

Lemma 1.4.4. *Let the group G act on the connected topological space X, and suppose that there is an open subset U such that $X = \bigcup_{g\in G} gU$. Then G is generated by $S := \{g \in G \mid gU \cap U \neq \emptyset\}$.*

Proof. Let $H := \langle S\rangle$. Then the two sets HU and $(G-H)\,U$ are both open. In addition, they are disjoint: Suppose we had $hU \cap fU \neq \emptyset$ for some $h \in H$ and $f \in G - H$. Then $f^{-1}h \in H$ and therefore $f \in H$, which contradicts our assumption.

As X is connected, the set $(G-H)\,U$ is empty as the other one cannot be empty. Hence $G - H = \emptyset$. **q.e.d.**

There are variations of this lemma, e.g., on group presentations. We will encounter them later.

Corollary 1.4.5. *A finite index subgroup of a finitely generated group is finitely generated.*

Proof. Let G be a group with a finite generating set Σ and let H be a subgroup of finite index. Let $r_1, \ldots, r_r$ be a list containing a representative for each H-coset in G. Hence every vertex of the Cayley graph Γ^G_Σ lies in the orbit of one of the r_i. Let U be the union of open stars of the r_i. The hypotheses of (1.4.4) are obviously satisfied. Hence

$$\{h \in H \mid hU \cap U \neq \emptyset\} = \{h \in H \mid hr_i = gr_j \text{ for some } g \in \Sigma \cup \{1\}\}$$

generates H. But that set is finite. **q.e.d.**

Exercise 1.4.6. Let $\Gamma = \Gamma_\Sigma^G$ be a Cayley graph for the infinite, finitely generated group G with respect to the finite generating system Σ. Show that Γ contains a bi-infinite short-lex geodesic (defined below).

Any edge path in Γ reads a word in $\Sigma \uplus \Sigma^{-1}$: while you are moving along the path, you pick up the labels of the edges you are going along, when you move with the direction of the edge you read the label, when you are going against the directed edge in Γ you read its inverse.

Fix an order on the set Σ. This induces an ordering on the set of word with letters from Σ: shorter words precede longer words and you use the lexicographic order to break ties. Regarding inverses as lower case variants of the capital letters in Σ, we actually have an order on words in $\Sigma \uplus \Sigma^{-1}$. Every group element is represented by a unique short-lex minimal word. Hence any two vertices in Γ are joined by a unique short-lex minimal edge path. We call those paths short-lex geodesic segments. Note that they are, in fact, geodesic segments.

Now a (bi-infinite) short-lex geodesic is a (bi-infinite) edge path such that every finite sub path is a short-lex geodesic.

Hint: First prove that Γ contains a bi-infinite geodesic.

1.4.1 Ends

Definition 1.4.7. A diagram (of sets and maps) is a directed graph D whose vertices v are labeled by sets M_v and whose edges $\vec{e}$ are labeled by maps $f_{\vec{e}} : M_{\iota(\vec{e})} \to M_{\tau(\vec{e})}$. The inverse limit of D is the set

$$\varprojlim D := \left\{ (m_v \in M_v)_{v \in \mathcal{V}_D} \;\middle|\; f_{\vec{e}}^{m_{\iota(\vec{e})}} = m_{\tau(\vec{e})} \text{ for all } \vec{e} \in \mathcal{E}(D) \right\}$$

Note that there are natural maps $\varprojlim D \to M_v$ for all vertices $v \in D$ and all triangles

$$\begin{array}{ccc} \varprojlim D & \to & M_{\tau(\vec{e})} \\ \downarrow & \nearrow f_{\vec{e}} & \\ M_{\iota(\vec{e})} & & \end{array}$$

commute.

Definition 1.4.8. Let X be a topological space. For any two nested compact subsets, $C \subseteq D \subseteq X$, we have a natural map

$$\pi_0(X - D) \to \pi_0(X - C)\,.$$

As compact subsets in X form a directed set, we can write the inverse limit

$$\partial_\infty X := \varprojlim_{C \subseteq X} \pi_0(X - C)$$

The elements of the set $\partial_\infty X$ are called the ends of X.

Example 1.4.9. The two Cayley graphs of C_∞ both have precisely two ends.

Observation 1.4.10. *This construction is functorial, so homeomorphisms of X induce bijections of $\partial_\infty X$ and we have a group homomorphism*

$$\mathrm{Homeo}(X) \to \mathrm{Perm}(\partial_\infty X)\,.$$

In particular, if $X = \Gamma_\Sigma^G$ is a (left) Cayley graph for a group G, there is a natural action of G on $\partial_\infty \Gamma_\Sigma$ turning the set of ends into a G-set.

Exercise 1.4.11. The number of ends in a Cayley graph is 0, 1, 2, or ∞: Let $\Gamma := \Gamma_\Sigma^G$ be the Cayley graph for the group G with with respect to the finite generating set Σ. Show that if Γ has finitely many ends, then the number of ends is ≤ 2. Hint: Assume Γ has three ends. Then there should be a central region where these ends get tied up. But a Cayley graph looks homogeneous as there is a vertex transitive group action, hence there cannot be a distinguished region.

Exercise 1.4.12. Given the same setup as in (1.4.11), show that the number of ends (0, 1, 2, or ∞) is independent of the choice of the finite generating system Σ.

The following theorem relates the geometry of a Cayley graph to a purely algebraic property of a group. In this respect it is like Gromov's theorem 1.4.25. But it is way simpler, and it is about C_∞.

Theorem 1.4.13. *A group has two ends if and only if it is virtually C_∞.*

Proof. That a group which is virtually C_∞ has two ends is easy. We only prove the converse. So let Σ be a generating set for G such that the corresponding Cayley graph Γ has two ends.

Our group G acts on $\partial_\infty \Gamma$ and the kernel of this action has index ≤ 2 in G. Therefore, by passing to a subgroup of index 2 if necessary, we may assume without loss of generality that G fixes both ends of Γ.

Since Γ has two ends, there is a compact subset C such that $\Gamma - C$ has exactly two infinite components W_- and W_+. We add all finite components of $\Gamma - C$ and can henceforth assume that the two infinite components are all there is in $\Gamma - C$. Since our space is infinite and C is contained in a ball of finite radius, there is an element $g \in G$ that moves C into one of the two infinite components. Let us assume $gC \subseteq W_+$.

First, we show that g has infinite order. Note that $gW_+ \subsetneq W_+$ for otherwise g would swap the ends of Γ.

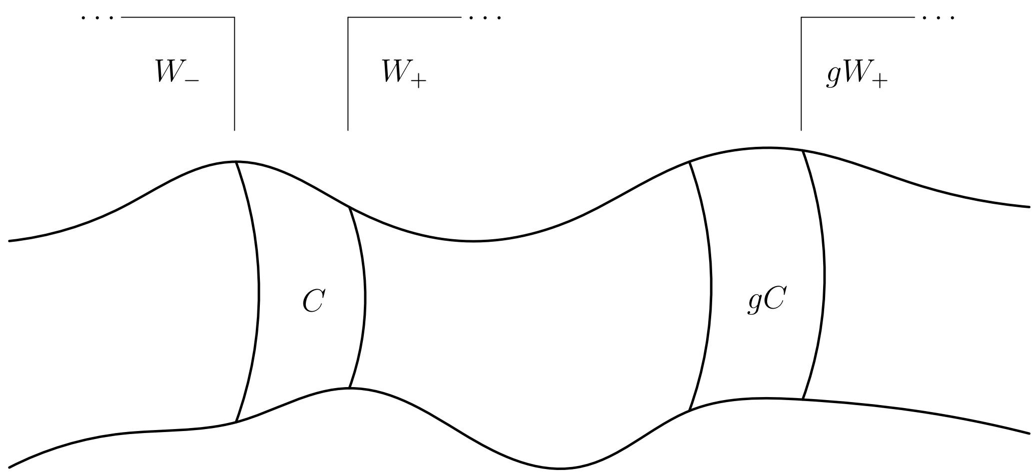

Hence we find

$$W_+ \supsetneq gW_+ \supsetneq g^2W_+ \supsetneq \cdots$$

and it follows that g has infinite order.

Let D be a compact subset containing C and its translate gC such that $\Gamma - D$ has exactly two components both of which are infinite. It follows that

$$\Gamma = W_- \cup D \cup gW_+.$$

We infer that $\Gamma = \bigcup_{i \geq 0} g^i W_-$ whence $\bigcap_{i \geq 0} g^i W_+ = \emptyset$. Similarly, $\bigcap_{i \leq 0} g^i W_- = \emptyset$. Moreover, for any $i > 0$,

$$\Gamma = g^{-i}W_- \cup \bigcup_{-i \leq j < i} g^j D \cup g^i W_+$$

whence

$$\Gamma = \bigcup_{s \in \mathbb{Z}} g^s D.$$

This, however, implies that D contains a representative for each coset in $\langle g \rangle \backslash G$. Hence $\langle g \rangle$ has finite index in G. **q.e.d.**

1.4.2 Growth

Definition 1.4.14. The growth function β_Σ of G relative to Σ is defined by

$$\beta_\Sigma(n) := \mathrm{vol}\bigl(\mathbb{B}_n^{1_G}\bigr)$$

where $\mathbb{B}_n^{1_G}$ is the ball of radius n in the Cayley graph centered at $1_G \in G$ and volume is measured by counting vertices.

Example 1.4.15. For the two Cayley graphs of C_∞, we find:

$$\beta_{\{1\}}(n) = 2n + 1$$

and

$$\beta_{\{2,3\}}(n) = \begin{cases} 1 & n = 0 \\ 5 & n = 1 \\ 6n + 1 & n \geq 2. \end{cases}$$

Note that any generating set for C_∞ will yield an ultimately linear growth function.

Definition 1.4.16. Let β and β' be two functions defined on $\mathbb{N}$. We say β' weakly dominates β if there are constants L and K such that

$$\beta(n) \leq L\beta'(Ln + K) + K.$$

We write $\beta \preceq \beta'$. We say that β and β' are weakly equivalent if they weakly dominate one another:

$$\beta \sim \beta' :\Longleftrightarrow \beta \preceq \beta' \ \ \& \ \ \beta' \preceq \beta.$$

Remark 1.4.17. Weak domination is transitive, and weak equivalence is an equivalence relation.

Observation 1.4.18. *If Σ and Ξ are two finite generating sets for G, then the growth functions β_Σ and β_Ξ are weakly equivalent. To see this, write the elements of Ξ as words in Σ. Let l be the maximum length that occurs in this list of words. Then any n-ball in Γ_Ξ is contained in the ln-ball in Γ_Σ which proves $\beta_\Xi \preceq \beta_\Sigma$.* **q.e.d.**

Exercise 1.4.19. Let G be finitely generated and H be a subgroup of finite index. By (1.4.5), H is finitely generated, too. Show that the growth functions for these two groups are weakly equivalent.

Observation 1.4.20. *For an infinite group with generating set Σ, we have*

$$n \leq \beta_\Sigma(n) \leq \sum_{i=0}^{n} (2\,|\Sigma|)^i.$$

Since there are at most as many elements in the n-ball as there are words of length $\leq n$ in the generators (and their inverses!). The lower bound follows from the fact that the Cayley graph is connected and infinite. **q.e.d.**

So growth in groups is somewhere between linear and exponential. Accordingly, one distinguishes three cases:

Definition 1.4.21. A finitely generated group is of polynomial growth if its growth function is weakly dominated by a polynomial. It is of exponential growth if it weakly dominates an exponential function. Otherwise it is of intermediate growth.

Exercise 1.4.22. Show that a finitely generated group with infinitely many ends has exponential growth.

Proposition 1.4.23. *Groups of subexponential growth are amenable.*

Proof. We will show that a group G of subexponential growth has a Følner sequence consisting of balls $\mathbb{B}_n^{1_G}$. For if this is not the case, then there is an $\varepsilon > 0$ such that for any ball $\mathbb{B}_n$ we have

$$\frac{|x\mathbb{B}_n \triangle \mathbb{B}_n|}{|\mathbb{B}_n|} > \varepsilon$$

for someelement $x \in \Sigma$. Since $\mathbb{B}_{n+1}$ contains $\mathbb{B}_n$ as well as $x\mathbb{B}_n$, we find:

$$\beta_\Sigma(n+1) = |\mathbb{B}_{n+1}| \geq |x\mathbb{B}_n \cup \mathbb{B}_n| \geq \left(1 + \frac{\varepsilon}{2}\right)|\mathbb{B}_n| = \left(1 + \frac{\varepsilon}{2}\right)\beta_\Sigma(n).$$

From this, it is obvious that G has exponential growth. **q.e.d.**

Exercise 1.4.24. Show: If a finitely generated group G has exponential growth, then the sequence of balls does not form a Følner sequence.

By (1.4.23), groups of polynomial growth are amenable. However, a very deep theorem says that we already knew that:

Theorem 1.4.25 (Gromov [Grom81]). *A finitely generated group has polynomial growth if and only if it is virtually nilpotent.*

This is way too deep for this exposition, and fortunately it is not a statement about the infinite cyclic group. However, there is a characterization of C_∞ by means of its growth:

Theorem 1.4.26. *A group has linear growth if and only if it is virtually C_∞.*

Before we prove this, we need a geometric lemma on growth rates.

Lemma 1.4.27. *Let H be an ∞-index subgroup of G. Let Σ be a finite generating system for H and Ξ a finite generating set for G that contains Σ. Then:*

$$\beta_\Xi(2n) \geq (n+1)\,\beta_\Sigma(n) \qquad \textit{for all } n \in \mathbb{N}$$

In particular, if H has polynomial growth of degree d, then G has at least polynomial growth of degree $d+1$.

Proof. Let $X := H\backslash\Gamma^G_\Xi$ be the space of orbits of vertices in Γ^G_Ξ under the action of H. We turn this into a metric space by defining the distance between two orbits to be the minimum distance between two representatives. Note that we can choose one of them freely. Choosing a pair of representatives realizing the distance shows that two points of distance n in X are joined by a path of length n. It follows that the ball $\mathbb{B}^H_n \subset X$ contains at least $n+1$ points $x_0, \ldots, x_n$. Each of these vertices has a representative $g_i \in \mathbb{B}^{1_G}_n \subseteq \Gamma_\Xi$. Now we consider translates of the balls in Γ^H_Σ. We find

$$\mathbb{B}^{1_G}_{2n} \supseteq \bigcup_{i=0}^{n} \mathbb{B}^{1_H}_n g_i$$

and the inequality follows from the fact that the union is disjoint. **q.e.d.**

Proof of Theorem 1.4.26. Since finite groups have "constant growth" which is not linear, we only have to consider infinite groups. One direction is obvious: If a group has an infinite cyclic subgroup of finite index, it has linear growth. So we have to show the converse: Any infinite group G of linear growth contains an infinite cyclic subgroup of finite index. By the preceding lemma, it suffices to show that there is an element of infinite order in G since the infinite cyclic subgroup it generates cannot have infinite index in G.

Let us fix a finite generating set Σ for G. We know by (1.4.6) that there is a bi-infinite short-lex geodesic inside the Cayley graph $\Gamma := \Gamma^G_\Sigma$.

First we prove that this geodesic is ultimately periodic at its "right end": By linear growth, there is a constant L such that, for infinitely many n, we have:

$$\mathrm{vol}(\mathbb{B}_n) - \mathrm{vol}(\mathbb{B}_{n-1}) \leq L \tag{1.1}$$

Consider a finite subset of vertices $W = \{v_1, v_2, \ldots, v_r\}$ on the geodesic with $r > L$. For any n that satisfies (1.1), there is a pair of distinct vertices $v, w \in W$ such that the geodesic segments of length n starting at these vertices and extending to the right both read the same word – this is the box principle: First observe that there are only $\mathrm{vol}(\mathbb{B}_n) - \mathrm{vol}(\mathbb{B}_{n-1}) \leq L$ many different group elements that these geodesic elements could read. Then note that the short-lex order implies that these group elements have unique words representing them. Hence there are at most L different words we are reading along these segments. Hence two of them agree.

Since there are infinitely many n satisfying (1.1), there is a pair of vertices for which this happens infinitely many times – this, again, is the box principle. Hence there are two vertices in W such that the corresponding infinite segments extending to the right read identical infinite words. It follows that the geodesic is ultimately right periodic.

The group element represented by the period obviously has infinite order. So we have our desired cyclic subgroup of finite index. **q.e.d.**

Exercise 1.4.28. Show that a torsion free group G that contains an infinite cyclic subgroup of finite index is infinite cyclic. Remark: There is no hint for this problem because I want you to find a short, elegant (and probably new) solution.

Chapter 2

Free Groups of Finite Rank

Let M be a set. The free group F_M generated by M is a group that contains M as a subset and that is uniquely determined up to unique isomorphism by the universal property that any map $f : M \to G$ from M to any group G extends to a unique group homomorphism $\varphi_f : F_M \to G$.

The elements of F_M are reduced words in the alphabet $M \uplus M^{-1}$. Multiplication is concatenation of words followed by reduction, i.e., cancellation of subwords mm^{-1} until no longer possible. The empty word serves as the trivial element. Of course there are some claims here to be proved, but you are supposed to have done this already in some other class.

Let F be a non-abelian, finitely generated free group. We will show:

[2.2.6] F is linear.

[2.2.6] F is residually finite.

[2.2.10] F is Hopfian.

[2.3.1] F is not amenable.

[2.4.1] A non-trivial finitely generated group is free if and only if its cohomological dimension is 1.

[2.5.2] F has the finite intersection property.

[2.6.1] F has a solvable word problem.

[2.6.3] F has a solvable conjugacy problem.

2.1 Free Constructions

Definition 2.1.1. Let D be a diagram of groups G_v and homomorphisms $\varphi_{\vec{e}} : G_{\iota(\vec{e})} \to G_{\tau(\vec{e})}$. The <u>direct limit</u> of D is a group $\varinjlim D$ together with homomorphisms $\iota_v : G_v \to \varinjlim D$ such that

1. all triangles

$$\begin{array}{ccc} \varinjlim D & \leftarrow & G_{\tau(\vec{e})} \\ \uparrow & \nearrow f_{\vec{e}} & \\ G_{\iota(\vec{e})} & & \end{array}$$

   commute.

2. Given any other group H together with a family of homomorphism $\varphi_v : G_v \to H$ making the corresponding triangles (as in 1) commutative, there is a unique homomorphism $\pi : \varinjlim D \to H$ such that all triangles

$$\begin{array}{ccc} G_v & \xrightarrow{\varphi_v} & H \\ \downarrow \iota_v & \nearrow \pi & \\ \varinjlim D & & \end{array}$$

   commute.

Exercise 2.1.2. The usual category theoretic nonsense proves uniqueness of direct limits for free. Show that direct limits exist in the category of groups and homomorphisms.

Definition 2.1.3. Let $C \hookrightarrow A$ and $C \hookrightarrow B$ be two monomorphisms. The <u>amalgamated product</u> $A *_C B$ is the direct limit of the diagram

$$A \leftarrow C \rightarrow B$$

The <u>free product</u> $G * H$ of two groups is their amalgamated product along the trivial group:

$$G * H := \varinjlim (G \leftarrow 1 \rightarrow H)$$

These cases arise naturally in topology.

Example 2.1.4 (van Kampen). Let X be a path connected topological space with base point. Assume we are given a open cover $X = U \cup V$ such that U, V, and $X \cap V$ are path connected subsets of X that contain the base point. Then

$$\pi_1(X) = \pi_1(U) *_{\pi_1(U \cap V)} \pi_1(V)$$

Example 2.1.5. The free group on n generators is the free product of n copies of C_∞:

$$F_n := \mathop{\ast}_{i=1}^{n} C_\infty$$

In general, the free group generated by the set M is

$$F_M := \mathop{\ast}_{m \in M} C_\infty$$

Observation 2.1.6. *As a consequence of van Kampen's theorem, we see that a group is free if and only if it is the fundamental group of a graph.*

Corollary 2.1.7. *Subgroups of free groups are free.*

Proof. Let F be a free group. Then F is the fundamental group of a graph. Its subgroups occur as fundamental groups of covers. However, any cover of a graph is a graph. Hence any subgroup of F is the fundamental group of a graph and, therefore, free. **q.e.d.**

Observation 2.1.8. *Obviously, we can construct very large covers. In particular, F_2 contains a copy of $F_{\mathbb{N}}$.*

Observation 2.1.9. *From their geometric realization, we can read off a presentation for free groups:*

$$F_n = \langle x_1, \ldots, x_n \rangle$$

Each generator corresponds to a loop in a wedge of n circles. The Cayley graph of F_n corresponding to this system of free generators is the universal cover of the wedge of circles. It is a tree.

Corollary 2.1.10. *Non-abelian free groups have exponential growth.* **q.e.d.**

Exercise 2.1.11 (Schreier's Index Formula). Let G be a subgroup of F_n of finite index s. Prove that G is isomorphic to $F_{s(n-1)+1}$.

2.2 How to Detect Free Groups

Lemma 2.2.1 (Ping Pong Lemma). *Let G be a group acting on a set X. Suppose H_1 and H_2 are two subgroups of G with cardinalities at least 3 and 2, respectively. Let H be the subgroup generated by H_1 and H_2.*

Assume that there are two non-empty subsets Y_1 and Y_2 in X such that

$Y_2 \not\subseteq Y_1$

$gY_2 \subseteq Y_1$ *for all* $g \in H_1, g \neq 1$

$gY_1 \subseteq Y_2$ *for all* $g \in H_2, g \neq 1$

Then H is isomorphic to the free product $H_1 * H_2$.

Proof. We have to show that a product whose factors are all non-trivial and alternately taken from the groups H_1 and H_2 is non-trivial. We start by considering a product of odd length

$$w = a_1 b_1 a_2 b_2 \cdots b_{r-1} a_r$$

wherein $a_i \in H_1 - \{1\}$ and $b_i \in H_2 - \{1\}$. We have

$$wY_2 = a_1 b_1 a_2 b_2 \cdots b_{r-1} a_r Y_2 \subseteq a_1 b_1 a_2 b_2 \cdots b_{r-1} Y_1 \cdots \subseteq a_1 Y_2 \subseteq Y_1$$

whence w acts non-trivially as $Y_2 \not\subseteq Y_1$.

For a word that starts and ends with a letter from H_2, we conjugate it by a non-trivial element from H_1. As conjugation preserves being trivial or non-trivial, we are reduced to the first case.

For a word of even length, only one boundary letter is in H_1. Let this letter be $a \in H_1$. Conjugation by an element of $H - \{1, a\}$ reduces us to the first case. Here, we need that H_1 has at least three elements. **q.e.d.**

Example 2.2.2. Let T be a tree. An automorphism of T is called hyperbolic if it stabilizes a bi-infinite geodesic in T upon which it acts as a non-trivial shift. Then, this geodesic C_φ is unique and called the axis of the automorphism.

Let φ and ψ be two hyperbolic automorphisms of T with disjoint axes. Then $\langle \varphi, \psi \rangle$ is free.

Proof. We will study the action on the set of ends $\partial_\infty T$. Note that each oriented edge $\vec{e}$ defines a decomposition $\partial_\infty T = \partial_\infty^+(\vec{e}) \uplus \partial_\infty^-(\vec{e})$ as in the picture:

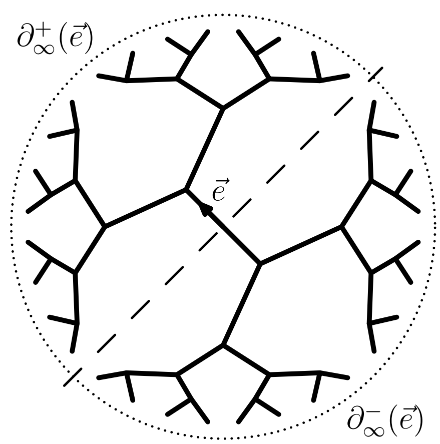

Let $\vec{e}$ be an edge on the geodesic joining C_φ and C_ψ.

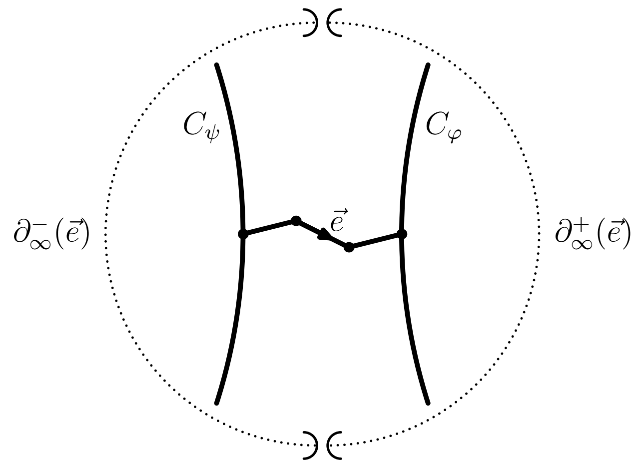

Then any non-trivial power of φ will suck all of $\partial^-_\infty(\vec{e})$ into $\partial^+_\infty(\vec{e})$ and any non-trivial power of ψ will take $\partial^+_\infty(\vec{e})$ into $\partial^-_\infty(\vec{e})$. Hence

$$\langle \varphi, \psi \rangle = \langle \varphi \rangle * \langle \psi \rangle$$

which is a non-abelian free group. **q.e.d.**

Exercise 2.2.3. Suppose φ and ψ are two hyperbolic automorphisms of a tree T whose axes have a finite intersection. Show that sufficiently high powers φ^k and ψ^l generate a free group.

Exercise 2.2.4. Show that F_2 embeds into $C_2 * C_3$. Here, C_n is the cyclic group of order n.

Example 2.2.5. The two matrices $\begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$ and $\begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}$ generate a free group inside $\mathrm{SL}_2(\mathbb{Z})$.

Proof. Each of the two matrices generates an infinite cyclic subgroup inside $\mathrm{SL}_2(\mathbb{Z})$. So we have to consider the group generated by the subgroups

$$H_1 := \left\{ \begin{pmatrix} 1 & 2s \\ 0 & 1 \end{pmatrix} \;\middle|\; s \in \mathbb{Z} \right\}$$

and

$$H_2 := \left\{ \begin{pmatrix} 1 & 0 \\ 2s & 1 \end{pmatrix} \;\middle|\; s \in \mathbb{Z} \right\}$$

Put

$$Y_1 := \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \middle| \, |x| > |y| \right\}$$

and

$$Y_2 := \left\{ \begin{pmatrix} x \\ y \end{pmatrix} \middle| \, |x| < |y| \right\}$$

To verify the hypotheses of the Ping Pong Lemma, we compute

$$\begin{pmatrix} 1 & 2s \\ 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} x + 2sy \\ y \end{pmatrix}$$

whence we have to show that for $|x| < |y|$ it follows that $|x + 2sy| > |y|$. This is obvious from the triangle inequality:

$$|x + 2sy| > |2sy| - |x| = |y| + ((2s - 1)\,|y| - |x|)$$

The other hypothesis is checked analogously. **q.e.d.**

Corollary 2.2.6. *Finitely generated free groups are linear and residually finite.*

Proof. As we have seen, F_2 embeds into $\mathrm{SL}_2(\mathbb{Z})$, which is residually finite as every non-trivial element survives in a factor $\mathrm{SL}_2(\mathbb{Z}/p\mathbb{Z})$ where p is a sufficiently large prime number. **q.e.d.**

Exercise 2.2.7. Prove that any finitely generated group

(a) has only finitely many normal subgroups of index 2007.

(b) has only finitely many subgroups of index 2007.

Definition 2.2.8. A group G is Hopfian if every surjective endomorphism $G \twoheadrightarrow G$ is an automorphism.

Exercise 2.2.9. Show that any finitely generated residually finite group is Hopfian.

Corollary 2.2.10. *Finitely generated free groups are Hopfian.* **q.e.d.**

The following celebrated theorems, due to J. Tits, are way more sophisticated application of the Ping Pong Lemma (or more precisely: a slight variation of it).

Theorem 2.2.11 (Tits [Tits72]). *Over a field of characteristic* 0*, a linear group either is virtually solvable (i.e., it is small) or has a non-abelian free subgroup (i.e., it is big).*

Theorem 2.2.12 (Tits [Tits72]). *A finitely generated linear group either is virtually solvable or has a non-abelian free group.*

These results (and many others that followed) motivate:

Definition 2.2.13. A group G satisfies the Tits Alternative if each finitely generated subgroups either is virtually solvable or contains a non-abelian free group.

Remark 2.2.14. Tits' result states that linear groups satisfy the Tits-Alternative. Another example would be $\mathrm{Out}(F_n)$.

Exercise 2.2.15. Show that a virtually solvable group cannot contain a non-abelian free group.

Exercise 2.2.16. By Tits' theorem, the group $\mathrm{SO}_3(\mathbb{R})$ has a non-abelian free subgroup. Find an embedding of $F_2 \hookrightarrow \mathrm{SO}_3(\mathbb{R})$.

2.3 Kazhdan's Property (T) and Amenability

We already observed (1.3.7) that non-abelian free groups do not have Kazhdan's property (T). However, although property (T) and amenabilty are mutually exclusive for infinite groups, they are not complementary.

Theorem 2.3.1. *Non-abelian free groups are not amenable.*

Proof. We only do the argument for $F_2 = \langle x, y \rangle$ Consider the following bounded functions on F_2:

$$f_1 \quad : \quad w \mapsto \begin{cases} 1 & \text{if } w = 1 \\ 0 & \text{otherwise} \end{cases}$$

$$f_x \quad : \quad w \mapsto \begin{cases} 1 & \text{if } w \text{ starts with } x \\ 0 & \text{otherwise} \end{cases}$$

$$f_{\bar{x}} \quad : \quad w \mapsto \begin{cases} 1 & \text{if } w \text{ starts with } \bar{x} \\ 0 & \text{otherwise} \end{cases}$$

$$f_y \quad : \quad w \mapsto \begin{cases} 1 & \text{if } w \text{ starts with } y \\ 0 & \text{otherwise} \end{cases}$$

$$f_{\bar{y}} \quad : \quad w \mapsto \begin{cases} 1 & \text{if } w \text{ starts with } \bar{y} \\ 0 & \text{otherwise} \end{cases}$$

Obviously,

$$1 = f_1 + f_x + f_{\bar{x}} + f_y + f_{\bar{y}}.$$

Now we form

$$h := f_1 + \bar{x} f_x + x f_{\bar{x}} + \bar{y} f_y + y f_{\bar{y}}$$

We have

$$h^w = \begin{cases} 5 & \text{if } w = 1 \\ 3 & \text{otherwise} \end{cases}$$

This clearly rules out the possibility of an invariant measure on this group. **q.e.d.**

The trick in this proof motivates

Definition 2.3.2. A paradoxical partition of unity is a finite partition of unity

$$1 = f_1 + \cdots + f_r$$

together with an r-tupel of group elements $g_1, \ldots, g_r$ such that:

- Each f is a bounded function.
- There is an $\varepsilon > 0$ satisfying

$$g_1 f_1 + \cdots + g_r f_r \geq 1 + \varepsilon$$

Observation 2.3.3. *No amenable group admits a paradoxical partition of unity.*

2.3.1 Equivalent Formulations for Amenability

We have been using graphs for quite a while. Maybe, it would be good to provide a definition. What we call a graph will be an unoriented multi-graph.

Definition 2.3.4. A graph Γ is a map $\tau : \overrightarrow{\mathcal{E}}_\Gamma \to \mathcal{V}_\Gamma$ where $\mathcal{V}_\Gamma$ is a set (its elements are called vertices) and $\overrightarrow{\mathcal{E}}_\Gamma$ is a free $\mathbb{Z}^2$-set, i.e., a set together with a fixpoint free involution $\mathrm{op} : \overrightarrow{\mathcal{E}}_\Gamma \to \overrightarrow{\mathcal{E}}_\Gamma$. The elements of $\overrightarrow{\mathcal{E}}$ are (oriented) edges. The elements of $\mathcal{E}_\Gamma := \overrightarrow{\mathcal{E}}_\Gamma/\mathrm{op}$ are called (geometric) edges.

An orientation on Γ is a section $\mathrm{o} : \mathcal{E} \to \overrightarrow{\mathcal{E}}$.

We define another map $\iota : \overrightarrow{\mathcal{E}}_\Gamma \to \mathcal{V}_\Gamma$ by $\iota(\vec{e}) := \tau(\mathrm{op}(\vec{e}))$. The map τ assigns to each edge its terminal vertex whereas ι provides the initial vertex of the edge.

A graph is called locally finite if every vertex has only finitely many edges attached to it.

Definition 2.3.5. Let Γ be a locally finite graph. A flow on Γ is a map $\Phi : \overrightarrow{\mathcal{E}}\Gamma \to \mathbb{R}$ on the oriented edges satisfying

$$\Phi(\mathrm{op}(\vec{e})) = -\Phi(\vec{e})$$

For any vertex v of Γ, the net production is

$$P_\Phi(v) := \sum_{v=\iota(\vec{e})} \Phi(\vec{e})$$

A vertex is called a source if its net production is > 0, a sink if its net production is < 0 and it is called balanced if its net production is 0. A flow Φ is called ε-productive if every vertex has net production $\geq \varepsilon$.

A capacity on Γ is a map $C : \overrightarrow{\mathcal{E}}\Gamma \to \mathbb{R}$ satisfying

$$-C^{\vec{e}} \leq C^{\mathrm{op}(\vec{e})}.$$

Note that every non-negative real number defines a capacity by assigning this number to each oriented edge. When using a real number where a capacity should be expected, we silently make this identification. The flow Φ is bounded by the capacity C if $\Phi(\vec{e}) \leq C_{\vec{e}}$ for each oriented edge $\vec{e}$. Note that a geometric edge can have different capacities in its two directions.

Let v and w be two vertices in Γ. A cut is a set of oriented edges in Γ such that any path from v to w has to pass through at least one edge in the cut thereby respecting the given orientation of that edge. (It may pass through other edges in the cut in any representation, but at least one edge has to be crossed "in the right direction".)

For any set of vertices $W \subseteq \Gamma$, its boundary ∂W is the set of edges with one endpoint inside W and the other endpoint outside W.

For a finite set W of vertices, we define the net production by

$$P_\Phi(W) := \sum_{v \in W} P_\Phi(v) = \sum_{\iota(\vec{e}) \in W} \Phi_{\vec{e}}.$$

For any set of oriented edges $\overrightarrow{\mathcal{E}}_0$, we define the total capacity by

$$C_{\overrightarrow{\mathcal{E}}_0} := \sum_{\vec{e} \in \overrightarrow{\mathcal{E}}_0} C_{\vec{e}}.$$

Theorem 2.3.6 (Max-Flow-Min-Cut Theorem). *Let Γ be a finite graph, C a capacity and $P : \mathcal{V}_\Gamma \to \mathbb{R}$ an arbitrary function. Then there is a flow bounded by the capacity C and satisfying $P_\Phi = P$ if and only if for every finite set of vertices W we have*

$$P(W) \leq C_{\partial(W)}.$$

Proof. The condidition is clearly necessary.

The space of all flows Φ obeying C is compact. Define the defect of a flow to be

$$\sum_{v \in \Gamma} |P(v) - P_\Phi(v)| .$$

The defect is obviously continuous. The let Φ be a flow with minimal defect. We have to show that this defect is 0.

So assume by contradiction, the defect is > 0. Then there is a vertex v that does not produce what is prescribed by P. Assume it does not produce enough. Consider the set W of all vertices that can be reached from v by a non-saturated path. The boundary of this set consists of saturated edges. Hence

$$P_\Phi(W) = C_{\partial(W)} \geq P(W) .$$

Since v produces strictly less than prescribed by P, this inequality implies that another vertex w, also in W, does produce to much (or does not consume enough). Since both vertices are in W we can use the path connecting them to improve the defect of Φ.

The case, where the vertex v produces too much if handled analogously. **q.e.d.**

Remark 2.3.7. In a finite graph, the condition $P(W) \leq C_{\partial(W)}$ applied to the complement of a set of vertices implies

$$-C_{\partial((\Gamma - W))} \leq P(W) .$$

The above version of the max-flow-min-cut theorem has the advantage that it generalizes nicely to infinite graphs:

Theorem 2.3.8 (Max-Flow-Min-Cut for Infinite Graphs). *Let Γ be a locally finite graph, C a capacity and $P : \mathcal{V}_\Gamma \to \mathbb{R}$ an arbitrary function. Then there is a flow bounded by the capacity C and satisfying $P_\Phi = P$ if and only if*

$$-C_{\partial(\Gamma - W)} \leq P(W) \leq C_{\partial(W)}$$

for any finite set of vertices W.

Proof. Again, it is clear that the condition is necessary. To prove the other direction, we employ the following strategy: First we use the min-flow-max-cut theorem for finite graphs to prove that for any finite set of vertices W, we can find flow Φ_W that is bounded by 1 yet "looks" ε-productive on W. In a second step, we use an ultrafilter construction, to patch these flows together.

Let W be a finite set of vertices. We collapse the complement to one vertex, i.e., we pretend all edges in $\partial(W)$ are connected to one vertex v_*. This way, we obtain a finite graph Γ_W. The function P induces a production function on Γ_W by setting $P(v_*) := -P(W)$.

We claim that there is a flow Φ_W on Γ_W that realizes P. To apply the max-flow-min-cut theorem for finite graphs, we have to check the inequality. So let U be a set of vertices in Γ_W. If it does not contain the new vertex v_*, then we have

$$P(U) \leq C_{\partial(U)}$$

by our hypotheses – we only have to reinterpret these inequalities in Γ and in Γ_W. If $v_* \in U$, we use

$$-C_{\partial(\Gamma-U)} \leq P(U)$$

to get

$$-C_{\partial(\Gamma_W-U)} \leq P(U)$$

and then

$$P(U) \leq C_{\partial(U)}.$$

Hence (2.3.6) applies.

The second step is to construct a flow that has net production P everywhere. To do this, pick your favorite ultrafilter $\mathcal{U}$ refining the coinitial filter on the directed set D of finite vertex sets in Γ. For each element W in this index set D, already have a flow Φ_W. So if we want to define the global flow Φ on an oriented edge $\vec{e}$, we put

$$\Phi(\vec{e}) := \mathcal{U}\text{-}\lim \Phi_W(\vec{e})$$

It is a routine matter to check that this does the job. **q.e.d.**

Corollary and Definition 2.3.9. *Let Γ be a graph and $\varepsilon \geq 0$. Then the following are equivalent:*

1. *For any finite set of vertices $W \subseteq \Gamma$,*

$$\frac{|\partial W|}{|W|} \geq \varepsilon.$$

2. *There is an ε-productive flow Φ bounded by capacity* 1.

If Γ satisfies the condition 1 above for an $\varepsilon > 0$, we say Γ satisfies a <u>*strong isoperimetric inequality*</u>. **q.e.d.**

Definition 2.3.10. A generating set Σ for a group G is symmetric if $\Sigma = \Sigma^{-1}$. For symmetric generating sets, one usually uses the reduced (right) Cayley graph wherein the edges $g \xrightarrow{x} xg$ and $xg \xrightarrow{x^{-1}} g$ are considered as opposite orientations of one underlying geometric edge.

Corollary 2.3.11. *Let G be a finitely generated group with finite symmetric generating set $\Sigma \not\ni 1$. Then the following are equivalent:*

1. *G is amenable.*
2. *G has a Følner sequence.*
3. *The reduced Cayley graph Γ does not satisfy a strong isoperimetric inequality.*
4. *There is no productive bounded flow on the reduced (right) Cayley graph.*
5. *G does not have a paradoxical partition of unity.*

In fact, one could use ordinary (unreduced, left) Cayley graphs instead. We confine ourselves to reduced right Cayley graphs only to avoid technical issues.

Proof. We already proved (2)$\Longrightarrow$(1) and (1)$\Longrightarrow$(5). The implication (3)$\Longrightarrow$(2) is immediate, and (4)$\Longrightarrow$(3) follows from (2.3.9).

The remaining implication (5)$\Longrightarrow$(4) is done by re-interpreting a flow as a paradoxical partition: Let Φ be an ε-productive flow bounded by 1 on Γ. Define

$$f_\sigma^g := \begin{cases} -\Phi\Big(g \xrightarrow{x} xg\Big) & \text{if } \Phi\Big(g \xrightarrow{x} xg\Big) < 0 \\ 0 & \text{otherwise} \end{cases}$$

be the outbound flow from g to σg. Moreover, put

$$f_1^g := |\Sigma| - \sum_{\sigma\in\Sigma} f_\sigma^g.$$

Obviously, we have a partition of the constant function $|\Sigma|$.

Now compute

$$\begin{aligned} f_1^g + \sum_{\sigma\in\Sigma} \sigma^{-1} f_{\sigma^{-1}}^g &= |\Sigma| + \sum_{\sigma\in\Sigma} f_{\sigma^{-1}}^{\sigma g} - f_\sigma^g \\ &= |\Sigma| + \sum_{\sigma\in\Sigma} \Phi\Big(g \xrightarrow{\sigma} \Sigma g\Big) \\ &\geq |\Sigma| + \varepsilon \end{aligned}$$

Division by $|\Sigma|$ yields a paradoxical partition of unity. **q.e.d.**

Remark 2.3.12. The usual way of proving these equivalences establishes separately

$$(1) \iff (5) \iff (4)$$

and

$$(1) \iff (2) \iff (3)$$

This way, however, no relation in the numerical constants of the isoperimetric inequality and the productivity of the flow is established. Moreover, the implication $(1) \Longrightarrow (2)$ involves functional analysis.

Exercise 2.3.13. A paradoxical decomposition of a group G is a partition

$$G = S_1 \uplus \cdots \uplus S_r \uplus T_1 \uplus \cdots \uplus T_s$$

such that there are group elements $g_1, \ldots, g_r$ and $h_1, \ldots, h_s$ such that

$$G = g_1 S_1 \uplus \cdots \uplus g_r S_r$$

and

$$G = h_1 T_1 \uplus \cdots \uplus h_s T_s.$$

Prove that F_2 has a paradoxical decomposition.

Remark 2.3.14. As a matter of fact, a group has a paradoxical decomposition if and only if it is not amenable. As a criterion to check this, however, using flows or a paradoxical partition of unity is easier.

2.4 Stallings' Theorem

The goal of this section is the following characterization of free groups

Theorem 2.4.1 (Stallings [Stal68]). *A finitely generated group G is free if and only if it has cohomological dimension* 1.

The cohomological dimension $\operatorname{cd} G$ of a group G is an element of $\mathbb{N} \cup \{\infty\}$, and we will define this number in (2.4.1). But we shall outline the proof right away to motivate the exposition. We will argue by induction on the size of a minimal generating set.

Definition 2.4.2. The rank of a finitely generated group G is the size $\operatorname{rk}(G)$ of a generating set of minimal size.

We need the following facts:

1. [2.4.21] $\operatorname{cd} G = 0$ if and only if $G = 1$.

2. [2.4.20] If $\operatorname{cd} G < \infty$, then G is torsion free.

3. [2.4.19] If $H \leq G$, then $\operatorname{cd} H \leq \operatorname{cd} G$.

4. [2.4.22] If $\operatorname{cd} G = 1$, then $\mathrm{e}(G) \geq 2$.

5. [1.4.28] If $\mathrm{e}(G) = 2$ and G is torsion free, the G is the infinite cyclic group.

6. [2.4.27] If $\mathrm{e}(G) = \infty$ and G is torsion free, then $G = A * B$ for two non-trivial subgroups A and B.

7. [2.4.34] For finitely generated groups A and B,

$$\mathrm{rk}(A * B) = \mathrm{rk}(A) + \mathrm{rk}(B)$$

   This is known as Grushko's Theorem.

Proof of Stallings' Theorem. First observe that the group G is torsion free because of (2). We induct on $\mathrm{rk}(G)$. If $\mathrm{rk}(G) = 1$, we have a cyclic group which must be C_∞ since this is the only torsion free cyclic group.

So assume $\mathrm{rk}(G) > 1$. From (4) we know that G has at least two ends. If it had two ends, it would be virtually cyclic by (1.4.13). Since G is torsion free, by (5), it had to be C_∞ which has rank 1. Hence G has infinitely many ends. Now, (6) implies that

$$G = A * B$$

for some finitely generated, non-trivial subgroups A and B.

By (3), $\operatorname{cd} A \leq 1$, and by (1), $\operatorname{cd} A = 1$. Finally Grushko's Theorem implies $\mathrm{rk}(A) < \mathrm{rk}(G)$ whence we can infer by induction that A is a free group.

Analogously, B is free. Hence G is a free product of free groups and, therefore, free. **q.e.d.**

Exercise 2.4.3. Prove that $\mathrm{rk}(F_n) = n$.

Corollary 2.4.4. *Every minimal set of generators for F_n is a set of free generators.*

Proof. Let Σ be a minimal generating set for F_n. The inclusion $\Sigma \hookrightarrow F_n$ extends to a group homomorphism $F_\Sigma \to F_n$ which is onto because Σ generates F_n. On the other hand, by exercise (2.4.3), $|\Sigma| = \mathrm{rk}(F_n) = n$. Hence we have a surjection of free groups of the same rank. Since these groups are Hopfian (2.2.10), the map $F_\Sigma \to F_n$ is an isomorphism. **q.e.d.**

2.4.1 Cohomology of Groups and the Eilenberg-Ganea Problem

Fact and Definition 2.4.5. Let G be a group, R a commutative ring with unity $1 \neq 0$, and M a (left) RG-module – note that the involution $g \mapsto g^{-1}$ allows us to regard any left RG-module as a right module and vice versa. For any projective resolution

$$\cdots \to P_2 \to P_1 \to P_0 \to R \to 0$$

the homology groups

$$\mathrm{H}_*(G; M) := \mathrm{H}_*(P_* \otimes_{RG} M)$$

and cohomology groups

$$\mathrm{H}^*(G; M) := \mathrm{H}^*(\mathrm{Hom}_{RG}(P_*, M))$$

are independent of the chosen projective resolution.

It is, of course, crucial to find nice resolutions.

Example 2.4.6. For the finite cyclic group, we can cook up a very nice periodic resolution. Observe that

$$RC_n = R\,[t] \,/\, \langle t^n \rangle \,.$$

With this identification, we can write down the following resolution:

$$\cdots \xrightarrow{\times\left(t^{n-1}+\cdots+1\right)} RG \xrightarrow{\times(t-1)} RG \xrightarrow{\times\left(t^{n-1}+\cdots+1\right)} RG \xrightarrow{\times(t-1)} RG \xrightarrow{\varepsilon} R \to 0$$

From this resolution, we get:

$$\begin{aligned} \mathrm{H}_i(C_n; \mathbb{Z}) &= \begin{cases} \mathbb{Z} & i = 0 \\ \mathbb{Z}^n & i \text{ is odd} \\ 0 & i \text{ is even and } > 0 \end{cases} \\ \mathrm{H}^i(C_n; \mathbb{Z}) &= \begin{cases} \mathbb{Z}^n & i \text{ is even} \\ 0 & i \text{ is odd} \end{cases} \end{aligned}$$

Example 2.4.7. Let K be a contractible simplicial complex upon which G acts freely. That is, no simplex is fixed by any group element. Then the action induces an action of G on the simplicial chain complex $\mathrm{C}_*(K; R)$ which thereby turns into a chain complex of RG-modules. Theses modules are free since G acts freely on K. Since K

is contractible, the simplicial chain complex is exact. Appending the augmentation map (sum up the coefficients on vertices) yields a free resolution

$$\cdots \mathrm{C}_2(K;R) \xrightarrow{\partial} \mathrm{C}_1(K;R) \xrightarrow{\partial} \mathrm{C}_0(K;R) \xrightarrow{\varepsilon} R$$

which can be used to compute the (co)homology of G.

How to come up with a good candidate for K? A generic choice would be the infinite join

$$K := \mathop{\ast}_{\mathbb{N}} G$$

with the diagonal group action. In particular, we can always find a resolution of the trivial RG-module R by free modules.

Example 2.4.8. Instead of a simplicial complex, one could use a cell complex provided the group acts freely on cells. One way to construct a contractible, free G-complex is to start with a Cayley graph for G. Glue in free G-sets of 2-cells to kill loops. This might introduce non-trivial π_2. Glue in free G-sets of 3-cells to get rid of this, and continue to kill all fundamental groups.

The result is a contractible CW-complex with a free G-action. The advantage is that the 1-skeleton still looks like the Cayley graph.

Definition 2.4.9. The <u>cohomological dimension over R</u> of a group G is is the last dimension for which the cohomology functor is non-trivial, i.e, it is the least element $\operatorname{cd} G \in \mathbb{N} \cup \{\infty\}$ satisfying

$$\mathrm{H}^i(G;-) = 0 \text{ for all } i > \operatorname{cd}_R G.$$

Definition 2.4.10. The <u>geometric dimension</u> of a group G is the smallest dimension $\operatorname{gd} G$ of a contractible simplicial complex upon which G can act freely.

Observation 2.4.11. *Using the resolution of (2.4.7), we see*

$$\operatorname{cd}_R G \leq \operatorname{gd} G$$

Observation 2.4.12. *For $H \leq G$, we have* $\operatorname{gd} H \leq \operatorname{gd} G$.

Observation 2.4.13. *Only the trivial group has geometric dimension* 0. *A group is free if and only if it has geometric dimension* 1.

Hence, we can restate Stallings' Theorem as follows.

Theorem 2.4.14. *Any finitely generated group G of cohomological dimension* 1 *has geometric dimension* 1.

This is more in line with the following:

Fact 2.4.15 (Eilenberg-Ganea [EiGa57]). *For any group G,*

$$\operatorname{gd} G \leq \max(\operatorname{cd} G, 3).$$

Remark 2.4.16. Settling the case $\operatorname{cd} G = 2$ is the Eilenberg-Ganea Problem. It is generally believed that there are groups of cohomological dimension 2 with geometric dimension 3. In fact, some particular groups are conjectured to have this property. However, as of now, all methods of estimating the geometric dimension of groups are based on homological machinery. Hence, we do not have a proof that one of the alleged examples actually does the trick.

Fact 2.4.17. *If $m = \operatorname{cd}_R G < \infty$, and*

$$P_{m-1} \to \cdots \to P_1 \to P_0 \to R \to 0$$

is a partial projective resolution, then the kernel

$$P_m := \ker(P_{m-1} \to P_{m-2})$$

is projective. In particular, there is a finite projective resolution

$$0 \to P_m \to \cdots \to P_1 \to P_0 \to R \to 0$$

for R.

Observation 2.4.18. *Let H be a subgroup of G. Any free RG-module is a free RH-module. Hence any projective RG-module, begin a direct summand of a free RG-module, is a fortiori a projective RH-module.*

Corollary 2.4.19. *For $H \leq G$, we have $\operatorname{cd}_R H \leq \operatorname{cd}_R G$.* **q.e.d.**

Since a group with torsion has a finite cyclic subgroup, we can infer from (2.4.6):

Corollary 2.4.20. *If $\operatorname{cd} G < \infty$, the group G is torsion free.* **q.e.d.**

Proposition 2.4.21. *If $\operatorname{cd} G = 0$, then G is trivial.*

Proof. By (2.4.17), $\mathbb{Z}$ is a projective $\mathbb{Z}G$-module. Hence the augmentation map $\mathbb{Z}G \to \mathbb{Z}$ splits. The image of $1 \in \mathbb{Z}$ under the split must be a G-invariant non-trivial element of $\mathbb{Z}G$. Hence it has constant coefficients for all group elements. This can only happen for finite G as only finitely many coefficients can be non-zero. Hence G is finite. On the other hand, G is torsion free. **q.e.d.**

Exercise 2.4.22. Let G be finitely generated. Show that $\operatorname{cd} G = 1$ implies that G has at least two ends. Hint: Let Γ be a Cayley graph for G. Relate several cohomology theories of Γ and G with coefficients in $\mathbb{Z}^2$ and $\mathbb{Z}^2 G$.

2.4.2 Tree Actions and Free Products

We want to detect a non-trivial splitting of a group as a free product. For this purpose the following geometric criterion comes in handy.

Proposition 2.4.23. *Let G be a group acting on a tree T without terminal vertices and at least some branch points. Assume that the induced action on the set of geometric edges is free and transitive. Then G splits non-trivially as a free product.*

Proof. Since the group acts transitively on the edges, there are either two orbits of vertices or the action on the vertices is transitive, too. The first case is dealt with in (2.4.25) and the second case is done in (2.4.24). **q.e.d.**

Lemma 2.4.24. *Let G act on a tree T such that the following conditions are satisfied:*

1. *G acts freely and transitively on the set of geometric (unoriented) edges.*
2. *G does act transitively on the set of vertices.*
3. *T has no terminal vertices and is not isomorphic to a line.*

Let e be an edge in T that connects the vertices v and w. Let G_v denote the stabilizer of v and let g be an element in G such that $gv = w$. Then $G = G_v * \langle g \rangle$.

Proof. Let G_w denote the stabilizer of w. Let U be an open neighborhood of e. The tree T is connected and covered by the G translates of U. Hence G is generated by

$$\{g \in G \mid gU \cap U \neq \emptyset\} = G_v g^{-1} \cup G_v \cup \{g\} \cup G_w \cup G_w g$$

On the other hand $G_w = gG_vg^{-1}$. Hence

$$G = \langle G_v, g \rangle$$

Note that g has infinite order and shifts e to a neighboring edge. Hence we can construct a bi-infinite geodesic C upon which g acts as a unit shift.

Now, consider the action of G on $\partial_\infty T$. We define two subsets. Let $\mathcal{E}_v$ be the set of ends represented by geodesic paths starting at v avoiding C, and let $\mathcal{E}_e$ be the complement of $\mathcal{E}_v$. Obviously, every non-trivial power of g moves $\mathcal{E}_v$ into $\mathcal{E}_e$.

On the other hand, a non-trivial element of G_v cannot move e to $g^{-1}e$ since otherwise G would flip the orientation of an edge and therefore act with non-trivial edge stabilizers. Hence non-trivial elements of G_v take $\mathcal{E}_e$ into $\mathcal{E}_v$.

Since G acts transitively on the set of edges and every vertex has degree ≥ 3, the stabilizer G_v is non-trivial. On the other hand $\langle g \rangle$ is infinite cyclic. Thus, the Ping Pong Lemma (2.2.1) applies. Hence $G = \langle G_v, g \rangle = G_v * \langle g \rangle$. **q.e.d.**

Exercise 2.4.25. Let G act on a tree T such that the following conditions are satisfied:

1. G acts freely and transitively on the set of geometric (unoriented) edges.

2. G does not act transitively on the set of vertices.

3. T has no terminal vertices and is not isomorphic to a line.

Let e be an edge in T that connects the vertices v and w. Let G_v and G_w denote the stabilizers of these vertices. Then $G = G_v * G_w$.

2.4.3 Stallings' Structure Theorem

This section is devoted to the proof of the following

Theorem 2.4.26 (Stallings' Structure Theorem). *Let G be a finitely generated group with infinitely many ends. Then there is a tree T upon which G acts edge-transitively without inversions and such that all edge stabilizers are finite.*

Before we emabark on this big theorem, let us note the following consequence which is of immediate concern to us in our disscussion of torsion free groups:

Corollary 2.4.27. *Let G be a finitely generated torsion free group with infinitely many ends. Then $G = A * B$ for some non-trivial subgroups A and B.*

Proof. By the Structure Theorem there is a tree upon which G acts with finite edge stabilizers and one edge-orbit. Since G is torsion free, the stabilizers are actually trivial. **q.e.d.**

The idea of the proof is to replace a Cayley graph for G by a tree that somewhat interpolates between the ends of the Cayley graph. The group G will act on this tree, and then we can use (2.4.23) to ensure a splitting. The argument is based on the proof in [DiDu89].

So our goal is to find a nice tree for G to act on. We will actually find the set of edges first (they correspond to splittings of the Cayley graph) – more precisely, we will find the oriented edges of the tree. Then we will have to make up the vertices. This process is completely formal and motivates the definition of tree sets.

Definition 2.4.28. A tree set is a set T together with a fixpoint free involution $\overline{(-)} : T \to T$ and a binary relation $\rightarrowtail$ satisfying the following axioms:

1. The relation $\rightarrowtail$ is a partial ordering.

2. $t \rightarrowtail t' \Longleftrightarrow \overline{t'} \rightarrowtail \bar{t}$.

3. For any two elements $t, t' \in T$ exactly one of the following six cases occurs:
$$t = t', \quad \bar{t} = t', \quad t \rightarrowtail t', \quad \bar{t} \rightarrowtail t', \quad t \rightarrowtail \overline{t'}, \quad \bar{t} \rightarrowtail \overline{t'}.$$

4. For any element $t \in T$, the set $T_t := \{t' \in T \mid t' \rightarrowtail t\}$ contains no infinite chain $t'_1 \rightarrowtail t'_2 \rightarrowtail t'_3 \rightarrowtail \cdots$.

Remark 2.4.29. It follows from the axioms (1)–(3) alone that, for any two elements $t_1, t_2 \in T$, the interval

$$[t_1, t_2] := \{t' \in T \mid t_1 \rightarrowtail t' \rightarrowtail t_2\}$$

is totally ordered. Hence axiom (4) just states that intervals are finite.

Exercise 2.4.30. Show that intervals in tree sets are totally ordered.

Exercise 2.4.31. Given a tree set T, we construct a graph as follows: the set of oriented edges is T, the vertex set is $\mathcal{V} := T/\sim$ where

$$t \sim t' :\Longleftrightarrow t \rightarrowtail \overline{t'} \text{ and } \left[t, \overline{t'}\right] = \emptyset,$$

the endpoint map $\tau : T \to \mathcal{V}$ is the canonical projection $T \to T/\sim$, and the initial vertex map $\iota : T \to \mathcal{V}$ is given by $\iota(t) := \tau\left(\bar{t}\right)$. Show that this graph is a tree – here we use the convention that two opposite oriented edges form one geometric edge, thus avoiding bigons.

Given this equivalence of trees and tree-sets, we can restate the Structure Theorem in the form suitable for proof:

Theorem 2.4.32 (Structure Theorem, Tree-set Version). *Let G be a finitely generated group with infinitely many ends. Then there is a tree set T upon which G acts with finite stabilizers and at most two orbits.*

Lemma 2.4.33. *Let Γ be a Cayley graph for G with respect to a finite generating system. If Γ has more than three ends, any infinite connected subgraph Δ with finite boundary comprises at least two ends of Γ.*

Proof. Let C be a finite connected subgraph containing $\partial(\Delta)$ with at least three infinite complementary components. Chose a translate gC inside Δ – this exists as Δ is infinite: pick a translate as far away as to make sure it does not intersect C nor the finite complementary regions (of which there are only finitely many as C has finite boundary).

As the Cayley graph is homogeneous, gC will split the space of ends into at least three non-empty subsets. As $\partial(\Delta) \subseteq C$ and C is connected, at least two of these are covered by Δ. **q.e.d.**

Proof of the Structure Theorem (2.4.32). Let Γ be a Cayley graph for G over a finite generating set, and define

$$\mathcal{H}_i := \left\{ U \subset \Gamma \,\middle|\, |U| = \infty = \left|\overline{U}\right| \text{ and } |\partial(U)| < i \right\}.$$

Since Γ has more than one end, these sets are non-empty for sufficiently large i. Let m be the least index for which $\mathcal{H} := \mathcal{H}_m \neq \emptyset$. Note that, because of minimality, the vertex collections in $\mathcal{H}$ are "connected", i.e., for each $U \in \mathcal{H}$, any two points in U can be joined by a path that passes through vertices in U only.

Claim A. Every infinite descending chain $U_1 \supsetneq U_2 \supsetneq \cdots$ in $\mathcal{H}$ has empty intersection.

Proof. Assume the chain had non-empty intersection U_∞. Pick an index i_1. As U_{i_1} is connected, there is an edge e_{i_1} connecting U_∞ to $U_{i_1} - U_\infty$. Hence there is an index i_2 such that, for any $j \geq i_2$, the edge e actually bridges between U_j and $U_{i_1} - U_j$. In particular, $e \in \partial(U_j)$.

Now replace i_1 by i_2 and argue in exactly the same way, to find a new edge e_{i_2} and an index i_3 as above. Observe that $e_{i_2} \neq e_{i_1}$. Once you constructed $e_{i_{m+1}}$, observe that $\partial\left(U_{i_{m+2}}\right)$ has more than m edges. This contradicts or definition of $\mathcal{H}$. □

Let $U_0 \in \mathcal{H}$ be minimal with $1 \in U_0$.

Claim B. For any group element $V \in \mathcal{H}$, at least one of the following set is finite:

$$U_0 \cap V, \quad \overline{U_0} \cap V, \quad U_0 \cap \overline{V}, \quad \overline{U_0} \cap \overline{V}. \tag{2.1}$$

Proof. Assume by contradiction, all of these intersections are infinite. First, we claim

$$4m \leq \partial(U_0 \cap V) + \partial\left(\overline{U_0} \cap V\right) + \partial\left(U_0 \cap \overline{V}\right) + \partial\left(\overline{U_0} \cap \overline{V}\right) \leq 2\,|\partial(U_0)| + 2\,|\partial(V)| = 4m.$$

The first inequality follows from the definition of m. The second one follows from this picture:

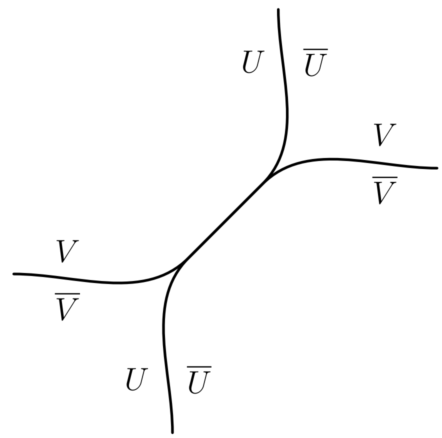

Consider the intersection that contains 1, say $U_0 \cap \overline{V}$. Obviously, $1 \in U_0 \cap \overline{V} \subseteq U_0$. But on the other hand, $U_0 \cap \overline{V} \in \mathcal{H}$. By minimality of U_0, we have $U_0 \cap \overline{V} = U_0$, whence $V \cap U_0 = \emptyset$. Thus, not all intersections are infinite. $\square$

Claim C. If two of the intersections in (2.1) are finite, it is either the pair

$$U_0 \cap V \quad \overline{U_0} \cap \overline{V}$$

or the pair

$$\overline{U_0} \cap V \quad U_0 \cap \overline{V}.$$

Proof. Consider the diagram

$$\begin{array}{ccccc} (U_0 \cap V) & \cup & \left(U_0 \cap \overline{V}\right) & = & U_0 \\ \cup & & \cup & & \\ \left(\overline{U_0} \cap V\right) & \cup & \left(\overline{U_0} \cap \overline{V}\right) & = & \overline{U_0} \\ \shortparallel & & \shortparallel & & \\ V & & \overline{V} & & \end{array}$$

Since U_0, $\overline{U_0}$, V, and $\overline{V}$ are all infinite, the claim follows. $\square$

Claim D. Let $U, V \in \mathcal{H}$ such that $\overline{U} \cap U_0$ is finite, then there are only finitely many $g \in G$ such that

$$\overline{U} \cap V \text{ and } \overline{V} \cap U_0 \text{ are finite.}$$

Proof. Because of (2.4.33), there is a compact subset C in the Cayley graph Γ of G such that the intersections

$$\overline{U} \cap \overline{C}$$

and

$$\overline{U_0} \cap \overline{C}$$

contain at least two infinite components.

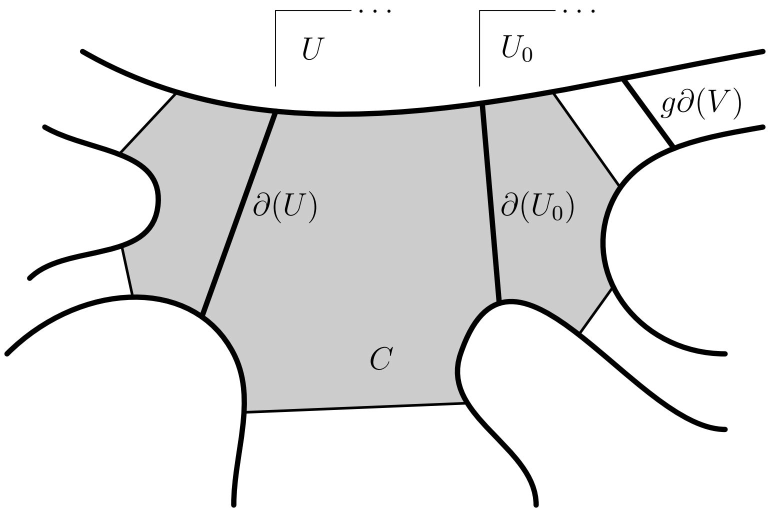

Choose a finite connected subgraph Δ containing $\partial(V)$. Then for all but finitely many $g \in G$, the translate $g\partial(V) \subseteq g\Delta$ is contained in an infinite complementary component of C. For these g, it is impossible that $\overline{U} \cap V$ and $\overline{V} \cap U_0$ are finite. □

We define an equivalence relation on $\mathcal{H}$ by

$$U \cong V :\Longleftrightarrow \overline{U} \cap V \text{ and } U \cap \overline{V} \text{ are finite.}$$

In addition, we define a partial order relation $\rightarrowtail$ on $\mathcal{H}$ by

$$U \rightarrowtail V :\Longleftrightarrow \overline{U} \cap V \text{ is finite but } U \ncong V.$$

Put

$$T := \{gU\} \cup \{g\overline{U}\} \big/ \cong$$

where

$$U \cong V :\Longleftrightarrow (U \rightarrowtail V \text{ and } V \rightarrowtail U).$$

Obviously $\rightarrowtail$ descends and defines a partial order on T which we will also denote by $\rightarrowtail$.

