UNIVERSITY OF SOUTHAMPTON

Faculty of Social Sciences
Mathematical Sciences

Uncountably many quasi-isometry classes of groups of type FP via graphical small cancellation theory

by

Thomas M Brown

ORCiD: 0000-0001-8713-6223

A thesis for the degree of Doctor of Philosophy

August 2021

University of Southampton

Abstract

Faculty of Social Sciences
Mathematical Sciences

Doctor of Philosophy

Uncountably many quasi-isometry classes of groups of type FP via graphical small cancellation theory

by Thomas M Brown

This thesis presents a construction of a new class of groups that are type FP but are not finitely presentable. This is the first such construction that does not rely on Morse theory on cubical complexes and so reinforces the rift between the algebraic property and its geometric counterpart. The central tool used here is small cancellation theory which allows us a comparatively simple way to prove the above claim and also allows access to further results regarding these groups.

Contents

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. None of this work has been published before submission.

Signed:.. Date:..26 - 08 - 2021..

Acknowledgements

I owe a great amount of thanks to my supervisor Ian Leary for providing guidance, patience, mental stimulation and friendship. I feel extremely lucky to have played a part in this project with you.

Thank you also to Peter, Ashot, Armando, Nansen, Hector, Ana, Illaria, Motiejus, Hollis, Naomi, Vlad, Sam, Amina and Matthew for your company on our various mathematical escapades.

I wish to thank Fabio, Illaria, Abi, Xin, Burfitt, Ana, Hector, Simos, Ruth, Martina, James, Conrad, Motiejus and Hollis for inducting me into my PhD; Karl, Vlad, Naomi, Megan, Holly, Simon, Laurie, Pete, Aaron and Laura for sharing the journey with me; and Sam, George, Lily and Matt for helping to push me out the other side. Whether via a useful academic insight, a stress inducing game of mega-bounce (and its variants) or a stress relieving pub trip; each of you played a huge role in my experience.

Thank you to my parents: Mum for being positive, supporting and always there for a phone call (except when I **finally** got a job) and Dad for always encouraging and for driving me to achieve the best I possibly can. Finally, thank you to Leonie for providing a constant supply of belief, a necessary level of distraction, unbounded amounts of joy and humouring me when I presented you with yet **another** colourful diagram.

Chapter 1

Introduction

In this thesis we present a new construction of uncountably many groups of type FP.

A group is finitely presentable if there is some finite presentation describing it. A group is type FP if $\mathbb{Z}$ admits a finite resolution by finitely generated projective modules over the corresponding group ring. Until 1997, when Bestvina and Brady produced the first groundbreaking counterexamples (2), it was unknown whether a group of type FP could ever be not finitely presentable.

When L is an acyclic flag complex, the Bestvina-Brady group BB_L is defined to be the kernel of the map

$$\phi : A_\Gamma \to \mathbb{Z},$$

where Γ is the 1-skeleton of L, A_Γ is the corresponding Right-angled Artin group and ϕ maps each generator to $1 \in \mathbb{Z}$. This group is always type FP and, when L has non-trivial fundamental group, it is not finitely presentable.

In 2015 Leary pushed this result and produced uncountably many such counterexamples (18). This work was motivated by the observation that there is a presentation for BB_L such that, for each $n \in \mathbb{Z}$, there are a disjoint set of relators corresponding to the fundamental group of L. The underlying idea is that perhaps we could make a choice:

- for some $n \in \mathbb{Z}$, use relators corresponding to the fundamental group;
- for any other $n \in \mathbb{Z}$, use relators corresponding to attaching a cone on L.

Since the homology of L is the same as the homology of the cone on L (both trivial), the corresponding groups $G_L(S)$ are still type FP. They depend on some subset $S \subseteq \mathbb{Z}$ and $G_L(S)$ is finitely presentable if and only if S is finite.

The main tool utilised in both of these constructions is Morse theory on cubical complexes. The motivation for this thesis was to find similar examples of groups of type FP that are not finitely presentable without the involvement of these methods.

The groups discussed here have a similar flavour to those in (18). We use a combinatorial object, dubbed a *suitable complex*, along with an integer $k \geq 2$ to construct a family of groups

$$\{G(S) \mid S \subseteq \mathbb{N}_0\},$$

such that each $G(S)$ is type FP and $G(S)$ is finitely presentable if and only if S is finite. The underlying tool used here is graphical small cancellation theory which allows us to apply a similar trick to Leary, for $n \in S$ we have relations corresponding to the cone on K and for $n \notin S$ we have relations corresponding to the fundamental group of K.

In comparison to the constructions of Bestvina-Brady and Leary, there are heavy restrictions on the combinatorial object used as an input and a large portion of the project involved proving the existence of such an object. Although this means we have less freedom when constructing examples, it allows us to make stronger statements about the groups in general. We will give a presentation for $G(S)$ and show that, given some conditions on K and k, we are able to use van Kampen-reduced disk diagrams to prove the following:

- the set S is finite if and only if $G(S)$ is finitely presentable;
- for all $S \subseteq \mathbb{N}_0$, the cohomological dimension of $G(S)$ is 2 and $G(S)$ is type FP;
- when $S = \emptyset$, $|S| = 1$, and $S = \{0, 1\}$, the geometric dimension of $G(S)$ is 2;
- the set $\{G(S) \mid S \subseteq \mathbb{N}_0\}$ contains uncountably many quasi-isometry classes of groups;
- for all $S \subseteq \mathbb{N}_0$, the homological Dehn function of $G(S)$ is exponential;
- for all $S \subseteq \mathbb{N}_0$, $G(S)$ is acylindrically hyperbolic;
- when $S \subseteq \mathbb{N}_0$ is periodic, $G(S)$ is non-Hopfian;
- for all $S \subseteq \mathbb{N}_0$, the first sigma invariant $\Sigma^1(G(S))$ is empty.

This thesis is split into two parts. Chapter 2 deals with the existence of suitable complexes, providing some concrete examples, and presenting ideas on how to find more. Chapter 3 tackles the construction and properties of the groups $G(S)$.

Chapter 2

Suitable complexes

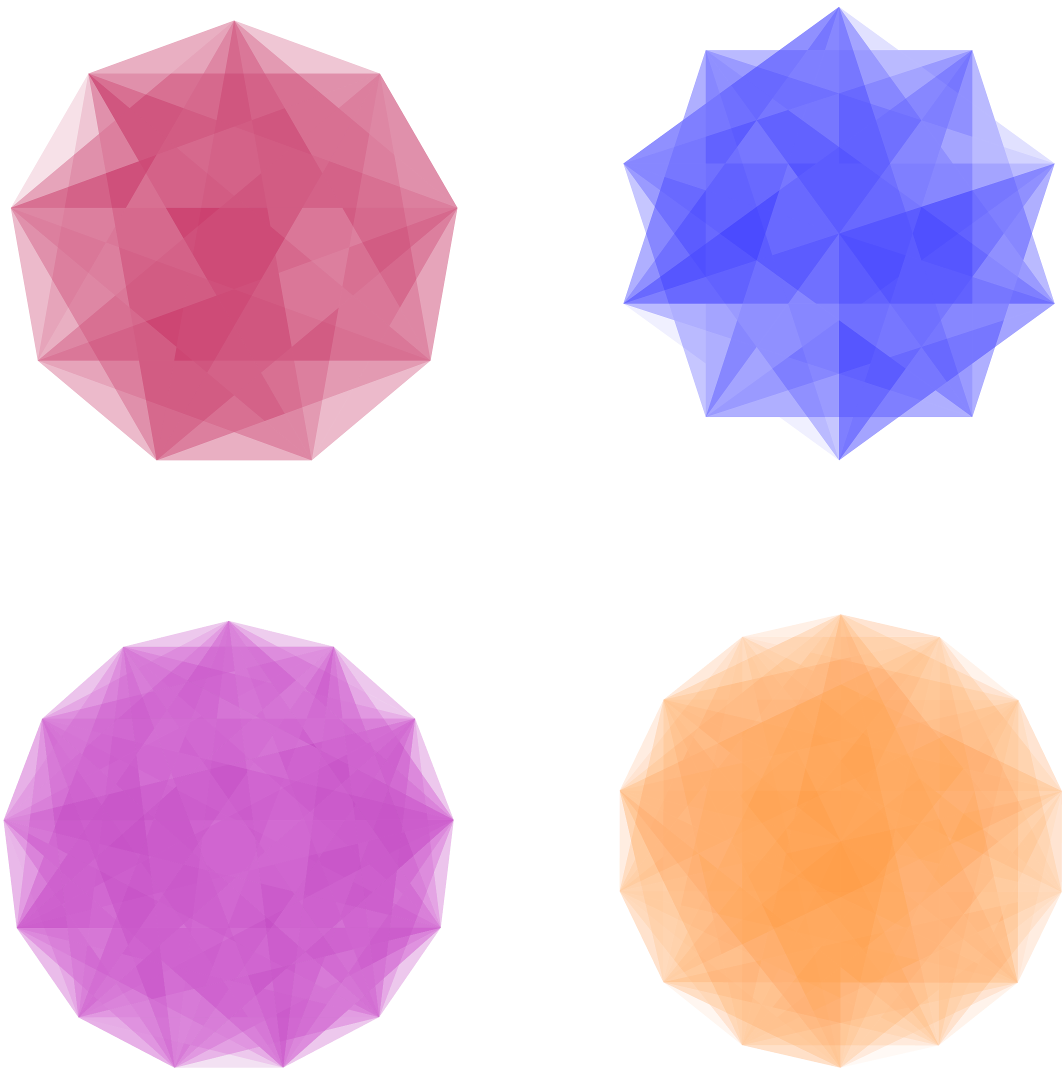

1 Definitions

In this section we will give the definition of a suitable complex, along with its cousins: pre-suitable complexes and very-suitable complexes, discussing the motivation behind the properties in each case.

We will begin with the definition of a suitable complex, as initially described in (8), but first we need to clarify some terms.

Definition 1.1. A 2-**complex** is a CW-complex with no cells of dimension $n > 2$.

Definition 1.2. A **simplicial graph** is a graph such that if there is an edge e between two vertices u and v then $u \neq v$ and e is the only such edge.

In other words we have no *self-loops* and no *multi-edges*.

Definition 1.3. The **girth** of a simplicial graph Γ is the length of the shortest loop in Γ. If there are no loops in Γ then we say that the girth of Γ is infinite.

The girth of a simplicial graph is always at least 3.

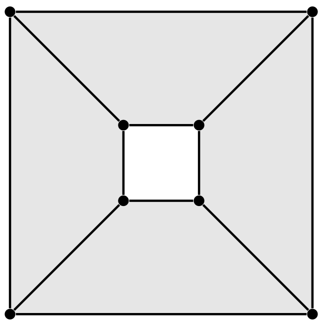

FIGURE 2.1: A 2-complex on a simplicial graph with girth 4.

Definition 1.4. A space is **acyclic** if it has the same homology as a point.

Any contractible space is certainly acyclic. Suitable complexes, as we will see, are more interesting examples.

Definition 1.5. A **suitable complex** K is a finite 2-complex with the following properties.

1. The 1-skeleton K^1 is a simplicial graph;
2. The attaching map for each polygon P of K is an embedding of the boundary circle ∂P into K^1;
3. The girth g of the graph K^1 is at least 13;

4. The perimeter l_P of each polygon P satisfies $l_P > 2g$;

5. The polygons of K satisfy a $C'(\frac{1}{6})$ condition: for each pair $P \neq Q$ of polygons of K, each component of the intersection $\partial P \cap \partial Q$ of their boundaries contains strictly fewer than $\min\{l_P/6, l_Q/6\}$ edges;

6. K is acyclic.

We call the 0-cells **vertices**, 1-cells **edges** and 2-cells **polygons**. If P is a polygon then we let ∂P denote the unique loop that makes up the boundary of P. We let l_P be the length of this boundary loop.

Remark. Properties 1, 2, 3 and 5 are required to allow the graph $\Gamma(S)$ as defined in Definition 2.21 to satisfy $C'(\frac{1}{6})$. Property 6 is the key tool in showing that the kernel $K_{S,T}$ from Theorem 3.30 is acyclic and property 4 tells us that the fundamental group of K is non-trivial and so each $K_{S,T}$ is non-trivial.

Prior to this project, no suitable complexes had been constructed. Each of the constructions that will be discussed here will involve the construction of a pre-suitable complex L.

Definition 1.6. A **pre-suitable complex** L is a finite 2-complex satisfying every property of being suitable, except we relax property 3 as follows.

3*. The girth g of the graph L^1 is at least 3.

The reason we are able to do this can be summarised as follows.

Lemma 1.7. *Let L be a pre-suitable complex. There is a map*

$$\phi : E(L) \to \mathbb{N}$$

such that subdividing each edge a in L into $\phi(a)$ edges will result in a suitable complex.

Proof. Let ϕ map every edge to some $d \geq 5$ and let K be the resulting 2-complex. One can check that each of the properties of being pre-suitable is unaffected except 3*. The girth of K^1 is at least 5 times that of L^1 and so K is suitable. ☐

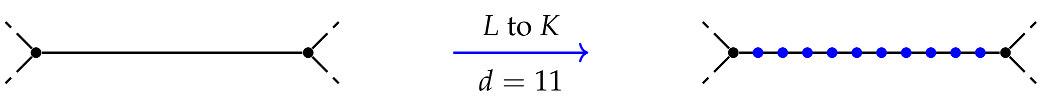

FIGURE 2.2: We replace each edge in a pre-suitable L by $d \geq 5$ edges to construct a suitable complex K.

One could consider more exotic maps that the one above but should be careful to check whether the resulting complex satisfies properties 4 and 5. For example, if an edge shared by two polygons is mapped to some integer sufficiently larger than any other edge then the resulting complex will fail to be $C'(\frac{1}{6})$.

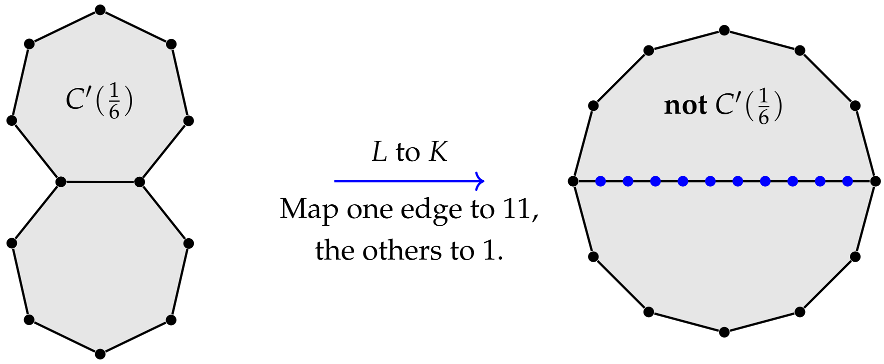

FIGURE 2.3: Carelessly choosing ϕ can cause problems.

For many of our later results we are required to consider a strengthening of being suitable in order to make claims via structure within disk diagrams.

Definition 1.8. Let L be a pre-suitable complex and let d be some integer with $d \geq 11$. Let K be the 2-complex constructed in Lemma 1.7 defined by the map

$$\phi : a \mapsto d$$

for each edge a in L. We say K is **very-suitable**, d is the **subdivision constant** and the set L^0 is the **pre-subdivision vertex set**.

Corollary 1.9. *A very-suitable complex is a suitable complex.*

Proof. This follows immediately from Lemma 1.7 since $11 \geq 5$. $\square$

Together these allow us to focus our search on pre-suitable complexes as they are combinatorially the most simple.

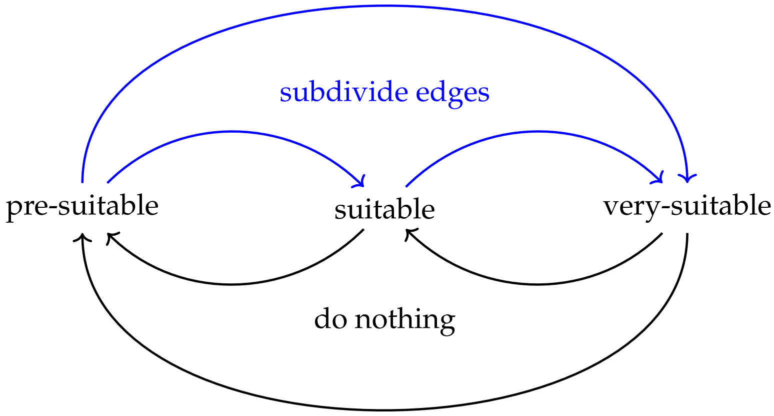

FIGURE 2.4: We can easily move between various strengths of being suitable.

Currently no suitable or very-suitable complexes have been constructed without first constructing a pre-suitable complex of girth 3.

2 Constructions

In this section we discuss some successful constructions of pre-suitable complexes. We progress in chronological order and attempt to give some insight into method and motivation.

2.1 Properties of pre-suitable complexes on a complete graph

Given we are able to consider 1-skeleta of girth 3, a natural place to begin is with complete graphs.

Proposition 2.1. *If L is a pre-suitable complex and L^1 is a complete graph with $n = |L^0|$ vertices then L contains*

$$\frac{(n-1)(n-2)}{2} = 1 + \frac{n(n-3)}{2}$$

polygons.

Proof. Since the Euler characteristic of an acyclic complex is 1 this follows immediately. □

Proposition 2.2. *If L is a pre-suitable complex and L^1 is a complete graph then $|L^0| \geq 7$.*

Proof. For the polygons of L to satisfy $C'(\frac{1}{6})$ they must have at least 7 edges and since they must each embed into L^1 there must be at least 7 vertices. □

This can be strengthened to require that $|L^0|$ is at least 9 via a brute force computer search.

Proposition 2.3. *If L is a pre-suitable complex and L^1 is a complete graph then there is a polygon P in L such that l_P is odd.*

Proof. Suppose every polygon in L has even length. Then every loop in L^1 that represents the trivial element in $H_1(L)$ has even length since it can be decomposed as finite sum of polygon boundaries. Since L^1 is complete, there is a loop of length 3 which must be non-trivial in $H_1(L)$, contradicting the property that L is acyclic. □

We will use K_n to denote the complete graph on n vertices.

2.2 Construction via rotating families

The first attempt was motivated by the following construction given by taking the 2-skeleton of the Poincaré homology sphere.

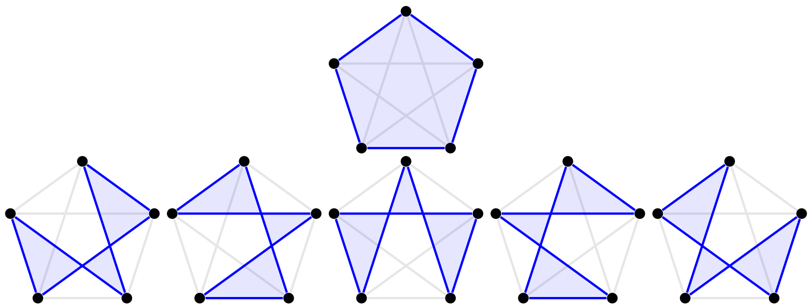

FIGURE 2.5: We present the boundary cycles of each of the 6 polygons.

Attaching the 5-gons described above to a K_5 defines a 2-complex with all the properties of being suitable except 4 and 5.

The polygons above are suggestively presented to draw attention to the fact that we have a single polygon with boundary traversing the 'outside' of the graph and 5 polygons which can be viewed as rotations of some initial polygon. We make this rigorous via the following definition.

Definition 2.4. Consider the complete graph K_n as having its vertices drawn at the n unique points on a circle with angles some multiple of $2\pi/n$.

There is a unique loop of length n through K_n that can be seen as traversing the circle. We call a polygon attached via this loop the **outside polygon**. Given an arbitrary polygon attached to K_n, we say the corresponding **rotating family** of polygons is exactly the set of polygons such that rotating the complex by $2\pi/n$ is an automorphism of 2-complexes.

A rotating family will always contain $m|n$ polygons. We will be interested only in rotating families that contain exactly n polygons.

We will consider larger K_n and longer polygons. Proposition 2.2 tells us we need at least 7 vertices and, since the polygons all have the same length as the number of vertices and at least one of them must have odd length by Proposition 2.3, we must consider complete graphs on $n \geq 7$ vertices where n is odd.

We also want our polygons want to fit the form in Figure 2.5. Since we require

$$1 + n\frac{(n-3)}{2}$$

polygons, we are looking for $\frac{n-3}{2}$ rotating families, each containing n polygons, such that the resulting complex is pre-suitable.

We will use the following algorithm.

Algorithm 2.5. Start with some K_n and search as follows.

1. Generate all possible rotating families, ruling out any families with polygons sharing more than 1 edge in a row. We also rule out any with a polygon sharing more than one edge in a row with the outside polygon.
2. In those that remain, search for $\frac{n-3}{2}$ rotating families that pairwise do not contain polygons sharing more than 1 edge in a row.
3. For each of these combinations, compute the homology. If any are acyclic then we are done.

We observe that for $n > 12$ we are checking for a stronger condition than $C'(\frac{1}{6})$ but for smaller n this is equivalent.

Example 2.6. We begin with $n = 7$ so we need 2 rotating families. There are only 6 rotating families that make it to Step 2.

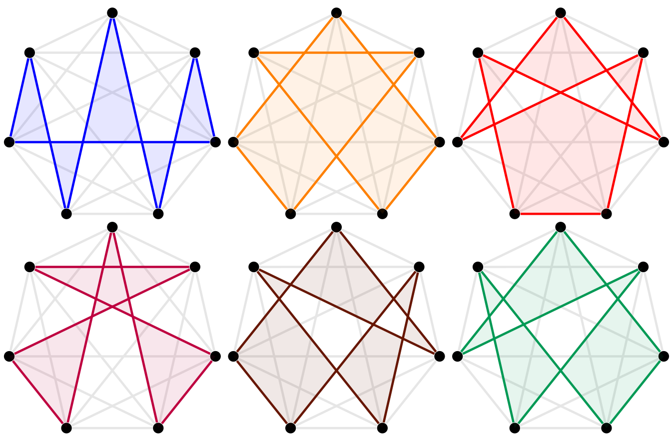

FIGURE 2.6: Each polygon above defines a rotating family that is $C'(\frac{1}{6})$ with itself and the outside path.

We observe that each pair of these fails to be $C'(\frac{1}{6})$. Moreover, one can check that no pair of these can be combined to make an acyclic complex; the first homology is always non-trivial.

Example 2.7. We next consider $n = 9$ and are looking for triples of rotating families. This time there are 522 families that make it to Step 2 and of these there are many that are all pairwise $C'(\frac{1}{6})$. Computing the homology however we do not find any that are acyclic.

One of the examples stands out as having a particularly 'nice' first homology group:

$$\mathbb{Z}_{17}^{7}.$$

It can be described by the following rotating families of polygons.

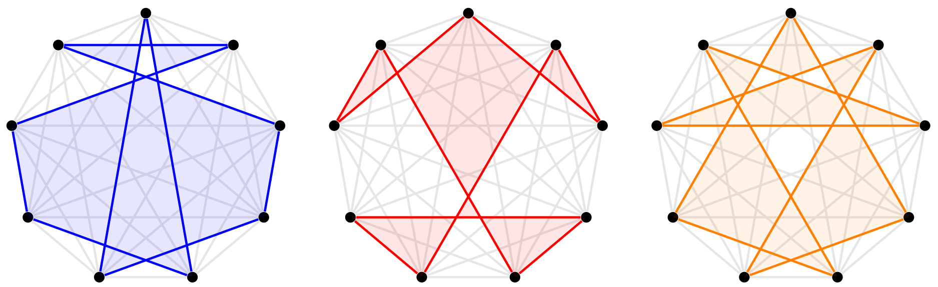

FIGURE 2.7: These polygons define a complex with a significant first homology group.

This complex interestingly reappears in Section 2.3.

We observe that each of these polygons is reflectively symmetrical and since we are aiming for a particularly 'symmetrical' first homology group (the trivial group), we alter our algorithm to simplify the search space.

Algorithm 2.8. Start with some K_n and search as follows.

1. Generate all possible rotating families **that have a reflective symmetry**, ruling out any families with polygons sharing more than 1 edge in a row. We also rule out any with a polygon sharing more than one edge in a row with the outside polygon.
2. In those that remain, search for $\frac{n-3}{2}$ rotating families that pairwise do not contain polygons sharing more than 1 edge in a row.
3. For each of these combinations, compute the homology. If any are acyclic then we are done.

For $n = 11$ we also fail to find any pre-suitable complexes.

Example 2.9. When $n = 13$ we find the following set of 5 rotating families that are all pairwise $C'(\frac{1}{6})$.

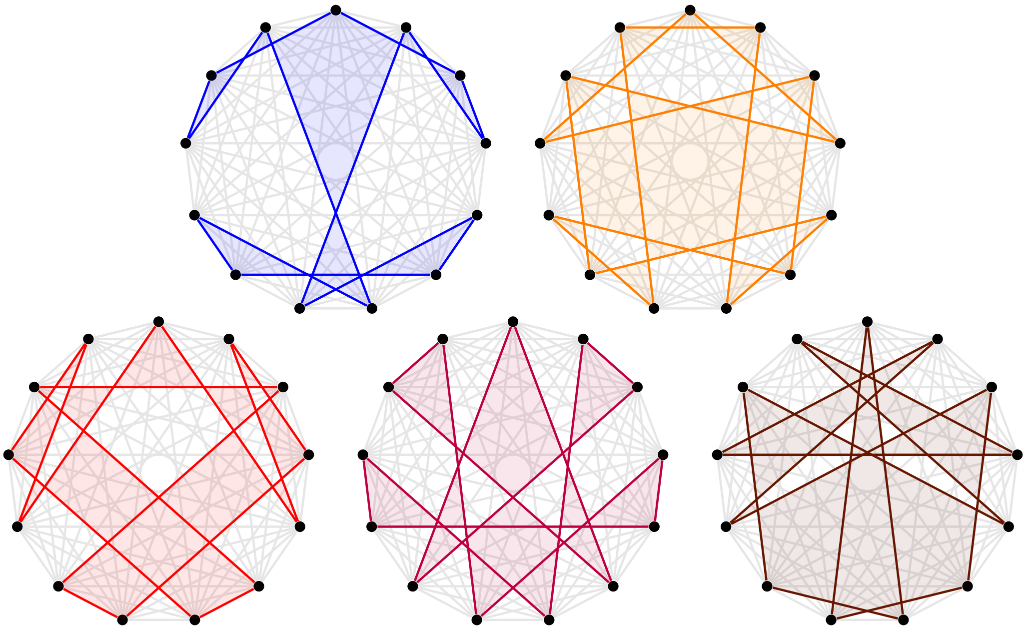

FIGURE 2.8: These paths define a complex with cyclic first homology group.

The resulting complex has first homology group $\mathbb{Z}_{281}$. This is the first time so far a cyclic group has appeared. Moreover, this is the first time where we are permitted to have a pair of polygons sharing two edges in a row whilst remaining $C'(\frac{1}{6})$.

We utilise this freedom by altering the outside polygon. Replacing it by any of the following defines a pre-suitable complex.

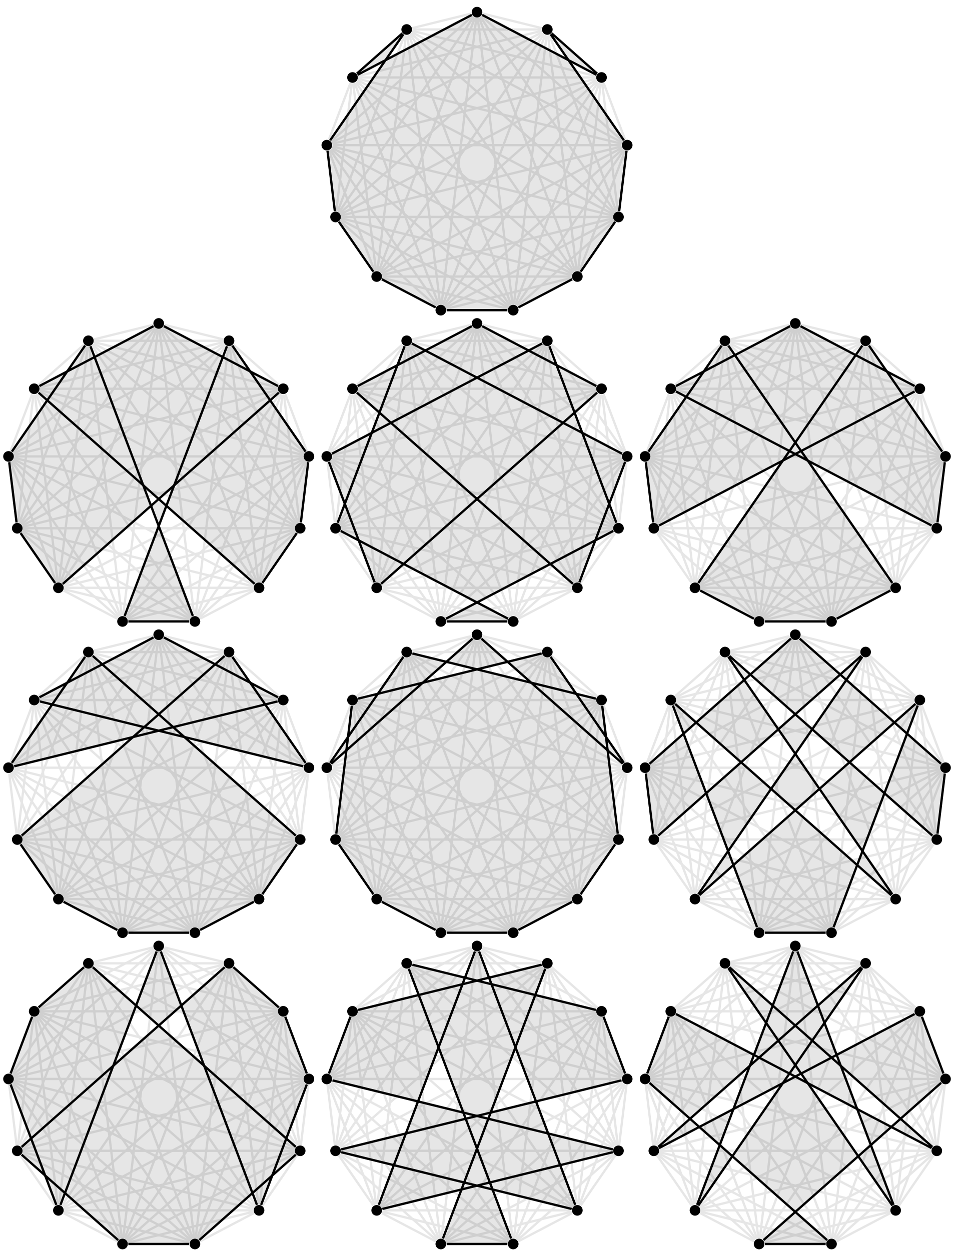

FIGURE 2.9: Replacing the outside path with any of these results in a pre-suitable complex.

This complex is very large with 13 vertices, 78 edges and 66 polygons. Particularly when we start to subdivide its edges to make it suitable and very-suitable. However, it allowed the project to take off as it was the first evidence that such an object did exist.

This idea should generalise to other and possibly all odd values of $n \geq 13$ but no evidence has been found to confirm this as yet. We rephrase this as an open question.

Open Question 2.10. For what other odd $n \geq 13$ does this construction work?

2.3 Construction via projective general linear groups

The second attempt was motivated by (1) where group actions on finite 2-dimensional acyclic simplicial complexes are studied, we note that these complexes cannot be directly pre-suitable since each 2-cell has boundary length 3. Interestingly the Poincaré example reappears in this section; we will first describe a construction and then discuss how the complex given in Figure 2.5 fits in.

We will let g^G represent the unique conjugacy class of some group G that contains g.

Algorithm 2.11. Choose some prime power q and let g^G be some conjugacy class in the projective general linear group $PGL_2(q)$. Let the complex $L := L(q, g)$ be defined by Step 2 and search for a pre-suitable complex as follows:

1. Let L^1 be the complete graph on $q + 1$ vertices.
2. The finite group $PGL_2(q)$ acts 3-transitively on L^1, so to each element $h \in g^G$ we can assign a set of loops in L^1 by considering the loops mapped out by its action. If there are any loops of length at least 3 that are not already the boundary of a polygon, then attach a polygon to each of these loops.
3. If the resulting complex satisfies $C'(\frac{1}{6})$ and has trivial first homology then search for a pre-suitable subcomplex.

This method is a different flavour to the previous attempt since here we add (potentially) too many polygons initially and then correct our overkill. In order to shed more light on this idea we will consider some small examples.

Example 2.12. Let $q = 2$ and consider the action of $PGL_2(q)$ on K_3.

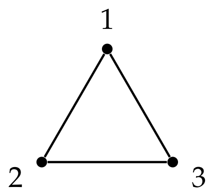

FIGURE 2.10: $PGL_2(2)$ acts 3-transitively on K_3.

This group has three conjugacy classes:

- $L(2, ())$ is the graph above since no polygons are defined by the trivial element.
- $L(2, (1, 2))$ is also the graph above since each element in the conjugacy class has order 2 and so we define no loops.

- $L(2, (1,2,3))$ is acyclic (in fact contractible) as the conjugacy class is $(1,2,3)^G = \{(1,2,3),(1,3,2)\}$ and so we attach a single polygon around the unique path of length 3. Although the resulting complex is acyclic it is clearly not pre-suitable.

Example 2.13. Let $q = 3$ and consider the action of $PGL_2(q)$ on K_4.

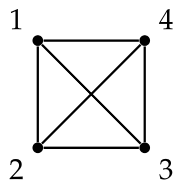

FIGURE 2.11: The complete graph on 4 vertices.

This group has five conjugacy classes, two of which result in complexes with at least one polygon:

- $L(3, (1,2,3))$ is the 2-skeleton of a 4-simplex via 4 polygons of boundary length 3.
- $L(3, (1,2,3,4))$ is gives a complex with three squares attached to the 1-skeleton via the loops are shown below.

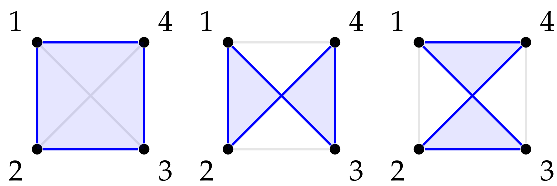

FIGURE 2.12: A 2-complex with 3 squares attached to a K_4.

The first homology group of this complex is $\mathbb{Z}_2$, this is not a surprise following Proposition 2.3. In fact this space is exactly the projective plane defined by identifying antipodal points of the cube.

We now discuss how this construction relates to Figure 2.5.

Example 2.14. Consider the action of $PGL_2(4)$ on K_5 and consider the conjugacy class $C = (1,2,3,4,5)^G$. This defines 6 polygons of length 5 and so $L(4, (1,2,3,4,5))$ is exactly the 2-complex described in Figure 2.5.

This is the first real hope that for some large q we may find a pre-suitable complex. Previous to this, any complex constructed with trivial first homology has been made of triangles and, for all q, $PGL_2(q)$ contains a unique conjugacy class of elements of order 3. The resulting complex contains every possible triangle and will therefore always have trivial first homology.

The other vital idea used here is summarised as follows.

Proposition 2.15. *Let P and Q be distinct polygons in $L(q,g)$. Each connected component of $\partial P \cap \partial Q$ contains at most one edge.*

Proof. Since $PGL_2(q)$ acts 3-transitively on K_{q+1}, it acts transitively on the set of ordered triples of its vertices.

We say a polygon P contains an ordered triple of vertices (u, v, w) if it contains the edges $\{u, v\}$ and $\{v, w\}$. A polygon of length d contains $2d$ triples of vertices. There are $(q+1)q(q-1)$ distinct ordered triples and $(q+1)q(q-1)/2d$ polygons in $L(q,g)$.

So for each ordered triple (u, v, w) there is exactly one polygon containing it. □

This means, so long as we are attaching only polygons of length at least 7, we will always result in a complex satisfying $C'(\frac{1}{6})$. This small cancellation condition was a major obstacle in the previous section.

The following table records the homology of each $L(q,g)$ for $q \leq 16$ and g of order at least 7 (in order to try and find $C'(\frac{1}{6})$ complexes) to give a flavour of the complexes we are constructing.

q	g	$H_1(L(q,g))$
7	$(1,5,3,7,6,8,2)$	$\mathbb{Z}^6$
7	$(1,7,5,4,6,3,8,2)$	$\mathbb{Z}_2 \times \mathbb{Z}_4^8$
7	$(1,4,8,7,6,2,5,3)$	$\mathbb{Z}_2^3 \times \mathbb{Z}_6^6$
8	$(1,8,9,4,5,3,2)$	0
8	$(1,9,5,2,8,4,3)$	0
8	$(1,4,2,9,3,8,5)$	0
8	$(1,6,4,8,5,9,7,3,2)$	$\mathbb{Z}_{17}^7$
8	$(1,4,5,7,2,6,8,9,3)$	$\mathbb{Z}_{17}^7$
8	$(1,5,2,8,3,4,7,6,9)$	$\mathbb{Z}_{17}^7$
9	$(1,5,6,9,4,7,8,2)$	$\mathbb{Z}_2^{12}$
9	$(1,9,8,5,4,2,6,7)$	$\mathbb{Z}_2^{12}$
9	$(1,6,7,3,9,10,5,4,8,2)$	$\mathbb{Z}_2^{19} \times \mathbb{Z}_4$
9	$(1,3,5,2,7,10,8,6,9,4)$	$\mathbb{Z}_2^{19} \times \mathbb{Z}_4$
11	$(1,3,10,5,7,6,12,4,8,2)$	$\mathbb{Z}_2^{15} \times \mathbb{Z}_{10}^9 \times \mathbb{Z}_{20}$
11	$(1,5,12,2,10,6,8,3,7,4)$	$\mathbb{Z}_2^{23} \times \mathbb{Z}_{10} \times \mathbb{Z}^{11}$
11	$(1,12,6,4,9,11,3,5,10,8,2)$	$\mathbb{Z}^{20}$
11	$(1,9,5,4,10,7,12,3,6,11,8,2)$	$\mathbb{Z}_2^{14} \times \mathbb{Z}_6 \times \mathbb{Z}_{18}^{10}$
11	$(1,7,8,4,6,9,12,2,10,11,5,3)$	$\mathbb{Z}_2^{25} \times \mathbb{Z}_8^{10}$
13	$(1,6,4,3,5,11,14)(2,8,10,9,7,12,13)$	0
13	$(1,3,14,4,11,6,5)(2,9,13,10,12,8,7)$	0
13	$(1,5,12,6,13,3,2)(4,7,9,11,14,10,8)$	0
13	$(1,13,9,4,6,5,8,11,10,12,7,3)$	$\mathbb{Z}_2^{16} \times \mathbb{Z}_4^2 \times \mathbb{Z}_{28}^{12}$
13	$(1,5,7,4,10,13,8,3,6,12,9,11)$	$\mathbb{Z}_2^2 \times \mathbb{Z}_4^{28}$
13	$(1,9,12,8,11,5,4,6,13,10,7,14,2)$	$\mathbb{Z}^{36}$
13	$(1,8,6,10,4,9,3,7,5,12,11,13,14,2)$	$\mathbb{Z}_2^{23} \times \mathbb{Z}_{26}^5 \times \mathbb{Z}^{14}$
13	$(1,10,3,12,14,8,4,7,11,2,6,9,5,13)$	$\mathbb{Z}_2^{30} \times \mathbb{Z}_{58}^{11} \times \mathbb{Z}_{116}$
13	$(1,9,11,8,3,13,6,7,14,10,5,2,4,12)$	$\mathbb{Z}_2^{14} \times \mathbb{Z}_4^{27} \times \mathbb{Z}_8$
16	$(1,12,6,5,15,4,13,16,14,10,8,11,3,9,2)$	$\mathbb{Z}_{239}^{15}$
16	$(1,6,15,13,14,8,3,2,12,5,4,16,10,11,9)$	$\mathbb{Z}_{239}^{15}$
16	$(1,15,14,3,12,4,10,9,6,13,8,2,5,16,11)$	$\mathbb{Z}_{239}^{15}$
16	$(1,16,2,13,9,4,3,15,11,5,8,6,10,12,14)$	$\mathbb{Z}_{239}^{15}$
16	$(1,14,11,13,15,12,8,7,17,5,3,10,16,6,4,9,2)$	$\mathbb{Z}_{1087}^{15}$
16	$(1,11,15,8,17,3,16,4,2,14,13,12,7,5,10,6,9)$	$\mathbb{Z}_{1087}^{15}$
16	$(1,13,8,5,16,9,14,15,7,3,6,2,11,12,17,10,4)$	$\mathbb{Z}_{239}^{15}$
16	$(1,15,17,16,2,13,7,10,9,11,8,3,4,14,12,5,6)$	$\mathbb{Z}_{1087}^{15}$
16	$(1,12,3,9,13,17,6,14,8,10,2,15,5,4,11,7,16)$	$\mathbb{Z}_{239}^{15}$
16	$(1,8,16,14,7,6,11,17,4,13,5,9,15,3,2,12,10)$	$\mathbb{Z}_{239}^{15}$
16	$(1,7,4,15,10,14,17,9,12,16,11,5,2,8,6,13,3)$	$\mathbb{Z}_{239}^{15}$
16	$(1,17,2,7,9,8,4,12,6,15,16,13,10,11,3,14,5)$	$\mathbb{Z}_{1087}^{15}$

FIGURE 2.13: A summary of attempts to find a 2-complex containing a pre-suitable subcomplex.

There are several rows of interest in this table. The first are denoted in light blue above, they highlight the fact that we have recovered the same homology as the complex constructed in Example 2.7. Moreover, one can check that each of these is exactly this complex.

The second, and most important, rows of interest are highlighted in darker blue. These complexes have trivial first homology and so there is a possibility that they contain a pre-suitable subcomplex. One can check that in each case that there are too many polygons to be pre-suitable already. It remains to attempt to remove polygons.

Example 2.16. We begin with the case $q = 8$ and $g = (1, 8, 9, 4, 5, 3, 2)$. The conjugacy class g^G contains 72 elements but is closed under taking inverses, so the complex $L(q, g)$ contains 36 polygons of length 7. In order to be acyclic we need remove 8 polygons as

$$\frac{q(q-1)}{2} = 28.$$

In this case there is a surprisingly simple trick to find a pre-suitable complex. Choose any vertex v in $L(q, g)$ and remove any polygon that is not incident to this vertex. This leaves 28 polygons and results in a pre-suitable complex.

This pre-suitable complex is much smaller than that described in Example 2.9 with 9 vertices, 36 edges and 28 polygons. This is the pre-suitable complex with the smallest number of vertices known to date.

This trick works for each of the conjugacy classes of elements of order 7 as highlighted. Moreover each of the $L(q, g)$ are the same complex up to relabelling the vertices. This follows from the fact that $Aut(PGL_2(8))$ acts transitively on these conjugacy classes.

The trick does not work for the cases when $q = 13$.

Example 2.17. Let $q = 13$ and $g = (1, 6, 4, 3, 5, 11, 14)(2, 8, 10, 9, 7, 12, 13)$. Here g^G contains 156 elements, is closed under taking inverses and each element defines two polygons. So, $L(q, g)$ has 156 distinct polygons. In order to be pre-suitable we are required to remove 78 of these.

The algorithm described below succeeds in finding such a complex.

We wish to explore the space of all subcomplexes of $L(q, g)$ with the same 1-skeleton. One can consider a natural graph structure on this space where we draw an edge between two subcomplexes if we can obtain one from the other by removing a polygon. We also arrange the space by considering a height function based on the number of polygons.

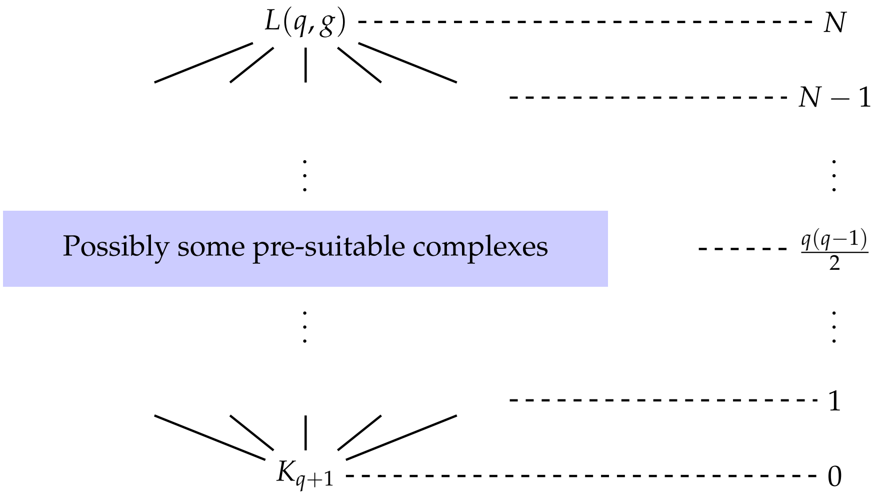

FIGURE 2.14: The space of subcomplexes of $L(q,g)$ with K_{q+1} as their 1-skeleton.

This space contains 2^N complexes, where N is the number of polygons in $L(q,g)$, which is far to large to brute force search when q is large.

A natural idea to search this space effectively is given below.

Algorithm 2.18. Let $L_N = L(q,g)$ and construct a sequence of complexes $L_N, L_{N-1}, ..., L_{\frac{q(q-1)}{2}}$ as follows:

- **Step** 1**:** If there is a subcomplex at level $N-1$ with trivial first homology that we have not already visited then set L_{N-1} be this complex and continue.
- **Step** j**:**
  - **IF** there is a subcomplex at level $N-j$ that, in our search space, is connected to $L_{N-(j-1)}$ via an edge, has trivial first homology and we have not visited before then set L_{N-j} and continue;
  - **ELSE** record that we have visited $L_{N-(j-1)}$ and return to **Step** $j-1$.
- **Step** $\frac{q(q-1)}{2}$**:** Terminate the algorithm and return $L = L_{\frac{q(q-1)}{2}}$.

If there is a pre-suitable complex in the search space then this algorithm will find it since, if there is a complex with non-trivial first homology then all complexes connected 'below' it will also have non-trivial homology because you cannot reduce homology by removing polygons.

This algorithm has the form of a depth first search and works fast enough to find pre-suitable complexes for each of the $L(13,g)$ above (where g has order 7) without having to search too much of the space.

In Appendix A we present MAGMA (4) code for construction $L(q,g)$ and for searching this space.

As things stand we have found pre-suitable complexes of three forms:

- 13-gons on a K_{13};
- 7-gons on a K_9;
- 7-gons on a K_{14}.

Since in each of the K_{13} cases involved at least one pair of polygons that shared two edges. We do not have any pre-suitable complexes that satisfy $C'(\frac{1}{p})$ for any $p > 6$. Although we do not have any initial applications of such a strengthening it would be interesting if there are no such examples as it would show that we are working within a tight band of possibility.

Open Question 2.19. Is there a pre-suitable complex that satisfies $C'(\frac{1}{p})$ for any integer $p > 6$?

In an attempt to answer this question we can extend the table in Figure 2.13, focusing only on the cases where the homology is trivial.

q	Polygon size	Number of polygons in $L(q,g)$	Number of polygons needed
8	7	36	28
8	7	36	28
8	7	36	28
13	7	156	78
13	7	156	78
13	7	156	78
17	9	272	136
19	9	380	171
19	9	380	171
25	13	600	300
25	13	600	300
25	13	600	300
25	13	600	300
27	7	1404	351
27	7	1404	351
27	7	1404	351
29	7	1740	406
29	7	1740	406
29	7	1740	406
32	11	1488	496
32	11	1488	496
32	11	1488	496
32	11	1488	496
32	11	1488	496
37	9	2812	666
37	9	2812	666
37	9	2812	666
41	7	4920	820
41	7	4920	820
43	7	5676	903
43	7	5676	903
43	7	5676	903
43	11	3612	903
43	11	3612	903
43	11	3612	903
43	11	3612	903
43	11	3612	903
43	21	1892	903

FIGURE 2.15: We record the occurrences of $L(q,g)$ with trivial first homology group. To save space we just give the order of the element g that defines the complex, observing that this does not mean that all conjugacy classes containing elements of this order give a trivial $H_1(L(q,g))$. We also draw attention to how many polygons we are required to cut down to in order to find a pre-suitable complex.

We observe many cases where we are attaching polygons with lengths longer than 7. These are candidates for containing a pre-suitable subcomplexes with a stronger small cancellation conditions but none have been found yet.

This table opens up more questions.

Open Question 2.20. Are there infinitely many values of q and g for which $H_1(L(q,g))$ is trivial?

There are certainly values of q for which there is no such g; for example, $q = 11$. However, as q increases there are more choices for g and the table above seems to suggest that this correlates to more occurrences of $H_1(L(q,g)) = 0$. As q grows large however the computations take much longer.

The next question is more approachable.

Open Question 2.21. Does the complex $L(q,g)$ with trivial $H_1(L(q,g))$ always contain a pre-suitable subcomplex?

This would be easiest to approach negatively and the natural first place to look for a non-existence would be within the $L(17,g)$ as highlighted above.

2.4 Minimality of pre-suitable complexes

As mentioned above, one can check by brute force that there are no pre-suitable complexes with 1-skeleton K_7 or K_8. So, on a complete graph, the pre-suitable complex from Example 2.16 is minimal in this sense.

The groups we construct from our (pre-)suitable complexes have presentations whose size depends on the number of edges and polygons. So if we can minimalise these at the cost of more vertices we could argue that we have a 'smaller' pre-suitable complex.

Example 2.22. Let L^1 be the following graph.

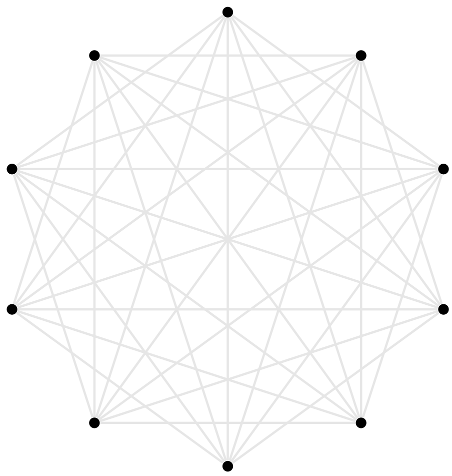

FIGURE 2.16: A graph on 10 vertices with degree 7.