Claim E. T is a tree set.

PROOF. Obviously, $\rightarrowtail$ is a partial order, and taking complements is an order reversing involution. The condition (3) is satisfied by claim (B). Finally, intervals are finite by (D). □

So finally, we have constructed the tree set. What about the action of G? The number of orbits is clearly bounded by two since we used the orbits of U_0 and $\overline{U_0}$ to define T. That the stabilizers are finite, follows from (D) for the case $U = U_0$. **q.e.d.**

2.4.4 Grushko's Theorem

Theorem 2.4.34. *If a free product $A * B$ has a set of n generators, then it has a separated set of n generators, i.e., a generating set contained in $A \cup B$. In particular,*

$$\mathrm{rk}(A * B) = \mathrm{rk}(A) + \mathrm{rk}(B)\,.$$

We give a version of Stallings' proof based on [Stal88]. We will, however, avoid the use of two-complexes and work with folds (also invented by Stallings). This makes the proof more combinatorial. It actually yields an algorithm for the construction of a separated generating set.

Proof. Let us start with an arbitrary set of n generators for $A * B$. Recall that the elements are essentially words over the alphabet $A \cup B - 1$. We will devise an algorithm that takes this generating set as an input and has a separated generating set of at most equal size as its output.

Our main bookkeeping device is a graph with edge labels in $A \cup B$. Here is the precise data structure that we use:

Definition 2.4.35. An A, B-labeling is a connected graph Γ together with a map $\phi : \overrightarrow{\mathcal{E}}\Gamma \to A \cup B - 1$ satisfying

$$\phi(\mathrm{op}(\vec{e})) = \phi(\vec{e})^{-1}\,.$$

The values of ϕ are called labels. We color all edges red that have their label in A. The remaining edges are colored blue. A vertex in Γ is monochromatic if all edges in its link have the same color. A path in Γ is called monochromatic if it contains edges of one color only.

Since any directed edge "reads" an element of $A * B$, any directed path "reads" the product. We call a path null-homotopic if it reads the trivial element in $A * B$.

Clearly, we have a homomorphism

$$\phi : \pi_1(\Gamma, v) \to A * B$$

for any vertex $v \in \Gamma$.

An A, B-labeling is generating if the induced homomorphism is surjective.

The start for our algorithm is a rose:

Observation 2.4.36. *We can realize any finite generating set for $A * B$ by a generating A, B-labeling modeled on a rose with subdivided loops.*

The goal of the algorithm is a rose too:

Observation 2.4.37. *Any generating $A{*}B$-labeling modeled on a rose with undivided loops, i.e., a graph with only one vertex, represents a separated generating set for $A{*}B$.*

Our algorithm modifies A, B-labelings. The goal is to reduce the number of vertices. Eventually, we have a graph with only one vertex, which represents a separated generating set.

For the reduction step, assume our labeling Γ has at least two vertices. Then Γ contains path that is not a loop. This path can read any element of $A * B$, but as $\phi : \pi_1(\Gamma, v) \to A * B$ is surjective, we can append a loop so that we obtain a null-homotopic path that connects two different vertices.

The next step in the reduction is to come up with a monochromatic, null-homotopic path connecting two different vertices. For those, who do not care about a construction: Among all null-homotopic non-loops, chose a shortest one and prove that this must be a monochromatic path.

A more constructive approach runs like this: We already have a null-homotopic non-loop. First remove from this path all monochromatic, null-homotopic loops. This will shorten the path and we will still have a null-homotopic non-loop. The edges in this path come in runs of edges of the same color. Since we remove null-homotopic, monochromatic loops, each of these runs is either not null-homotopic or not a loop. Each run gives an element in A or B. The product of these has to be trivial (the whole path is null-homotopic), but by the definition of runs, the factors are taken alternatingly from A and B. An alternating product of non-trivial elements, however, cannot be trivial in the free product $A * B$. Hence one of the runs reads the identity element. Then, this run is not a loop. Hence we found a monochromatic, null-homotopic path connecting two different vertices.

Note that our path has length ≥ 2 since we do not allow for trivial edge labels. Since it reads the trivial element, the first edge reads the inverse of the rest. Therefore, adjusting the orientations, we get the following picture:

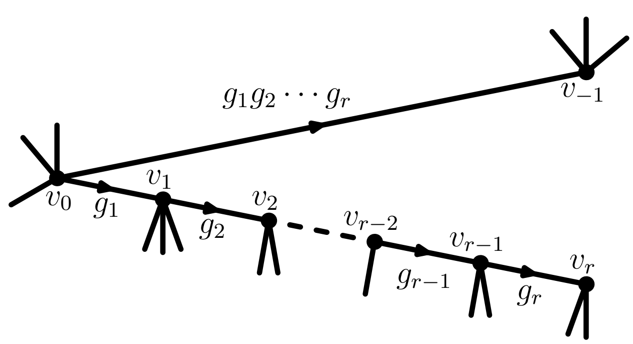

Now, we perform a <u>fold</u>, i.e., we replace the picture above by

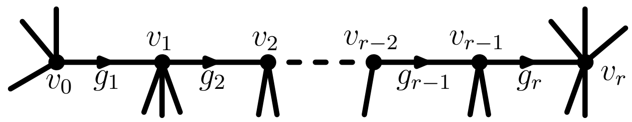

Clearly, the fundamental group of Γ did not change: we removed one edge and reduced the number of vertices by one. Moreover, the induced homorphism ϕ is still surjective: for every closed loop before the fold we can find a closed loop after the fold reading the same group element. So we still have a generating A, B-labeling. But the number of vertices dropped by one. Keep going, until there is only one vertex left. This completes the proof.

Wait a minute. Pictures can be so misleading, and in this prove, there is a serious gap: We want to remove the initial edge of our path. The justification is that we do not need it as the complementary segment of the path already reads the same group element. But what if this segment passes through the initial segment? Then deleting it would cut the path – ouch.

Hm. Here is a patch. First let us observe that there is another process that reduces the number of vertices in Γ. Suppose v is a monochromatic vertex, then we can remove v. To do this, we replace

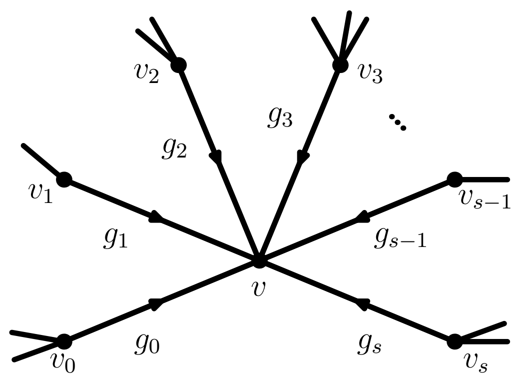

by

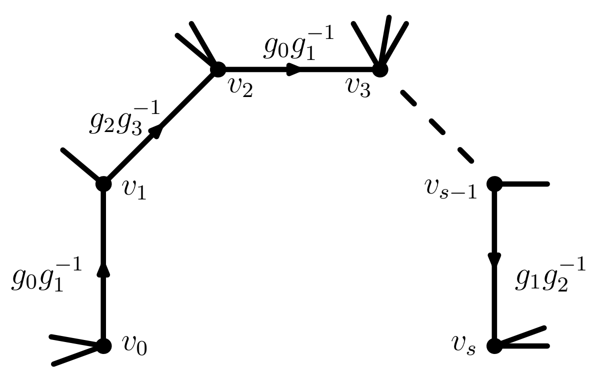

So let us have a closer look at the initial edge e of our monochromatic path. We assume that our path is red. Let v_{-1} be the first and v be the second vertex in the path. Our plan is to do an unfold along e first to split v_{-1} into two vertices, one of which keeps the connections to blue edges whereas the other one stays connected to the red edges. So we replace the picture

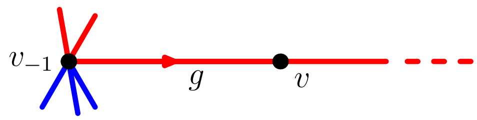

by

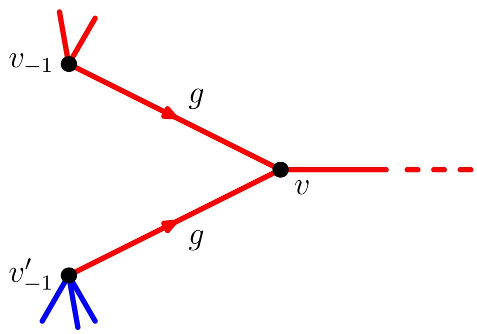

Note that we increased the number of vertices by one. On the other hand, we have a path starting at the new vertex v'_{-1} which is monochromatic and null-homotopic and whose first edge is not used a second time. Hence we can immediately follow the unfold by a fold which restores the number of vertices. So what was the progress?

By the unfold, the vertex v_{-1} became monochromatic. Hence we can get rid of one of them, thereby reducing the number of vertices. So we did one step back and two steps forward, and this does complete the proof. **q.e.d.**

Remark 2.4.38. There is a topological interpretation of this proof, which explains why we call paths that read the trivial element "null-homotopic": Realize A and B as a fundamental groups of base pointed spaces. Then, by van Kampen's Theorem, their wedge X has fundamental group $A * B$. A generating set can be realized as a map from the n-rose Y to X. A path in Y is that maps to a null-homotopic loop in X is what we called null-homotopic.

One of Stallings proofs for Grushko's Theorem takes place in this setting: The preimage of the wedge-point is probably not connected. So we want to reduce the number of its components. This is done by glueing two cells onto Y without changing the homotopy type. This procedure corresponds to the folds above: we just collapsed the two cell immediately after the gluing. The patch is only needed because we insist on this immediate return to a graph. For this reason, the topological proof is, in fact, shorter and more elegant.

2.5 The Hanna Neumann Conjecture

Definition 2.5.1. A group G has the finite intersection property if the intersection of any two finitely generated subgroups in G is finitely generated.

Theorem 2.5.2 (Howson [Hows54]). *Free groups enjoy the finite intersection property.*

By now, there are many proofs of this theorem. The one given here is stolen from [Shor90].

Definition 2.5.3. Let G be a group with finite generating set Σ. A subgroup $H \leq G$ is quasi-convex with respect to Σ if there is a constant $R \geq 0$ such that every geodesic path in the Cayley graph Γ_{Σ}^{G} joining two points in H lies in an R-neighbourhood of H. That is, every point on such a path has distance $\leq R$ to at least one point in $H \subset \Gamma$.

Example 2.5.4. Any finitely generated subgroup H of a free group is quasi-convex with respect to the standard generators: Let B be a ball in the Cayley tree Γ centered at 1 containing all generators of H. The union

$$HB = \bigcup_{h \in H} hB$$

of H-translates of B is connected and hence a subtree. Any geodesic joining two point of H in Γ actually lies in HB. The constant R therefore can be chosen to be the radius of B.

Proposition 2.5.5. *Quasi-convex subgroups are finitely generated.*

Proof. Let G, H, Σ, and R be as in the definition (2.5.3), and let B be the open ball in $\Gamma = \Gamma_{\Sigma}^{G}$ of radius $R + 1$. It is easy to see that

$$X = HB \subseteq \Gamma$$

is connected and that

$$\Xi := \{h \in H \mid hB \cap B \neq \emptyset\}$$

is finte. By (1.4.4) this finite set generates H. **q.e.d.**

Proposition 2.5.6. *The intersection of two quasi-convex subgroups is quasi-convex. More precisely, let G be a group with finite generating set Σ and let A and B be two subgroups that are both quasi-convex with respect to Σ. Then $A \cap B$ is quasi-convex with respect to Σ.*

Proof. Let R_A and R_B be the quasi-convexity constants for the two subgroups. Let $\mathcal{P}_R$ be the set of paths in $\Gamma = \Gamma_\Sigma^G$ starting at 1 of length R. For any such path p let $\overline{p}$ denote its end point – note that this is a vertex in Γ and therefore an element of the group G. Consider the finite set

$$\Pi := \{(p,q) \in \mathcal{P}_{R_A} \times \mathcal{P}_{R_B} \mid \overline{p}a = \overline{q}b \text{ for some } a \in A,\ b \in B\}$$

For any pair $(p,q) \in \Pi$, pick two paths $\gamma_{(p,q)}$ (with $\overline{\gamma_{(p,q)}} \in A$) and $\delta_{(p,q)}$ (with $\overline{\delta_{(p,q)}} \in B$) such that $\overline{p\gamma_{(p,q)}} = \overline{q\delta_{(p,q)}}$. Let R be the maximum length that occurs as a path of the form $p\gamma_{(p,q)}$. We will show that any point that has distance $\leq R_A$ to A and distance $\leq R_B$ to B must have distance $\leq R$ to $A \cap B$. From this, quasi-convexity follows since the statement of course applies to points on geodesic paths joining points in $A \cap B$.

So let $g \in G$ be in the R_A-neighbourhood of A and in the R_B-neighbourhood of B. Then there is a path p of length $\leq R_A$ from g to some element of A. Similarly there is a short path q connecting g to B. Hence $(p,q) \in \Pi$. Observe that $\overline{gp} \in A$ whence $\overline{gp\gamma_{(p,q)}} \in A$. Similarly, $\overline{gq\delta_{(p,q)}} \in B$, but because $\overline{p\gamma_{(p,q)}} = \overline{q\delta_{(p,q)}}$ it follows that $\overline{gp\gamma_{(p,q)}} \in A \cap B$. Hence g has distance $\leq R$ to $A \cap B$. **q.e.d.**

Proof of (2.5.2). By (2.5.4), finitely generated subgroups of free groups are quasi-convex. By (2.5.6), the intersection of two of these is quasi-convex itself. Finally, by (2.5.4), the intersection is finitely generated. **q.e.d.**

Remark 2.5.7. Let G and H be two finitely generated subgroups of the free group F. As G and H are free groups of finte rank, one might ask how the rank of $G \cap H$ relates to the ranks of G and H. Hanna Neumann [Neum55] showed:

$$\mathrm{rk}(G \cap H) - 1 \leq 2\,(\mathrm{rk}(G) - 1)\,(\mathrm{rk}(H) - 1)$$

She asked whether

$$\mathrm{rk}(G \cap H) - 1 \leq (\mathrm{rk}(G) - 1)\,(\mathrm{rk}(H) - 1)$$

This problem is still open and know as the Hanna Neumann conjecture.

Exercise 2.5.8. Show that $F_2 \times C_\infty$ does not have the finite intersection property. That is, find two finitely generated subgroups whose intersection is not finitely generated.

2.6 Equations in Free Groups and the Conjugacy Problem

Reduced words provide normal forms for elements in F_n. Hence it is easy to decide if two words in the free generators represent the same group element. We can phrase this as

Observation 2.6.1. *Free groups have a solvable word problem.*

This is only one of many algorithmic problems one might study for a group with a fixed generating set:

Definition 2.6.2. The word problem for a group $G = \langle \Sigma \rangle$ is to decide algorithmically for any two words in $\Sigma \uplus \Sigma^{-1}$ whether they represent the same group element.

The conjugacy problem is to decide algorithmically whether two given words w_1 and w_2 represent conjugated group elements, i.e., if the equation

$$Xw_1X^{-1} = w_2$$

has a solution in the ambient group.

The subgroup membership problem is to decide algorithmically whether a given word represents an element of the subgroup generated by a finite list of given elements.

The subgroup conjugacy problem is to decide algorithmically for two finite sets of words if the two subgroups they generate are conjugate.

Theorem 2.6.3. *The conjugacy problem in F_n is solvable. More precisely, if the equation*

$$XuX^{-1} = w$$

has at least one solution in F_n, then there is a solution v such that

$$|v| < \frac{|u| + |w|}{2}.$$

Proof. Suppose v is a solution of minimal length. So we have $vuv^{-1} = w$. Assuming that the words v, u, and w are reduced, we observe that cancellations on the left had side occur only on the boundaries of u. We have to keep track of these cancellations.

First, there might be cancellations on both sides of u. In this case, we are undoing a conjugation. Since we cannot conjugate something nontrivial into the empty word, the number of those cancellations is $< \frac{|u|}{2}$.

Afterwards, cancellations can only occur on one side of u', where u' is whatever is left of u. Note that for each letter that cancels in front of u', a copy of this letter appears in the back. Hence it does not make sense to have $|u'|$ cancellations or more – we could spare these superflous letters.

The total number of cancellations, therefore, is $< \frac{|u|-|u'|}{2} + |u'| \leq |u|$.

On the other hand, we know that after all cancellations are done, the right hand side equalls w. Hence

$$2|v| + |u| < |w| + 2|u|$$

and the claim follows. **q.e.d.**

Corollary 2.6.4. *Different generators of F_n are not conjugate.*

Exercise 2.6.5. Find an *efficient* algorithm to solve the conjugacy problem in finitely generated free groups.

Exercise 2.6.6. Modify the graphs-and-folds technique used in proving Grushko's Theorem to devise an algorithm that does the following:

> The input it takes is a finite set $\{g_1, \ldots, g_r\}$ of elements in F_n given as reduced words in the standard generators.
>
> The output is a list of free generators $\{h_1, \ldots, h_s\}$ for the subgroup $\langle g_1, \ldots, g_r \rangle$ generated by the g_i.

Exercise 2.6.7. Find an algorithm that solves the subgroup membership problem for finitely generated free groups.

Remark 2.6.8. The subgroup conjugacy problem is related to the recognition problem for groups: Suppose we had a machine that could tell us if the standard generator x_1 is has a conjugate in $\langle g_1, \ldots, g_r \rangle$, then we could run the test on all the generators and see if the normal closure of $\langle g_1, \ldots, g_r \rangle$ is all of the free group. Given this machine, we have an easy way of deciding if a finite presentation acutally presents the trivial group. This problem, however is undecidable.

Conjecture 2.6.9. *The subgroup conjugacy problem is unsolvable for non-abelian free groups.*

For free groups, there has been a lot of research about decidability questions. We list the most famous results:

Theorem 2.6.10 (Makanin [Maka82]). *There is an algorithm that, given an equation in a free group, decides whether the equation has a solution.*

Theorem 2.6.11 (Razborov [Razb84]). *There is an algorithm that, given a system of equations in a free group, decides whether it has a solution.*

In both cases, the basis for the algorithm generalizes what we did for the conjugacy problem: Find an a priori bound on the length of a minimal solution.

The ultimate theorem along those lines is a recent solution to Tarski's problem. To state the problem (and the solution) we need to understand the "elementary theory" of a free group.

Definition 2.6.12. Let F_n be a free group of rank n. Consider statements of first order logic over an alphabet containing a multiplication operation, the identity relation, infinitely many variables, and one constant symbol for any of the n free generators of F_n. We can interpret those statements over F_n in an obvious way.

The elementary theory of F_n is the set of all true statements.

The elementary theory encodes everything that can be said about F_n with "finite linguistic means", i.e., you are not allowed to use the langugage of sets or phrases like "and so forth". Obviously, it would be nice if the elementary theories of free groups were decidable. This would generalize the theorems of Makanin and Razborov to a large extent – e.g., we could deal with systems of equations and inequalities.

By means of the standard inclusions

$$F_2 \leq F_3 \leq F_4 \leq \cdots$$

we can interpret any statement over a free group on n generators as a statement about all free groups of higher rank, as well. Tarski asked whether the elementary theories of non-abelian finite rank free groups are all equal, i.e, if there is no statement that has different truth values when interpreted in different free groups.

Both problems, decidability of elementary theories and Tarski's problem, have positive solutions. The priority for these results, which turn out to be strongly related, is still unsettled. O. Kharlampovich and A. Myasnikov have their proofs spread out in [KM98a], [KM98b], [KM98c], [KM00a], [KM00b], and [KM00c]; Z. Sela has presented his account in [Se01a], [Se01b], [Se01c], [Se01d], [Se01e], and [Se01f]. Both proofs are several hundred pages each. The upshot is:

Theorem 2.6.13 (Kharlampovic-Myasnikov, Sela). *The elementary theories of all non-abelian free groups of finite rank coincide and are decidable.*

Remark 2.6.14. It is also possible to describe the set of all solutions to a system of equations (or more generally the set of all solutions to an open sentence) by means of "parametrized words". This also comes out of the heavy machinery used to solve Tarski's problem.

Part II

Arithmetic Groups

Chapter 3

$\mathrm{SL}_2(\mathbb{Z})$ and the Hyperbolic Plane

The group $\mathrm{SL}_2(\mathbb{Z})$ is the most simple arithmetic group. Many of its properties extend to arithmetic groups or even S-arithmetic groups in general. On the other hand, it has some features that are special and do not generalize to higher rank arithmetic groups.

We observed in the proof of (2.2.6):

- $\mathrm{SL}_2(\mathbb{Z})$ is linear.
- $\mathrm{SL}_2(\mathbb{Z})$ is residually finite.

Here, we will show:

[3.3.3] $\mathrm{SL}_2(\mathbb{Z})$ acts cocompactly, with finite vertex stablizers on a tree.

[3.2.7] $\mathrm{SL}_2(\mathbb{Z})$ has only finitely many conjugacy classes of finite subgroups.

[3.2.9] $\mathrm{SL}_2(\mathbb{Z})$ is virtually torsion free.

[3.3.5] $\mathrm{SL}_2(\mathbb{Z})$ is virtually free. Thus:

(a) $\mathrm{SL}_2(\mathbb{Z})$ is virtually of type F.
(b) $\mathrm{SL}_2(\mathbb{Z})$ is of type F_∞.
(c) $\mathrm{SL}_2(\mathbb{Z})$ is not amenable.

[3.4.3] The conjugacy problem in $\mathrm{SL}_2(\mathbb{Z})$ is solvable.

[3.5.6] $\mathrm{SL}_2(\mathbb{Z})$ contains finite index normal subgroups that do not contain a congruence subgroup.

[3.6.13] $\mathrm{SL}_2(\mathbb{Z})$ is SQ-universal.

We will start, however, by describing the action of $\mathrm{SL}_2(\mathbb{Z})$ on the hyperbolic plane.

3.1 The Symmetric Space of $\mathbf{SL_2}(\mathbb{R})$

The group $\mathrm{SL}_2(\mathbb{R})$ acts on the complex projective line $\mathbb{P}^1(\mathbb{C})$ in an obvious way. The complex projective line is the Riemann sphere, and since the coefficients of matrices in $\mathrm{SL}_2(\mathbb{R})$ are real, the equator of the Riemann sphere is invariant under this action. Moreover, the action does not swap the northern and southern hemispheres. Hence, there is an induced action on the northern hemisphere – the north pole is the imaginary unit i. This action is given by Möbius transformations:

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} z := \frac{az+b}{cz+d}.$$

The kernel of this action is the center of $\mathrm{SL}_2(\mathbb{R})$:

$$\{\mathbb{I}_2, -\mathbb{I}_2\}.$$

The northern hemisphere is a well known model for the hyperbolic plane $\mathbb{H}^2$.

Exercise 3.1.1. Prove that $\mathbb{H}^2$ has constant curvature -1.

Exercise 3.1.2. Show that Möbius transformations are isometries of $\mathbb{H}^2$.

Exercise 3.1.3 (extra credit). Show that any orientation preserving isometry of $\mathbb{H}^2$ is given by a Möbius transformation.

Exercise 3.1.4. Show that geodesics in $\mathbb{H}^2$ are vertical lines or half circles orthogonal to the real axis.

Definition 3.1.5. A horizontal line or a circle tangent to the real axis in $\mathbb{H}^2$ is called a horocircle.

Exercise 3.1.6. Show that the action of $\mathrm{SL}_2(\mathbb{R})$ takes horocircles to horocircles.

Observation 3.1.7. *Elements of the form* $\begin{pmatrix} 1 & b \\ 0 & 1 \end{pmatrix}$ *act as translations that preserve the imaginary part and shift the real part by b. These elements are called translations.*

Elements of the form $\begin{pmatrix} a & 0 \\ 0 & a^{-1} \end{pmatrix}$ *act act transitively on the imaginary line. They move lines in the upper half plane to parallel lines. These elements are dilations.*

Corollary 3.1.8. *The action of* $\mathrm{SL}_2(\mathbb{R})$ *on* $\mathbb{H}^2$ *is transitive.* **q.e.d.**

We determine the stabilizer of i: First, we have

$$\begin{aligned} \begin{pmatrix} a & b \\ c & d \end{pmatrix} \mathrm{i} &= \mathrm{i} \\ a\mathrm{i} + b &= -c + d\mathrm{i} \\ a = d \qquad & b = -c. \end{aligned}$$

Now, the determinant gives:

$$a^2 + b^2 = 1.$$

It follows that the stabilizer of i is precisely the group $\mathrm{SO}(2)$. This is a compact Lie group of rotations. Since these rotations act transitively on the unit tangent vectors at i, we obtain the following strengthening of (3.1.8)

Observation 3.1.9. $\mathrm{SL}_2(\mathbb{R})$ *acts transitively on the following sets:*

- *The set of isometric embeddings* $\mathbb{R} \longrightarrow \mathbb{H}^2$.
- *The unit sphere bundle of the Riemannian manifold* $\mathbb{H}^2$.

Theorem 3.1.10. *Any compact subgroup of* $\mathrm{SL}_2(\mathbb{R})$ *is conjugate to a subgroup of* $\mathrm{SO}(2)$

Proof. A compact subgroup has bounded orbits. The hyperbolic plane is negatively curved and simply connected. Hence, compact subsets have unique centers – the *center* of a bounded subset is the center of a minimal covering disk. Since the group acts by isometries (3.1.2), the compact subgroup fixes the center of any of its orbits. Thus, any compact subgroup fixes a point of the hyperbolic plane.

By transitivity of the action (3.1.8), we find an element that moves the fixed point to i. This element conjugates the compact subgroup into $\mathrm{Stab}\,\mathrm{i}$. **q.e.d.**

Corollary 3.1.11. $\mathbb{H}^2 = \mathrm{SL}_2(\mathbb{R}) \,/\, \mathrm{SO}(2)$.

Definition 3.1.12. An element $M \in \mathrm{SL}_2(\mathbb{R})$ is called

- elliptic if $|\mathrm{tr}(M)| < 2$,
- parabolic if $|\mathrm{tr}(M)| = 2$, and
- hyperbolic if $|\mathrm{tr}(M)| > 2$.

Observation 3.1.13. *The characteristic polynomial of*

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{R})$$

is

$$\begin{vmatrix} 1-xa & xb \\ xc & 1-xd \end{vmatrix} = x^2 - (a+d)\,x + 1.$$

Hence the matrix $M \in \mathrm{SL}_2(\mathbb{R})$ has two conjugate complex eigenvalues if it is elliptic. This is to say, M fixes a point in $\mathbb{H}^2$. In particular, M is conjugate to a rotation. If M is hyperbolic, it has two real eigenvalues whence it has two fixed points on $\partial(\mathbb{H}^2) = \mathbb{P}^1(\mathbb{R})$. Finally, if M is parabolic, it has one fixed point on the boundary $\partial(\mathbb{H}^2)$. **q.e.d.**

3.2 A Fundamental Domain for $\mathrm{SL}_2(\mathbb{Z})$

Definition 3.2.1. Let G act on $\mathbb{H}^2$ by isometries. A strong fundamental domain for the action is a subset $D \subseteq \mathbb{H}^2$ such that every G-orbit has precisely one point in D.

Observation 3.2.2. *Let $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{R})$. Then we have:*

$$\begin{aligned} M\frac{-d\pm 1}{c} &= \begin{pmatrix} a & b \\ c & d \end{pmatrix}\frac{-d\pm 1}{c} \\ &= \frac{-ad \pm a + bc}{-cd \pm c + dc} \\ &= \frac{-1 \pm a}{\pm c} \\ &= \frac{a \mp 1}{c}. \end{aligned}$$

This computation shows that M takes the half circle of radius $\frac{1}{|c|}$ centered at $\frac{-d}{c}$ to the half circle of the same radius centered at $\frac{a}{c}$. Since M preserves the orientation, we see that

$$M\left\{ z \in \mathbb{H}^2 \,\middle|\, \left| z - \frac{-d}{c} \right| < \frac{1}{|c|} \right\} = \left\{ z \in \mathbb{H}^2 \,\middle|\, \left| z - \frac{a}{c} \right| > \frac{1}{|c|} \right\}$$

and

$$M\left\{ z \in \mathbb{H}^2 \,\middle|\, \left| z - \frac{-d}{c} \right| > \frac{1}{|c|} \right\} = \left\{ z \in \mathbb{H}^2 \,\middle|\, \left| z - \frac{a}{c} \right| < \frac{1}{|c|} \right\}.$$

Corollary 3.2.3. *Fix* $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{R})$. *The* $\langle M \rangle$*-orbit of any point intersects*

$$\left\{ z \in \mathbb{H}^2 \;\middle|\; \left| z - \frac{-d}{c} \right| \geq \frac{1}{|c|} \text{ or } \left| z - \frac{a}{c} \right| > \frac{1}{|c|} \right\}$$

in at most one point.

Put

$$\overline{D} := \left\{ z \in \mathbb{H}^2 \;\middle|\; |z| \geq 1 \text{ and } -\frac{1}{2} \leq \Re(z) \leq \frac{1}{2} \right\}$$

and

$$\Sigma := \left\{ M \in \mathrm{SL}_2(\mathbb{Z}) \;\middle|\; M\overline{D} \cap \overline{D} \neq \emptyset \right\}.$$

We shall first determine Σ. Let $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in \mathrm{SL}_2(\mathbb{Z})$. If there are two points $x, y \in \overline{D}$ with

$$Mx = y$$

then (3.2.3) implies $c \in \{-1, 0, 1\}$.

$\underline{c = 0}$: Since $\mathrm{Det}(M) = 1$, we have $a = d = \pm 1$. It follows that M either acts as the identity or as a translation. Now $b \in \{-1, 0, 1\}$ follows, and we obtain

$$M \in \left\{ \pm \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \pm \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \pm \begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix} \right\}.$$

It is obvious that all these matrices belong to Σ.

$\underline{c = \pm 1}$: There are only three half circles of radius $1 = \frac{1}{|c|}$ centered at integer points that intersect $\overline{D}$ non-trivially. Thus (3.2.3) implies $a, d \in \{-1, 0, 1\}$. Once we pick a and d the last entry b is determined by $\mathrm{Det}(M) = 1$. Thus, we have the following candidates:

$$\pm \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \quad \pm \begin{pmatrix} -1 & 0 \\ 1 & -1 \end{pmatrix}, \quad \pm \begin{pmatrix} 1 & -2 \\ 1 & -1 \end{pmatrix}, \quad \pm \begin{pmatrix} -1 & -2 \\ 1 & 1 \end{pmatrix},$$

$$\pm \begin{pmatrix} \pm 1 & -1 \\ 1 & 0 \end{pmatrix}, \quad \pm \begin{pmatrix} 0 & -1 \\ 1 & \pm 1 \end{pmatrix}, \quad \pm \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

By inspection, one establishes that the following matrices actually belong to Σ:

$$\pm \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \pm \begin{pmatrix} -1 & 0 \\ 1 & -1 \end{pmatrix}, \pm \begin{pmatrix} \pm 1 & -1 \\ 1 & 0 \end{pmatrix}, \pm \begin{pmatrix} 0 & -1 \\ 1 & \pm 1 \end{pmatrix}, \pm \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}.$$

Hence we have

$$\Sigma = \left\{ \begin{array}{c} \pm\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \pm\begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}, \pm\begin{pmatrix} 1 & -1 \\ 0 & 1 \end{pmatrix}, \pm\begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix}, \\ \pm\begin{pmatrix} -1 & 0 \\ 1 & -1 \end{pmatrix}, \pm\begin{pmatrix} \pm 1 & -1 \\ 1 & 0 \end{pmatrix}, \pm\begin{pmatrix} 0 & -1 \\ 1 & \pm 1 \end{pmatrix}, \pm\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} \end{array} \right\}$$

Exercise 3.2.4. Show that, for any $M \in \mathrm{SL}_2(\mathbb{Z})$,

$$M\overline{D} \cap \overline{D}$$

does not contain a non-empty open subset of $\mathbb{H}^2$.

Let us define a subset D of $\overline{D}$ by excluding the right boundary and the open right half of the bottom boundary. Thus D is given by the following picture:

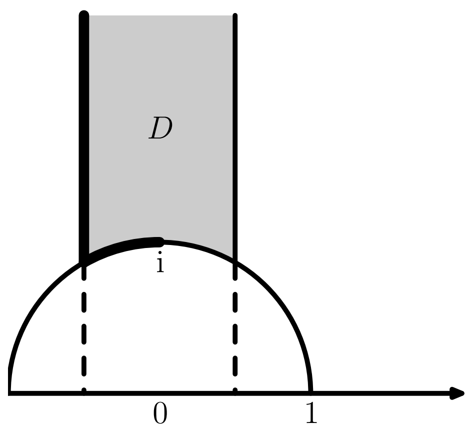

Lemma 3.2.5. $\mathbb{H}^2 = \mathrm{SL}_2(\mathbb{Z})\, D$.

Proof. Since

$$\overline{D} \subset D \cup \begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} D \cup \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} D,$$

it suffices to prove the claim for $\overline{D}$ instead of D.

We claim that the following algorithm eventually moves every $z \in \mathbb{H}^2$ into $\overline{D}$. Put $z_0 := z$ and define two sequences of points by the following rules:

- Let

  $$z_i' := z_i + n$$

  where $n \in \mathbb{Z}$ is chosen such that

  $$-\frac{1}{2} \leq \Re(z_i') < \frac{1}{2}.$$

- Put

$$z_{i+1} := \begin{cases} -\frac{1}{z_i'} & \text{for } |z_i'| < 1 \\ z_i' & \text{otherwise.} \end{cases}$$

Note that $z_s' \in \mathrm{SL}_2(\mathbb{Z})\, z_i$ and $z_{s+1} \in \mathrm{SL}_2(\mathbb{Z})\, z_i'$. Hence it suffices to prove that, eventually, $z_i \in \overline{D}$.

First observe that

$$\Im(z_i') \leq \frac{1}{2}$$

implies

$$|z_{i+1}| \geq 2\,|z_i'|$$

since $|\Re(z_i')| \leq \frac{1}{2}$. Thus, we have $\Im(z_i') = \Im(z_i) > \frac{1}{2}$ for i large enough. for such an i, let $z_i' = x + \mathrm{i}y$ and assume $|z_i'| < 1$. We claim $z_{i+1} \in \overline{D}$. This is apparent from the following picture:

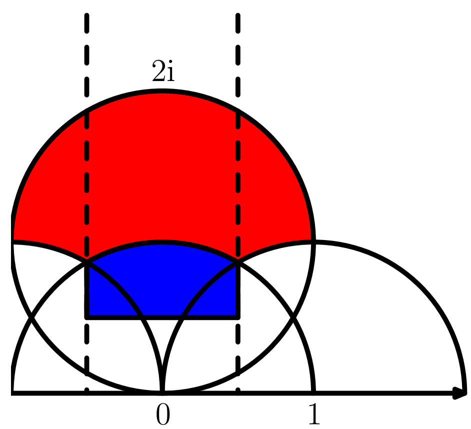

which is valid by (3.1.6). **q.e.d.**

Our discussion so far can be summarized as follows:

Proposition 3.2.6. *The collection of closed subsets $M\overline{D}$ where $M \in \mathrm{SL}_2(\mathbb{Z})$ forms an $\mathrm{SL}_2(\mathbb{Z})$-invariant tiling of $\mathbb{H}^2$ with ideal triangles.* **q.e.d.**

From this, we can derive a good deal of information about $\mathrm{SL}_2(\mathbb{Z})$.

Theorem 3.2.7. *The group $\mathrm{SL}_2(\mathbb{Z})$ has only finitely many conjugacy classes of finite subgroups.*

Proof. Let $F \leq \mathrm{SL}_2(\mathbb{Z})$ be finite. Then there is a point $x \in \mathbb{H}^2$ such that

$$Fx = x.$$

Chose $M \in \mathrm{SL}_2(\mathbb{Z})$ such that $Mx \in \overline{D}$. Then

$$MFM^{-1} \subseteq \Sigma.$$

But Σ is finite. **q.e.d.**

Exercise 3.2.8. Let G be a residually finite group that has only finitely many conjugacy classes of finite subgroups. Show that G is virtually torsion free.

Corollary 3.2.9. $\mathrm{SL}_2(\mathbb{Z})$ *is virtually torsion free.*

Theorem 3.2.10. *The group* $\mathrm{SL}_2(\mathbb{Z})$ *is finitely presented.*

Proof. Let U be a contractible open neighborhood of $\overline{D}$ contained in the union of all tiling triangles that intersect $\overline{D}$. Then (A.1.10) applies. **q.e.d.**

We can improve (3.2.6) a little. This strengthenening is, however, not needed for applications. Therefore, we omit all details of the proof.

Proposition 3.2.11. D *is a strong fundamental domain for* $\mathrm{SL}_2(\mathbb{Z})$ *in* $\mathbb{H}^2$.

Proof. By (3.2.5), half of the claim is already proved. Thus we only have to show that no two points in D are in the same $\mathrm{SL}_2(\mathbb{Z})$-orbit. This is done by inspection of the elements in the finite set Σ. **q.e.d.**

3.3 The Tree of $\mathbf{SL_2}(\mathbb{Z})$

Let us draw the tiling (3.2.6) in the unit disc model:

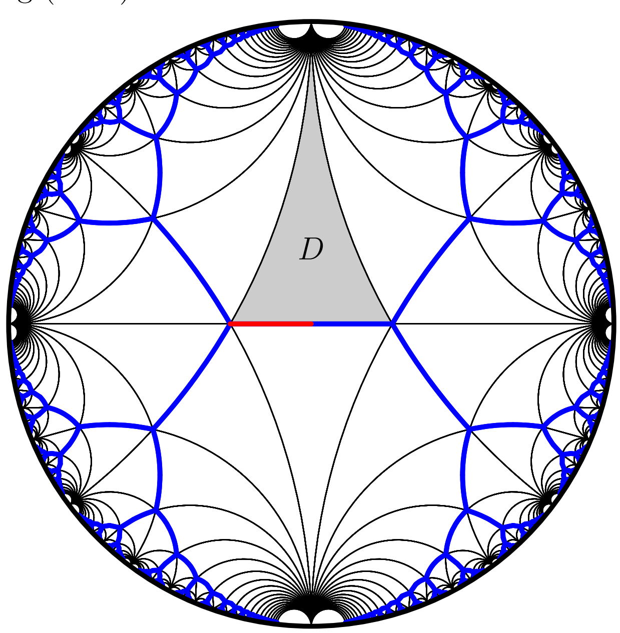

We see that the finite edges of the tiling triangles form a trivalent tree.

Proposition 3.3.1. *We may trust our eyes, the finite edges of the tiling triangles do, indeed, form a tree.*

Proof. There is no doubt that these edges form a graph. We have to argue that this graph is contractible. Thus it suffices to recognize this graph a a deformation retract of $\mathbb{H}^2$.

In the fundamental triangle $\overline{D}$, we define a retraction by moving every point x down along the geodesic from ∞, the ideal triangle vertex, through x at hyperbolic unit speed until the point hits the bottom boundary edge of $\overline{D}$. This retraction, being defined entirely in terms of the hyperbolic metric and the geometry of $\overline{D}$, obviously extends $\mathrm{SL}_2(\mathbb{Z})$-equivariantly to all of $\mathbb{H}^2$. **q.e.d.**

Definition 3.3.2. A group G acts properly discontinuously on a topological space X if, for every compact subset $C \subseteq X$, the set

$$\{g \in G \mid gC \cap C \neq \emptyset\}$$

is finite.

Corollary 3.3.3. $\mathrm{SL}_2(\mathbb{Z})$ *acts cocompactly and properly discontinously on a tree.*

Proof. Cocompactness is obvious. For the compact subset given by the fundamental edge e, we already established that the action is properly discontinuous. The general case follows: Let C be a compact subset of the tree. Then it is covered by finitely many edges $M_i e$. Thus,

$$MC \cap C \neq \emptyset$$

implies

$$MM_i e \cap M_j e \neq \emptyset$$

which in turn yields

$$M = M_j x M_i^{-1}$$

for some $x \in \Sigma$. It is apparent that only finitely many elements arise that way. **q.e.d.**

Corollary 3.3.4. $\mathrm{SL}_2(\mathbb{Z})$ *is of type* F_∞.

Proof. This follows from the action on the tree in view of (D.1.2). **q.e.d.**

Exercise 3.3.5. Show that $\mathrm{SL}_2(\mathbb{Z})$ contains a non-abelian free subgroup of finite index.

Exercise 3.3.6. Let G act on a tree T such that the following conditions are satisfied:

1. G acts transitively on the set of geometric (unoriented) edges.
2. G does not act transitively on the set of vertices.
3. T has no terminal vertices and is not isomorphic to a line.

Let e be an edge in T that connects the vertices v and w. Let G_v and G_w denote the stabilizers of these vertices and let G_e denote the stabilizer of the edge e. Then $G = G_v *_{G_e} G_w$.

Infer that

$$\mathrm{SL}_2(\mathbb{Z}) = C_6 *_{C_2} C_4.$$

3.4 The Conjugacy Problem

Definition 3.4.1. A group G is combable with respect to the generating set Σ, if there is a constant C and distinguished paths from 1_G to each vertex $v \in \Gamma_\Sigma^G$ that have the C-fellow traveler property. That is, whenever we have two distinguished paths p and q starting at $1_G \in \Gamma_\Sigma^G$ whose endpoints have distance ≤ 1, the following inequality holds along the paths:

$$d(p_t, q_t) \leq C. \tag{3.1}$$

Here the paths are traversed with unit speed. If all the combing paths can be chosen to be geodesics, the group G is called geodesically combable. If any two geodesic paths whose endpoints have distance ≤ 1 satisfy the inequality (3.1), we say that G has the C-fellow traveler property.

Exercise 3.4.2. Prove that for some finite generating set and some constant C, geodesic paths in the Cayley graph of $\mathrm{SL}_2(\mathbb{Z})$ have the C-fellow traveler property. Hint: Use the action on the tree.

Proposition 3.4.3. *Geodesically combable groups have solvable conjugacy problem.*

Proof. Let u and v be two words. We show that we can restrict the search for a conjugating element w with $wu = vw$ to words of a length bounded by a constant that depends only on $|u|$ and $|v|$.

So let w be a shortest conjugating element – if there is one. We have to give an upper bound for $|w|$. This is done in two pictures:

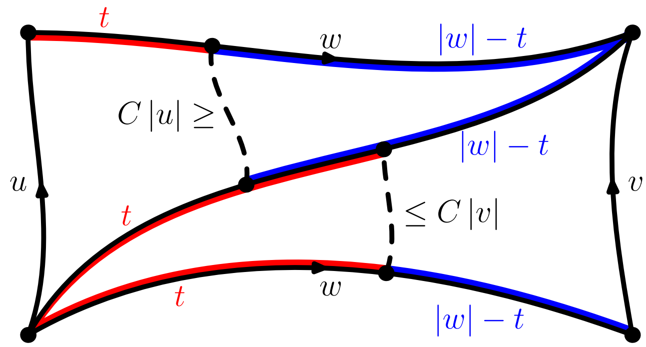

This picture proves that the upper and lower geodesic are fellow travelers for a constant essentially proportional to $|u| + |v|$. To see this, we just observe that the two-colored segment that connects the dashed paths is short by the triangle inequality: the diagonal has length between $|w| - |v|$ and $|w| + |v|$.

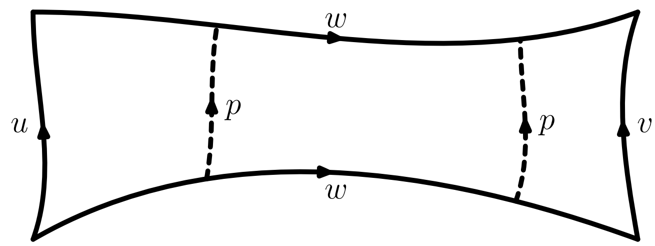

This picture shall remind you of the pigeon hole principle. If w is too long, then there will be two points such that the short vertical geodesic connections read identical group elements. We can then cut out the middle rectangle and shorten the conjugation rectangle. **q.e.d.**

3.5 Finite Quotients and Congruence Subgroups

We note an immediate consequence of (3.3.6).

Observation 3.5.1. *$C_3 * C_2$ is a quotient of $\mathrm{SL}_2(\mathbb{Z})$. In fact, it follows from (3.3.6) that $C_3 * C_2 = \mathrm{PSL}_2(\mathbb{Z})$.* **q.e.d.**

It follows that every group that can be generated by two elements of orders two and three is a quotient of $\mathrm{SL}_2(\mathbb{Z})$.

Exercise 3.5.2. Show that the alternating group A_{11} is a quotient of $\mathrm{SL}_2(\mathbb{Z})$.

Remark 3.5.3. More is true: For $n \geq 9$, the alternating group A_n and the symmetric group S_n are quotients of $\mathrm{PSL}_2(\mathbb{Z})$. See [Magn74, page 119] for references.

There is a class of obvious finite quotients of $\mathrm{SL}_2(\mathbb{Z})$ and $\mathrm{PSL}_2(\mathbb{Z})$, namely the groups $\mathrm{PSL}_2(\mathbb{Z}^m)$.

Definition 3.5.4. A normal subgroup of $\mathrm{SL}_2(\mathbb{Z})$ is a <u>congruence subgroup</u> if it is the kernel of the canonical homomorphism

$$\mathrm{SL}_2(\mathbb{Z}) \to \mathrm{PSL}_2(\mathbb{Z}^m)$$

for some m.

Exercise 3.5.5. Prove that for every $m \in \mathbb{N}$, the homomorphism

$$\mathrm{SL}_2(\mathbb{Z}) \to \mathrm{PSL}_2(\mathbb{Z}^m)$$

is onto.

Proposition 3.5.6. *Not all finite index normal subgroups of* $\mathrm{SL}_2(\mathbb{Z})$ *contain a congruence subgroup.*

This result is remarkable because it fails in higher ranks: Every finite index normal subgroup in $\mathrm{SL}_n(\mathbb{Z})$ contains a congruence subgroup provided $n \geq 3$. (4.0.17)

Proof. In view of (3.5.2) and (3.5.5), it suffices to show that A_{11} is not a quotient of $\mathrm{PSL}_2(\mathbb{Z}^m)$. Since A_{11} is simple, we only have to prove that it does not occur as a factor in a decomposition series for $\mathrm{PSL}_2(\mathbb{Z}^m)$. We will, in fact, prove the slightly stronger statement, that A_{11} does not occur as a factor in the decomposition series of $\mathrm{SL}_2(\mathbb{Z}^m)$.

As a first step, we prove

$$\mathrm{SL}_2(\mathbb{Z}^m) = \bigtimes_q \mathrm{SL}_2(\mathbb{Z}^q)$$

where q ranges over the prime powers in a decomposition of m. Let $\begin{pmatrix} a & b \\ c & d \end{pmatrix}$ be an integer matrix of determinant 1 representing a given element $N \in \mathrm{SL}_2(\mathbb{Z}^m)$. Fix an enumeration $q_1, \ldots, q_r$ of these factors. By the Chinese remainder theorem, there are integers a_i, b_i, c_i, and d_i satisfying the congruences

$$a \cong a_i b \cong b_i c \cong c_i d \cong d_i \mod q_i$$

and

$$M_j := \begin{pmatrix} a_j & b_j \\ c_j & d_j \end{pmatrix} \cong \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \mod q_i \text{ for } i \neq j.$$

Moreover, these numbers are unique $\mod m$. Observe that the product M of these matrices is congruent to $M_i \mod q_i$ regardless of the ordering of the factors. It represents the element N in $\mathrm{SL}_2(\mathbb{Z}^m)$ that we picked in the first place. In addition, the matrices M_i define elements in $\mathrm{SL}_2(\mathbb{Z}^m)$ that commute pairwise – to see this, just verify that the commutators $M_iM_j - M_jM_i$ vanish $\mod q_k$ for any k. Finally, the matrices M_i descend to matrices in $\mathrm{SL}_2(\mathbb{Z}^{q_i})$. It follows that

$$\mathrm{SL}_2(\mathbb{Z}^m) = \bigtimes_i \mathrm{SL}_2(\mathbb{Z}^{q_i}).$$

Thus we reduced our task to proving that A_{11} does not occur as a factor in the decomposition series of $\mathrm{SL}_2(\mathbb{Z}^q)$ where q is a prime power $q = p^{k+1}$.

Check that $\mathrm{PSL}_2(\mathbb{Z}^p)$ has order $\frac{p(p-1)(p+1)}{2}$. It follows that p is the largest prime factor of this order. Thus A_{11} is not isomorphic to any of the groups $\mathrm{PSL}_2(\mathbb{Z}^p)$ because $p = 11$ would be the only chance, and here orders do not match.

However the following exercise (3.5.7) implies by an easy induction that the only non-abelian factors of a decomposition series of $\mathrm{SL}_2(\mathbb{Z}^q)$ are isomorphic to $\mathrm{PSL}_2(\mathbb{Z}^p)$.

q.e.d.

Exercise 3.5.7. Fix a prime p and a natural number k. Show that

$$N_{p,k} := \left\{ \begin{pmatrix} 1+xp^k & yp^k \\ zp^k & 1-xp^k \end{pmatrix} \mod p^{k+1} \;\middle|\; 0 \le x,y,z < p \right\} \le \mathrm{SL}_2\left(\mathbb{Z}^{p^{k+1}}\right)$$

is a normal abelian p-subgroup of rank 3 in $\mathrm{SL}_2\left(\mathbb{Z}^{p^{k+1}}\right)$ such that

$$\mathrm{SL}_2\left(\mathbb{Z}^{p^{k+1}}\right)/N_{p,k} \cong \mathrm{SL}_2\left(\mathbb{Z}^{p^k}\right).$$

3.6 Small Cancellation Theory for Free Products

3.6.1 Van Kampen Diagrams for Presentations

Definition 3.6.1. A disc map / spherical map is a graph, embedded in the plane / 2-sphere such that the bounded complementary regions are discs. A map is proper if each vertex has degree ≥ 3.

Let D be a proper map and denote by F_i the number of regions bounded by i edges. The number of vertices is V the number of edges is E, and $F = \sum F_i$ is the number of two cells. Obviuosly,

$$\begin{aligned} 2E &= \sum iF_i \\ 3V &\le 2E \end{aligned}$$

Thus, we can estimate the Euler characteristik by

$$6\,(V - E + F) \leq -2E + 6F = \sum (6 - i)\, F_i$$

Since the Euler characteristic of a disc is 1, we have proved the main lemma of small cancellation theory:

Lemma 3.6.2. *There are no proper disc maps all of whose two cells have at least six edges.* **q.e.d.**

We will use this idea, eventually, to prove that certain maps are injective. However, the key points are probably more easy to comprehend in a model case. Consider the presentation

$$\mathcal{P} := \left\langle x_1, y_1, \ldots, x_g, y_g \,\middle|\, x_1 y_1 x_1^{-1} y_1^{-1} \cdots x_g y_g x_g^{-1} y_g^{-1} \right\rangle$$

and let $\Gamma_{\mathcal{P}}$ be the Cayley-2-complex associated with $\mathcal{P}$. We shall use spherical maps to prove

Proposition 3.6.3. *If $g \geq 2$, then $\Gamma_{\mathcal{P}}$ is contractible.*

Proof. Being a Cayley complex, $\Gamma_{\mathcal{P}}$ is simpliy connected and of dimension 2. Thus, by Hurewicz theorem, it suffices to prove that $\pi_2(\Gamma_{\mathcal{P}})$ is trivial.

Every element of $\pi_2(\Gamma_{\mathcal{P}})$ can be represented by a *van Kampen diagram D*, i.e., a map whose oriented edges are labelled by generators such that each two cell boundary reads a cyclic permutation of the relation or its inverse. We can assume that spherical diagrams are *reduced*, i.e., no edge connects mirror images of two cells.

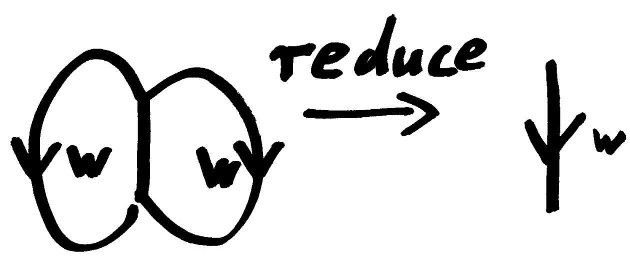

The key observation is that the relation has the following strange property: In any reduced diagram, two cells cannoc have long commong segments along their boundaries. The combinatorics of this particular rules out that a path of two edges is common to two relator discs. Thus, the diagram D is actually a proper map and every cell has at least 8 neighbors. Then the Euler characteristic estimate implies that D is not homeomorphic to a sphere. Hence all spheres in $\Gamma_{\mathcal{P}}$ are 0-homotopic. **q.e.d.**

3.6.2 Van Kampen Diagrams for Free Products

Let us consider the free product

$$G = G_1 * G_2 * \cdots * G_r$$

and a set of cyclicly reduced words R closed under taking inverses and cyclic permutation of words. Let N be the normal subgroup of G spanned by R.

Definition 3.6.4. A van Kampen diagram over R is a disc map D whose oriented edges are labeled by elements of

$$\biguplus_i G_i - \{1\}$$

such that for each two cell f one of the following holds:

1. The boundary $\partial(f)$ reads a word that representes 1 and all of whose letters belong to one of the factors G_i. In this case, we call f an auxiliary cell.
2. The boundary $\partial(f)$ evaluates to an element of R. Here, we say that an edge path evaluates to a word w, if the word that the edge path ready reduces to w.

Two diagrams whose boundaries evaluate to the same words are equivalent.

Observation 3.6.5. *The boundary word of a van Kampen diagram representes the trivial element in G/N. Conversely, for any element in N, there is a van Kampen diagram whose boundary represents that element.* **q.e.d.**

Definition 3.6.6. Every interior edge e in a van Kampen diagram has two neighboring two cells which provide two alternative paths equivalent to e. If these two paths evaluate to identical words, we say that the edge is unreduced. A diagram without unreduced edges is reduced.

Lemma 3.6.7. *For every van Kampen diagram, there is an equivalent reduced van Kampen diagram.*

Proof. Let e be an unreduced edge. The easy case is that the two neighboring cells are different. Then we reduce:

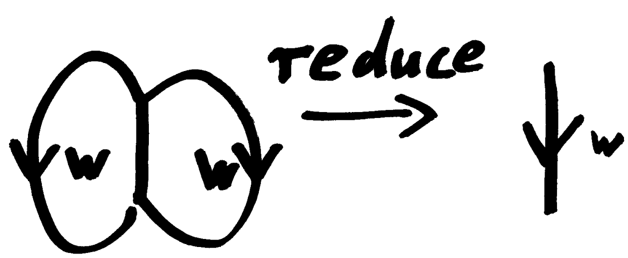

If the unreduced edge arises from a self-idenitfication along the boundary of a two cells, we can also reduce:

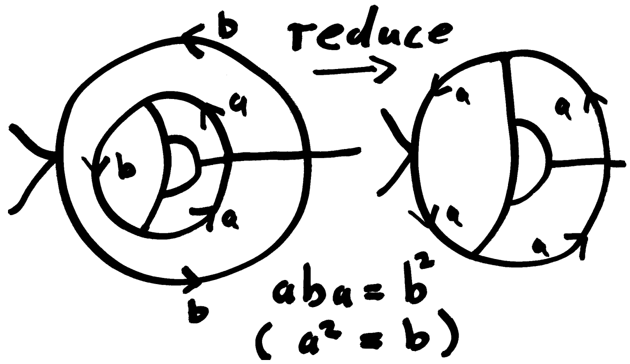

These processes decrease the number of two cells in a diagram and, therefore, must stop eventually. **q.e.d.**

Lemma 3.6.8. *If at least one of the canonical maps $G_i \to G/N$ is not injective, then there is a reduced diagram without auxiliary cells whose boundary represents a non-trivial element in one of the factors G_i.*

Proof. By hypothesis, there are non-trivial elements in some factors that belong to N. For these elements, there are van Kampen diagrams. Among those diagrams, let D be one with a minimal number of two cells. Then D is obviously reduced. Thus, we have to prove that it does not contain axiliary cells.

To see this, assume f was an auxiliary cell. We will reduce the number of boundary edges, thereby eliminating this cell eventually.

First let us show that all auxiliary cells are embedded. So suppose not.

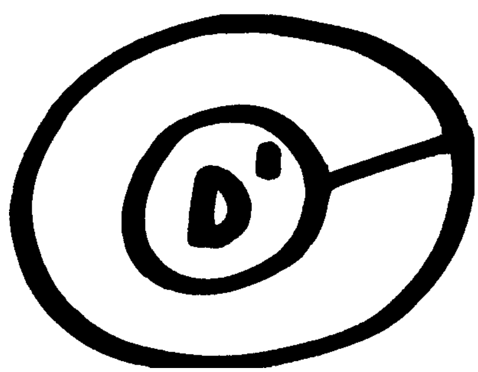

If the outer boundary evaluates to 1, so does the inner boundary. Then D' would be a smaller witness. If the outer boundary does not evaluate to 1, we delete the interior and create a new auxiliary cell.

So from now on, we assume that auxiliary cells are embedded. First consider the case that xx^{-1} occurs along the boundary. Then this picture shows how to shrink the f.

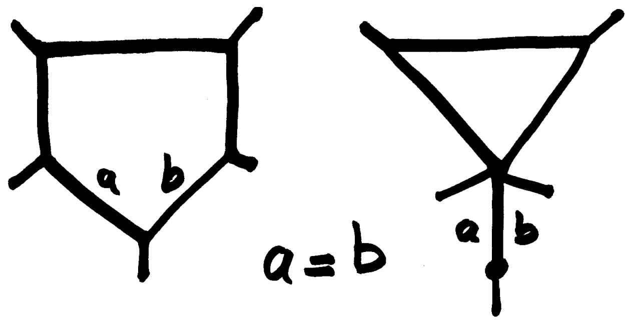

Now we deal with the remaining case:

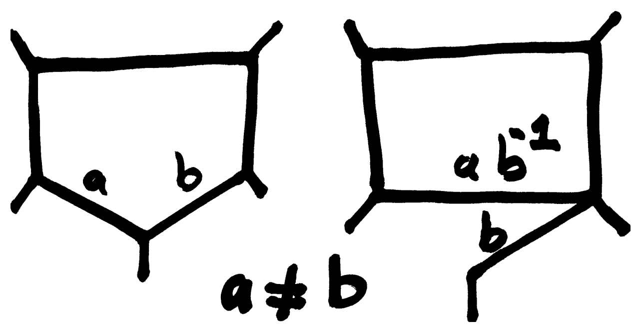

Since the two cell will eventually disappear, we have a contradiction to the minimality of D. **q.e.d.**

We give a slightly simplified version of small cancellation theory. For a full treatment see [LSch77, Chapter V]

Definition 3.6.9. A piece is a word that occurs as an initial segment of two elements of R.

Lemma 3.6.10. *Suppose for every word $w \in R$ and every piece b in w, we have*

$$6\,(|b| + 2) \leq |w|$$

then all maps $G_i \to G/N$ are injective.

Proof. Suppose not and let D be a minimal witness for non-injectivity. In particular, D is reduced and contains non auxiliary cells. Eliminate all redundant vertices – vertices of degree two that connect edges labelled by letters from the same factor G_i.

The van Kampen diagram carries a proper map M whose edges correspond to edge paths in D. Observe that these edge paths are pieces after chopping of the first and last edge if necessary. Thus M consists entirely of cells whose boundaries have length ≥ 6. This is a contradiction to (3.6.2) **q.e.d.**

Here is the application, we are headed for:

Definition 3.6.11. A group G is SQ-universal if every countable group embedds into a quotient of G.

Lemma 3.6.12. *The group $G = G_3 * G_2$ with $|G_3| \geq 3$ and $|G_2| \geq 2$ is SQ-universal.*

Proof. We slightly modify the proof in [LSch77, Theorem 10.3]. Let $H = \{h_0, h_1, h_2, \ldots\}$ be a countable group and fix two non-trivial distinc element $x_1, x_2 \in C_3$ and the non-trivial element $y \in C_2$. Define N to be the normal subgroup in $H * G$ spanned by the relations

$$h_i = (xy)^{200i+1} x_2 y (xy)^{200i+2} x_2 y \cdots (xy)^{200(i+1)} x_2 y.$$

It is easy to check that pieces for this presentation satisfy the hypothesis of (3.6.10) **q.e.d.**

Corollary 3.6.13. $\mathrm{SL}_2(\mathbb{Z})$ *is SQ-universal.* **q.e.d.**

Exercise 3.6.14. Show that SQ-universal groups have uncountably many non-isomorphic quotients.

Corollary 3.6.15. $\mathrm{SL}_2(\mathbb{Z})$ *has uncountably many non-isomorphic quotients. In particular, there are uncountably many non-isomorphic groups that are generated by two elements.*

Remark 3.6.16. More is known. See [BBW79] and the references therein.

- $\mathrm{SL}_2(\mathbb{Z})$ has uncountably many non-isomorphic simple quotients.
- $\mathrm{SL}_2(\mathbb{Z})$ has uncountably many non-isomorphic solvable quotients.
- $\mathrm{SL}_2(\mathbb{Z})$ has uncountably many non-isomorphic residually alternating quotients.

Chapter 4

$\mathbf{SL}_n(\mathbb{Z})$

The main series of examples coming from semi-simple linear algebraic groups are the groups $\mathrm{SL}_n(\mathbb{Z})$. Moreover by definition, any arithmetic subgroup embeds into $\mathrm{SL}_n(\mathbb{Z})$ for n large enough.

Proposition 4.0.17. *Every subgroup of finite index in* $\mathrm{SL}_n(\mathbb{Z})$ *contains a congruence subgroup.*

Chapter 5

General Arithmetic Groups

Arithmetic groups are subgroups of linear algebraic groups. So we have to start with some remarks on these.

Definition 5.0.18. A linear algebraic group defined over a field k is a set of polynomials that cut out a subgroup of $\mathrm{SL}_n(K)$ for any superfield $K \supseteq k$. We also call these things k-groups.

Remark 5.0.19. You might wonder why definition (5.0.18) requires the set of polynomials to define a subgroup for every superfield of k. The reason is that for finite fields a certain set of polynomials might happen accidentally to define a subgroup of $\mathrm{SL}_n(k)$. What we really want, however, is that the polynomials are kind of a group like object: It should be provable by algebraic manipulations alone that the set of solutions to the equations is closed with respect to taking quotients. In fact, it suffices to require that the subgroup in $\mathrm{SL}_n\left(\overline{k}\right)$ is a group.

Exercise 5.0.20. Show that the polynomial $(a_{21} + 1)^3 = 1$ defines a subgroup of $\mathrm{SL}_2(\mathbb{F}_2)$, but it does not define a subgroup of $\mathrm{SL}_2(\mathbb{F}_4)$.

Remark 5.0.21. You also might wonder why we require the determinant to be 1 and do not leave the freedom to have this equation, which is a polynomial equation anyway, among our defining equations or not. Here, the reason is that we want the coefficients of the inverse of an element g to be polynomials in the coefficients of that matrix g.

Again, it turns out that this does not result in a loss of generality because $\mathrm{GL}_n(k)$ embeds into $\mathrm{SL}_{n+1}(k)$.

Remark 5.0.22. This is the more down to earth definition of linear algebraic groups. In a more scary way, we could define them as group objects in the category of affine

k-varieties without referring to a particular representation in SL. It turns out that we do not really loose generality that way since affine algebraic groups always have faithful linear representations and thus are isomorphic to linear algebraic groups.

Example 5.0.23. Let $\mathfrak{G}$ be a nilpotent $\mathbb{Q}$-Lie algebra. Then the Campbell-Baker-Hausdorff formula

$$\mathfrak{g} \bullet \mathfrak{h} := \mathfrak{g} + \mathfrak{h} + \frac{1}{2}[\mathfrak{g}, \mathfrak{h}] + \frac{1}{12}([\mathfrak{g}, [\mathfrak{g}, \mathfrak{h}]] + [\mathfrak{h}, [\mathfrak{h}, \mathfrak{g}]]) + \cdots$$

defines a group law on the set $\mathfrak{G}$. Recall that the Campbell-Baker-Hausdorff series is made to satisfy

$$\mathrm{e}^{\mathfrak{g} \bullet \mathfrak{h}} = \mathrm{e}^{\mathfrak{g}} \mathrm{e}^{\mathfrak{h}}$$

for non-commuting $\mathfrak{g}$ and $\mathfrak{h}$.

Since $\mathfrak{G}$ is nilpotent, this is a polynomial group law. Observe that $-\mathfrak{g}$ is the inverse of $\mathfrak{g}$ with respect to this weird multiplication. Thus, $\mathfrak{G}$ is an affine algebraic group with multiplication $\bullet$ whose underlying variety is just an affine space: the vector space underlying $\mathfrak{G}$.

We turn the class of linear algebraic groups into a category by saying what morphisms ought to be.

Definition 5.0.24. Let G and H be two linear algebraic groups defined over k. A k-homomorphism from G to H is a polynomial map with coefficients taken from k that defines a group homomorphism $G^K \to H^K$ for any superfield K.

Example 5.0.25. The group $\mathrm{SL}_2()$ defined over $\mathbb{Q}$ is isomorphic to the subgroup of $\mathrm{SL}_3()$ defined by

$$\begin{pmatrix} \star & \star & 0 \\ \star & \star & 0 \\ 0 & 0 & 1 \end{pmatrix} = MXM^{-1}$$

where $M \in \mathrm{GL}_3(\mathbb{Q})$ is fixed. The isomorphism is given by

$$\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix} \mapsto M^{-1} \begin{pmatrix} a_{11} & a_{12} & 0 \\ a_{21} & a_{22} & 0 \\ 0 & 0 & 1 \end{pmatrix} M.$$

Let us consider $k = \mathbb{Q}$ and $K = \mathbb{R}$. In this case, a linear algebraic group defines a subset of $\mathrm{SL}_n(\mathbb{R})$ by polynomial equations. Hence the set of solutions is a differentiable manifolds, at least at the non-critical points, which form a dense subset. However, since the group acts transitively on its set of points, actually all points look alike and are non-critical. Hence a linear algebraic group defined over $\mathbb{Q}$ is always a Lie group.

Remark 5.0.26. Although it might not be obvious, the same holds true in general: A linear algebraic group defined over any field is always a smooth variety whose irreducible components and connected components coincide and are finite in number. In fact, this not only true but it is true for the same reason as in the case $k = \mathbb{Q}$ provided the reason is formulated carefully.

Definition 5.0.27. Let G be a linear algebraic group over k. A Borel-subgroup $B \leq G$ is a maximal connected solvable subgroup.

The radical $\mathrm{Rad}(G)$ of G is defined to be the connected component of the identity of the intersection of all Borel-subgroups of G:

$$\mathrm{Rad}(G) := \left(\bigcap_{B \in \mathcal{B}} B \right)^0 .$$

The group G is called semi-simple if $\mathrm{Rad}(G) = 1$.

Remark 5.0.28. The radical of G is a connected solvable normal subgroup of G. The quotient $G/\,\mathrm{Rad}(G)$ is semi-simple. Hence the study of linear algebraic groups naturally centers around the study of connected solvable groups on the one hand and semi-simple groups on the other hand.

Example 5.0.29. Groups of upper triangular matrices are solvable. The group $\mathrm{SL}_n()$ is semi-simple.

Definition 5.0.30. Let G be a linear algebraic group defined over $\mathbb{Q}$. We denote by $G^{\mathbb{Z}}$ the subgroup of $G^{\mathbb{Q}}$ that consists of matrices with integer coefficients only.

Definition 5.0.31. Two subgroups of a common supergroup are called commensurable if their intersection has finite index in both of them.

Exercise 5.0.32. Let $\varphi : G \to H$ be a $\mathbb{Q}$-isomorphism of $\mathbb{Q}$-groups. Show that $\varphi\big(G^{\mathbb{Z}}\big)$ and $H^{\mathbb{Z}}$ are commensurable.

Exercise 5.0.33. Let $\varphi : G \to H$ be a $\mathbb{Q}$-epimorphism of $\mathbb{Q}$-groups. Show that $\varphi\big(G^{\mathbb{Z}}\big) \cap H^{\mathbb{Z}}$ has finite index in $\varphi\big(G^{\mathbb{Z}}\big)$.

Exercise 5.0.34. Let $\varphi : G \to H$ be a $\mathbb{Q}$-morphism of $\mathbb{Q}$-groups. Show that $\varphi^{-1}\left(H^{\mathbb{Z}}\right) \cap G^{\mathbb{Z}}$ has finite index in $G^{\mathbb{Z}}$.

Definition 5.0.35. Let G be a linear algebraic groups over $\mathbb{Q}$. An arithmetic subgroup of G is a subgroup of $G^{\mathbb{Q}}$ commensurable with $G^{\mathbb{Z}}$.

An abstract group is called arithmetic if it is isomorphic to an arithmetic subgroup of some $\mathbb{Q}$-group.

Exercise 5.0.36. Let k be a finite extension of $\mathbb{Q}$, let $\mathcal{O} \subseteq k$ be the subring of algebraic integers in k, and let G be a k-group. Prove that $G^{\mathcal{O}}$ is an arithmetic group. Hint: Represent k as a matrix algebra over $\mathbb{Q}$ and use this to find a rational representation of $G^{\mathcal{O}}$.

Example 5.0.37. The group of diagonal 2×2-matrices with determinant 1

$$\left\{ \begin{pmatrix} a & 0 \\ 0 & b \end{pmatrix} \middle| \; ab = 1 \right\}$$

is isomorphic to the multiplicative group. Hence its arithmetic subgroup is $\{\pm 1\}$ which is the group with two elements. In particular, the trivial group is arithmetic.

Exercise 5.0.38. Show that any finite group is arithmetic.

Example 5.0.39. The infinite cyclic group is arithmetic. In fact, it is the arithmetic subgroup of the $\mathbb{Q}$-group

$$\begin{pmatrix} 1 & \star \\ 0 & 1 \end{pmatrix}.$$

Fact 5.0.40 (Mal'cev). *Any torsion free finitely generated nilpotent group is arithmetic.*

Fact 5.0.41 (from rational homotopy theory). *The mapping class group of any simply connected finite CW-complex is arithmetic. Conversely, any torsion free arithmetic group arises this way.*

5.1 Preliminary Observations

Arithmetic groups are groups of integer matrices with determinant 1. This already implies some properties:

Observation 5.1.1. *Every finitely generated subgroup of an arithmetic group has a solvable word problem.*

Observation 5.1.2. $\mathrm{SL}_n(\mathbb{Z})$ *is residually finite.* **q.e.d.**

Corollary 5.1.3. *Arithmetic subgroups are residually finite.* **q.e.d.**

Exercise 5.1.4. Prove that $\mathrm{SL}_n(\mathbb{Z})$ is generated by elementary matrices, i.e., matrices that have 1s in the diagonal and precisely one additional 1 in an off-diagonal slot.

Exercise 5.1.5 (extra credit). Prove that $\mathrm{SL}_n(\mathbb{Z})$ is generated by two elements for $n \geq 5$. Remark: The statement holds for $n \geq 2$. However, a proof of the more general statement distinguishes between n even and n odd.

Exercise 5.1.6 (Minkowski (1887)). Show that the kernel of the map

$$\mathrm{SL}_n(\mathbb{Z}) \to \mathrm{SL}_n(\mathbb{Z}^p)$$

is torsion free for any odd prime p. Hint: If M is a torsion element of $\mathrm{SL}_n(\mathbb{Z})$ then the roots of its characteristic polynomial are roots of unity. This tells you something about the way it factors over $\mathbb{Z}$.

Corollary 5.1.7. *Arithmetic groups are virtually torsion free.* **q.e.d.**

Part III

Mapping Class Groups

Chapter 6

Out(F_n) and Aut(F_n)

See [Vogt00].

6.1 Topological Representatives for Automorphisms

Definition 6.1.1. Let X be a topological space with base point P. A self-homotopy equivalence is a base point preserving map $f : X \to X$ such that there is a base point preserving homotopy inverse, i.e., a base point preserving map $h : X \to X$ such that

$$f \circ h \sim \mathrm{id}_X \text{ rel } P$$

and

$$h \circ f \sim \mathrm{id}_X \text{ rel } P.$$

The mapping class group of (X, P) is the group

$$M(X, P) := \{[f]_P \mid f : X \to X \text{ is a homotopy equivalence rel. } P\}.$$

This is the group of self-homotopy equivalences modulo homotopy relative to the base point.

Observation 6.1.2. *The mappging class group is a group.* **q.e.d.**

Observation 6.1.3. *The map*

$$\begin{array}{rcl} \nu : M(X, P) & \to & \mathrm{Aut}(\pi_1(X, P)) \\ [f] & \mapsto & \alpha_{[f]} : [\gamma] \mapsto [f \circ \gamma] \end{array}$$

is a group homomorphism.

If X is the rose R_n on n petals, this homomorphism has an inverse given as follows: An automorphism of $\pi_1(R_n, P)$ assigns to each petal a loop in R_n. We can think of this loop as a map from its petal to R_n. Since the base point is preserved, the maps we obtained for the individual petals agree at the base point. This way, we defined a map $R_n \to R_n$. **q.e.d.**

Corollary 6.1.4. $\mathrm{Aut}(F_n) = M(R_n, P)$. **q.e.d.**

6.1.1 Stallings Folds

Let Γ and Δ be two graphs. A fold is a map

$$f_{\vec{e},\vec{e}_\bullet} : \Gamma \to \Delta$$

that identifies two oriented edges $\vec{e}$ and $\vec{e}_\bullet$ that have the same initial vertex. A fold is called singular if the two edges also share their terminal vertices, it is called non-singular otherwise. Here are some examples:

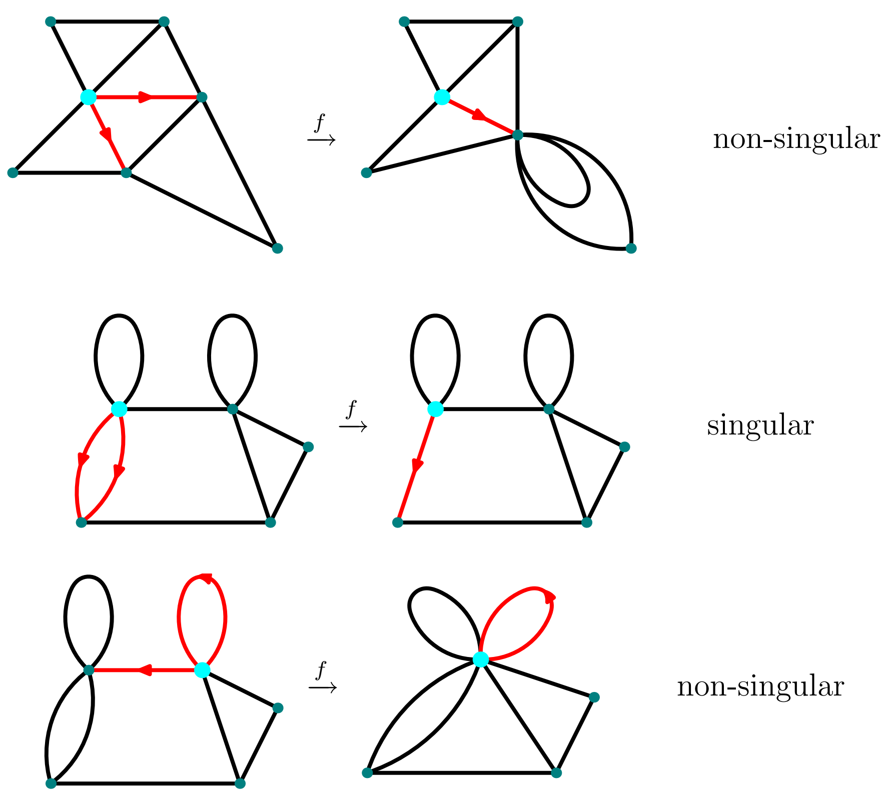

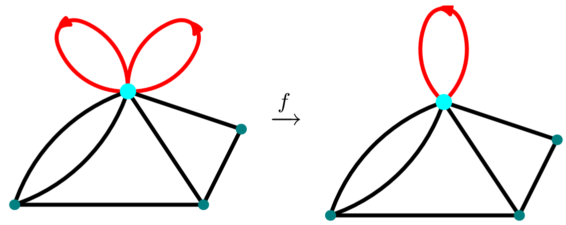

Observation 6.1.5. *A non-singular fold is a homotopy equivalence. A singular fold induces a non-injective map in homotopy. Its kernel is the normal subgroup generated by the loop* $\vec{e}\mathrm{op}(\vec{e}_\bullet)$. **q.e.d.**

Observation 6.1.6. *Let* $\varphi : \Gamma \to \Delta$ *be a graph morphism. If* φ *is not locally injective, then there is a vertex* v *from which two edges* $\vec{e}$ *and* $\vec{e}_\bullet$ *issue that are identified by the map* φ*. Thus,* φ *factors through the fold* $f_{\vec{e},\vec{e}_\bullet}$*:*

$$\varphi = \varphi' f_{\vec{e},\vec{e}_\bullet}.$$ **q.e.d.**

Proposition 6.1.7 (Stallings). *Let* Γ *be a finite graph and let* $\varphi : \Gamma \to \Delta$ *be a graph morphism. Then there is a finite sequence of folds*

$$\Gamma = \Gamma_0 \xrightarrow{f_1} \Gamma_1 \xrightarrow{f_2} \Gamma_2 \xrightarrow{f_3} \cdots \xrightarrow{f_r} \Gamma_r$$

and a graph morphism

$$\psi : \Gamma_r \to \Delta$$

such that

$$\varphi = \psi \circ f_r \circ f_{r-1} \circ \cdots \circ f_1$$

and such that ψ *is locally injective.*

Proof. Since every fold decreases the number of edges, every sequence of folds must terminate. So you try to write $\varphi = \varphi_1 f_1$. If this succeds, you try the same on φ_1. Continue until you do not find a way of factoring through a fold. We have observed in (6.1.6) that in this case the map is locally injective. **q.e.d.**

Let $\alpha : F_n \to F_n$ be an automorphism. We realize α as a graph morphism: Subdivide R_n such that petal l_i has as many segments as needed to write the word representing $\alpha(x_i)$. Call the subdivided rose $\overline{R}_n$. These words then define a map

$$\varphi : \overline{R}_n \to R_n.$$

Let us factor out a maximal sequence of folds:

$$\varphi = \psi \circ f_r \circ f_{r-1} \circ \cdots \circ f_1$$

such that ψ is locally injective. By (6.1.5), all the folds are non-singular, for otherwise we would not induce an isomorphism of fundamental groups.

Observation 6.1.8. *A locally injective map takes non-backtraking paths to non-backtracking paths. Thus, since φ is onto in π_1, the map $\psi : \Gamma_r \to R_n$ is an isomorphism of graphs: Look at the vertex in R_n and consider which paths in Γ_r give rise to the simple loops based at the vertex. It is immediate that for each such simple loop in R_n there has to be a corresponding loop in Γ_r based at its base point. Then, however, local injectivity rules out the existence of any other edges.* **q.e.d.**

Corollary 6.1.9. *Any assignment $x_i \mapsto w_i$ of words to generators determines a homomorphism $F_n \to F_n$. This homomorphism is an automorphism if and only if the topological representative, realized as a graph morphism $\varphi : \overline{R}_n \to R_n$, decomposes as*

$$\varphi = \psi \circ f_r \circ f_{r-1} \circ \cdots \circ f_1$$

such that all folds are non-singular and $\psi : \Gamma_n \to R_n$ is an isomorphism of graphs. This criterion can be checked algorithmically. **q.e.d.**

Corollary 6.1.10. *Every generating set of F_n that consists of precisely n elements is a free generating set.*

Proof. We only needed surjectivity to argue that ψ is an isomorphism of graphs. Since a non-singular fold will ensure non-surjectivity (check this in homology, if you consider it non-obvious), we infer from surjectivity alone that all folds are non-singular and the final locally injective graphmorphism is an isomorphism. Thus, every surjection $F_n \twoheadrightarrow F_n$ is an isomorphism. Compare also Grushko's Theorem (2.4.34), which is also proved using Stallings folds. **q.e.d.**

Example 6.1.11. Let us consider $F_3 = \langle \mathbf{b}, \mathbf{g}, \mathbf{r} \rangle$.

1. The assignment

$$\begin{array}{rcl} \mathbf{b} & \mapsto & \mathbf{b} \\ \mathbf{g} & \mapsto & \mathbf{brgr}^{-1} \\ \mathbf{r} & \mapsto & \mathbf{rgr}^{-1} \end{array}$$

   has the following topological representative and crucial stages in the folding sequence:

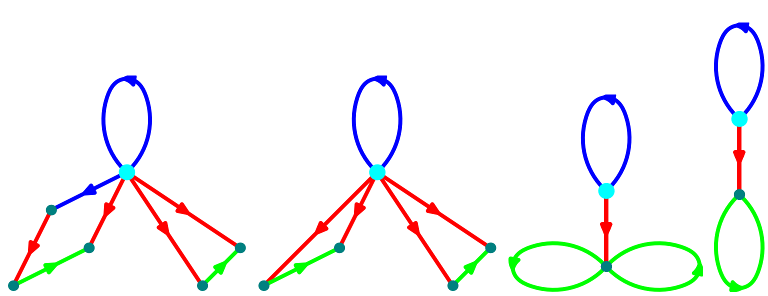

The non-singular fold at the end detects a failure of injectivity. The final picture also indicates a failure of surjectivity. Indeed, we can read off that the image of the homomorphism is the subgroup generated by $\mathbf{b}$ and $\mathbf{rgr}^{-1}$.

2. The assignment

$$\begin{array}{rcl} \mathbf{b} & \mapsto & \mathbf{gb} \\ \mathbf{g} & \mapsto & \mathbf{rgb} \\ \mathbf{r} & \mapsto & \mathbf{g}^{-1}\mathbf{bg} \end{array}$$

yields:

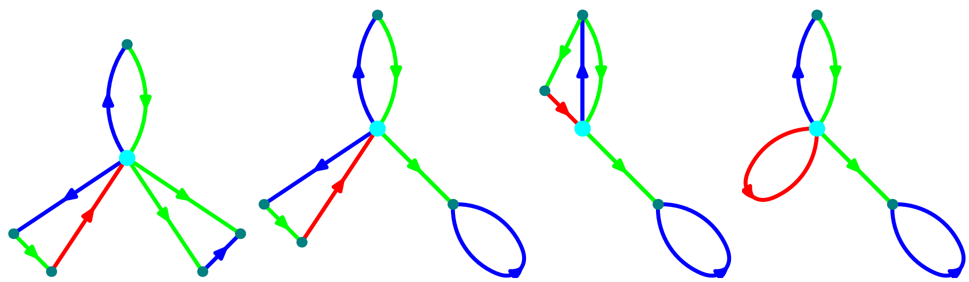

and we see that surjectivity fails.

3. The assignment

$$\begin{array}{rcl} \mathbf{b} & \mapsto & \mathbf{bbrb}^{-1}\mathbf{g} \\ \mathbf{g} & \mapsto & \mathbf{brb}^{-1}\mathbf{g} \\ \mathbf{r} & \mapsto & \mathbf{brb}^{-1} \end{array}$$

has the following topological representative and crucial stages in the folding sequence:

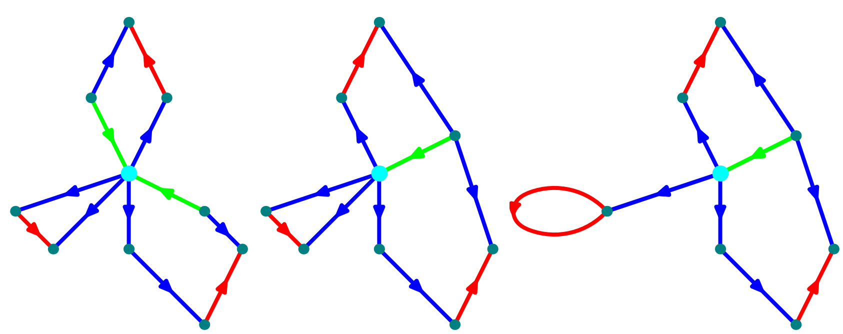

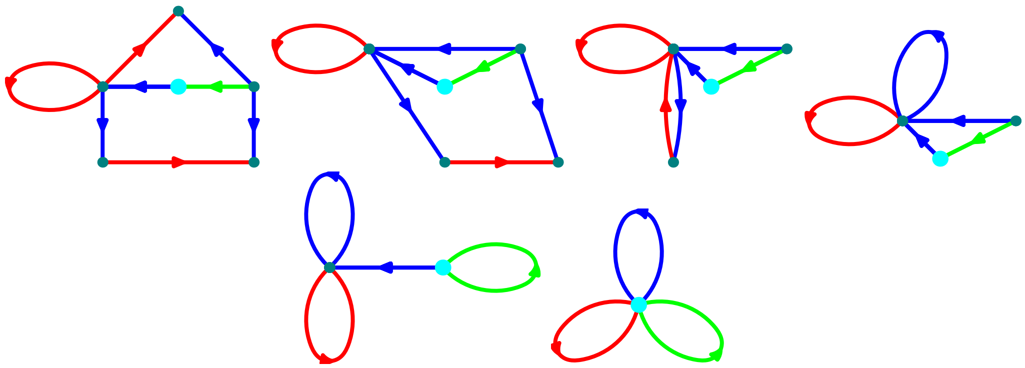

In this case, we actually have an automorphism.

6.1.2 Bounded Cancellation and the Fixed Subgroup

Definition 6.1.12. Let $\varphi : F \to F'$ a homorphism of free groups. A cancellation bound for φ is a number N such that, for any two elements $w, u \in F$ with $|wu| = |w| + |u|$, we have

$$|\varphi(wu)| \geq |\varphi(w)| + |\varphi(u)| - 2N.$$

(Note that letters always cancel in pairs.)

More generally, a graph morphism $\varphi : \Gamma \to \Delta$ has a cancellation bound N if, for any two reduced paths p and q in Γ whose concatenation is also reduced, we have the inequality

$$|\varphi_*(pq)| \geq |\varphi_*(p)| + |\varphi_*(q)| - 2N$$

of path-lengths. Here $|\varphi_*(p)|$ is the reduced path homotopic relative to endpoints to $\varphi(p)$.

Theorem 6.1.13. *If $\varphi : F \to F'$ is an injective homomorphism of finitely generated free groups, then φ has a cancellation bound.*

Observation 6.1.14. *If $\varphi : \Gamma \to \Delta$ is locally injective, then it has cancellation bound 0.*

Lemma 6.1.15. *Let N_1 be a cancellation bound for $\varphi_1 : \Gamma_0 \to \Gamma_1$, and let N_2 be a cancellation bound for $\varphi_2 : \Gamma_1 \to \Gamma_2$. Then $N_1 + N_2$ is a cancellation bound for $\varphi_2 \circ \varphi_1$.*

Proof. Let p and q be reduced paths in Γ_0 whose concatenation is also reduced. Since N_1 is a cancellation bound for φ_1, there are reduced paths p_1, q_1 and r_1 in Γ_1 such that:

$$\begin{aligned} \varphi_{1*}(p) &= p_1 r_1 \\ \varphi_{1*}(q) &= r_1^{\mathrm{rev}} q_1 \\ \varphi_{1*}(pq) &= p_1 q_1 \\ |r_1| &\leq N. \end{aligned}$$

Now p_1 and q_1 are reduced and have a reduce concatenation; and since N_2 is a cancellation bound for φ_2, we conclude that

$$\begin{aligned} |\varphi_2 \circ \varphi_{1*}(pq)| &= |\varphi_{2*}(p_1 q_1)| \\ &\geq |\varphi_{2*}(p_1)| + |\varphi_{2*}(q_1)| - 2N_2 \\ &\geq |\varphi_2 \circ \varphi_{1*}(p)| - N_1 + |\varphi_2 \circ \varphi_{1*}(q)| - N_1 - 2N_2 \end{aligned}$$

which is what was claimed. **q.e.d.**

Lemma 6.1.16. *Every non-singular fold $f : \Gamma \to \Delta$ has cancellation bound* 1.

Proof. Let $\vec{e}$ and $\vec{e}_\bullet$ both start at the vertex v, but suppose these edges have different end points. We think of these two edges as spanning a "V" (this might not be true, as one of the edges might be a loop; but in this argument it does not matter.) The only way to obtain a non-reduced path in Δ from a reduced path in Γ is by traversing the "V". However, if p, q and pq are reduced paths in Γ at most one traversal of the "V" can take place at the point where p and q are concatenated. **q.e.d.**

Proof of Theorem (6.1.13). We observed already that an injective homomorphism gives rise to a folding sequence without singular folds. Since the last map in the folding sequence is locally injective, it does has cancellation bound 0. **q.e.d.**

Definition 6.1.17. Let $\varphi : G \to G$ be an injective homomorphism. The fixed subgroup of φ is

$$!!!...!!!(FIXME)$$

6.2 A Generating Set for Aut(F_n)

Theorem 6.2.1 (Nielsen). *The following automorphisms of F_n generate* Aut(F_n)*:*

1. *Transposition of two free generators.*
2. *Inversion of a free generator.*
3. *The autmorphisms $\alpha_{i,j}$ defined as follows:*

$$x_k \mapsto \begin{cases} x_k & \text{if } k \neq i \\ x_i x_j & \text{if } k = i. \end{cases}$$

6.2.1 Proof of Nielsen's Theorem

The idea is to use Stallings folds. So let $\varphi : \overline{R}_n \to R_n$ be a graph morphism representing the automorphism $\alpha : F_n \to F_n$. We decompose

$$\varphi = \psi \circ f_r \circ \cdots \circ f_1$$

with ψ locally injective and $f : \Gamma_{i-1} \to \Gamma_i$ a fold. We know that ψ is actually an isomorphism of graphs. Thus this gives a permutation of the genertators, some of which are possibly inverted. We will want to recognize the folds as being related to the generators $\alpha_{i,j}$. However, this is not straight forward since there is no canonical identification of F_n with the fundamental groups of the intermediate graphs Γ_i.

The way to fix this, is to consider spanning trees in these graphs. Since permuting and inverting generators is covered by the generating set, we will need a data structure that just keeps track of an unodered set of free generators up to inversion (we could call a two element subset $\{g, g^{-1}\} \subset F_n$ an unsigned element).

Definition 6.2.2. So let Γ be a graph together with

1. a labeling, i.e., a graph morphism $\rho : \Gamma \to R_n$ (one should think of this as an assignment of free generators or their inverses to the oriented edges such that swapping the orientation of an edge corresponds to inverting the generator),
2. a base vertex v_0,
3. and a spanning tree T.

We will associate to this the following set of unsigned elements in F_n: For each edge e not in T, there is a reduced cyclic edge path, unique up to orientation, starting at v_0 traveling along a geodesic in T going through e and heading back to v_0 within T. Collect all these elements. Let $S(\Gamma, \rho, v_0, T)$ denote this set.

We have to study how this set changes with respect to the following transformations:

1. Change of the spanning tree.
2. Folding the graph.

Let us do change of spanning trees first. To simplify matters, we will want to change spanning trees only a little bit, say replacing one edge at a time.

Definition 6.2.3. Let Γ be a graph. The complex of forests is the simplicial complex $\mathcal{F}(\Gamma)$ whose vertices are the non-loop edges in Γ and whose simplices are the subforests in Γ. Note that all maximal simplices in $\mathcal{F}(\Gamma)$ have the same dimension. Such complexes are called chamber complexes—the chambers are the maximal simplices. In the case of $\mathcal{F}(\Gamma)$ the chambers are precisely the spanning trees.

Two chambers are called adjacent if they share a codimension-1 face. A gallery is a sequence of chambers such that neighboring terms in the sequence are adjacent chambers.

Lemma 6.2.4. *$\mathcal{F}(\Gamma)$ is gallery connected, i.e., any two chambers are joined by a gallery.*

Proof. Let T and T' two spanning trees, and let e be an edge of T' that does not occur in T. Adding this edge to T will create a circle. So we are to remove an edge from this circle. Note that any edge will do. At least one of the edges along this circle does not belong to T' since this tree does not contain circles. So we can exchange an edge that is in T but not in T' by the edge e that is in T' but not in T and form a new spanning tree. However, this tree differs from T' in fewer places. So, we keep doing this until we have removed all differences (which are only finite in number). **q.e.d.**

Proposition 6.2.5. *Let T and T' be two adjacent spanning trees in the labelled and basepointed graph Γ. Then $S_T := S(\Gamma, \rho, v_0, T)$ and $S_{T'} := S(\Gamma, \rho, v_0, T')$ are related as follows:*

> *There is an element $g \in S_T \cap S_{T'}$, and for every other element $h \in S_{T'}$ there is an element $h' \in S_T$ such that one of the following holds:*

$$\begin{aligned} h &= h' \\ h &= h'g \\ h &= h'g^{-1} \\ h &= gh' \\ h &= g^{-1}h' \end{aligned}$$

Note that we can realize a transition of this type as a product of Nielsen generators.

Proof. !!! PICTURE !!! **q.e.d.**

The very same picture also yields our first result about how folds change the generating set:

Proposition 6.2.6. *Let $f : \Gamma \to \Delta$ be a fold compatible with the labeling ρ that identifies an edge in the spanning tree T with a loop. Then Δ has a spanning tree T' induced by T and a labeling τ induced by the labeling ρ. and $S_\Gamma := S(\Gamma, \rho, v_0, T)$ is related to $S_\Delta := S(\Delta, \tau, v_0, T')$ in the same way as described in (6.2.5).*

Proof. !!! PICTURE !!! **q.e.d.**

Observation 6.2.7. *A fold of two edges in the spanning tree does not affect* $S(\Gamma, \rho, v_0, T)$. **q.e.d.**

We put everything together. If a fold identifies two edges none of which is a loop, then we can change the spanning tree to contain both of these edges. The change of the spanning tree is taken care of by (6.2.5). If one of the edges is a loop, we can at least put the other edge in the spanning tree (it cannot be a loop itself, since we do not have singular folds). Afterwards, we are done by (6.2.6). Therefore, along our chain of folds, we can use Nielsen generators to realize each fold.

6.2.2 The Homotopytype of the Complex of Forests

Theorem 6.2.8. *Let* Γ *be a finite graph with* $m+2$ *vertices. Then the complex* $\mathcal{F}(\Gamma)$ *is homotopy equivalent to a wedge of* m*-spheres.*

Proof. Induct on the number of edges. Starting point is the case of a bridge edge which serves as a cone point. The Induction step is that removing a non-bride gives you a complex of the same type with fewer edges which is m-spherical by induction. Now, the relative link of the vertex corresponding to the removed edge is a the forest complex for the graph obtained by collapsing this edge. This is $(m-1)$-spherical. **q.e.d.**

Exercise 6.2.9. Find a recursive way to compute the number of spanning trees in a graph Γ.

Exercise 6.2.10. Let M be a finite set. The partition complex X is the flag complex whose vertices are non-trivial partitions of M into two disjoint subsets (nontrivial means, none of the subsets is empty). There is an edge between $\{S_0, T_0\}$ and $\{S_1, T_1\}$ if these two partitions are nested, i.e, one of the four possible inclusions hold:

$$S_0 \subsetneq S_1, \quad T_0 \subsetneq S_1, \quad S_0 \subsetneq T_1, \quad T_0 \subsetneq T_1.$$

Prove that X is a wedge of spheres.

6.3 Outer Space and its Relatives

See: [CuVo86], [BrVo01], [CuMo87], [HaVo96], [Skor??].

6.3.1 Categories Based on Graphs

We already used a category of graphs above when we used Stallings folds. In that case, the morphisms were not allowed to crush edges. We shall have a need for at least two other categories based on the same class of objects (graphs) but with different sets of morphisms. But first, let us review a little bit about graphs.

We call a graph Γ trivalent if every vertex has degree at least 3. A graph Γ is called n-connected if simultaneously removing strictly fewer than n edges will not disconnect Γ. (Removing an edge does not change the vertex set, we remove the interior only.)

Exercise 6.3.1. Let Γ be n-connected, and let v and w be two vertices in Γ. Show that there are n edge-disjoint paths from v to w.

A 2-connected graph is called a core graph.

Exercise 6.3.2. Show that a graph is a core graph if and only if it is a union of reduced loops. (Recall that an edge path is reduced if it does not contain a subpath of the form $\vec{e} \rightarrowtail \mathrm{op}(\vec{e})$.

Corollary 6.3.3. *Every connected graph Γ has a core, i.e., a maximal core graph, which is unique; in fact, it is the union of all reduced closed edge paths in Γ.* **q.e.d.**

Definition 6.3.4. Two edges e_0 and e_1 in a core graph Γ are equivalent if one of the following equivalent conditions is satisfied:

1. Either, removing both edges simultaneously disconnects Γ; or $e_0 = e_1$.
2. Every reduced loop that passes through e_0 also passes through e_1.

From the first wording, it is apparent that equivalence is a reflexive and symmetric relation. From the second formulation, we infer that equivalence is transitive.

Exercise 6.3.5. Show that the conditions (1) and (2) are equivalent.

Based on the class of graphs, there is the category of collapses whose objects are graphs and whose morphisms are given by collapsing subforests: A collaps $c : \Gamma_0 \to \Gamma_1$ is an isotopy class of a a map that sends edges either homeomorphically to edges or crushes them to points such that the preimage of the 0-skeleton of Γ_1 is a subforest of Γ_0.

Observation 6.3.6. *Every collaps is a homotopy equivalence.*

Observation 6.3.7. *Vertex valency and graph connectivity can never decrease during a collaps. In particular any collaps of a trivalent core graph is a trivalent core graph.*

6.3.2 Marked Graphs, Labelled Graphs, and Metric Trees

Exercise 6.3.8. Let Γ_0 and Γ_1 be two finite graphs with base points P_0 and P_1. Prove that every isomorphism

$$\phi : \pi_1(\Gamma_0, P_0) \to \pi_1(\Gamma_1, P_1)$$

is induced by a base point preserving homotopy equivalence $f : \Gamma_0 \to \Gamma_1$. Moreover, show that this homotopy equivalence f is unique up to homotopy trough base point preserving maps.

Exercise 6.3.9. Show that a map $f : \Gamma_0 \to \Gamma_1$ between graphs is a homotopy equivalence if and only if it induces an isomorphism of fundamental groups.

Exercise 6.3.10. Let $f : \Gamma \to \Gamma$ be a self-homotopy equivalence of the base pointed finite graph Γ. Show that f is homotopic (not relative to the base point!) to the identity if and only if f induces an inner automorphism of $\pi_1(\Gamma)$.

Let R be the rose with n loops.

Definition 6.3.11. A metric graph is a finite graph Γ together with an assignment of strictly positive real numbers (lengths) to its unoriented edges. The sum of of all these lengths is the volume of Γ.

A marking of a Γ is a homotopy equivalence

$$\mu : R \to \Gamma.$$

A labelling of Γ is a homotopy equivalence

$$\lambda : \Gamma \to \Gamma.$$

Markings and labellings of base pointed graphs are supposed to preserve base points.

Two markings $\mu_0 : R \to \Gamma_0$ and $\mu_1 : R \to \Gamma_1$ of (metric) graphs are equivalent if there is an isomorphism (isometry) $\zeta : \Gamma_0 \to \Gamma_1$ such that the diagram

$$\begin{tikzcd} \Gamma_0 \arrow[rr, "\zeta"] & & \Gamma_1 \\ & R \arrow[ul, "\mu_0"] \arrow[ur, "\mu_1"'] & \end{tikzcd}$$

commutes up to homotopy (relative to base points if needed).

Similarly, two labellings $\lambda_0 : \Gamma_0 \to R$ and $\lambda_1 : \Gamma_1 \to R$ are equivalent, if there is an isomorphism (isometry) $\zeta : \Gamma_0 \to \Gamma_1$ such that the diagram

$$\begin{tikzcd} \Gamma_0 \arrow[rr, "\zeta"] \arrow[dr, "\lambda_0"'] & & \Gamma_1 \arrow[dl, "\lambda_1"] \\ & R & \end{tikzcd}$$

commutes up to homotopy (again base point preserving if there are base points involved).

Observation 6.3.12. *There is a bijective correspondence of markings and labellings given by passage to a homotopy inverse. This correspondence is compatible with equivalence.*

Observation 6.3.13. *The group* $\mathrm{Aut}(F_n)$*, regarded as the mapping class group of* R *acts by composition on the set of marked graphs as well as on the set of labelled graphs. Both actions are compatible with equivalence. Markings and labellings, however, do not form isomorphic* $\mathrm{Aut}(F_n)$*-sets: one of them is a right* $\mathrm{Aut}(F_n)$*-set and the other one is a left-*$\mathrm{Aut}(F_n)$ *set, and switching from one side to the other involves inverting the group element.*

6.3.3 Metric Trees and $\mathbb{R}$-Trees

Definition 6.3.14. A CAT$(-\infty)$-space is called an $\underline{\mathbb{R}\text{-tree}}$ T. This is a geodesic metric space wherein every geodesic triangle degenerates to a tripod. Let us expand this: For any two points $x, y \in T$, put

$$[x,y] := \{z \in T \mid d(x,y) = d(x,z) + d(z,y)\}.$$

Then, T is an $\mathbb{R}$-tree if the following conditions hold:

1. For all pairs (x,y), the set $[x,y]$ is isometric to a segment in $\mathbb{R}$.

2. Whenever $[x,y] \cap [y,z] = \{y\}$, we have $[x,z] = [x,y] \cup [y,z]$.

3. For any three point x, y, and z, there is a point c such that

$$[x,y] \cap [x,z] = [x,c].$$

The $\underline{\text{link}}$ $\mathrm{Lk}(x)$ of a x is the set of infinitesimal geodesics issuing from x. More precisely, call two geodesic segments starting at x equivalent if their intersection consists of more that $\{x\}$. Note that in this case, the intersection contains a whole non-trivial geodesic segment. A point in $\mathrm{Lk}(x)$ is an equivalence class of geodesic segments starting at x. Note that we have a canonical map:

$$\begin{aligned} T - \{x\} &\rightarrow \mathrm{Lk}(x) \\ y &\mapsto [x, y] \end{aligned}$$

A point is called a vertex if its link does not contain precisely two elements. A vertex is called terminal, if its link is empty or contains one element.

Exercise 6.3.15. True or false: The elements of the link $\mathrm{Lk}(x)$ correspond bijectively to the components of $T - \{x\}$.

Definition 6.3.16. Let G be a group that acts on T by isometries. The action is called minimal if T does not contain a proper G-invariant subtree. The action is called reducible if one of the following holds:

1. Every element of G is elliptic.
2. T has an end that is stabilized by all of G.
3. The action of G on the space of ends of T has an invariant set consisting of precisely two ends. In this case, the action is hyperbolic.

Otherwise, the action is called irreducible. An action is semi-simple if it is trivial, hypercolic, or irreducible.

The translation length

$$\begin{aligned} \underline{\tau} : G &\rightarrow \mathbb{R}^+ \\ g &\mapsto \min_{x \in T} d(x, gx) \end{aligned}$$

of a given action is a length function because of the triangle inequality.

6.3.4 The Definition of Auter Space and Outer Space

Outer space, $\mathcal{X} := \mathcal{X}_n$ is a space acted upon by $\mathrm{Out}(F_n)$. Its construction is based on metric graphs with a marking. Adding a base point to the graph, we will obtain a construction for auter space, $\mathcal{Y} := \mathcal{Y}_n$, upon which the group $\mathrm{Aut}(F_n)$ acts. Let us describe the construction of $\mathcal{X}$.

The set $\mathcal{X} := \mathcal{X}_n$ of all equivalence classes of R_n-marked small volume 1 graphs (with base point) is outer space (auter space). We have to put a topology on this set. The idea is, of course, that changing the length of edges slightly should not move you far in $\mathcal{X}$. However, we have to discuss the case where the length of an edge goes to 0 along a path: if this happens to a loop we run to infinity; if it happens to a non-loop, we move towards the collaps.

Before we embark on the topology of $\mathcal{X}$, let us give an alternative description of the underlying set. Let Γ be a small metric graph of volume 1, and let $\mu : R \to \Gamma$ be a marking. The universal cover of Γ is a metric tree T. The marking induces an action of $F_n = \pi_1(R)$ on T by isometries.

Definition 6.3.17. A length function $\ell : G \to \mathbb{R}^+$ is a function satisfying

$$\ell(gh) \leq \ell(g) + \ell(h)\,.$$

6.4 Proofs of Contractibility

6.4.1 Proof by Continous Folding (the Trees Proof)

This is based on the work of M. Steiner and D. Skora.

Definition 6.4.1. A map $\varphi : \underline{\Gamma} \to \underline{\Delta}$ between metric graphs is a piecewise isometry if there is a decomposition of $\underline{\Gamma}$ into line segments such that f is an isometry on each segment.

Let Γ be a graph and $\underline{\Delta}$ be a metric graph. Every graph morphism $f : \Gamma \to \underline{\Delta}$ is homotopic relative to vertices to a piecewise isometry φ for some suitably chosen metric on Γ provided f. (Recall that graph morphisms do not crush edges.) The piecewise isometry φ is called tight if it minimizes the length of each edge in Γ, i.e., every piecewise isometry ψ that is homotopic to φ relative to vertices does not allow for smaller edge lengths assigned to Γ. (Note that we can consider each edge separately since we are dealing with homotopies that leave all vertices fixed.)

Observation 6.4.2. *If a piecewise isometry of metric graphs is not locally injective, there is a continuous metric fold.*

6.4.2 Proof by Sophisticated Low-Dimensional Topology (the Spheres Proof)

This is based on the work of A. Hatcher and K. Vogtmann.

Note that $\pi_1(\mathbb{S}^1 \times \mathbb{S}^2)$ is infinte cyclic. Let M be a connected sum of n copies of $\mathbb{S}^1 \times \mathbb{S}^2$. Then $F_n = \pi_1(M)$. Note that M consists of two n-handlebodies that are glued via the identity along their boundary surfaces. It is, therefore, easy to draw half of M.

Definition 6.4.3. A sphere set is a finite set of disjoint, embedded spheres in M. Such a set is simple if every complementary region is simply connected. Two sphere sets are compatible if their union is a sphere set.

A sphere system is a sphere set wherein no two spheres are isotopic and no sphere is trivial.

A Dehn twist in M along a sphere S with a specified axis a diffeomorphism that is the identity outside a tubular neighborhood $S \times [0,1]$ of S and that acts as a rotation by $2\pi t$ about the given axis in the slice $S \times \{t\}$. Dehn twists change sphere systems only up to homotopy.

Fact 6.4.4 (Laudenbach). *Homotopic sphere systems are isotopic.*

Corollary 6.4.5. *Dehn twists act trivially on Y.*

Fact 6.4.6 (Laudenbachs Dehn-Nielsen Theorem). *The group* $\mathrm{Out}(F_n)$ *is isomorphic to*

$$\pi_0(\mathrm{Diff}^\circ(M)) \Big/ \textit{subgroup generated by Dehn twists}$$

and the isomorphism is given by the natural action of $\mathrm{Diff}^\circ(M) = F_n$ *on* $\pi_1(M)$*. In particular, there is a natural action of* $\mathrm{Out}(F_n)$ *on* Y*.*

Fix a simple sphere system $\underline{\mathcal{S}} = \{\underline{S}_1, \ldots, \underline{S}_s\}$.

A. Hatcher figured out that two sphere systems can be isotoped as to minimize their intersections.

Definition 6.4.7. A sphere system $\mathcal{S}$ is in normal form with respect to $\underline{\mathcal{S}}$, if for every sphere S in $\mathcal{S}$, one of the following holds:

1. The sphere S is contained in $\underline{\mathcal{S}}$.
2. The sphere $S \cup \underline{\mathcal{S}}$ is a sphere system.
3. The sphere S has non-empty, transverse intersection with $\underline{\mathcal{S}}$ and, for each component W of $M - \underline{\mathcal{S}}$ the following two conditions are both satisfied:
   (a) Each component of $S \cap W$ has at most one boundary circle in each boundary sphere of W.
   (b) No component of $S \cap W$ is isotopic in W to a disk in the boundary of W.

Two sphere systems $\mathcal{S}_0$ and $\mathcal{S}_1$ in normal form are equivalent if there is a homotopy from $\mathcal{S}_0$ to $\mathcal{S}_1$ such that:

1. The common spheres of S_0 and $\underline{\mathcal{S}}$ stay fixed pointwise during the homotopy.

2. The homotopy is transverse to $\underline{\mathcal{S}}$ on the other spheres at all times.

3. The circles in $S_t \cap \underline{\mathcal{S}}$ vary only by an isotopy inside $\underline{\mathcal{S}}$. In particular, there is a well defined notion of innermost circles: A circle component C in $\mathcal{S}_0 \cap \underline{\mathcal{S}}$ is innermost if it bounds a disc D in $\underline{\mathcal{S}}$ such that $D \cap \mathcal{S}_0 = C$.

Fact 6.4.8 (Hatcher). *Every sphere system can be isotoped as to be in normal form with respect to $\underline{\mathcal{S}}$; and any two isotopic sphere systems in normal form are equivalent.*

6.4.3 A Continuous Contracting Flow

Let $\overline{\mathcal{X}}$ be the simplicial complex whose m-simplices are isotopy classes of $(m+1)$-sphere systems. Note that these systems are not required to be simple. Let Y be the simplicial complex whose vertices are isotopy classes of simple sphere systems and wherein a set of those systems forms a simplex if it is a chain (totally ordered) with respect to inclusion.

Proposition 6.4.9. *The simplicial closure of outer space is isomorphic to $\overline{\mathcal{X}}$ and the spine of outer space is isomorphic to Y.*

Proof. !!! fix me !!! **q.e.d.**

Construction 6.4.10 (Innermost Surgery). Let S be a sphere set that intersects $\underline{\mathcal{S}}$ transversally except for common spheres. Let C be an innermost circle component in $S \cap \underline{\mathcal{S}}$, and let $S \in S$ be the sphere containing C. Consider a parallel copy of S that intersects $\underline{\mathcal{S}}$ in an even smaller circle and perform surgery along the disc D, i.e., cut this copy along the disc and glue in discs to close the holes that are parallel to D.

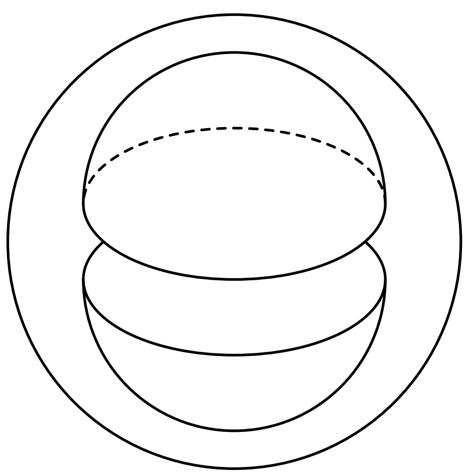

This way, we obtain to spheres S_+ and S_-. Note that $S \cup S_+ \cup S_- - S$ is a sphere set, compatible with S, that intersects $\underline{\mathcal{S}}$ in fewer circle components.

Obviously, we can surger simultaneously along severyl disjoint discs that are attached to the sphere S *from the same side.*

Observation 6.4.11. *Performing surgery in S has the following effects on the complementary regions in $M - S$: The component that contains D is cut along D into two pieces. The component not containing D is changed by attaching a 2-handle. Thus, if the complementary components were simply connected, they stay that way. In other words: If S is simple, than innermost surgery yields simple sphere sets.* **q.e.d.**

A point in $\overline{\mathcal{X}}$ is a formal convex combination of spheres that form a sphere system. The coefficient of a sphere can be thought of as a thickness assigned to the sphere. Spheres of width 0 are deleted from the picture. We will treat isotopy classes of sphere systems like sphere sets. Weights of parallel spheres are added up to give the weight of their isotopy class. We want to employ a continuous version of surgery.

Construction 6.4.12 (Continuous Parallel Innermost Surgery). Let $\mathcal{S} = S_1 \cup \cdots \cup S_r$ with weights $t_1, \ldots, t_r$ satisfying $\sum t_i = 1$. Thicken those spheres S_i a little bit that are not contained in the fixed system $\underline{\mathcal{S}}$. We think of the thickenned spheres as embedded annuli $S_i \times [0, t_i]$. Let $\overline{S}$ the sphere set obtained from $\mathcal{S}$ by replacing S_i by $S_i \times \{0\} \cup S_i \times \{t_i\}$. In $\overline{S}$, the weight t_i is distributed evenly to both of these spheres. (Here, we use the convention that weights of parallel spheres add up.)

For each sphere $\underline{S}$ in the fixed simple system $\underline{\mathcal{S}}$, the intersection $\overline{S} \cap \underline{S}$ consists of disjoint circles, at most one for each sphere $S \in \overline{S}$. Let $T_{\underline{S}}$ be the metric tree whose vertices correspond to complementary regions in $\underline{S} - \overline{S}$ and whose edges correspond to those spheres in $\overline{S}$ that intersect $\underline{S}$. The lenght of an an edge in $T_{\underline{S}}$ is given be the weight of its corresponding sphere in $\overline{S}$.

We now perform innermost surgery simultaneously on all terminal points in the trees $T_{\underline{S}}$. We are doing continuous surgery and the weight is transfered to from the surgered spheres to their successors so that the terminal edges in the trees $T_{\underline{S}}$ shrink with unit speed.

This defines a continuous flow on $\overline{\mathcal{X}}$.

Observation 6.4.13. *The endpoint of each flow-line is a sphere system all of whose spheres are disjoint from the fixed system $\underline{\mathcal{S}}$. Thus, if $\underline{\mathcal{S}}$ is maximal, all spheres will be parallel to spheres in $\underline{\mathcal{S}}$ at the end of the flow. In this case, the flow-lines all end in the simplex defined by $\underline{\mathcal{S}}$, and the flow visibly defines a contraction.* **q.e.d.**

Observation 6.4.14. *It follows from (6.4.11) that the flow restricts to a flow on $\mathcal{X} \subset \overline{\mathcal{X}}$. It follows that Outer Space is contractible.* **q.e.d.**

6.4.4 A Discrete Version of the Argument and an Exponential Isoperimetric Inequality for Out(F_n)

Sphere sets and isotopy classes of sphere systems are partially ordered by inclusion. Let X and Y be the geometric realizations of these posets. There is a natural projection $X \to Y$ that turns a sphere set S into a sphere system (up to isotopy) by deleting all trivial spheres in S and, aditionally, deleting all but one sphere in each isotopy class of S.

We will construct a contraction of the spine of Outer Space. To this end, we use non-continuous surgery, i.e., we do not talk about slowly transfering the weight; instead surgered spheres are removed at once and replaced by their successors. Thus, we are hopping from one vertex to a neighbor in X. However, we will retain the notion that a sphere has an inside and an outside, and that surgery cannot be done on both sides at the same time.

Construction 6.4.15 (Parallel Innermost Surgery). A sphere set S is $\underline{\mathcal{S}}$-oriented if S intersects no sphere of $\underline{\mathcal{S}}$ in a single circle, and all innermost surgery discs for the sphere $S \in S$, lie on the same side of S. (Note that because of the first condition, there is no ambiguity as to what the surgery discs are.)

Parallel innermost surgery is the operation of performing surgery along all innermost discs in a $\underline{\mathcal{S}}$-oriented sphere set at once.

Lemma 6.4.16. *For any sphere set S, let $\overline{S}$ be a sphere system obtained by sticking in parallel copies for all spheres in S not contained in $\underline{\mathcal{S}}$. Then $\overline{S}$ is $\underline{\mathcal{S}}$-oriented, and so is any sphere set obtain from it by a finite sequence of parallel innermost surgeries.*

Proof. Let Γ_S be the finite graph whose vertices are the components of $\underline{\mathcal{S}} - S]$ and whose edges are the circle components of $\underline{\mathcal{S}} \cap S$. Note that Γ_S is a finite union of trees, one for each sphere in $\underline{\mathcal{S}}$ that intersects S transversally. Observe, that $\Gamma_{\overline{S}}$ is the barycentric subdivision of Γ_S. Orient twin spheres on opposite ways. This induces an orientation on $\Gamma_{\overline{S}}$. The innermost discs correspond to extremal edges in $\Gamma_{\overline{S}}$. The corresponding edges are all oriented the same way (say outward), and the surgery discs are all on the same side. Thus, $\overline{S}$ is $\underline{\mathcal{S}}$-oriented.

Performing the surgery prunes the extremal edges from the graph. Now all extremal edges are oriented inward, and the new surgery discs are again all on the same side of their discs. This continues. The orientations of extremal edges swap sign every turn, but they remain in sync. **q.e.d.**

Now we can give a continuous version of surgery. A point in $\overline{\mathcal{X}}$ is a convex combination of spheres in a sphere system. We think of the weights as thicknesses of the sphere. This is our interpretation for the two parallel copies of the sphere used in parallel innermost surgery. Instead of combinatorially hopping from one sphere system to the next, observe that all the spheres are compatible. Thus, we can think of surgery as decreasing the weight for a sphere and icreasing correspondingly the weights of its surgery children.

Definition 6.4.17. For a simple sphere system $\mathcal{S}$ in normal form, we call the sequence $\overline{\mathcal{S}} = \mathcal{S}_0, \mathcal{S}_1, \mathcal{S}_2, \ldots$ obtained by iterated parallel innermost surgery the canonical path. Since the number of components in $\mathcal{S}_i \cap \underline{\mathcal{S}}$ decreases in each step, the canonical path is finite and terminates with a sphere set S_r compatible with $\underline{\mathcal{S}}$. We extend the canonical path by two steps: $\mathcal{S}_r, \mathcal{S}_r \cup \underline{\mathcal{S}}, \underline{\mathcal{S}}$. The combing path of $\mathcal{S}$ is the projection of the extended canonical path into Y.

Lemma 6.4.18. *Equivalent sphere systems $\mathcal{S}_0$ and $\mathcal{S}_1$ have identical combing paths.*

Proof. The homotopy proving the equivalence is an isotopy in $\mathcal{S}_t \cap \underline{\mathcal{S}}$. Moreover, $\mathcal{S}_t$ moves inside components of $M - \underline{\mathcal{S}}$. Therefore, the surgeries correspond bijectively. **q.e.d.**

Lemma 6.4.19. *Let $\mathcal{S}$ and $\mathcal{S}'$ be two simple sphere systems, and suppose $\mathcal{S} \subset \mathcal{S}'$. Then the combing paths for $\mathcal{S}$ and $\mathcal{S}'$ are close.*

Proof. It suffices to consider the case where $\mathcal{S}$ is obtained from $\mathcal{S}'$ by deleting one sphere. **q.e.d.**

Let z be the complex whose m-simplices are $(m+1)$-systems of spheres. Let z' be the union of all those simplices in z that correspond to simple sphere systems. You should think of a point in z as a sphere set where the spheres have a thickness and these weights add up to 1.

Proposition 6.4.20. *z is contractible. The contraction induces a contraction of z'.*

Proof. Now the edges in the trees have a thickness.You shrink them at unit speed from the terminal points. A sphere whose thickness becomes 0 is surged.

Finally, this contraction restricts to z' because the 1-connectedness of the complement is preserved in the surgery process. **q.e.d.**

Consider a simplex $\underline{\mathcal{S}}_1 \subset \underline{\mathcal{S}}_2 \subset \cdots \subset \underline{\mathcal{S}}_m$ in X. Let us fix a sphere set S and let $\overline{S}$ be a sphere set that contains two parallel copies for each sphere in S each assigned an orientation pointing away from the other sphere. We assume $\overline{S}$ to be normal with respect to $\underline{\mathcal{S}}_m$. This implies that $\overline{S}$ is normal with respect to all $\underline{\mathcal{S}}_j$.

Construction 6.4.21. Put $\Delta_1 := \underline{\mathcal{S}}_1$ and, for $j > 1$, put $\Delta_j := \underline{\mathcal{S}}_j - \underline{\mathcal{S}}_{j-1}$. Note that the Δ_j are disjoint. For each multi-index $I := (i_1, \ldots, i_m)$, let S_I be the sphere set obtained from $\overline{S}$ by doing the first i_j steps of parallel innermost surgery in the trees $T_{\underline{S}}$ with $\underline{S} \in \Delta_j$. Since the $\mathcal{S}_j$ are disjoint, the order of these surgeries does not matter.

Observation 6.4.22. *If $I = (i_1, \ldots, i_m)$ has all even or all odd entries, then the spheres in S_I corresponding to innermost discs are all oriented the same way: outward if the entries are even, inward if the entries are odd.* **q.e.d.**

Observation 6.4.23. *Let $I = (i_1, \ldots, i_m)$ be an all-odd multi-index. Then for any $J = (k_1, \ldots, k_m) \in \{0, 1\}^m$, the vertices*

$$S_{(i_1,i_2,i_3\ldots,i_m)}, S_{(i_1+k_1,i_2,i_3,\ldots,i_m)}, S_{(i_1+k_1,i_2+k_2,i_3,\ldots,i_m)}, \ldots, S_{(i_1+k_1,i_2+k_2,i_3+k_3,\ldots,i_m+k_m)}$$

form a simplex in Y. The reason is that these successive surgeries use disjoint surgery discs all on the same side of S_I. **q.e.d.**

We can push this a little:

Lemma 6.4.24. *Let $I = (i_1, \ldots, i_m)$ be an all-odd multi-index. Then for any $J = (k_1, \ldots, k_m) \in \{-1, 0, 1\}^m$, the vertices*

$$S_{(i_1,i_2,i_3\ldots,i_m)}, S_{(i_1+k_1,i_2,i_3,\ldots,i_m)}, S_{(i_1+k_1,i_2+k_2,i_3,\ldots,i_m)}, \ldots, S_{(i_1+k_1,i_2+k_2,i_3+k_3,\ldots,i_m+k_m)}$$

form a simplex in Y.

Proof. The new feature are the negative entries of $J = (k_1, \ldots, k_m)$. Here, we are asked to undo a surgery along a disc. Note that this is performing a surgery along an arc transverse to the innermost disc that was surgered away. Yet again, the lemma hinges upon the observation that we can perform all these operations independently in any order always obtaining compatible sphere sets becauce all involved surgery arcs and discs are disjoint and on the same side of $\underline{\mathcal{S}}$. **q.e.d.**

Corollary 6.4.25. *Let the positive orthant in $\mathbb{R}^m$ be triangulated so that the integer lattice points are the vertices and a !!! fix me !!!. Then, the map $I \mapsto S_I$ induces a simplicial map from the triangulation of the positive orthant to Y.* **q.e.d.**

Observation 6.4.26. *The vertices $(i_1, \ldots, i_m)$ with*

$$i_m \leq i_{m-1} \leq \cdots \leq i_2 \leq i_1 \leq C$$

form a big simplex σ_C of dimension $m + 1$ whose vertices are

$$(0, 0, 0, \ldots, 0, 0), (C, 0, 0, \ldots, 0, 0), (C, C, 0, \ldots, 0, 0), \ldots (C, C, \ldots, C, 0), (C, C, \ldots, C, C).$$

q.e.d.

Let C_i be the number of surgery steps in the canonical path from S to $\underline{\mathcal{S}}_i$. Then

$$C_1 \leq C_2 \leq \cdots \leq C_m.$$

Observation 6.4.27. *All vertices $S_{(C_1,\ldots,C_j,i_{j+1},\ldots,i_m)}$, are contained in the star of $\underline{\mathcal{S}}_j$.*

Put $C := C_1$.

Construction 6.4.28. Observe that the path

$$S_{(C,C,\ldots,C,C)}, S_{(C+1,C+1,\ldots,C+1,C+1)}, \ldots, S_{(C_m,C_m,\ldots,C_m,C_m)}, \underline{\mathcal{S}}_m$$

is an edge path.

All the vertices with indices $\geq C_{m-1}$ are in the star of $\underline{\mathcal{S}}_{m-1}$. So we cone that off.

Inductively, we cone off more and more. Eventually, we constructed a big canonical simplex T with vertices $S, \underline{\mathcal{S}}_1, \ldots, \underline{\mathcal{S}}_m$.

Corollary 6.4.29. *The spine of outer space is contractible.* **q.e.d.**

Lemma 6.4.30. *Let $\underline{\mathcal{S}} \subseteq \underline{\mathcal{S}}'$ be two simple sphere systems adjacent in Y, and let C and C' be the lengths of the canonical paths from S to $\underline{\mathcal{S}}$ and $\underline{\mathcal{S}}'$. Prove that*

$$C \leq C' \leq 2nC + n.$$

Proof. The numbers $2C$ and $2C'$ are the maximum diameters of the trees $T_{\underline{S}}$ where the sphere $\underline{S'}$ ranges over $\underline{\mathcal{S}}$ for C and over $\underline{\mathcal{S}}'$ for C'. Thus, $2C$ and $2C'$ measure the length of a longest nested chain of circles in $\overline{S} \cap \underline{S'}$ for $\underline{S'} \in \underline{\mathcal{S}}$ and $\underline{S'} \in \underline{\mathcal{S}}'$, respectively.

So fix $\underline{S'} \in \underline{\mathcal{S}}'$ and let $C_1, C_2, \ldots, C_{2C'}$ be a nested chain of circles of maximum length. Note that if $\underline{S'} \in \underline{\mathcal{S}}$, this chain contributes to C, as well; and we have $C = C'$.

So from now on, we assume $\underline{S'} \in \underline{\mathcal{S}}' - \underline{\mathcal{S}}$. Let P be the component of $M - \underline{\mathcal{S}}$ containing $\underline{S'}$. Recall that P is a simply connected three manifold whose boundary consists of disjoint non-parallel spheres.

To count the circles $C_1, C_2, \ldots$ in a maximal nested chain, we consider for each of these circles the component W_i of $S \cap P$ determined by C_i. There are two kinds of circles, which we count seperately:

1. The component W_i might be disjoint from $\partial(P)$. In this case, W_i is a sphere. By normality, W_i has to separate the boundary spheres of P into two non-empty subsets. Since the length of any chain of such partitions is bounded by the number of boundary spheres of P, the number of those circles in our chain is at most $2n$.

2. The component W_i intersects the boundary of P. Since the circles C_i are nested, the circles in $W_i \cap \underline{S}$ are nested for any boundary sphere $\underline{S} \subset \partial(P)$. The length of a chain intersecting $\underline{S}$ non-trivially is therefore bounded by $2C$ as $\underline{S}$ belongs to $\underline{\mathcal{S}}$. The number of boundary spheres, again, is bounded by $2n$.

Adding the two counts, we obtain

$$2C' \leq 2n2C + 2n,$$

which yields the desired result. **q.e.d.**

Exercise 6.4.31. Show that T contains $C_1^m + C_2^{m-1} + \cdots + C_m^1 + 1$ simplices of top-dimension m. [Hint: Use induction.]

Corollary 6.4.32. $\mathrm{Out}(F_n)$ *satisfies an exponential isoperimetric inequality.* **q.e.d.**

Proof. Fix a loop with N edges. Let $\mathcal{S}$ be a vertex on that loop and consider the contraction towards that vertex $\mathcal{S}$. Since each vertex on the loop is of distance $\frac{1}{2}N$ to $\mathcal{S}$, we infer from (6.4.30) that the combing paths have length bounded from above by C^N. Thus, the number of triangles used in the contraction for one edge is bounded by $C^{C'N}$. Thus the total number of triangles is bounded by $NC^{C'N}$. **q.e.d.**

6.4.5 The Graphs Proof

6.5 !!! FIXME !!!

M. Bridson proved in his thesis that Outer Space does not allow a CAT(0) metric. S. Gersten strengthened this result and showed that there is no CAT(0) space for $\mathrm{Out}(F_n)$ to act upon:

Theorem 6.5.1 (Gersten). *For $n \geq 3$, the group $\mathrm{Aut}(F_n)$, and for $n \geq 4$, the group $\mathrm{Out}(F_n)$ cannot act properly discontinuously on any CAT(0) space.*

We will reproduce his proof.

Remark 6.5.2. Since any CAT(0) group has a quadratic isoperimetric inequality, it would be nice to have a proof that $\mathrm{Out}(F_n)$ does not satisfy a quadratic isoperimetric inequality.