This can be described as a copy of K_{10} where we have removed an embedded loop of edges of length 10. We attach 26 polygons of length 7 to construct L as follows.

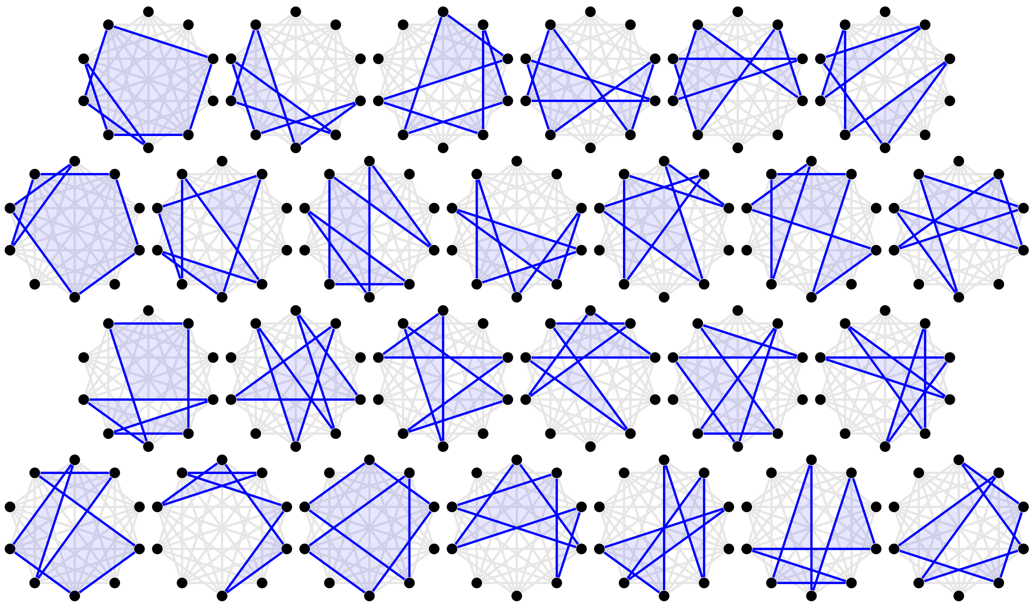

FIGURE 2.17: A collection of 26 polygons that define a pre-suitable complex.

This complex has less edges and polygons than any pre-suitable complex found so far.

We observe that one can trivially construct a pre-suitable complex without a complete 1-skeleton by joining two pre-suitable complexes together via an edge. However this certainly fails to be minimal and is somehow non-satisfactory.

2.5 Catalogue of pre-suitable complexes

Here we collect and compare the pre-suitable complexes we have found so far.

Reference	Vertices	Edges	Polygons	Polygon size
Example 2.22	10	35	26	7
Example 2.16	9	36	28	7
Example 2.9	13	78	66	13
Example 2.17	14	91	78	7

FIGURE 2.18: A catalogue of pre-suitable complexes.

This table is sorted by number of edges.

3 A complex that is not suitable.

The unique selling point of this project is that the groups constructed are type FP. The key property of being suitable that allows this to be true is the property of being acyclic. After showing that each $G(S)$ is type FP we will move on to proving many other properties, in many cases these properties are interesting because $G(S)$ is type FP but do not rely on this finiteness property in order to be true.

Motivated by the intention to allow the reader to more easily follow future arguments, we present an additional construction of a 2-complex that satisfies all properties of being suitable except Property 6 (K is acyclic).

As with suitable complexes with start with a complex of girth 3.

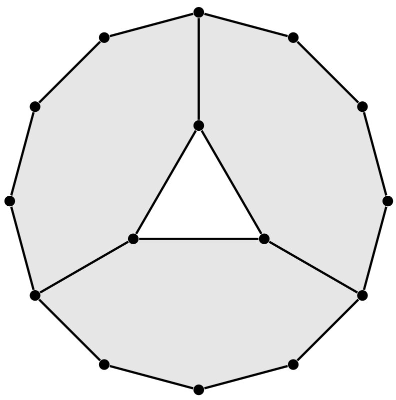

FIGURE 2.19: An analogue to being a pre-suitable complex that is not acyclic.

This complex is considerably smaller than any pre-suitable complex we have found. For analogues to suitable and very-suitable complexes we subdivide as previously.

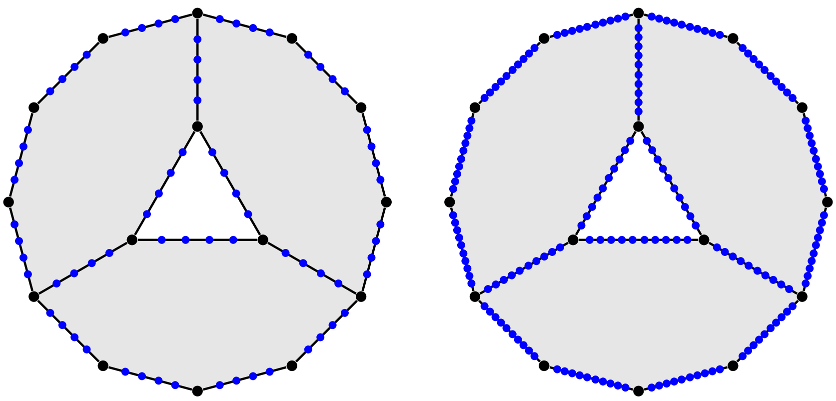

FIGURE 2.20: Analogues to being suitable and (very)-suitable. These will be used in examples where a being acyclic is not necessary.

Chapter 3

Groups of type FP

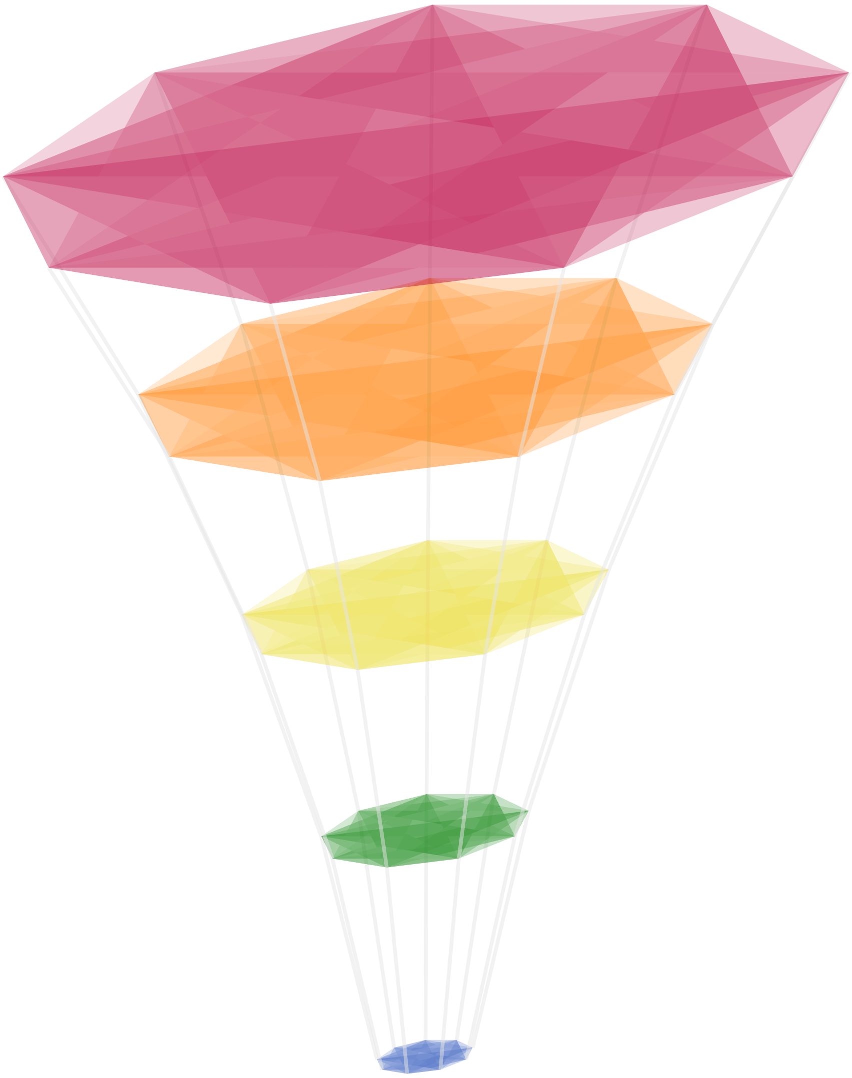

1 Disk diagrams in a presentation for $G(S)$

In this section we will give our first definition of the group $G(S)$ via a presentation. We will then analyse van Kampen-reduced disk diagram in this presentation to generate some theory that will be repeatedly used in later sections.

1.1 A presentation for $G(S)$

Initially we will allow K to be an arbitrary 2-complex. A useful example to consider is the not-suitable complex from Section 3. We also choose some integer $k > 1$. We will define a family of presentations

$$\{\mathcal{P}(S) \mid S \subseteq \mathbb{N}_0\}$$

such that $\mathcal{P}(S)$ is a presentation for the group $G(S)$.

Definition 1.1. Assign to every edge in K an arbitrary direction and then assign a unique label to each directed edge and each polygon. We then take $\mathcal{C}$ to be the finite set of all cycles in K^1 that do not self-intersect.

The presentation $\mathcal{P}(S)$ is defined as follows. We have two types of generator:

- **Edge generators:** a_i for each edge a_i in K.
- **Polygon generators:** t_P for each polygon P in K.

We have three types of relator:

- **Polygon relators:** for each polygon P with boundary word $a_1 a_2 \dots a_{l_P}$ and for each $n \in \mathbb{N}_0 - S$ we have $a_1^{k^n} a_2^{k^n} \cdots a_{l_P}^{k^n}$. We say this relator has degree n and may refer to it as P^n.
- **Closed path relators:** for each closed path γ in $\mathcal{C}$ with boundary word $a_1 a_2 \dots a_{l_\gamma}$ and for each $n \in S$ we have $a_1^{k^n} a_2^{k^n} \cdots a_{l_\gamma}^{k^n}$. We say this relator has degree n and may refer to it as γ^n.
- **t_P-relators:** for each a_i in the boundary of the polygon P we have $t_P a_i t_P^{-1} = a_i^k$.

We may refer to the combined set of polygon relators and closed path relators as **K-relators**.

Example 1.2. Considering the 2-complex K from Figure 2.19, we label K and observe that $\mathcal{C}$ contains 14 cycles.

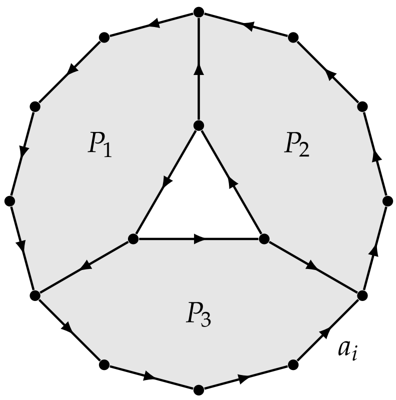

FIGURE 3.1: We label the edges and polygons as well as directing the edges. For brevity we only present a label for one edge, we assume that all the other edges have their own unique labels.

We then observe that we have:

- 18 edge generators;
- 3 polygon generators;
- 3 polygon relators for each $n \in \mathbb{N}_0 - S$;
- 14 closed path relators for each $n \in S$;
- 21 t_P-relators (7 for each polygon).

So, the only difference between $G(S)$ and some $G(T)$ are the closed path relators.

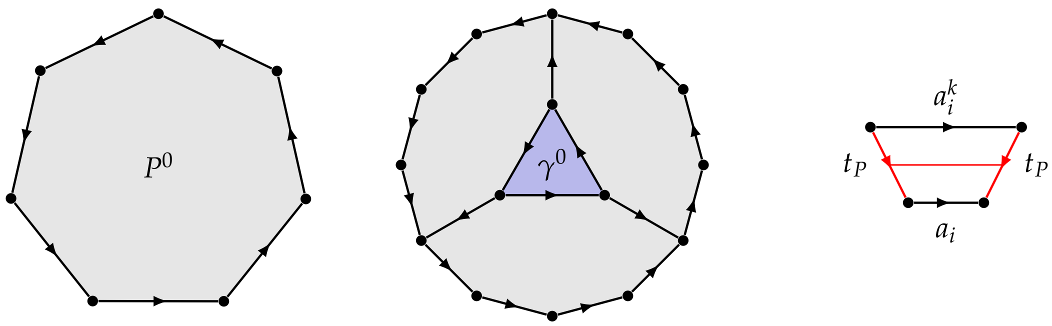

FIGURE 3.2: We represent a polygon relator of degree 0, a closed path relator of degree 0 and a t_P-relator. The red line and edges will be used to show the structure of t_P-relators in disk diagrams later.

The presentation $\mathcal{P}(S)$ is not very efficient in terms of the number of relators. For example one can realise a polygon relator P^n using some other polygon relator P^m and a collection of t_P-relators. We also note that, for $n \in S$, the polygon relator P^n is included in both the polygon relators and the closed path relators. Finally we observe that not all of the closed path relators are needed; in Example 1.2 we could suffice with just 4, or even 1.

The benefit of using this presentation lies in its induced subpresentations.

1.2 Induced presentations of subgroups

In Section 3 we see $G(S)$ defined via a graph of groups. The groups involved in the graph are

- $H(S)$ the central vertex group;
- H_P the edge groups;
- G_P the outer vertex groups.

Definition 1.3. Given a presentation $\mathcal{P} = \langle A \mid R \rangle$ and a subset $B \subseteq A$ we can consider the **induced subpresentation**

$$\langle B \mid R_B \rangle$$

where R_B is defined to be the subset of R that only contains words written in elements of B.

For a generic group presentation we have the obvious trivial induced subpresentations:

- $B = \emptyset$ will induce the empty presentation which gives the trivial group;
- $B = A$ will induce the original presentation $\mathcal{P}$.

Definition 1.4. We realise presentations for the groups above via subpresentations of $\mathcal{P}(S)$.

- Let $\mathcal{P}_H(S)$ be the subpresentation generated by only the edge generators of $\mathcal{P}(S)$, this is a presentation for $H(S)$.
- Let $\mathcal{P}_{H_P}$ be the subpresentation generated by only the edge generators of $\mathcal{P}(S)$ that correspond to edges in the boundary of the polygon P, this is a presentation for H_P.
- Let $\mathcal{P}_{G_P}$ be the subpresentation generated by the edge generators of $\mathcal{P}(S)$ that correspond to edges in the boundary of the polygon P along with the single polygon generator t_P, this is a presentation for G_P.

In the latter two cases we have used suggestive notation to portray the fact that H_P and G_P do not depend on the choice of $S \subseteq \mathbb{N}_0$.

Example 1.5. Expanding on Example 1.2, we observe that $H(S)$ has the 18 edge generators, 3 polygon relators for each $n \in \mathbb{N}_0 - S$ and 14 closed path relators for each $n \in S$.

Letting P be one of the polygons and assuming the boundary word is written $a_1a_2\ldots a_7$, the group H_P has the following presentation:

$$\left\langle a_1, a_2, \ldots, a_7 \,|\, a_1^{k^n} a_2^{k^n} \cdots a_7^{k^n} \;:\; \forall n \in \mathbb{N}_0 \right\rangle.$$

Finally, with the same polygon P, the group G_P has the following presentation:

$$\left\langle a_1, a_2, \ldots, a_7, t_P \,|\, a_1^{k^n} a_2^{k^n} \cdots a_7^{k^n} \;:\; \forall n \in \mathbb{N}_0, \;\; t_P a_i t_P^{-1} = a_i^k \;:\; \forall a_i \in P \right\rangle.$$

When we construct $G(S)$ using a suitable complex K we will realise presentations for H_P and G_P that are almost identical to the ones given above; the only difference will be that the polygons may have longer boundaries.

1.3 Disk diagrams

We wish to discuss disk diagrams in the presentations above. First we recall some background material on disk diagrams and classical metric small cancellation theory.

Let G be a group with presentation

$$\mathcal{P} = \langle A \,|\, R \rangle.$$

Definition 1.6. A **diagram** D in $\mathcal{P}$ is a 2-complex such that each edge is directed and labelled by an element of X and each 2-cell has boundary reading a word in R. Given a 2-cell r in D, we let ∂r be the cycle graph making up the boundary of r. The collection of boundary components of D is denoted ∂D and the number of 2-cells in D is $|D|$.

Example 1.7. Let $G = \mathbb{Z}^2$ with the standard presentation $\langle a, b \,|\, aba^{-1}b^{-1} \rangle$. Below is an example of a corresponding diagram.

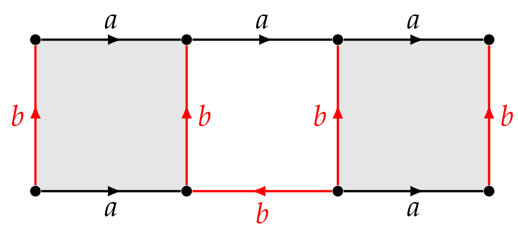

FIGURE 3.3: Here we have a diagram D with a pair of 2-cells and an empty cycle in the middle. The boundary ∂D is the entire 1-skeleton and $|D| = 2$.

Definition 1.8. A **disk diagram** is a diagram that is homeomorphic to a disk.

The diagram in Figure 3.3 is not a disk diagram.

Remark. The key fact that allows disk diagrams to be useful is as follows. A word w written in the generators of a presentation is trivial in the corresponding group G if and only if there is a disk diagram D with boundary $\partial D = w$.

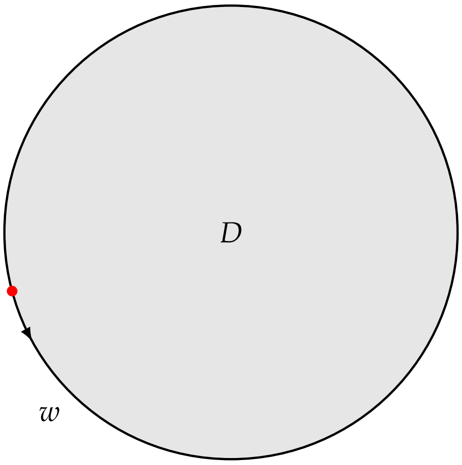

FIGURE 3.4: A disk diagram D with boundary word reading w. If we started reading the boundary at a different vertex we would read a different word, but only up to conjugation (cyclic permutation).

Definition 1.9. A disk diagram D is **van Kampen-reduced** if there are no pairs of 2-cells in D such that the word read around the boundary of their union freely reduces to the empty word.

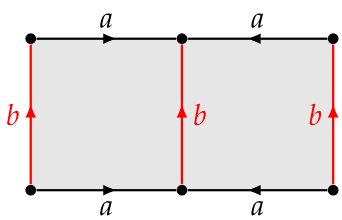

FIGURE 3.5: The word read around this disk diagram is $baa^{-1}b^{-1}aa^{-1}$ which freely reduces to the empty word. It is therefore not van Kampen-reduced, moreover, any diagram containing this disk diagram is not van Kampen-reduced.

Definition 1.10. A disk diagram D satisfies the **classical metric small cancellation condition** $C'(\lambda)$ if, whenever a pair of 2-cells r_1 and r_2 with $|\partial r_1| \leq |\partial r_2|$ share a path of edges p in their boundaries, then

$$|p| < \lambda |\partial r_1|.$$

We culminate this section with a useful lemma that ties our definitions together. We first define a **ladder** to be a van Kampen-reduced disk diagram such that the **dual graph** (a vertex for each 2-cell joined by an edge if their boundaries share at least one edge) is a path graph.

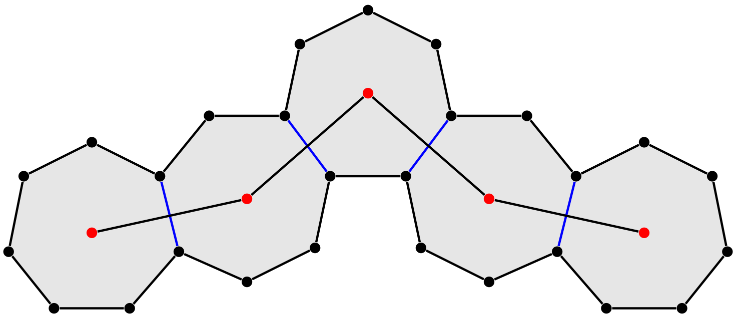

FIGURE 3.6: Here we present a ladder made up of 7-gons that satisfies $C'(\frac{1}{6})$ along with its dual graph. We assume for the sake of the example that each edge has a unique label.

We say a diagram D is **non-trivial** if it contains more than one 2-cell.

Lemma 1.11 (Greendlinger). *Let D be a non-trivial van-Kampen reduced disk diagram satisfying $C'(\lambda)$ for some $\lambda \leq \frac{1}{6}$. There are at least two relators in D such that $\partial r \cap \partial D$ contains a connected component q with*

$$|q| > (1 - 3\lambda)|\partial r| \geq \frac{1}{2}|\partial r|.$$

Moreover, if there are exactly two such relators, then D is a ladder.

Greendlinger's Lemma is a central tool in small cancellation theory. A good background for classical small cancellation theory is (17) and for a direct proof of the above statement see (19).

Definition 1.12. If every diagram in a given presentation satisfies $C'(\lambda)$ then we say the **presentation satisfies** $C'(\lambda)$.

Proposition 1.13. *If $l_P \geq 13$ then $\mathcal{P}_{H_P}$ satisfies $C'(\frac{1}{6})$.*

Proof. Let P^m and P^n be arbitrary distinct relators in $\mathcal{P}_{H_P}$ with $m < n$. The longest shared subpath between these relators is of the form $a_i^{k^m} a_j^{k^m}$ with length $2k^m$. Since the shortest of these relators has length $l_P k^m \geq 13k^m > \frac{1}{6}2k^m$, these relators are pairwise $C'(\frac{1}{6})$. As these relations were chosen arbitrarily we deduce that $\mathcal{P}_{H_P}$ satisfies $C'(\frac{1}{6})$. □

In general we can deduce that $\mathcal{P}_{H_P}$ satisfies $C'(\frac{2}{l_P - 1})$.

When a presentation $\mathcal{P}$ satisfies $C'(\frac{1}{6})$, supposing there is an algorithm that can determine the set of relators in $\mathcal{P}$ (this may not always be the case with these groups since S may be chosen undecidably), then there is a corresponding algorithm solving the word problem. Such an algorithm must take, as input, a word w and return whether or not w represents the trivial word in G.

Algorithm 1.14 (Dehn's Algorithm)**.** Let w be a word written in generators of a presentation $\mathcal{P}$ that satisfies $C'(\frac{1}{6})$. We generate a sequence of words $w_0, w_1, w_2, \ldots$ as follows:

- **Step** 0**:** Set $w_0 = w$.
- **Step** j**:** Consider w_{j-1};
  - **IF** there is a subword u of w_{j-1} such that u makes up more than half of some relator uv in $\mathcal{P}$ then replace u by v^{-1} in w_{j-1} and set the resulting word to be w_j.
  - **ELSE** terminate the algorithm with output w_{j-1}.

By Lemma 1.11, this will terminate with the empty word if and only if w represents the trivial word in $\mathcal{P}$.

Greendlinger's Lemma and Dehn's Algorithm are often accompanied by Corollary 1.18. We will first need a handful more definitions.

Definition 1.15. A diagram is **spherical** if it is homeomorphic to a disk.

Definition 1.16. A presentation $\mathcal{P}$ for a group is said to be **aspherical** if there are no van Kampen-reduced spherical diagrams in $\mathcal{P}$.

Example 1.17. Consider the following presentation for $\mathbb{Z}^3$:

$$\langle x, y, z \mid [x, y], [y, z], [z, x] \rangle .$$

There is a van Kampen-reduced spherical diagram that can be observed as the boundary of a 3-dimensional cube.

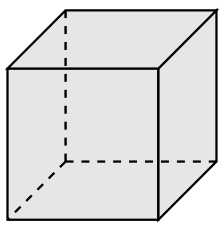

The canonical presentation for $\mathbb{Z}^2$ is aspherical, and for $n \geq 3$ the canonical presentation for $\mathbb{Z}^n$ is not, this is analogous to the idea that we can think of aspherical presentations as being 2-dimensional, we will make this more rigorous in Section 3.

Corollary 1.18. *If a presentation satisfies $C'(\frac{1}{6})$ then it is aspherical.*

Proof. Suppose there exists a van Kampen-reduced spherical diagram $\mathcal{S}$. Remove a relator r' from $\mathcal{S}$ to produce a disk diagram D. By Lemma 1.11, D contains a relator r with

$$|\partial r \cap \partial D| > \frac{1}{2}|\partial r|.$$

This means the subdiagram of $\mathcal{S}$ consisting of r and r' does not satisfy $C'(\frac{1}{6})$, a contradiction. $\square$

1.4 van Kampen-reduced disk diagrams in $G(S)$

From now on we require K to be a very-suitable complex with subdivision constant d and pre-subdivision vertex set L^0, fix some integer $k > 1$ and some $S \subseteq \mathbb{N}_0$. Let Γ be the Cayley graph corresponding to $\mathcal{P}(S)$.

The aim of this section is to carefully construct disk diagrams for loops in Γ that satisfy certain properties. This will culminate in the following theorem.

Theorem 1.27. Let γ be a loop in the Cayley graph for $\mathcal{P}(S)$. There is a disk diagram D in $\mathcal{P}(S)$ with $\partial D = \gamma$ such that

$$|D| \leq \left(\frac{1}{2}|\gamma|^2 + |\gamma|\right) C^{\frac{|\gamma|}{2}},$$

where $C = |L^0|d$. Moreover, if $k \geq C$ then D does not contain any K-relators of degree $n > |\gamma|$.

We first restrict our attention to $H(S)$. Let Λ be the Cayley graph of $H(S)$ corresponding to the presentation $\mathcal{P}_H(S)$. In (8) a graphical small cancellation theoretical property is discussed of $H(S)$. Here we will strengthen this claim without adhering to the language of graphical small cancellation theory.

Proposition 1.19. *Let γ be a loop in Λ. There is a van Kampen-reduced disk diagram D in $\mathcal{P}_H(S)$, with $\partial D = \gamma$, such that D satisfies $C'(\frac{1}{6})$.*

Moreover, within D:

1. *if two relators r and r' have respective degrees $n < n'$ then the longest connected component of $\partial r \cap \partial r'$ has length at most $2k^n$;*
2. *if r is a closed path relator then it does not share any edges with a relator of the same degree.*

Proof. Let D' be some van Kampen-reduced disk diagram in $\mathcal{P}_H(S)$ with $\partial D' = \gamma$. Each relator in D' is a K-relator, so we can label each relator by its degree. We will construct the diagram D above in stages.

Consider any pairs of relators in D' that bear the same label and whose boundaries share at least an edge. If the boundary of the union of this pair corresponds to a closed path in K^1 then we say these K-relators are **aligned**, otherwise we say they are

misaligned. If they are aligned, then we can alter the diagram D' such that the word read around the boundary of the union is reduced (by performing **diamond moves** as coined in (9)), we perform this process for all such pairs in D'.

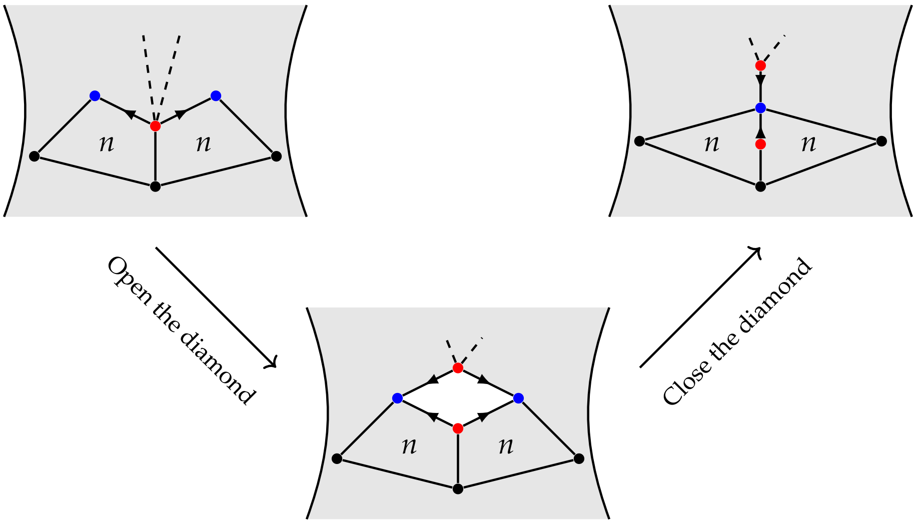

FIGURE 3.7: Pairs of edges bearing the same arrows are taken to share the same word, so in the top left we have a pair of aligned K-relators sharing the label n such that the boundary of the union is not freely reduced. By splitting the red vertex into two vertices, pulling them apart and then pushing the blue vertices together we perform a diamond move. We see that the boundary of the region has not changed, yet the boundary of the union of the K-relators is now reduced.

Now consider pairs of aligned closed path relators in D' labelled some $n \in S$. In each case, since the boundary of their union corresponds to a reduced closed path in K^1, we replace the pair by a single 'relator' of degree n. We repeat this process across D', the resulting diagram has no aligned closed path relators.

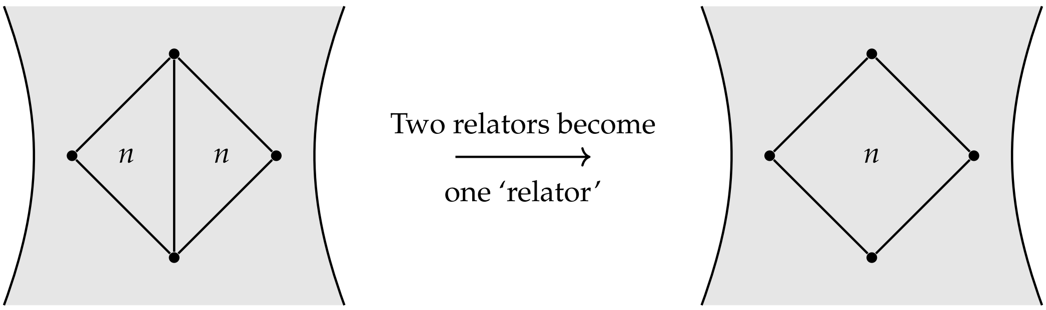

FIGURE 3.8: Two aligned closed path relators are replaced by a single closed path 'relator'. We use quotations as this may not be included in our original set of relators, we will deal with this issue next.

In the previous process we may have created a 'relator' that corresponds to a closed path in K^1 that self-intersects and so does not correspond to a relator in $\mathcal{P}_H(S)$. Given such a 'relator', we will pinch the self-intersection out as in Figure 3.9. We repeat until our diagram, which we now denote D, returns to being a disk diagram in $\mathcal{P}_H(S)$

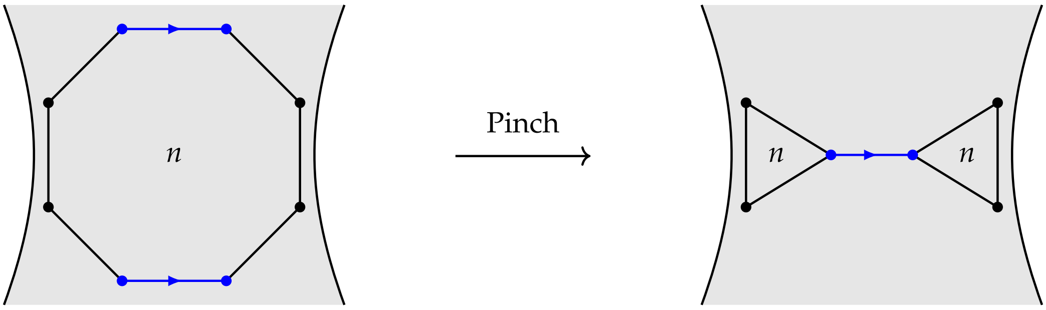

FIGURE 3.9: On the left is a K-relator corresponding to a closed loop in K^1 such that the blue paths are identified in K^1. The right is the result of pinching the red paths together; a pair of relators of degree two, neither of which contain more than one copy of any vertex.

We now show that D has the required properties. First we observe that at no stage did we alter the boundary of D', so $\partial D = \gamma$. Suppose D is non-trivial. Given any two relators r and r' in D, with labels $n \leq n'$ respectively, that share an edge, let q be a connected component of $\partial r \cap \partial r'$. We have one of the following:

- $n = n'$ and r and r' are misaligned, so $|q| < k^n$;
- $n = n$ and r and r' are aligned, so r and r' correspond to the boundaries of distinct polygons P, Q of K and so $|q| < \frac{1}{6}\min\{|\partial r|, |\partial r'|\}$;
- $n < n'$, so $|q| \leq 2k^n$.

We also observe that in D any relator r of degree n has boundary length at least $3dk^n \geq 39k^n$, as such we deduce that D satisfies $C'(\frac{1}{6})$. □

Before we prove our first corollary we recall the presentation of a polygon subgroup H_P:

$$\mathcal{P}_{H_P} = \langle a_1, a_2, \ldots, a_{l_P} \mid a_1^{k^n} a_2^{k^n} \cdots a_{l_P}^{k^n} = 1 \ : \ n \in \mathbb{N}_0 \rangle$$

where $a_1 a_2 \ldots a_{l_P}$ is the word read around ∂P.

Corollary 1.20. *Let γ be a loop in Λ such that, for some polygon P in K, every edge in γ appears in ∂P. There is a van Kampen-reduced disk diagram D in $\mathcal{P}_{H_P}$ with $\partial D = \gamma$.*

Moreover, if the word read around γ is written in $\{a_1^k, a_2^k, \ldots, a_{l_P}^k\}$ then there are no relators of degree 0 in D.

Proof. Let D be the van Kampen-reduced disk diagram that satisfies $C'(\frac{1}{6})$ as in Proposition 1.19. Let r be one of the relators in D such that $\partial r \cap \partial D$ contains a connected component q with

$$|q| > \frac{1}{2}|\partial r|.$$

If r is a polygon relator Q^n then $Q = P$ by the small cancellation condition on polygons of K.

Suppose instead that r is not a polygon relator, then it is some closed path relator of degree $n \in S$. So there is a path $a_1a_2\ldots a_d$ in K^1 that does not lie in ∂P and corresponds to a subpath of ∂r of the form $p_r = a_1^{k^n}a_2^{k^n}\cdots a_d^{k^n}$. This subpath does not lie in ∂D since it consists of edges that are not in ∂P.

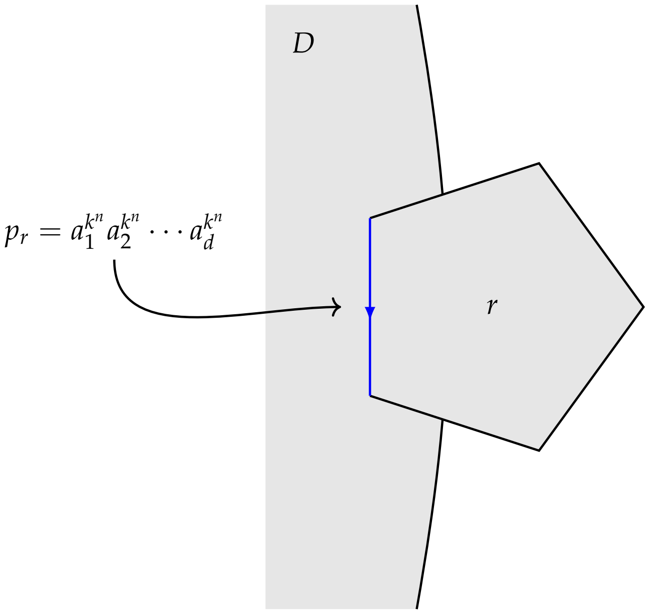

FIGURE 3.10: A closed path relator with over half of its boundary in ∂D. The blue path represents a string of edges that do not appear in ∂P that corresponds to a path in K^1 that does not pass through any vertices from L^0.

Since none of these edges are in the boundary of D they must each lie between two relators. Since r is a closed path relator, by Proposition 1.19, it can share at most $2k^n$ with any other relator and we deduce that, since $d \geq 11$ the number of relators in D incident to r via the path p_r is at least 6.

Construct a graph as follows:

- **Step** 0**:** Start with a single vertex v_r corresponding to the relator r.
- **Step** 1**:** For each relator r' incident to r via an edge in $p_r = a_1^{k^n}a_2^{k^n}\cdots a_d^{k^n}$ add a vertex $v_{r'}$ and an edge between v_r and $v_{r'}$.
- **Step** 2**:** Since none of these relators can be aligned with another relator via their own subpath $p_{r'} = a_1^{k^m}a_2^{k^m}\cdots a_d^{k^m}$, else they could not have been added to the graph in the first place, and none of the edges in $p_{r'}$ can appear in ∂D, each of these r' are incident to at least 5 other relators via $p_{r'}$. Add a new vertex for each of these that are not already in the graph and connect via an edge to $v_{r'}$.
- **Step** j**:** Repeat until we find there are no new vertices to consider.

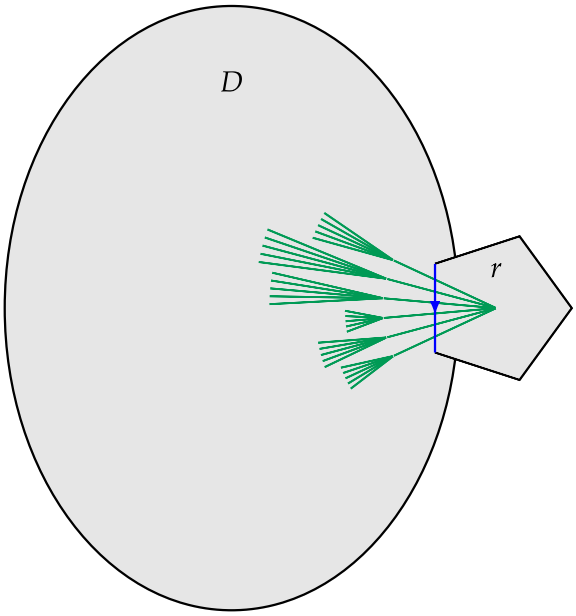

FIGURE 3.11: We use the structure of the relator r to construct a planar graph where every vertex has degree at least 6. Vertices represent relators and green edges represent two relators sharing an edge in their boundary. This represents the first steps in constructing the graph described above.

The resulting graph is finite planar since it can be seen as the dual of a finite disk diagram and since there are no finite planar graphs with average degree at least 6, we have a contradiction.

So r must be a polygon relator P^n. Removing r from D gives a new diagram with a shorter boundary where we can repeat the argument. Repeating until the empty diagram gives our first claim.

For the second we observe that r certainly cannot be P^0. So removing r from D will give us a new shorter diagram with boundary still written in $\{a_1^k, a_2^k, \ldots, a_{l_P}^k\}$. Repeating until the empty diagram gives our second claim. □

We will later find that diagrams of this form regularly occur when we consider diagrams in $\mathcal{P}(S)$. Diagrams in $\mathcal{P}(S)$ differ only from diagrams in $\mathcal{P}_H(S)$ via the addition of edges labelled t_P and t_P-relators. We will first discuss the structure that relators posses.

Let D be a van Kampen-reduced disk diagram in $\mathcal{P}(S)$. A t_P**-corridor** T is a connected subdiagram of D, consisting only of t_P-relators such that any pair of adjacent 2-cells in T are attached via an edge labelled t_P and, if there is an edge in the boundary of T labelled t_P, then it lies in the boundary of D. If a t_P-corridor T has an edge labelled t_P in ∂D then we say T is **open**, otherwise we say T is **closed**. Open t_P-corridors are easily detected on observation of ∂D. Closed t_P-corridors are harder to detect but they do hold some accessible structure.

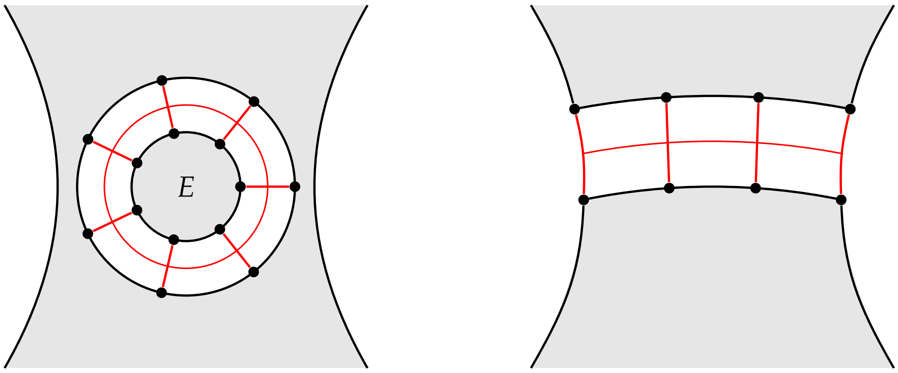

FIGURE 3.12: On the left we have a closed t_P-corridor T where we have used E to label the enclosure (defined below) of T and on the right we have an open t_P-corridor. In each case the red edges represent edges labelled t_P and the red line is used to exaggerate the 'corridor-like' behaviour.

We distinguish the possible orientations of t_P-corridors. Let T be a t_P-corridor in a diagram D. There are two strings of edges in ∂T that do not feature the label t_P (in a closed t_P-corridor this will be exactly the two components of ∂T). One of these strings is k times as long as the other, we will call this the **long side** of T. The shorter string we call the **short side** of T. Given a closed t_P-corridor T there is a unique subdiagram D' of D such that D' is a disk diagram and $\partial D'$ is exactly one of the sides of T. We call D' the **enclosure** of T.

Proposition 1.21. *Let γ be a loop in Γ. There is a van Kampen-reduced disk diagram D in $\mathcal{P}(S)$ such that D contains no closed t_P-corridors.*

Proof. Let D' be an arbitrary van Kampen-reduced disk diagram in $\mathcal{P}(S)$. We partially order the closed t_P-corridors in D' such that $T > T'$ if T' appears in the enclosure of T. The enclosure of a minimal such T will have all its boundary edges in ∂P and contain no edges labelled by some t_Q. So this can be considered as a disk diagram in $\mathcal{P}_H(S)$ and, since each edge in the boundary is in the boundary of the polygon P, it can be considered as a diagram in $\mathcal{P}_{H_P}$ by Corollary 1.20.

We consider the two possible orientations of T.

- If the enclosure of T is incident to the short side of T then we replace the union of T and its enclosure by a new diagram in H_P, similar to the enclosure but each edge has been replaced by k copies of itself and so each polygon relator P^n is replaced by P^{n+1}.

- If the enclosure of T is incident to the long side of T then, by Corollary 1.20, it contains no copies of P^0. We therefore replace the union of T and its enclosure by a new diagram in H_P, similar to the enclosure but each path of edges of the form

a_i^k (as written in $\{a_1^k, a_2^k, \ldots, a_{l_P}^k\}$) is replaced by a_i and so each polygon relator P^n is replaced by P^{n-1}.

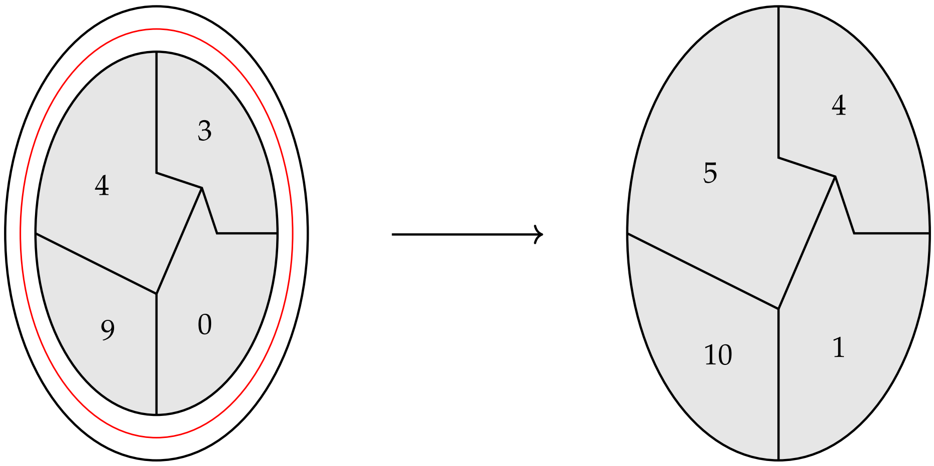

FIGURE 3.13: Here we have assumed that the t_P-corridor has its short side incident to its enclosure.

After performing this we have one less closed t_P-corridor and we repeat the process until there are no closed t_P-corridors. □

We are left with open t_P-corridors and subdiagrams that can be considered as disk diagrams in $\mathcal{P}_H(S)$, will refer to these subdiagrams by D_i and assume that each of them is as described in the proof of Proposition 1.19. If a subdiagram D_i is incident to a single open t_P-corridor T then we will call it a **leaf**. We will use the convention that if there is a t_P-corridor with a side entirely contained within the boundary of a disk diagram D, then we will say there is an empty leaf D_i along its side.

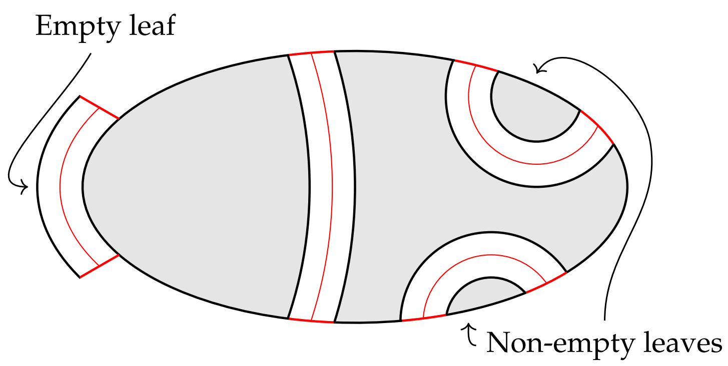

FIGURE 3.14: This diagram has three leaves, one of which is empty. In total it has five subdiagrams D_i and four open t_P-corridors. There will always be one more D_i than there are open t_P-corridors.

If there are no open t_P-corridors in D then we say $D_i = D$ is the only leaf. We next consider a technical proposition that, after we have chosen a leaf D_i, allows us to move complications away from D_i, to be dealt with elsewhere.

Proposition 1.22. *Let γ be a loop in Γ such that there is no subword of γ reading over half of some K-relator. If there is a t_P edge in γ then there is a van Kampen-reduced disk diagram D in $\mathcal{P}(S)$ such that D contains an open t_P-corridor T incident to a leaf D_i where, for every relator r in D_i, no connected component q of $\partial r \cap \partial T$ has*

$$|q| > \frac{1}{2} l_P k^n.$$

Proof. Let D be a van Kampen-reduced disk diagram in $\mathcal{P}(S)$ with $\partial D = \gamma$. We may assume that D contains no closed t_P-corridors by Proposition 1.21 and that each subdiagram D_i is set up as in Proposition 1.19. Let T be an open t_P-corridor in D that is incident to a leaf D_i, such a T exists as there is an edge labelled t_P in the boundary.

Suppose this D_i contains a polygon relator r of degree n such that there is a connected component q of $\partial r \cap \partial T$ with

$$|q| > \frac{1}{2} l_P k^n.$$

In this case, since our polygons satisfy $C'(\frac{1}{6})$ and this polygon follows more than half of the boundary of P, r is a copy of P^n. We alter D by replacing the union of T and r by a new t_P-corridor and a new polygon relator of degree $n \pm 1$ (depending on which side of T the leaf is incident to). We call this a **slide** out of D_i.

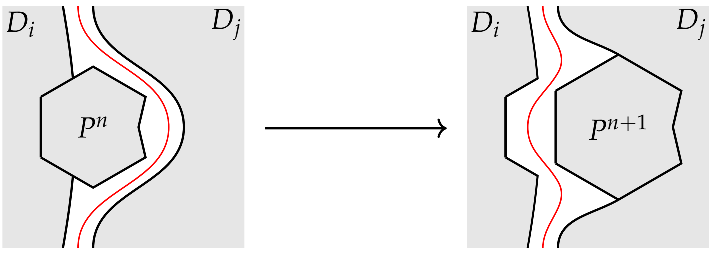

FIGURE 3.15: Assuming the short side of T is incident to D_i, we slide a polygon relator out of D_i into some D_j.

When the new polygon relator is added to some D_j on the other side of T, it may create a pair of relators that are not van Kampen-reduced. If this is the case then remove both relators in the pair and, since the word around their boundary freely reduces to the empty word, close up the hole.

Suppose instead there is a closed path relator r in D_i, of degree n, such that there is a connected component q of $\partial r \cap \partial T$ with

$$|q| > \frac{1}{2} l_P k^n.$$

The path q consists of over half the boundary of P^n, so replace the relator r by a shorter closed path relator r', also of degree n, and a copy of the polygon relator P^n. This new polygon relator certainly has a subpath of over $\frac{1}{2}l_P k^n$ edges incident to T so we slide it out of D_i as below.

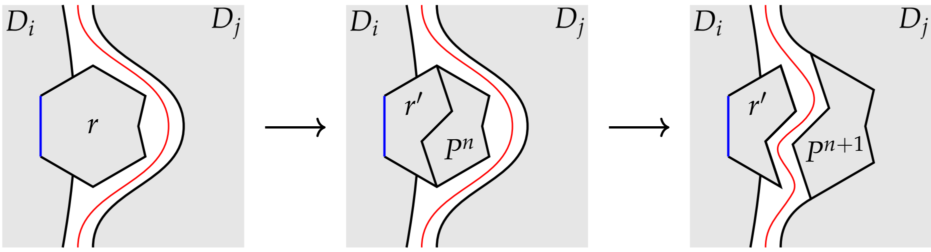

FIGURE 3.16: Again assuming the short side of T is incident to D_i, we alter r such that we have a polygon relator to slide and then we perform such a slide. The blue edge represents a path of edges that do not appear in ∂P.

We repeat this process until we have neither of the cases discussed. At this point D has the properties described above. □

The leaf D_i described above is now very restricted. We know it satisfies $C'(\frac{1}{6})$, yet there seems to be little room for relators with large components in ∂D_i. The following proposition ties these ideas together where these can appear.