The central idea of Gersten's proof is to find a subgroup in $\mathrm{Aut}(F_3)$ that visibly cannot act properly discontinuously on a CAT(0) space. The group he finds has the following presentation:

$$H := \left\langle a, b, c, s \,\middle|\, sas^{-1} = a, sbs^{-1} = ba, scs^{-1} = ca^2 \right\rangle.$$

Lemma 6.5.3. *Every word with vanishing s-exponent sum is equivalent to a word that is s-free. Every word that has non-zero s-exponent sum is not trivial.*

Proof. If the exponents for s add up to 0, we can find a subword sus^{-1} where u does not involve s. However, the relations allow us to conjugate all the letters in u. Since the right hand sides of the relations do not involve s, we got rid of two occurrences of s. We continue until all occurrences are gone. **q.e.d.**

Consider the homomorphism

$$\begin{array}{rcl} \iota : H & \rightarrow & \mathrm{Aut}(\langle a, b, c\rangle) \\ a & \mapsto & \iota_a \\ b & \mapsto & \iota_b \\ c & \mapsto & \iota_c \\ s & \mapsto & \varphi, \end{array}$$

where $\varphi : \langle a, b, c\rangle \rightarrow \langle a, b, c\rangle$ is the homomorphism

$$\begin{array}{rcl} a & \mapsto & a \\ b & \mapsto & ba \\ c & \mapsto & ca^2. \end{array}$$

It is immediate from the relations that ι is a well defined homomorphism.

Lemma 6.5.4. *The homomorphism ι is injective.*

Proof. Observe that the generators a, b, and c act trivially on the abelianization $\mathbb{Z}^3$ whereas the homomorphism φ descends to an automorphism of $\mathbb{Z}^3$ described by the matrix

$$\begin{pmatrix} 1 & 1 & 2 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}$$

which has infinite order. Thus every element in $\ker(\iota)$ has s-exponent sum 0. By (6.5.3), a kernel element will be in the subgroup generated by a, b, and c. However, their images form a free generating set for $F_3 \leq \mathrm{Aut}(F_3)$. **q.e.d.**

Lemma 6.5.5. *The group H cannot act properly discontinuously on any $CAT(0)$ space.*

Proof. We use the fact that any free abelian group of finite rank m acting properly discontinuously on a CAT(0) space X stabilizes a flat F of dimension m upon which it acts as a group of translations whose geometric rank is m. (This implies that the quotient of the flat modulo the abelian group is compact.)

Observe that we can rewrite

$$H = \langle a, s \mid as = sa \rangle * \langle b, c \rangle / b^{-1}sb = as, c^{-1}sc = a^2 s.$$

Thus, a and s span a free abelian group of rank 2 inside H. Let τ_a, τ_{as}, and τ_{a^2s} be the Euclidean translations on the flat F corresponding to the words a, as, and a^2s. Since the words are conjugate, the three translations have the same length. Hence they lie on a circle. On the other hand, they visibly lie on the line through τ_a with direction τ_s. Thus, we have a line that intersects a circle in three points which is impossible. **q.e.d.**

Proof of (6.5.1). Since $H \leq \mathrm{Aut}(F_3)$ embedds into $\mathrm{Aut}(F_n)$ for $n \geq 3$ and into $\mathrm{Out}(F_n)$ for $n \geq 4$, the claim follows from (6.5.5) **q.e.d.**

Chapter 7

The Mapping Class Group of a Closed Surface

Chapter 8

Braid Groups

Part IV

Miscellaneous Important Groups

Chapter 9

Coxeter Groups and Artin Groups

9.1 Euclidean Reflection Groups

Let

- $\mathbb{E}$ be a Euclidean space, and let
- $\mathcal{H}$ be a set of hyperplanes satisfying the following:
    1. $\mathcal{H}$ is locally finite, i.e., a set of hyperplanes such that any compact subset of $\mathbb{E}$ intersects only finitely many hyperplanes from $\mathcal{H}$.
    2. $\mathcal{H}$ is a W-invariant subset of $\mathbb{E}$ where W is the subgroup of $\mathrm{Isom}(\mathbb{E})$ generated by all reflections ρ_H with $H \in \mathcal{H}$.

Definition 9.1.1. Such a group W is called a Euclidean reflection group.

Exercise 9.1.2. Assume that $\mathcal{H}$ is finite. Show that $\bigcap_{H\in\mathcal{H}} H \neq \emptyset$. See (9.1.16) for a more elaborate statement.

We want to derive a presentation for W.

9.1.1 The Chamber Decomposition of $\mathbb{E}$

A chamber is a complementary component of $\mathcal{H}$, i.e., a component of $\mathbb{E} - \bigcup_{H\in\mathcal{H}} H$. Note that the closure of a chamber C is a convex polytope (possibly non-compact). The faces of this polytope span hyperplanes that belong to $\mathcal{H}$. We say that those hyperplanes from $\mathcal{H}$ are supporting C. For any chamber C, we denote by

- $|C|$ the set of hyperplanes in $\mathcal{H}$ supporting C.

  Two chambers, C and D, are called adjacent along H if $H \cap C = H \cap D$ is a CoDim-1-face. In this case, we write

  $$C|_H D.$$

  They are called adjacent if they are adjacent along some H. In this case, we write

  $$C|D.$$

  Note that adjacency and adjacency along H are symmetric and reflexive relations.

  A gallery is a sequence

  $$C_0|C_1|\cdots|C_r$$

  of chambers such that C_i is adjacent to C_{i+1} for all $i < r$. If $C_i|_H C_{i+1}$, we say that the gallery crosses H at this step.

  The last index r gives the length of the gallery, which henceforth is the number of hyperplanes that are crossed by the gallery. The distance

- $\delta\left(C, D\right)$ of the chambers C and D is the minimum length of a gallery connecting them. Note that two chambers are adjacent if and only if their distance is at most 1.

Exercise 9.1.3. Show that C and D are H-adjacent if and only if H supports both and $\{C, \rho_H C\} = \{D, \rho_H D\}$.

Observation 9.1.4. *Any two chambers are connected by a gallery of finite length.* **q.e.d.**

Exercise 9.1.5. Prove that a gallery from C to D has minimum length if and only if it does not cross any hyperplane twice. Moreover, the set of hyperplanes that are crossed by a minimum length gallery from C to D is precisely the set of those $H \in \mathcal{H}$ that separate C from D. In particular, this set is the same for all those minimum length galleries.

Note that W acts on the set $\mathcal{C}$ of chambers by distance preserving permutations.

Observation 9.1.6. *If the hyperplane H supports the chamber C, then $\rho_H C|_H C$.*

Let us fix an arbitrary chamber

- C^*, the fundamental chamber. Put
- $S := \{\rho_H \mid H \in |C^*|\}$.

Lemma 9.1.7. *W acts transitively on $\mathcal{C}$ and is generated by S.*

Proof. Let
$$C^* = C_0|C_1|\cdots|C_{r-1}|C_r$$
be any gallery starting at C^*. We will show that there are elements $w_i \in \langle S\rangle$ with $C_i = w_iC^*$. This is an easy induction: Suppose w_i has been found already. Let H be the hyperplane with $C_i|_H C_{i+1}$. Then $w_i^{-1}H$ is a hyperplane in $\mathcal{H}$ that supports C^*. Thus
$$\rho_H = w_i s w_i^{-1} \text{ for some } s \in S$$
and
$$C_{i+1} = \rho_H C_i = \rho_H w_i C^* = w_i s w_i^{-1} w_i C^* = w_i s C^*.$$
This way, we constructed an element $w_{i+1} = w_i s \in \langle S\rangle$.

Since every chamber can be connected to C^* by a gallery, the subgroup $\langle S\rangle$ already acts transitively on $\mathcal{C}$.

Consider $H \in \mathcal{H}$. Let $C = wC^*$ (where $w \in \langle S\rangle$) be a chamber supported by H. As we already have observered, there is an element $s \in S$ such that
$$\rho_H = wsw^{-1} \in \langle S\rangle\,.$$
Thus the generating set for W is contained in $\langle S\rangle$. **q.e.d.**

Lemma 9.1.8. *Let $\underline{s} = s_1s_2\cdots s_r$ be a word representing $w \in W$. If this word is a minimum length representative for w, then its length r equals $\delta(C^*, wC^*)$. Otherwise, one can obtain a shorter word representing w by deleting two of the letters, i.e., there are two indices $i < j$ such that*
$$w = s_1\cdots s_{i-1}s_{i+1}\cdots s_{j-1}s_{j+1}\cdots s_L.$$

Proof. Put

- $w_i := s_1\cdots s_i$,
- $C_i := w_iC^*$, and let
- H_i be the hyperplane satisfying $s_i = \rho_{H_i}$.

We claim that the corresponding gallery

$$C^* = C_0|_{w_0H_1}C_1|_{w_1H_2}\cdots C_{r-2}|_{w_{r-2}H_{r-1}}C_{r-1}|_{w_{r-1}H_r}$$

does not cross any hyperplane twice provided that $s_1 \cdots s_r$ is a minimum word length representative for w. Then the claim follows from (9.1.5).

So let us suppose that

$$w_{i-1}H_i = w_{j-1}H_j$$

for some $i < j$. We conclude

$$w_{i-1}s_iw_{i-1}^{-1} = w_{j-1}s_jw_{j-1}^{-1}$$

whence

$$s_1 \cdots s_{i-1}s_is_{i-1} \cdots s_1 = s_1 \cdots s_{j-1}s_js_{j-1} \cdots s_1.$$

Thus,

$$1 = s_i \cdots s_{j-1}s_js_{j-1} \cdots s_{i+1}.$$

Multiplying from the right, we obtain

$$s_{i+1} \cdots s_{j-1} = s_i \cdots s_j$$

which implies that we have a shorter word for w:

$$w = s_1 \cdots s_{i-1}s_{i+1} \cdots s_{j-1}s_{j+1} \cdots s_L$$

This is a contradiction. **q.e.d.**

Corollary 9.1.9. *The action of W on $\mathcal{C}$ is simply transitive.* **q.e.d.**

This corollary allows us to draw the Cayley graph of W with respect to S. Since all generators have order 2, we simplify matters by ommiting all the bi-gons that would arise that way. Thus, we define the reduced Caley graph

$$\Gamma := \Gamma_S(W)$$

of W to have a vertex for each group element and an edge (labelled by s) for each *unordered* pair $\{w, ws\}$. Note that W acts from the left.

Observation 9.1.10. *Pick a point inside the fundamental chamber. The W-orbit of this point can be identified with the vertex set of Γ. The edges of Γ correspond to* CoDim-*1-faces in the chamber decomposition of $\mathbb{E}$. In fact, we can connect the vertices by edges perpendicular to those faces. This way, the Cayley graph is W-equivariantly embedded in $\mathbb{E}$.*

Example 9.1.11. Here are the planar reflection groups whose fundamental chambers are triangles:

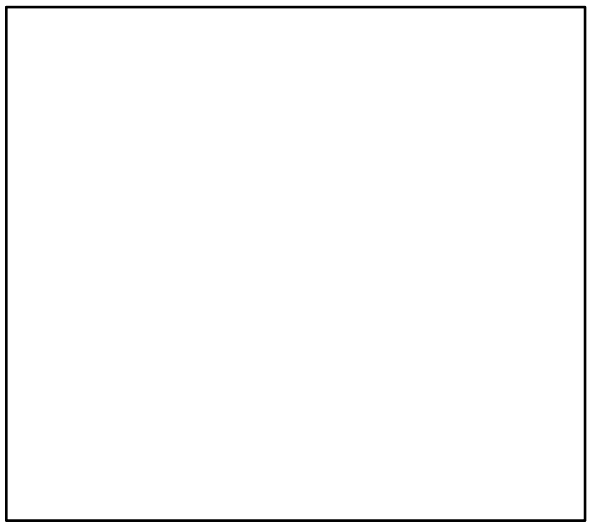

9.1.2 The Coxeter Matrix

The Coxeter Matrix of the pair (W, S) is the $S \times S$-matrix

$$M := \left(m_{s,t} := \mathrm{ord}_W(st)\right)_{s,t \in S}.$$

The entries are taken from $\{1, 2, 3, \ldots, \infty\}$. Note that M is symmetric and satisfies:

$$m_{s,t} = 1 \quad \text{if and only if} \quad s = t. \tag{9.1}$$

Theorem 9.1.12. *The group W has the presentation*

$$W = \left\langle s \in S \mid (st)^{m_{s,t}} = 1 \text{ for } m_{s,t} < \infty \right\rangle.$$

Proof. The given relations obviously hold. To deduce any given other relation, realize the relation as a closed loop in the Cayley graph. This graph lies in the ambient Euclidean space. Find a bounding disk that intersects the CoDim-2-skeleton of the chamber decomposition transversally. Now see the van Kampen diagram. **q.e.d.**

For each s let $\mathbf{u}_s$ be the unit vector perpendicular to the hyperplane inducing the reflection s. (There is a choice here: we use the vector that points away from the fundamental chamber.)

Exercise 9.1.13. Show that for any $s, t \in S$,

$$\langle \mathbf{u}_s, \mathbf{u}_t \rangle = \begin{cases} -\cos\left(\frac{\pi}{m_{s,t}}\right) & \text{for } m_{s,t} \text{ finite} \\ -1 & \text{for } m_{s,t} \text{ infinite.} \end{cases}$$

Now, we can settle the question, whether S is finite.

Proposition 9.1.14. *The fundamental chamber has finite support.*

Proof. Suppose otherwise. Then the set of unit vectors $\mathbf{u}_s$ had an accumulation point by compactness of the unit sphere. However, their pair-wise scalar products are negative. **q.e.d.**

Corollary 9.1.15. *The set $\mathcal{H}$ decomposes into fintely many parallelity classes.*

Proof. Suppose otherwise, then, by compactness, there would be hyperplanes that span arbitrary small angles. Take a point very close to their intersection that lies in a chamber. Since the angles around faces of chambers are bounded away from 0, we have a contradiction. **q.e.d.**

Exercise 9.1.16. Show that the following are equivalent:

1. $\mathcal{H}$ is finite.
2. W is finite.
3. W is torsion.
4. $\bigcap_{H\in\mathcal{H}} H \neq \emptyset$.

Corollary 9.1.17. *A Euclidean reflection group W is virtually free abelian.*

Proof. Consider the action of W upon the sphere at infinity. By (9.1.15), this sphere is decomposed into finitely many regions, upon which W acts by spherical isometries. The image of W in $\mathrm{Isom}(S)$ is a finite Euclidean reflecion group by (9.1.16). The kernel of the homomorphism consists of translations. **q.e.d.**

9.1.3 The Cocompact Case

In this section, we assume that the fundamental chamber has compact closure. All the result are valid in the general case, though. In deed, we will prove them for arbitrary Coxeter groups later.

Observation 9.1.18. *Every point of $\mathbb{E}$ is either contained in a chamber or belongs to the closures of at least two adjacent chambers. In the latter case, it has a translate in the closure of C^*. Thus, the closure of the fundamental chamber is a fundamental domain for the action of W, i.e., the translates of the closure cover $\mathbb{E}$ while the translates of C^* stay disjoint.*

Theorem 9.1.19. *W has only finitely many finite subgroups up to conjugacy.*

Proof. A finite subgroup fixes a point. This point is a translate of some point in $\overline{C^*}$. Thus any finite subgroup is conjugate to a subgroup of a stabilizer of a point in $\overline{C^*}$. There are only finitely many of those since C^* has only finitely many faces. **q.e.d.**

Theorem 9.1.20. *The conjugacy problem in W is solvable.*

Proof. !!! Do the CAT(0) proof !!! **q.e.d.**

9.2 Coxeter Groups

Definition 9.2.1. Let S be a set. A Coxeter matrix over S is a symmetric matrix $M = (m_{s,t})_{s,t\in S}$ with entries $m_{s,t}$ in $\{1, 2, 3, \ldots, \infty\}$ such that

$$m_{s,t} = 1 \text{ if and only if } s = t.$$

The Coxeter group defined by M is the group given by the presentation

$$W = \langle s \in S \mid (st)^{m_{s,t}} = 1 \text{ if } m_{s,t} \text{ finite} \rangle .$$

The pair (W, S) is called a Coxeter system.

Example 9.2.2. Every Euclidean reflection group is a Coxeter group.

Coxeter groups are defined by generators and relations. In general, it is hard to tell wheter a group given in this manner is trivial or not. So our first problem will be to see that Coxeter groups are not trivial.

Observation 9.2.3. *Every defining relation of W has even length. Thus, there is a well defined surjective homomorphism*

$$W \to C_2$$

sending each generator in S to the generator of C_2. In particular, none of the generators is trivial in W. **q.e.d.**

Thus, every generator generates a subgroup of order 2 inside W.

9.2.1 The Geometric Representation

To show that the generators have order 2, we used a representation of W. Now, we shall extend this method to show that the products st also have the orders that we would expect from the presentation.

Definition 9.2.4. Let (W, S) be a Coxeter system with Coxeter matrix M. Let $V := \bigoplus_{s \in S} \mathbb{R}e_s$ be the real vector space generated by S: To avoid confusion, we denote the basis vector corresponding to s by e_s.

Define a bilinear form on V by

$$\langle e_s, e_t \rangle_M := \begin{cases} -\cos\left(\frac{\pi}{m_{s,t}}\right) & \text{if } m_{s,t} < \infty \\ -1 & \text{if } m_{s,t} = \infty \end{cases}$$

and define an action of W on V where the generator s acts as the linear automorphism

$$\rho_s : e_t \mapsto e_t - 2 \langle e_s, e_t \rangle_M e_s.$$

This action defines the geometric representation

$$\rho : W \to \mathrm{Aut}(V).$$

Exercise 9.2.5. Check that the geometric representation does exist, i.e., check that the automorphisms ρ_s satisfy the defining relations of W.

Lemma 9.2.6. *The order of st in M is given by the entry $m_{s,t}$ of the Coxeter matrix.*

Proof. Note that the action of the subgroup $\langle s, t \rangle$ leaves the subspace $V_{s,t} := \langle e_s, e_t \rangle$ invariant.

$m_{s,t} = \infty$: The action hits e_t as follows:

$$e_t \xrightarrow{\rho_s} e_t + 2e_s \xrightarrow{\rho_t} 3e_t + 2e_s \xrightarrow{\rho_s} 3e_t + 4e_s \xrightarrow{\rho_t} 5e_t + 4e_s \xrightarrow{\rho_s} \cdots$$

Thus, the product $\rho_t \rho_s$ has infinite order.

$m_{s,t} < \infty$: In this case, the bilinear form $\langle -, - \rangle_M$ restricts to a positive definite bilinear form on $V_{s,t}$, and a direct computation shows that the product $\rho_t \rho_s$ is a rotation about of order $m_{s,t}$. **q.e.d.**

Corollary 9.2.7. *Thus, the generators s and t span a copy of the dihedral group $D_{m_{s,t}}$ inside W.* **q.e.d.**

Exercise 9.2.8. Show that W is finite if the bilinear form $\langle -, - \rangle_M$ is positive definite.

Exercise 9.2.9. Show that if W is finite, then there is a unique bilinear form $\langle -, - \rangle$ on V characterized by the following properites

1. $\langle -, - \rangle$ is positive definite.
2. All basis vectors e_s have unit length.
3. The action of W preserves $\langle -, - \rangle$.

Moreover, this bilinear form is $\langle -, - \rangle_M$.

Corollary 9.2.10. *Finite Euclidean reflection groups and finite Coxeter groups are the very same thing.*

Remark 9.2.11. The classification of finite Coxeter groups is done by classifying all Coxeter matrices that are positive definite.

Exercise 9.2.12. A Coxeter system is called irreducible if there is no non-trivial partition of the generators into two sets S_1 and S_2 such that each generator from S_1 commutes with each generator from S_2. Classify all irreducible Coxeter systems over three generators whose Coxeter groups are finite. (Hint: You should recover descriptions of the Platonic solids along the way; in fact, the existence of the Platonic solids can be derived from this classification.)

9.2.2 The Geometry of a Coxeter System

We studied Euclidean reflection groups by means of the assiciated Chamber system upon which the group acts. To study general Coxeter groups, we will construct the geometry from the group. So, we will construct a chamber system from the (reduced) Cayley complex $\Gamma_S(W)$ for the Coxeter presentation. The vertices of the Cayley complex are the chambers, and two chambers are s-adjacent if they are joined by an edge with labes s. Of course an edge path in the Cayley complex is a gallery in the chamber system. We will see that this chamber system allows reflections and half spaces.

Definition 9.2.13. Two edges e and $e_\bullet$ in $\Gamma_S(W)$ are opposite if they are contained in a relator disc and have maximal distance in this circle. We write $e \longleftrightarrow e_\bullet$. The edges e and $e_\bullet$ are parallel if $e = e_\bullet$ or if there is a finite sequence

$$e = e_0 \longleftrightarrow e_1 \longleftrightarrow e_2 \longleftrightarrow \cdots \longleftrightarrow e_r = e_\bullet.$$

We write $e \parallel e_\bullet$.

Parallelity is an equivalence relation. Its equivalence classes are called walls.

It is useful to extend the notion of parallelism to oriented edges. Let us consider opposite edges first. Inside a relator disc, an oriented edge induces an orientation of the boundary circle of the disc. We call two oriented edges of a relator disc opposite, if they induce opposite orientations of the boundary circle and their underlying geometric edges are opposite. As above, parallelism is defined as the transitive closure of opposition. Equivalence classes of oriented edges under parallelism are called oriented walls

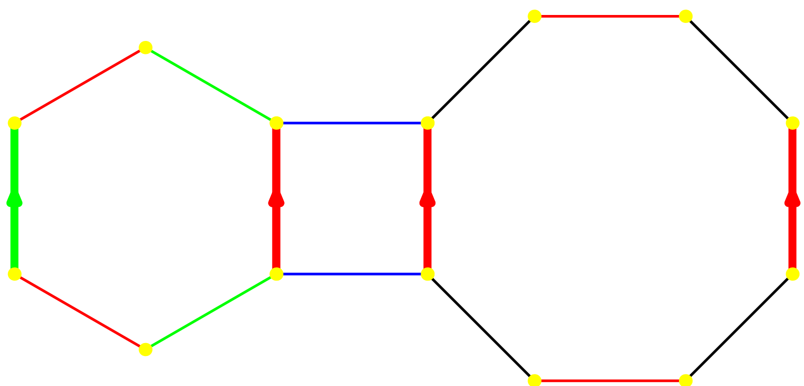

Observation 9.2.14. *Let $\vec{e}_0$ and $\vec{e}_1$ are opposite oriented edges in a relator cell. Then removing these two edges cuts the boundary circle of the relator disc into two arcs; and the arc from $\iota(\vec{e}_0)$ to $\iota(\vec{e}_1)$ reads the same word as the arc from $\tau(\vec{e}_0)$ to $\tau(\vec{e}_1)$. (Here, we use that relator discs are not "crushed".)*

By induction, it follows that if $\vec{e}$ and $\vec{e}_\bullet$ are parallel, then there is a group element $w \in W$ such that

$$\iota(\vec{e}) = \iota(\vec{e}_\bullet)\, w \text{ and } \tau(\vec{e}) = \tau(\vec{e}_\bullet)\, w.$$

Note that we are multiplying from the right, which in general will tear edges apart.

Observation 9.2.15. *Let us observe a "local converse": Each relator disc is in its own right the Cayley graph of a finite dihedral group. Let $\vec{e}$ and $\vec{e}_\bullet$ be two oriented edges in this cell. If there is a group element w in the dihedral group such that*

$$\iota(\vec{e}) = \iota(\vec{e}_\bullet)\, w \text{ and } \tau(\vec{e}) = \tau(\vec{e}_\bullet)\, w$$

then $\vec{e}$ and $\vec{e}_\bullet$ are either opposite or identical.

Corollary 9.2.16. *An (oriented) wall either avoids a relator cell or intersects it in a pair of opposite (oriented) edges.* **q.e.d.**

Proof. Let $\vec{e}$ be an oriented edge in a relator cell. We have to show that the only parallel edge in this cell is the opposite one. So let $\vec{e}_\bullet$ be any other parallel edge in this relator cell. We know that there is an element $w \in W$ such that

$$\iota(\vec{e}) = \iota(\vec{e}_\bullet)\, w \text{ and } \tau(\vec{e}) = \tau(\vec{e}_\bullet)\, w.$$

Since both edges belong to the relator cell, the element w actually belongs to the dihedral subgroup generated by the two labels around the relator cell. Now, it follows from the local converse that $\vec{e}$ and $\vec{e}_\bullet$ are opposite or identical. **q.e.d.**

The same reasoning actually yields:

Corollary 9.2.17. *Let e be an edge and let $\vec{e}$ and $\mathrm{op}\vec{e}$ denote the two corresponding oriented edges. Then $\vec{e}$ and $\mathrm{op}\vec{e}$ are not parallel.*

In particular, every wall can be oriented in precisely two ways. **q.e.d.**

Let H be a wall. The boundary $\partial(H)$ of H is the set of vertices (chambers) that are incident with at least one edge in H – recall that a wall is an equivalence class of edges.

Example 9.2.18. Here is the Cayley graph for the group $\left\langle \mathbf{b}, \mathbf{g}, \mathbf{r} \,\middle|\, \mathbf{b}^2 = \mathbf{g}^2 = \mathbf{r}^2 = (\mathbf{br})^3 = (\mathbf{bg})^2 = 1 \right\rangle$. drawn in the hyperbolic plane. Oberserve how the axes for the reflections intersect groups of two or four edges perpendicularly. These are precisely the walls. The shaded regions are the relator discs.

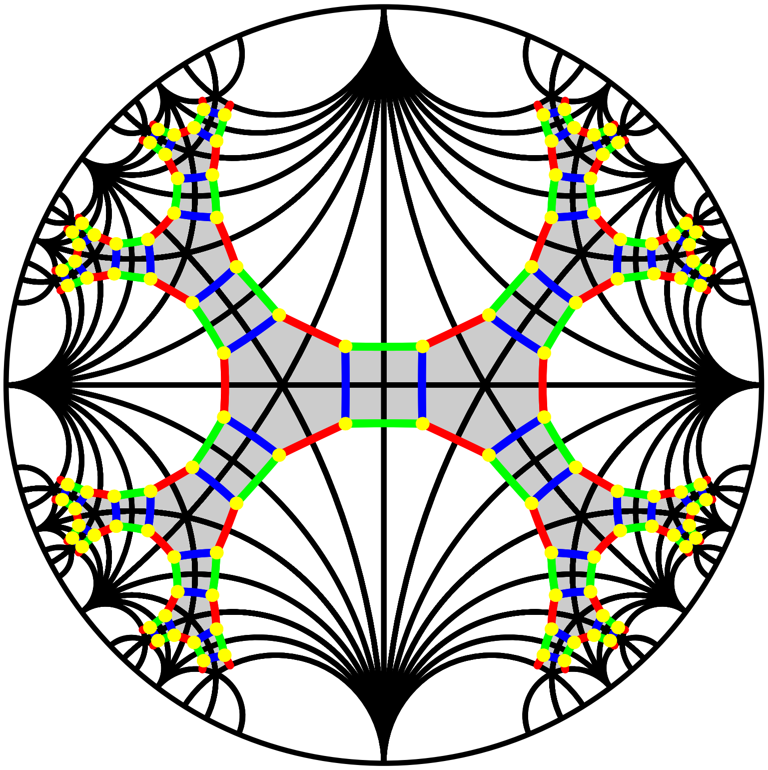

Let H be a wall and let $\mathbf{g}$ be an edge path (a gallery). Let

- hits $(H, \mathbf{g})$ denote the number of times that the edge path $\mathbf{g}$ passes through the wall H.

Definition 9.2.19. An elementary homotopy of an edge path in the Cayley graph of a Coxeter group is one of the following two types of moves:

1. Replacing a subpath reading part of a relator disc by the complementary part of the relation.

2. Adding or removing a backtracking edge.

Two paths in the Cayley graph are called homotopic if one can be obtained from the other by a finite sequence of elementary homotopy.

Observation 9.2.20. *Since elementary homotopies correspond to substitutions in words waranted by the defining relations, two galleries are homotopic if and only if they connect the same end points.*

Observation 9.2.21. *Given a wall and a path, the number of crossings between the wall and the gallery changes by an even number during any elementary homotopy of the gallery. This follows from (9.2.16) Thus for a given wall H and two galleries $\mathbf{g}_0$ and $\mathbf{g}_1$, we have*

$$\operatorname{hits}(H, \mathbf{g}_0) \equiv \operatorname{hits}(H, \mathbf{g}_1) \mod 2.$$

In particular, the endpoints of an edge inside H cannot be connected by a gallery that does not cross H.

Corollary 9.2.22. *Every wall separates Γ.*

Observation 9.2.23. *If $\vec{e}$ and $\vec{e}_\bullet$ are parallel their terminal vertices can be joined by a path that does not intersect the wall they belong to. This follows by induction from the corresponding statement about opposite oriented edges, which is obvious.*

Corollary 9.2.24. *Each wall separates the Cayley graph into precisely two half spaces.*

Proof. We already know that walls separate. That there are not more than two components follows from (9.2.23). **q.e.d.**

Lemma 9.2.25. *Associated to each wall, there is a unique element in W that acts like a reflection along the wall.*

Proof. Let e be an edge in the wall. Then there is a unique element in W that interchanged its endpoints. (This is, indeed, true for any edge: There is a unique group element taking the initial point to the terminal point. But then, it has to swap the two points, because the action of W preserves the labelling of edges by generators.)

Now, just check that this swap condition extends to edges that are opposite in a relator disc. **q.e.d.**

Corollary 9.2.26. *Half spaces are convex, i.e., if two chambers lie in a given half space, then so does every minimal chamber between them.* **q.e.d.**

Corollary 9.2.27. *The gallery distance of two chambers is the number of walls seperating them.* **q.e.d.**

Definition 9.2.28. A morphism of graphs is a distance non-increasing map from the vertices of graph to the vertices of another graph. A folding of a graph is an idempotent graph endomorphism $f : \Gamma \to \Gamma$ such that the preimage of each vertex v is either empty or contains precisely two vertices (one of which is v). The image

α_f of a folding is called a half space or a root. Two foldings f and f' are opposite if their images are disjoint and the following hold:

$$\begin{aligned} f &= f \circ f' \\ f' &= f' \circ f. \end{aligned}$$

Any two opposite foldings f and f' induce a reflection

$$\begin{aligned} \rho : \Gamma &\rightarrow \Gamma \\ v &\mapsto \begin{cases} f(v) & \text{if } v \in \alpha_{f'} \\ f'(v) & \text{if } v \in \alpha_f \end{cases} \end{aligned}$$

Exercise 9.2.29. Show that a (locally finite) graph is the Cayley graph of a (finitely generated) Coxeter group if and only if the following conditions holds:

1. For each oriented edge $\vec{e}$ there is a unique folding $f_{\vec{e}}$ of Γ satisfying $f_{\vec{e}}(\iota(\vec{e})) = \tau(\vec{e})$.
2. If $\vec{e}$ and $\mathrm{op}\vec{e}$ are opposite orientations of the same underlying geometric edge, then $f_{\vec{e}}$ and $f_{\mathrm{op}\vec{e}}$ are opposite foldings.

Above, we introduced the geometric representation of W on the vector space V spanned by $\{e_s \mid s \in S\}$. Let V^* be the dual of V. It turns out that the induced action of W on V^*,

$$\begin{aligned} \tau : W &\rightarrow \mathrm{Aut}(V^*) \\ w : \lambda &\mapsto \lambda \circ \rho_w, \end{aligned}$$

gives another description of the chamber system: For any s, define the posite and negative halfspace in V^* by

$$\begin{aligned} U_s^+ &:= \{\lambda \in V^* \mid \lambda(e_s) > 0\} \\ U_s^- &:= \{\lambda \in V^* \mid \lambda(e_s) < 0\}. \end{aligned}$$

The Tits cone

$$C := \{\lambda \in V^* \mid \lambda(e_s) > 0 \text{ for all } s \in S\}$$

is the intersection of the positive cones.

Exercise 9.2.30. Show that for every $w \in W$,

$$\tau_w(C) \subseteq U_s^+ \text{ if and only if } |sw| = |w| + 1$$

and

$$\tau_w(C) \subseteq U_s^- \text{ if and only if } |sw| = |w| - 1.$$

Exercise 9.2.31. Infer from (9.2.30) that the geometric representation is faithful.

Corollary 9.2.32. *Finitely generated Coxeter groups are linear.* **q.e.d.**

9.2.3 The Deletion Condition

In (9.1.8), we have seen, that the pair (W, S) for a Euclidean reflection group satisfies the Deletion Condition:

Definition 9.2.33 (Deletion Condition). Let (W, S) be a pair where W is a group and S is a generating set for W consisting entirely of elements of order 2. We say that this pair satisfies the Deletion Contition if:

> For any non-reduced word $s_1 \cdots s_r$ over S there are two indices i and j such that
> $$s_1 \cdots s_r =_W s_1 \cdots \hat{s}_i \cdots \hat{s}_j \cdots s_r.$$
> The carets indicate ommision.

This is, one can delete two letters from any non-minimum-length word to obtain a shorter representative for the same element of W.

In this section, we will recognize (W, S) as a Coxeter system using the Deletion Condition.

Lemma and Definition 9.2.34 (Exchange Condition). *The pair (W, S) satifies the Exchange Condition, i.e.:*

> *Let $s_1 \cdots s_r$ and and $t_1 \cdots t_r$ be two reduced words over S representing the same element $w \in W$. If $s_1 \neq t_1$, then there is an index $i \in \{2, \ldots, r\}$ such that*
> $$w =_W s_1 t_1 \cdots \hat{t}_i \cdots t_r.$$

Proof. This is a formal consequence of the Deletion Condition: From

$$s_1 \cdots s_r =_W t_1 \cdots t_r,$$

we obtain

$$s_2 \cdots s_r =_W s_1 t_1 \cdots t_r$$

where the right hand is longer than the left hand whence there must be a pair of letters that can be dropped without changing the value of the product. However, one of the two letters must be the leading s_1: Otherwise, we had

$$s_2 \cdots s_r =_W s_1 t_1 \cdots \hat{t}_i \cdots \hat{t}_i \cdots t_r$$

whence

$$s_1 \cdots s_r =_W t_1 \cdots \hat{t}_i \cdots \hat{t}_i \cdots t_r$$

contradicting the minimality of the initial words.

Thus, we have

$$s_2 \cdots s_r =_W t_1 \cdots \hat{t}_i \cdots \cdots t_r$$

whence

$$s_1 \cdots s_r =_W s_1 t_1 \cdots \hat{t}_i \cdots \cdots t_r. \qquad \textbf{q.e.d.}$$

The Coxeter Matrix of the pair (W, S) is the $S \times S$-matrix

$$M := (m_{s,t} := \mathrm{ord}_W(st))_{s,t \in S}.$$

The entries are taken from $\{1, 2, 3, \ldots, \infty\}$. Note that M is symmetric and satisfies:

$$m_{s,t} = 1 \quad \text{if and only if} \quad s = t. \tag{9.2}$$

Any symmetric matrix satisfying (9.2) is called a Coxeter matrix.

An elementary M-reduction is one of the following moves:

1. Delete a subword ss.
2. Replace a subword $\underbrace{sts\cdots}_{m_{s,t}\text{ letters}}$ by $\underbrace{tst\cdots}_{m_{s,t}\text{ letters}}$.

Theorem 9.2.35 (Tits). *Let $\underline{s} = s_1 \cdots s_{|\underline{s}|}$ be a reduced word over S. Then $\underline{s}$ can be obtained from any word $\underline{t} = t_1 \cdots t_{|\underline{t}|}$ by a sequence of elementary M-reductions.*

Proof. This is also a purely formal consequence of the Deletion Condition. Let us first prove the theorem under the additional hypothesis that $\underline{t}$ is reduced, as well. In this case, $|\underline{s}| = |\underline{t}|$ and only moves of type (2) are possible. We induct on the length of the words.

Assume first that $s_1 = t_1$. Then $s_2 \cdots s_{|\underline{s}|}$ and $t_2 \cdots t_{|\underline{s}|}$ are two reduced words representing the same group element. By induction, we can pass from one to the other by elementary M-reductions.

So assume $s_1 \neq t_1$. So we can apply the exchange condition both ways and obtain

$$\begin{aligned} s_1 \cdots s_{|\underline{s}|} &=_W s_1 t_1 \cdots \hat{y}_i \cdots t_{|\underline{s}|} \\ t_1 \cdots t_{|\underline{s}|} &=_W t_1 s_1 \cdots \hat{x}_i \cdots s_{|\underline{s}|} \end{aligned}$$

Note that both equations actually can be realized by M-reduction since the words start with identical letters. Thus, we only have to realize an M-reduction to pass from $s_1 t_1 \cdots \hat{y}_i \cdots t_{|\underline{s}|}$ to $t_1 s_1 \cdots \hat{x}_i \cdots s_{|\underline{s}|}$. If $m_{s,t} = 1$, we are done. Otherwise we apply the exchange condition again:

!!! ... !!! (Finish this)

Now let us drop the assumption that $\underline{t}$ is reduced. It suffices to prove that $\underline{t}$ can be shortened by M-reductions. We induct on the length of $\underline{t}$. If $\underline{t}_2 \cdots \underline{t}_{|\underline{t}|}$ is not reduced, we apply the induction hypothesis to this subword.

So we assume that $\underline{t}_2 \cdots \underline{t}_{|\underline{t}|}$ is reduced. Then we find

$$\underline{t}_1 \cdots \underline{t}_{|\underline{t}|} =_W \underline{t}_2 \cdots \hat{t}_i \cdots \underline{t}_{|\underline{t}|}$$

whence $\underline{t}_2 \cdots \underline{t}_{|\underline{t}|}$ can be transformed into $t_1\underline{t}_2 \cdots \hat{t}_i\underline{t}_{|\underline{t}|}$ by M-reductions. (Both of these words are reduced, so we are in the case that we have discussed already.) Now, we can shorten:

$$\underline{t}_1 \cdots \underline{t}_{|\underline{t}|} \xrightarrow{M} \underline{t}_1 t_1 \underline{t}_2 \cdots \hat{t}_i \underline{t}_{|\underline{t}|} \xrightarrow{M} \underline{t}_2 \cdots \hat{t}_i \underline{t}_{|\underline{t}|}.$$

The final step is an operation of type (1). **q.e.d.**

Corollary 9.2.36. *The pair (W, S) is a Coxeter system.*

Proof. A relation is a word that evaluates to 1 in W. Therefore, any relation can be transformed into the empty word by M-reductions. However, these correspond to the relations of the Coxeter presentation. **q.e.d.**

9.2.4 The Moussong Complex

The goal in this section is to describe a piecewise Euclidean CAT(0) complex upon which the Coxeter group W acts cocompactly, properly, and discontinuously. We can find such a complex, provided the generating set S of reflections is finite. The existence of such a complex settles a lot of questions at once:

Corollary 9.2.37. *For every finitely generated Coxeter group the following hold:*

1. *W has solvable conjugacy problem.*
2. *W has only fintely many conjugacy classes of finite subgroups.*

The construction start with the Cayley graph. For any subset $J \subseteq S$, we define J-residues to be the components of the Cayley graph after all edges whose labels are *not* in J have been removed. So we restrict ourselves to edges with labels in J and look at the connected components of the resulting graph. By (9.2.10), every finite Coxeter group W_J is a Euclidean reflection group acting on some Euclidean space $\mathbb{E}$. The reflections are induced by finitely many hyperplanes that all pass through a common point. The hyperplanes chop up $\mathbb{E}$ into chambers. In one of these chambers find a point that has distance $\frac{1}{2}$ from all the walls. The orbit of this points spans a convex polyhedron P_J all of whose edges have length 1. Indeed, the 1-skeleton of this convex polyhedron is a Cayley graph for the finite Coxeter group W_J. The following exercise justifies all these claims.

Exercise 9.2.38. For any subset $J \subseteq S$, let M_J be the submatrix of M whose rows and columns have indices in J, and let (W_J, J) be the Coxeter system defined by M_J. Prove:

1. The inclusion $J \hookrightarrow S$ induces an injective group homomorphism $W_J \to W$ that identifies the group W_J with the subgroup of W generated by $J \subseteq S$.

2. in view of the preceeding result, we regard W_J as a subgroup of W. These subgroups are called special parabolic subgroups. Prove that

$$W_{J \cap I} = W_J \cap W_I$$

   for any two subsets $J, I \subseteq S$.

3. The J-residues in Γ_W are in bijective correspondence to the left cosets of W_J.

4. Every J-residue is isomorphic to the Cayley graph of W_J with respect to the generating set J.

Example 9.2.39. Here is the polyhedron for

$$\left\langle \mathbf{b}, \mathbf{g}, \mathbf{r} \,\middle|\, \mathbf{b}^2 = \mathbf{g}^2 = \mathbf{r}^2 = (\mathbf{bg})^3 = (\mathbf{br})^3 = (\mathbf{gr})^2 = 1 \right\rangle,$$

which is the symmetric group on four letters:

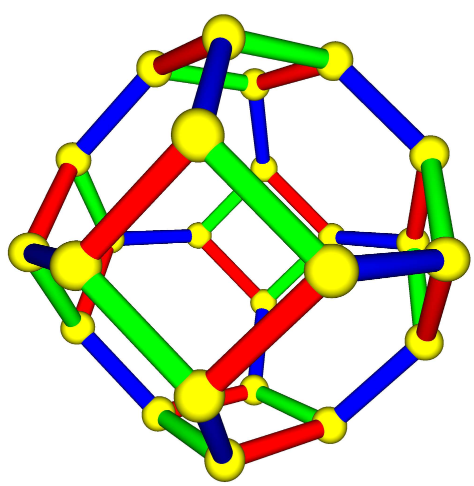

Note how the faces correspond to cosets of special parabolic subgroups.

The Moussong complex X for the Coxeter system (W, S) is defined by the following procedure:

- Start with with the Cayley graph, and declare all edges to be of length 1. Observe that the edges with label s correspond precisely to the $\{s\}$-residues.

- Construct the 2-skeleton by glueing in polygons $P_{\{s,t\}}$ for any pair $\{s,t\}$ of generators that generate a finite subgroup. More precisely, if the $\{s,t\}$-residues are finite then $P_{\{s,t\}}$ is a polygon whose boundary is isomorphic to these residues. The isomorphism induces attaching maps that we use to glue in one copy of $P_{\{s,t\}}$ for each residue.

- The 3-skeleton is defined similarly. For every $J \subseteq S$ of size three, we glue in copies of P_J if the J-residues are finite. Note that the boundary sphere of P_J consists of polygons that are isomorphic to the cells P_I for strict subsets $I \subset J$. Thus, we find the boundary spheres of our 3-cells in the 2-skeleton that has been constructed already.

- Proceed on higher skeleta until every finite residue is geometrically realized.

Observation 9.2.40. *The Moussong complex carries a natural piecewise Euclidean structure: all its cells are convex polyhedra in Euclidean space, and all attaching maps identify lower dimensional cells isometrically with faces of higher dimensional cells.*

Observation 9.2.41. *The* 1*-skeleton of* X *is the Cayley graph. The* 2*-skeleton is the* Cayley complex *of the Coxeter presentation for* W*: The* 2*-cells in* X *are precisely the cells whose boundaries read valid relations in* W*. It follows that* X *is simply connected.*

Corollary 9.2.42. *To prove* X *to be* $CAT(0)$ *it suffices to show that vertex links in* X *are* $CAT(1)$ *since* X *is piecewise Euclidean and simply connected.*

Example 9.2.43. Here is a (distorted picture of) the Moussong complex for the group

$$\left\langle \mathbf{b}, \mathbf{g}, \mathbf{r} \,\middle|\, \mathbf{b}^2 = \mathbf{g}^2 = \mathbf{r}^2 = (\mathbf{br})^3 = (\mathbf{bg})^2 = 1 \right\rangle .$$

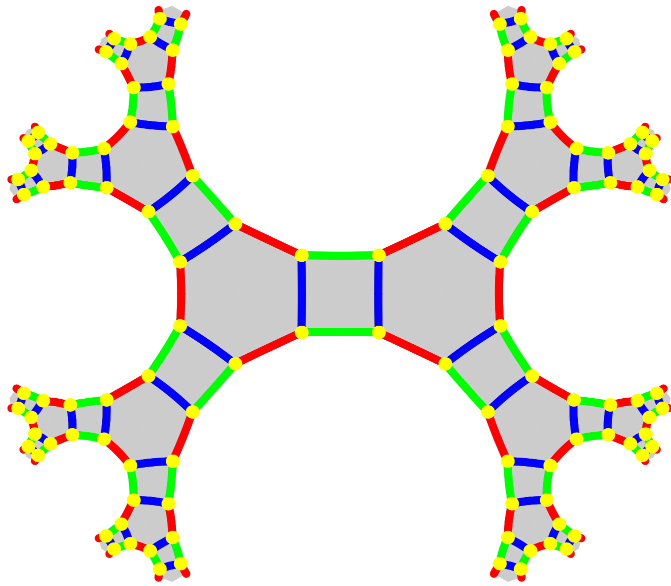

The grey shaded area consists of hexagons and squares that are glued in. All these polygons are regular Euclidean polygons.

Observation 9.2.44. *Since W acts transitively on the set of vertices, all vertex links are isometric.*

Theorem 9.2.45. *The vertex link L of X is CAT(1).*

Proof. We give an explicit description of the link as a piecewise spherical complex: The vertex set of L is S. For every subset $J \subseteq S$ that generates a finite subgroup in W, we glue in a spherical simplex whose edge lengths are

$$d(s,t) = \pi - \frac{\pi}{m_{s,t}}.$$

We have to show that the resulting complex is metrically flag, i.e., if we find a subset $J \subseteq S$ such that all elements are joined by an edge (equivalently, $m_{s,t}$ is finite), then this subset should generate a finite subgroup if the edge lengths can be realized by a spherical simplex.

So suppose $\{\mathbf{u}_s \mid s \in J\}$ is a collection of unit vectors whose distances realize the edge lengths. Then

$$\langle \mathbf{u}_s, \mathbf{u}_t \rangle = \cos\left(\pi - \frac{\pi}{m_{s,t}}\right) = -\cos\left(\frac{\pi}{m_{s,t}}\right)$$

which is precisely the coefficient in the bilinear form $\langle -, - \rangle_{M_J}$, which therefore is positive definite. By (9.2.8), this implies that W_J is finite as required. **q.e.d.**

9.3 Artin Groups

!!! FIXME: Rewrite this. More pictures, include Garside structures and the recent Eilenberg-Maclane spaces based on Bestvina's greedy normal forms. !!!

Let M be a Coxeter matrix with index set S. The Artin group defined by M is given by the presentation:

$$A_M := \left\langle s \in S \;\middle|\; \underbrace{sts\cdots}_{m_{s,t}\text{ factors}} = \underbrace{tst\cdots}_{m_{s,t}\text{ factors}} \right\rangle .$$

The Coxeter matrix M defines a Coxeter group W_M at the same time. The canonical homomorphism

$$A_M \rightarrow W_M$$

is surjective. An Artin group is said to be of finite type if the associated Coxeter group is finite.

Remark 9.3.1. Sometimes a group G is called of finite type or of type F if it has a finite Eilenberg-Maclane complex. Therefore the statement

> Artin groups of finite type are of finite type.

is actually meaningfull. It happens to be true.

9.3.1 The Braid Group

!!! This whole section needs PICTURES !!!

Configuration Spaces as Hyperplane Arrangements

The labeled configuration space of n points in the plane is

$$\tilde{C}_n := \{(z_1, \ldots, z_n) \in \mathbb{C}^n \mid z_i \neq z_j \text{ for } i \neq j\} .$$

An element in this space is a set of n points in the plane that are labeled so that we can tell them apart. The symmetric group Perm_n on n letters acts on these configurations by permuting the labels. Hence the quotient

$$C_n := {}_{\mathrm{Perm}_n}\backslash^{\tilde{C}_n}$$

is the configuration space of n-point subsets in the plane.

Definition 9.3.2. The braid group B_n is the fundamental group of C_n. The pure braid group P_n is the fundamental group of $\tilde{C}_n$.

Observation 9.3.3. *The projection*

$$\pi : \tilde{C}_n \to C_n$$

is a covering map with Perm_n *acting as its group of deck transformations. Consequently, we have a short exact sequence*

$$P \hookrightarrow B_n \twoheadrightarrow \mathrm{Perm}_n$$

of groups. In particular, the pure braid group is a finite index normal subgroup of the braid group.

Our first goal is to prove that configuration spaces are Eilenberg-Maclane spaces for braid groups. Later, we will find smaller Eilenberg-Maclane spaces.

Theorem 9.3.4 (Fadell-Neuwirth 1962 [FaNe62, Corollary 2.2]). *The space* $\tilde{C}_n$ *is a* $\mathrm{K}(P_n, 1)$*. Consequently,* C_n *is a* $\mathrm{K}(B_n, 1)$*.*

We will follow the proof in [FaNe62].

For any finite set $P \subset \mathbb{C}$ of punctures, put

$$\tilde{C}_{P,n} := \{(z_1, \ldots, z_n) \mid z_i \notin P \text{ and } z_i \neq z_j \text{ for } i \neq j\}.$$

This is the configuration space of n labeled points in a plane with $m := |P|$ punctures. Note that up to homeomorphism, the position of the punctures does not matter since all m-punctures planes are homeomorphic.

Fact 9.3.5. *The map*

$$\begin{array}{rcl} \pi : \tilde{C}_{P,n} & \to & \mathbb{C} - P \\ (z_1, \ldots, z_n) & \mapsto & z_1 \end{array}$$

is a fibre bundle whose fibre over $z \in \mathbb{C} - P$ *is* $\tilde{C}_{P\cup\{z\},n-1}$*.*

This fact allows us to "freeze" the points of the configuration one by one: Since fibre bundles are fibrations, we have a long exact sequence of homotopy groups

$$\cdots \to \pi_m\left(\tilde{C}_{P\cup\{z\},n-1}\right) \to \pi_m\left(\tilde{C}_{P,n}\right) \to \pi_m(\mathbb{C} - P) \to \pi_{m-1}\left(\tilde{C}_{P\cup\{z\},n-1}\right) \to \cdots$$

which proves

$$\pi_m\left(\tilde{C}_{P\cup\{z\},n-1}\right) = \pi_m\left(\tilde{C}_{P,n}\right) \text{ for } m \geq 2$$

since $\mathbb{C} - P$ has trivial homotopy groups in dimension 2 and above. Applying this observation repeatedly, we conclude that for $m \geq 2$:

$$0 = \pi_m\left(\tilde{C}_{\{z_1,\ldots,z_n\},0}\right) = \pi_m\left(\tilde{C}_{\{z_1,\ldots,z_{n-1}\},1}\right) = \cdots = \pi_m\left(\tilde{C}_{\{z_1\},n-1}\right) = \pi_m\left(\tilde{C}_{\emptyset,n}\right).$$

This proves (9.3.4) as $\tilde{C}_{\emptyset,n} = \tilde{C}_n$.

Shrinking the Eilenberg-Maclane Space

The space of all configuration deformation retract onto the subspace $\tilde{C}_n^0$ of all those configuration whose center of gravity is 0. Note that the symmetric group Perm_n acts on $V = \{(t_1, \ldots, t_n) \mid \sum_i t_i = 0\} \leq \mathbb{R}^n$ by permuting the coordinates. This is, in fact, the geometric representation of Perm_n as a finite reflection group. Decomposing the n-tuples in $\tilde{C}_n^0$ into real and imaginary parts, we obtain

$$\tilde{C}_n^0 = V \times V - \bigcup_{H \in \mathcal{H}} H \times H$$

where $\mathcal{H}$ is the set of walls defining S as a finite reflection group on V.

Let X be the Moussong comlpex associated to Perm_n. Recall that this is a convex polyhedrong in V given as the convex hull of a point chosen in a sector such that it has distance $\frac{1}{2}$ to all walls bounding its sector. Shrinking configurations if necessary by rescaling them using a real scalar yields a deformation retraction of $\tilde{C}_n^0$ onto

$$Y_n := X \times X - \bigcup_{H \in \mathcal{H}} H \times H.$$

Note that Y_n is an Eilenberg-Maclane space for the pure braid group. Let us define a poset

$$\mathcal{A}_n := \{(c, v) \mid c \text{ cell in } X,\ v \text{ vertex in } c\}$$

where the order is given by

$$(c, v) \preceq (d, w) \text{ if and only if } c \leq d \text{ and } v = \pi_c(w)\,.$$

We will prove

Lemma 9.3.6. *There is a cover $Y_n = \bigcup_{\alpha \in \mathcal{A}_n} U_\alpha$ by convex open sets indexed by the element of $\mathcal{A}_n$ such that for any subset $\sigma \subset \mathcal{A}_n$,*

$$U_\sigma := \bigcap_{\alpha \in \sigma} U_\alpha \neq \emptyset \text{ if and only if } \sigma \text{ is a chain in } \mathcal{A}_n.$$

Corollary 9.3.7. *The geometric realization of $\mathcal{A}$ is an Eilenberg-Maclane space for the pure braid group.*

Proof. For any closed cell c in the Moussong complex, let $\mathcal{H}_c$ denote the set of walls cutting through c. Note that removing these walls chops up the Moussong comlex into convex open subsets. The set of these subsets is in 1-1-correspondence to the vertices of C: Each vertex of c pick the convex open set $C_{(c,v)}$ that contains v.

On the other hand, let D_c be the open star of the barycenter of c in the barycentric subdivision of X. Then D_c is, again, a convex open subset of X. Finally, put

$$U_{(c,v)} := D_c \times C_{(c,v)}.$$

This is a cover of X by convex open sets.

!!! finish this !!! **q.e.d.**

Corollary 9.3.8. *The geometric realization $|\mathcal{A}_n|$ is an Eilenberg-Maclane space for the pure braid group B_n.*

Remark 9.3.9. All of this is Perm_n-equivariant. Thus

$$\mathrm{Perm}_n \backslash |\mathcal{A}_n|$$

is an Eilenberg-Maclane space for the braid group.

As a consequence, we can actually work out a presentation for the braid group B_n. Let us consider the case of B_3 first. Here, the underlying Coxeter group is the symmetric group on 3 letters with standard genrating set given by two transpositions. Our Eilenberg-Maclane complex has precisely one 2-cell, which is a hexagon, two edges, and one vertex. The tricky part is to figure out, how the 2-cell is attached.

!!! ... !!!

It turns out, that we get the following presentation of the braid group B_n:

$$B_n = \left\langle s_1, \ldots, s_n \;\middle|\; \begin{array}{ll} s_i s_j s_i = s_j s_i s_j & \text{for } |i-j| \geq 2 \\ s_i s_j = s_j s_i & \text{for } |i-j| = 1 \end{array} \right\rangle.$$

Exercise 9.3.10. Show that the $\mathrm{H}_3\left(\mathrm{Perm}_n \backslash |\mathcal{A}_4|\right)$ is non-trivial. Infer that B_4 does not have an Eilenberg-Maclane complex of dimension ≤ 2.

Exercise 9.3.11. Prove that $B_3 = \langle a, b, c \mid ab = bc = ca \rangle$.

Exercise 9.3.12. More generally, prove that

$$B_n = \left\langle x_{[i,j]} \;\; (i \neq j) \;\middle|\; \begin{array}{ll} x_{[i,j]} x_{[j,k]} = x_{[j,k]} x_{[k,i]} & \text{if } [i,j,k] \\ x_{[i,j]} x_{[k,l]} = x_{[k,l]} x_{[i,j]} & \text{if } [i,j,k,l] \end{array} \right\rangle$$

where we put a cyclic ordering on $\{1, 2, \ldots, n\}$ and $[a, b, \ldots]$ denotes the fact that the listed elements form a cycle in their given order. In particular, the generators are indexed by cycles of length 2.

Exercise 9.3.13. Prove that $B_3 = \langle a, b, c, s \mid ab = bc = ca = s \rangle$. Moreover, show that the Cayley 2-complex (i.e., the universal cover of the canonical 2-complex associated to this presentation) admits a CAT(0) metric. (This implies that the presentation 2-complex for this presentation is an Eilenberg-Maclane space for B_3.)

Exercise 9.3.14. Decide whether the presentation 2-complex for the presentation

$$B_3 = \langle a, b, c \mid ab = bc = ca \rangle$$

is an Eilenberg-Maclane complex for B_3.

CAT(0)-Structures

9.3.2 General Artin Groups

Fact 9.3.15 (van Lek). *Let M be a Coxeter matrix over S, and let $J \subseteq S$ be a set of generators with restricted Coxeter matrx M_J. Then the canonical homorphism*

$$A_{M_J} \to A_M$$

is injective. The image is the subgroup generated by J.

Fact 9.3.16. *The space*

$$X \times X - \bigcup_H H \times H$$

is homotopy equivalent to the poset

$$\mathcal{A}_M := \{(c, v) \mid c \text{ cell in } X,\ v \text{ vertex in } c\}$$

where the order is given by

$$(c, v) \preceq (d, w) \text{ if and only if } c \leq d \text{ and } v = \pi_c(w)\,.$$

The fundamental group of these spaces is the pure Artin group.

This space is conjectured to be an Eilenberg-Maclane space.

Fact 9.3.17 (Charney-Davis). *The poset $\mathcal{A}_M$ is an Eilenberg-Maclane space for the pure Artin group P_M, provided any two Artin generators generate a finite subgroup, i.e., the Coxeter matrix M is 2-spherical. One obtains an Eilenberg-Maclane space for the corresponding Artin group A_M by modding out the group action of WM]. In particular, Artin groups of finite type have a finite Eilenberg-Maclane complex.*

9.3.3 Artin Groups of Finite Type

Fact 9.3.18 (Brieskorn-Saito, Deligne). *Artin groups of finite type have solvable word and conjugacy problem.*

Fact 9.3.19 (Charney). *Artin groups of finite type are biautomatic.*

Fact 9.3.20 (Bestvina). *Artin groups of finite type have the look and feel of $CAT(0)$-groups: Let A be an Artin group of finite type. Then the following hold:*

1. *The group A contains only finitely many conjugacy classes of finite subgroups.*

2. *Every solvable subgroup of A is finitely generated and virtually abelian.*

3. *The set of translation lengths is bounded away from 0. (Note that Artin groups of finite type have a finite Eilenberg-Maclane complex by (9.3.17) and are, therefore, torsion free.)*

Fact 9.3.21 (Squier). *An Artin group of finite type over the generating set S is a duality group of dimension $|S|$.*

Fact 9.3.22 (Krammer, Cohen-Wales). *Artin groups of finite type are linear.*

9.3.4 Right-Angled Artin Groups and the Example of M. Bestvina and N. Brady

Right-angled Artin groups are also known as graph groups since the data determining the presentation can most easyly be visualized as a graph: To any graph Γ with vertex set $\mathcal{V}$, we associate the group

$$G_\Gamma := \langle v \in \mathcal{V} \mid vw = wv \text{ if there is an edge } v — w \text{ in } \Gamma \rangle \,.$$

Note that there is a canonical homomorphism

$$\begin{array}{rcl} \varphi : G_\Gamma & \rightarrow & \mathbb{Z} \\ v & \mapsto & 1 \end{array}$$

whose kernel will be denoted by K_Γ.

In this section, we also identify Γ with its associated flag complex, i.e., the simplicial complex that shares the vertices with Γ and whose simplices are cliques in Γ: A set of vertices forms a simplex if the vertices are pairwise connected by edges.

Theorem 9.3.23 (Bestvina-Brady [BeBr97]). *If Γ is a finite flag complex then the following hold:*

1. *K_Γ is of type F_m if and only if Γ is $(m-1)$-connected.*
2. *K_Γ is of type FP_m if and only if Γ is $(m-1)$-acyclic.*

This section is devoted to a proof of this result. Note that the theorem allows one to construct groups with prescribed finiteness properties. In particular, we could take Γ to be 1-acyclic but not simply connected and infer:

Corollary 9.3.24. *There is a group of type* FP2 *that is not finitely presented.* **q.e.d.**

First, we construct an Eilenberg-Maclane space for G_Γ. Let $T_{\mathcal{V}}$ a product of a family of circles C_v indexed by the vertices in $\mathcal{V}$. We assume that all these circles have a basepoint so that we can regard them as subspaces in $T_{\mathcal{V}}$. For any subset σ of $\mathcal{V}$ we regard the torus $T_\sigma = \times_{v\in\sigma} C_v$ as a subtorus of $T_{\mathcal{V}}$. We put

$$Q_\Gamma := \bigcup_{\sigma \text{ simplex}} T_\sigma$$

and let

$$X_\Gamma := \tilde{Q}_\Gamma$$

denote its universal cover.

Observation 9.3.25. *The complex Q_Γ has precisely one vertex P, and the link of this vertex is*

$$\mathrm{Lk}(P) = S(\Gamma) = \bigcup_{\sigma \text{ simplex}} S_\sigma \subset \mathbb{R}^{\mathcal{V}}$$

where S_σ denotes the unit sphere in $\mathbb{R}^\sigma \subseteq \mathbb{R}^{\mathcal{V}}$. The cuibical structure on Q_Γ induces the triangulation on $S(L)$ given by

$$S_\sigma = \mathop{\Large\ast}_{v\in\sigma} S_{\{v\}}.$$

Note that X_Γ is a piecewise Euclidean cube complex and all its vertex links are isomorphic to $S(\Gamma)$.

Exercise 9.3.26. Show that $S(\Gamma)$ is a flag complex.

Corollary 9.3.27. *X_Γ is $CAT(0)$ and, therefore, contractible.* **q.e.d.**

The canonical homomorphism φ has a topological representative

$$h : Q \to C$$

that is piecewise linear and restricts to the degree 1 map on each C_v. It lifts to a piecewise linear map

$$h : X \to \mathbb{R}$$

which is affine on each cube in X.

Definition 9.3.28. A combinatorial Morse function on a piecewise Euclidean complex is a real valued function h that is affine on closed cells, non-constant on edges, and has a discrete set of critical values, i.e., the image of the 0-skeleton is discrete in $\mathbb{R}$.

The descending (ascending) link of a vertex v is that part of its link spanned by those cells for which v is a maximum (minimum) for h.

The s-level set is the h-preimage of the real number s. The s-sublevel set is the preimage of $(-\infty, s]$. For any closed interval I, we call its h-preimage the I-slice.

Lemma 9.3.29. *Let $r < s$ be two real numbers such that there are no critical values in $[r, s]$. Then the r-sublevel and s-sublevel sets are homotopy equivalent. Similarly, any two slices whose difference does not contain vertices are homotopy equivalent.*

Proof. Observe that the level set cuts through the polyhedral cells of the complex. Thereby, the upper level set creates a free face in each cell. You can collapse the top-dimensional material in the affected cells away. This defines a deformation retraction. Now induct on lower dimensional material. **q.e.d.**

!!! Finish this !!!

Chapter 10

Grigorchuk's First Group

The "first Grigorchuk Group" $\mathcal{G}_1$ comes close to a "universal counterexample" as far as finitely generated groups are concerned. It is the group theoretic analogue of a fractal.

[10.1.4] $\mathcal{G}_1$ is finitely generated.

[10.7.10] $\mathcal{G}_1$ is not finitely presented.

[10.5.15] $\mathcal{G}_1$ has intermediate growth.

[10.6.1] $\mathcal{G}_1$ is amenable.

[10.6.8] $\mathcal{G}_1$ is not elementary amenable.

[10.4.16] Every finite two group embedds into $\mathcal{G}_1$.

[10.3.10] $\mathcal{G}_1$ is a general Burnside group. Thus:

- [10.3.5] Every element of $\mathcal{G}_1$ has finite order.
- [10.3.7] $\mathcal{G}_1$ is not virtually solvable and does not contain a non-abelian free group.

[10.4.21] $\mathcal{G}_1$ is residually finite.

[10.4.25] $\mathcal{G}_1$ is just infinite.

[10.4.24] $\mathcal{G}_1$ has the congruence subgroup property.

[10.2.33] The word problem is solvabel in $\mathcal{G}_1$.

[10.5.10] The conjugacy problem is solvabel in $\mathcal{G}_1$.

[10.2.13] $\mathcal{G}_1$ is an automaton group.

[10.4.10] $\mathcal{G}_1$ is commensurable to $\mathcal{G}_1 \times \mathcal{G}_1$.

[10.3.13] $\mathcal{G}_1$ is a 2-group.

[10.7.7] $\mathcal{G}_1$ is not co-Hopfian.

[10.7.8] $\mathcal{G}_1$ is Hopfian.

10.1 The Infinite Binary Rooted Tree

The First Grigorchuk Group $\mathcal{G}_1$ is a group of automorphisms of the infinite binary rooted tree. So let us consider the automorphism group of this gadget first.