Proposition 1.23. *Under the conditions described in Proposition 1.22 the leaf D_i is a ladder such that, if r is a relator of degree n in D_i, we have*

$$dk^n \leq |\partial r \cap \partial D| \leq \frac{|L^0|}{2} dk^n$$

and

$$dk^n \leq |\partial r \cap \partial T| \leq \frac{|L^0|}{2} dk^n,$$

where T is the unique open t_P-corridor incident to D_i.

Moreover, up to possibly performing some extra slides, we can deduce the word read around ∂D_i via observation of $\partial D_i \cap \partial D$.

Proof. We will consider D_i as an overlapping union of subdiagrams $D_{i,j}$. These are defined to be the smallest subdiagrams whose union is D_i and the intersection of $D_{i,j}$ and $D_{i,j+1}$ is a single relator that has edges in both ∂D and ∂T. If we can prove that each $D_{i,j}$ is a ladder then D_i is also a ladder.

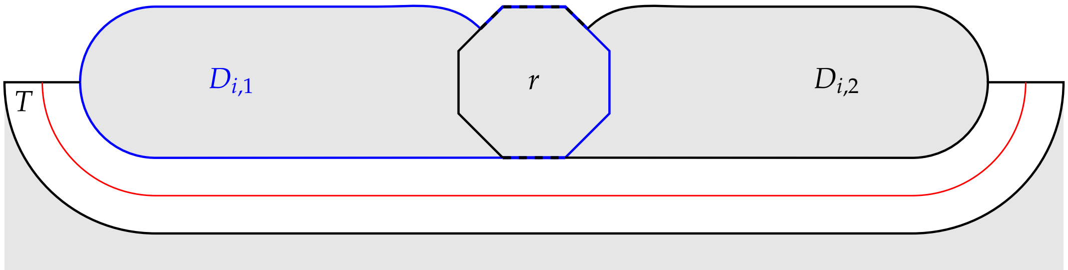

FIGURE 3.17: We represent the decomposition of a leaf D_i into two subdiagrams: $D_{i,1}$ with blue boundary and $D_{i,2}$ with black boundary. We observe that they both include the relator labelled r which has edges in ∂D and ∂T.

Consider some $D_{i,j}$ in D_i. Since $D_{i,j}$ satisfies $C'(\frac{1}{6})$ we use Lemma 1.11 and let r be one of the relators with a connected component $q \subseteq \partial r \cap \partial D_{i,j}$ such that $|q| > \frac{1}{2}|\partial r|$. This q cannot appear entirely in ∂D by our supposition. Suppose q appears entirely in ∂T, if r is a polygon relator then it corresponds to the polygon P and we have already slid it out of D_i as in Proposition 1.22.

Suppose instead that r is not a polygon relator. We construct a graph as in the proof of Corollary 1.20, observing that this time the degree of each vertex need not be 6, since we may have relators r' with $p_{r'} = a_1^{k^m} a_2^{k^m} \cdots a_d^{k^m}$ intersecting ∂D.

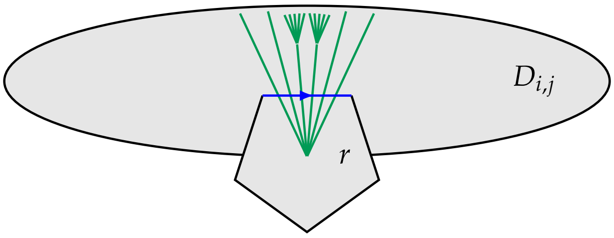

FIGURE 3.18: We construct a graph as in Figure 3.11 except here we may have vertices of degree less than 6. Each of these vertices corresponds exactly to a relator r' such that a portion of the path $p_{r'}$ in ∂D.

We will denote this graph Γ_r. Since the relator r has no edges in ∂D our graph must have at least 7 vertices and at least 6 of them must correspond to a relator r' with some $p_{r'}$ in ∂D. Let D'_r be the minimal subdiagram of D_r that contains every relator corresponding to a vertex in Γ_r, and, by removing D'_r from D, we are left with connected disk diagram. Finally, let D''_r be the maximal subdiagram of D'_r where we remove all relators r from D'_r with $|\partial r \cap \partial D'_r| > 0$ except those where $\partial r \cap \partial D'_r \subseteq \partial D$.

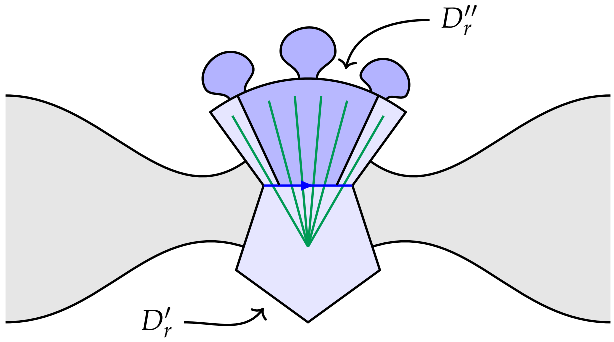

FIGURE 3.19: The entire blue subdiagram is D'_r and the darker subdiagram is D''_r.

The subdiagram $D'_r - D''_r$ is a ladder where every relator r' has exactly $|\partial r'| - |p_{r'}|$ edges in $\partial D'_r - \partial D''_r$. Moreover each relator r'' of degree n'' in D''_r is incident to at most two relators in $D'_r - D''_r$, sharing at most $2k^{n''}$ edges with each. So we have $|\partial r'' \cap \partial D'_r| \leq 4k^n$ and, since D''_r satisfies $C'(\frac{1}{6})$, we apply Greendlinger's Lemma (1.11).

If D''_r is a ladder then each relator shares less than $\frac{1}{6}$ of its edges with any other relator and at most $4k^n$ edges with $\partial D'_r$. Since

$$4k^n < dk^n \leq \frac{1}{7}|\partial r|$$

we see that each relator in the ladder has a connected component q of $\partial r \cap \partial D$ with

$$|q| > \left(\frac{1}{6} + \frac{1}{6} + \frac{1}{7}\right)|\partial r| = \frac{20}{42}|\partial r| < \frac{1}{2}|\partial r|,$$

a contradiction.

If D''_r is not a ladder then there are three relators in D''_r with a connected component of $\partial r \cap \partial D''_r$ of length over half $|\partial r|$. None of these components can appear entirely in $\partial D'_r$ so at least one of them must appear entirely in ∂D, a contradiction.

We deduce that $D_{i,j}$ can only have two relators fitting Greendlinger's claim and so, since each $D_{i,j}$ must be a ladder, so is D_i.

Next we consider the claimed inequalities. The upper bounds follow immediately since otherwise we would have a relator with over half its boundary in either ∂T or ∂D. For the lower bound we let r be come K-relator in the ladder and consider the following cases.

- If r is a polygon relator then, considering either ∂D or ∂T, it has length at least $7dk^n$ and since less than $\frac{5}{6}$ of its edges are already accounted for we find dk^n as a loose lower bound.
- If r is a closed path relator then, as argued above, the whole of the path p_r (of length dk^n) must appear in ∂D giving $dk^n \leq |\partial r \cap \partial D|$. For the intersection with

∂T we observe that r can only share at most $2k^n$ edges with each neighbour and less than $\frac{1}{2}|\partial r|$ with ∂D. Since $|\partial r| \geq 3dk^n$ and $d \geq 11$ we have

$$|\partial r \cap \partial T| > |\partial r| - \frac{1}{2}|\partial r| - 4k^n = \frac{1}{2}3dk^n - 4k^n \geq dk^n.$$

Finally we argue that we can see enough of each relator in the ladder to deduce the word read along $\partial D_i \cap \partial T$. We will do this relator by relator.

Let r be some relator in D_i. If r is a polygon relator then it has at least $\frac{1}{6}$ of its boundary incident to ∂T, so, by the small cancellation condition on polygons, r is a copy of P^n. On the other hand, if r is a closed path relator then it has at least dk^n edges in ∂D, shares at most $2k^n$ with each neighbour and has at least dk^n edges that appear in ∂P. The boundary of a r can be considered as corresponding to a path that follows a portion of ∂P (we choose this to be maximal and containing the dk^n edges that appear in ∂T) and some other path through the 1-skeleton K^1.

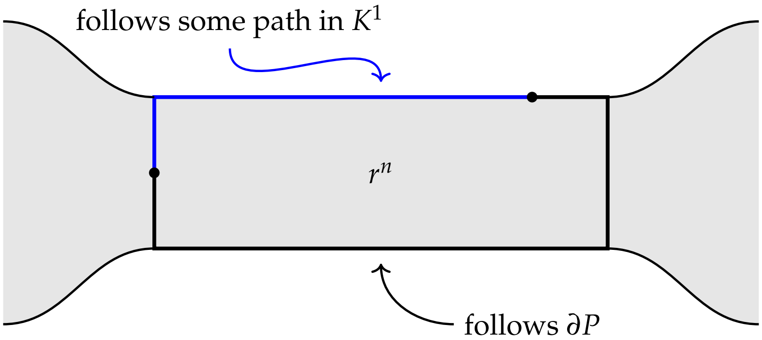

FIGURE 3.20: A decomposition of a closed path relator (via its corresponding closed path) into a path following ∂P and a path that traverses K^1.

Now, consider the unique pair of vertices that lie between these two paths. Since one of the edges incident to each of these does not appear in ∂P, they must either appear in ∂D, or lie between r and a neighbour in D_i. Since there are only $2k^n$ edges between r and each neighbour, there can be only one vertex that corresponds to a vertex in L^0 and so we know exactly the blue path of edges in Figure 3.20. So, the boundary of r is defined up to potentially choosing which way around ∂P we traverse between these vertices and we are free to make either choice since, up to performing a slide of some P^n, the diagram is the same.

We can tell the difference between polygon relators and closed path relators by the existence (or not) of the path p_r in ∂D. □

Corollary 1.24. *Let γ be a loop in Γ. At least of the following holds:*

- *There is a subword of γ reading over half some K-relator.*

- *There are a pair of t_P edges, with opposite orientation, such that the word between them makes up one side of a ladder as in Proposition 1.23.*

We now have enough information to present an algorithm on van Kampen-reduced disk diagrams in $\mathcal{P}(S)$. One may be able this into an algorithm solving the word problem for $G(S)$ when S is recursively presented.

Algorithm 1.25. Let D be a van Kampen-reduced disk diagram in $\mathcal{P}(S)$. Construct a sequence of van Kampen-reduced disk diagrams $D_0, D_1, \ldots$ as follows:

- **Step** 0**:** Set $D_0 = D$.
- **Step** j**:** Consider D_{j-1};
  - **IF** there is a K-relator r in D_{j-1} such that $\partial r \cap \partial D$ contains a connected component q with $|q| > \frac{1}{2}|\partial r|$ then remove r from D_{j-1} and set D_j to be this new diagram.
  - **ELIF** there is a leaf adjacent to a t_P-corridor T as in Proposition 1.23 then remove both T and the corresponding leaf region from D_{j-1} and set D_j to be this new diagram.
  - **ELSE** D_{j-1} is the empty diagram and we terminate the algorithm.

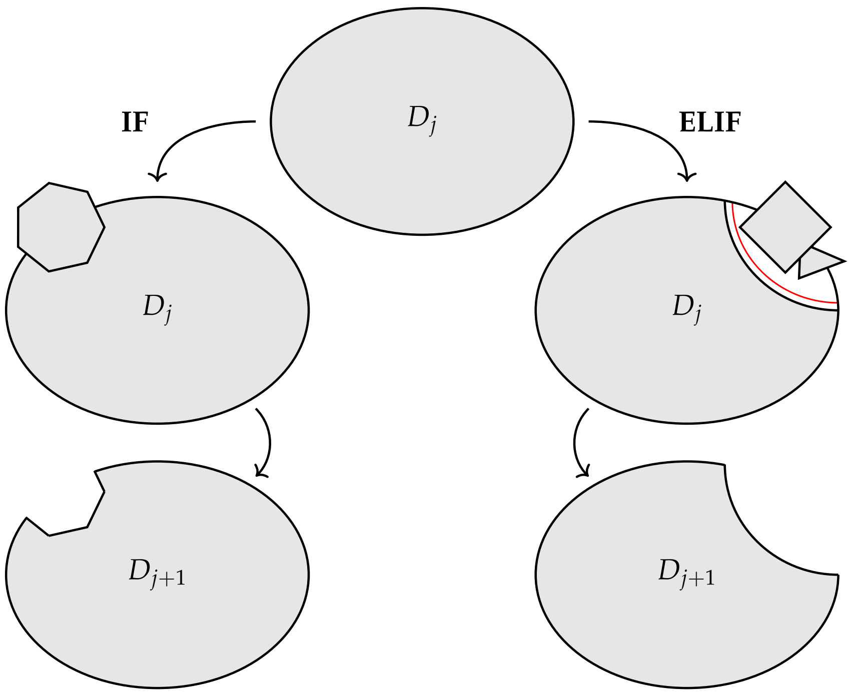

FIGURE 3.21: We either remove a K-relator or a t_P-corridor and a leaf.

We wish to show that this algorithm will terminate in a finite number of steps depending on $|\partial D|$.

Proposition 1.26. *Algorithm 1.25 terminates in at most*

$$\left(\frac{1}{2}|\partial D|^2 + |\partial D|\right) C^{\frac{|\partial D|}{2}}$$

steps, where $C = |L^0|k$.

Proof. Suppose we are at step j considering D_j. If we have the first case then $|\partial D_j| < |\partial D_{j-1}|$. In the second case the boundary could get up to $C = |L^0|k$ times longer, however we have reduced the number of open t_P-corridors. As such our algorithm will terminate in at most

$$(t(D)+1)C^{t(D)}|\partial D|$$

steps, where $t(D)$ is the number of open t_P-corridors in D. This bound is certainly not sharp, however it does agree with the classical bound of $|\partial D|$ for $H(S)$ when there are no open t_P-corridors. We can loosely bound the number of open t_P-corridors by $|\partial D|/2$ since any open t_P-corridor must have at least two edges in the boundary of D and we have our bound. □

We have enough to prove the main theorem of this section.

Theorem 1.27. *Let* γ *be a loop in the Cayley graph for* $\mathcal{P}(S)$. *There is a disk diagram* D *in* $\mathcal{P}(S)$ *with* $\partial D = \gamma$ *such that*

$$|D| \leq \left(\frac{1}{2}|\gamma|^2 + |\gamma|\right) C^{\frac{|\gamma|}{2}},$$

where $C = |L^0|d$. *Moreover, if* $k \geq C$ *then* D *does not contain any* K*-relators of degree* $n > |\gamma|$.

Proof. For the first claim we observe that at each step we have $|D_{j+1}| < |D_j|$ so we can bound $|D|$ above by the number of steps in Algorithm 1.25.

For the second claim let D be some van Kampen-reduced disk diagram containing a K-relator r of degree n. In the process of Algorithm 1.25 we must at some step, say j, remove r from D. In order to do this we must have

$$|\partial D_j| > dk^n$$

as, at worst, r appears in some ladder with at least this many edges in ∂D.

Now, in the process of Algorithm 1.25, the boundary of the diagram can only become larger in at most $\frac{|\partial D|}{2}$ many steps. In each of these steps it may become at most $|L^0|k$

times larger. So we have, for all j,

$$|\partial D_j| \leq |\partial D| \left(|L^0|k\right)^{\frac{|\partial D|}{2}}.$$

Now, supposing that $|\partial D| \leq n$, we have

$$|\partial D_j| \leq n \left(|L^0|k\right)^{\frac{n}{2}}.$$

So, in order for D to contain the relator r, we must satisfy

$$dk^n < n\left(|L^0|k\right)^{\frac{n}{2}} \iff d\left(\frac{k}{|L^0|}\right)^{\frac{n}{2}} < n \implies d^{\frac{n+2}{2}} < n,$$

where the last implication follows from $k \geq |L^0|d$. Since $d \geq 11$, this last statement is always false and we complete our claim. □

Lemma 1.28. *Let D be a van Kampen-reduced disk diagram D with $\partial D = \gamma$, D does not contain any closed path relators of degree $n > |\gamma|$.*

Proof. To construct the diagram D discussed in Theorem 1.27 we start with an arbitrary diagram D' with boundary γ.

Following the construction of D from D' it is clear that, if D' contains a closed path relator of some degree m, then so does D. Since D cannot contain any closed path relators of degree $n > |\gamma|$ then neither could D'. □

1.5 Spherical diagrams in $G(S)$

Here we will discuss the existence of aspherical presentations for $G(S)$ for some choices of S. Proving this for $G(\varnothing)$ is vital for showing all $G(S)$ are of type FP in Section 3. We will also discuss some obstructions as to why we are not able to do this for all S.

We first describe a new presentation for $G(S)$. We require the following definition which we will discuss more carefully in Section 2.

Definition 2.5. Let Γ be a graph. A **standard set of cycles** in Γ is a set of closed paths generating the fundamental group of Γ, such that there is some spanning forest (a tree for each connected component of Γ) $T \subseteq \Gamma$ such that, when T is contracted to a point, the set of generating cycles is exactly the set of edges that remain.

Let $\mathcal{P}'(S)$ be the subpresentation of $\mathcal{P}(S)$ where, instead of taking $\mathcal{C}$ to be all closed paths in K^1 that do not self-intersect, we take a standard set of cycles in K^1.

Proposition 1.29. *The induced subpresentation of $\mathcal{P}'(S)$ for $H(S)$ is aspherical.*

Proof. The induced subpresentation for $H(S)$ is exactly the presentation that gives us the results regarding properties of $H(S)$ as described in Corollary 2.23; one of which is the asphericity of the presentation. $\square$

Proposition 1.30. *Any van Kampen-reduced spherical diagram in $\mathcal{P}'(S)$ contains a closed t_P-corridor.*

Proof. Let D be a spherical diagram in $\mathcal{P}'(S)$. There are certainly no open t_P-corridors in D since a sphere has no boundary component. Suppose there are also no closed t_P-corridors, then D is a van Kampen-reduced spherical diagram in the induced subpresentation for $H(S)$ which, as discussed above, is aspherical and gives our contradiction. $\square$

This allows us to significantly simplify our problem.

Lemma 1.31. *Let $\mathcal{Q}$ be a subpresentation of $\mathcal{P}'(S)$ and let $\mathcal{Q}_P$ be the induced subpresentation for G_P. If $\mathcal{Q}_P$ is aspherical for each polygon P in K then $\mathcal{Q}$ is aspherical.*

Proof. Suppose there is a van Kampen-reduced spherical diagram in $\mathcal{Q}$, then by Proposition 1.30 there is a closed t_P-corridor and so, via Corollary 1.20, both enclosures of this t_P-corridor can be considered as diagrams in $\mathcal{Q}_P$. Since the corridor itself is also made up of relators in $\mathcal{Q}_P$ we are done. $\square$

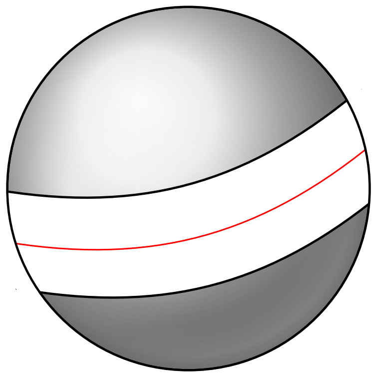

FIGURE 3.22: A spherical diagram in some subpresentation $\mathcal{Q}$.

We will first consider the case $S = \emptyset$. Let $\mathcal{Q}^0(\emptyset)$ be the subpresentation of $\mathcal{P}'(\emptyset)$ consisting of all generators, all polygon relators of degree 0 and all t_P-relators. The corresponding induced subpresentation $\mathcal{Q}^0_P(\emptyset)$ is

$$\left\langle a_1, a_2, \ldots, a_{l_P} \,\middle|\, a_1 a_2 \cdots a_{l_P},\ a_i^{t_P} = a_i^k \ :\ i = 1, 2, \ldots, l_P \right\rangle$$

where $a_1 a_2 \cdots a_{l_P}$ is the word read around ∂P. The case $n = 0$ in Proposition 1.32 is enough to show that, via Lemma 1.31, the presentation $\mathcal{Q}^0(\emptyset)$ for $G(\emptyset)$ is aspherical.

Proposition 1.32. *For all $n \in \mathbb{N}_0$, the presentation*

$$\left\langle a_1, a_2, \ldots, a_{l_P} \mid a_1^{k^n} a_2^{k^n} \cdots a_{l_P}^{k^n},\, a_i^{t_P} = a_i^k \,:\, i = 1, 2, \ldots, l_P \right\rangle$$

is aspherical.

Proof. Suppose D is a van Kampen-reduced spherical diagram in the presentation above. Since there are no open t_P-corridors in D we can construct a diagram D' in H_P where we replace each closed t_P-corridor and one of its enclosures by a copy of the enclosure where we alter the degrees of the polygon relators as in Proposition 1.21.

Since the induced subpresentation for H_P is aspherical, this D' must not be van Kampen-reduced. Let r and r' be a pair of relators in D' that are not van Kampen-reduced.

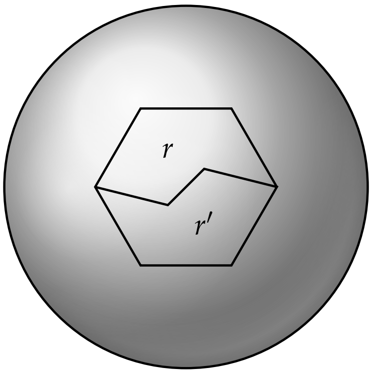

FIGURE 3.23: A pair of not van Kampen-reduced relators in a spherical diagram.

Each is exactly the polygon relator P^m for some m. In the construction of D' from D each of these relators must have been constructed from a polygon relator of degree n surrounded by $|m-n|$ closed t_P-corridors. If $m = n$ then there are a pair of polygon relators in D that are not van Kampen-reduced. If $m \neq n$ then there is at least one pair of t_P-relators in D that are not van Kampen-reduced. Either way no such D can exist and our presentation for G_P is aspherical. $\square$

One extra conclusion of the statement above tells us that the presentation $\mathcal{Q}^n(S)$ (defined in the canonical way by instead taking the polygon relators of degree n) is also an aspherical presentation for $G(\emptyset)$, we therefore have infinitely many aspherical presentations for $G(\emptyset)$. Moreover each polygon in K could be assigned a different degree if desired, this however does not increase the cardinality mentioned above. A more useful application of the generality of Proposition 1.32 is that is allows us to consider the cases where $S = \{n\} \subseteq \mathbb{N}_0$.

Choose some $n \in \mathbb{N}_0$ and let $\mathcal{Q}^n(\{n\})$ be the subpresentation of $\mathcal{P}'(\{n\})$ consisting of all generators, the closed path relators of degree n and all t_P-relators. The corresponding subpresentation $\mathcal{Q}_P^n(\{n\})$ is exactly the presentation in Proposition 1.32.

Corollary 1.33. *When* $|S| \leq 1$ *we have an aspherical presentation for each* $G(S)$.

Proof. The group $G(\emptyset)$ has aspherical presentation $\mathcal{Q}^0(\emptyset)$ and $G(\{n\})$ has aspherical presentation $\mathcal{Q}^n(\{n\})$. □

When $|S| \geq 2$ we run into an immediate problem regarding the van Kampen-reduced spherical diagram in $\mathcal{P}(S)$ that follows.

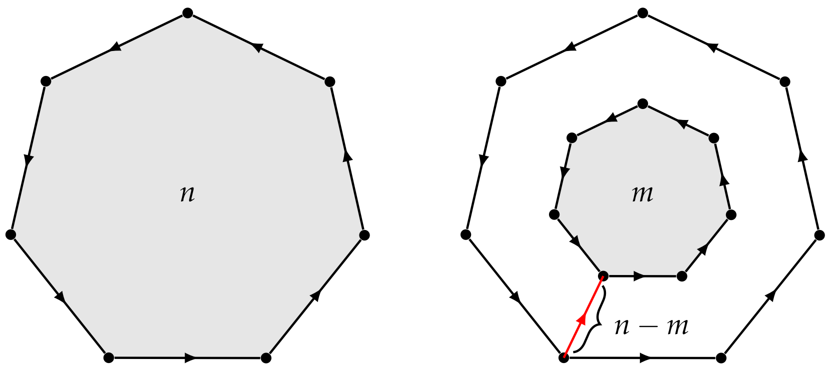

FIGURE 3.24: Choosing some $m, n \in S$ with $m < n$ we form a sphere by attaching the two disks diagrams above via their equivalent boundaries. We observe that there are exactly $\frac{k^n - k^m}{k-1}$ of each t_P-relator in this sphere.

A choice of S where we could hope to avoid such a sphere is $S = \{0, 1\}$. In this case the corresponding sphere above only contains one of each of the t_P-relators, so removing one t_P-relator from our presentation for each polygon P in K gives a candidate for an aspherical presentation.

Let $\mathcal{R}(\{0,1\})$ be the subpresentation of $\mathcal{P}'(\{0,1\})$ where we include the closed path relators of degree 0 and 1 and, for each polygon P, we include all but one of the t_P-relators.

Proposition 1.34. *Suppose* D *is a van Kampen-reduced spherical diagram in the induced subpresentation* $\mathcal{R}_P(\{0,1\})$ *for some polygon* P*. Every enclosure of a closed* t_P*-corridor that does not contain a closed* t_P*-corridor must contains a polygon relator of degree* 0 *and a polygon relator of degree* 1.

Proof. Let e be the edge in ∂P corresponding to the 'missing' t_P-relator.

Choose some 'minimal' closed t_P-corridor T in D as above and let D' be an enclosure of T that does not contain any t_P-relators. Since D' is made up of polygon relators the edge e certainly appears in D', however e cannot appear in the boundary $\partial D'$ since it is not in any of our t_P-relators.

Suppose there are no polygon relators of degree 1 in D'. Consider some polygon relator of degree 0, it has a single copy of the edge e in its boundary. Since this edge is not in the boundary of D it must be incident to some other polygon relator of degree 0. However, if two polygon relators of degree 0 share an edge then, since they each only contain one copy of e, they are not van Kampen-reduced, a contradiction.

Next suppose there are no polygon relators of degree 0. Let Γ_e be a connected subgraph of the 1-skeleton of D' such that Γ_e is maximal in the following sense:

- each edge in Γ_e is labelled e;
- there are no edges labelled e in D' that are incident to Γ_e.

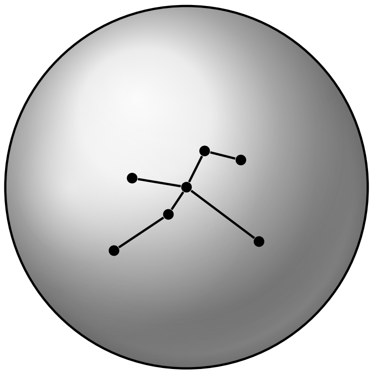

FIGURE 3.25: A tree of edges labelled e in a spherical diagram.

The graph Γ_e must be a tree else the group element e would have finite order in $H(S)$, which is torsion free by Corollary 2.23, and Γ_e must be finite since D' is finite. Consider some edge in Γ_e that is attached to a leaf (vertex of degree 1) in Γ_e. The two polygon relators incident to this edge must be aligned, else the tree would not be maximal, and so they are not van Kampen-reduced, a contradiction.

Since D' cannot be empty we see that D' must contain polygon relators of both degrees 0 and 1. $\square$

Corollary 1.35. *Any subdiagram D' as above must be incident to the short side of the corresponding closed t_P corridor.*

Proof. If it were incident to the long side then, by Corollary 1.20, it cannot contain any polygon relators of degree 0. $\square$

We now have an understanding regarding what can happen at the short side of a closed t_P-corridor. We now move on to consider what can happen regarding the long side of a t_P-corridor.

Proposition 1.36. *Let T be a closed t_P-corridor in D. There are no polygon relators incident to the long side of T.*

Proof. From the subdiagram of D enclosed by the long side of T one can construct a disk diagram D' in H_P as in the proof of Proposition 1.32. By Corollary 1.20, D' cannot contain any polygons of degree 0.

Now, suppose r is a polygon relator incident to the long side of T in D. The boundary ∂r contains the edge e. Recreate the graph Γ_e as above and again consider the edge incident to some leaf as in the proof of Proposition 1.34. The two polygon relators incident to this edge must have degree 1 in D' and, since there is no t_P-relator with an edge labelled e, both of these polygons must correspond to aligned polygon relators in D. Hence D is not van Kampen-reduced. $\square$

Proposition 1.37. *The presentation $\mathcal{R}_P(\{0,1\})$ is aspherical.*

Proof. Suppose D is a van Kampen-reduced spherical diagram in $\mathcal{R}_P(\{0,1\})$ and let D' be the corresponding spherical diagram in H_P as constructed in Proposition 1.32. Again, since the corresponding presentation for H_P is aspherical, this D' must be not van Kampen-reduced. Let r and r' by a pair of relators that are no van Kampen-reduced in D', they both correspond to some polygon relator of degree m.

In D each of these relators corresponds to either a polygon of degree 0 or a polygon of degree 1, potentially with some closed t_P-corridors around it. If they both have some closed t_P-corridors around them then there is a pair of t_P-relators in D that are not van Kampen-reduced. If neither have any closed t_P-relators around them then they correspond to a pair of polygon relators that are not van Kampen-reduced. Finally, if one of them corresponds to a polygon relator of degree 1 and the other corresponds to a polygon relator of degree 0 with a single t_P-corridor around it then we have a contradiction via Proposition 1.36. $\square$

Corollary 1.38. *When $S = \{0,1\}$ we have an aspherical presentation for $G(S)$.*

Proof. The presentation $\mathcal{R}(\{0,1\})$ is aspherical and, since we are able to reconstruct the missing t_P-relator using the remainder of the sphere as in Figure 3.24, it presents $G(S)$. $\square$

For other choices of S we cannot use this trick, since the sphere in Figure 3.24 used above to rebuild the missing t_P-relator contains more than one copy of each t_P-relator. Moreover, we currently have no grasp on whether an aspherical presentation is possible for other choices of S. We note that these other presentations may have completely different generating sets, so proving non-existence is extremely difficult.

Open Question 1.39. For what other choices of $S \subseteq \mathbb{N}_0$ are there aspherical presentations for $G(S)$?

When S is finite we are able to show that $G(S)$ is at least finitely presented.

Proposition 1.40. *If K is suitable and S is finite then $G(S)$ is finitely presentable.*

Proof. As mentioned, we only need P^0 and all the t_P-relators in order for the word read around some P^n to be trivial.

Removing all polygon relators of degree $n > 0$ from $\mathcal{P}'(S)$ results in a presentation for $G(S)$ containing finitely many generators and exactly

$$\sum l_P + 1 + |S|$$

relators, where the sum is over the polygons in K. The l_P is the t_P-relators, the 1 is the polygon relator of degree 0 and the $|S|$ is the degrees of some closed path relator. The $|S|$ follows from the fact that, since K is acyclic, any standard set of cycles is the same size as the set of polygons.

It follows that when S is finite we have a finite presentation for $G(S)$. $\square$

2 Graphical small cancellation theory

In this section we will give an overview of graphical small cancellation theory and its application to our groups. This is a generalisation of classical small cancellation theory and is the main idea behind the inception of the groups $G(S)$.

2.1 A group from a graph

There are many ways in geometric group theory to pass between graphs and groups preserving various data. In this section we will describe a way to pass from a labelled directed graph to a group via a group presentation.

Definition 2.1. A **labelled directed graph** Γ is a directed graph such that each edge is labelled by some element of a labelling set X.

Example 2.2. Let Γ be the following labelled directed graph labelled by $X = \{a, b, c\}$.

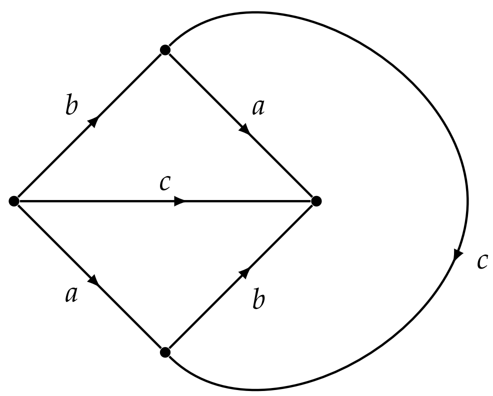

We will consider words read around closed loops in the graph, we therefore use the convention that if we read a directed edge labelled x backwards then we read the letter x^{-1}.

Definition 2.3. Given a labelled directed graph as above we define the group H_Γ. It is generated by the set of labels that appear in Γ and has relators defined by the set of words read around closed paths in Γ.

Example 2.4. From the graph in Example 2.2 we find the group presentation

$$\langle a, b, c \mid ab = c, ba = c, ca = b \rangle$$

which presents the group $\mathbb{Z}_2 \times \mathbb{Z}$.

Above we observe a presentation with three relators which we have deduced from a graph with certainly more than three cycles. The following definition helps to formalise this idea.

Definition 2.5. Let Γ be a graph. A **standard set of cycles** in Γ is a set of closed paths generating the fundamental group of Γ, such that there is some spanning forest (a tree for each connected component of Γ) $T \subseteq \Gamma$ such that, when T is contracted to a point, the set of generating cycles is exactly the set of edges that remain.

We note that the may be multiple choices of a standard set of cycles, each depending on the choice of spanning forest, however each will contain the same number of cycles: the number of connected components of Γ minus its Euler characteristic.

Example 2.6. A standard set of cycles for the graph above can be represented by the boundaries of the grey regions below. It can be realised by choosing the spanning tree to consist of the blue edges.

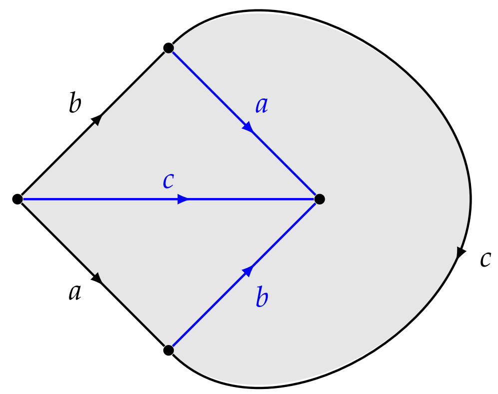

FIGURE 3.26: The resulting standard set of cycles contains cycles reading bac^{-1}, abc^{-1} and cba^{-1}.

Lemma 2.7. *Any group presentation can be defined via some labelled directed graph using Definition 2.3.*

Proof. For each relator r in the presentation take a cycle graph reading the word r. The disjoint union is a labelled directed graph that defines the corresponding group. □

Lemma 2.7 suggests that perhaps this construction is too general to be of any use. However, given some restrictions on the graphs we are able to prove some strong results.

2.2 Graphical small cancellation theory

In Section 1 we discuss the classical metric small cancellation condition for group presentations. In this section we describe how one can generalise this condition via the construction in Definition 2.3.

We first need to check that we have chosen our graph carefully.

Definition 2.8. A labelled directed graph Γ is **reduced** if there is no vertex v such that there is a pair of edges whereby they read the same letter when starting at v.

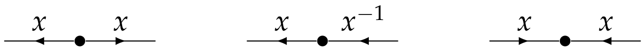

FIGURE 3.27: Each of these situations is not permitted in a reduced graph.

Definition 2.9. Let Γ be a reduced labelled directed graph. A **piece** in Γ is a path that embeds in Γ with two different initial vertices. We will often write the word read along a piece to denote it.

Example 2.10. The graph Γ in Example 2.2 is certainly reduced. We will now consider pieces in Γ:

- Each of a, b and c are clearly pieces as each of them appear more than once in Γ.

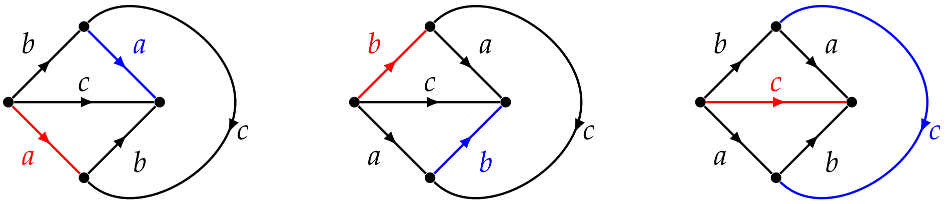

FIGURE 3.28: Pieces of length 1 in Γ.

- We observe that there are no pieces of length greater than 1 that contain the label b.
- There are several pieces that consist of edges labelled only by a and c: ac^{-1} and $a^{-1}c$ of length 2 (observing that we need not count the inverse of each piece), $ac^{-1}a$ and $ca^{-1}c$ of length 3 and $ac^{-1}ac^{-1}$ and $ca^{-1}ca^{-1}$ of length 4.

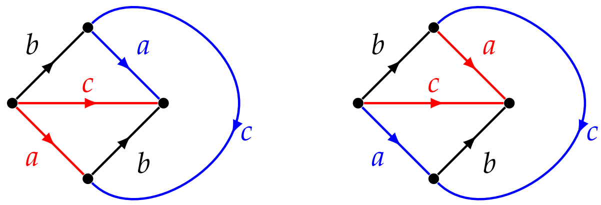

FIGURE 3.29: Pieces of length 2 in Γ.

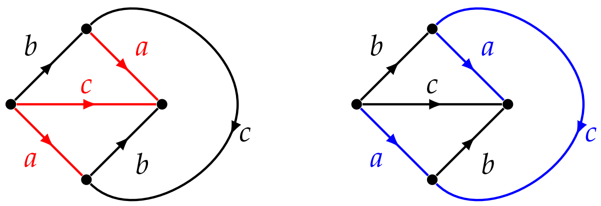

FIGURE 3.30: The two embeddings of $ac^{-1}a$ in Γ.

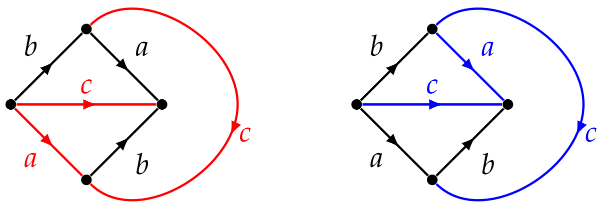

FIGURE 3.31: The two embeddings of $ca^{-1}c$ in Γ.

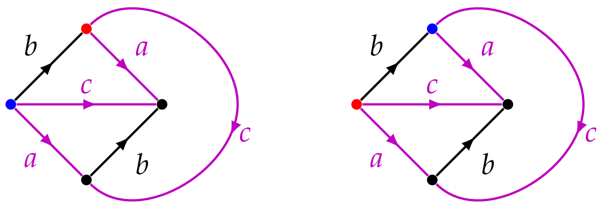

FIGURE 3.32: On the left we represent the two embeddings of $ac^{-1}ac^{-1}$, one starting at the red vertex, the other at the blue. On the right we similarly represent the embeddings of $ca^{-1}ca^{-1}$.

- There can be no pieces of length greater than 4 since there are no cycles of such lengths that embed in Γ.

Definition 2.11. A reduced labelled directed graph Γ satisfies $C'(\lambda)$ if whenever p is a piece contained within some cycle γ then

$$|p| < \lambda|\gamma|.$$

The condition is known as a graphical small cancellation condition and we get the following useful theorem of Gromov (12) which can be compared with the classical small cancellation theory results. Ollivier (21) should be used as a background resource on graphical small cancellation theory.

Definition 2.12. We say a graph is **non-filamentous** if there are no edges that do not appear in any embedded cycle.

Theorem 2.13 (Gromov). *Let Γ be a reduced non-filamentous labelled directed graph. Let X be the set of labels that appear on Γ and let R be the set of words read around a standard set of cycles in Γ. Let g be the girth of Γ (the length of the shortest cycle).*

If Γ satisfies $C'(\frac{1}{6})$ then the presentation $\mathcal{P} = \langle A \,|R\rangle$ presents the group H_Γ with the following properties:

1. *It is a direct limit of hyperbolic groups (if Γ is finite then H_Γ is hyperbolic) and is torsion-free.*
2. *The presentation $\mathcal{P}$ is aspherical and so H_Γ has cohomological dimension at most 2.*
3. *If $|R| > |X|$ then H_Γ is infinite and not quasi-isometric to $\mathbb{Z}$.*
4. *The shortest relation in H_Γ has length g.*
5. *For any reduced word w representing the identity in H_Γ, some cyclic permutation of w contains a subword of some word read on a cycle embedded in Γ, of length at least $(1 - 3\lambda) > \frac{1}{2}$ times the length of this cycle.*
6. *The natural maps from each connected component of Γ into the Cayley graph of H_Γ are isometric embeddings.*

The key ideas in the proof of this statement are utilised in the proof of Proposition 1.19 where we align and perform diamond moves on relators until we have a disk diagram which we can attack with classical small cancellation theory.

We observe that this theorem is a strengthening of classical metric small cancellation theory since, given a finite presentation satisfying $C'(\frac{1}{6})$, we can construct a finite graph satisfying $C'(\frac{1}{6})$ via Lemma 2.7.

Example 2.14. We present the following graph that satisfies $C'(\frac{1}{6})$.

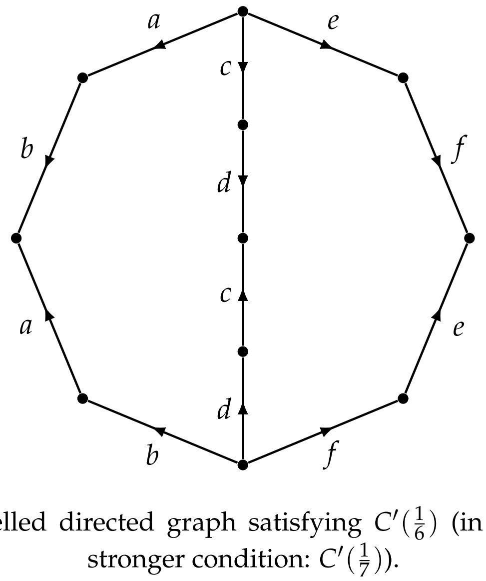

FIGURE 3.33: A labelled directed graph satisfying $C'(\frac{1}{6})$ (in fact this satisfies the stronger condition: $C'(\frac{1}{7})$).

On observation this graph has no pieces of length greater than 1 and moreover every path of length 1 is a piece. Since the shortest cycle is of length 8 we certainly satisfy $C('\frac{1}{6})$.

Next we consider the group H_Γ. We must first choose some spanning tree, say:

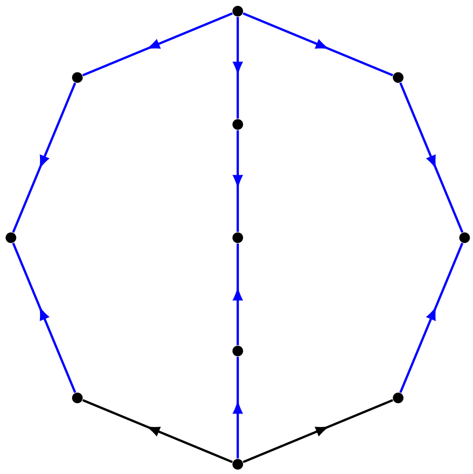

FIGURE 3.34: The blue edges represent a spanning tree in Γ.

This spanning tree results in the following presentation:

$$\left\langle a,b,c,d,e,f \mid [a,b][c,d]^{-1},[e,f][c,d]^{-1}\right\rangle.$$

This example neatly demonstrates the fact that this is a strengthening of classical small cancellation theory and shows the way in which the standard set of cycles allow us to be aspherical; if we took all three canonical relations (the cycles of length 8), we would easily be able to construct a spherical diagram.

2.3 $H(S)$ via graphical small cancellation theory

We will construct a graph $\Gamma(S)$ such that $H_{\Gamma(S)}$ is exactly $H(S)$; this is the original way in which $H(S)$ was defined. The graph will be a disjoint union

$$\bigcup_{n\in\mathbb{N}_0}\Gamma_n,$$

the structure of Γ_n depending on whether or not $n \in S$. If $n \notin S$ then the graph will reflect the polygons of K and if $n \in S$ the graph will reflect the 1-skeleton of K. This is analogue to the 'polygon relators' and 'closed path relators' as in Section 1.

We will define two sequences of graphs, $(\Gamma_n^\gamma)_{n\in\mathbb{N}_0}$ and $(\Gamma_n^P)_{n\in\mathbb{N}_0}$. The graph Γ_n will either be Γ_n^γ or Γ_n^P, depending on whether or not $n \in S$. We must first consider the

1-skeleton K^1. We make this graph a labelled directed graph by assigning to each edge an arbitrary direction and a unique label.

Definition 2.15. The labelled directed graph Γ_n^γ is the graph defined by replacing every edge in K^1 by a path of length k^n, each edge in which bearing the same label and direction as the initial edge.

The superscript γ refers to analogy between this and the closed path relators γ^n.

Example 2.16. We consider the following (not-suitable) 2-complex from Section 3 to demonstrate the idea.

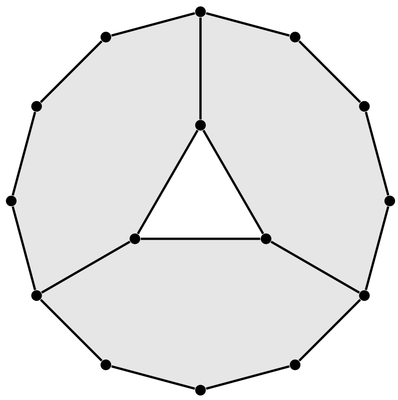

Letting $k = 2$, the following graphs represent Γ_0^γ and Γ_3^γ.

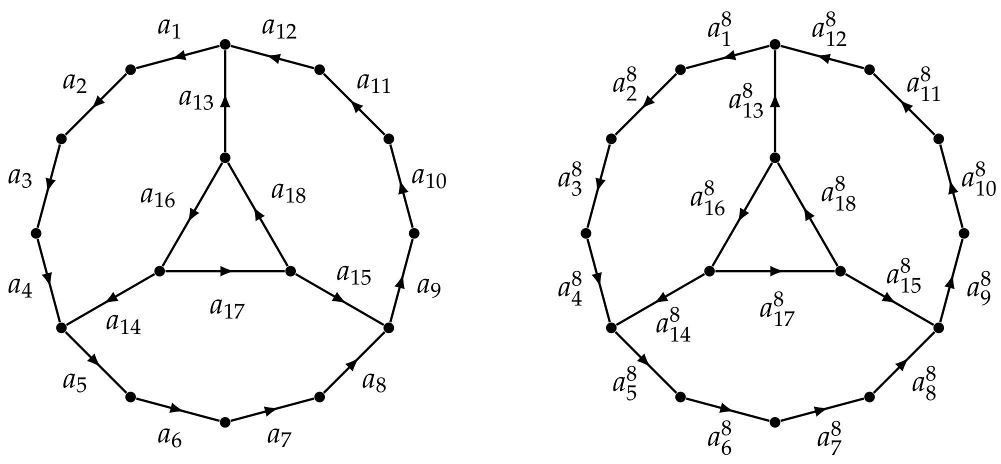

FIGURE 3.35: We compare Γ_0^γ and Γ_3^γ. The edges labelled a_i^8 are assumed to be paths of length 8 with each edge labelled a_i.

We first consider the small cancellation condition of a single graph Γ_n^γ.

Proposition 2.17. *If K is suitable then the graph Γ_n^γ satisfies $C'(\frac{1}{g})$ where g is the girth of K^1.*

Proof. The longest pieces are of the form $a_i^{k^n-1}$ and so have length $k^n - 1$. The shortest closed path γ' in Γ_n^γ has length gk^n, so we always have $|p| < \frac{1}{g}|\gamma'|$. $\square$

When K is suitable we have $g \geq 13$. So Γ^{γ}_{n} certainly satisfies $C'(\frac{1}{13})$ and so the critical $C'(\frac{1}{6})$.

Definition 2.18. The labelled directed graph Γ^{P}_{n} is a disjoint union of cycle graphs: one for each polygon in K. The word read around the cycle corresponding to P is exactly the word read around the closed path in K^1 defined by the boundary of P. We then replace every edge in each cycle by a path of length k^n, each edge in which bearing the same label and direction as the initial edge.

Example 2.19. Using the same complex as in Example 2.16 above we present the following graph Γ^{P}_{2}.

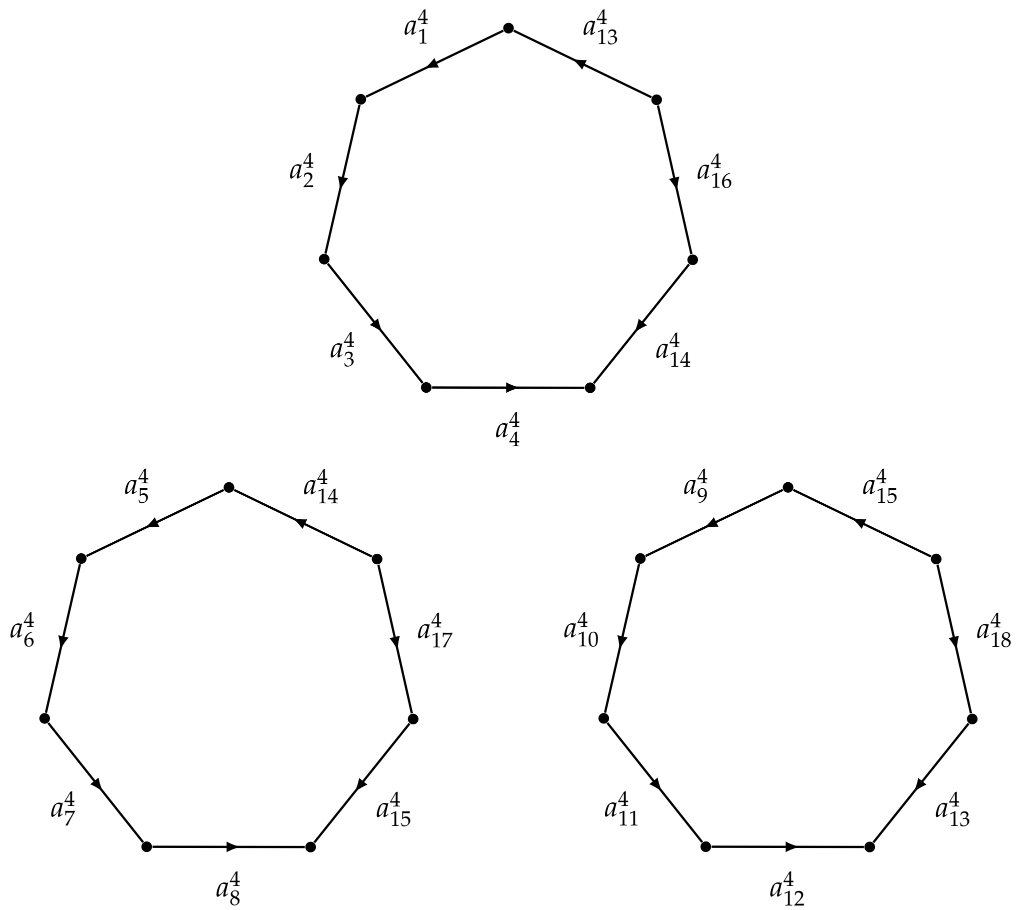

FIGURE 3.36: The graph Γ^{P}_{2} is a disjoint union of three cycle graphs.