Let $T_2^\bullet$ be the rooted binary tree without terminal vertices. Any vertex in $T_2^\bullet$ can be reached from the root by a minimal path. Along this path, you have to make binary decisions whether you want to go left or right. Hence the vertices in $T_2^\bullet$ are finite word over $\{-1, 1\}$ with the empty word as the root. The vertices come in levels indexed by natural numbers: the root is the unique vertex at level 0, its children are at level 1, and so on. The level of a vertex is its path-metric distance to the root. Let $T_2^\bullet(n)$ denote the sub-tree spanned by the vertices of level $\leq n$.

let $\mathcal{T}_2^\bullet := \mathrm{Aut}(T_2^\bullet)$ denote the automorphism group of $T_2^\bullet$. Note that any automorphism preserves the root (as this is the only vertex of valency 2). Hence the sets $T_2^\bullet(n)$ are invariant under automorphism, and we have, for any n, a canonical homomorphism

$$\pi_n : \mathcal{T}_2^\bullet \to \mathrm{Aut}(T_2^\bullet(n))\,.$$

In fact, $\mathcal{T}_2^\bullet$ is easily seen to be the inverse limit of the system

$$\mathrm{Aut}(T_2^\bullet(0)) \leftarrow \mathrm{Aut}(T_2^\bullet(1)) \leftarrow \mathrm{Aut}(T_2^\bullet(2)) \leftarrow \cdots\,.$$

It follows that $\mathcal{T}_2^\bullet$ is pro-finite. So it is a compact topological group.

Let us have a look at the kernels of these homomorphism. We define

$$\mathcal{T}_2(n) := \ker(\pi_n : \mathcal{T}^\bullet \to \mathrm{Aut}(T_2^\bullet(n)))\,.$$

All of these subgroups are normal. The first one deserves our utmost attention: Note that $T_2^\bullet$ contains two copies of itself as subtrees – the vertices at level 1 serve as roots for these subtrees. Let us call these subtrees the left subtree T^{l} and the right subtree T^{r}. The subgroup $\mathcal{T}_2^\bullet$ is the group of automorphisms taking T^{l} to T^{l} and T^{r} to T^{r}. Moreover, this subgroup is clearly isomorphic to the square of $\mathcal{T}_2^\bullet$:

$$\mathcal{T}_2^\bullet \times T_2^\bullet \cong \mathcal{T}_2(1)\,.$$

For this reason, we can interpret the direct product of subgroups of $\mathcal{T}_2^\bullet$ as a subgroup of $\mathcal{T}_2(1) \leq \mathcal{T}_2^\bullet$:

Notation 10.1.1. For two subgroups $A, B \leq \mathcal{T}_2^\bullet$, let

$$A \otimes B$$

denote the group of those elements in $\mathcal{T}_2(1)$ that act as an element of A on the left subtree and as an element of B on the right subtree. Note that $A \otimes B$ is isomorphic to $A \times B$ as an abstract group.

On the other hand, the short exact sequence

$$\mathcal{T}_2(1) \hookrightarrow \mathcal{T}^\bullet \twoheadrightarrow C_2$$

splits since the swap $\sigma \in \mathcal{T}^\bullet$ has order two. This is the automorphism that interchanges the left and right subtree. The formal definition makes use of the representation of vertices as words over $\{\pm 1\}$: The swap σ just flips the sign in the first slot. Hence

$$\mathcal{T}^\bullet \cong \mathcal{T}_2(1) \rtimes C_2$$

Putting things together, we obtain a strange isomorphism:

$$\mathcal{T}_2^\bullet \cong (\mathcal{T}_2^\bullet \times T_2^\bullet) \rtimes C_2 = \mathcal{T}_2^\bullet \wr C_2$$

This allows us to define automorphisms recursively. As an example, consider the tree automorphism φ defined by

$$\varphi = (1, \varphi)\sigma.$$

This equation has a unique solution. Indeed, the right hand side tells us first how φ acts on level 1 vertices. Then we can plug this information back into the right hand side. Now the right hand side is defined on all vertices up to level 2. We can continue in this fashion. Similarly we have

Observation 10.1.2. *Let φ_i be variables with values in $\mathcal{T}_2^\bullet$, w_i and u_i given words in these variables and some fixed given automorphisms (like the swap), and ε_i given elements of $\{1, \sigma\}$. Then any system of equations*

$$\begin{array}{rcl} \varphi_1 & = & (w_1, u_1)\varepsilon_1 \\ \vdots & \vdots & \vdots \\ \varphi_n & = & (w_n, u_n)\varepsilon_n \end{array}$$

has a unique solution. **q.e.d.**

Exercise 10.1.3. Show that the φ defined by

$$\varphi = (1, \varphi)\sigma$$

has infinite order.

Now we can define cool automorphisms!

Definition 10.1.4 (Grigorchuk [Grig80]). Let the automorphisms β, γ, and δ be defined by

$$\begin{aligned} \beta &= (\sigma, \gamma) \\ \gamma &= (\sigma, \delta) \\ \delta &= (1, \beta) \end{aligned}$$

The First Grigorchuk Group $\mathcal{G}_1$ is the subgroup of $\mathcal{T}^\bullet$ generated by $\{\sigma, \beta, \gamma, \delta\}$:

$$\mathcal{G}_1 := \langle \sigma, \beta, \gamma, \delta \rangle .$$

Our observation (10.1.2) is quite powerful. In fact, we can derive a complete multiplication table for the generators β, γ, and δ as follows:

Proposition 10.1.5. *The set $\{1, \beta, \gamma, \delta\}$ is a subgroup of $\mathcal{G}_1$ isomorphic to Klein's Vierergruppe $C_2 \times C_2$.*

Proof. From the defining equations, we get

$$\begin{aligned} \beta^2 &= (1, \gamma^2) \\ \gamma^2 &= (1, \delta^2) \\ \delta^2 &= (1, \beta^2) \end{aligned}$$

Regarding this as a system of equations in β^2, γ^2, and δ^2, we conclude that the unique solution is

$$\beta^2 = \gamma^2 = \delta^2 = 1.$$

We turn this trick into a method, i.e., we use it twice. So let us write down a system of equations in the products of length 3. We find:

$$\begin{aligned} \beta\gamma\delta &= (1, \gamma\delta\beta) \\ \gamma\delta\beta &= (1, \delta\beta\gamma) \\ \delta\beta\gamma &= (1, \beta\gamma\delta) \end{aligned}$$

and it follows that

$$\beta\gamma\delta = \gamma\delta\beta = \delta\beta\gamma = 1.$$

So we established that the defining relations of Klein's Vierergruppe hold. Hence we only have to check that β, γ, and δ are non-trivial. But this follows from the swaps that occur in the defining set of equations. **q.e.d.**

Let us call an automorphism of $T_2^\bullet$ is recursive if it can be defined by a finite system of equations. If you invert all the equations, you see that the inverse of a recursive automorphism is recursive, too. Moreover, it is easy to see that products of recursive automorphisms are recursive. Hence the recursive automorphisms form a group. We shall study a subgroup of this momentarily.

The exposition in this section owes a lot to [Harp00] and [BGS01]

10.2 Automaton Groups and the Word Problem

Definition 10.2.1. A simplistic finite state automaton over the alphabet $\mathcal{A}$ is a finite directed graph with a distinguished start vertex whose edges and vertices are labeled by elements of $\mathcal{A}$ such that for any vertex v and any letter $a \in \mathcal{A}$ there is precisely one edge starting at v labeled with a.

A transformation is a length preserving map $\mathcal{A}^* \to \mathcal{A}^*$ where $\mathcal{A}^*$ is the set of all words over the alphabet $\mathcal{A}$, including the empty word.

Remark 10.2.2. There is an obvious way to use a finite state automaton over $\mathcal{A}$ to define a transformation $\mathcal{A}^* \to \mathcal{A}^*$. Given a sequence of letters, there is a unique directed path starting at the start vertex that reads this sequence of letters. The output is given by reading the vertex labels along this path, starting with the vertex after the start vertex.

Definition 10.2.3. A sophisticated finite state automaton over the alphabet $\mathcal{A}$ is a finite directed graph with a distinguished start vertex together with to labelings. The vertices are labeled by elements of $\mathrm{Perm}(\mathcal{A})$ and the edges carry labels taken from the alphabet $\mathcal{A}$ such that for any vertex v and any letter $a \in \mathcal{A}$ there is precisely one edge starting at v labeled with a.

Remark 10.2.4. There is also an obvious way to use a sophisticated finite state automaton over $\mathcal{A}$ to define a transformation $\mathcal{A}^* \to \mathcal{A}^*$. Given a sequence of letters, take the unique directed path starting at the start vertex that reads this sequence of letters. The output is given by applying the permutations you read along this path to the letters in your sequence.

Exercise 10.2.5. Show that a transformation $\mathcal{A}^* \to \mathcal{A}^*$ can be defined by a simplistic finite state automaton if and only if it can be realized by a sophisticated finite state automaton.

Remark 10.2.6. Maybe, one should introduce the even more convenient notion of a finite state automaton deluxe where the vertices carry labels in $\mathrm{Map}(\mathcal{A}, \mathcal{A})$. This generalizes both notions: For simplistic automata, use constant maps; and for sophisticated automata, use permutations. However, it does not add to the computational power of these devices.

Definition 10.2.7. A vertex v in a finite state automaton is accessible if there is a directed path from the start vertex to v. An automaton is sophomoric if it has inaccessible vertices.

From now on, all automata will be simplistic but not sophomoric. Let us have a look at some small automata over the alphabet with two letters $\{\mathbf{L}, \mathbf{R}\}$.

Example 10.2.8. Here is an automaton, that maps everything to a string of $\mathbf{R}$s of the same length.

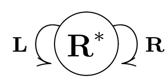

Example 10.2.9. The identity can be realized with two vertices:

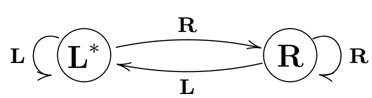

Example 10.2.10. And here is the swap:

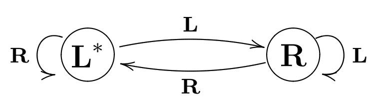

Example 10.2.11. The twist automorphism

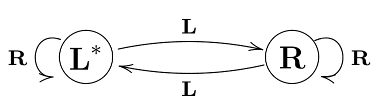

needs two states but has a fairly complicated dynamic.

The picture shows that this automorphism is just a twisted swap, whence the name!

Exercise 10.2.12. Prove or disprove: The twist has infinite order.

Example 10.2.13. Here is a picture that displays all the generators of $\mathcal{G}_1$ at once. You just have to pick the right vertex as start.

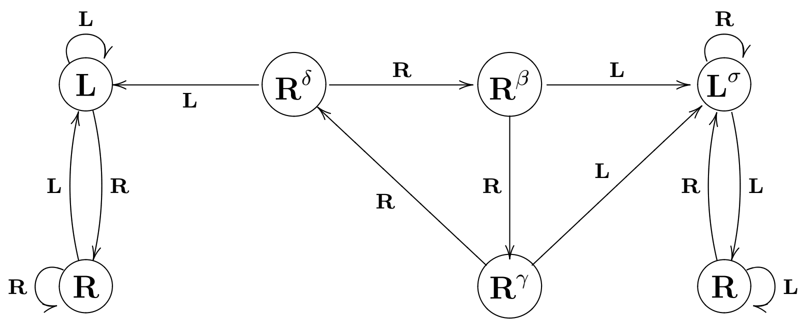

Definition 10.2.14. A map $\varphi : \mathcal{A}^* \to \mathcal{A}^*$ is finitary if it can be realized by a finite state automaton.

Observation 10.2.15. *Of course, bijective transformations $\varphi : \mathcal{A}^* \to \mathcal{A}^*$ are tree-automorphisms of a rooted tree where each vertex has precisely $|\mathcal{A}|$ children. Hence we can speak of finitary tree-automorphisms.*

Moreover, any finitary automorphism is recursive. You construct the defining equations from the automation A as follows: For each vertex introduce a variable, and the defining equation will have the children of this vertex in the pair followed by a swap of the identity, depending on the local labels. Rule: the child that prints **R** *is in the right slot. As the preceding sentence is incomprehensible for those who are not in the know, let us consider an example. Here is the system for the twist (10.2.11) where t corresponds to the start at* **R** *and y corresponds to the start at* **L**.

$$\begin{aligned} x &= (y, x)\sigma \\ y &= (y, x) \end{aligned}$$

Exercise 10.2.16. Is there a recursive system of equations defining a tree automorphism that cannot be realized by a finite state automaton?

Exercise 10.2.17. This is an automaton over $\{0, 1\}$. What does it do? Is the transformation invertible? Is it of infinite order? If the transformation is invertible, find a recursive definition as small as possible.

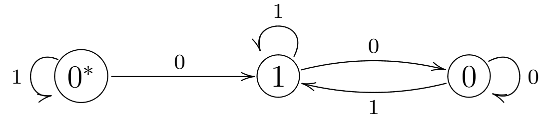

We will show that finitary bijections form a group. Obviously, we have to construct automata for inverses and products. So let us start with inverses.

Key Idea 10.2.18. *The basic idea of inverting an automaton is this: if you are at the start vertex and the input sequence gives you a letter, you do not have to look at the edge labels but at the labels of the neighboring vertices. If there is precisely one with that label, you know where to go. While you are on the way, you print out the label of the edge. So we would like to construct the inverse of an automaton just by exchanging the label of an oriented edge with the label of its terminal vertex.*

However, this simple minded procedure might not be allowed! Have a look at example (10.2.11). The trouble comes from the fact that some vertices have edges with different label pointing to them.

Definition 10.2.19. A finite state automaton is <u>tidy</u> if every vertex has all its incoming edges given identical labels.

The main tool for inverting an automaton is the blow-up construction:

Proposition 10.2.20 (Blow-Up). *For every finite state automaton, there is an equivalent tidy finite state automaton.*

Proof. We use a cover. So let A be a finite state automaton over the alphabet $\mathcal{A}$. We define a new automaton B on the vertex set $\mathcal{V}_A \times \mathcal{A}$. For an edge $\vec{e}$ in A with label a pointing from S to T, we glue in $|\mathcal{A}|$ edges in B. They point from the vertices $(S, -)$ to the vertex (T, a).

It is obvious that the result B is tidy and equivalent to A. To understand this, let us give the pictures for the twist-automaton. Here is the untidy version again:

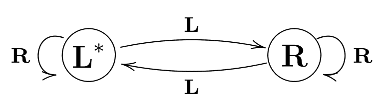

And this is the way to tidy it up:

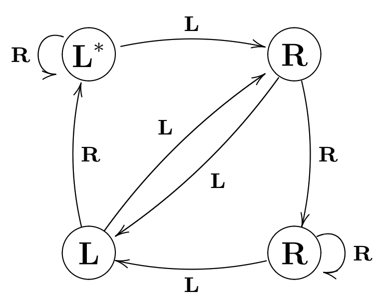

Play and see!! **q.e.d.**

For tidy automata, finding the inverse is no problem.

Proposition 10.2.21. *Let A be a tidy finite state automaton. Define a graph B with edge and vertex labels on the same vertex set by exchanging the label of each oriented edge with the label of its terminal vertex. The result is a finite state automaton if and only if A realizes an invertible transformation. In this case B is tidy and realizes the inverse.*

Proof. Look at the picture for the twist-automorphism. Here is the tidy twist:

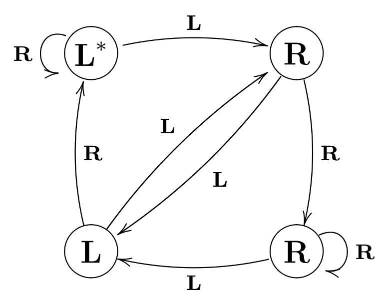

And here is the inverse (the untwist):

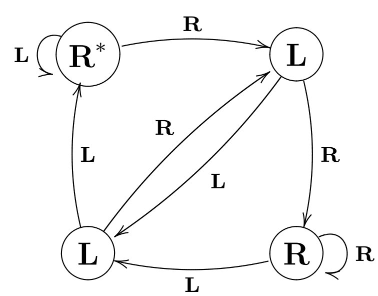

Now, play with it and see how the key idea (10.2.18) works. **q.e.d.**

Corollary 10.2.22. *It is decidable whether a finite state automaton defines an invertible transformation, and if it does an automaton realizing the inverse can be constructed effectively.*

Problem 10.2.23. Is there an algorithm that takes a finite state automaton as input and decides whether it describes a transformation of finite order?

Proposition 10.2.24. *Suppose the transformations $\varphi : \mathcal{A}^* \to \mathcal{A}^*$ and $\psi : \mathcal{A}^* \to \mathcal{A}^*$ are finitary, i.e., they can both be realized by finite state automata A and B. Then the composition $\psi\varphi$ (second factor acts first on the input!) is finitary, too. Moreover, an automaton realizing the product can be effectively constructed from A and B.*

Proof. The product automaton is constructed on the vertex set $A \times B$. The start vertex is the pair of start vertices. From (S_1, T_1) we have an edge labeled with a to (S_2, T_2) if there is an a-edge from S_1 to S_2 and an edge from T_1 to T_2 that has the same label as v_2. The point is, that this label is the output of A at this stage and therefore guides the computational path in B. **q.e.d.**

Corollary 10.2.25. *The set of all those bijections that can be realized by finite state automata forms a group.*

Definition 10.2.26. An automata group over $\mathcal{A}$ is a group of bijective finitary transformations $\mathcal{A}^* \to \mathcal{A}^*$.

Example 10.2.27. From (10.2.13) it follows that $\mathcal{G}_1$ is an automata group.

The following lemma introduces an idea that is ubiquitous in the study of finite state automata. The key observation is that a path in a finite graph has to contain a loop if it becomes too long.

Lemma 10.2.28. *If two finite state automata A and B over the same alphabet are inequivalent, then there is an input of length $\leq |A| \times |B|$ for which their outputs differ.*

Let us first note an immediate consequence.

Corollary 10.2.29. *There is an algorithm that, taking two finite state automata A_1, and A_2 over the same alphabet is its input, decides if these automata define the same transformation.*

Proof. Let us assume the shortest input sequence that proves the two automata to be inequivalent has length $> |A| \times |B|$. Follow the computational paths in A and B for this sequence. As there are only $|A| \times |B|$ many pairs of states $(S, T) \in A \times B$, one of these pairs is visited twice. Then, however, the part of the input sequence between the two times can be cut out without affecting the rest of the computation. **q.e.d.**

Exercise 10.2.30. Prove: A finite state automaton takes ultimately periodic inputs to ultimately periodic outputs.

Exercise 10.2.31. The decision procedure for equivalence of finite state automata based on checking inputs of length $\leq |A| \times |B|$ is exponential. Find an *effective* algorithm that decides if two automata are equivalent.

Let us fix some consequences that pertain to Grigorchuk's Group.

Corollary 10.2.32. *The word problem for finitely generated automata groups is solvable.*

Corollary 10.2.33. *The word problem in $\mathcal{G}_1$ is solvable.*

This solution to the word problem closely parallels the solution to the word problem in finitely generated linear groups. The idea is as follows: Given a word in the generators, just multiply the corresponding matrices and check if the result is the identity matrix. Of course, this presupposes that you can multiply matrices. The problem is that, say, complex numbers do not have finite representations.

To overcome this problem, observe that finitely many matrices have only finitely many coefficients. So we start with matrix coefficients like π, e or $\sqrt{-5}$ and in the end, we have a martix whose entries are rational functions in the original coefficients. How can we detect, if this is the identity matrix? Well, this is precisely what a computer algebra system is supposed to do. The way is based on the observtion that all computations really take place in a finite extension of the prime field. Now, there is a little lemma to be proved that says you can always do this extension in two steps:

(a) Pass to a purely transcendental extension. This field obviously has a computationally effective arithmetic.

(b) Move on to an algebraic extension of finite degree – this can be done since finitely generated algebraic extensions are finite. Those extensions can be represented as matrix algebras over their base field. Hence they, too, are computationally effective.

These considerations have two consequences:

Theorem 10.2.34. *A finitely generated linear group has a solvable word problem.* **q.e.d.**

Theorem 10.2.35 (Mal'cev). *A finitely generated linear group has a faithful representation over a pure transcendental extension of the prime field.*

Proof. There is only a finite extension missing. This, however, can be realized as a matrix algebra. **q.e.d.**

This technique is know as "restriction of scalars".

10.3 Burnside's Problem

Definition 10.3.1. A group G is <u>periodic</u> if there is a number $n \in \mathbb{Z}$ such that for all $g \in G$.

$$g^n = 1.$$

The number n is called the <u>exponent</u> of G.

General Burnside Problem 10.3.2. *Are there no finitely generated, infinite torsion groups?*

Burnside Problem 10.3.3. *Are there no finitely generated, infinite periodic groups?*

Restricted Burnside Problem 10.3.4. *Are there only finitely many finite groups with a given exponent and a given bound on the number of generators?*

The answers are "no", "no", and "yes". Grigorchuk's Group provides a example for the General Burnside Problem.

Definition 10.3.5. A <u>general Burnside group</u> is a finitely generated, infinite torsion group.

It is not at all easy to come up with examples of finitely generated, infinite torsion groups since these are subject to rather strong restrictions.

Theorem 10.3.6. *Finitely generated, virtually solvable torsion groups are finite.*

Proof. Let us ignore the word "virtually" for a while. First observe that finitely generated abelian torsion groups are finite. Being abelian is the same as being step-1-solvable. Of course, we proceed by induction.

Let $N \hookrightarrow G \twoheadrightarrow Q$ be a short exact sequence with Q abelian, N step-s-solvable, and G finitely generated and torsion. It follows that Q is finitely generated and torsion and, therefore, finite. Hence N has finite index and is finitely generated, too. As a subgroup of G, it is clearly torsion. By induction, N is finite. Hence G is finite-by-finite whence finite. This completes the induction.

Now, we know that finitely generated solvable torsion groups are finite. So suppose G is torsion, finitely generated, and virtually solvable. Then the solvable subgroup of finite index is finitely generated and torsion. Hence it is finite. But if G has a finite subgroup of finite index, G is finite. **q.e.d.**

Corollary 10.3.7. *A general Burnside group does not satisfy the Tits alternative, i.e., it neither contains a non-abelian free group, nor is it virtually solvable.*

Proof. Since a torsion group cannot contain a non-abelian free subgroup, we are reduced to the virtually solvable case. **q.e.d.**

Corollary 10.3.8 (Burnside-Schur-Kaplanski). *Finitely generated linear torsion groups are finite.*

Remark 10.3.9. Of course, this is sort of a a mock proof of the Burnside-Schur-Kaplanski theorem: Tits' proof that linear groups satisfy the Tits alternative uses results from representation theory that come close to give the Burnside-Schur-Kaplanski result directly.

Theorem 10.3.10. *The First Grigorchuk group $\mathcal{G}_1$ is a general Burnside group. In particular it is neither linear nor virtually solvable.*

As $\mathcal{G}_1$ is finitely generated by construction, we have to prove that it is infinite [10.3.12] and torsion [10.3.13].

Lemma 10.3.11. *Let $\mathcal{G}_1(1) := \ker(\mathcal{G}_1 \to \mathrm{Aut}(T_2^\bullet(1)))$ be the stabilizer of the level 1 subtree. The inclusion*

$$\mathcal{G}_1(1) \hookrightarrow \mathcal{T}_2(1) = \mathrm{Aut}(T_2^\bullet) \otimes \mathrm{Aut}(T_2^\bullet)$$

induces an inclusion

$$\mathcal{G}_1(1) \hookrightarrow \mathcal{G}_1 \otimes \mathcal{G}_1$$

which is "surjective in each coordinate", i.e., for each element $\xi_0 \in \mathcal{G}_1$ there is an element $\xi_1 \in \mathcal{G}_1$ such that $(\xi_0, \xi_1) \in \mathcal{G}_1(1)$.

Proof. $\mathcal{G}_1(1)$ contains (in fact it is generated by) the elements

$$(\sigma, \beta), \quad (\sigma, \gamma), \quad (\sigma, \delta), \quad (\beta, \sigma), \quad (\gamma, \sigma), \quad (\delta, \sigma)$$

and therefore surjects onto $\mathcal{G}_1$, e.g., by projection onto the first coordinate. **q.e.d.**

Corollary 10.3.12. *$\mathcal{G}_1$ is infinite.*

Proof. $\mathcal{G}_1(1)$ is a proper subgroup of $\mathcal{G}_1$. However, a finite group cannot have a proper subgroup that surjects onto the bigger group. **q.e.d.**

Proposition 10.3.13. *$\mathcal{G}_1$ is a 2-group, i.e., every element has finite order which is a power of 2.*

Proof. This is done by induction on a "complexity", a refined version of word length. First observe that, since all generators of $\mathcal{G}_1$ are involutions, we can represent each element as a word in the generators without using inverses. Moreover, (10.1.5) implies that we can use a word in which the swap alternates with the other generators. Finally, since the order of an element is unaffected by conjugation, we may assume that the word length is ≤ 1 or even. Let us call those words *cyclicly reduced*.

For each cyclicly reduced word w let $C_\sigma(w)$ denote the number of occurrences of σ in w. Define C_β, C_γ, and C_δ analogously. The *complexity* of w to be the tuple

$$C(w) := [C_\sigma(w)\,; C_\beta(w)\,, C_\gamma(w)\,, C_\delta(w)]\,.$$

We order complexities lexicographically. Note that for cyclicly reduced words the σ-count closely reflects the length of the word. And for $C_\sigma(w) \geq 1$, we have $C_\sigma(w) = C_\beta(w) + C_\gamma(w) + C_\delta(w)$ unless $w = \sigma$.

The induction starts with

$$C(w) \in \{[0; 0, 0, 1]\,, [0; 0, 1, 0]\,, [0; 1, 0, 0]\,, [1; 0, 0, 0]\}\,.$$

These cases correspond to the generators which are involutions.

Now assume that w has a higher complexity. We distinguish two cases:

<u>$C_\sigma(w)$ is even</u>: Then $w = (w_1, w_2)$ for two word whose σ-count is at most $\frac{C_\sigma(w)}{2}$. Hence induction applies.

<u>$C_\sigma(w)$ is odd</u>: Now we will consider w^2. This word has an even σ-count whence we have

$$w^2 = (w_1, w_2)$$

and our aim is to show that w_1 and w_2 have a smaller complexity that w.

We have subcases:

<u>$C_\delta(w) > 0$</u>: In this case, $C_\sigma(w_i) < C_\sigma(w)$. The reason is $\delta = (1, \beta)$ and $\sigma\delta\sigma = (\beta, 1)$. The 1-components do the shortening.

<u>$C_\delta(w) = 0$</u>: In this case, no cancellations occur. The word w_i will be cyclicly reduced right away, and we can track where the letters come from. It transpires that β-letters in w give rise to γ-letters in the w_i and γ-letters in w will provide δ-letters in the w_i. Hence, we have

$$C(w_i) = [C_\sigma(w)\,; 0, C_\beta(w)\,, C_\gamma(w)] < [C_\sigma(w)\,; C_\beta(w)\,, C_\gamma(w)\,, 0]\,.$$

This completes the induction. **q.e.d.**

Remark 10.3.14. Let us explicitly write out the low complexity cases:

- $\sigma\delta$ has order 4. Hence $\langle\sigma,\delta\rangle$ is a dihedral group of order 8.
- $\sigma\gamma$ has order 8. Hence $\langle\sigma,\gamma\rangle$ is a dihedral group of order 16.
- $\sigma\beta$ has order 16. Hence $\langle\sigma,\beta\rangle$ is a dihedral group of order 32.

10.4 Subgroup Structure

Definition 10.4.1. Let us define the following elements:

$$\begin{array}{rcl} t & := & \sigma\beta\sigma\beta = (\gamma\sigma,\sigma\gamma) \\ v & := & (\beta\sigma\delta\sigma)^2 = (t,1) \\ w & := & (\sigma\beta\sigma\delta)^2 = (1,t) \end{array}$$

Darn Technical Computation 10.4.2. *We have the following conjugacy identities:*

$$\begin{array}{lll} \beta t\beta = \sigma\beta\sigma\beta = t^{-1} & \beta w\beta = t^{-1}\delta t\delta = t^{-1}w^{-1}t & \beta v\beta = v^{-1} \\ \sigma t\sigma = t^{-1} & \sigma w\sigma = v & \sigma v\sigma = w \\ \delta t\delta = w^{-1}t & \delta w\delta = w^{-1} & \delta v\delta = \beta\sigma\delta\sigma\delta\beta\delta\sigma\delta\sigma = v. \end{array}$$

In the last identity, we used $\sigma\delta\sigma\delta = \delta\sigma\delta\sigma$.

In particular, $\langle t,v,w\rangle$ is normal in $\mathcal{G}_1$.

Moreover, we have

$$\begin{array}{rcl} \beta\beta\beta & = & \beta \\ \sigma\beta\sigma & = & t\beta \\ \delta\beta\delta & = & \beta \end{array}$$

In particular $\langle\beta,t,v,w\rangle$ is normal in $\mathcal{G}_1$.

Definition 10.4.3. We put

$$\begin{array}{rcl} B & := & \langle\beta,t,v,w\rangle \\ K & := & \langle t,v,w\rangle \end{array}$$

Lemma 10.4.4. *Let $\pi_s : \mathcal{G}_1 \to \mathrm{Aut}(T_2^\bullet(s))$ be the canonical projection. Then*

$$\begin{array}{rcl} |\pi_3(\mathcal{G}_1)| & = & 128 \\ |\pi_3(B)| & = & 16 \\ |\pi_3(K)| & = & 8 \end{array}$$

In particular, the index of B in $\mathcal{G}_1$ is at least 8 and the index of K in B is at least 2.

Proof.

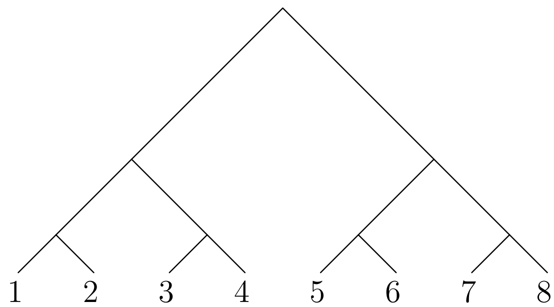

We consider the canonical homomorphism from $\mathrm{Aut}(T_2^\bullet(3))$ to $\mathrm{Perm}(1, \dots, 8)$ and compute the images of the generators σ, β, and δ:

$$\begin{array}{rcl} \sigma & \mapsto & (1,5)(2,6)(3,7)(4,8) \\ \beta & \mapsto & (1,3)(2,4)(5,6) \\ \delta & \mapsto & (5,6) \end{array}$$

It requires a finite amount of work to check that the image is isomorphic to

$$(C_2 \wr C_2) \wr C_2.$$

The order of this group is 128.

For the elements t, v, and w, we have

$$\begin{array}{rcl} t & \mapsto & (1,4,2,3)(5,7,6,8) \\ v & \mapsto & (1,2)(3,4) \\ w & \mapsto & (5,6)(7,8) \end{array}$$

Now the amount of work for determining the images of K and B is finite, too. **q.e.d.**

Another Technical Computation 10.4.5. *We write v, and w as products of conjugates of t and t^{-1}.*

$$\begin{array}{rcl} w & = & t\delta t^{-1}\delta \\ v & = & \sigma w \sigma \end{array}$$

So v and w lie in the normal span of t.

Proposition 10.4.6. *The subgroup B has index 8 in $\mathcal{G}_1$ and is the normal subgroup generated by β.*

Proof. We know that B is normal and therefore contains the normal closure of β. To prove the other inclusion, we observe that t lies clearly in the normal span of β. We already saw (10.4.5) that u and v lie in the normal closure of t.

Now, we determine the index of B. Since B is normal, this amount to compute the size of the group $\mathcal{G}_1/B$. Since β dies in this quotient, δ and γ become equal whence $\mathcal{G}_1/B$ is actually a factor of $\langle\sigma,\delta\rangle$ which has order 8. Hence the index of B is at most 8. We saw in (10.4.4) that the index is at least 8. **q.e.d.**

Exercise 10.4.7. Put $D := \langle(\sigma,\delta),(\delta,\sigma)\rangle \leq \mathcal{G}_1 \otimes \mathcal{G}_1$. Show that the image of $\mathcal{G}_1(1)$ in $\mathcal{G}_1 \otimes \mathcal{G}_1$ is

$$(B \otimes B) \rtimes D$$

where the action is conjugation in each component. Infer that $\mathcal{G}_1(1)$ has index 8 in $\mathcal{G}_1 \otimes \mathcal{G}_1$.

Proposition 10.4.8. *The subgroup K has index 16 in $\mathcal{G}_1$ and is the normal closure of t. Moreover, K contains $\mathcal{G}_1(3)$ as a subgroup of index 8.*

Proof. It follows from (10.4.5) that K is the normal closure of t. In (10.4.4) we saw that its index in B is at least 2. Hence it suffices to show that B/K has order ≤ 2. This, however, is clear as the quotient is generated by the image of β which has order 2.

It follows that K is actually the preimage of its order 8-image in $\mathrm{Aut}(T_2^\bullet(3))$ which implies that $\mathcal{G}_1(3)$ is a normal subgroup of K of index 8. **q.e.d.**

Corollary 10.4.9. *For any element $\xi \in \mathcal{G}_1$, the coset $K\xi$ can be effectively computed.*

Proof. Since the coset $\mathcal{G}_1(3)\,\xi$ depends only on the action of ξ on the vertices of level 3, the proposition follows from the fact that $\mathcal{G}_1(3)$ is a finite index subgroup in K. **q.e.d.**

The subgroup K is very important because this is the fractal part of $\mathcal{G}_1$ whose self-similarity is at the heart of virtually all theorems about the First Grigorchuk Group.

Proposition 10.4.10. *$K \otimes K$ is a subgroup of K of index 4.*

Proof. First, we observe

$$w = \sigma\beta\sigma\delta\sigma\beta\sigma\delta = (\gamma,\sigma)(1,\beta)(\gamma,\sigma)(1,\beta) = (1,t).$$

From this we get $1 \otimes K \leq K$ as follows: For any $\xi \in \mathcal{G}_1$, there is at least one partner $\zeta \in \mathcal{G}_1$ such that $(\zeta, \xi) \in \mathcal{G}_1(1)$. Hence, we have

$$(\zeta, \xi)(1, t)(\zeta, \xi)^{-1} = (1, \xi t \xi^{-1}) \in K.$$

As conjugation by σ swaps the coordinates, we also have $K \otimes 1 \leq K$. Hence $K \otimes K \leq K$.

As for the index, we quote (10.4.7). We know that K has index 8 in $\mathcal{G}_1(1)$ which has index 8 in $\mathcal{G}_1 \otimes \mathcal{G}_1$. This accounts for an index of 64. On the other hand, $K \otimes K$ has index $256 = 16^2$ in $\mathcal{G}_1 \otimes \mathcal{G}_1$. Hence it is of index 4 in K. **q.e.d.**

Corollary 10.4.11. *For any four elements $\xi_1, \xi_2, \zeta_1, \zeta_2 \in \mathcal{G}_1$, consider the pairs (ξ_1, ζ_1) and (ξ_2, ζ_2) in $\mathcal{G}_1 \otimes \mathcal{G}_1$. If $\xi_1 \equiv \xi_2 \mod K$ and $\zeta_1 \equiv \zeta_2 \mod K$ then*

$$(\xi_1, \zeta_1) \in \mathcal{G}_1 \iff (\xi_2, \zeta_2) \in \mathcal{G}_1.$$

Proof. Clear since $K \otimes K \leq K \leq \mathcal{G}_1$ and $K \otimes K \leq \mathcal{G}_1 \otimes \mathcal{G}_1$. **q.e.d.**

Remark 10.4.12. The statements (10.4.9) and (10.4.11) allow us to decide algorithmically if a pair (ξ, ζ) of elements in $\mathcal{G}_1$ defines another element of $\mathcal{G}_1$.

Remark 10.4.13. The group K has the strange property of containing its square as a subgroup of finite index. Amazingly, one can do even better: Not only does there exists a finitely generated group G that is isomorphic to $G \times G$, but every finitely generated group can be embedded into such a self-similar finitely generated group [Meie82].

The self-similar structure of K rules out polynomial growth:

Corollary 10.4.14. *$\mathcal{G}_1$ does not have polynomial growth.*

Proof. Clearly $\mathcal{G}_1$ and K have the same growth. On the other hand, K contains an isomorphic copy of itself of infinite index. This contradicts polynomial growth by (1.4.27). **q.e.d.**

10.4.1 Finite 2-groups

Proposition 10.4.15. *Let $H \leq K$ be a subgroup. Then K contains an isomorphic copy of $H \wr C_2$.*

Proof. Here is the idea: Since $H \otimes H \leq K \otimes K \leq K$ we have a copy of $H \wr C_2$ inside $\mathcal{G}_1$, namely

$$(H \otimes H) \rtimes \langle \sigma \rangle .$$

Unfortunately, $\sigma \notin K$. This is the problem we have to fix.

We need an element in K that acts somewhat like σ. So, note that K contains the following element:

$$\begin{aligned}
\tilde{\sigma} &:= t^4 \\
&= (\gamma\sigma, \sigma\gamma)^4 \\
&= ((\sigma\delta, \delta\sigma), (\delta\sigma, \sigma\delta))^2 \\
&= (((\beta, \beta), (\beta, \beta)), ((\beta, \beta), (\beta, \beta))) \\
&= ((((\sigma, \gamma), (\sigma, \gamma)), ((\sigma, \gamma), (\sigma, \gamma))), (((\sigma, \gamma), (\sigma, \gamma)), ((\sigma, \gamma), (\sigma, \gamma)))) \in \mathcal{G}_1(4) .
\end{aligned}$$

This is an element of order 2 and spans a copy of C_2.

Now let us work on the subgroup H. Using $K \otimes K \leq K$, we descend to $K^{32} \leq K$ and find two isomorphic copies of H, namely,

$$H_0 := H \otimes 1 \otimes 1 \otimes \cdots \otimes 1 \leq K^{32}$$

and

$$H_1 := 1 \otimes H \otimes 1 \otimes \cdots \otimes 1 \leq K^{32}.$$

These copies commute, hence $H_0 H_1 \cong H \times H$. Moreover,

$$\tilde{\sigma} H_0 \tilde{\sigma} = H_1.$$

Hence we have

$$H \wr C_2 \cong (H_0 H_1) \rtimes \langle \tilde{\sigma} \rangle \leq K$$

as desired. **q.e.d.**

Corollary 10.4.16. *Every finite 2-group embeds into $\mathcal{G}_1$. In particular, $\mathcal{G}_1$ is not periodic.*

Proof. Every finite 2-group embeds into an iterated wreath product $((C_2 \wr C_2) \wr \cdots \wr C_2) \wr C_2$. This follows by induction from the following two facts:

1. [10.4.17] Every finite 2-group surjects onto C_2.

2. [10.4.18] Every extension G of N by a finite group Q injects into $N \wr Q$. **q.e.d.**

Lemma 10.4.17. *Let G be a finite 2-group. Then there is a surjective homomorphism $G \twoheadrightarrow C_2$.*

Proof. We have $|G| = 2^n$. Consider the set

$$M := \left\{ S \subset G \,\middle|\, 2^{n-1} = |S| \right\}$$

and let G act on it by

$$gS = \{gs \mid s \in S\}.$$

It is obvious that G cannot fix an element of M. An easy induction proves that

$$|M| = \binom{2^n}{2^{n-1}} \equiv 2 \mod 4.$$

Since powers of 2 are the only possible sizes for orbits, the congruence shows that one of them has length 2. This induces the desired homomorphism. **q.e.d.**

Lemma 10.4.18. *Let $N \hookrightarrow G \twoheadrightarrow Q$ be a short exact sequence of groups with Q finite. Then there is an injective homomorphism $G \hookrightarrow N \wr Q$.*

Proof. Choose a set-theoretic section $\sigma : Q \to G$. It is easy to check that

$$g \mapsto \left(\left(\sigma^q g \sigma^{q\overline{g}} \right)_{q \in Q}, \overline{g} \right)$$

defines an injective homomorphism. Here $\overline{g}$ denotes the image of g in the quotient Q. **q.e.d.**

Now, we can prove that $\mathcal{G}_1$ is not linear without quoting a big theorem.

Proposition 10.4.19. *A finitely generated linear torsion group G is periodic.*

Proof. Using Mal'cev's observation (10.2.35), we assume that $G \leq \mathrm{GL}_n(k)$ where k is a purely transcendental extension of its prime field k_0.

Consider $g \in G \leq \mathrm{GL}_n(k)$ and let $\mu_g(t)$ be its minimal polynomial. We aim to show that only finitely many polynomials arise this way. This will imply the proposition since the order of an endomorphism only depends on its minimal polynomial.

Since g is torsion, we have, say, $g^n = 1$. Since every root of $\mu_g(t)$ is an eigenvalue for g, we find, that roots of $\mu_g(t)$ are roots of unity. Hence these roots are algebraic integers over k_0. The coefficients of μ_g are elementary symmetric functions of the roots. Since k is a purely transcendental extension of k_0, these coefficients belong to k_0. For finite k_0 we are done now.

For $k_0 = \mathbb{Q}$, everything takes place inside $\mathbb{C}$. Thus, we can talk about absolute values. Roots of unity have absolute value 1. This implies a bound on the absolute values of coefficients of $\mu_g(t)$. On the other hand, these coefficients are integers, as we have seen. Therefore, we have only finitely many numbers from which to chose our coefficients. **q.e.d.**

Corollary 10.4.20. *$\mathcal{G}_1$ is not linear.*

10.4.2 Congruence Subgroups

Put $\mathcal{G}_1(s) := \ker(\mathcal{G}_1 \to \mathrm{Aut}(T_2^\bullet(s)))$.

Observation 10.4.21. *Since every non-trivial tree automorphism has to act non-trivially on some finite subtree, we have*

$$\bigcap_{s\geq 1} \mathcal{G}_1(s) = 1.$$

In particular, $\mathcal{G}_1$ is residually finite.

Definition 10.4.22. A congruence subgroup of $\mathcal{G}_1$ is a group that contains the groups $\mathcal{G}_1(s)$. for s large enough.

Remark 10.4.23. This definition is reminiscent of arithmetic groups. The congruence subgroups of $\mathrm{SL}_n(\mathbb{Z})$ are those normal subgroups that contain a kernel of a congruence-homomorphism $\mathrm{SL}_n(\mathbb{Z}) \to \mathrm{SL}_n(\mathbb{Z}^m)$. It turns out that for $n \geq 3$, every non-central normal subgroup in $\mathrm{SL}_n(\mathbb{Z})$ is a congruence subgroup.

We aim to prove

Theorem 10.4.24 (Congruence Subgroup Property). *Every non-trivial normal subgroup of $\mathcal{G}_1$ is a congruence subgroup.*

Remark 10.4.25. This is to say that any quotient of $\mathcal{G}_1$ is a quotient of one of the finite groups $\mathrm{Aut}(T_2^\bullet(s))$. In particular, all proper quotients of $\mathcal{G}_1$ are finite, i.e., it is just infinite.

Definition 10.4.26. A group is just infinite if it is infinite but all its proper quotients are finite.

Lemma 10.4.27. *The canonical homomorphism*

$$\mathcal{G}_1(2) \to (\mathcal{G}_1 \otimes \mathcal{G}_1) \otimes (\mathcal{G}_1 \otimes \mathcal{G}_1) = {\mathcal{G}_1}^4$$

is onto in each coordinate.

Proof. Check

$$\begin{aligned}
((\sigma,\gamma),(\gamma,\sigma)) &= \gamma\sigma\delta\sigma\beta \\
((\gamma,\sigma),(\sigma,\gamma)) &= \sigma\gamma\sigma\delta\sigma\beta\sigma \\
((\beta,\beta),(\beta,\beta)) &= \sigma\gamma\sigma\gamma\sigma\gamma\sigma\gamma
\end{aligned}$$

and recall that $\mathcal{G}_1 = \langle \sigma, \beta, \gamma \rangle$. **q.e.d.**

Lemma 10.4.28. *For any level s, the action of $\mathcal{G}_1$ on the set of vertices of level s is transitive.*

Proof. This is an easy induction. Since $\sigma \in \mathcal{G}_1$, we are done for $s = 1$.

On level $s \geq 2$, we have a left half and a right half. The swap σ interchanges theses halves. Hence we only have to see that the action on the left half is transitive. This follows by induction from the fact that for any $\xi \in \mathcal{G}_1$, there is a partner $\zeta \in \mathcal{G}$ such that $(\xi, \zeta) \in \mathcal{G}_1(1)$. So we can act transitively on the left half by going into the left subtree and choosing an appropriate ξ. **q.e.d.**

Proposition 10.4.29. *For the commutator subgroup of K, we have*

$$\mathcal{G}_1(5) \leq [K, K] \leq \mathcal{G}_1(3)\,.$$

Proof. Note that $[K, K]$ is normal in $\mathcal{G}_1$ since inner automorphism take commutators to commutators. As a normal subgroup, it is generated by the commutators of the three generators t, v, and w. So let us compute these commutators first:

$$\begin{aligned}
[v, w] &= [(t, 1), (1, t)] \\
&= 1 \\
[v, t] &= [(t, 1), (\gamma\sigma, \sigma\gamma)] \\
&= (w, 1) \\
&= ((1, t), (1, 1)) \\
\xi := \left[w, t^{-1}\right] &= [(1, t), (\sigma\gamma, \gamma\sigma)] \\
&= (1, w) \\
&= ((1, 1), (1, t))
\end{aligned}$$

From this computation we infer:

- $[K, K] \leq \mathcal{G}_1(3)$ since $t \in \mathcal{G}_1(1)$.
- $(1 \otimes 1) \otimes (1 \otimes K) \leq [K, K]$. This follows from (10.4.27) and the fact that K is the normal closure of t.

On the other hand, we have

$$\begin{aligned}
\sigma\xi\sigma &= ((1, t), (1, 1)) \\
\beta\sigma\xi\sigma\beta &= ((t, 1), (1, 1)) \\
\sigma\beta\sigma\xi\sigma\beta\sigma &= ((1, 1), (t, 1))
\end{aligned}$$

whence we have

$$\begin{array}{rcl} (1 \otimes K) \otimes (1 \otimes 1) & \leq & [K, K] \\ (K \otimes 1) \otimes (1 \otimes 1) & \leq & [K, K] \\ (1 \otimes 1) \otimes (K \otimes 1) & \leq & [K, K] \end{array}$$

It follows that $K \otimes K \otimes K \otimes K \leq [K, K]$. Thus

$$\mathcal{G}_1(5) \leq \mathcal{G}_1(4) \otimes \mathcal{G}_1(4) \leq \mathcal{G}_1(3)^4 \leq K^4 \leq [K, K]$$

which proves the claim. **q.e.d.**

Exercise 10.4.30. Prove that

$$K \otimes K \otimes K \otimes K = [K, K]$$

and determine the index of $[K, K]$ in K.

Lemma 10.4.31. *Fix $\xi = (\xi_1, \ldots, \xi_{2^t}) \in \mathcal{G}_1(t) - \mathcal{G}_1(t+1)$ such that $\xi_1 \notin \mathcal{G}_1(1)$. Then, for any two elements $\kappa_1, \kappa_2 \in K$, we have*

$$[[\xi, \kappa_1'], \kappa_2'] = \left(\left(\left[\kappa_1^{-1}, \kappa_2\right], 1\right), 1, \ldots, 1\right)_{2^t}$$

where

$$\kappa_i' := ((\kappa_i, 1), (1, 1), \ldots, (1, 1))_{2^t} \in K.$$

Proof. We observe $\xi_1 = (\zeta_0, \zeta_1)\sigma$, and compute

$$[\xi, \kappa_1'] = \left((\kappa_1^{-1}, \zeta_1 \kappa_1^{-1} \zeta_1^{-1}), 1, \ldots, 1\right)_{2^t}$$

and the claim follows. **q.e.d.**

Corollary 10.4.32. *$\mathcal{G}_1$ has trivial center.*

Proof. As $\mathcal{G}_1$ acts transitively on the vertices of any fixed level (10.4.28), any non-trivial element in $\mathcal{G}_1$ is conjugate to an element satisfying the hypotheses of (10.4.31). For a central element, however, the double commutator had to be trivial which contradicts $\mathcal{G}_1(5) \leq [K, K]$ (10.4.29). **q.e.d.**

Now we are ready to prove that $\mathcal{G}_1$ has the congruence subgroup property.

Proof of Theorem (10.4.24). Let N be a normal subgroup and $\xi \in N$ be a non-trivial element.

As $\bigcap \mathcal{G}_1(s) = 1$, we find t such that $\xi \in \mathcal{G}_1(t) - \mathcal{G}_1(t+1)$. Since the action of $\mathcal{G}_1$ on the vertices of any fixed level is transitive (10.4.28), we can conjugate ξ such that it satisfies the hypotheses of (10.4.31). For arbitrary elements κ_1 and κ_2 of K, we have

$$[[\xi, \kappa_1'], \kappa_2'] = \big(([\kappa_1^{-1}, \kappa_2], 1), 1, \ldots, 1\big)_{2^t}.$$

This element belongs to the normal closure of ξ. Since $\mathcal{G}_1(5) \leq [K, K]$, it follows that

$$\mathcal{G}_1(5) \otimes 1 \otimes \cdots \otimes 1 \leq N$$

where we have 2^{t+1} factors. Using (10.4.28) again, it follows that

$$\mathcal{G}_1(5) \otimes \cdots \otimes \mathcal{G}_1(5) \leq N.$$

This implies $\mathcal{G}_1(t+1+5) \leq N$. **q.e.d.**

10.5 The Weight of Elements

To study Grigorchuk's group $\mathcal{G}_1$, it is useful to use a modified metric on the Cayley graph instead of the ordinary word metric.

Definition 10.5.1. Assign weights to the generators as follows:

$$\sigma \mapsto 3, \quad \beta \mapsto 5, \quad \gamma \mapsto 4, \quad \delta \mapsto 3.$$

The weight of a word w in the generators is the sum of the weights of its letters. The weight $\|\xi\|$ of an element $\xi \in \mathcal{G}_1$ is the weight of a minimum weight word representing ξ.

Observation 10.5.2. *We can regard the generators as one letter words or as group elements. A priori the weight as words might be different from their weights as elements. However, since every two letter word has weight at least* 6*, the two notions of weight coincide on the generators.* **q.e.d.**

Observation 10.5.3. *Weights are sub-multiplicative:*

$$\|\xi\zeta\| \leq \|\xi\| + \|\zeta\| \text{ for all } \xi, \zeta \in \mathcal{G}_1. \qquad \textbf{q.e.d.}$$

Observation 10.5.4. *The weight function defines a metric on the Cayley graph that is bi-lipschitz equivalent to the word metric.* **q.e.d.**

Definition 10.5.5. We call a word in the letters $\{\sigma, \beta, \gamma, \delta\}$ reduced if the letter σ alternates with the other letters.

Observation 10.5.6. *Any minimum weight word is reduced.* **q.e.d.**

Remark 10.5.7. Also, any minimum length word is reduced.

The importance of the weight stems from the following:

Lemma 10.5.8 (Weight Reduction). *Let* $\xi = (\xi_0, \xi_1) \in \mathcal{G}_1(1)$ *, then*

$$\|\xi_0\| + \|\xi_1\| \leq \frac{7}{8} \|\xi\| + 3.$$

In particular,

$$\|\xi_0\| + \|\xi_1\| < \|\xi\| \qquad \text{for } \|\xi\| > 24,$$

which applies to all but finitely many elements.

Proof. Let w be a reduced four-letter word. Since σ occurs twice in w it represents a pair (w_0, w_1) where w_i is a two letter word or a single letter. It is easy to check that

$$\|w_0\| + \|w_1\| \leq \frac{7}{8} \|w\| .$$

(The worst case is the word $\sigma\beta\sigma\beta = (\gamma\sigma, \sigma\gamma)$.)

If w is reduced, has an even number of σ-occurences and length at most three, we can still check that

$$\|w_0\| + \|w_1\| \leq \frac{7}{8} \|w\| + 3.$$

To finish the proof, let u be a minimum weight word representing ξ. Subdivide u into blocks of length four possibly with a shorter block at the end. Apply the preceeding considerations block by block. The claim follows. **q.e.d.**

Observation 10.5.9. *Minimum weight representatives can be algorithmically computed, i.e., given a word* w*, we can compute a minimum length word* u *representing the same group element as* w*.* **q.e.d.**

10.5.1 The Conjugacy Problem

Theorem 10.5.10. *The conjugacy problem for $\mathcal{G}_1$ is solvable.*

We simplify the proof given in [GrWi00]. Put

$$Q(\xi, \zeta) := \{K\chi \mid \xi^\chi = \zeta\}.$$

Obviously, two elements ξ and ζ are conjugate if and only if the finite set $Q(\xi, \zeta)$ of K-cosets is non-empty. We will show that this set can be computed.

Observation 10.5.11 (mixed case). *If $\xi = (\xi_0, \xi_1)$ and $\zeta = (\zeta_0, \zeta_1)\sigma$, then ξ and ζ are not conjugate. In this case*

$$Q(\xi, \zeta) = Q(\zeta, \xi) = \emptyset.$$ **q.e.d.**

The following two lemmas show how to compute $Q(\xi, \zeta)$ in the other two (pure) cases: recall that by (10.4.11), we can effectively check whether a pair (χ_0, χ_1) defines an element of $\mathcal{G}_1$.

Lemma 10.5.12. *If $\xi = (\xi_0, \xi_1)$ and $\zeta = (\zeta_0, \zeta_1)$, then*

$$Q(\xi, \zeta) = Q_1 \cup Q_\sigma$$

with:

$$\begin{array}{rcl} Q_1 & := & \{K(\chi_0, \chi_1) \mid K\chi_i \in Q(\xi_i, \zeta_i) \ \text{ and } (\chi_0, \chi_1) \in \mathcal{G}_1\} \\ Q_\sigma & := & \{K(\chi_0, \chi_1)\sigma \mid K\chi_i \in Q(\xi_i, \zeta_{1-i}) \ \text{ and } (\chi_0, \chi_1) \in \mathcal{G}_1\} \end{array}$$

Proof. First, we compute the elements $\chi = (\chi_0, \chi_1)$ in $Q(\xi, \zeta)$. We find

$$\begin{array}{rl} & \xi^\chi = \zeta \\ \iff & (\xi_0, \xi_1)^{(\chi_0, \chi_1)} = (\zeta_0, \zeta_1) \\ \iff & \left| \begin{array}{l} \xi_0^{\chi_0} = \zeta_0 \\ \xi_1^{\chi_1} = \zeta_1 \end{array} \right| \text{ and } (\chi_0, \chi_1) \in \mathcal{G}_1 \end{array}$$

Similarly, we compute the elements $\chi = (\chi_0, \chi_1)\sigma$ in $Q(\xi, \zeta)$. Here, we find

$$\begin{array}{rl} & \xi^\chi = \zeta \\ \iff & (\xi_0, \xi_1)^{(\chi_0, \chi_1)\sigma} = (\zeta_0, \zeta_1) \\ \iff & (\xi_0, \xi_1)^{(\chi_0, \chi_1)} = (\zeta_1, \zeta_0) \\ \iff & \left| \begin{array}{l} \xi_1^{\chi_0} = \zeta_0 \\ \xi_0^{\chi_1} = \zeta_1 \end{array} \right| \text{ and } (\chi_0, \chi_1) \in \mathcal{G}_1 \end{array}$$

From these two computations, the claim follows. **q.e.d.**

Lemma 10.5.13. *If $\xi = (\xi_0, \xi_1)\sigma$ and $\zeta = (\zeta_0, \zeta_1)\sigma$, then*

$$Q(\xi, \zeta) = Q_1 \cup Q_\sigma$$

with:

$$\begin{aligned} Q_1 &:= \left\{ K(\chi_0, \xi_0^{-1}\chi_0\zeta_0) \,\middle|\, K\chi_0 \in Q(\xi_0\xi_1, \zeta_0\zeta_1) \ \textit{and}\ (\chi_0, \xi_0^{-1}\chi_0\zeta_1) \in \mathcal{G}_1 \right\} \\ Q_\sigma &:= \left\{ K(\chi_0, \chi_1)\sigma \,\middle|\, K\chi_0 \in Q(\xi_0\xi_1, \zeta_1\zeta_0) \ \textit{and}\ (\chi_0, \xi_0^{-1}\chi_0\zeta_1) \in \mathcal{G}_1 \right\} \end{aligned}$$

Proof. This proof is similar. However, the computation is a little more tricky. First we compute the elements $\chi = (\chi_0, \chi_1)$ in $Q(\xi, \zeta)$. We find

$$\begin{aligned} & \xi^\chi = \zeta \\ \iff\ & (\xi_0, \xi_1)\sigma^{(\chi_0,\chi_1)} = (\zeta_0, \zeta_1)\sigma \\ \iff\ & (\chi_0^{-1}, \chi_1^{-1})(\xi_0, \xi_1)\sigma(\chi_0, \chi_1) = (\zeta_0, \zeta_1)\sigma \\ \iff\ & (\chi_0^{-1}, \chi_1^{-1})(\xi_0, \xi_1)(\chi_1, \chi_0)\sigma = (\zeta_0, \zeta_1)\sigma \\ \iff\ & \left| \begin{array}{c} \chi_0^{-1}\xi_0\chi_1 = \zeta_0 \\ \chi_1^{-1}\xi_1\chi_0 = \zeta_1 \end{array} \right| \text{ and } (\chi_0, \chi_1) \in \mathcal{G}_1 \\ \iff\ & \left| \begin{array}{c} \chi_0^{-1}\xi_0\chi_1 = \zeta_0 \\ \chi_0^{-1}\xi_0\xi_1\chi_0 = \zeta_0\zeta_1 \end{array} \right| \text{ and } (\chi_0, \chi_1) \in \mathcal{G}_1 \\ \iff\ & \left| \begin{array}{c} \chi_1 = \xi_0^{-1}\chi_0\zeta_0 \\ \chi_0^{-1}\xi_0\xi_1\chi_0 = \zeta_0\zeta_1 \end{array} \right| \text{ and } (\chi_0, \chi_1) \in \mathcal{G}_1 \\ \iff\ & K\chi_0 \in Q(\xi_0\xi_1, \zeta_0\zeta_1) \text{ and } (\chi_0, \xi_0^{-1}\chi_0\zeta_0) \in \mathcal{G}_1 \end{aligned}$$

The case $\chi = (\chi_0, \chi_1)\sigma$ is similar and left to the reader. **q.e.d.**

Exercise 10.5.14. Verify that $K\xi = K\zeta$ implies $K(\xi, \chi) = K(\zeta, \chi)$, $K(\chi, \xi) = K(\chi, \zeta)$, $K(\xi, \chi)\sigma = K(\zeta, \chi)\sigma$, and $K(\chi, \xi)\sigma = K(\chi, \zeta)\sigma$.

Proof of (10.5.10). We will show that the set $Q(\xi, \zeta)$ can be effectively computed. Here the two elements are given as words in the generators. Since we can alogrithmically compute minimum weight representatives, we will always assume that group elements are given as minimum weight words in the generators. Now the algorithm is as follows:

1. If $\|\xi\| \leq 27$ and $\|\zeta\| \leq 27$, return the value for

$$Q(\xi, \zeta)$$

   from a finite table.

2. If ξ and ζ act differently on the first level of the infinite binary tree (i.e., one swaps, the other does not), then return the empty set for $Q(\xi, \zeta)$. This condition is easy to check: The parity of σ-occurences in reduced words determines whether an element acts trivially or as the swap on the first level.

3. If ξ and ζ both act trivially on the first level, rewrite
$$\xi = (\xi_0, \xi_1)$$
and
$$\zeta = (\zeta_0, \zeta_1)$$
which can be done algorithmically. Recurse and compute Q_1 and Q_σ as in (10.5.12). Return
$$Q(\xi, \zeta) := Q_1 \cup Q_\sigma.$$

4. If ξ and ζ both act non-trivially on the first level, rewrite
$$\xi = (\xi_0, \xi_1)\sigma$$
and
$$\zeta = (\zeta_0, \zeta_1)\sigma$$
which can be done algorithmically. Recurse and compute Q_1 and Q_σ as in (10.5.13). Return
$$Q(\xi, \zeta) := Q_1 \cup Q_\sigma.$$

The algorithm terminates because of the Weight Reduction Lemma (10.5.8). **q.e.d.**

10.5.2 Intermediate Growth

Theorem 10.5.15 (Grigorchuk). *$\mathcal{G}_1$ has intermediate growth.*

Proof. We already observed that $\mathcal{G}_1$ does not have polynomial growth (10.4.14). Thus, we have to prove that its growth is subexponential. This will be done in (10.5.17). **q.e.d.**

Grigorchuk's group $\mathcal{G}_1$ was the first group of intermediate growth; and to this day, most constructions for groups of intermediate growth are based on the principles used in the design of $\mathcal{G}_1$.

Let G be a group with fixed finite generating set Σ. The growth function $\beta_\Sigma(n)$, defined in (1.4.14), counts the number of group elements that can be connected to the trivial element in the Cayley graph by a path of length at most n. Thus it depends on the generating set. We have seen in (1.4.18) that different generating sets still yield weakly equivalent (1.4.16) growth functions. In particular, it is independent of the generating set whether a group has exponential, intermediate, or polynomial growth; and in the case of polynomial growth, the degree of the polynomial also does not depend on the generating set.

This argument extends to growth functions based on the weight-metric instead of the word-metric:

Observation 10.5.16. *Let $\beta'(n)$ be the number of elements in $\mathcal{G}_1$ of weight $\leq n$, and let $\beta(n)$ be the number of elements in $\mathcal{G}_1$ of word-length $\leq n$. Then*

$$\beta(n) \leq \beta'(5n) \text{ for all } n \in \mathbb{Z}_0$$

since every group element of length n has weight at most $5n$: recall that the maximum weight assigned to a generator is 5. **q.e.d.**

Thus, to prove sub-exponential growth for Grigorchuk's group, it suffices to show:

Proposition 10.5.17. *The function β' is sub-exponential.*

Since we have to rule out exponential growth, we need a better understanding of this particular growth type. We will see that exponential growth functions in groups can be characterized by a real number, their growth-rate.

Observation 10.5.18. *The weight growth function β' is sub-multiplicative, i.e.,*

$$\beta'(n_1 + n_2) \leq \beta'(n_1)\,\beta'(n_2) \text{ for all } n_1 \text{ and } n_2.$$

The reason is, of course, that any word of weight $\leq n_1 + n_2$ is a product of two shorter words of weights at most n_1 and n_2, respectively. **q.e.d.**

Exercise 10.5.19. Let $a_1, a_2, a_3, \ldots$ be a positive sub-additive sequence, i.e, suppose $a_{i+j} \leq a_i + a_j$. Show that the Cesaro-limit $\lim_{i\to\infty} \frac{a}{i}$ exists and satisfies

$$\lim_{i\to\infty} \frac{a}{i} = \inf_i \frac{a_i}{i}.$$

Corollary 10.5.20. *Since β' is positive and sub-multiplicative, its logarithm is sub-additive. Hence, we can infer that*

$$\lambda := \lim_{n\to\infty} \sqrt[n]{\beta'_\Sigma(n)} = \inf_{n\geq 1} \sqrt[n]{\beta'(n)}$$

exists. We call this number, the growth rate of β'. Obviously, $\mathcal{G}_1$ has exponential growth if only if the growth rate of β' is strictly bigger than 1.

Observation 10.5.21. *Note that λ satisfies*

$$\lambda^n \leq \beta'(n)\,.$$

Moreover, any $\lambda' > \lambda$ yields an upper bound:

$$\beta'(n) \leq \lambda'^n \text{ for } n \text{ large enough.}$$

Multiplying by some constant, we can always turn this into a universal upper bound:

$$\beta'_\Sigma(n) \leq C\lambda^n \text{ for all } n.$$

q.e.d.

This discussion shows that the growth in a group cannot oscilate between to bounds: it is impossible that we have, say, the inequalities

$$2^n \leq \beta'(n) \leq 3^n$$

and that each bound is sharp infinitely many times.

Remark 10.5.22. The reasoning applies verbatim to growth functions based on word-metrics, as well. Be warned, however, that the growth-rate is not an intrisic invariant of the group: different generating sets (and more generally, different length functions on the group) will yield different growth rates.

Now we can complete the proof of Theorem (10.5.15).

Proof of (10.5.17). We have to show that $\beta'(n)$ is sub-exponential. The reason is that there is no consistent choice for a growth rate. Thus, we obtain a contradiction to (10.5.20).