Proposition 2.20. *If K is suitable then the graph Γ^{P}_{n} satisfies $C'(\frac{1}{6})$.*

Proof. The longest pieces are exactly the length of the longest pieces when considering the interaction between the polygons in K. So the graph Γ^{P}_{n} satisfies the same small cancellation condition as the complex K. □

Now we define the graph $\Gamma(S)$.

Definition 2.21. Let $S \subseteq \mathbb{N}_0$ be some subset. We define the graph $\Gamma(S)$ to be the infinite disjoint union of graphs as follows:

$$\Gamma(S) = \bigcup_{n \in \mathbb{N}_0} \begin{cases} \Gamma_n^{\gamma} & n \in S \\ \Gamma_n^{P} & n \notin S \end{cases}$$

Theorem 2.22. *If K is suitable then the graph $\Gamma(S)$ satisfies $C'(\frac{1}{6})$.*

Proof. By Proposition 2.17 and Proposition 2.20 we see that when searching for pieces that appear twice within the disjoint subgraphs of $\Gamma(S)$ we satisfy $C'(\frac{1}{6})$. It remains to compare two arbitrary subgraphs Γ_m and Γ_n, we drop the superscripts without loss of generality and assume that $m < n$. The path $a_i^m a_j^m$, of length $2m$ appears in both of these graphs and is the longest such piece that does so. Since the girth of these graphs is certainly at least gk^m, where g is the girth of K, we observe that $|p| < \frac{2}{g}|\gamma|$ and so, since $g \geq 13$, $\Gamma(S)$ always satisfies $C'(\frac{1}{6})$. □

We now plug $\Gamma(S)$ into Theorem 2.13 to summarise results about $H(S)$.

Corollary 2.23. *If K is suitable then $H(S)$ with presentation falling from $\Gamma(S)$ has the following properties:*

1. *$H(S)$ is a direct limit of hyperbolic groups and is torsion free.*
2. *We have an aspherical presentation for $H(S)$ and so $H(S)$ has cohomological dimension 2.*
3. *$H(S)$ is infinite and not quasi-isometric to $\mathbb{Z}$.*
4. *The shortest relation in $H(S)$ is of length g (the girth of K^1).*
5. *For any reduced word representing the identity in $H(S)$, some cyclic permutation of w contains a subword of some word read on a cycle embedded in Γ, of length at least $\frac{1}{2}$ times the length of this cycle.*
6. *The natural map from each connected component of $\Gamma(S)$ into the Cayley graph of $H(S)$ are isometric embeddings.*

Proof. The cohomological dimension of a group is 1 if and only if G is a free group. Since $H(S)$ is clearly never free the cohomological dimension of $H(S)$ is 2. Since $\Gamma(S)$ satisfies $C'(\frac{1}{6})$ the rest of these claims follow immediately from Theorem 2.13. □

Example 2.24. Consider the not-suitable complex from Section 4 with girth 15. Let $S = 2\mathbb{N}_0 = \{0, 2, 4, \dots\}$. The first four components of $\Gamma(S)$ can be viewed (up to subdivisions) as the graph below.

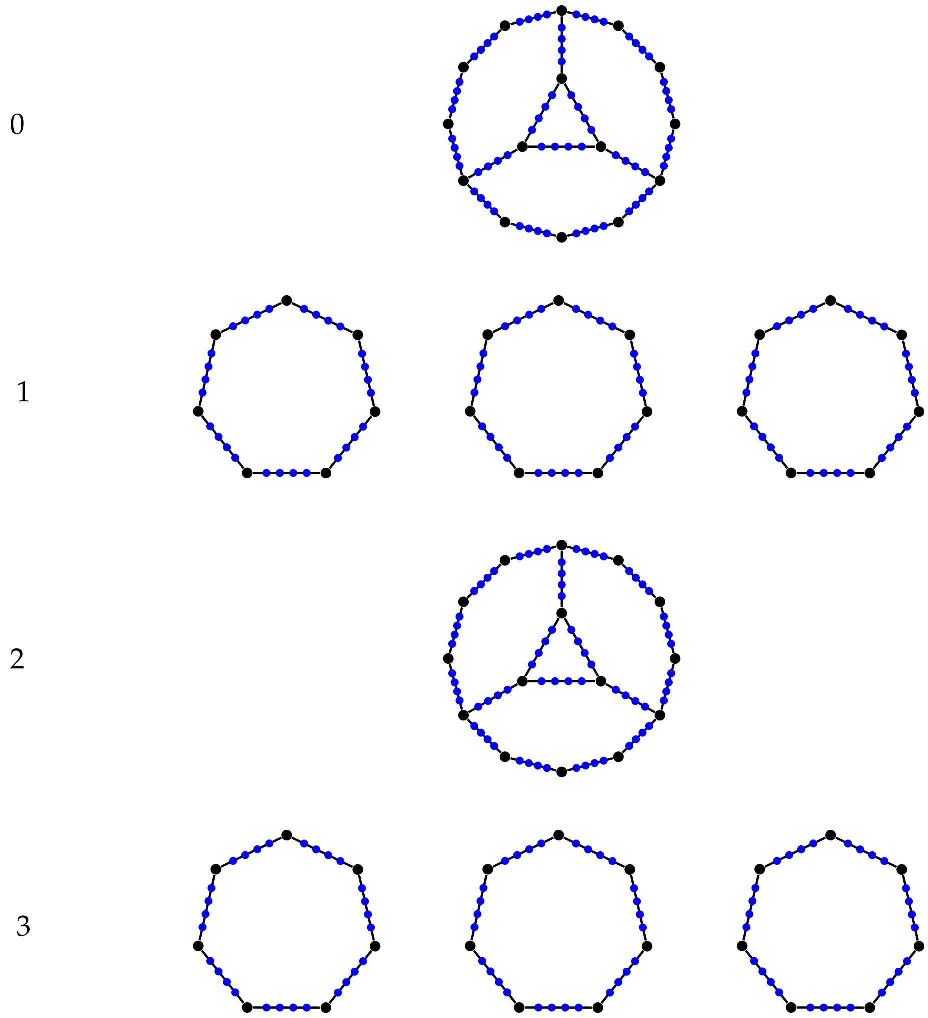

FIGURE 3.37: A portion of the graph $\Gamma(S)$.

We observe here that the only difference between choosing whether or not some $n \in \mathbb{N}_0$ is in S is the addition of the relator of degree n that corresponds to the central cycle of length $15k^n$.

When K is suitable, and so more complicated than our example here, the difference between Γ_n^γ and Γ_n^p will be more drastically different. We do observe that, since the boundary of each polygon is a cycle in the 1-skeleton of K, the difference in relations is down to words read around the cycles in K^1 that generate the fundamental group of K. This is why we require that $\pi_1(K)$ is non-trivial in constructing our suitable complexes; else the group $H(S)$, and in turn $G(S)$, would be invariant under changing S and so would all be finitely presented.

3 Finiteness properties of $G(S)$

In this section we discuss some interesting finiteness properties of the group $G(S)$. This was the main motivation for the inquest into these groups. We will culminate with the following theorems.

Theorem 3.32. The group $G(S)$ is type FP, and has cohomological dimension 2, for all $S \subseteq \mathbb{N}_0$.

Theorem 3.34. The group $G(S)$ is finitely presentable if and only if S is finite.

3.1 Eilenberg-MacLane spaces

Eilenberg-MacLane spaces can be associated to any group and witness properties of the group in their topology.

Definition 3.1. An **Eilenberg-MacLane space** for a group G is a connected CW-complex with fundamental group isomorphic to G and contractible universal cover. We use $K(G, 1)$ to denote an Eilenberg-MacLane space for a group G.

Example 3.2. We construct an Eilenberg-MacLane space for the free group on n generators as follows. Start with a single vertex v and attach an edge from v to v for each generator. This has fundamental group isomorphic to F_n and universal cover an infinite tree of valence $2n$, which is contractible.

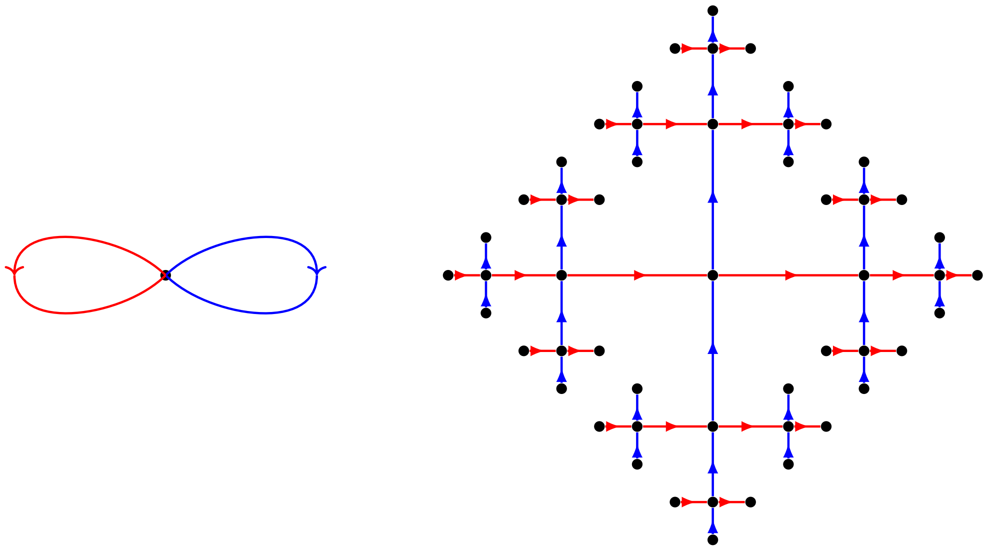

FIGURE 3.38: Here we observe an Eilenberg-Mac Lane space for F_2 aside a portion of its universal cover. As we will see later this space is actually a presentation 2-complex for F_2 aside a portion of the corresponding Cayley 2-complex.

When $n = 1$ we have $F_1 = \mathbb{Z}$ and we find that the circle is a $K(\mathbb{Z}, 1)$.

Example 3.3. The infinite dimensional real projective space $\mathbb{R}P^\infty$ admits a CW-structure with a cell in every dimension. This is an Eilenberg-MacLane space for the cyclic group of 2 elements.

In fact any Eilenberg-MacLane space for a non-trivial finite group has at least one cell in infinitely many dimensions.

3.2 Presentation 2-complexes

A useful source of Eilenberg-MacLane spaces are presentation 2-complexes.

Definition 3.4. Let $\mathcal{P} = \langle A \mid R \rangle$ be a presentation. We construct the corresponding **presentation 2-complex** as follows:

- Start with a single vertex v;
- For each generator $a \in A$ attach a directed edge from v to v labelled a;
- For each relator $r \in R$ attach a 2-cell along the unique path of edges that reads the word r.

Not all presentation 2-complexes are Eilenberg-MacLane spaces. For example, take some presentation of a non-trivial finite group, the resulting presentation 2-complex has no cells of dimension $n > 2$, but, as discussed, any Eilenberg-MacLane space for a finite group has cells in infinitely many dimensions.

We do have a criterion for when such a space exists. We recall Definition 1.16 which tells us a presentation $\mathcal{P}$ is **aspherical** if there are no van Kampen-reduced spherical diagrams in $\mathcal{P}$.

Proposition 3.5. *If a group presentation $\mathcal{P}$ for a group G is aspherical then the corresponding presentation 2-complex is an Eilenberg-MacLane space for G.*

Proof. Let X be the presentation 2-complex of $\mathcal{P}$ and let $\tilde{X}$ be its universal cover. The fundamental group is certainly G so we need only check that $\tilde{X}$ is contractible. We observe that there are no subspaces of $\tilde{X}$ that are homeomorphic to a 2-sphere, else they would define van Kampen-reduced spherical diagrams in $\mathcal{P}$. Hence $\tilde{X}$ is contractible and so X is an Eilenberg-Mac Lane space for G. $\square$

3.3 Graphs of groups

One can carefully combine Eilenberg-MacLane spaces in order to produce Eilenberg-MacLane spaces for new groups. Here we will discuss how graphs of groups can be used to do this.

Definition 3.6. Let Γ be a directed graph. Assign to every vertex $v \in \Gamma$ a group G_v and assign to every edge, incident to vertices u and v, a group $H_{u,v}$ such that $H_{u,v}$ surjects into both G_u and G_v. Such a structure is called a **graph of groups**.

$$G_u \xrightarrow{H_{u,v}} G_v$$

FIGURE 3.39: A simple graph of groups with two vertices and one edge.

The **fundamental group of a graph of groups** is defined as follows. Let T be a spanning tree for some graph of groups Γ. Let $\pi_1(\Gamma)$ be the group generated by the vertex groups G_v along with a generator t_e for each directed edge $e \in \Gamma$. We have the following additional relations:

- if e is directed from u to v then we have

  $$t_e \phi_u(h) t_e^{-1} = \phi_v(h)$$

  for all $h \in H_{u,v}$, where ϕ_u and ϕ_v are the surjections from $H_{u,v}$ into G_u and G_v respectively;

- $t_e = 1$ if the edge e is in T.

We present examples that highlight two common forms of graphs of groups.

Example 3.7. If Γ has a single vertex and a single edge then it is called an **HNN-extension**.

For example, let $G_v \cong \mathbb{Z}$ and let $H_{v,v} \cong \mathbb{Z}$ with embeddings $\phi_v(h) = g^2$ and $\psi_v(h) = g^3$.

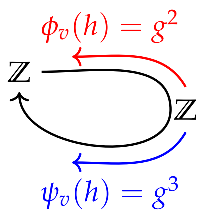

FIGURE 3.40: The coloured lines are not part of the graph, here they are used to draw attention to the embeddings.

In this case $\pi_1(\Gamma)$ has presentation

$$\left\langle g, t \mid tg^2t^{-1} = g^3 \right\rangle$$

and so is isomorphic to the Baumslag-Solitar group $BS(2,3)$.

Example 3.8. If Γ consists of two vertices joined by a single edge then it is called a **free product with amalgamation**.

This time let G_u, G_v and $H_{u,v}$ each be a copy $\mathbb{Z}$ and let $H_{u,v}$ embed via the maps $\phi_u(h) = g_u^2$ and $\psi_v(h) = g_v^3$.

$$\mathbb{Z} \xleftarrow{\phi_u(h) = g_u^2} \mathbb{Z} \xrightarrow{\psi_v(h) = g_v^3} \mathbb{Z}$$

The fundamental group of this graph of groups has presentation

$$\langle x, y \mid x^2 = y^3 \rangle$$

and is therefore trefoil knot group (the fundamental group of $\mathbb{R}^3 - K$ where K is the trefoil knot).

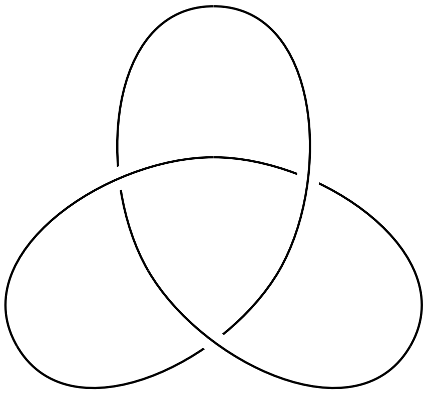

FIGURE 3.41: A representation of the trefoil knot where line breaks represent overlappings.

Example 3.8 uses the same groups and embeddings as in Example 3.7 but the resulting groups are different due to the different underlying graphs.

3.4 Graphs of spaces

We consider an equivalent route to getting hold of the fundamental group of a graph of groups.

Definition 3.9. Let Γ be a graph of groups. One can construct a **graph of spaces** by taking a $K(G_v, 1)$ for each vertex in Γ and attaching them together using the direct product $K(H_{u,v}, 1) \times [0, 1]$ for the corresponding edge group $H_{u,v}$. We use the embeddings of $H_{u,v}$ into G_u and G_v to define the attachment. The resulting space is an Eilenberg-MacLane space for the fundamental group of Γ.

Example 3.10. We can realise a $K(G,1)$ for the trefoil group by expanding on Example 3.8. For a $K(\mathbb{Z},1)$ we take the circle. So for G_u and G_v we have a circle and for $H_{u,v}$ we take the direct product of the circle and an interval: a cylinder.

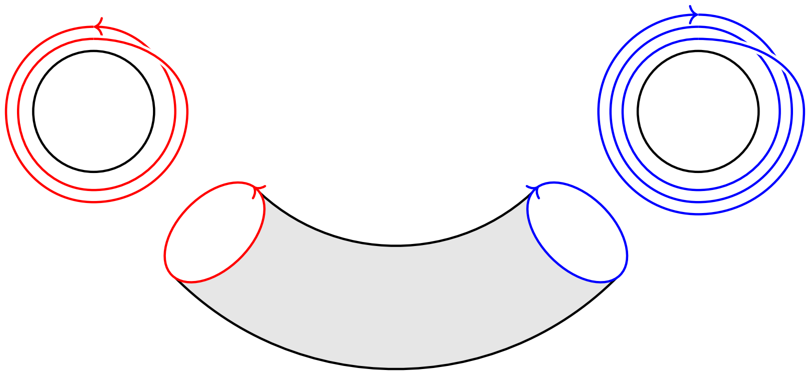

FIGURE 3.42: We represent the construction of a $K(G,1)$ for the trefoil knot group. We attach the ends of a cylinder to two different circles in two different ways: the red end we wrap twice around the left circle and the blue end we wrap three times around the right circle.

We use the maps ϕ_u and ϕ_v to tell us how to attach the cylinder to the two circles. By contracting some path through the cylinder to a point we actually realise the presentation 2-complex corresponding to the group above. This shows that the presentation $\langle x, y \mid x^2 = y^3 \rangle$ is aspherical.

3.5 Eilenberg-MacLane spaces for H_P, G_P, $H(S)$ and $G(S)$

We start with H_P, recalling the presentation

$$\mathcal{P}_{H_P} = \left\langle a_1, a_2, \ldots, a_{l_P} \mid a_1^{k^n} a_2^{k^n} \cdots a_{l_P}^{k^n} \ : \ \forall n \in \mathbb{N}_0 \right\rangle .$$

Proposition 3.11. *There is a* 2*-dimensional, infinite* $K(H_P, 1)$*.*

Proof. Via Proposition 1.13 we observe that the infinite presentation above satisfies $C'(\frac{1}{3d})$ and so $C'(\frac{1}{6})$. This presentation is therefore aspherical and so the presentation 2-complex X is a $K(H_P, 1)$.

This X has a single vertex, l_P edges and a 2-cell for every $n \in \mathbb{N}_0$. $\square$

Next we consider G_P, recalling the presentation

$$\left\langle a_1, a_2, \ldots, a_{l_P}, t_P \mid a_1^{k^n} a_2^{k^n} \cdots a_{l_P}^{k^n} \ : \ \forall n \in \mathbb{N}_0, \ \ t_P a_i t_P^{-1} = a_i^k \ : \ \forall a_i \in P \right\rangle ,$$

which clearly does not satisfy $C'(\frac{1}{6})$.

Proposition 3.12. *There is a* 2*-dimensional, finite* $K(G_P, 1)$.

Proof. In Proposition 1.32 we present a family of aspherical presentations for G_P. Each offers a 2-dimensional, finite $K(G_P, 1)$ via a presentation 2-complex. □

In Section 2 we discuss an aspherical presentation for $H(S)$.

Proposition 3.13. *There is a* 2*-dimensional, infinite* $K(H(S), 1)$.

Proof. There is an infinite aspherical presentation for $H(S)$ as defined via graphical small cancellation theory in Corollary 2.23. This allows us to construct a 2-dimensional, infinite $K(H(S), 1)$ via a presentation 2-complex. □

Eilenberg-MacLane spaces for $G(S)$ are more complicated and utilise a graph of groups. We will use $H(S)$ and G_P as vertex groups and each H_P as an edge group.

Proposition 3.14. *The group* H_P *embeds in both* $H(S)$ *and* G_P.

Proof. The set of generators of H_P is a subset of the set of generators of both $H(S)$ and G_P. So it suffices to show, in each case, that H_P is exactly the subgroup generated by this subset.

We begin with $H(S)$. In Corollary 1.20 we see that any word in $H(S)$ written in the generators of H_P is trivial in H_P. This completes our claim.

For G_P we consider Proposition 1.21. This tells us that for any trivial word in $G(S)$ that is written in the generators of $H(S)$ is trivial in $H(S)$. Then, since we are actually working within G_P, this word is actually written in generators of H_P and so we can again use Corollary 1.20 to show that this word is trivial in H_P. □

Now consider the following graph of groups.

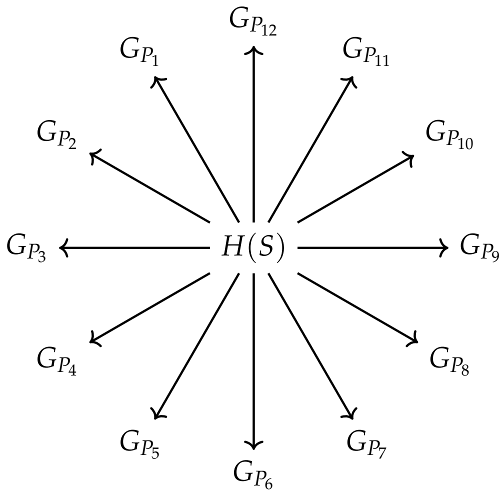

FIGURE 3.43: A graph of groups defining $G(S)$, each edge is assumed to be labelled H_{P_i} for the corresponding P_i.

Proposition 3.15. *The graph of spaces corresponding to this graph of groups is a 3-dimensional $K(G(S), 1)$ for all $S \subseteq \mathbb{N}_0$.*

Proof. We read the presentation for the fundamental group of this graph of groups. The graph is a tree so we have no new generators and further this becomes a series of free products with amalgamation. This presentation is exactly $\mathcal{P}(S)$ as in Section 1.

The graph of spaces corresponding to this graph is made up of 2-dimensional Eilenberg-MacLane spaces. When we construct the space $H_P \times [0, 1]$ to attach G_P to $H(S)$ we introduce 3-cells, so the resulting space is 3-dimensional. $\square$

3.6 Geometric finiteness properties

Above we mentioned whether or not certain groups can have finite Eilenberg-MacLane spaces. Here we go into more details on the various properties that a $K(G, 1)$ have. For a full background see (7), (3) and (2, Section 1).

Definition 3.16. A group G is **type** F if there is a $K(G, 1)$ with finitely many cells.

By Proposition 3.12 we see that G_P is type F.

Definition 3.17. A group G is **type** FH if it acts freely on an acyclic CW-complex with only finitely many orbits of cells.

Type F and type FH are known as topological finiteness properties.

Proposition 3.18. *A group of type F is of type FH.*

Proof. Let X be the universal covering of a finite $K(G,1)$. This space is acyclic since it is contractible and the action of G has only finitely many orbits of cells since the $K(G,1)$ has only finitely many cells. □

Definition 3.19. The **geometric dimension** of a group G is the minimal dimension of a $K(G,1)$.

As above, $\mathbb{Z}$ has geometric dimension 1 and all finite groups have infinite geometric dimension.

3.7 Geometric dimension of $G(S)$

Theorem 3.20. *If $S = \emptyset$, $|S| = 1$ or $S = \{0,1\}$ then $G(S)$ has geometric dimension 2 and is type F.*

Proof. We construct presentation 2-complexes for $G(S)$ in each of these cases via the aspherical presentations described in Section 1.5. Since each of these presentations is finite the group is type F. □

For other choices of S the geometric dimension of $G(S)$ is unknown.

3.8 Algebraic finiteness properties

The section is suggestively brief since we will not rely on any of the definitions later. For a comprehensive background see (7) or (18).

Definition 3.21. Let G be a group. The **group ring** of G, denoted $\mathbb{Z}G$, is the ring consisting of finite sums of the form

$$\sum_i a_i g_i$$

where $a_i \in \mathbb{Z}$ and $g_i \in G$. Multiplication and addition are defined canonically.

Definition 3.22. A group G is **type** FL if the trivial module $\mathbb{Z}$ for the group ring $\mathbb{Z}G$ admits a finite resolution by finitely generated free $\mathbb{Z}G$-modules.

Proposition 3.23. *A group of type FH is of type FL.*

Proof. The cellular chain complex corresponding to any free acyclic G-CW-complex provides a finite resolution as required. □

Definition 3.24. A group G is **type** FP if the trivial module $\mathbb{Z}$ for the group ring $\mathbb{Z}G$ admits a finite resolution by finitely generated projective $\mathbb{Z}G$-modules.