Assume the function β' is exponential with growth rate $\lambda > 1$. Fix a number L slightly bigger than λ. Then we have

$$\lambda^n \leq \beta'(n) \leq O(L^n)\,.$$

Since every element of $\mathcal{G}_1 - \mathcal{G}_1(1)$ can be pushed into $\mathcal{G}_1(1)$ by multiplication by σ, we find:

$$\begin{array}{rcll}
\lambda^n & \leq & \beta'(n) & \\
& \leq & \left|\{\xi \in \mathcal{G}_1(1) \mid \|\xi\| \leq n+3\}\right| & \\
& \leq & \sum_{i+j<\frac{7}{8}(n+3)+3} \beta'(i)\,\beta'(j) & \text{by (10.5.8)} \\
& \leq & \left(\frac{7}{8}\,(n+3)+3\right)^2 O\left(L^{\frac{7}{8}(n+3)+3}\right) & \text{by counting terms} \\
& = & O\left(L^{\frac{7}{8}n}\right). &
\end{array}$$

This says that L is not just slightly bigger than λ. A real contradiction can be obtained by making the notion of "slightly bigger" more precise. **q.e.d.**

10.6 Amenability

Since $\mathcal{G}_1$ has intermediate growth, we have the following:

Corollary 10.6.1. *$\mathcal{G}_1$ is amenable.*

Proof. Follows from (10.5.15) and (1.4.23). **q.e.d.**

Definition 10.6.2. The class EG of elementary amenable groups is the smallest class containing all finite groups and the integers that is closed with respect to

1. taking subgroups
2. taking quotients
3. forming extensions
4. forming direct unions

The class NF consists of all groups that do not contain a non-abelian free subgroup.

We denote the class of all amenable groups by AG.

Remark 10.6.3. From what we know already, it is immediate that

$$EG \subseteq AG \subseteq NF.$$

We will prove that $\mathcal{G}_1$ is not elementary amenable. The standard proof quotes the fact that elementary amenable groups cannot have intermediate growth [Chou80]. Our approach, however, will be based on congruence subgroups.

Observation 10.6.4. *All abelian groups are elementary amenable.* **q.e.d.**

This observation gives rise to the following stratification of the class EG. Let EG_0 be the class of groups that are finite or abelian. For any ordinal $\alpha > 0$ define

$$EG_\alpha := \left\{ G \;\middle|\; G \text{ is an extension or a direct union of groups in } \bigcup_{\beta<\alpha} EG_\beta \right\}.$$

Proposition 10.6.5. *Each stratum EG_α is closed with respect to subgroups and quotients.*

Proof. Let $N \hookrightarrow G \twoheadrightarrow F$ be a short exact sequence of groups, and let G_0 be a subgroup of G. Then the following is a short exact sequence:

$$N_0 \hookrightarrow G_0 \twoheadrightarrow F_0$$

where $N_0 := N \cap G_0$ and $F_0 \leq F$ is the image of G_0 in F.

Moreover, if G_0 is normal in G, then N_0 is normal in N and F_0 is normal in F and the following diagram has an exact botton row by the 3×3 lemma

$$\begin{array}{ccccc} N_0 & \longrightarrow & G_0 & \longrightarrow & F_0 \\ \downarrow & & \downarrow & & \downarrow \\ N & \longrightarrow & G & \longrightarrow & F \\ \downarrow & & \downarrow & & \downarrow \\ N/N_0 & \longrightarrow & G/G_0 & \longrightarrow & F/F_0 \end{array}$$

From this, it follows that if G is included in EG_α because it is an extension of groups in lower strata, then all its subgroups and quotients will be in EG_α, too.

Dealing with direct unions is easy and we will leave it to the reader. **q.e.d.**

Corollary 10.6.6. *$EG = \bigcup_\alpha EG_\alpha$.*

Proof. The key observation is that $\bigcup_\alpha EG_\alpha$ is closed with respect to subgroups, quotients, extensions, and direct unions. Subgroups and quotients are immediate from (10.6). Extensions are easy: you might have to go up one stratum. The only case that requires thought is how to deal with a direct union. So let $G = \bigcup_i G_i$ where each $G_i \in EG_{\alpha_i}$ for some α_i. This is, where set theory kicks in and tells you that there is an ordinal at least as big as each of these α_i. So everything actually takes place in one (sufficiently high) stratum. Then you have to move up once more, and you will find G. **q.e.d.**

After these preparations, let us start with a simple consequence of (10.4.31).

Lemma 10.6.7. *Fix $\xi = (\xi_1, \ldots, \xi_{2^t}) \in \mathcal{G}_1(t) - \mathcal{G}_1(t+1)$, and let N be the normal closure of ξ in $\mathcal{G}_1(t)$. Then, N contains an embedded copy of $\mathcal{G}_1(5)$.*

Proof. Let us assume for a moment that $\xi_1 \notin \mathcal{G}_1(1)$. In this case, (10.4.31) says that for any two $\kappa_1, \kappa_2 \in K$, there are $\kappa_1', \kappa_2' \in \mathcal{G}_1(t+1)$ such that

$$[[\xi, \kappa_1'], \kappa_2'] = \big(([\kappa_1^{-1}, \kappa_2], 1), 1, \ldots, 1\big)_{2^t}.$$

Thus,

$$[K, K] \otimes 1 \otimes \cdots \otimes 1 \subseteq N.$$

Since $\mathcal{G}_1(5)$ embeds into $[K, K]$, the claim follows.

Now let us deal with the case $\xi_1 \in \mathcal{G}_1(1)$: By (10.4.28)we can conjugate ξ inside $\mathcal{G}_1$ such that $\xi_1 \notin \mathcal{G}_1(1)$. Since $\mathcal{G}_1(t)$ is normal in $\mathcal{G}_1$, this operation puts us in the situation above. **q.e.d.**

Theorem 10.6.8. *For any s, the congruence subgroup $\mathcal{G}_1(s)$ is not elementary amenable. In particular, $\mathcal{G}_1$ is amenable but not elementary amenable.*

Proof. Let α be minimal with $\mathcal{G}_1(s) \in EG_\alpha$ for some level s. Since the stratum EG_α is closed with respect to subgroups, we may assume $s \geq 5$.

Note that $\mathcal{G}_1(s)$ is not a direct union of groups from lower strata since it is finitely generated: after finitely many steps, all the generators are cought whence the direct union is a finite ascending union. Thus, $\mathcal{G}_1(s)$ is an extension

$$N \hookrightarrow \mathcal{G}_1(s) \twoheadrightarrow F$$

where N and F belong to lower strata.

Since N is a non-trivial normal subgroup in $\mathcal{G}_1(s)$, (10.6.7) implies that there is an embedding $\mathcal{G}_1(s) \leq \mathcal{G}_1(5) \hookrightarrow N$. Thus, N cannot belong to a lower stratum since α was minimal. **q.e.d.**

10.7 Presentations

Here, we will outline a strategy for obtaining a presentation of $\mathcal{G}_1$. The starting point is the group

$$\Gamma := \left\langle \sigma, \beta, \gamma, \delta \,\middle|\, \sigma^2 = \beta^2 = \gamma^2 = \delta^2 = \beta\gamma\delta = 1 \right\rangle \cong C_2 * \mathbb{V}$$

which has an obvious homomorphism onto $\mathcal{G}_1$. Let Γ_1 and Γ_K denote the preimages of $\mathcal{G}_1(1)$ and K, respectively, and meditate on the following diagram:

$$\begin{array}{ccccccc}
N & \hookrightarrow & \Gamma & \xrightarrow{\pi} & \mathcal{G}_1 & & \\
 & & \uparrow & & \uparrow & & \\
\Gamma \times \Gamma & \xleftarrow{\psi = (\psi_{\mathbf{L}}, \psi_{\mathbf{R}})} & \Gamma_1 & \twoheadrightarrow & \mathcal{G}_1(1) & \rightarrow & \mathcal{G}_1 \otimes \mathcal{G}_1 \\
\uparrow & & \uparrow & & \uparrow & & \uparrow \\
\Gamma_K \times \Gamma_K & & \Gamma_K & \twoheadrightarrow & K & \leftarrow & K \otimes K
\end{array}$$

Our goal is do determine the kernel N of the epimorphism π.

Let $N_0 := 1 \leq \Gamma_1 \leq \Gamma$, and put

$$N_{i+1} := \psi^{-1}\left(N_i \times N_i\right).$$

Note that N_1 is the kernel of ψ. Put

$$N_\infty := \bigcup_i N_i.$$

Lemma 10.7.1. $N = N_\infty$.

Proof. The inclusion $N_\infty \leq N$ is clear. To see the other inclusion, let w be a reduced word in N. Note that for $|w| \geq 3$, we have

$$|\psi_{\mathbf{L}}(w)| \leq \frac{|w|}{2} \quad \text{and} \quad |\psi_{\mathbf{R}}(w)| \leq \frac{|w|}{2}.$$

Since a non-trivial reduced word of length ≤ 2 in Γ is not in N, we infer that

$$\psi^{\lceil \log_2(|w|) \rceil}(w) = \left(1_1, \ldots, 1_{2^{\lceil \log_2(|w|) \rceil}}\right).$$

This implies $w \in N_{\lceil \log_2(|w|) \rceil} \leq N_\infty$. **q.e.d.**

Observation 10.7.2. *Since $\psi_{\mathbf{L}}(w) = \psi_{\mathbf{R}}(\sigma w \sigma)$ and $\beta, \gamma, \delta \in \Gamma_1$, all the N_i are normal in Γ and we have $N_{i+1} = \overline{\psi_{\mathbf{R}}^{-1}(N_i)}^{\Gamma}$.* **q.e.d.**

Our next goal is to find a subset R that generates N as a normal subgroup of Γ, so that we can turn the identity $\mathcal{G}_1 = \Gamma/N$ into a presentation for $\mathcal{G}_1$. Our starting point will be a finite set $R_1 \subset \Gamma$ that generates N_1 as a normal subgroup in Γ: To find R_1, note that the image of ψ is a finitely presented group since it has finite index in the finitely presented group $\Gamma \times \Gamma$. Hence we can find a finite presentation by the Reidemeister-Schreier method (A.1.12). Thus N_1 is finitely generated as a normal subgroup of Γ_1 and we can find explicitly a finite set R_0 that generates N_1 as a normal subgroup of Γ_1. Improving upon this, you can construct a finite, and possibly smaller, set that generates N_1 as a normal subgroup of Γ. We will not carry this out, but a good choice for R_1 is known.

Fact 10.7.3 ([Gri98]). *$R_1 = \left\{(\sigma\delta)^4, (\sigma\gamma)^8, (\sigma\delta\sigma\gamma\sigma\gamma)^4\right\}$ generates N_1 as a normal subgroup of Γ.*

The next step is to determine, in terms of R_1, our normal generating set R for N. The problem is: N_{i+1} is defined as a preimage, thus we have everything backwards. Taking iterated preimages does not lend itself to a nice description of a set R. The following fact will allow us to overcome this obstacle.

Miracle 10.7.4. *The letter substitution*

$$\begin{array}{rcl} \sigma & \mapsto & \sigma\gamma\sigma \\ \beta & \mapsto & \delta \\ \gamma & \mapsto & \beta \\ \delta & \mapsto & \gamma \end{array}$$

defines an endomorphism $\Phi : \Gamma \to \Gamma$ *which satifies*

$$\Phi(w) = (1, w) \qquad \textit{for all } w \in \Gamma_K.$$

To see this, just evaluate Φ *on* t, v, *and* w *and check that you obtain* $(1,t)$, $(1,v)$, *and* $(1,w)$.

It follows that N_{i+1} is the normal closure of $\Phi(N_i)$. Hence, N is the normal closure of $R := \bigcup_i \Phi^i \operatorname{im}(R_1)$. Thus, the observation that $\Phi\big((\sigma\delta)^4\big) = (\sigma\gamma)^8$ completes our sketchy proof of

Theorem 10.7.5 ([Lys85],[Gri98]). *The First Grigorchuk Group* $\mathcal{G}_1$ *has a presentation*

$$\mathcal{G}_1 = \left\langle \sigma, \beta, \gamma, \delta \,\middle|\, \sigma^2 = \beta^2 = \gamma^2 = \delta^2 = \beta\gamma\delta = \Phi^i\big((\sigma\delta)^4\big) = \Phi^i\big((\sigma\delta\sigma\gamma\sigma\gamma)^4\big) = 1,\ (i \geq 0) \right\rangle.$$

Definition 10.7.6. A group is co-Hopfian if every injective endomorphism is onto.

Exercise 10.7.7. The endomorphism $\Phi : \Gamma \to \Gamma$ descends to an injective endomorphism of $\mathcal{G}_1$. The index of the image is infinite. Thus $\mathcal{G}_1$ is not co-Hopfian.

Remark 10.7.8. As $\mathcal{G}_1$ is finitely generated and residually finite, it is automatically Hopfian (2.2.9).

Exercise 10.7.9. Show that, for all $i \geq 1$,

$$\Phi^i\big((\sigma\delta)^4\big) \in N_{i+1} - N_i.$$

Corollary 10.7.10. *For all* i, *the inclusion* $N_i \leq N_{i+1}$ *is strict. In particular,* $\mathcal{G}_1$ *is not finitely presented.*

Proof. If $\mathcal{G}_1$ was finitely presented, all but finitely many relators in the presentation (10.7.5) would be redundant. Hence the inclusions $N_i \leq N_{i+1}$ would stabilize. **q.e.d.**

Chapter 11

Thompson's Group F

!!! FIXME: This chapter is currently being revised. !!!

Thompson's Group F is a discrete subgroup of the homeomorphism group of the Cantor set. It has the following infinite presentation:

$$F := \langle x_0, x_1, \ldots \mid x_m x_i = x_i x_{m+1} \text{ for } i < m \rangle .$$

[11.3.4] The abelianization of F is $C_\infty \times C_\infty$. Thus, F is infinite.

[11.4.4] F acts homeomorphically on $\mathbb{R}$, $\mathbb{R}^\infty_{\geq 0} = [0, \infty)$ and $[0, 1]$.

[11.7.3] F is of type F_∞.

[11.4.23] F has a solvable word problem.

[11.8.1] The conjugacy problem in F is solvable.

[11.4.24] F has exponential growth.

[11.4.25] F has trivial center.

[11.5.1] The commutator subgroup of F is simple and all proper quotients of F are abelian.

[11.5.7] F is not residually finite.

[11.5.8] F is torsion free.

[11.6.1] F is not elementary amenable.

The standard reference for this group and some of its relatives is [CFP96].

11.1 Associativity, Trees, and Homeomorphisms

Recall first that there is a correspondence of finite rooted planar binary trees and parenthesized expressions. The word planar here means that the two children of a node are distinguished: one of them is the right child, the other one is the left child. Since all binary trees in this chapter will be rooted and planar, we will drop these attributes from now on.

We just give an example that should make the correspondence clear:

$A\,((BC)\,D)$ corresponds to 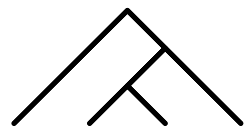 .

Note that we do not label the leaves of the tree with the letters from the expression since we do not consider the letters an essential part of the expression, we are only interested in the arrangement of parentheses. Thus the letters in a parenthesized expression are insignificant.

The law of associativity

$$(AB)\,C = A\,(BC)$$

is usually regarded as an equation that is supposed to hold in a given algebraic structure for all values of A, B, and C. However, we can also regard it as a transformation rule for parenthesized expressions:

ξ_0 : $(AB)\,C \quad \rightarrow \quad A\,(BC)$

 $\rightarrow$ 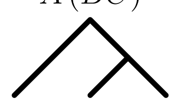

Of course, there is also the inverse transformation:

ξ_0^{-1} : $A\,(BC) \quad \rightarrow \quad (AB)\,C$

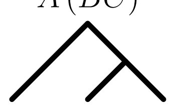 $\rightarrow$ 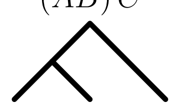

An easy induction on the length of expressions is used in every class on elementary group theory to prove:

Theorem 11.1.1. *Any parenthesized expression can be put in left-parenthesized normal form by a sequence of applications of the transformation ξ_0 and its inverse ξ_0^{-1}. As a consequence, we can move parentheses around at will.* **q.e.d.**

This theorem is not as precise as it could be: In the course of the sequence, we might have to apply the transformation ξ_0 and the inverse transformation to some subexpression deeply nested inside of the expression that we want to change. In fact, if we restricted the rules and **allow only transformations on the top-level expansion** of a parenthesized expression, the theorem would be false.

Example 11.1.2. The transformation

$$\xi_1 \; : \; \left\{ \begin{array}{ccc} A\,((BC)\,D) & \rightarrow & A\,(B\,(CD)) \\ & \rightarrow & \end{array} \right.$$

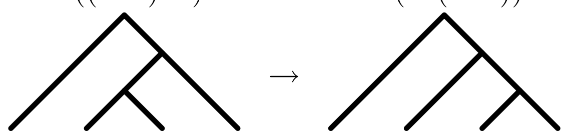

cannot be realized as a sequence of applications of the transformations ξ_0 and ξ_0^{-1} on the top-level expansion of parenthesized expressions.

Theorem 11.1.3. *Suppose we are allowed to apply the transformation rules $EffGenerator[0]$ and ξ_1 as well as their inverses freely to the* ***top level expansion*** *of every parenthesized expression, then we can put any expression into left-parenthesized normal form. Consequently, we can move parentheses around at will.*

Proof. The procedure is algorithmic. In a first step, we isolate the first letter as shown in this example:

$$\begin{array}{rl} & (((ab)(cd))e)f \\ \xrightarrow{x_0} & ((ab)(cd))(ef) \\ \xrightarrow{x_0} & (ab)((cd)(ef)) \\ \xrightarrow{x_0} & a(\underbrace{b((cd)(ef))}_{\text{right tail}}) \end{array}$$

For the second step, note that ξ_1 acts like ξ_0 applied to the right tail, i.e., we can apply ξ_0 to the right tail subexpression by applying ξ_1 to the top-level expression. If we had some magic sequence that, when applied to the top-level, acts as ξ_1 applied to the right tail, we could argue by induction on the length of the expression that the tail can be put in normal form.

Thus, we construct a substitution rule ξ_2 that acts like ξ_1 in the tail:

$$\xi_2 : \left\{ \begin{array}{rl} & a(\underbrace{A\,((BC)\,D)}_{\text{tail}}) \\ \xrightarrow{\xi_0^{-1}} & (aA)\,((BC)\,D) \\ \xrightarrow{\xi_1} & (aA)\,(B\,(CD)) \\ \xrightarrow{\xi_0} & a\,(A\,(B\,(CD))) \end{array} \right.$$

This completes the proof. **q.e.d.**

We have already illustrated the transformations ξ_0 and ξ_1 by pairs of binary trees. Such a pair actually describes a homemorphism of the Cantor set as follows: Recall that the Cantor set is the space of ends of the infinite binary tree.

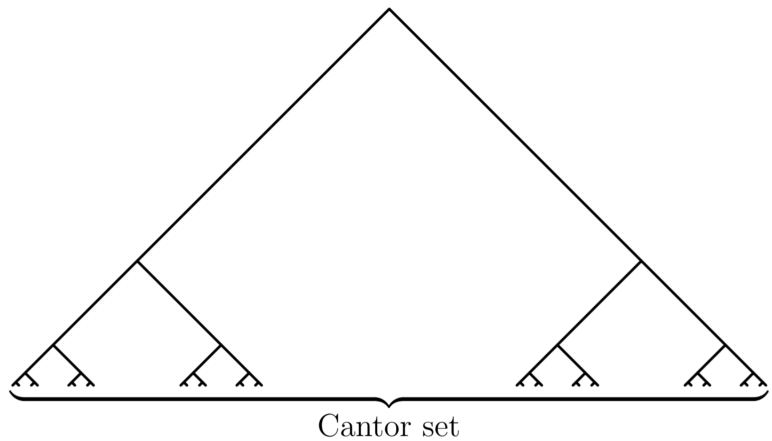

Every finite binary tree canonically embeds into the infinite binary tree: the root maps to the root, and for every node already embedded its right child maps to the right child of the image node and its left child maps to the left child of the image node. The leaves of a finite binary tree, therefore, correspond to subtrees in the inifite binary tree, which in turn determine subsets of the Cantor set, all of which are homemorphic copies of the Cantor set.

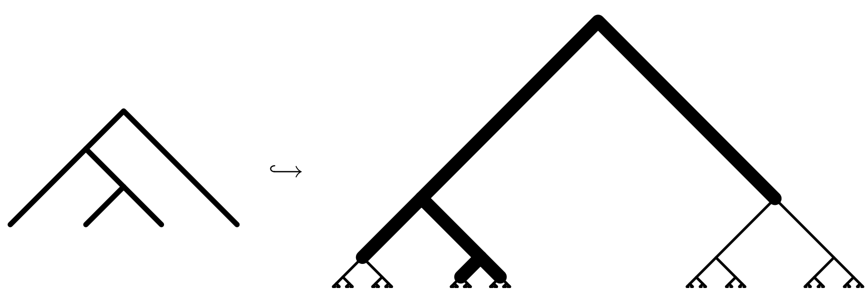

A pair of binary trees that have the same number of leaves induces two decompositions of the Cantor set into an equal number of Cantor subsets. The induces homeomorphism just identifies matching subsets in an order preserving way. The following picture illustrates the rule. Given a pair of finite binary trees T^{top} and T^{bot}, one of them represents the domain of the homeomorphism (this is the top tree) and the other one, representing the range of the homeomorphism, is drawn upside down (this is the bottom tree). Such a stack is called a tree diagram. This way, vertices match so that corresponding chunks of the cantor set are aligned. The right hand of the picture illustrates how the homeomorphism actually works whereas the left hand shows the allignment of leaf-vertices.

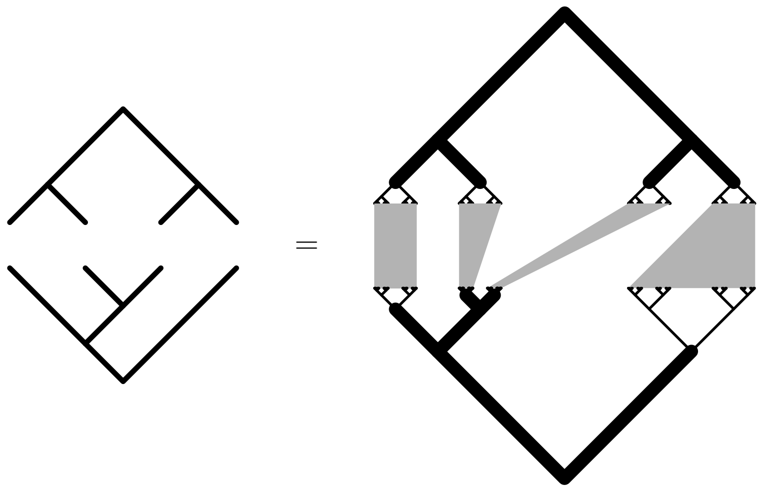

Definition 11.1.4. A homeomorphism of the Cantor set $\mathcal{C}$ that can be given by a finite tree diagram is called a finite $\mathcal{C}$-homeomorphism.

Example 11.1.5. We already met the finite $\mathcal{C}$-homeomorphisms

$\xi_0 =$ 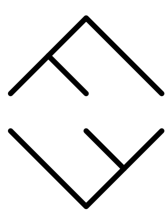 , $\xi_1 =$ 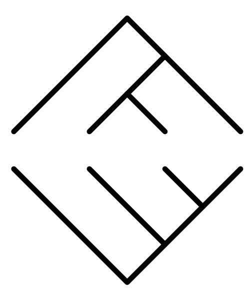

Note that there are many tree diagrams representing the same finite $\mathcal{C}$-

homeomorphism. The following picture shows how this happens.

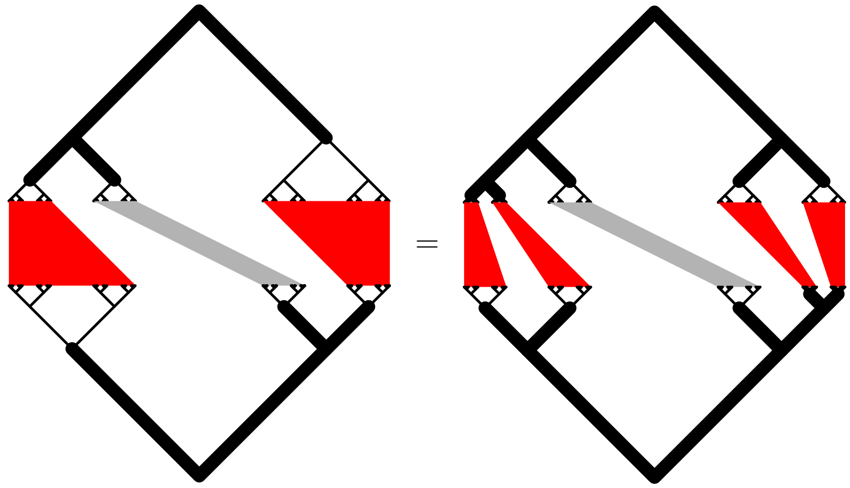

=

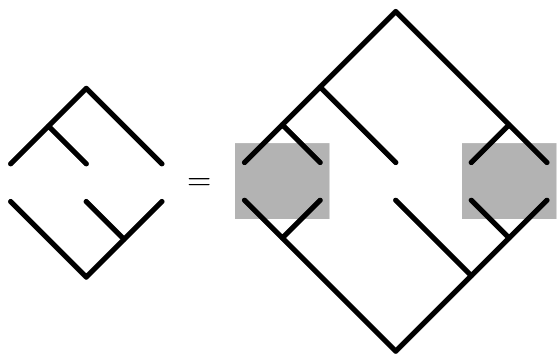

=

The diagram on the right is unreduced. We also call it a blow up of the left diagram. Every diagram that has a pair of opposing carets is unreduced and removing such a pair is called a reduction step. Every diagram can be reduced by a finite number of reduction steps.

Theorem 11.1.6. *Every finite $\mathcal{C}$-homeomorphism is represented by a unique reduced diagram.*

The proof introduces a good deal of useful terminology.

Proof. We only have to argue uniqueness. Define standard dyadic subsets of the Cantor set $\mathcal{C}$ recursively as follows: The set $\mathcal{C}$ is called the generation 0 standard dyadic subset of $\mathcal{C}$. A generation $s+1$ standard dyadic subset is a generation s standard dyadic subset of either the left or the right half of $\mathcal{C}$. We call the canonical homeomorphism identifying two standard dyadic subsets of generations s and t a standard map of degree $t-s$.

Given a finite $\mathcal{C}$-homeomorphism ξ, we say that a standard dyadic subset S of $\mathcal{C}$ is ξ-maximal if ξ restricts to a standard map on S (in particular, the ξ-image of S is also standard dyadic).

Note that every point of $\mathcal{C}$ is contained in precisely one of those ξ-maximal sets. Moreover, since ξ is finite, there are only finitely many ξ-maximal sets. They form a partition of $\mathcal{C}$.

Note that the vertices of the infinite binary rooted tree corresesopond to the standard dyadic subsets of $\mathcal{C}$. Now it is easy to see that there is a forest diagram whose top forest corresponds to the dyadic decomposition of $\mathcal{C}$ into ξ-maximal subsets and that every forest diagram for ξ reduces to this unique forest diagram. **q.e.d.**

Theorem and Definition 11.1.7. *The set F of finite $\mathcal{C}$-homeomorphisms forms a group. This group is called Thompson's group.*

Proof. First note that F is closed with respect to taking inverses: Interchanging the domain and the range tree of a tree diagram defines the inverse homeomorphism. It remains to show that F is closed under composition of homeomorphisms.

In order to multiply finite $\mathcal{C}$-homeomorphisms, we start with two tree diagrams $\mathfrak{t}_1 = (T_1^{\mathrm{top}}, T_1^{\mathrm{bot}})$ and $\mathfrak{t}_2 = (T_2^{\mathrm{top}}, T_2^{\mathrm{bot}})$. Let us first consider the case where $T_1^{\mathrm{bot}} = T_2^{\mathrm{top}}$. In this case, the bottom tree of the first factor (representing the image of the homeomorphism) determines the same decomposition of the Cantor set $\mathcal{C}$ as the top tree of the second factor (representing the domain of that homeomorphism). Thus the composition of the two homeomorphisms is represented by the diagram $(T_1^{\mathrm{top}}, T_2^{\mathrm{bot}})$. So if $T_1^{\mathrm{bot}} = T_2^{\mathrm{top}}$, the composition of the homeomorphisms is itself a finite $\mathcal{C}$-homeomorphism. The follwoing diagram illustrates what is happening here:

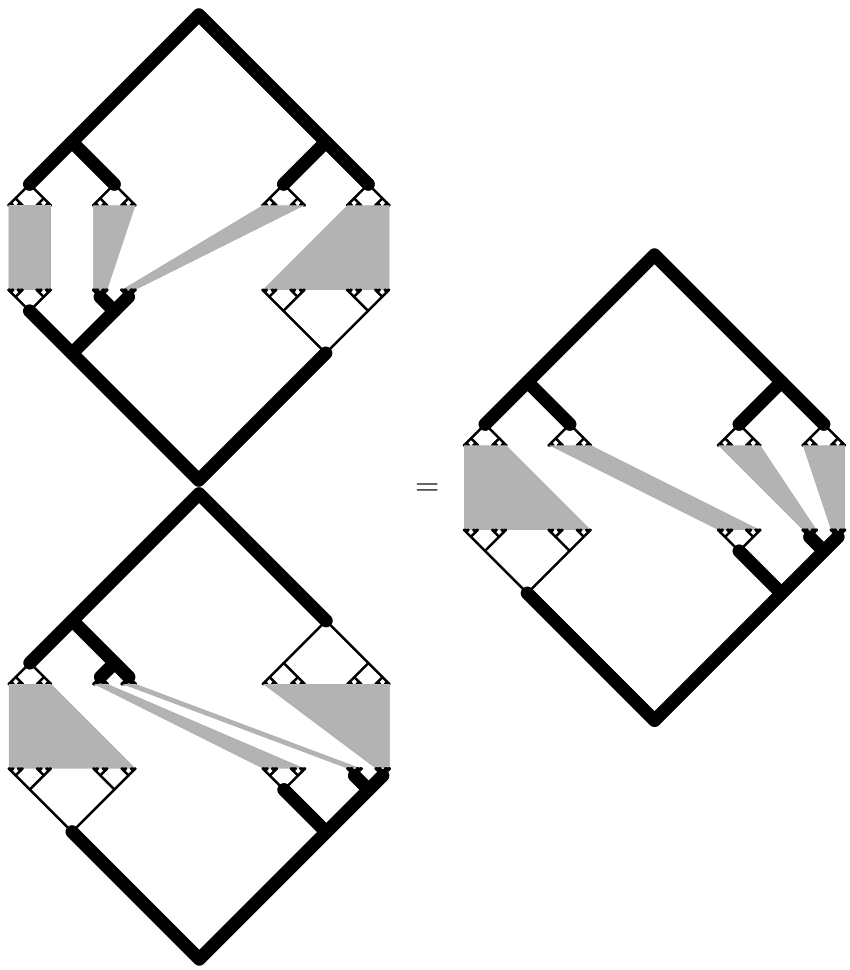

The general case now follows from our discussion of reducing and unreducing tree diagrams: Any two finite binary trees have a common blow up. The minimal common blow up is given as the union of their images in the infinite binary tree under their canonical embeddings. Thus, any two tree digrams can be unreduces so that the bottom tree of the first matches the top tree of the other. Again, we give a picture:

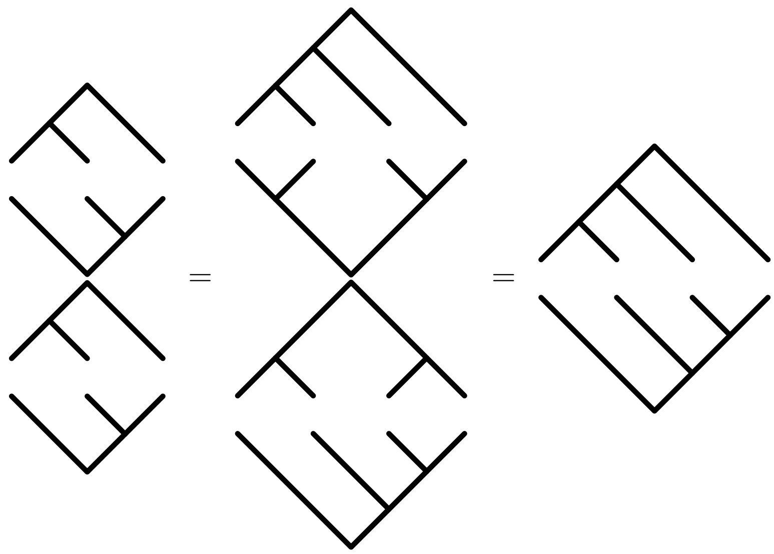

Please note that this proof is constructive and gives an algorithm for doing computations in Thompson's group F. **q.e.d.**

Remark 11.1.8. It follows immediately from Theorem 11.1.3 that the homeomorphism group generated by ξ_0 and ξ_1 contains all finite homeomorphisms. Thus, Thompson's group F is generated by the elements ξ_0 and ξ_1.

Remark 11.1.9. The way (un)reducing relates to multiplication suggests that Thompson's group F ought to e the group of fractions of a monoid. This is in fact true: For $i \geq 1$, put

$$\xi_i := \xi_0^{-i} \xi_1 \xi_0^{i}.$$

The tree diagram for ξ_i is as follows:

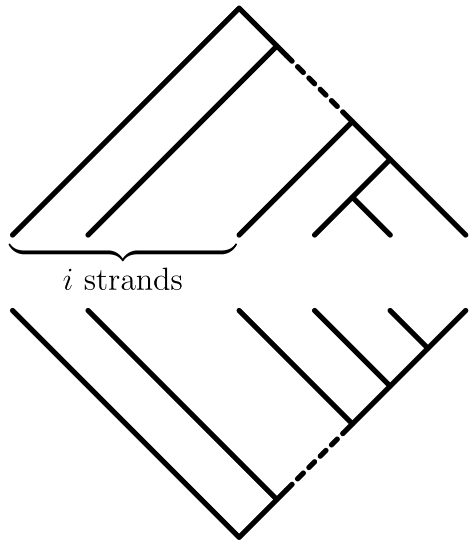

It will turn out that F is the group of fractions of the monoid generated by the ξ_i. (**??**).

Exercise 11.1.10. Show that the following relations hold among the elements ξ_i:

$$\xi_m \xi_i = \xi_i \xi_{m+1} \qquad \text{for } i < m.$$

11.2 The Positive Monoid and Forest Diagrams

Finite $\mathcal{C}$-homeomorphisms preserve the left-to-right order on the Cantor set. In particular both endpoints 0 and 1 are fixed. We will remove the right endpoint. Then we still have the complete left half of $\mathcal{C}$, which is a copy of $\mathcal{C}$. The right half is not quite complete. We only have its left quarter, which is another copy of $\mathcal{C}$. The right quarter has a missing point, and so it goes. Thus $\mathcal{C} - \{1\}$ is a countable infinite union of copies of $\mathcal{C}$:

$$\mathcal{C} - \{1\} = \bigcup_{i=0}^{\infty} \mathcal{C}_i = \mathbb{N}[0] \times \mathcal{C}$$

where $\mathcal{C}_i$ is a homeomorphic copy of $\mathcal{C}$. The copy $\mathcal{C}_i$ is to the left of $\mathcal{C}_j$ if $i < j$. The set $\mathcal{C} - \{1\}$ is the space of ends of the following forest with countably many binary trees:

We can now rewrite tree diagrams into forest diagrams. Given a finite $\mathcal{C}$-homeomorphism f and a tree diagram representing it, the right-most vertices in the tree diagram represent f-matching chunks of $\mathcal{C}$ that contain 1. Removing the right endpoint 1 turns these chunks into countable unions. Diagrammatically, the transformation of a tree diagram into a forest diagram it looks as follows:

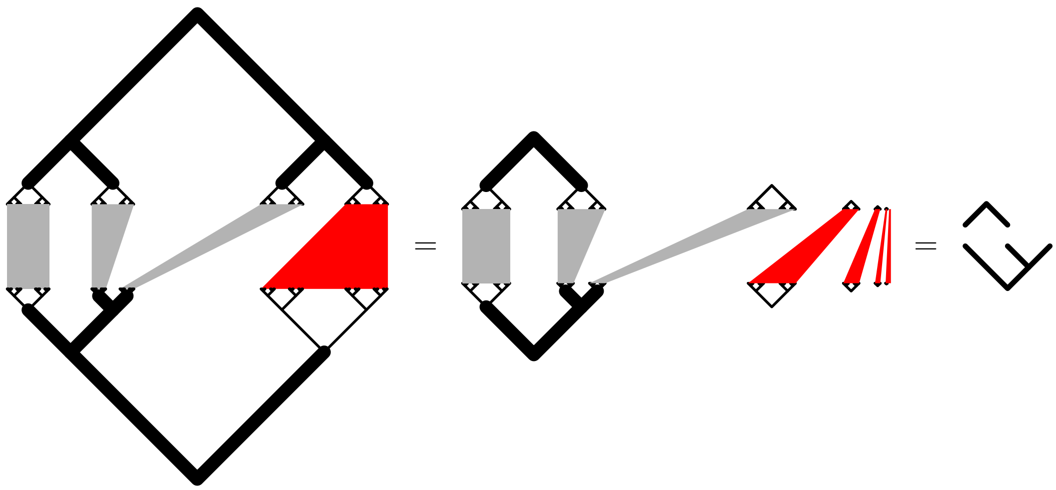

Note that the right-most material is redundant. Thus, we delete it from the schematic forest diagram far on the right.

Blow-up, reduction and multiplication of forest diagrams works precisely the way it does for tree diagrams:

Definition 11.2.1. A forest diagram that does not contain a matching pair of symmetric trees is called reduced.

Exercise 11.2.2. Show that every finite $\mathcal{C}$-homeomorphism is represented by a unique forest diagram.

An element of F is called positive if its reduced forest diagram has a 0-hight bottom forest, i.e., the bottom forest consits entirely of trivial trees.

Exercise 11.2.3. Show that the homeomorphism ξ_i is represented as follows:

$\xi_0 =$

$\xi_1 =$

$\xi_2 =$

$\vdots$

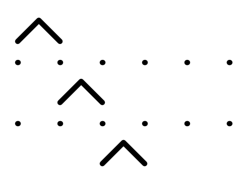

In particular, all the ξ_i are positive elements of F.

Observation 11.2.4. *In multiplying two positive elements, we have to blow up only the first factor. Thus the bottom forest of the second factor remains trivial. We conclude that the positive elements for a monoid inside F.* **q.e.d.**

We can actually be more precise and work out an easy rule for multiplying positve elements:

Observation 11.2.5. *Let (T_1, B_1) and (T_2, B_2) be two reduced forest diagrams of positive elements. Then the reduced forest diagram for the product has a trivial bottom forest and its top forest is given as follows:*

> *Number the leaves of T_1 from left to right. Number the roots of T_2 from left to right. Identitify vertices that have been assigned equal numbers.*

The reason is that this is precisely how to construct the blow up of the first tree diagram whose bottom forest matches the top forest T_2. **q.e.d.**

To put this as a slogan: In multiplication of positive element, you put the left factor on top of the right factor.

Lemma 11.2.6. *The reduced forest diagram for the element*

$$\xi_0^{a_0} \xi_1^{a_1} \cdots \xi_r^{a_r}$$

has trivial bottom forest and its top forest is given by the following criterion:

> *The leaf at position i has a right-ascending edge path of length a but not of length $a + 1$.*

Proof. First, let us illustrate the criterion by one example:

$\xi_2 =$ 

$\xi_2^2 =$ 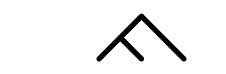

$\xi_0 \xi_2^2 =$ 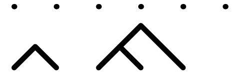

$\xi_0^2 \xi_2^2 =$ 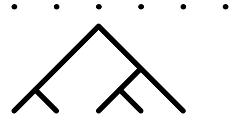

This already is a proof: The criterion holds for the elements ξ_i and since we are multiplying from the left by elements whose indices do not exceed the indices that have been dealt with, we do not mess up the criterion in each step: recall that multiplication from the left by ξ_i is dropping a caret whose left foot connects to the ith root of the forest. **q.e.d.**

Proposition 11.2.7 (Normal Form for Elements of the Positive Monoid). *Every positive element of F can be expressed uniquely as a product*

$$\xi_0^{a_0}\xi_1^{a_1}\cdots\xi_r^{a_r}$$

with non-negative exponents a_i.

Proof. This follows immediately from (11.4.18) since we can encode every forest diagram with trivial bottom forest by the associated sequence of exponents; and vice versa, every sequence of exponents defines such a forest diagram for a positive element. **q.e.d.**

Corollary 11.2.8 (Existence of Normal Forms in F). *Every element $f \in F$ can be expressed as*

$$\xi_0^{a_0}\xi_1^{a_1}\cdots\xi_r^{a_r}\xi_s^{-b_s}\cdots\xi_1^{-b_1}\xi_0^{-b_0}$$

such that the following conditions are met:

1. *$a_r \neq 0$ and $a_s \neq 0$.*

2. *$r \neq s$.*

3. *Whenever $a_i \neq 0$ and $b_i \neq 0$, then at least one of a_{i+1} or b_{i+1} is 0.*

Proof. Realize f by a reduced forest diagram, and let

$$\xi_0^{a_0}\xi_1^{a_1}\cdots\xi_r^{a_r}$$

and

$$\xi_0^{b_0}\xi_1^{b_1}\cdots\xi_s^{b_s}$$

be the normal forms for the top and bottom forest, respectively. Then

$$f = \xi_0^{a_0}\xi_1^{a_1}\cdots\xi_r^{a_r}\xi_s^{-b_s}\cdots\xi_1^{-b_1}\xi_0^{-b_0}.$$

The three conditions are restating the fact that our forest diagram is reduced: any violation would indicate a pair of opposing carets.

Conversely, any way of writing f in normal form translates into a reduced forest diagram for f, whence normal forms are unique as reduced forest diagrams are unique. **q.e.d.**

11.3 Presentations for Thompson's Group

The goal of this section is to find various presentations for Thompson's group F. The main result is:

Theorem 11.3.1. *Thompson's group has the following infinite presentation:*

$$F = \langle x_0, x_1, \ldots \mid x_m x_i = x_i x_{m+1} \textit{ for } i < m \rangle .$$

Proof. For the sake of comparison, we need to distinguish the two groups the theorem claims are equal. Thus, we put

$$G := \langle x_0, x_1, \ldots \mid x_m x_i = x_i x_{m+1} \text{ for } i < m \rangle .$$

By (11.1.10), the map

$$x_i \mapsto \xi_i$$

induces a homomorphism

$$G \to F.$$

The key observation is that these relations are sufficient to put every element into its normal form: Note that, for $i < m$, the following relations all hold in G:

$$\begin{aligned} x_m x_i &= x_i x_{m+1} \\ x_i^{-1} x_m^{-1} &= x_{m+1}^{-1} x_i^{-1} \\ x_i^{-1} x_m &= x_{m+1} x_i^{-1} \\ x_m^{-1} x_i &= x_i x_{m+1}^{-1} \end{aligned}$$

We want to write a given word in the form

$$x_0^{a_0} \cdots x_r^{a_r} x_r^{-b_r} \cdots x_0^{-b_0}$$

where

1. $a_i, b_i \geq 0$.
2. $0 \in \{a_r, b_r\} \neq \{0\}$.
3. If $a_i > 0 < b_i$, then $\{a_{i+1}, b_{i+1}\} \neq \{0\}$.

Note that we can use the last two relations to move the positive powers of generators to the left and the negative powers to the right. In a second step, we use the first two relations to sort the indices within the left and right half in increasing and decreasing orders respectively. We now achieved a rough normal form, wherein the positive and negative powers occur in the right order. Then, we do all possible cancellations at the center. This will ensure the second requirement. Finally, observe that the conjugacy relation

$$x_i x_{m+1}^{\pm 1} x_i^{-1} = x_m^{\pm 1}$$

allows us to ensure the final third requirement of normal forms.

We are now in a position to prove injectivity. So suppose that some word w in the generators x_i is mapped to the trivial element in F. Using the defining relations, we can put w into normal form. But the normal of the trivial element is the empty word. This implies that the relation $w = 1$ follows from the defining relations. **q.e.d.**

Exercise 11.3.2. Show that the following presentations also define Thompson's Group F:

$$\begin{aligned} F_{\text{finite,a}} &:= \left\langle a, b \,\middle|\, b^{aa} = b^{ab},\, b^{abb} = b^{abb^a} \right\rangle \\ F_{\text{finite,b}} &:= \left\langle c, d \,\middle|\, d^{cc} = d^{cd},\, d^{cdc} = d^{cdd} \right\rangle. \end{aligned}$$

Here, we use the convention $g^h := h^{-1} g h$.

Corollary 11.3.3. *Thompson's Group F is finitely presented.*

Proposition 11.3.4. *F is infinite. In fact, the abelianization is $F^{ab} = C_\infty \times C_\infty$.*

Proof. From the presentation, we have

$$\begin{aligned} F^{\text{ab}} &= \langle x_0, x_1, \ldots \mid x_i x_m = x_m x_i = x_i x_{m+1} \text{ for } i < m \rangle \\ &= \langle x_0, x_1, \ldots \mid x_0 x_1 = x_1 x_0,\ x_i = x_{i+1} \text{ for } 1 \leq i \rangle \end{aligned}$$

which is a presentation of $C_\infty \times C_\infty$. **q.e.d.**

11.4 An Action of Thompson's Group F

Definition 11.4.1. Let $\mathbb{R}_{\geq 0}^{\infty} := [0, \infty)$ be the positive real numbers including infinity. The Group $F(\frac{1}{2}, \mathbb{R}_{\geq 0}^{\infty})$ is the group of piecewise linear self-homeomorphism f of $\mathbb{R}_{\geq 0}^{\infty}$ satisfying:

1. Slopes change only finitely many times.
2. Slopes changes only a places in $\mathbb{Z}\left[\frac{1}{2}\right]$.
3. Slopes are in $2^{\mathbb{Z}}$.

4. There is an integer s such that
$$f^t = t + s$$
for all sufficiently large t.

Remark 11.4.2. Since there is a piecewise linear homeomorphisms $\mathbb{I} \to \mathbb{R}^{\infty}_{\geq 0}$ identifying the unit interval an $\mathbb{R}^{\infty}_{\geq 0}$, we have an isomorphism

$$F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0}) = F(\frac{1}{2}, \mathbb{I}) := \mathrm{PL}_{\mathbb{Z}[\frac{1}{2}], 2^{\mathbb{Z}}}(\mathbb{I})$$

of $F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ and the group $F(\frac{1}{2}, \mathbb{I})$ of piecewise linear self-homeomorphism of the unit-interval with finitely many break points all in $\mathbb{Z}\left[\frac{1}{2}\right]$ and slopes in $2^{\mathbb{Z}}$.

Note that $F(\frac{1}{2}, \mathbb{I})$ also is the strict subgroup of $F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ of those homeomorphism that are the identity on $[1, \infty)$.

Let us have a closer look at those piecewise linear self-homeomorphisms of $\mathbb{R}^{\infty}_{\geq 0}$:

Example 11.4.3. The following maps are in $F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$:

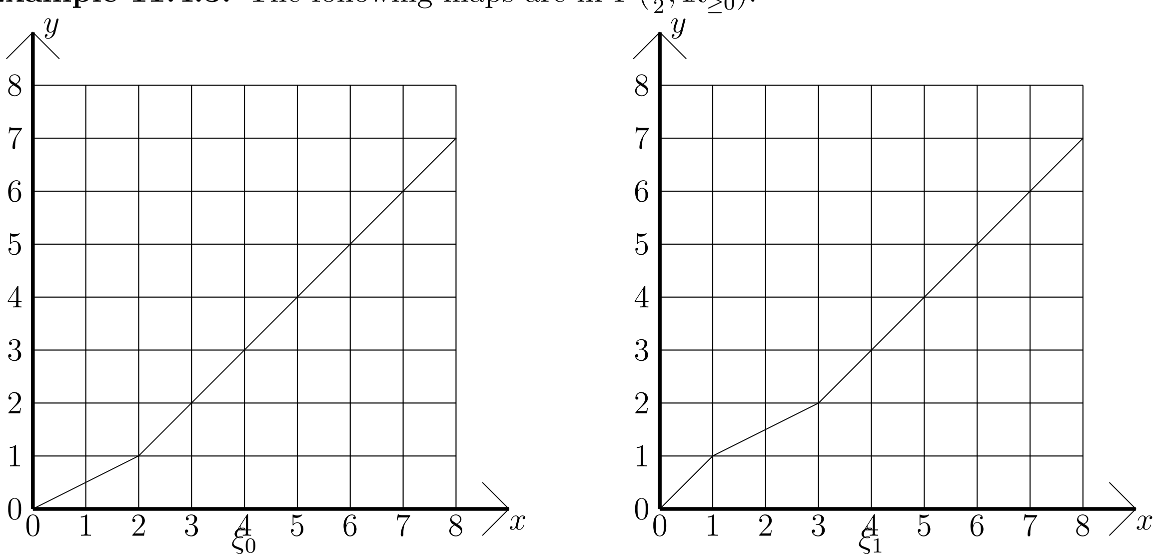

In general, we define ξ_i to be the function that has slope 1 on $[0, i]$ and $[i + 2, \infty]$ and has slope $\frac{1}{2}$ in $[i, i + 2]$.

Exercise 11.4.4. Prove that the map
$$x_i \mapsto \xi_i$$
extends to a group homomorphism
$$F \to F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0}).$$

11.4.1 Forest Diagrams

Definition 11.4.5. A standard dyadic interval is an interval of the form $\left[\frac{n}{2^m}, \frac{n+1}{2^m}\right]$ where n and m are non-negative integers. Let us call a standard dyadic interval of length 1 a standard unit interval. Note that this implies that its boundary points are integer.

A finitary dyadic decomposition of $\mathbb{R}^{\infty}_{\geq 0}$ is a set $\{I_0, I_1, \ldots\}$ of standard dyadic intervals satisfying:

1. $\mathbb{R}^{\infty}_{\geq 0} = \biguplus_{i \geq 0} I_i$.

2. For all i sufficiently large, I_i has length 1.

Let $\mathcal{I}$ and $\mathcal{J}$ be two finitary dyadic decompositions of $\mathbb{R}^{\infty}_{\geq 0}$. We say that $\mathcal{J}$ is a refinement of $\mathcal{I}$ – and that $\mathcal{I}$ is coarser than $\mathcal{J}$ – if every interval in $\mathcal{J}$ is contained in an interval from $\mathcal{I}$.

The trivial finitary dyadic decomposition is the set of standard unit intervals in $\mathbb{R}^{\infty}_{\geq 0}$. Every finitary dyadic decomposition is a refinement of the trivial one.

Observation 11.4.6. *Note that if you pick a standard dyadic interval in a finitary dyadic decomposition of $\mathbb{R}^{\infty}_{\geq 0}$ and replace it by its two halves, you obtain another finitary decomposition. We call this operation an elementary split. The result of a split is always a refinement. Vice versa, every refinement of a decomposition can be obtained by finitely many splits.*

Definition 11.4.7. An infinite binary forest is a directed tree where each vertex has two outgoing edges one of which is labeled as left whereas the other is labeled as right. The endpoints of these two edges are the children of the vertex. Each vertex is the parent of its children. A finite binary forest is a subgraph of an infinite binary forest all of whose vertices either have two children or none.

Example 11.4.8. The standard dyadic intervals form a forest $\mathcal{F}^{(2)}$. Each standard dyadic interval I is a vertex in $\mathcal{F}^{(2)}$ and its two children are the left and right half of I, respectively.

Observation 11.4.9. *Every finitary dyadic decomposition $\mathcal{I}$ of $\mathbb{R}^{\infty}_{\geq 0}$ defines a subforest $\mathcal{F}^{(2)}_{\mathcal{I}} \subseteq \mathcal{F}^{(2)}$: The intervals in $\mathcal{I}$ are vertices of $\mathcal{F}^{(2)}$ and $\mathcal{F}^{(2)}_{\mathcal{I}}$ is the minimal subforest containing all of them. Since a standard dyadic interval of length 2^{-m} corresponds to a vertex of distance m to a root vertex, you can actually read off the decomposition $\mathcal{I}$ from $\mathcal{F}^{(2)}_{\mathcal{I}}$ pretty easily by looking at the "spaces between the leafs" as indicated in this example:*

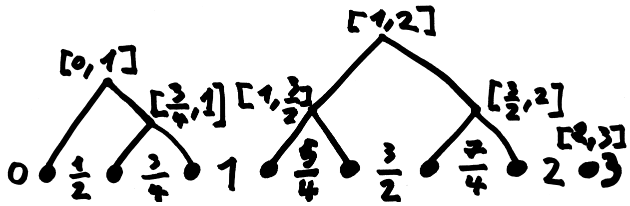

We are interested in finitary dyadic decompositions because they provide a convenient way of describing elements of $F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$. The reason is that if you map a standard dyadic interval linearly to another standard dyadic interval, the slope of the map is a power of 2. Hence a pair $(\mathcal{I}, \mathcal{J})$ of finitary dyadic decompositions describes an element $\varphi_{(\mathcal{I},\mathcal{J})} \in F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ as follows: The left-most interval in $\mathcal{I}$ is mapped linearly to the left-most interval in $\mathcal{J}$, and from there you proceed from left to right matching intervals in $\mathcal{I}$ with intervals in $\mathcal{J}$.

Definition 11.4.10. A pair of finitary standard dyadic decompositions of $\mathbb{R}^{\infty}_{\geq 0}$ is called a <u>forest diagram</u>.

Observation 11.4.11. *If $(\mathcal{I}, \mathcal{J})$ is a forest diagram representing φ, then $(\mathcal{J}, \mathcal{I})$ is a forest diagram that represents φ^{-1}.*

Because of the way the intervals are supposed to match up, we draw the forest diagram $(\mathcal{I}, \mathcal{J})$ as a top forest with leafs pointing down and a bottom forest with leafs pointing up such that leafs corresponding to matching intervals are aligned.

Example 11.4.12. The diagram

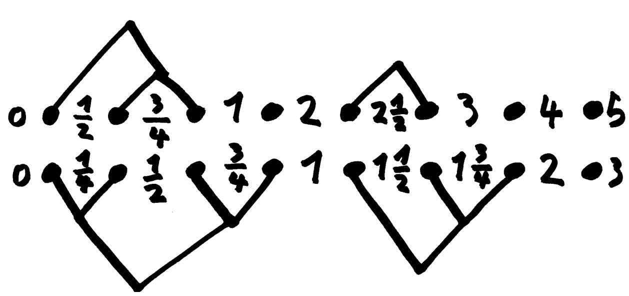

represents the piecewise linear map that interpolates the chart

$$\begin{array}{rcl}
0 & \mapsto & 0 \\
\frac{1}{4} & \mapsto & \frac{1}{2} \\
\frac{1}{2} & \mapsto & \frac{3}{4} \\
\frac{3}{4} & \mapsto & 1 \\
1 & \mapsto & 2 \\
1\frac{1}{2} & \mapsto & 2\frac{1}{2} \\
1\frac{3}{4} & \mapsto & 3 \\
2 & \mapsto & 4
\end{array}$$

and has slope 1 thereafter.

There are, of course, different forest diagrams that define the same element of $F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ – for instance, any symmetric forest diagram $(\mathcal{I}, \mathcal{I})$ defines the identity. More general, any forest diagram can be expanded without changing the homeomorphism it represents by simultaneously splitting matching top and bottom leafs. Conversely, deleting of matching symmetric subtrees will also not alter the homeomorphism.

Definition 11.4.13. A forest diagram that does not contain a matching pair of symmetric trees is called reduced.

Note that any non-reduced forest diagram contains a matching pair of carets:

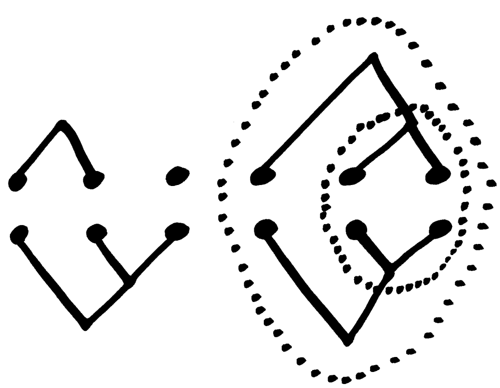

The importance of forest diagrams derives from the following

Theorem 11.4.14. *Every element $\varphi \in F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ is represented by a unique reduced forest diagram.*

Proof. Let I be a closed interval with dyadic rational endpoints such that φ is linear on I. Since the slope of $\varphi|_I$ is a power of 2, the image $\varphi\,\mathrm{im}(I)$ also has dyadic rational endpoints. Since a sufficiently high power of $\frac{1}{2}$ divides evenly into the endpoints of I and $\varphi\,\mathrm{im}(I)$, we can find a decomposition of I into standard dyadic intervals that maps to a decomposition of $\varphi\,\mathrm{im}(I)$ by standard dyadic intervals. Since φ has only finitely many breakpoints, existence of a forest diagram representing φ follows from the fact that for sufficiently large t, the map φ acts as an integer shift, which implies that "far to the right" we can describe φ by matching standard unit intervals.

Given existence of forest diagrams representing φ, we obtain a reduced one by, well, reducing any unreduced forest diagram for φ.

As for uniqueness, consider the maximum standard intervals that have the following properties:

1. The homeomorphism φ restricts to a linear map.

2. The image of the interval under φ is a standard dyadic interval.

Observe that these intervals form a finitary dyadic decomposition $\mathcal{I} := \mathcal{I}_\varphi$ of $\mathbb{R}^{\infty}_{\geq 0}$ which in turn defines a subforest $\mathcal{F}^{(2)}_\varphi$ of $\mathcal{F}^{(2)}$ that is contained in the bottom forest of any forest diagram for φ. On the other hand, it is obvious from the definition that φ takes the decomposition $\mathcal{I}$ to a finitary dyadic decomposition $\mathcal{J}$. It follows that $(\mathcal{I}, \mathcal{J})$ is a reduced forest diagram for φ to which any other forest diagram for φ reduces. **q.e.d.**

11.4.2 Normal Forms

We have yet to prove that normal forms are unique. Eventually, this will follow from (11.4.14), but we have to exhibit the interplay between reduced forest diagrams and normal forms first.

Example 11.4.15. The elements ξ_i is represented by the following forest diagram:

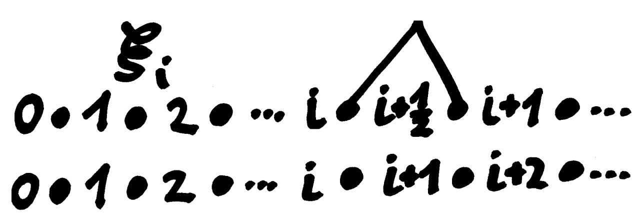

Definition 11.4.16. An element $\varphi \in F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ is called positive if its reduced forest diagram has a depth 0 bottom forest, i.e., the homeomorphism φ takes every standard unit interval standard dyadic interval.

Observation 11.4.17. *Positive elements in $F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ form a submonoid, i.e., the product of two positive elements is positive. In fact, it is easy to multiply the corresponding forest diagrams. Let φ and ψ be two positive elements with top forests $\mathcal{I}_\varphi$ and $\mathcal{I}_\psi$, respectively. As the roots of $\mathcal{I}_\psi$ correspond to standard unit intervals, they match with the vertices of the bottom forest for φ. Hence the top tree for the product $\varphi\psi$ is obtained from $\mathcal{I}_\varphi$ and $\mathcal{I}_\psi$ by identifying the leafs of $\mathcal{I}_\varphi$ with the roots of $\mathcal{I}_\psi$. Here is an example:*

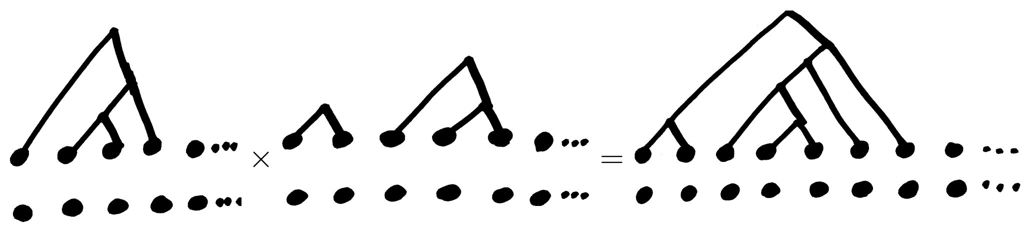

To put this as a slogan: In multiplication of positive element, you put the left factor on top of the right factor. An immediate consequence is the

Proposition 11.4.18. *The reduced forest diagram for the element*

$$\xi_0^{a_0} \cdots \xi_r^{a_r}$$

has trivial bottom forest and its top forest is given by the following criterion:

> *The leaf at position i has a right-ascending edge path of length a but not of length $a + 1$.*

Proof. First, let us illustrate the criterion by one example:

$\xi_2 =$ 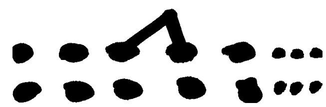

$\xi_2^2 =$ 

$\xi_0\xi_2^2 =$ 

$\xi_0^2\xi_2^2 =$ 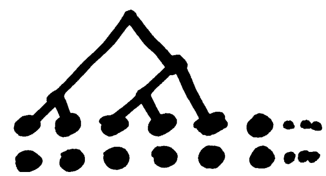

This already is a proof: The criterion holds for the elements ξ_i and since we are multiplying from the left by elements whose indices do not exceed the indices that have been dealt with, we do not mess up the criterion in each step. **q.e.d.**

Remark 11.4.19. Note that every forest is uniquely determined by the numbers a_i which give the maximum length of a right ascending edge path starting at leaf i. Conversely, any sequence of non-negative integers with finite support determines a forest. In particular, the set of positive elements is the submonoid generated by the ξ_i.

Observation 11.4.20. *In general, we have to think of a forest diagram as a quotient of two positive elements. Since we are not in an abelian setting, we cannot easily multiply quotient. The rule for multiplying two forest diagrams, therefore, is to unreduce both by splitting top an bottom leafs as to make the bottom forest of the left factor equal to the top forest of the of the right factor. Once this is done, the isomorphic forest cancel:*

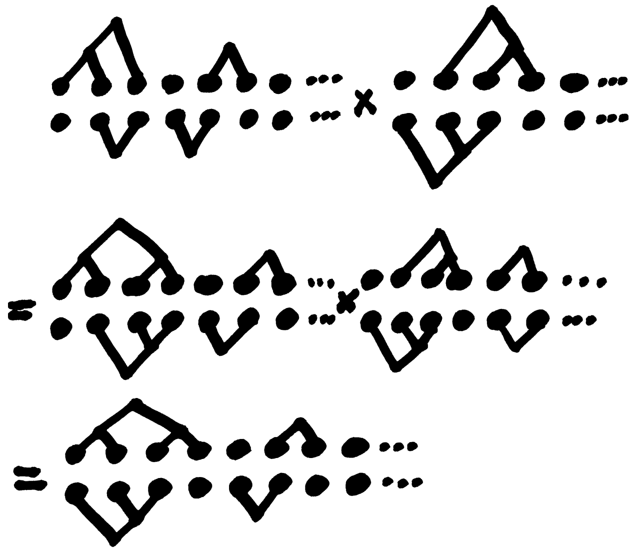

Corollary 11.4.21. *The homomorphism $F \to F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ is surjective.*

Theorem 11.4.22. *The normal forms of (??) are unique.*

Proof. First, recall that a normal form is a word of the form

$$x_0^{a_0} \cdots x_r^{a_r} x_r^{-b_r} \cdots x_0^{-b_0} = \left(x_0^{a_0} \cdots x_r^{a_r}\right)\left(x_0^{b_0} \cdots x_r^{b_r}\right)^{-1}$$

where

1. $a_i, b_i \geq 0$.
2. $0 \in \{a_r, b_r\} \neq \{0\}$.
3. If $a_i > 0 < b_i$, then $\{a_{i+1}, b_{i+1}\} \neq \{0\}$.

Condition (1) implies that both parts of the word give rise to a forest. Hence a normal form defines a forest diagram in an obvious way. The other two restrictions imply that this forest diagram is reduced. Since a reduced diagram is uniquely determined by the homeomorphism it defines (11.4.14), the claim follows. **q.e.d.**

This theorem has many consequences.

Corollary 11.4.23. *Thompson's group F has solvable word problem.* **q.e.d.**

Corollary 11.4.24. *Thompson's group F has exponential growth.*

Proof. x_0^{-1} and x_1 generate a free monoid inside F. To see this, write a word, say

$$x_0^{-1}x_0^{-1}x_0^{-1}x_1x_1x_0^{-1}x_0^{-1}x_1x_0^{-1}x_0^{-1}x_1x_1$$

and pull all the x_0^{-1} to the right as to obtain the normal form:

$$x_4x_4x_6x_8x_8x_0^{-7}.$$

As the subscripts remember how many letters x_0^{-1} passed by, two different words in the monoid have different normal forms and, therefore, represent different elements of F. **q.e.d.**

Corollary 11.4.25. *F has trivial center.*

Proof. Let $f = x_0^{a_0} \cdots x_r^{a_r} x_r^{-b_r} \cdots x_0^{-b_0}$ be a non-trivial element of F. We will show that it does not commute simultaneously with x_0 and x_1.

$\underline{a_0 > 0 < b_0}$: In this case, $x_0^{-1} f x_0$ is in normal form and visibly not equal to f.

$\underline{r > 0 \text{ and } b_0 = 0}$: Now, we have

$$\begin{aligned} x_0^{-1} f x_0 &= x_0^{-1} x_0^{a_0} x_1^{a_1} \cdots x_r^{a_r} x_r^{-b_r} \cdots x_1^{-b_1} x_0 \\ &= x_0^{a_0} x_2^{a_1} \cdots x_{r+1}^{a_r} x_{r+1}^{-b_r} \cdots x_2^{-b_1} \\ &\neq f. \end{aligned}$$

$\underline{r > 0 \text{ and } a_0 = 0}$: This is dealt with like the previous case.

$\underline{r = 0}$: Now, f is a power of x_0 and, therefore, does not commute with x_1. **q.e.d.**

Exercise 11.4.26. Show that the centralizer of x_1 in F is isomorphic to $F \times C_\infty$.

Exercise 11.4.27. Does the group

$$\langle \ldots x_{-2}, x_{-1}, x_0, x_1, x_2, \ldots \mid x_m x_i = x_i x_{m+1} \text{ for } i < m \rangle$$

embed into F?

11.4.3 A Remark on Tree Diagrams

As noted in (11.4.2), we have an isomorphism

$$F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0}) = F(\frac{1}{2}, \mathbb{I})$$

whence we can identify the three groups F, $F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$, and $F(\frac{1}{2}, \mathbb{I})$. The notion of forest diagram carries over to $F(\frac{1}{2}, \mathbb{I})$ as this group can be viewed as a the subgroup of $F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ whose elements have support in $[0, 1]$, i.e, these homeomorphism are the identity outside $\mathbb{I}$. In fact, since $\mathbb{I}$ corresponds to a root in $\mathcal{F}^{(2)}$, the forest diagrams for elements in $F(\frac{1}{2}, \mathbb{I})$ are somewhat degenerate in that the top an bottom forest are almost trivial – only the first subtree can be non-trivial. Hence, when using $F(\frac{1}{2}, \mathbb{I})$ as a model for F, one often uses tree diagrams instead of forest diagrams. There is, however, not much of a difference.

11.5 Subgroups and Quotients

We already saw (11.4.25) that F has trivial center. In this section we will prove the following:

Theorem 11.5.1. *The commutator subgroup of group F is simple, and all proper quotients of F are abelian.*

These two claims are strongly related as shown by the following:

Lemma 11.5.2. *If a group G has trivial center and its commutator subgroup $[G, G]$ is simple, then every proper quotient of G is abelian.*

Proof. Let $N \trianglelefteq G$ be a non-trivial normal subgroup, and let $n \in N - \{1\}$ be a non-trivial element. As G has trivial center, there is an element $g \in G$ that does not commute with n. Hence $[n, g] \in [G, G] \cap N$ proves that $N \cap [G, G]$ is non-trivial. This is obviously a normal subgroup of G and hence a normal subgroup of $[G, G]$. As the latter group is simple, we have $N \cap [G, G] = [G, G]$ whence $[G, G] \leq N$. Hence G/N is abelian. **q.e.d.**

Thus, we only have to show that the commutator subgroup $[F, F]$ is simple. For this reason, we shall examine the commutator subgroup in various models. We start with the presentation.

Proposition 11.5.3. *An element $f = x_0^{a_0} \cdots x_r^{a_r} x_r^{-b_r} \cdots x_0^{-b_0}$ is in the commutator subgroup if and only if*

$$a_0 - b_0 = 0 = a_1 + \cdots + a_r - b_r - \cdots - b_1.$$

Proof. The canonical projection $F \to F^{\mathrm{ab}} = \mathbb{Z} \times \mathbb{Z}$ sends x_0 to $(1,0)$ and x_i to $(0,1)$ for $i \geq 1$. Hence the claim follows from the fact that

$$a_0 - b_0 = 0 = a_0 + \cdots + a_r - b_r - \cdots - b_0$$

is equivalent to

$$a_0 - b_0 = 0 = a_1 + \cdots + a_r - b_r - \cdots - b_1$$

which is easily seen. **q.e.d.**

This characterization translates into the other models for F straightforwardly.

Corollary 11.5.4. *A homeomorphism $\varphi \in F(\frac{1}{2}, \mathbb{R}^{\infty}_{\geq 0})$ is in the commutator subgroup if and only if it is the identity close to 0 and close to ∞.*

Proof. Write $\varphi = \xi_0^{a_0} \cdots \xi_r^{a_r} \xi_r^{-b_r} \cdots \xi_0^{-b_0}$. The slope at 0 translates into $a_0 - b_0$ whereas $a_0 + \cdots + a_r - b_r - \cdots - b_0$ is the shift realized by φ for large arguments. **q.e.d.**

The same routine translation from one model to the next gives:

Corollary 11.5.5. *A homeomorphism $\varphi \in F(\frac{1}{2}, \mathbb{I})$ is in the commutator subgroup if and only if it is the identity close to 0 and close to 1.* **q.e.d.**

Lemma 11.5.6. *Let $f = x_0^{a_0} \cdots x_r^{a_r} x_r^{-b_r} \cdots x_0^{-b_0}$ be an element of F with $s := a_0 + \cdots + a_r - b_r - \cdots - b_0 > 0$. Then the normal closure of f contains the commutator subgroup $[F, F]$.*

Proof. Let N be the normal closure of f. Then, for sufficiently large M, we have

$$\begin{aligned} f^{-1} x_M f &= \left(x_0^{a_0} \cdots x_r^{a_r} x_r^{-b_r} \cdots x_0^{-b_0} \right)^{-1} x_M x_0^{a_0} \cdots x_r^{a_r} x_r^{-b_r} \cdots x_0^{-b_0} \\ &= x_{M + a_0 + \cdots + a_r - b_r - \cdots - b_0} \\ &= x_{M+s}. \end{aligned}$$

Hence we have

$$x_M \equiv x_{M+s} \mod N.$$

Conjugating this congruence by a high power of x_0, we infer:

$$x_1 \equiv x_{1+s} \mod N$$

Conjugating by a high power of x_1, we finally obtain:

$$x_1 \equiv x_2 = x_0^{-1} x_1 x_0 \mod N.$$

Hence $x_0 x_1 \equiv x_1 x_0 \mod N$ which implies that the factor F/N is abelian since x_0 and x_1 generate F. **q.e.d.**

Proof of Theorem (11.5.1). As we have already seen (11.5.2), we only have to prove that the commutator subgroup $[F, F]$ is simple. So let $N \trianglelefteq [F, F]$ be a non-trivial normal subgroup and let n be a non-trivial element in N.

Working in the unit interval model for $F = F(\frac{1}{2}, \mathbb{I})$, we draw n:

$n =$ 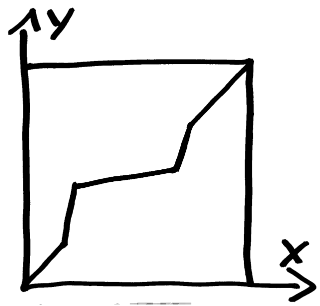

.

Let $J \subset \mathbb{I}$ be the support of n. Then the picture becomes this:

$n =$ 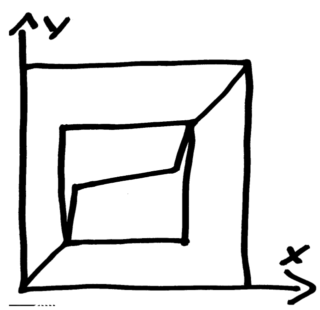

.

Note that $F(\frac{1}{2}, J)$ is an isomorphic copy of $F(\frac{1}{2}, \mathbb{I})$ inside $\left[F(\frac{1}{2}, \mathbb{I}), F(\frac{1}{2}, \mathbb{I})\right]$. Moreover, n satisfies the hypothesis of (11.5.6) inside $F(\frac{1}{2}, J)$. Hence

$$\left[F(\frac{1}{2}, J), F(\frac{1}{2}, J)\right] \leq N \cap F(\frac{1}{2}, J).$$

On the other hand, any element of $\left[F(\frac{1}{2}, \mathbb{I}), F(\frac{1}{2}, \mathbb{I})\right]$ is conjugate inside $\left[F(\frac{1}{2}, \mathbb{I}), F(\frac{1}{2}, \mathbb{I})\right]$ to an element of $\left[F(\frac{1}{2}, J), F(\frac{1}{2}, J)\right]$. Hence the normal closure of $\left[F(\frac{1}{2}, J), F(\frac{1}{2}, J)\right]$ in $\left[F(\frac{1}{2}, \mathbb{I}), F(\frac{1}{2}, \mathbb{I})\right]$ is the whole commutator subgroup $\left[F(\frac{1}{2}, \mathbb{I}), F(\frac{1}{2}, \mathbb{I})\right]$. Hence $\left[F(\frac{1}{2}, \mathbb{I}), F(\frac{1}{2}, \mathbb{I})\right] = N$. **q.e.d.**

Corollary 11.5.7. *Thompson's group F is not residually finite.*

Proof. Since F is infinite, any finite quotient is proper whence abelian. Therefore, every co-finite subgroup contains the commutator subgroup which is non-trival. **q.e.d.**

Exercise 11.5.8. Show that F is torsion-free.

Theorem 11.5.9. *Every non-abelian subgroup of F contains a copy of $C_\infty \times C_\infty$.*

Proof. Let f and g be two elements that do not commute. We have to find a copy of $C_\infty \times C_\infty$ in the subgroup $\langle f, g \rangle$.

Since f has only finitely many break points, the closed set

$$\{t \mid f(t) = t\}$$

decomposes into finitely many intervals. As the same holds for g, it follows that the open set

$$\{t \mid f(t) \neq t \text{ or } g(t) \neq t\}$$

decomposes into finitely many open intervals $I_0, \ldots, I_r$.

We consider the commutator $f_0 := [f, g]$. We will find a conjugate of f_0 inside $\langle f, g \rangle$ that commutes with f_0. The key idea is to conjugate such that the support of f_0 in I_0 is moved off itself, which is possible since this homemomorphism f_0 is the identity in small neighborhoods of the endpoints of I_0. So we only have to prove the existence of this conjugating element.

Let t be any point in I_i. Obviously, $\inf \langle f, g \rangle \cdot t$ is a global fix point for the action of $\langle f, g \rangle$ in the closure of I_i. Clearly, it has to be $\inf I_i$. Hence this points is a limit point of the orbit of t. Hence the problem in the preceding paragraph has a solution and we can conjugate f_0 as to move its support off itself inside I_0. Call this conjugate g_0.

Now, consider $f_1 := [f_0, g_0]$. By construction, this element is the identity on I_0. So we run the same argument as above in the interval I_1. This way, we can take care of all intervals. Eventually, we have two elements f_r and g_r that commute. **q.e.d.**

Corollary 11.5.10. *F does not contain non-abelian free groups.* **q.e.d.**

11.6 Amenability

It is not known whether F is amenable. However, if F is amenable, it is at least not obviously so.

Theorem 11.6.1. *Thompson's group F is not elementary amenable.*

Proof. We prove $F \notin EG_\alpha$ for all α. And, of course, this is done by transfinite induction or contradiction. We proceed by contradiction. So assume there was an ordinal α with $F \in EG_\alpha$. Then we can take α to be minimal with this property.

Since F is neither abelian nor finite, $\alpha > 0$ and so F sneaked into EG_α as an extension or as a direct union of groups in lower strata. The direct union possibility is ruled out since F is finitely generated. Thus, we have to deal with ways of representing F as an extension.

Let $N_\beta \hookrightarrow F \twoheadrightarrow Q_0$ be a short exact sequence where Q_0 is a proper quotient. By (11.5.1), Q_0 is abelian. Hence $N_\beta \in EG_\beta$ contains $[F, F]$. Hence $[F, F] \in EG_\beta$. However, since $[F, F]$ contains a copy of F, it follows that $F \in EG_\beta$. Thus, α was not minimal. **q.e.d.**

11.7 Finiteness Properties

We will construct a very nice (i.e., contractible) cube complex upon which F acts. The construction is due to K.S. Brown and R. Geoghegan [BrGe84].

This construction starts from the Cayley complex for the infinite presentation

$$\tilde{F} := \langle x_0, x_1, \ldots \mid x_m x_i = x_i x_{m+1} \text{ for } i < m \rangle .$$

Thus, we have a vertex for each element in F. From each vertex countably many edges issue, and the relations give rise to squares. This complex $\Gamma_{\tilde{F}}$ is simply connected. Our goal is to fill in cubes as to kill off higher homotopy groups. To do this, we fill in all the obvious cubes: For any square whose edges are labeled by j, k and $k+1$ with $j < k$, and any $i < j$, we have a cube

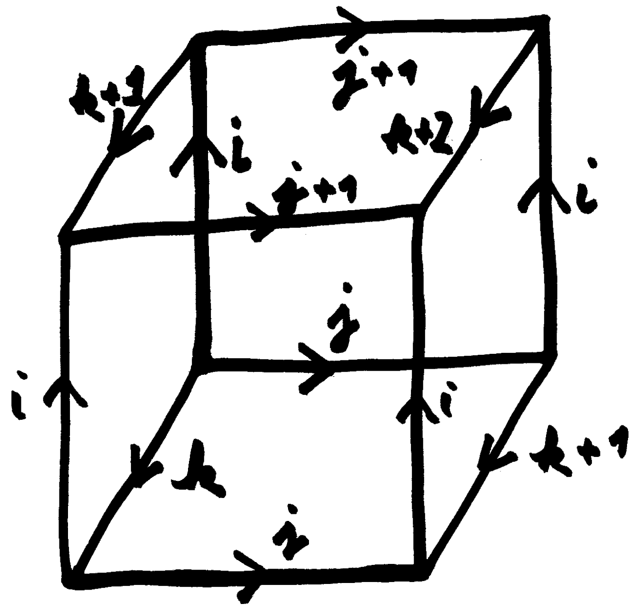

We continue in higher dimensions. Thus for each vertex in $\Gamma_{\tilde{F}}$ and each tupel $(q_0 < q_2 < \cdots < q_r)$, we have a cube issuing from that vertex whose edges are labeled by the indices q_i. Note that each of these cubes has a unique vertex with all edges pointing away (the source) and a unique vertex with all edges coming in (the sink). Let Y denote the cube complex constructed this way, and put $X := F\backslash Y$.

Theorem 11.7.1 ([BrGe84]). *Y is contractible.*

Proof. Let P be the monoid of positive elements in F and let Y_+ be the subcomplex of Y spanned by $P \subset Y$. This is the complex formed by all those cubes whose sink is in P.

Claim A. Y_+ is contractible.

PROOF. We will do Morse theory. There is an obvious height function on the vertices: the length of the positive word – or alternatively, the number of non-leaf nodes in the forest representing that vertex. Moving along an outgoing edge increases the height by 1. Thus, the height extends linearly to the cubes. This is a combinatorial Morse function as it is non-constant on edges. Since the minimal set is just one point (the identity vertex), contractibility of Y_+ would follow from contractibility of descending links (C.2.11).

The descending link is the part of the link in a vertex v spanned by its incoming edges. These correspond to splitting off a generator on the right, i.e., deleting a terminal caret in the forest. Thus we can tell which incoming edges issue from vertices in Y_+. More importantly, we see that we can delete all terminal carets and find a source in Y_+ which gives rise to a simplex in the link of v that connects all the descending vertices. Thus, we proved that the descending link is a simplex and therefore contractible. The claim follows. □

Now we consider the cover of Y by translates fY_+.

Claim B. For any finite set of elements $f_0, \ldots, f_m \in F$, there is an element $f \in F$ such that

$$f_0 Y_+ \cap \cdots \cap f_m Y_+ = f Y_+.$$

PROOF. This follows from the corresponding claim about translates fP which in turn follows by induction from the claim that for any $f \in F$, there is a $g \in P$ such that

$$P \cap fP = gP.$$

Writing f as a pair of trees, we see that in order to find a positive right multiple of f, we first have to annihilate the bottom tree. Afterward, we will end up with a right multiple of the top tree. Thus, g is the element represented by the top tree of f. □

From this claim, it follows at once, that the translates fY_+ form a cover of Y by contractible subcomplexes such that any intersection of these subcomplexes is contractible. The nerve of the cover is the simplex with vertex set F. Hence, by (C.2.17), the space Y has the homotopy type of a big simplex. **q.e.d.**

Theorem 11.7.2. *Y is CAT(0).*

Proof. The edges around each vertex v come in pairs: For each generator x_i, there is one incoming edge and one outgoing edge, both labeled by i. We denote the incoming edge by the pair $(i,1)$ and the outgoing edge by $(i,1)$. These pairs represent the vertices in $\mathrm{Lk}(v)$. In order to determine the simplices, we have to know which of these edges span a cube. Obviously, there are no cubes that involve both, $(i,0)$ and $(i,1)$. Hence we can represent cubes in the star of v by tupels

$$((q_1,\varepsilon_1),\ldots,(q_1,\varepsilon_1))$$

with $q_i < q_{i+1}$.

The following picture assumes $i < j$ and shows the edge labels of all possible squares involving these to labels.

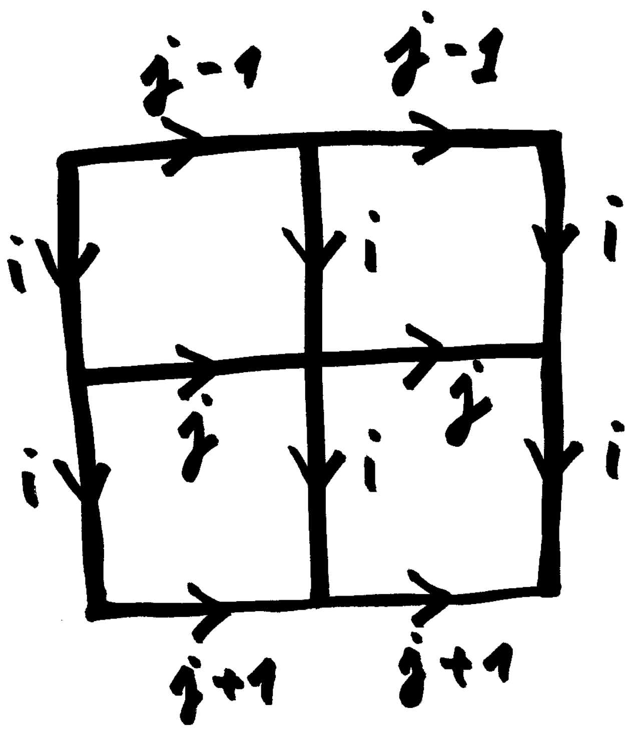

The incoming edges only can form a square if $i < j-1$. In general, a tupel

$$((q_1,\varepsilon_1),\ldots,(q_1,\varepsilon_1))$$

of edges spans a cube if $q_i < q_{i+1} - \varepsilon_i$.

This condition, however, is satisfied for a tupel if and only if it is satisfied for all of its subtupels of size two. This is, a set of vertices in $\mathrm{Lk}(v)$ spans a simplex if and only if any two of them are joined by an edge. Thus, $\mathrm{Lk}(v)$ is a flag complex. It follows from Gromov's lemma that Y is CAT(0) provided it is simply connected. That, in turn follows from the fact that the 2-skeleton of Y is a Cayley complex for F. **q.e.d.**

We give a description of the quotient X. This is also a cube complex. X has one vertex. The edges in the 1-skeleton are indexed by non-negative integers. We index higher dimensional cubes by the indices of the edges pointing toward the sink of the cube. Then, an ascending sequence

$$\mathbf{q} := (q_1, \ldots, q_m)$$

describes a cube if and only if $0 \leq q_1$ and $q_i + 2 \leq q_{i+1}$. The $2m$ faces of the cube $\mathbf{q}$ are given by the boundary operators

$$\begin{aligned} \partial_i^+ \mathbf{q} &:= (q_1, \ldots, q_{i-1}, q_{i+1}, \ldots q_m) \\ \partial_i^- \mathbf{q} &:= (q_1, \ldots, q_{i-1}, q_{i+1} - 1, \ldots q_m - 1)\,. \end{aligned}$$

Theorem 11.7.3. *X is homotopy equivalent to a CW-complex that has one vertex and two cells in each dimension ≥ 1.*

Proof. We distinguish three types of cells. For each dimension ≥ 1, we have two *essential* cells $(0, 3, \ldots, 3m-3)$ and $(1, 4, \ldots, 3m-2)$. We call an m-cube $\mathbf{q} = (q_1, \ldots, q_m)$ *collapsible* if there is an index i such that $q_{i+1} \neq q_i + 3$ and such that, for the last index j with that property, $q_{j+1} = q_j + 2$. In this case

$$\partial_j^+ \mathbf{q} = (q_1, \ldots, q_{j-1}, q_{j+1}, \ldots q_m)$$

is called the *free face* of the collapsible cube $\mathbf{q}$. Cells that are neither essential nor collapsible are considered *redundant*.

Observe that the free face $\partial_j^+ \mathbf{q}$ of a collapsible cube $\mathbf{q}$ is redundant since $q_{j+1} \geq q_{j-1} + 4$ and $q_{i+1} = q_i + 3$ for all $i > j + 1$.

On the other hand, each redundant cell $\mathbf{p} = (p_1, \ldots, p_m)$ is the free face of a unique collapsible cell since for the last index i with $p_{i+1} \neq p_i + 3$, we have $p_{i+1} \geq p_i + 4$ whence

$$(p_1, \ldots, p_i, p_{i+1} - 2, p_{i+1}, \ldots, p_m)$$

is a collapsible cell with free face $\mathbf{p}$ – note that we have to allow $i = 0$ here because of cubes like $(2, 5, 8, \ldots)$.

Now consider the collapsible cube $\mathbf{q} = (q_1, \ldots, \ldots q_m)$ with free face $\partial_j^+ \mathbf{q}$ which is to say that we have

$$q_i = q_j + 3\,(i - j) - 1$$

for all $i > j$. The following statements are easy to see:

1. For $i < j$, the face $\partial_i^- \mathbf{q}$ is collapsible.

2. For $i > j$, the tupel $\partial_i^- \mathbf{q}$ precedes the free face $\partial_j^+ \mathbf{q}$ in the lexicographic order.

3. Also, $\partial_j^- \mathbf{q}$ precedes the free face in the lexicographic order.

We define $X_{-0} := X_{+0} := X_0$, and inductively

$$\begin{aligned} X_{-m} &:= X_{+m-1} \cup \text{ essential cells in dimension } m \\ X_{+m} &:= X_{-m} \cup \text{ redundant cells in dimension } m \cup \text{ collapsible cells in dimension } m+1. \end{aligned}$$

Obviously, we have

$$X_{+m-1} \subseteq X_{-m} \subseteq X_m \subseteq X_{+m}.$$

Moreover, we have $X_{-m} \simeq X_{+m}$ since the latter is obtained from the former by a transfinite sequence of elementary expansions which are performed according to the lexicographic order. Since the lexicographic order is a well ordering, this actually makes sense.

Now, we are ready to construct the space Z which is homotopy equivalent to X but uses at most two cells per dimension. Put $Z_0 := X_0$. Higher skeleta are constructed by induction together with the homotopy equivalence. So assume we have already a homotopy equivalence

$$\pi_{m-1} : X_{+m-1} \to Z_{m-1}.$$

Put

$$Z_m := Z \cup B_0^m \cup B_1^m$$

where the attaching maps are induced by the maps that embed the boundaries of the two essential m-cubes in X_{+m-1}. Thus we obtain a homotopy equivalence

$$\pi_m : X_{-m} \to Z_m$$

and $X_{-m} \simeq X_{+m}$ gives us the equivalence

$$\pi_m : X_{+m} \to Z_m$$

which completes the proof. **q.e.d.**

11.8 The Conjugacy Problem

Theorem 11.8.1. *The conjugacy problem in F is solvable.*

Proof. ... **q.e.d.**

Part V

Appendices

Appendix A

Combinatorial Group Theory: Nuts and Bolts

A.1 Generators and Relations

A.1.1 Generating Sets / Cayley Graphs

Definition A.1.1. Let S be a set of elements in the group G. The intersection of all subgroups $H \leq G$ containing S is a subgroup of $\langle S \rangle \leq G$. It is called the subgroup generated by S. The subset S is called a generating set for $\langle S \rangle$.

Definition A.1.2. A graph Γ is a map $\tau : \overrightarrow{\mathcal{E}}_\Gamma \to \mathcal{V}_\Gamma$ where $\mathcal{V}_\Gamma$ is a set (its elements are called vertices) and $\overrightarrow{\mathcal{E}}_\Gamma$ is a free $\mathbb{Z}^2$-set, i.e., a set together with a fixpoint free involution $\mathrm{op} : \overrightarrow{\mathcal{E}}_\Gamma \to \overrightarrow{\mathcal{E}}_\Gamma$. The elements of $\overrightarrow{\mathcal{E}}$ are (oriented) edges. The elements of $\mathcal{E}_\Gamma := \overrightarrow{\mathcal{E}}_\Gamma/\mathrm{op}$ are called (geometric) edges.

An orientation on Γ is a section $\mathrm{o} : \mathcal{E} \to \overrightarrow{\mathcal{E}}$.

We define another map $\iota : \overrightarrow{\mathcal{E}}_\Gamma \to \mathcal{V}_\Gamma$ by $\iota(\vec{e}) := \tau(\mathrm{op}(\vec{e}))$. The map τ assigns to each edge its terminal vertex whereas ι provides the initial vertex of the edge.

Definition A.1.3. A G-labeling of a graph Γ is a map $\phi : \overrightarrow{\mathcal{E}}_\Gamma \to G$ satisfying

$$\phi(\mathrm{op}(\vec{e})) = \phi(\vec{e})^{-1}.$$

Definition A.1.4. Let G be a group with finite generating system Σ. The (left) Cayley graph $\Gamma := \Gamma_\Sigma^G$ is constructed as follows: The set $\mathcal{V}_\Gamma$ is G. The set of oriented edges is $\overrightarrow{\mathcal{E}}_\Gamma := G \times \Sigma \times \{\pm 1\}$. We have to specify a fixpoint free involution op :

$\overrightarrow{\mathcal{E}}_\Gamma \to \overrightarrow{\mathcal{E}}_\Gamma$ and an endpoint map τ. For any oriented edge $\vec{e} = (g, \sigma, \varepsilon)$

$$\begin{aligned} \operatorname{op}(\vec{e}) &:= \big(g\sigma, \sigma^{-1}, -\varepsilon\big) \\ \tau(\vec{e}) &:= g \end{aligned}$$

The set of geometric edges is obviously isomorphic to $G \times \Sigma$ and the map

$$(g, \sigma) \mapsto (g, \sigma, 1)$$

defines an orientation on Γ. We regard this as the standard or positive orientation. Note that G acts from the left on Γ^G_Σ, and this action preserves the orientation.

There is a corresponding notion of a right Cayley graph upon which G acts from the right.

For an oriented edge $\vec{e} = (g, \sigma, \varepsilon)$, we call σ is the label of $\vec{e}$. We denote the label of $\vec{e}$ by $\sigma_{\vec{e}}$. When we talk about labels of geometric edges, we either (silently) identify them with the positively oriented edges or there is an implied orientation for the edge (e.g., when the edge is part of a directed path).

Any directed path in Γ reads a word in $\Sigma \uplus \Sigma^{-1}$: while you are moving along the path, you pick up the labels of the oriented edges you are traveling. Note that all G-translates of a given path read the same word. On the other hand, any word w over $\Sigma \uplus \Sigma^{-1}$ defines a G-orbit of paths: For any vertex in $v \in \Gamma$ there is a unique directed path starting at v that reads w.

A.1.2 Defining Relations / Cayley Complexes

Definition A.1.5. Let $G = \langle \Sigma \rangle$ be a group and $\Gamma = \Gamma^G_\Sigma$. A Relation in G over Σ is a word w over $\Sigma \uplus \Sigma^{-1}$ that evaluates to $1 \in G$ by multiplication of its letters.

Note that relations in G correspond to loops in Γ. Therefore, a relation r over Σ determines a G-orbit of loops in Γ. We can G-equivariantly glue in a family $\mathcal{D}_r$ of 2-cells killing theses loops: The 2-cells are polygons whose boundaries are reading r.

A set R of relations defines G if glueing in $\bigcup_{r\in R} \mathcal{D}_r$ kills the fundamental group of Γ. In this case, the resulting space

$$\Gamma_{\Sigma,R}(G) := \Gamma \cup \bigcup_{r\in R} \mathcal{D}_r$$

is called the Cayley complex of the presentation $\mathcal{P} = \langle \Sigma \mid R \rangle$ which is said to present the group G.

Note that G acts freely (from the left) on $\Gamma_{\Sigma,R}$ and the quotient

$$K_{\mathcal{P}} := G \backslash \Gamma_{\Sigma,R}(G)$$

is called the presentation complex or standard two-complex of the presentation $\mathcal{P}$. Moreover, $\Gamma_{\Sigma,R}(G)$ is the universal cover of $K_{\mathcal{P}}$ and G acts as the group of deck transformations. In particular, every group is the fundamental group of a 2-complex since

$$G = \pi_1(K_{\mathcal{P}}) .$$

Two presentations are equivalent if they present the same group (or more precisely, if they present isomorphic groups).

Fact A.1.6. *Let Σ generate G. A set R of words over $\Sigma \uplus \Sigma^{-1}$ defines G if and only if it generates the kernel of the canonical projection*

$$F_{\Sigma \uplus \Sigma^{-1}} \to G$$

as a normal subgroup of $F_{\Sigma \uplus \Sigma^{-1}}$.

Fact A.1.7. *Let $\langle \Sigma \mid R \rangle$ be a presentation for G. Then, a word w is a relation in G over Σ if and only if, when interpreted as an element of the free group $F_{\Sigma \uplus \Sigma^{-1}}$, it represents an element of the normal subgroup generated by R.*

Fact A.1.8. *The "calculus of presentation" is given by the following rules.*

1. *Let $\langle \Sigma \mid R \rangle$ present G, and let R' be a set of relations in G over Σ. Then*

   $$\langle \Sigma \mid R \cup R' \rangle$$

   also presents G. This process is called adding redundant relations. There is an obvious inverse process of deleting redundant relations.

2. *Let $\langle \Sigma \mid R \rangle$ present G, and let w be a word over $\Sigma \uplus \Sigma^{-1}$ that represents the element $g \in G$, then*

   $$\left\langle \Sigma \cup \{g\} \;\middle|\; R \cup \left\{wg^{-1}\right\} \right\rangle$$

   also presents G. This process is called adding a generator and a defining equation. You can also add many generators with their defining equations at the same time – even infinitely many.

   The inverse passage from a presentation of the form

   $$\left\langle \Sigma \cup \{g\} \;\middle|\; R \cup \left\{wg^{-1}\right\} \right\rangle$$

   to $\langle \Sigma \mid R \rangle$ is the deletion of a redundant generator. Again, you can delete many generators at the same time.

3. *Two presentations are equivalent if and only if you can pass from one to the other by a finite chain of the following processes: adding redundant relations, deleting redundant relations, adding generators and defining equations, deleting redundant generators.*

Remark A.1.9. A relation of the form wg^{-1} is often written (and always read) as $w = g$ for this is what it says.

Let us discuss the question of how to come up with presentations. In addition to ad hoc methods (which are at times unavoidable and can be very powerful too), there are two geometric principles.

Theorem A.1.10. *Let G act by homeomorphisms on the 1-connected space X, and let U be an open 0-connected subset in X such that $X = GU$. Put*

$$\Sigma := \{g \in G \mid U \cap gU \neq \emptyset\}$$

and

$$R := \{xy = (xy) \mid x, y \in \Sigma \text{ and } U \cap xU \cap xyU \neq \emptyset\}.$$

Then $\mathcal{P} := \langle \Sigma \mid R \rangle$ is a presentation for G.

Proof. Let $\tilde{G}$ be the group presented by $\mathcal{P}$. Define $\tilde{X} = \biguplus_{g \in \tilde{G}} gU/\sim$ where

$$g_0u_0 \sim g_1u_1 :\Longleftrightarrow g_1 = g_0x,\ u_0 = xu_1 \in U \cap xU \text{ for some } x \in \Sigma.$$

Let us call x a witness for the equivalence.

It is not hard to show that $\sim$ is an open equivalence relation. For transitivity, you have to use the defining relations for $\tilde{G}$: If

$$g_0u_0 \sim g_1u_1$$

with witness $x \in \Sigma$ and

$$g_1u_1 \sim g_2u_2$$

with witness $y \in \Sigma$, then

$$u_0 = xu_1 = xyu_2 \in U \cap xU \cap xyU \neq \emptyset$$

whence $(xy) \in \Sigma$ and R contains the relation $xy = (xy)$. Then

$$g_0u_0 \sim g_2u_2$$

with witness (xy).

Since U is path-connected, so is $\tilde{X}$. This is clear from the picture:

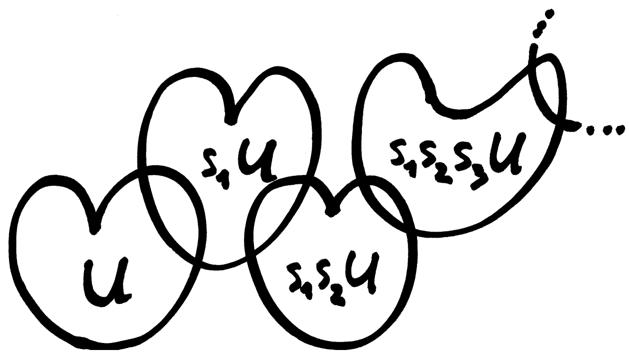

There is an obvious group homomorphism $\tilde{G} \to G$ which induces a map $\tilde{X} \to X$. This map is a covering map, and the group of deck-transformations is $\ker\left(\tilde{G} \to G\right)$. Since X is 1-connected and $\tilde{X}$ is a connected cover, the projection is an isomorphism, the group of deck-transformations is trivial, and $\tilde{G} = G$. **q.e.d.**

The second geometric source of presentations are 2-complexes. Their fundamental groups have presentations that can be read off the complex easily. This methods goes way back to Poincare, was described in a more precise fashion by Tietze, and works because of the Seifert-Van-Kampen theorem.

Fact A.1.11. *Given a 2-complex X, a presentation for $\pi_1(X)$ can be read off as follows:*

1. *Choose a spanning tree T for the 1-skeleton.*
2. *Introduce a formal generator for each edge in X.*
3. *For each 2-cell read off a relation along its boundary.*
4. *Finally declare trivial each generator that comes from an edge in T.*

Let $H \leq G = \langle \Sigma \mid R \rangle$. Then H is the fundamental group of a cover of the presentation 2-complex associated to the presentation $\mathcal{P} = \langle \Sigma \mid R \rangle$. In this case (A.1.11) is know as the Reidemeister-Schreier method for finding presentations of subgroups of groups defined by generators and relations.

Fact A.1.12. *Let $H \leq G = \langle \Sigma \mid R \rangle$, and fix a set T of words $w_1, \ldots, w_r$ in the free group F_Σ representing the right cosets of H in G. (Such a set is called a Schreier transversal.) For each $g \in G$, let $\overline{g}$ be the word w_i representing the coset Hg.*

For each $w \in T$ and $x \in \Sigma$, define

$$y_{w,x} := wx\overline{(wx)}^{-1}$$

Then H is generated by the $y_{w,x}$ with defining relations wrw^{-1} where $w \in T$ and $r \in R$.

This needs an illustration.

Example A.1.13. Let H be the kernel of the homomorphism

$$\left\langle x_1, \ldots, x_r \,\middle|\, {x_1}^2 \cdots {x_r}^2 = 1 \right\rangle \to C_2$$

sending each x_i to the non-trivial element.

There are two cosets, represented by 1 and x_r. We have the following generators for H:

$$\begin{aligned} y_i &:= 1 x_i x_r^{-1} && \text{for } i \leq r-1 \\ z_i &:= x_r x_i && \text{for } i \leq r. \end{aligned}$$

Now, we write the two relations first in the generators x_i:

$$\begin{aligned} {x_1}^2 \cdots {x_r}^2 &= 1 \\ x_r {x_1}^2 \cdots {x_r}^2 x_r^{-1} &= 1. \end{aligned}$$

Now we rewrite these:

$$\begin{aligned} x_1 x_r^{-1} x_r x_1 x_2 x_r^{-1} x_r x_2 \cdots x_r x_r^{-1} x_r x_r &= 1 \\ x_r x_1 x_1 x_r^{-1} x_r x_2 x_2 x_r^{-1} \cdots x_r x_r x_r x_r^{-1} &= 1. \end{aligned}$$

Finally, we rewrite these in the generators for H:

$$\begin{aligned} y_1 z_1 \cdots y_{r-1} z_{r-1} z_r &= 1 \\ z_1 y_1 \cdots z_{r-1} y_{r-1} z_r &= 1 \end{aligned}$$

An easy Tietze transformation eliminates z_r, and we have:

$$H = \langle y_1, \ldots, y_{r-1}, z_1, \ldots, z_{r-1} \mid y_1 z_1 \cdots y_{r-1} z_{r-1} = z_1 y_1 \cdots z_{r-1} y_{r-1} \rangle .$$

Exercise A.1.14. Prove (A.1.12).

Definition A.1.15. A group is finitely presented if it has at least one finite presentation, i.e., a presentation that employs only finitely many generators and finitely many relations.

Observation A.1.16. *It follows immediately from (A.1.12) that a subgroup of finite index in a finitely presented group is finitely presented.*

Exercise A.1.17. Prove that a virtually finitely presented group is finitely presented.

Exercise A.1.18. Let G be finitely generated. Show that every generating set for G contains a finite subset that already generates G.

Exercise A.1.19. Let G be finitely presented. Show that any finite generating set for G set can be extended to a finite presentation for G.

Exercise A.1.20. Let

$$G = \langle x_1, \ldots, x_r \mid r_1, r_2, \ldots \rangle$$

be a finitely presented group. Show that in the above presentation all but finitely many relations are redundant.

Exercise A.1.21. Prove that finitely presented by finitely presented groups are finitely presented.

Exercise A.1.22. The fundamental group of English is generated by the 26 letters $\mathsf{a}, \ldots, \mathsf{z}$, subject to all relations $w = u$ where w and u are letter sequences that are pronounced identically in at least some contexts. E.g., gh is trivial since it is silent in some words, and $\mathsf{ir} = \mathsf{er}$ because of the words "bird" and "her".

Prove that the fundamental group of English is trivial.

In passing from the Cayley graph Γ_Σ^G to the Cayley complex $\Gamma_{\Sigma,R}(G)$ we killed the fundamental group by gluing in 2-cells. This might introduce non-trivial π_2. Of course, we could go on and kill this π_2 by gluing in 3-cells in a G-equivariant way, at the cost of, maybe, introducing non-trivial π_3. We can continue and kill all homotopy groups. We get a contractible free G-complex. Given this construction, what will the presentation 2-complex turn into? The answer is given in the following definition.

Definition A.1.23. An Eilenberg-Maclane-complex for G is a CW-complex with fundamental group G and contractible universal cover.

See section 2.4.1, in particular (2.4.7) and (2.4.8), for more on Eilenberg-Maclane spaces and their relation to group cohomology.

A.1.3 Van Kampen Diagrams

Let

$$\mathcal{P} := \langle x_1, \ldots, x_r \mid r_1, \ldots, r_s \rangle$$

be a finite presentation for the group G. The associated presentation 2-complex $K_\mathcal{P}$ is the two dimensional CW-complex constructed as follows: There is precisely one 0-cell (the vertex). For each generator x_i, we attach a loop, labeled by its generator x_i. So far, we have constructed a graph, and every word w over the alphabet $\{x_1, x_1^{-1}, \ldots, x_r, x_r^{-1}\}$ defines a loop based at the vertex. We attach a 2-cell along this path.

Fact A.1.24. *The fundamental group of $K_{\mathcal{P}}$ is presented by $\mathcal{P}$.*

A Van Kampen disc diagram is a pictorial proof that some word is trivial. Let us discuss examples before providing a more rigorous definition.

Example A.1.25. We consider $\langle x, y, z \mid xy = yz = zx \rangle$. This is a picture proof that $xyx = yxy$:

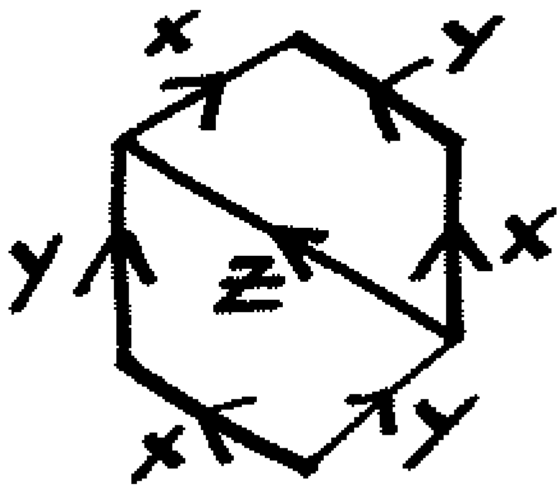

Example A.1.26. Here, we prove that $zyzx = yxzy^{-1}zy$ follows from $yx = zy$:

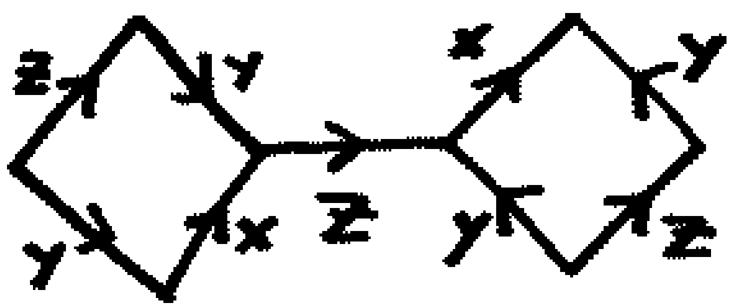

Definition A.1.27. A Van Kampen disc diagram for the word w is a planar graph Γ, embedded in the 2-sphere all of whose complementary regions are simply connected together with

1. a distinguished complementary region (represented in the drawings as the region containing ∞) whose boundary (with the orientation induced from the 2-sphere reads w,

2. a labeling of its oriented edges by the generators, i.e., a map

$$\lambda : \overrightarrow{\mathcal{E}}\Gamma \to \{x_1, x_1^{-1}, \ldots, x_r, x_r^{-1}\}$$

satisfying

$$\lambda(\vec{e}) = \lambda(\mathrm{op}\vec{e})^{-1},$$

and

3. a labeling of its non-distinguished complementary regions by the relators r_j (or their inverses) such that the boundary of each region reads the corresponding relator word or its inverse, respectively.

It is clear from the examples we discussed above that a Van Kampen diagram for w proves $w =_G 1$.

Fact A.1.28. *Every word representing the trivial element has a Van Kampen disc diagram. This diagram can be chosen to be reduced: Think of the relations as (regular) polygonal discs whose boundaries read the r_j. A Van Kampen diagram provides a way of manufacturing a planar region (the complement of the distinguished region) by glueing together copies of these polygons. Some of these discs need to be flipped over before glueing them onto the graph; this happens if they are labeled by the inverse of a relation. A diagram is unreduced if you two polygons sharing an edge such that one if the flipped image of the other obtained by reflecting the polygon along the very edge that they share. Any unreduced diagram can be reduced.*

Example A.1.29. This picture illustrates schematically, how reduction works:

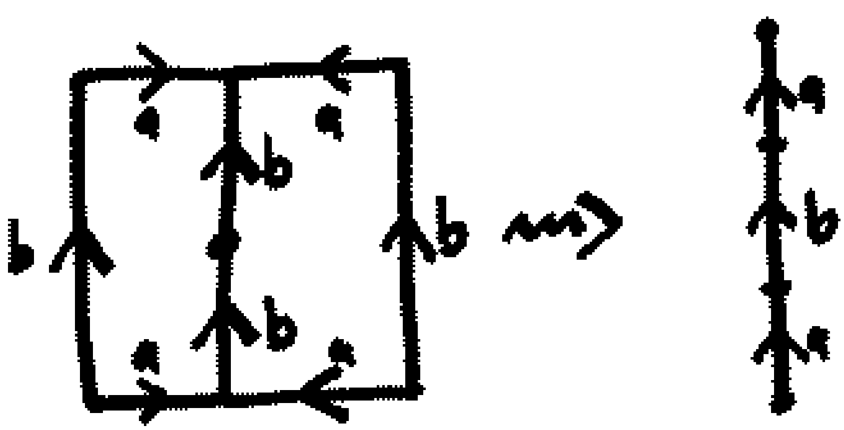

A.1.4 Isoperimetric Inequalities

Definition A.1.30. The area of a van Kampen disc diagram for the word w is its number of non-distinguished regions. The length of w is the perimeter of the diagram. An isoperimetric function for the presentation $\mathcal{P}$ is a function $\delta : \mathbb{N} \to \mathbb{N}$ such that every word of length $\leq n$ representing 1 in G has a van Kampen diagram of area $\leq \delta(n)$. A minimal isoperimetric function is called a Dehn function for the group G – note that this notion is only well defined up to weak equivalence (1.4.16) if we do not specify the generating set.

Observation A.1.31. *If a group has a computable isoperimetric function, then it has a solvable word problem. (In fact, it was shown in [SBR02] that every Dehn function for G is equivalent to the time function for a non-deterministic Turing machine solving the word problem in G.)*

Fact A.1.32. *Every group with a sub-linear isopermetric function has a linear isoperimetric function. The groups with linear isoperimetric functions are precisely the word hyperbolic groups.*

Fact A.1.33. *The weak equivalence class of an isoperimetric function for G can be computed from any simply connected cocompact G-CW-complex that is quasi-isometric to the universal cover of the presentation 2-complex.*

A.1.5 Small Cancellation Theory

Definition A.1.34. A vertex of Γ is essential if has degree $\neq 2$. The rational is that we can delete degree-2-vertices: A segment is an edge path whose terminal vertices are essential and whose inner vertices are not.

Observation A.1.35. *If all relations r_j and the word w are cyclic reduced, then each essential vertex has degree ≥ 3. We shall henceforth make this assumption. Moreover, note that if each vertex has degree ≥ 3, we can use an easy Euler characteristic count to prove:*

> *The average number of segments around each non-distinguished 2-cell is strictly less than 6.*

So let F be the number of non-distinguished regions, E be the number of segments, and V be the number of essential vertices. Then

$$1 = V - E + F.$$

On the other hand

$$\begin{aligned} 2E &\geq (6+3\varepsilon)\,F \\ 2E &\geq 3V \end{aligned}$$

Thus

$$(6+2\varepsilon)\,E \geq (6+3\varepsilon)\,(F+V) = (6+3\varepsilon)\,(E+1)$$

which is a contradiction. **q.e.d.**

Small cancellation theory is based upon the idea that one can write down combinatorial conditions to ensure that segments are "short". For the moment, we shall not precisely spell out those conditions, but we shall examine the consequences of the following:

Hypothesis A.1.36. *In each reduced Van Kampen diagram for w, the length of any segment is at most $\frac{1}{7}$ of the perimeter of its neighboring non-distinguished regions.*

So this is what we mean by "short". Note that an immediate consequence of (A.1.36) is that an interior region has at least seven segments in its boundary. A region is interior if it is does not share an edge with the distinguished region.

Observation A.1.37. *At most 95% of all regions are interior. Since they have at least seven segments in their boundary, too high a percentage of interior regions will drive the average number of segments in the boundary of regions above 6. We have already observed that this is impossible.*

Corollary A.1.38. *A small cancellation presentation satisfies a linear isoperimetric inequality. Thus, the word problem is solvable.*

Another application of the small cancellation hypothesis is the construction of 2-dimensional Eilenberg-Maclane spaces.

Definition A.1.39. A spherical Van Kampen diagram is a van Kampen disc diagram whose distinguished region also reads a relator word. Thus, we can drop the marker for the distinguished region and consider them all equal. Therefore, a spherical diagram is a planar graph embedded in the sphere all of whose complementary regions are simply connected and identified with the relators so that the boundaries of the regions read the corresponding relations. The notion of reduction and reduced diagrams carry over without difficulty.

Fact A.1.40. *If there is no reduced spherical diagram, then $\pi_2(K_{\mathcal{P}}) = 0$.*

Observation A.1.41. *Every small cancellation presentation is aspherical: Our Euler characteristic count showed that a sphere cannot be tiled in such a way that each tile has at least six edges, if each vertex is of degree ≥ 3.* **q.e.d.**

Remark A.1.42. The presentation 2-complex of any aspherical presentation is an Eilenberg-Maclance space for the group presented, i.e., the universal cover of the presentation 2-complex is contractible. Hence any finite aspherical presentation presents a torsion free group of type F.

Proof. The universal cover of an aspherical 2-complex is a simply connected 2-complex whose homotopy group in dimension 2 vanishes. By Hurewicz's Theorem, the first non-trivial homology and homotopy group of the cover occur in the same dimension. But beyond dimension 2 all homology vanishes since there are no cells to generate the chains. Thus all homology groups and all homotopy groups of the universal cover are trivial. **q.e.d.**

A.2 Constructions

Here is a brief overview on various way of building groups on top of other groups.

A.2.1 Direct, Semidirect, and Wreath Products

Definition A.2.1. Let $(G_i)_{i\in I}$ be a family of groups. The cross product of the G_i is the group

$$\bigtimes_{i\in I} G_i := \mathrm{Map}(I, G)$$

where multiplication is defined pointwise.

The direct product of the G_i is the subgroup of the product consisting of those families (g_i) where $g_i = 1$ for all but at most finitely many $i \in I$. We denote the direct product by

$$\prod_{i\in I} G_i.$$

Note that for finite index sets I, the notion of cross product and direct product coincides.

Definition A.2.2. A homomorphism $\pi : G \to Q$ is a retraction if there is a homomorphism $\sigma : Q \to G$ such that the composition $Q \xrightarrow{\sigma} G \xrightarrow{\pi} Q$ is the identity. Note that in this case π has to be surjective. The homomorphism σ is called a (group theoretic) section.

A short exact sequence of groups is a sequence of groups and homomorphisms

$$N \xrightarrow{\iota} G \xrightarrow{\pi} Q$$

such that N embeds into G as a normal subgroup with quotient Q – i.e., ι is injective, π is surjective, and the image of ι is the kernel of π. In this case G is called an extension of N by Q.

A short exact sequence splits if π is a retraction. In this case G is called a split extension.

Definition A.2.3. Let $N \xrightarrow{\iota} G \xrightarrow{\pi} Q$ be a split extension with section $Q \xrightarrow{\sigma} G$. Then we have a homomorphism

$$\begin{aligned} \varphi : Q &\rightarrow \mathrm{Aut}(N) \\ q &\mapsto \alpha_{\sigma^q} \end{aligned}$$

where α_g is conjugation by g:

$$\begin{aligned} \alpha_g : N &\rightarrow N \\ n &\mapsto gng^{-1} \end{aligned}$$

We call φ the homomorphism induced or realized by the split extension.

Given a pair of groups Q and N, it turns out, any homomorphism $Q \rightarrow \mathrm{Aut}(N)$ can be realized by a split extension. This is called the semi-direct product:

Definition A.2.4. Let $\varphi : Q \rightarrow \mathrm{Aut}(N)$ be as above. The semi-direct product of Q and N along φ is the group $N \rtimes_{\varphi} Q$ which has $N \times Q$ as its underlying set with multiplication defined by

$$(n_1, q_1)(n_2, q_2) := (n_1 \varphi_{q_1}^{n_2}, q_1 q_2)$$

The way to memorize this is: We really want to think of the n_i and q_i as element in an ambient group. So we should be able to write

$$n_1 q_1 n_2 q_2 = n_1 \left(q_1 n_2 q_1^{-1} \right) q_1 q_2$$

And then, we want conjugation by q_1 correspond to φ.

The semi-direct product is a split extension of N by Q which induces φ. The proof of this fact is usually left to the student as an exercise.

Remark A.2.5. The topological analogue of the direct product of two groups is the Cartesian product of two spaces. The analogue of a semi-direct product of groups is a fibration.

Definition A.2.6. The wreath product of G and H is the semi-direct product

$$G \wr H := \left(\prod_{h \in H} G \right) \rtimes H$$

where H acts on $\prod_{h \in H} G$ by left-multiplication on the indices.

A.2.2 Free Products / Amalgamated Products / HNN-Extensions

!!!...!!!

A.2.3 Graphs of Groups and Spaces

!!!...!!!

A.3 Faithful Representations

!!!...!!!

Appendix B

Geometry: Nuts and Bolts

B.1 Metric Spaces

Definition B.1.1. A metric space is proper if all closed balls are compact.

The length pseudo metric of a metric space X is given by

$$(x, y) \mapsto \inf_{p:x \longrightarrow y} |p| .$$

If the metric and the induced length pseudo metric coincide, the space X is called a length space.

A geodoesic (segment) in a metric space X is a distance preserving map

$$\gamma : [0, |\gamma|] \to X$$

whose domain is an interval. The length of the domain is the length of the geodesic. A bi-infinite geodesic or a geodesic line is a distance preserving map

$$\gamma : \mathbb{R} \to X,$$

and a geodesic ray is a distance preserving map

$$\gamma : \mathbb{R}^+ \to X.$$

A geodesic space is a a metric space wherein any two points are joined by a geodesic. A Hadamard space is a complete geodesic space.

Exercise B.1.2. Let X be a complete metric space. Show that X is geodesic if "it has midpoints", i.e., for every pair $\{x, y\}$ there is a point z such that $d(x, z) = d(y, z) = \frac{1}{2}d(x, y)$.

Exercise B.1.3. Let X be a complete metric space. Show that X is a length space if "it has approximate midpoints", i.e., for every pair $\{x, y\}$ and every $\varepsilon > 0$ there is a point z such that $d(x, z), d(y, z) \leq \varepsilon + \frac{1}{2}d(x, y)$.

Definition B.1.4. A sequence $f_i : X \to Y$ of maps between metric spaces is equicontinuous if for any $\varepsilon > 0$ there is a $\delta > 0$ such that, for every i: if $d(x, y) \leq \delta$, then $d(f_i(x), f_i(y)) \leq \varepsilon$.

Fact B.1.5 (Arzelà-Ascoli). *Let X be a compact metric space and Y be a separable metric space, then every sequence of equicontinuous functions $f_i : X \to Y$ has a subsequence that converges uniformly on compact subsets to a continuous map $f : X \to Y$.*

Corollary B.1.6. *In a complete geodesic space with unique geodesics, those geodesics vary continuously with their endpoints.*

Fact B.1.7 (Hopf-Rinow). *Let X be a complete, locally compact, length space. Then X is a proper geodesic space.*

Let M_κ^m be the simply connected Riemannian manifold of constant curvature κ of dimension m. (There is only one up to isometry.)

Definition B.1.8. A geodesic triangle Δ inside a metric space X is called κ-admissible if there is a triangle in M_κ^2 that has the same side length. Such a triangle is called a κ-comparison triangle.

Note that every point of Δ has a corresponding point in the comparison triangle. The triangle Δ is called κ-thin if distances of points in Δ are bounded from above by the distances of their corresponding points in a κ-comparision triangle.

A complete geodesic space X is CAT(κ) if any κ-admissible triangle in X is κ-thin. The space X is said to have curvature $\leq \kappa$, if it is locally CAT(κ), i.e., every point has an open ball around it so that this ball is a CAT(κ) space.

Example B.1.9. Any Riemannian manifold with non-positive sectional curvature is CAT(0).

Exercise B.1.10. Show that any two points in a CAT(0) space are connected by a unique geodesic segment.

Fact B.1.11 (Cartan-Hadamard). *Let X be a connected complete metric space of curvature $\leq \kappa \leq 0$. Then the universal cover of X with the induced length metric is CAT(κ). Moreover, geodesic segments in the universal cover are unique and vary continuously with their endpoints.*

Definition B.1.12. Let X be a CAT(0) space. A flat strip in X is a convex subspace that is isometric to a strip in the Euclidean plane bounded by two parallel lines.

Exercise B.1.13 (Flat Strip Theorem). Let X be CAT(0) and let $\gamma : \mathbb{R} \to X$ and $\gamma' : \mathbb{R} \to X$ be two geodesic lines. Show that the convex hull of these two lines is a flat strip provided that the geodesic lines are asymptotic, i.e., the function $d(\gamma(t)\,,\gamma'(t))$ is bounded.

Exercise B.1.14. Let X be a connected complete metric space of curvature $\leq \kappa \leq 0$. Show that every free homotopy class has a representative that is a closed geodesic. Moreover, any two such representatives bound a “flat annulus”, i.e., they lift to bi-infinite geodesics in the universal cover that bound a flat strip.

Definition B.1.15. A subset $X' \subseteq X$ is convex if with any pair of points in X it contains all geodesic segments joining them.

Fact B.1.16. *Let X' be a convex complete subspace of the proper CAT(0) space X. Then there is a nearest-point projection*

$$\pi : X \to X'$$

that takes every point $x \in X$ to the point in X' that is nearest to x. This image point exists by properness and is unique by CAT(0)-ness. For any point x outside X', the geodesic segment $[x, \pi(x)]$ is perpendicular to X'. Moreover π is a distance non-increasing map.

B.2 Piecewise Geometric Complexes

Recall that M^m_κ is the simply connected manifold of constant curvature κ of dimension m.

Definition B.2.1. A piecewise M_κ-complex is a CW-complex whose cells have the structure of convex polyhedra in some M_κ and whose attching maps are isometric identifications with faces.

Fact B.2.2 (Bridson). *If a connected piecewise M_κ complex has only finitely many shapes, then it is a complete geodesic space.*

Fact B.2.3. *Let X be a piecewise M_κ complex. For $\kappa \leq 0$, the following are equivalent:*

1. *X is CAT(κ).*

2. *X has unique geodesics.*

3. *All links in X are CAT(1) and X does not contain a closed geodesic.*

4. *All links in X are CAT(1) and X is simply connected.*

For $\kappa > 0$, the following are equivalent:

1. *X is CAT(κ).*

2. *For any two points $x, y \in X$ of distance $d(x, y) < \frac{\pi}{\sqrt{\kappa}}$, there is a unique geodesic segment joining x and y.*

3. *All links in X are CAT(1) and X does not contain a closed geodesic of length $< \frac{2\pi}{\sqrt{\kappa}}$.*

Fact B.2.4 (Gromov's Lemma). *A piecewise spherical simplicial complex all of whose edges have length $\frac{\pi}{2}$ is CAT(1) if and only if it is a flag complex, i.e., any collection of vertices that are pairwise joined by edges forms a simplex. (Whenever you see a possible 1-skeleton of a simplex, the simplex is actually there.)*

A piecewise spherical simplicial complex is called metrically flag if every collection of vertices that are pairwise joined by edges forms a simplex provided there is a non-degenerate spherical simplex with those edge lengths.

Fact B.2.5 (Moussong's Lema). *Let X be a piecewise spherical simplicial complex all of whose edges have length $\geq \frac{\pi}{2}$. Then X is CAT(1) if and only if X is metrically flag.*

Proof that metrically flag implies CAT(1). We follow the argument of D. Krammer [Kram95]. !!! ... !!! **q.e.d.**

Exercise B.2.6. Prove that a CAT(1) piecewise spherical simplicial complex is metrically flag.

B.3 Group Actions

Definition B.3.1. Let X be a metric space, and let $\lambda : X \to X$ be an isometry. The displacement function of λ is the map

$$D_\lambda : x \mapsto d(x, \lambda(x)) .$$

The displacement of λ is

$$D(\lambda) := \inf_{x \in X} D(x) .$$

The min-set of λ is the set

$$\mathrm{Min}(\lambda) := \{x \in X \mid D(x) = D(\lambda)\} .$$

The isometry is parabolic if its min-set is empty. It is semi-simple otherwise. A semi-simple isometry is called elliptic if its displacement is 0 and hyperbolic otherwise.

Observation B.3.2. *Let X be CAT(0). Then the min-set of any semi-simple isometry is a closed, convex, and complete subspace.*

Observation B.3.3. *Let X be CAT(0), let λ be a semi-simple isometry of X, and let X' be a closed, convex, complete, and λ-invariant subspace. Then*

$$D_X(\lambda) = D_{X'}(\lambda)$$

since the nearest-point projection to X' is λ-equivariant and distance non-increasing.

Proposition B.3.4. *Let X be a complete CAT(0) metric space and let λ be a hyperbolic isometry. Then the min-set of λ is a disjoint union of bi-infinite geodesics each of which is fixed by λ set-wise. Indeed, λ acts on each of these axes as a translation of amplitude $D(\lambda)$.*

Proof. Let x be in the min-set of λ. The axis through x is the union of the geodesic segments $\left[\lambda^{k-1}x, \lambda^k x\right]$. To see that this is a geodesic, assume that at one of the break points the angle was not π. The midpoints of the adjacent edges would have distance strictly less than $D(\lambda)$. **q.e.d.**

Corollary B.3.5. *For any semi-simple isometry λ, we have $D\big(\lambda^k\big) = kD(\lambda)$.* **q.e.d.**

B.3.1 Proper Actions

Definition B.3.6. Let G act by isometries on the metric space X. For any subset $Y \subset X$, we define the big stabilizer to be

$$\overline{\mathrm{Stab}}_G(S) := \{g \in G \mid gY \cap Y \neq \emptyset\}.$$

The action is said to be proper if every compact subset $C \subseteq X$ has a finite big stabilizer.

The action is properly proper if for every point $x \in X$ there is an $r > 0$ such that the closed ball $\bar{B}_r(x)$ has a finite big stabilzer.

Exercise B.3.7. Show that every properly proper action is proper, and that every proper action on a proper metric space is properly proper.

Theorem B.3.8. *Let G act properly properly by isometries on the metric space X. Then the following hold:*

1. *For every point $x \in X$ there is an $\varepsilon > 0$ such that*
$$\overline{\mathrm{Stab}}_G\big(\bar{B}_\varepsilon(x)\big) = \mathrm{Stab}_G(x).$$
2. *The action is discontinuous, i.e., the distance between orbits induce a metric (and not just a pseudo-metric) on the quotient space.*
3. *If G acts freely, then the projection $X \to G\backslash X$ is a covering projection and a local isometry.*

Proof. The first clain is easy and the other two follow. **q.e.d.**

B.3.2 Proper Cocompact Actions

Exercise B.3.9. Let G act properly and cocompactly on the length space X. Show that X is complete and locally compact.

By the Hopf-Rinow theorem (B.1.7), we infer:

Corollary B.3.10. *A length space that admits a cocompact proper action is a proper metric space.* **q.e.d.**

Proposition B.3.11. *Let G act properly and cocompactly by isometries on the proper metric space X. Then the following hold:*

1. *There are only finitely many conjugacy classes of point stabilizers.*
2. *Every element $g \in G$ acts by a semi-simple isometry.*

Proof. For both arguments, we fix a compact subset $C \subseteq X$ whose G-translates cover X.

(1) Choose a finite cover of C by finitely many balls $\bar{B}_1, \dots, \bar{B}_r$ whose big stabilizers are all finite. Then for every $x \in X$ there is an element $g \in G$ such that $gx \in \bar{B}_i$ for some i. Thus

$$g^{-1} \operatorname{Stab}(x)\, g \subseteq \overline{\operatorname{Stab}}\big(\bar{B}_i\big) \subseteq \bigcup_i \overline{\operatorname{Stab}}\big(\bar{B}_i\big)\,.$$

However the right hand is finite. Note that we only used strict properness of the action and could do away with properness of the space.

(2) Let (x_i) be a sequence of points in X such that $D_g(x_i) \to D(g)$ as $i \to \infty$. Choose elements g_i such that $y_i := g_i x_i \in C$.

Observe that

$$D_{g_i g g^{-1}} y_i = d\big(g_i^{-1} y_i, g g_i^{-1} y\big) = d(x_i, g x_i) \to D(g)\,.$$

Thus there is an ε such that the expression

$$d\big(x, g_i g g_i^{-1} x\big)$$

is bounded for all i and all $x \in C$ by $2\operatorname{diam}(C) + D(g) + \varepsilon$. Thus there is a closed ball such that all $g_i g g_i^{-1}$ are in its big stabilizer. Therefore, the sequence $g_i g g_i^{-1}$ traverses only finitely many different group elements. Passing through several subsequences, we may assume that $g_i g g_i^{-1}$ is constant and that y_i converges to some point $y \in C$. But then $g_i g g_i^{-1}$ is semi-hyperbolic with y in its min-set. Thus g is semi-hyperbolic with $g_i^{-1} y$ in its min-set. **q.e.d.**

Definition B.3.12. The translation distance of a group element $g \in G$ with respect to a given action on the metric space X is the limit

$$\tau(g) := \lim_{k \to \infty} \frac{d\big(x, g^k x\big)}{k}.$$

Exercise B.3.13. Show that translation distances exists and are independend of the choice of $x \in X$.

Exercise B.3.14. Let λ by a semi-simple isometry of a CAT(0) space X. Show that $\tau(\lambda) = D(\lambda)$.

Definition B.3.15. Let $G = \langle \Sigma \rangle$ be a finitely generated group. For every element $g \in G$, the translation length (with respect to Σ) is defined as

$$\underline{\tau}(g) := \lim_{k \to \infty} \frac{\big|g^k\big|}{k}.$$

Exercise B.3.16. Show that the translation length of a group element is well defined, i.e., the limit exists and is independent of the generating set Σ.

Corollary B.3.17. *If G acts isometrically, cocompactly, and properly on a proper $CAT(0)$ space, then the following hold:*

1. *The finite subgroups of G are precisely the subgroups of G that have a global fixed point in X. In particular:*
   (a) *The group G has only finitely many conjugacy classes of finite subgroups (B.3.11(1)).*
   (b) *An element of G has finite order if and only if it acts by an elliptic isometry.*
2. *The group G does not contain a Baumslag-Solitar group*
$$\left\langle a, b \,\middle|\, ba^q b^{-1} = a^p \right\rangle$$
   where $q \neq p$.
3. *The set of translation distances is discrete.*
4. *There is a strictly positive lower bound $\varepsilon > 0$ on the translation length of non-torsion elements of G.*

Proof.

(1) Since the action is proper, every point stabilizer is finite. It remains to prove that every finite subgroup has a global fixed point.

For any compact subset $C \subseteq X$ let $\bar{B}_C$ denote the smallest closed ball that contains C. (It follows from X being CAT(0) that this ball exists and is unique.) Let x_C be the center of $\bar{B}_C$. Note that since $\bar{B}_C$ is defined entirely in metric terms, we have

$$\bar{B}_{gC} = g\bar{B}_C$$

and

$$x_{gC} = gx_C$$

for any group element $g \in G$.

Let $F \leq G$ be a finite subgroup. Let C be the orbit of some point $y \in X$. Note that C is F-invariant. By the preceeding considerations, the point x_C is F-invariant, too.

(2) As a matter of fact, Baumslag-Solitar groups are torsion free. Thus, if G contained a copy, the elements inside the subgroup would be hyperbolic. Since a^q and a^p are conjugate, they have the same displacement. On the other hand, their displacements have the ratio $\frac{q}{p}$. It follows that the displacement of a is 0 which is a contradiction.

(3) Suppose we have a sequence g_i of group elements with different translation distances that converge to a limit L. Passing to a sequence of conjugates (which have the same translation lengths), we find a sequence of points $x_i \in C$ such that x_i is in the min-set of g_i. The contradiction is assumed at any accumulation point $x \in C$ as hit by the different g_i: The ball of radius $L + 3\varepsilon$ is not moved off itself by any g_i where $\tau(g_i) = D(g_i)$ is ε-close to L and x_i is ε-close to x.

(4) Observe that for any point $x \in X$, we have

$$\begin{aligned} d(x, g_1 \cdots g_r x) &\leq d(x, g_1 \cdots g_{r-1} x) + d(g_1 \cdots g_{r-1} x, g_1 \cdots g_r x) \\ &= d(x, g_1 \cdots g_{r-1} x) + d(x, g_r x) \\ &\leq \sum_{i=1}^{r} d(x, g_i x) \\ &\leq r \max_{i \in \{1,\ldots,r\}} d(x, g_i x)\,. \end{aligned}$$

Thus, fixing a generating set, we can find a constant C such that for any element $g \in G$,

$$\frac{d\big(x, g^k x\big)}{k} \leq C \frac{\big|g^k\big|}{k}.$$

Passing to the limit, we find

$$\underline{\tau}(g) \geq \frac{\tau(g)}{C} \geq \varepsilon$$

for some $\varepsilon > 0$ that exists by (3). **q.e.d.**

B.3.3 Abelian and Solvable Subgroups

Exercise B.3.18. Let G be virtually $\mathbb{Z}^n$, and let H be a subgroup of G that is isomorphic to $\mathbb{Z}^n$. Show that H has finite index in G. (Hint: Consider a finite index $\mathbb{Z}^n$ inside G and the action of H on $G/\mathbb{Z}^n$.)

Exercise B.3.19. Let G be finitely generated. Show that G is virtually abelian provided the commutator subgroup $[G, G]$ is finite. (Hint: Let H be the centralizer of $[G, G]$ in G and show that the center of H has finite index in H and that H has finite index in G.)

Appendix C

Topology: Nuts and Bolts

C.1 Topological Categories

There are way too many topological spaces – most of them pathological. Hence real topology (as opposed to point set topology) takes place in smaller categories. The additional structure makes their objects amenable to stronger methods. Fortunately, all spaces we are interested in, are very nice.

C.1.1 Paracompact Spaces

A Hausdorff space X is paracompact if every open cover has a locally finite subcover. Any space you would want to meet is paracompact. It is the minimum requirement for a space to be considered "nice" in any sense. The technical importance of paracompactness is

Fact C.1.1. *If X is paracompact and $\mathcal{U}$ is an open cover, then there is a partition of unity subject to $\mathcal{U}$.*

Example C.1.2. All manifolds are paracompact as they have a countable basis. All CW-complexes are paracompact.

C.1.2 Complexes

Cell Complexes come in different flavors. The most general kind is build from a set of vertices by successively glueing in higher dimensional cells. The cells of dimension $m + 1$ are balls whose boundary spheres are mapped via attaching maps into the m-skeleton which has been constructed already. The generality lies in the fact, that we do not make any assumptions about the attaching maps.

Piecewise Euclidean cell complexes have cells that have the additional structure of convex polyhedra in Euclidean space. Here the attaching maps are supposed to identify a boundary cell of an $(m+1)$-cell isometrically with a cell in the m-skeleton, which then is called a face of the $(m+1)$-cell. The category of piecewise Euclidean cell complexes is suitable for geometric methods because broken straight line paths have lengths: you can measure the lengths of the pieces inside the cells.

Generalized simplicial complexes are piecewise Euclidean cell complexes all of whose cells are regular unit length simplices. They differ from ordinary simplicial complexes only in that two simplices might share more than one boundary simplex. For instance you can realize a sphere simply by gluing two simplices along their boundary. The result is a generalized simplicial complex which is not a simplicial complex. Thus, we call a piecewise Euclidean cell complex combinatorial if the intersection of any two closed cells is empty or consists of just one closed cell. Note that links in combinatorial complexes are combinatorial.

A simplicial complex is called a flag complex if it does not have "hollow simplices", i.e., if you see the one skeleton of a triangle, there is a 2-simplex filled in; if you see the 2-skeleton of a tetrahedron, there is a 3-simplex filled in; and so on. Observe that any barycentric subdivision of a combinatorial cell complex is a flag complex.

Another useful specialization of piecewise Euclidean complexes are cube complexes whose cells look like unit cubes. Combinatorial cube complexes are more combinatorial than topological objects – very much like simplicial complexes. Compared to simplicial complexes they have the advantage that the cross product of cube complexes is a cube complex. Moreover, you can define links in cube complexes in two ways: in one way the links are simplicial complexes in the other way the links are cube complexes themselves.

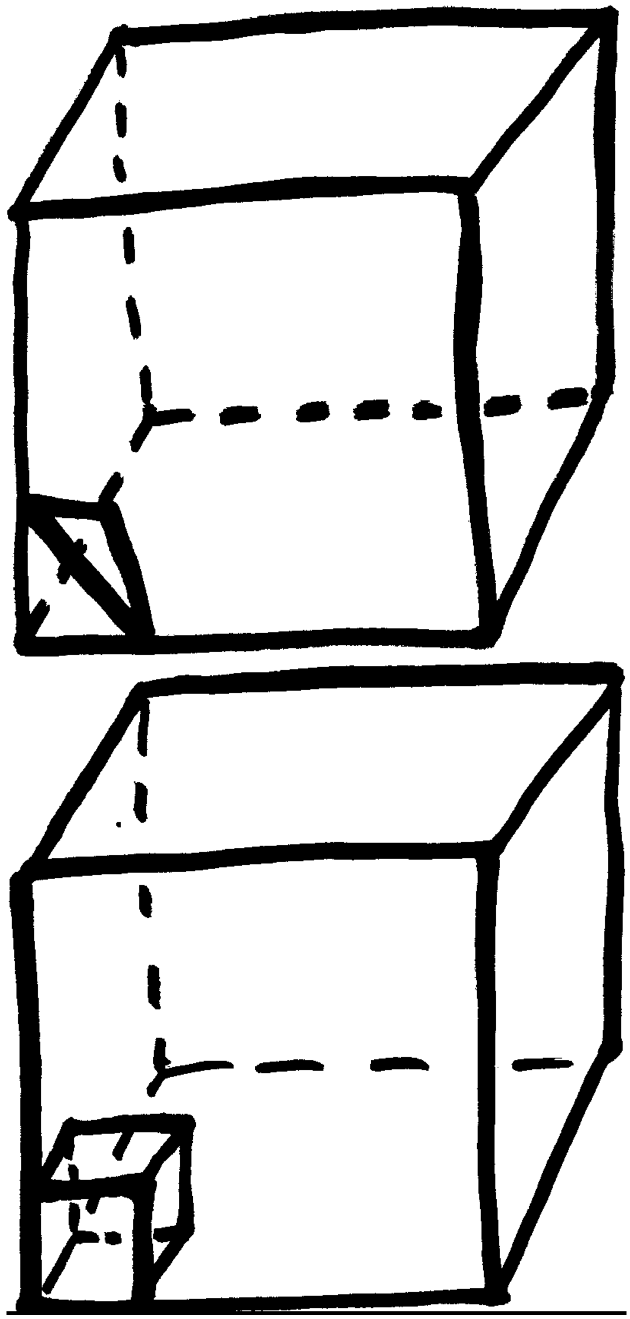

It turns out that for geometry, the links a better regarded as simplicial complexes – in fact, the simplices should be viewed metrically as spherical simplices.

C.1.3 Posets

A combinatorial substitute for cell complexes is the notion of a poset (partially ordered set). Every cell complex gives rise to a poset: The elements of the posets are the cells of the complex and the order is given by the face relation. Employing this analogy, let us agree to call the elements of a poset (open) cells. Moreover, if we have two cells p and q satisfying $p \preceq q$, we say that p is a face of q or, equivalently, that q is a coface of p. If $p \preceq q$ with $p \neq q$, we say that p is a strict face of q. The boundary of a cell is the subposet of all its strict faces, and the link is the subposet of all its strict cofaces. Thus, the link should properly be called the coboundary.

For each poset, the finite totally ordered subsets (chains) form a simplicial complex, whose geometric realization is considered the geometric realization of the poset. This way, all topological notions apply to posets. The analog of a closed cell in a poset is the subposet formed by a cell together with its boundary. The analog of a closed subset is a closed subposet, that is, a subposet that contains all faces of any cell it contains. I never encountered someone using the dual notion of a coclosed subposet.

Let us denote by P^{op} the poset P with the order relation reversed. Observe that P and P^{op} have the same geometric realization as a finite totally ordered subset of one is also a totally ordered subset of the other.

A morphism $f : P \to Q$ of posets is just a map that preserves the partial ordering. Given a cell $q \in Q$, we have define the fibre of f over q to be the subposet

$$f/q := \{p \in P \mid f(p) \preceq q\}$$

and the cofibre to be

$$q\backslash f := \{p \in P \mid q \preceq f(p)\}.$$

C.2 Computing Homotopy Groups

A space X is called m-connected if, for every $i \leq m$ any map from $S_i \to X$ extends to a map $D_{i+1} \to X$. Note that -1-connected means non-empty and that 0-connected is the same as path-connected and non-empty.

C.2.1 Fibrations and Fibre Bundles

For $i \geq 0$, let

$$\nu : D_{i-1} \hookrightarrow D_i$$

be the embedding of D_{i-1} in D_i as the northern hemisphere and let

$$\sigma : D_{i-1} \hookrightarrow D_i$$

be the embedding as the southern hemisphere.

Definition C.2.1. A map $\pi : E \to X$ has the homotopy lifting property in dimension m if for every commutative diagram

$$\begin{tikzcd} D_{m-1} \arrow[r, "\bar{f}"] \arrow[d, "\nu"'] & E \arrow[d, "\pi"] \\ D_m \arrow[r, "f"'] & X \end{tikzcd}$$

there exists a lift $\tilde{f} : D_m \to E$ such that

$$\begin{tikzcd} D_{m-1} \arrow[r, "\bar{f}"] \arrow[d, "\nu"'] & E \arrow[d, "\pi"] \\ D_m \arrow[r, "f"'] \arrow[ur, "\tilde{f}"] & X \end{tikzcd}$$

commutes.

A map that has the homotopy lifting property in all dimensions is called a fibration.

Example C.2.2. Every covering space projection is a fibration.

Definition C.2.3. A fibre bundle over X with fibre F is a map $f : E \to X$ such that:

1. There is an intersection closed open cover $\mathcal{U}$ of X and a familiy of local trivializations $(\phi_U : U \times F \to E)_{U \in \mathcal{U}}$ that are homeomorphism onto their image and preserve fibres, i.e., the following diagram

   $$\begin{tikzcd} U \times F \arrow[r, "\phi_U"] \arrow[d] & E \arrow[d] \\ U \arrow[r] & X \end{tikzcd}$$

   commutes.

2. For every inclusion $U \hookrightarrow V$, there is a fibre preserving change of coordinates $\phi : U \times F \to V \times F$ such that

$$\begin{tikzcd}
& U \times F \arrow[dl] \arrow[dr, "\phi_U"] \arrow[dd, "\phi"] & \\
X & & E \\
& V \times F \arrow[ul] \arrow[ur, "\phi_V"'] &
\end{tikzcd}$$

   commutes

Exercise C.2.4. Prove that fibre bundles are fibrations.

Suppose $\pi : E \to X$ has the homotopy lifting property in dimensions $\leq m+1$. Let F be the preimage of the base point in X. Recall that every element in $\pi_m(X)$ can be represented as a map $f : D_m \to X$ that takes the whole boundary sphere $\partial(D_m)$ to the base point of X. Let $\tilde{f}$ be the lift in

$$\begin{tikzcd}
D_{m-1} \arrow[r, "\bar{f}"] \arrow[d, "\nu"'] & E \arrow[d, "\pi"] \\
D_m \arrow[ur, "\tilde{f}"] \arrow[r, "f"'] & X
\end{tikzcd}$$

where D_{m-1} entirely sent to the basepoint of E. Then the composition

$$D_{m-1} \xrightarrow{\sigma} D_m \xrightarrow{\tilde{f}} E$$

maps D_{m-1} to F and sends the boundary sphere $\partial(D_{m-1})$ to the base point. Thus, it defines an element of $\pi_{m-1}(F)$.

Exercise C.2.5. Prove that this map is a well defined group homomorphism $\pi_m(X) \to \pi_{m-1}(F)$.

Exercise C.2.6. For $i \leq m$, show that the sequence

$$\pi_i(E) \to \pi_i(X) \to \pi_{i-1}(F)$$

is exact in the middle.

Exercise C.2.7. For $i \leq m$, show that the sequence

$$\pi_i(X) \to \pi_{i-1}(F) \to \pi_{i-1}(E)$$

is exact in the middle.

Exercise C.2.8. The aim of this exercise is to show that Eilenberg-Maclane complexes are unique up to homotopy equivalence. Let G be a group, and let X and Y be two cell complexes. Assume that G acts on both complexes freely and cellularly. Construct two G-equivariant maps $f : X \to Y$ and $h : Y \to X$ such that the composition $h \circ f$ is G-equivariantly homotopic to the identity on X and $f \circ h$ is G-equivariantly homotopic to the identity on Y. Conclude that

$$G \backslash X \simeq G \backslash Y.$$

(Hint: Use induction on skeleta for the construction of the maps as well as for the construction of the homotopies.)

Exercise C.2.9. Let C_n be a non-trivial finite cyclic group. Find a "nice" Eilenberg-Maclance complex for this group and compute its homology. In particular, demonstrate that there are non-vanishin homology groups in arbitrary high dimensions.

Conclude that no Eilenberg-Maclance space for C_n has finite dimension.

The following is immediate:

Corollary C.2.10. *The braid group B_n is torsion free.* **q.e.d.**

C.2.2 Combinatorial Morse Theory

Let X be a piecewise Euclidean cell complex. A Morse function on X is a map $h : X \to \mathbb{R}$ that is affine on closed cells and satisfies the following slope condition:

> There is an $\varepsilon > 0$ such that
>
> $$\varepsilon < |h(w) - h(v)|$$
>
> for all pairs of vertices v and w joined by an edge in X.

Note that h is, in particular, non-constant on edges.

We think of $h(x)$ as the height of the point $x \in X$. For each vertex v, we define the descending link $\mathrm{Lk}^{\downarrow}(v)$ with respect to h to be that part of $\mathrm{Lk}(v)$ spanned by all the cells that contain v as a vertex of maximum height. For each height s, define the sublevel set $X_s := \{x \in X \mid h(x) \leq s\}$.

Morse Lemma C.2.11. *For any two heights $s < t$ with $t - s < \varepsilon$*

$$X_t \simeq X_s \cup_D C$$

where

$$D := \biguplus_{s<h(v)\leq t} \mathrm{Lk}^{\downarrow}(v)$$

and

$$D := \biguplus_{s<h(v)\leq t} \mathrm{Cone}\big(\mathrm{Lk}^{\downarrow}(v)\big)\,.$$

In words: In order to get X_t from X_s up to homotopy equivalence, you have to cone off the descending links of the vertices that are in $X_t - X_s$.

Proof. Here is a sequence of pictures:

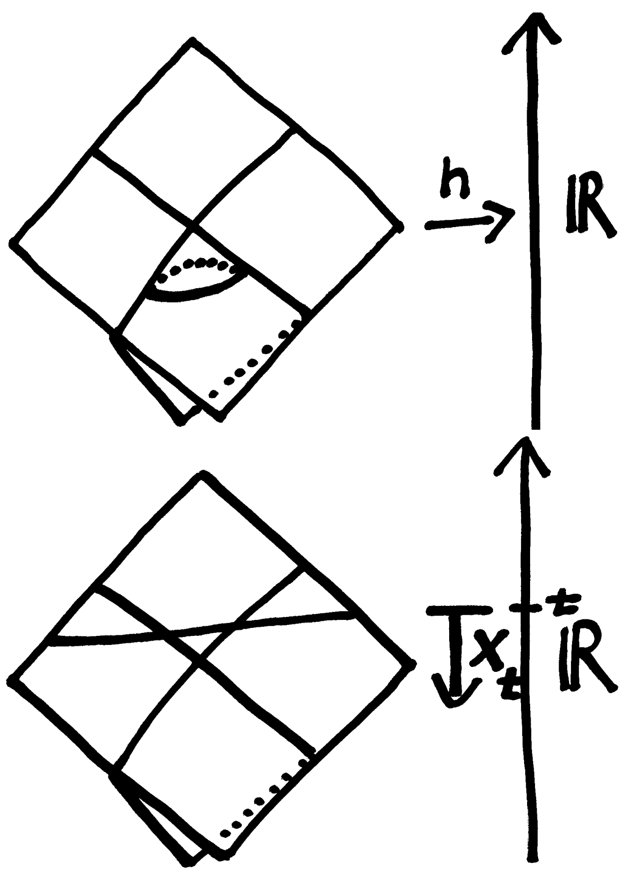

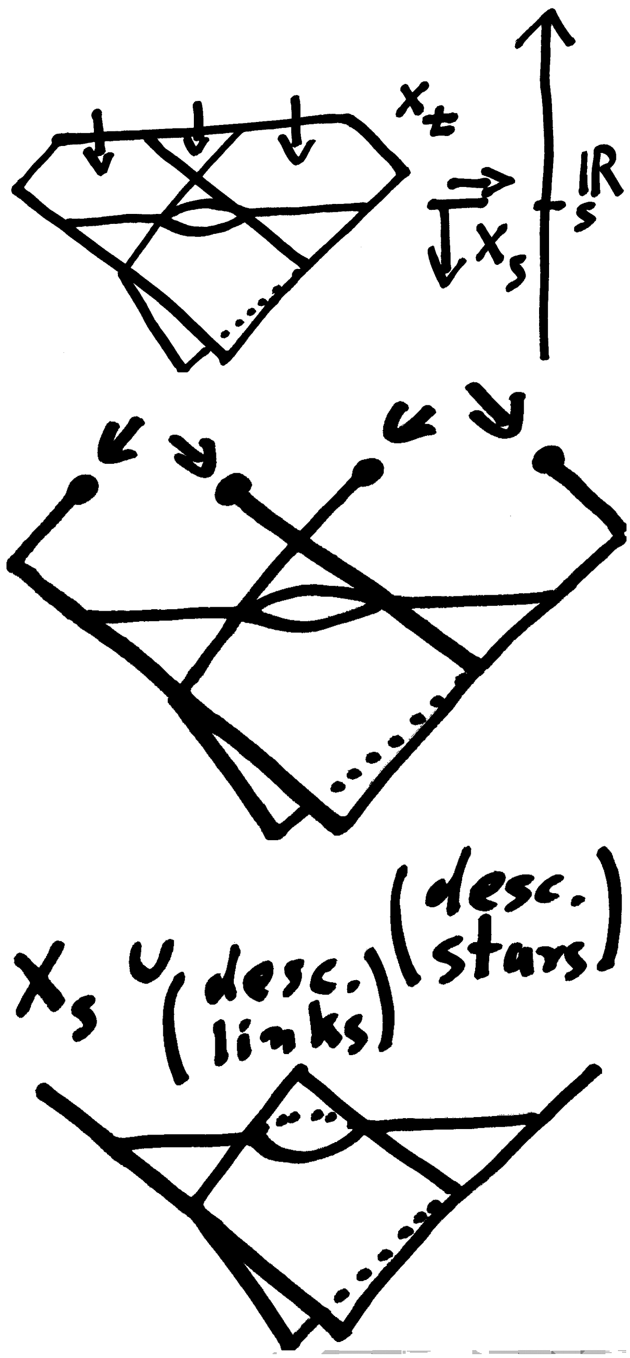

All we do is pushing in free faces. This defines a deformation retraction. **q.e.d.**

Corollary C.2.12. *If all descending links are contractible, then the homotopy type of sublevel sets does not change as the height increases. Hence the whole space has the homotopy type of any of its sublevel sets.* **q.e.d.**

C.2.3 The Vietoris-Smale-Quillen Argument

Lifting a theorem of Vietoris from homology to homotopy, S. Smale proved:

Fact C.2.13 ([Smal57]). *Let X and Y be 0-connected, locally compact, Hausdorff and metrizable. Let $f : X \to Y$ be proper and onto. Assume that X is locally m-connected, that Y is locally $(m-1)$-connected, and that $f^{-1}(y)$ is locally $(m-1)$-connected, for any point $y \in Y$. Then the following hold:*

1. *Y is locally m-connected.*

2. *The map f induces an isomorphism in homotopy groups*

$$\pi_i(X) \to \pi_i(Y)$$

for $0 \leq i < m$ and an epimorphism for $i = m$.

For geometric group theory the combinatorial setting is more suitable. The following version, due to D. Quillen, is stated in the language of posets. Note that it therefore applies to all piecewise Euclidean cell complexes.

Fact C.2.14 ([Quil78]). *Let $f : P \to Q$ be a morphism of posets. Assume that each fibre is m-connected. Then P is m-connected if and only if Y is m-connected.*

Moreover, if each fibre is contractible, then f is a homotopy equivalence.

This is a very powerful tool for computing the connectivity of a space.

C.2.4 Nerves

Definition C.2.15. The nerve of a cover $(U_i)_{i \in I}$ is the simplicial complex

$$N := \left\{ \sigma \subset I \;\middle|\; \bigcap_{i \in \sigma} U_i \neq \emptyset \right\}.$$

Fact C.2.16. *Let X be a paracompact topological space and $\mathcal{U} = (U_i)_{i \in I}$ be a cover by open subsets. If for every simplex σ in the nerve N, the subset*

$$X_\sigma := \bigcap_{i \in \sigma} U_i$$

is a contractible subspace of X, then X and N are homotopy equivalent.

Fact C.2.17. *Let X be a CW-complex and $\mathcal{U} = (U_i)_{i\in I}$ be a cover by closed subcomplexes. If for every simplex σ in the nerve N, the subcomplex*

$$X_\sigma := \bigcap_{i\in\sigma} U_i$$

is contractible, then X and N are homotopy equivalent.

Fact C.2.18. *Let P be a poset $\mathcal{U} = (U_i)_{i\in I}$ be a cover by closed subposets. If for every simplex σ in the nerve N, the subposet*

$$P_\sigma := \bigcap_{i\in\sigma} U_i$$

is contractible, then P and N are homotopy equivalent.

We give a proof of the poset-version to illustrate the connection with the Vietoris-Smale-Quillen argument.

Proof. Consider the map

$$\begin{array}{ccc} P & \xrightarrow{f} & N^{\mathrm{op}} \\ p & \mapsto & \{i \in I \mid p \in U_i\}\,. \end{array}$$

The fibres are easily seen to coincide with the subposets U_σ. Thus, the fibres are contractible, and f is a homotopy equivalence by (C.2.14). **q.e.d.**

Appendix D

Finiteness Properties

Definition D.0.19. A group G is of type F_m if it has an Eilenberg-MacLane complex $\mathrm{K}(G,1)$ with finite skeleta in dimensions $\leq m$.

Remark D.0.20. We do not simply require the m-skeleton to be finite so that the definition makes sense for $m = \infty$. In this case we require the Eilenberg-MacLane complex to have finitely many cells in each dimension.

Observation D.0.21. *A group G is of type F_m if and only if there is an $(m-1)$-connected simplicial complex X with a free and cocompact G-action.*

Corollary D.0.22. *Every group is of type F_0, a group is finitely generated if and only if it is of type F_1, and it is finitely presented if and only if it is of type F_2.* **q.e.d.**

Example D.0.23. Free groups are of type F_∞. So is Thompson's group F. Grigorchuk's group is finitely generated but not finitely presented.

Exercise D.0.24. Let G be of type F_m and let X be a $(m-2)$-connected simplicial complex of dimension $m-1$ with a free, cocompact G-action. Show that X embeds G-equivariantly into an $(m-1)$-connected simplicial complex of dimension m with a free, cocompact G-action.

Infer that a group is of type F_∞ if and only if it is of type F_m for all $m < \infty$.

Exercise D.0.25. Let $N \hookrightarrow G \twoheadrightarrow Q$ be a short exact sequence of groups where N is of type F_{m-1} and G is of type F_m. Then Q is of type F_m.

D.1 Brown's Criterion

The big lemma about finiteness properties is due to K.S. Brown [Bro87a, Theorems (2.2) and (3.2) and Remark (2) on page 48]

Theorem D.1.1. *Let G be a group and D a directed set. Let $(X_\alpha)_{\alpha\in D}$ be a directed system of G-CW-complexes upon which G acts cocompactly by cell-permuting homeomorphisms. Assume that $\varinjlim_{\alpha\in D} X_\alpha$ is $(m-1)$-connected and that for each $\alpha \in D$ and each cell p in the m-skeleton of X_α, the stabilizer $\mathrm{Stab}_G(p)$ is of type $\mathrm{F}_{m-\dim(p)}$. Then the following are equivalent:*

1. *G is of type F_m.*
2. *For each $i < m$, the directed system of homotopy groups $(\pi_i(X_\alpha))_{\alpha\in D}$ is essentially trivial.*

Here a directed system of groups $(H_\alpha)_{\alpha\in D}$ is called essentially trivial if for each $\alpha \in D$ there is an element $\beta > \alpha$ such that the homomorphism $H_\alpha \to H_\beta$ is trivial, i.e., maps everything to the trivial element.

Corollary D.1.2. *Suppose G act cocompactly by cell permuting homeomorphisms on an $(m-1)$-connected CW-complex X such that for every cell p, the stabilizer $\mathrm{Stab}(p)$ is of type $\mathrm{F}_{m-\dim(p)}$. Then G is of type F_m.*

Proof. Consider the directed system

$$X \xrightarrow{\mathrm{id}_X} X \xrightarrow{\mathrm{id}_X} X \xrightarrow{\mathrm{id}_X} \cdots$$

and check that it satisfies the hypotheses of Brown's Criterion. **q.e.d.**

Exercise D.1.3. Show that all finite groups are of type F_∞.

D.2 Applications of Brown's Criterion

Example D.2.1. Let G be a group and let $D := \{K \subseteq G \mid K \text{ is finite}\}$ be the set of finite subsets of G directed by inclusion. For $K \in D$, define the simplicial complex

$$X_K := \{\sigma \mid \sigma \subseteq gK \text{ for some } g \in G\}.$$

Obviously, G acts cocompactly on X_K. Simplex stabilizers conjugate into K and are, therefore, finite. Thus G is of type F_m if and only if $(\pi_i(X_\alpha))_{\alpha\in D}$ is essentially trivial for all $i < m$.

Let H be a subgroup of finite index in G. Then the induced action of H on X_K is still cocompact. Thus we can use the same directed system to detect finiteness properties of H. Thus, we have

Corollary D.2.2. *Let H be a subgroup of finite index in G. Then H is of type* F_m *if and only if G is of type* F_m.

With only little more effort, the same construction yields:

Exercise D.2.3. Let G be of type F_m and let H be a retract of G. Then H is of type F_m.

Proposition D.2.4. *Let $N \hookrightarrow G \twoheadrightarrow Q$ be a short exact sequence of groups where N and Q are of type* F_m*. Then G is of type* F_m.

Proof. Take any free Q complex that proves Q to be of type F_m. Consider this complex as a G-complex where the G-action is given via the projection $G \to Q$. Then, all cell stabilizers equal N and are of type F_m. Thus the same complex proves that G is of type F_m. **q.e.d.**

Definition D.2.5. Let X and Y be two metric spaces. A map $f : X \to Y$ is called Lipschitz if there are constants L and K such that

$$d(x, y) \leq Ld(f(x), f(y)) + K$$

for all $x, y \in X$.

A Lipschitz $f : X \to Y$ is a quasi-retraction if there exists a Lipschitz $h : Y \to X$ and a constant C such that

$$d(f(h(y)), y) \leq C$$

for all $y \in Y$. The map h is called a quasi-section for f. If there is a quasi-retraction $f : X \to Y$, the space Y is called a quasi-retract of X.

The map f is a quasi-isometry if there is a map $h : Y \to X$ such that f is a quasi-retraction with quasi-section h and h is a quasi-retraction with quasi-section f.

Two finitely generated groups are called quasi-isometric if they have quasi-isometric Cayley graphs for some finite generating sets. The notion of a quasi-retract carries over to groups in the same way.

Observation D.2.6. *Since compositions of Lipschitz maps are Lipschitz, quasi-isometry is an equivalence relation on the class of metric spaces.*

Exercise D.2.7. Show that X and Y are quasi-isometric if and only if there is a map $f : X \to Y$ and constants L, K, and C such that the following hold:

1. f is *bilipschitz*, i.e., $\frac{d(f(x), f(y))}{L} - K \leq d(x, y) \leq Ld(f(x), f(y)) + K$ for all $x, y \in X$.
2. f is quasi-surjective, i.e., $Y = \bigcup_{x \in X} \mathbb{B}^x_C$.

Exercise D.2.8. Let Σ and Ξ be two finite generating sets for G. Prove that Γ^G_Σ and Γ^G_Ξ are quasi-isometric.

Proposition D.2.9 ([Alon94]). *Let G be of type F_m. If H is a quasi-retract of G, then H is of type F_m.*

Proof. Our directed set will be $\mathbb{N}$. For $n \in \mathbb{N}$, define

$$\begin{aligned} X_n &:= \{\sigma \subseteq G \mid \operatorname{diam}(\sigma) \leq n\} \\ Y_n &:= \{\sigma \subseteq H \mid \operatorname{diam}(\sigma) \leq n\} \end{aligned}$$

The group G acts on X_n cocompactly and with finite stabilizers. The same holds for H and Y_n. Moreover, both directed systems converge to big simplices. Thus, we can use these directed systems to determine the finiteness properties of these groups.

Let $f : G \to H$ be a retraction with quasi-section $h : H \to G$ with constants L, K and C as in the definition. Then, we form

$$Y_m \xrightarrow{h} X_M \longrightarrow X_N \xrightarrow{f} Y_n$$

where $M \geq L\,(m + K)$ and N is chosen so that the middle map annihilate homotopy in dimensions $< m$. The number n, again, is derived from the Lipschitz constants. Now the composite map is induced by $f \circ h$ which is homotopic to the inclusion map into $Y_{n'}$ for any $n' \geq n + 2C$. Here, we use that fact that two simplicial maps f and h are homotopic if, for each simplex σ, the union $f(\sigma) \cup h\sigma$ is a simplex.

Hence the inclusion $Y_m \subseteq Y_{n'}$ annihilates homotopy groups in dimensions $< m$. This implies that H is of type F_m. **q.e.d.**

Corollary D.2.10. *Finiteness properties are <u>geometric</u>, i.e., they depend only on the quasi-isometry type of a group.*

D.3 The Stallings-Bieri Series

Finiteness properties are not yet well understood. We have, however, some series of groups for which finiteness properties have been established. In this section, we shall discuss the most accessible example of such a series, which is due to J.R. Stallings and R. Bieri.

Consider the exact homomorphism

$$F_{\{x_1,y_1\}} * F_{\{x_2,y_2\}} * \cdots * F_{\{x_n,y_n\}} \to \mathbb{Z}$$

that sends all generators x_i, y_i to 1 and let

$$G_n$$

be the kernel.

Proposition D.3.1. *The group G_n is of type* F_{n-1} *but not of type* F_n.

Proof. The free group $F_{\{x_i,y_i\}}$ acts freely and cocompactly on a regular tree T_i all of whose vertices have degree 4. The homomorphism

$$\begin{array}{rcl} F_{\{x_i,y_i\}} & \to & \mathbb{Z} \\ x_i,\, y_i & \mapsto & 1 \end{array}$$

induces an action of $F_{\{x_i,y_i\}}$ on $\mathbb{R}$ by translations. With respect to this action, we have a *height function*

$$h_i : T_i \to \mathbb{R}$$

that is $F_{\{x_i,y_i\}}$-equivariant. Note that at each vertex of T_i we have two ascending edges (labeled by the generators) and two descending edges (labeled by their inverses).

Put

$$X := \bigtimes_i T_i$$

and consider the height

$$\begin{array}{rcl} h : X & \to & \mathbb{R} \\ (t_i) & \mapsto & \sum_i h_i^t \end{array}$$

The ascending and descending links of vertices in X are spheres of dimension $n-1$. In fact, they arise as joins of ascending, respectively descending, links in the factors T_i. It follows by the Morse lemma (C.2.11) that slices

$$X_t := h^{-1}\left([-t,t]\right)$$

are $(n-2)$-connected: As t increases, we cone off $(n-2)$-connected subcomplexes, thereby not changing homotopy groups in dimensions $\leq n-2$. However, in the limit, we obtain the contractible space X. Thus, the homotopy groups in dimensions $\leq n-2$ were trivial all along.

We use X_t as a directed system of CW-complexes to apply Brown's Criterion. It is obvious that G_n acts freely and cocompactly on any X_t. Since the complexes X_t are already $(n-2)$-connected, it follows that G_n is of type F_{n-1}.

As for the other direction, we have to prove that the directed system $(X_t)_t$ is not essentially $(n-1)$-connected. Note that X, being a product of trees, is a metric space with unique geodesics. Thus, for each vertex $v \in X$, we have a geodesic retraction

$$X - \{v\} \to \mathrm{Lk}(v)\,.$$

It follows that a sphere in X is not 0-homotopic in $X - \{v\}$ unless it has a 0-homotopic image in $\mathrm{Lk}(v)$. Moreover, we have a retraction

$$\mathrm{Lk}(v) = \mathop{\ast}_i \{{x_i}^{\pm}, {y_i}^{\pm}\} \to S_{n-1} = \mathop{\ast}_i \{x_i, y_i\}$$

induced by ${x_i}^{\pm} \mapsto x_i$ and ${y_i}^{\pm} \mapsto y_i$. This way, we recognize ascending links of vertices as retracts of the links.

Now, we are ready to construct, for any specified number t, an $(n-1)$-sphere in X_0 that does not die in X_t. To do this, let us fix a vertex $v \in X$ whose height is $\leq t$. The ascending link $\mathrm{Lk}^{\uparrow}(v)$ is a sphere of dimension $n-1$. It is spanned by the ascending edges starting at v. Each of these edges lives in some component tree T_i and can be extended in this factor to a geodesic ray. This way, we can move the sphere up inside X until it reaches height 0. Let S be the sphere obtained that way. Since we used geodesic rays, S maps homeomorphically to $\mathrm{Lk}^{\uparrow}(v)$ under the geodesic retraction

$$X - \{v\} \to \mathrm{Lk}(v)$$

and is, by our previous considerations, not 0-homotopic in $X - \{v\} \supseteq X_t$.

It follows that the directed system X_t is not essentially $(n-1)$-connected whence G_n is not of type F_n. **q.e.d.**

D.4 Homological Finiteness Properties

Definition D.4.1. A group G is of type FP_m over the commutative ring R if R, regarded as an RG-module via the trivial G-action, has a projective resolution

$$\cdots \to P_i \to \cdots P_1 \to P_0 \twoheadrightarrow R$$

wherein the RG-modules P_i are finitely generated in dimension $i \leq m$.

Exercise D.4.2. Show that a group of type FP_m over $\mathbb{Z}$ is of type FP_m over any ring.

Appendix E

Dictionary

E.1 Properties

Let G be a group.

Definition E.1.1. G is finite if it has only finitely many elements.

Definition E.1.2. G is torsion if every element has finite order.

Definition E.1.3. G is periodic if there is a number $n \in \mathbb{N}$ such that for all $g \in G$.

$$g^n = 1.$$

The number n is called the exponent of G.

Definition E.1.4. G is torsion free if no element has finite order.

Definition E.1.5. G is cyclic if it is generated by one element.

Definition E.1.6. G is indicable if it admits an epimorphism onto the infinite cyclic group.

Definition E.1.7. G is abelian if the commutator subgroup is trivial.

Definition E.1.8. G is perfect if it equals its commutator subgroup. Equivalently, G is perfect if it does not have any non-trivial abelian quotients.

Definition E.1.9. G is s-nilpotent if it is trivial or the commutator subgroup is $(s-1)$-nilpotent.

Definition E.1.10. G is nilpotent if it is s-nilpotent for some s.

Definition E.1.11. G is simple if it has no proper non-trivial normal subgroups.

Definition E.1.12. G is pro-finite if it is the projective/inverse limit of finite groups.

Definition E.1.13. G is Hopfian if every surjective endomorphism $G \twoheadrightarrow G$ is an automorphism.

Definition E.1.14. G is co-Hopfian if every injective endomorphism $G \hookrightarrow G$ is an automorphism.

Definition E.1.15. G is linear if it has a faithfull representation of finite dimension over some field (not necessaryly of characteristic 0).

Definition E.1.16. G is of type F if there is a finite Eilenberg-Maclane space $\mathrm{K}(G, 1)$. An Eilenberg-Maclane space for G is a CW-complex with fundamental group G and a contractible universal cover (equivalentlty, one may require the higher fundamental groups to vanish).

Definition E.1.17. G is finitely generated if G has a finite generating system.

Definition E.1.18. G is finitely presented if G has a finite generating system and all relations among the generators follow from a finite set of relations.

Definition E.1.19. G has finiteness length $\leq m$ if there is an Eilenberg-Maclane space $\mathrm{K}(G, 1)$ with finite m-skeleton.

Definition E.1.20. G has solvable word problem if is is finitely generated and for a fixed generating set there is an algorithm that decides whether a given word in the generators represents the identity.

Definition E.1.21. G has solvable conjugacy problem if it is finitely generated and for a fixed generating set there is an algorithm that decides whether two given words represent conjugated elements.

Definition E.1.22. G has bounded generation if there is a finite sequence $g_1, \ldots, g_r$ of elements in G such that every element $g \in G$ can be written as a product

$$g = {g_1}^{k_1} \cdots {g_r}^{k_r}$$

with exponents $k_i \in \mathbb{Z}$.

Equivalently, G has bounded generation if it is an internal product of finitely many cyclic subgroups.

Definition E.1.23. G has Serre's Property FA if every action of G on a tree has a global fixed point.

Definition E.1.24. G is an automata group if it can be represented as a group of automorphisms of a finitary regular rooted tree all of whose elements can be realized as finite state automata. See (10.2.1) and (10.2.26).

Definition E.1.25. G has Kashdan's Property (T) if if every unitary representation that has almost invariant vectors has an invariant vector. Here a unitary representation of a group G on a Hilbert space $\mathcal{H}$ is said to have almost invariant vectors, if, for any finite subset $K \subseteq G$ and any $\varepsilon > 0$, there is a unit vector $\mathbf{u} \in \mathcal{H}$ satisfying

$$|g\mathbf{u} - \mathbf{u}| < \varepsilon.$$

Definition E.1.26. A group G satisfies the Tits Alternative if each finitely generated subgroups either is virtually solvable or contains a non-abelian free group.

Definition E.1.27. G is amenable if there is a left-invariant finitely additive probabilty measure defined on all subsets of G.

Definition E.1.28. G has polynomial growth if it is finitely generated and the volume of a ball in the Cayley graph grows polynomially with the radius.

Definition E.1.29. G has exponential growth if it is finitely generated and the volume of a ball in the Cayley graph grows expoentially with the radius.

Definition E.1.30. G has intermediate growth if it is finitely generated and has neither polynomial nor exponential growth. (Note that growth is at least polynomial and at most exponential.)

Definition E.1.31. G satisfies a quadratic isoperimetric inequality if if has a finite presentation $\mathcal{P}$ and there is a quadratic function $\delta : \mathbb{N} \to \mathbb{N}$ such that any loop of length l in the Cayley complex for $\mathcal{P}$ bounds a disc of area $\leq \delta_l$.

Definition E.1.32. G has the finite intersection property if the intersection of any two finitely generated subgroups in G is finitely generated.

Definition E.1.33. G is a Burnside group if it is finitely generated, infinite, and torison.

Definition E.1.34. G is SQ-universal if every countable group embedds into a quotient of G.

Definition E.1.35. G is combable with respect to the generating set Σ, if there is a constant C and distinguished paths from 1_G to each vertex $v \in \Gamma_\Sigma^G$ that have the C-fellow traveler property. That is, whenever we have two distinguished paths p and q starting at $1_G \in \Gamma_\Sigma^G$ whose endpoints have distance ≤ 1, the following inequality holds along the paths:

$$d(p_t, q_t) \leq C.$$

Here the paths are traversed with unit speed. If all the combing paths can be chosen to be geodesics, the group G is called geodesically combable.

Definition E.1.36. G is subgroup separable if for every subgroup $H \leq G$ and every element $g \in G - H$, there is a normal subgroup $N \trianglelefteq G$ of finite index such that $H \subseteq N$ but $g \notin N$.

E.2 Prefixes

Let "blah", "foo", and "bar" be properties, and let G be a group.

Definition E.2.1. G is foo–by–bar if there is a short exact sequence

$$1 \to H \to G \to F \to 1$$

where H is foo and F is bar.

Remark: You might find something like foo–by–bar–by–blah, which means (foo-by–bar)–by–blah.

Definition E.2.2. G is meta-blah if it is blah–by–blah.

Definition E.2.3. G is virtually blah if G has a blah-subgroup of finite index.

Definition E.2.4. A subgroup N of G is co-blah, if it is normal and the quotient G/N is blah.

Definition E.2.5. G is residually blah, if the intersection of all co-blah subgroups of G is trivial.

Definition E.2.6. G is locally blah if every finitely generated subgroup of G is blah.

Definition E.2.7. G is poly-blah if there is a subnormal series

$$1 = G_1 \trianglelefteq G_2 \trianglelefteq G_3 \trianglelefteq \cdots \trianglelefteq G_{n-1} \trianglelefteq G_n = G$$

with all quotients G_{i+1}/G_i being blah.

Definition E.2.8. G is just blah if it is blah but does not have any proper quotients that are blah.

Example E.2.9. Some well known concepts are just shorthands:

- solvable = poly-abelian
- coherent = locally finitely presented

E.3 Metaproperties

Definition E.3.1. Let "blah" be a property. It is called subgroup closed if every subgroup of a blah group is blah.

It is quotient closed if every epimorphic image of a blah group is blah.

It is called extension closed if every meta-blah group is blah.

Remark E.3.2. The following properties are subgroup closed:

- finite
- free
- cyclic
- abelian
- nilpotent
- solvable
- amenable

Exercise E.3.3. Prove: If foo and bar are subgroup closed, so is foo–by–bar.

Remark E.3.4. The following properties are quotient closed:

- finite
- abelian
- nilpotent
- solvable
- finitely generated

- Serre's Property FA

Remark E.3.5. The following properties are extension closed:

- finite
- solvable
- finitely generated
- finitely presented
- finiteness length $\leq m$
- finite type
- Serre's Property FA

E.4 The Unnamed Famous

There is a bunch of theorems that fit into the following schemes:

- If ..., then G has only finitely many conjugacy classes of finite subgroups.
- If ..., then all solvable subgroups of G are finitely generated and virtually abelian.
- If ..., then all elements of infinite order in G have translation length bounded away from 0.

Exercise E.4.1 (M. Bridson). Let G be polycylic and assume that all elements of infinite order have translation length bounded away from 0. Show that every solvable subgroup of G is finitely generated and virtually abelian.

Exercise E.4.2. Assume that G has only finitely many conjugacy classes of finite subgroups. Show that G is virtually torsion free provided it is residually finite.

Appendix F

Solutions to Selected Exercises

(1.4.28) Observe that it suffices to prove that G is abelian as we know precisely what the torsion free finitely generated abelian groups look like. Now suppose we had a non-trivial homomorphism $G \to \mathbb{R}$. Then the kernel cannot contain an element of infinite order as the infinite cyclic subgroup it generates had linear growth and infinite index which implies by (1.4.27) that G had at least quadratic growth. So all elements in the kernel are torsion for which reason the kernel is trivial as G is torsion free. Hence G would be abelian, and we would be done. Thus our aim is to find a non-trivial homomorphisms $\varphi : G \to \mathbb{R}$.

Let H be an infinite cyclic subgroup of finite index in G. Then, the kernel of the action of G on the finite set of cosets G/H is a normal subgroup N of finite index in G that is contained in H. Hence we may assume without loss of generality that H is normal. Let $\{g_1, \ldots, g_r\}$ be a generating set for G, and let t be a generator for H. Since G/H is finite, every generator has a power in H wherefore every generator commutes with a certain power of t. Passing to an even smaller subgroup, we may therefore assume that H is not only normal but, in fact, central in G.

Fix a generator t for H and coset representatives $g_1, \ldots, g_r$. Then every group element g can be written in a unique way as

$$g = g_i t^{k_g}.$$

It is easy to check that, for any finitely additive measure μ on G,

$$\varphi : h \mapsto \int_{g \in G} k_{hg} - k_g \, \mathrm{d}_\mu \, g$$

defines a homomorphism with $\varphi(t) = 1$.

(2.2.9) Let G be finitely generated and residually finite, and let $\lambda : G \twoheadrightarrow G$ be a surjective endomorphism. Suppose that λ is not an automorphism. Then, we can chose a non-trivial element $g \in G$ in the kernel of λ. Let $\varphi : G \to F$ be a homomorphism to a finite group such that $\varphi(g)$ is non-trivial.

We consider the set $\mathrm{Hom}(G; F)$ of all homomorphisms from G to F. As G is finitely generated, $\mathrm{Hom}(G; F)$ is a finite set. Note that

$$\begin{array}{ccc} \mathrm{Hom}(G;F) & \longrightarrow & \mathrm{Hom}(G;F) \\ \psi & \mapsto & \psi \circ \lambda \end{array}$$

is injective since λ is onto. It follows, that this operation is just a permutation of the finite set $\mathrm{Hom}(G; F)$. In particular, the operation has finite order, and for some number of factors, we have

$$\varphi = \varphi \circ \lambda \circ \lambda \circ \cdots \circ \lambda,$$

which implies that $\varphi(g)$ is trivial. This is a contradiction.

(2.4.25) Let U be an open neighborhood of e. The tree T is connected and covered by the G translates of U. Hence G is generated by

$$\{g \in G \mid gU \cap U \neq \emptyset\} = G_v \cup G_w$$

Hence G_v and G_w generate G.

Now, consider the action of G on $\partial_\infty T$. Note that cutting the edge e defines two subsets $\mathcal{E}_v := \partial^+_\infty(e) \subset \partial_\infty T$ and $\mathcal{E}_w := \partial^-_\infty(T) \subset \partial_\infty T$. Every non-trivial element of G_v moves $\mathcal{E}_w$ into $\mathcal{E}_v$ and vice versa. So the Ping Pong Lemma (2.2.1) applies provided G_v and G_w have enough elements.

Since G acts transitively on the set of edges, the stabilizer G_v acts transitively on the set of edges meeting in v. Since these are at least one of the two vertices v and w has degree 3 and the other one has at least degree 2, the Ping Pong Lemma (2.2.1) applies. Hence $G = \langle G_v, G_w \rangle = G_v * G_w$.

(10.2.16) Yes, there is. We will prove that

$$x = (x^2, x)\sigma$$

defines a tree automorphism that cannot be realized by a finite state automaton because it takes an infinite string of **L**s to an output that is not ultimately periodic.

It is easy to see that the output is given as follows. Define a map $f : \mathbb{Z} \to \mathbb{Z}$ by

$$f : n \mapsto \left\lceil \frac{3n}{2} \right\rceil = \begin{cases} f_{\text{odd}}(n) := \frac{3n+1}{2} & \text{if } n \text{ is odd} \\ f_{\text{even}}(n) := \frac{3n}{2} & \text{if } n \text{ is even} \end{cases}.$$

Define a sequence a_i by the recursion

$$\begin{aligned} a_1 &= 1 \\ a_{i+1} &= f(a_i) \end{aligned}$$

Then, the i^{th} letter in the output is an $\mathbf{L}$ if and only if a_i is even. Thus, we have to show that the sequence of parities of the a_i is not ultimately periodic.

We do this following an idea of Rodrigo Perez. Suppose, we see the sequence $\mathbf{R}\,\mathbf{R}\,\mathbf{L}\ldots$ starting at i in the output. What does this tell us about a_i? First, we can deduce from the first $\mathbf{R}$ that a_i is odd. Thus, $a_{i+1} = f_{\text{odd}}(a_i)$. The second $\mathbf{R}$, then implies that $a_{i+1} = \frac{3a_i+1}{2}$ is odd, too. Thus, $a_i \cong 3 \mod 4$ and $a_{i+2} = f_{\text{odd}}(f_{\text{odd}}(a_i)) = \frac{9a_i+5}{4}$. Finally, the $\mathbf{L}$ says that $a_{i+2} = \frac{9a_i+5}{4}$ is even, whence $a_i \cong 3 \mod 8$. Note that the congruence is alway uniquely determined since 3 is invertible modulo all powers of 2. Thus, we obtain an infinite system of congruences:

$$\begin{aligned} a_i &\cong 1 \mod 2 \\ a_i &\cong 3 \mod 4 \\ a_i &\cong 3 \mod 8 \\ \vdots\ &\ \vdots\ \ \vdots \end{aligned}$$

The key point is that an infinite system of such congruences determines its solution uniquely if it has a solution at all. Therefore, if the sequence of parities was ultimately periodic, we would find the same tail starting at two different positions i and j. The correponding numbers a_i and a_j satisfy the same congruences mod all powers of 2, whence $a_i = a_j$. Since the sequence $a_1, a_2, \ldots$ is visibly increasing, we infer $i = j$.

(10.2.31) We define a finite directed graph on the vertex set $\mathcal{V}_\Gamma := \mathcal{V}_A \times \mathcal{V}_B$. We have a directed edge from (S_1, S_2) to (T_1, T_2) if there is a letter $a \in \mathcal{A}$ such that there are directed edges labeled by a from S_i to T_i. Obviously, this graph can be constructed effectively.

The pair whose coordinates are the distinguished start vertices in A_1 and A_2 serves as a base point in Γ. There is a standard algorithm that computes the set $\mathcal{P} \subset \Gamma$ be the set of vertices that can be reached from the base point by a directed path. Hence the statement follows from the following claim:

> The two automata A_1 and A_2 define the same transformation $\mathcal{A}^* \to \mathcal{A}^*$ if and only if for each pair $(v_1, v_2) \in \mathcal{P}$ the labels of v_1 and v_2 agree.

This is, however, is clear: The directed paths in Γ correspond precisely to the paths taken in the individual automata when given identical inputs.

(D.2.3) Let $\pi : G \to H$ be the retraction. We do not label the section but regard H as a subgroup of G. Let D be the set of all those finite subsets of G that contain their π-image:

$$D := \{K \subseteq G \mid \pi\,\mathrm{im}(K) \subseteq K \text{ and } K \text{ is finite }\}.$$

For any $K \in D$, we have a retraction

$$X_K \to Y_K$$

induced by π where

$$\begin{aligned} X_K &= \{\sigma \mid \sigma \subseteq gK \text{ for some } g \in G\} \\ Y_K &= \{\sigma \mid \sigma \subseteq hK \text{ for some } h \in H\} \subseteq X_K. \end{aligned}$$

Since G is of type F_m, there is $L \in D$ with $K \subseteq L$ such that

$$X_K \hookrightarrow X_L$$

induces the trivial map in homotopy in dimensions $< m$. Thus

$$Y_K \hookrightarrow X_K \hookrightarrow X_L \to Y_L$$

induces the trivial map in homotopy in dimensions $< m$, as well. Thus the directed system

$$(\pi_i(Y_K))_{K \in D}$$

is essentially trivial for $i < m$. Hence H is of type F_m.

(D.4.2) Let

$$\cdots \to P_i \to \cdots P_1 \to P_0 \twoheadrightarrow \mathbb{Z}$$

be a projective resolution of $\mathbb{Z}$ by $\mathbb{Z}G$-modules wherein P_i is finitely generated for $i \leq m$. Recall that a $\mathbb{Z}G$-module can be regarded as an Abelian group with a G-action by automorphisms. Thus, $R \otimes_{\mathbb{Z}} P_i$ is a RG-module. Using the criterion that a module is projective if and only if it is a direct summand of a free module, it is easy to see that $R \otimes_{\mathbb{Z}} P_i$ is a projective RG-module. Clearly $R \otimes_{\mathbb{Z}} P_i$ is finitely generated if P_i is finitely generated.

It remains to show that

$$\cdots \to R \otimes_{\mathbb{Z}} P_i \to \cdots R \otimes_{\mathbb{Z}} P_1 \to R \otimes_{\mathbb{Z}} P_0 \twoheadrightarrow R \otimes_{\mathbb{Z}} \mathbb{Z} = R$$

is a resolution, i.e., an exact sequence. Exactness does not depend on the G-action and projective RG-modules are projective R-modules, Hence it suffices to prove the follwing claim:

> Let
> $$\cdots \to P_i \to \cdots P_1 \to P_0 \twoheadrightarrow \mathbb{Z}$$
> be an exact sequence of projective s-modules. Then the chain complex
> $$\mathrm{C}_*()(R) := \cdots \to R \otimes_{\mathbb{Z}} P_i \to \cdots R \otimes_{\mathbb{Z}} P_1 \to R \otimes_{\mathbb{Z}} P_0 \twoheadrightarrow R \otimes_{\mathbb{Z}} \mathbb{Z} = R$$
> is exact.

Note that the homology of $\mathrm{C}_*()(R)$ computes the homology of the trivial group with coefficients in R. This homology, however, is also clearly computed by

$$\cdots \to 0 \to R \twoheadrightarrow R$$

which has trivial homology.

Appendix G

References

[Alon94] *J. Alonso:* Finiteness conditions on groups and quasi-isometries. Journal of Pure and Applied Algebra 95 (1994) p. 121 – 129

[Alpe91] *R. Alperin:* Arboreal Group Theory. MSRI Publications 19 (Springer, 1991)

[BGS01] *L. Bartholdi, R.I. Grigorchuk, Z. Šunik´:* Branch Groups. preprint (2001)

[BeBr97] *M. Bestvina, N. Brady:* Morse theory and finiteness properties of groups. Inventiones mathematicae 129 (1997) p. 445 – 470

[BFH00a] *M. Bestvina, M. Feighn, M. Handel:* The Tits Alternative for $\mathrm{Out}(F_n)$ I: Dynamics of exponentially growing automorphisms. Annals of Mathematics 151 (2000) p. 517 – 623

[BFH00b] *M. Bestvina, M. Feighn, M. Handel:* The Tits Alternative for $\mathrm{Out}(F_n)$ II: A Kolchin Theorem. preprint

[BrVo01] *M. Bridson, K. Vogtmann:* The symmetries of Outer space. Duke Mathematical Journal 106 (2001), no. 2 p. 391 – 409

[Bro87a] *K.S. Brown:* Finiteness Properties of Groups. Journal of Pure and Applied Algebra 44 (1987) p. 45 – 75

[BrGe84] *K.S. Brown, R. Geoghegan:* An infinite-dimensional torsion-free FP_∞ group. Inventiones mathematicae 77 (1984) p. 367 – 381

[BBW79] *A.M. Brunner, R.G. Burns, J. Wiegold:* On the Number of Quotients, of One Kind or Another, of the Modular Group. Mathematical Scientist 4 (1979) p. 93 – 98

[CFP96] *J.W. Cannon, W.J. Floyd, W.R. Parry:* Introductory Notes to Richard Thompson's Groups. L'Enseignement Mathématique 42 (1996) p. 215 – 256

[Chou80] *C. Chou:* Elementary amenable groups. Illinois Journal of Mathematics 24 (1980) p. 396 – 407

[CuMo87] *M. Culler, K. Vogtmann:* Group Actions on $\mathbb{R}$-Trees. Proceedings of the London Mathematical Society 35 (1987) p. 571 – 604

[CuVo86] *M. Culler, K. Vogtmann:* Moduli of graphs and automorphisms of free groups. Inventiones mathematicae 84 (1986) p. 91 – 119

[DiDu89] *W. Dicks, M.J. Dunwoods:* Groups acting on graphs. Cambridge studies in advanced mathematics 17 (Cambridge University Press, 1989)

[EiGa57] *S. Eilenberg, T. Ganea:* On the Lusternik-Schnirelmann category of abstract groups. Annals of Mathematics 65 (1957) p. 517 – 518

[Epst62] *D.B.A. Epstein:* Ends. in [Fort62] p. 110 – 117

[FaNe62] *E. Fadell, L. Neuwirth:* Configuration Spaces. Mathematica Scandinavica 10 (1962) p. 111 – 118

[Fort62] *M.K. Fort:* Topology of 3-Manifolds. Proceedings of The University of Georgia Institute (1961), Prentice-Hall 1962

[GHV90] *E. Ghys, A. Haefliger, A. Verjovsky:* Groups from a geometric viewpoint. Proceedings of the workshop held in Trieste, March 26–April 6, 1990 p. World Sci. Publishing, 1991

[Grig80] *R.I. Grigorchuk:* On Burnside's Problem on periodic groups. Functional Analysis Applications 14 (1980) p. 41 – 43

[Gri98] *R.I. Grigorchuk:* An example of a finitely presented amenable group not belonging to the class EG. Math. USSR Sbornik 189 (1998) p. 75 – 95

[GrWi00] *R.I. Grigorchuk, J.S. Wilson:* The Conjugacy Problem for Certain Branch Groups. Proceedings of the Steklov Institute of Mathematics 231 (2000) p. 204 – 219

[Grom81] *M. Gromov:* Groups of Polynomial Growth and Expanding Maps. Publicationes Mathematicae IHES 53 (1981) p. 53 – 73

[Harp00] *P. de la Harpe:* Topics in Geometric Group Theory. University of Chicago Press (2000)

[HaVo96] *A. Hatcher, K. Vogtmann:* Isoperimetric Inequalities for Automorphism Groups of Free Groups. Pacific Journal of Mathematics 173 (1996) p. 425 – 441

[Hows54] *A.G. Howson:* On the intersection of finitely generated free groups. Journal of the London Mathematical Society 29 (1954) p. 428 – 434

[Inci97] *R. Incitti:* Regularities on the Cayley Graphs of Groups of Linear Growth. European Journal of Combinatorics 18 (1997) p. 175 – 178

[KM98a] *O. Kharlampovic, A. Myasnikov:* Tarski's problem about the elementary theory of free groups has a positive solution. ERA-AMS (1998) p. 101 – 108

[KM98b] *O. Kharlampovic, A. Myasnikov:* Irreducible affine varieties over a free group. I: Irreducibility of quadratic equations and Nullstellensatz. Journal of Algebra 200 (1998) p. 472 – 516

[KM98c] *O. Kharlampovic, A. Myasnikov:* Irreducible affine varieties over a free group. II: Systems in row-echelon form and description of residually free groups. Journal of Algebra 200,(1998) p. 517 – 570

[KM00a] *O. Kharlampovic, A. Myasnikov:* Implicit function theorem over free groups. preprint

[KM00b] *O. Kharlampovic, A. Myasnikov:* Equations over fully residually free groups. preprint

[KM00c] *O. Kharlampovic, A. Myasnikov:* Elementary theory of free nonabelian groups. preprint

[Kram95] *D. Krammer:* The conjugacy problem for Coxeter groups. PhD thesis (Utrecht 1995)

[LSch77] *R.C. Lyndon, P.E. Schupp:* Combinatorial Group Theory. Ergebnisse der Mathematik und ihrer Grenzgebiete 89 (Springer, 1977)

[Lys85] *I.G. Lysenok:* A System of Defining Relations for a Grigorchuk Group. Translations: Matematicheskie Zametki 38 (1985) p. 784 – 792

[Magn74] *W. Magnus:* Noneuclidean Tesselations and their Groups. Academic Press (New York, London 1974)

[Maka82] *G.S. Makanin:* Equations in a Free Group. Math. USSR Izvestiya 21 (1983) p. 483 – 546

[Meie82] *D. Meier:* Non-Hopfian Groups. Journal of the London Mathematical Society 26 (1982) p. 265 – 270

[Neum55] On the intersection of finitely generated free groups. *H. Neumann:* Publ. Math. Debrecen 4 (1955-56) p. 186 – 189 [Addendum: ibid. 5 (1957-58) p. 128]

[Quil78] *D. Quillen:* Homotopy Properties of the Poset of Nontrivial p-Subgroups of a Group. Advances in Mathematics 28 (1978) p. 101 – 128

[Razb84] *A.A. Razborov:* On Systems of Equations in a Free Group. Math. USSR Izvestiya 25 (1985) p. 115 – 162

[SBR02] *M.V. Sapir, J.C. Birget, E. Rips:* Isoperimetric and Isodiametric Functions of Groups. Annals of Mathematics 156 (2002) p. 345 – 466

[Se01a] *Z. Sela:* Diphantine Geometry over Groups I: Makanin-Razborov diagrams. preprint

[Se01b] *Z. Sela:* Diphantine Geometry over Groups II: Completions, closures and formal solutions. preprint

[Se01c] *Z. Sela:* Diphantine Geometry over Groups III: Rigid and solid solutions. preprint

[Se01d] *Z. Sela:* Diphantine Geometry over Groups IV: An iterative procedure for validation of a sentence. preprint

[Se01e] *Z. Sela:* Diphantine Geometry over Groups V: Quantifier elimination. preprint

[Se01f] *Z. Sela:* Diphantine Geometry over Groups VI: The elementary theory of a free group. preprint

[Shor90] *H. Short:* Quasiconvexity and a Theorem of Howson. in [GHV90] p. 168 – 176

[Skor??] *R. Skora:* Deformation of Length Functions in Groups. Preprint (????)

[Smal57] *S. Smale:* A Vietoris Mapping Theorem for Homotopy. Proceedings of the American Mathematical Society 8 (1957) p. 604 – 610

[Stal68] *J.R. Stallings:* On torsion-free groups with infinitely many ends. Annals of Mathematics 88 (1968) p. 312 – 334

[Stal88] *J. Stallings:* Foldings of G-trees. in [Alpe91] p. 355 – 368

[Tits72] *J. Tits:* Free Subgroups in Linear Groups. Journal of Algebra 20 (1972) p. 250 – 270

[Vogt00] *K. Vogtmann:* Automorphisms of Free Groups and Outer Space. To appear in the special issue of Geometriae Dedicata for the June, 2000 Haifa Conference on Geometric Group Theory