Type FL and FP are known as algebraic finiteness properties.

Proposition 3.25. *A group of type FL is of type FP.*

Proof. This follows immediately since a free module is itself projective. □

We have the amassed following chain of implications:

$$F \implies FH \implies FL \implies FP.$$

The groups $G(S)$ are not even finitely presented when S is infinite, let alone type F, so proving that all $G(S)$ are type FP is a surprising and significant result.

Definition 3.26. The **cohomological dimension** of a group G is the minimal length of a finite resolution for $\mathbb{Z}$ by finitely generated projective $\mathbb{Z}G$ modules.

Since the cell complex associated to the universal cover of a $K(G,1)$ gives us such a resolution we have

$$cd(G) \leq gd(G).$$

Eilenberg-Ganea showed that $cd(G) = gd(G)$ expect possibly $cd(G) = 2$ and $gd(G) = 3$.

This is expected to be false and the groups $G(S)$ do throw up candidates for counter examples as discussed in Open Question 3.33.

3.9 A criterion for being type FP

The main tool our results rely on is as follows. We say a group G is acyclic if it has an acyclic $K(G,1)$.

Proposition 3.27. *Suppose G is type F and N is an acyclic normal subgroup of G. Then the group G/N is type FP and the cohomological dimension of G/N is bounded above by the geometric dimension of G.*

Proof. Let X be the universal cover of a $K(G,1)$ and consider the quotient X/N. This is an Eilenberg-MacLane space for N, equipped with a free cellular action of G/N. Since N is acyclic so is X/N. The G/N-orbits of cells in X/N correspond to the G-orbits of cells in X and since G is of type F there are only finitely many of these orbits. So G/N is of type FH and so type FP.

If X has dimension n then so does X/N, the geometric dimension of X/N is an upper bound on the cohomological dimension of X/N. □

3.10 Maps between $H(S)$ and $H(T)$

We let $E(S)$ denote the universal cover of the $K(H(S),1)$ from Proposition 3.13.

For $S \subset T \subseteq \mathbb{N}_0$, let $K_{S,T}$ be the kernel of the surjection

$$H(S) \to H(T)$$

induced by the natural bijection between generating sets. We will compare the spaces $E(T)$ and $E(S)/K_{S,T}$ by constructing a third space F. In $E(S)/K_{S,T}$, for each $n \in T - S$ there is a free $H(T)$-orbit of subspaces homeomorphic to the complex K. We construct F from $E(S)/K_{S,T}$ by attaching a free $H(T)$-orbit of cones to each of these subspaces.

Proposition 3.28. *The inclusion $E(S)/K_{S,T} \to F$ is an $H(T)$-equivariant homology isomorphism.*

Proof. This follows immediately since the homology of K is trivial, so coning off subspaces homeomorphic to K will not change the homology. □

Proposition 3.29. *There is an inclusion $E(T) \to F$ that is an $H(T)$-equivariant deformation retraction.*

Proof. Any standard set of cycles on K is contractible and moreover is contained as a deformation retract in the cone on K. Performing this retraction simultaneously over each of the subspaces homeomorphic to K as above completes the claim. □

Theorem 3.30. *For each $S \subset T \subseteq \mathbb{N}_0$ the natural bijection between generating sets extends to a surjective group homomorphism $H(S) \to H(T)$, whose kernel $K_{S,T}$ is a non-trivial acyclic group.*

Proof. Since the generating sets are the same, the induced map is certainly a surjection. We observe that the kernel is non-trivial by considering Corollary 2.23 and observing that, for each $n \in T - S$ and come closed path in K^1 reading $a_1 a_2 \ldots a_l$, the element $a_1^{k^n} a_2^{k^n} \cdots a_l^{k^n}$ is non-trivial in $H(S)$ but is trivial in $H(T)$.

It remains to show that $K_{S,T}$ is acyclic. By Proposition 3.28, $E(S)/K_{S,T}$ is an Eilenberg-MacLane space for $K_{S,T}$ and has the same homology as F. Our claim then follows since F is contractible as it has the same equivariant homotopy as $E(T)$ by Proposition 3.29. □

3.11 Finiteness properties of $G(S)$

We are almost ready to apply Proposition 3.27; we have a collection of groups each with a host of normal acyclic subgroups. The remaining issue is that none of the $H(S)$ are type F since they are not even finitely presented. This is where $G(S)$ comes in to play.

Proposition 3.31. *For each $S \subset T \subseteq \mathbb{N}_0$, the kernel of the homomorphism $G(S) \to G(T)$ induced by the bijection on generating sets is acyclic.*

Proof. We will show that the kernel is a countable free product of copies of the acyclic group $K_{S,T}$. Let X_S and X_T be the Eilenberg-MacLane spaces for $G(S)$ and $G(T)$ as described above. The universal covers of these spaces are therefore both trees of spaces, with each edge and vertex spaces contractible.

Now, an Eilenberg-MacLane space for the kernel can be constructed by pulling back the universal cover of X_T along the map

$$F : X_S \to X_T.$$

The pullback along f also has the structure of a tree of spaces. Each vertex space corresponding to $H(S)$ is an Eilenberg-MacLane space for $K_{S,T}$ by Theorem 3.30 and each other vertex and edge space remains contractible. Hence, the kernel for the map $G(S) \to G(T)$ is a countable free product of acyclic groups and is thus acyclic. □

Theorem 3.32. *The group $G(S)$ is type FP, and has cohomological dimension 2, for all $S \subseteq \mathbb{N}_0$.*

Proof. Consider the surjection $G(\emptyset) \to G(S)$ induced by the bijection between generating sets. Let $N \leq G(\emptyset)$ be the kernel of this map, which is acyclic by Theorem 3.30. Then, Proposition 3.27, the group $G/N \cong G(S)$ is type FP and has cohomological dimension 2. □

Since $G(S)$ always has cohomological dimension 2 and we have only been able to show that the geometric dimension of $G(S)$ is 2 for a tiny portion of the choices of S, this leaves open many candidates for counter examples to the Eilenberg-Ganea conjecture.

Open Question 3.33. Is the geometric dimension of $G(S)$ always 2?

We next move on to consider when $G(S)$ is finitely presentable.

Theorem 3.34. *The group $G(S)$ is finitely presentable if and only if S is finite.*

Proof. Let $\mathcal{Q}(S)$ be the presentation described in Proposition 1.40.

If S is finite $\mathcal{Q}(S)$ is a finite presentation for $G(S)$.

Now, let S be infinite and suppose there is a finite presentation for $G(S)$, we may assume (via a re-writing map) that this presentation has the same generating set as $\mathcal{Q}(S)$. Let R be the finite set of relators in this presentation and let $r \in R$ be the relator of longest length. Consider some finite set of van Kampen-reduced disk diagrams that make up each relator in $\mathcal{Q}(S)$.

Since each of these diagrams has boundary length at most $|r|$, by Lemma 1.28 none of them contain any relators of degree $n > |r|$. Let $S_{|r|}$ be the subset of S consisting of $n \in S$ such that $n \leq |r|$. Every relator $r \in R$ is a consequence of some relators in $\mathcal{Q}(S_{|r|})$.

Now, let $a_1 a_2 \cdots a_l$ be a word read around a closed path in K^1 of minimal length. By Corollary 1.24 the word

$$a_1^{k^n} a_2^{k^n} \cdots a_l^{k^n}$$

is trivial in $G(S)$ if and only if $n \in S$. However, if $G(S)$ is finitely presented then we have shown that

$$a_1^{k^n} a_2^{k^n} \cdots a_l^{k^n}$$

is trivial if and only if $n \in S_{|r|}$, a contradiction. $\square$

4 Uncountably many quasi-isometry classes of groups

There are uncountably many choices of $S \subseteq \mathbb{N}_0$. A natural question is to consider how this reflects in the set

$$\{G(S) \mid S \subseteq \mathbb{N}_0\}.$$

In (8) a set valued invariant is discussed and is used to show that the set contains uncountably many isomorphism classes of groups. This section is motivated by the work in (23) where it is shown that Leary's families of groups (18) each contain uncountably many quasi-isometry classes of groups, the methods used there do not apply to the groups discussed here.

In this section we will strengthen this via the following theorem.

Theorem 4.20. If K is very-suitable and k is sufficiently large (depending on K) then the set $\{G(S) \mid S \subseteq \mathbb{N}_0\}$ contains uncountably many quasi-isometry classes.

The corresponding result for the uncountable family of groups of type FP in (18) was proved in (23). They use completely different methods that are not applicable here.

Proving two groups are not quasi-isometric is often not easy; quasi-isometries respect the large scale geometry of an object, ignoring small-scale details. In particular one can show that all finite groups are quasi-isometric to the trivial group. This highlights the fact that a collection of uncountably many quasi-isometry classes is particularly difficult to produce.

4.1 Quasi-isometries of graphs

We will be considering a group as a metric space via a Cayley graph, so we first consider quasi-isometries within the framework of graphs. Let Γ be a graph and let Γ^0 be its vertex set. Let d_Γ be the combinatorial distance function on Γ^0, so two vertices are distance d apart if the shortest path between them is of length d.

Definition 4.1. Let Γ, Λ be graphs and let $k \in \mathbb{N}$. A map $\phi : \Gamma^0 \to \Lambda^0$ is k**-Lipschitz** if, for all $u, v \in \Gamma^0$, we have

$$d_\Lambda(\phi(u), \phi(v)) \leq k d_\Gamma(u, v).$$

A 1-Lipschitz map between graphs must send pairs of adjacent vertices to either a pair of adjacent vertices or the same vertex. Higher values of k allow maps to send nearby vertices further and further apart. In particular, a k-Lipschitz map is k'-Lipschitz for all $k' > k$.

Definition 4.2. A pair of graphs Γ and Λ are ***k*-quasi-isometric** if there are k-Lipschitz maps $\phi : \Gamma^0 \to \Lambda^0$ and $\psi : \Lambda^0 \to \Gamma^0$ such that

$$d_\Gamma(u, \psi \cdot \phi(u)) \leq k$$

for all $u \in \Gamma^0$ and

$$d_\Lambda(v, \phi \cdot \psi(v)) \leq k$$

for all $v \in \Lambda^0$.

A pair of graphs are **quasi-isometric** if they are k-quasi-isometric for some $k \in \mathbb{N}$.

Proposition 4.3. *All finite connected graphs are quasi-isometric.*

Proof. It suffices to show that all finite connected graphs are quasi-isometric to a point.

Let Γ be a finite connected graph and let Λ be a single point. Let $\phi : \Gamma \to \Lambda$ map every vertex to the single point in Λ and let $\psi : \Lambda \to \Gamma$ map the single vertex to some arbitrary vertex v_Γ in Γ. Let k be the number of vertices in Γ.

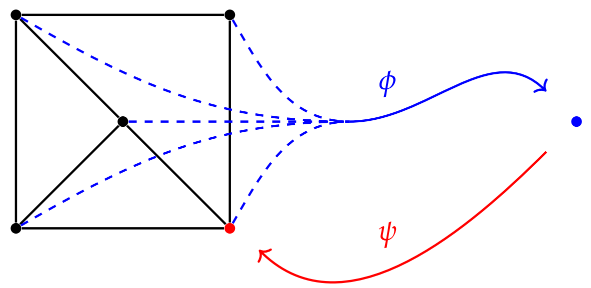

FIGURE 3.44: We map all vertices in a graph to a single point via ϕ and map the single point back arbitrarily via ψ.

Both ϕ and ψ are k-Lipschitz; in fact they are both 1-Lipschitz since the images of any pair of points are always the same point. Now, let $u \in \Gamma$,

$$d_\Gamma(u, \psi \cdot \phi(u)) = d_\Gamma(u, v_\Gamma) \leq k$$

as there are only k vertices in Γ. For the other claim let $v \in \Lambda$,

$$d_\Lambda(v, \phi \cdot \psi(v)) = 0 \leq k$$

as there is exactly one vertex in Λ.

So Γ and Λ are k-quasi-isometric and hence quasi-isometric. $\square$

4.2 Quasi-isometries of groups

We now reflect these ideas under the framework of groups.

Definition 4.4. A pair of finitely generated groups G and H are k**-quasi-isometric** if and only if there are k-quasi-isometric Cayley graphs Γ and Λ, each with respect to some finite generating set, corresponding to G and H respectively.

For this definition to be well defined one has to check that any pair of Cayley graphs for some group G are themselves quasi-isometric.

Proposition 4.5. *Let X and Y be finite generating sets for a group G. The Cayley graphs $\Gamma = Cay(G, X)$ and $\Lambda = Cay(G, Y)$ are quasi-isometric.*

Proof. Let $\phi : \Gamma \to \Lambda$ be a re-writing map that maps each generator $x \in X$ to some word in $Y \cup Y^{-1}$ that represents the same element in G. Let $\psi : \Lambda \to \Gamma$ be defined similarly. Let k be some integer greater than the length of the longest word in $\phi(X) \cup \psi(Y)$. By symmetry we need only consider one direction.

Let g, h be elements of G, and so vertices of Γ. The distance $d = d_\Gamma(g, h)$ is realised by a path of some length d, or otherwise, a word w of length d in $X \cup X^{-1}$ such that $g \cdot w = h$. When we map to Λ, the path of length d maps to a path of length at most kd in $Y \cup Y^{-1}$, so

$$d_\Lambda(\phi(g), \phi(h)) \leq k d_\Gamma(g, h).$$

To show that $d_\Gamma(g, \psi \cdot \phi(g)) \leq k$ we observe that $\psi \cdot \phi(g) = g$, so

$$d_\Gamma(g, \psi \cdot \phi(g)) = 0 \leq k$$

and we conclude that Γ and Λ are k-quasi-isometric. □

Our claim about finite groups follows immediately from Proposition 4.3 since any finite group has a finite Cayley graph.

4.3 Bowditch's taut loop spectra

In (5) the author presents an invariant with which to distinguish quasi-isometry classes of graphs.

Definition 4.6. Let Γ be a graph and let γ be a loop in Γ. A **net** for γ is a subgraph σ of Γ that can be embedded into a disk D with $\partial D = \gamma$.

We let $M(\sigma)$ be the length of the shortest loop in σ.

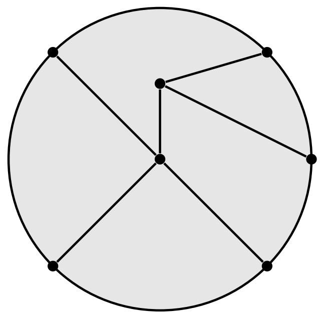

FIGURE 3.45: A graph embedded into a disk.

Definition 4.7. Let γ be a loop in Γ. We let $H(\gamma)$ be the smallest value of $M(\sigma)$ as σ ranges over all nets for γ. A loop γ is **taut** if $|\gamma| = H(\gamma)$.

Example 4.8. Let Γ_n be the standard Cayley graph for $\mathbb{Z}^n$ for some $n \geq 2$. The taut loops in Γ_n consist of every loop of length 4. This follows from the fact that there are no shorter loops and any longer loop is a consequence of these loops of length 4.

The taut loop spectrum, as defined in (5), associates a subset of the natural numbers to the taut loops in a graph.

Definition 4.9. The **taut loop spectrum** of a graph Γ is the set

$$H(\Gamma) = \{H(\gamma) \mid \gamma \text{ is a loop in } \Gamma\}.$$

Example 4.10. The embedded graph Γ from Figure 3.45 has $H(\Gamma) = \{3, 4\}$.

Example 4.11. The graphs Γ_n from Example 4.8 each have $H(\Gamma_n) = \{4\}$.

Any subset of the natural numbers can be realised by a graph.

Example 4.12. Let $H = \{h_1, h_2, \dots\} \subseteq \mathbb{N}$ be an arbitrary subset. Construct a graph Γ_H by attaching $|H|$ cycle graphs, each having length h_n respectively, in a chain attached via a single vertex.

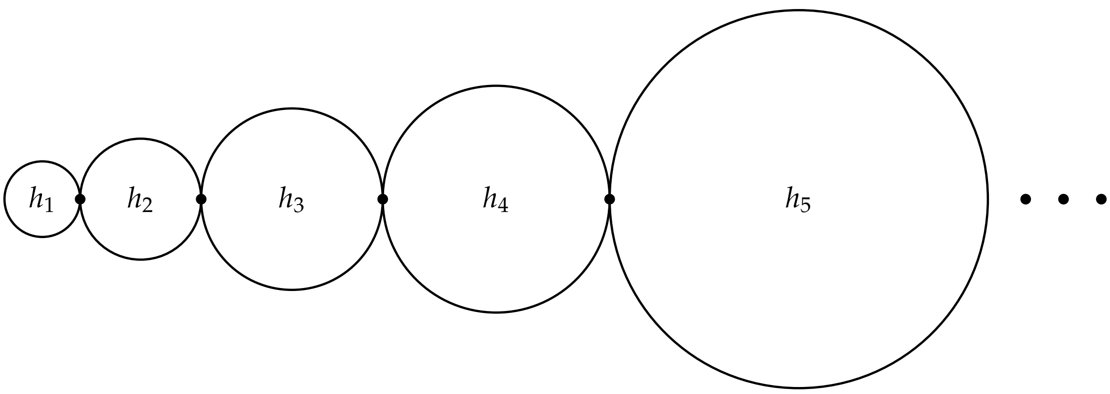

FIGURE 3.46: An infinite graph Γ with $H(\Gamma) = H$.

One can see that the set of taut loops is exactly the set of cycle graphs used in the construction and so $H(\Gamma_H) = H$.

One can construct a group realising a Cayley graph with the same condition by using relators of the prescribed lengths such that no generators appear twice. Moreover, so long as $\min(H) > 6$ one can force such a group to be finitely generated using small cancellation theory.

We wish to use these spectra to distinguish quasi-isometry classes. The following definition suggests the method of doing so.

Definition 4.13. We say that two subsets H and H' of the natural numbers are k**-related** if, given any $L \in H$ with $L > k^2 + 2k + 1$, there is some $L' \in H'$ with

$$L/k \leq L' \leq kL,$$

and conversely, swapping H and H'.

Example 4.14. Similarly to before, all finite subsets are k-related for some k.

The following lemma (5, Lemma 2) completes the link.

Lemma 4.15. *Suppose that connected graphs Γ and Λ are k-quasi-isometric, then the sets $H(\Gamma)$ and $H(\Lambda)$ are k-related.*

4.4 The taut loop spectra of $G(S)$

We are unable to compute the taut loop spectra of $G(S)$ precisely. Instead we present a set of intervals within $\mathbb{N}$ in which lengths of taut loops are shown to exist. Since the definition of k-related sets relies on inequalities, this gives us enough to show we have uncountably many quasi-isometry classes.

Here we fix some very-suitable complex K, with subdivision constant d and pre-subdivision vertex set L^0 and some $k \geq |L^0|d$. Let $S \subseteq \mathbb{Z}$ and let Γ be the Cayley graph of $G(S)$ corresponding to the presentation $\mathcal{P}(S)$. We begin by defining $|S| + 1$ intervals via a pair of sequences $\{x_n\}$, $\{y_n\}$ indexed by the set $\{-1\} \cup S$. We have

$$x_n = n$$

and

$$y_n = \begin{cases} k+3 & \text{if } n = -1, \\ |L^0|dk^n & \text{if } n \geq 0. \end{cases}$$

We define our intervals to be $[x_n, y_n]$ for each $n \in \{-1\} \cup S$. We note that some of our intervals may overlap. We also note that both x_n and y_n are strictly increasing.

We will first prove that we have a taut loop with length in $[x_n, y_n]$ for each $n \in \{-1\} \cup S$. Then we will show that there are no taut loops with lengths outside of these intervals.

Proposition 4.16. *Let γ be the boundary of a t_P-relator in Γ. Then $H(\gamma) \in [x_{-1}, y_{-1}]$.*

Proof. Since $|\gamma| = k + 3$, any net σ for γ must have $M(\sigma) \leq k + 3$ and so $H(\gamma) \in [x_{-1}, y_{-1}]$. □

Proposition 4.17. *Let γ be the boundary of a closed path relator γ of degree $n \in S$. Then $H(\gamma) \in [x_n, y_n]$.*

Proof. Suppose there is a net σ with $M(\sigma) < n$. Let γ' be a loop in σ that realises the value $M(\sigma)$. From σ we are able to construct a van Kampen-reduced disk diagram D in $\mathcal{P}(S)$, containing the relator r, with $\partial r = \gamma$ and with $\partial D = \gamma'$.

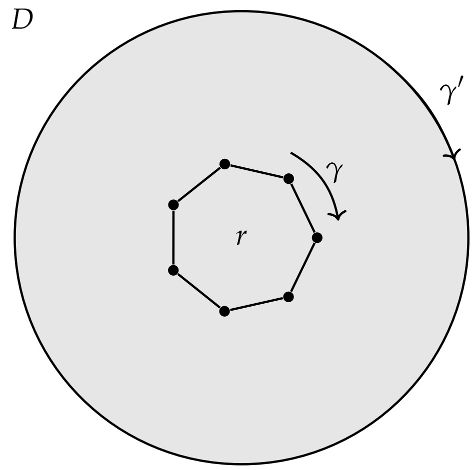

FIGURE 3.47: A van Kampen-reduced disk diagram with boundary length $M(\sigma)$ that contains a relator with boundary word γ.

By Lemma 1.28, since $|\partial D| < n$, D cannot contain any closed path relators of degree n, a contradiction.

For the upper bound we observe that, since $|\gamma| \leq |L^0|dk^n$, any net σ must have $M(\sigma) \leq |L^0|dk^n$ and so we are done. □

We observe that the upper bound can be easily strengthened by constructing a net consisting of as many t_P-corridors as one can attach to the boundary of γ and a loop in the middle with length some linear function of n depending on the structure of K. Since we have made no assumptions on the structure of K it is cleaner to leave our upper bound as it is.

It is also likely that the lower bound can be strengthened for some choices of K and k. Combined these may allow us to show that we have $x_n = y_n$ and compute the taut loop length exactly.

It remains to show that we have no taut loops outside of these intervals.

Proposition 4.18. *Let γ be a taut loop in Γ. Then $|\gamma| \in [x_n, y_n]$ for some $n \in \{-1\} \cup S$.*

Proof. Let D be an arbitrary disk diagram with $\partial D = \gamma$. If a closed path relator of degree n appears in D then, assuming n is the maximal such degree, $|\gamma| \in [x_n, y_n]$ as above. So assume there is no such relator. In this case we can construct a net using t_P-relators and polygon relators of degree 0. Since the length of these polygon relators is less than $|L^0|d \leq k$, the largest relator in D therefore has length at most $k+3$, and so we have constructed a net σ for γ with $M(\sigma) \in [x_{-1}, y_{-1}]$. □

4.5 Uncountably many quasi-isometry classes

We wish to show that

$$\{G(S) \mid S \subseteq \mathbb{N}_0\}$$

contains uncountably many quasi-isometry classes of groups. If we can show it is true for a subset then we are done. Let $Z \subset \mathbb{N}_0$ be an infinite subset of $\mathbb{N}_0$ such that, when $m < n \in Z$, we have $y_m^2 < x_n$. This is certainly possible because x_n and y_n tend to infinity and $x_n \leq y_n$.

Theorem 4.19. *Let $S, T \subseteq Z$ be sets. If the groups $G(S)$ and $G(T)$ are quasi-isometric then the symmetric difference of S and T is finite.*

Proof. If $G(S)$ and $G(T)$ are quasi-isometric then they are k-quasi-isometric for some $k \in \mathbb{N}$. We have computed bounds for the taut loop spectra above, let us compare these lengths in the language of being k-related. Let $L \in [x_m, y_m]$ and $L' \in [x_n, y_n]$ for some $m < n$. We have

$$L'/L \geq x_n/y_m \geq y_m.$$

If $G(S)$ and $G(T)$ are quasi-isometric then $H(G(S))$ and $H(G(T))$ are k-related for some $k \in \mathbb{N}$. By Definition 4.13, k must be at least $L'/L \geq y_m$ for all m in the symmetric difference of S and T. We deduce that the symmetric difference of S and T is finite. □

Theorem 4.20. *If K is very-suitable and k is sufficiently large (depending on K) then the set $\{G(S) \mid S \subseteq \mathbb{N}_0\}$ contains uncountably many quasi-isometry classes.*

Proof. Consider the set of all subsets of Z, this is an uncountable set. We can split this into equivalence classes by saying that subsets $S, T \subseteq Z$ are equivalent if and only if their symmetric difference is finite. Since each equivalence class contains only countably many sets, there are uncountably many equivalence classes and, by Theorem 4.19, this means there are uncountably many quasi-isometry classes in

$$\{G(S) \mid S \subseteq Z\}$$

which is, in turn, a subset of the required set.

Finally, since there are only countably many ways to construct a finite $K(G, 1)$ and groups that are not quasi-isometric are certainly not isomorphic, only countable many choices of S result in groups of type F. $\square$

4.6 Bounding presentation sizes via taut loop spectra

An open problem that was pursued during this project was the relation gap problem. Due to a lack of progress made in this area we will not delve too deeply into the details.

Loosely speaking the **relation gap** measures the difference between the rank of the relation module N^{ab} of a group G with a given generating set X and the minimal number of relators required in a corresponding presentation $\langle A \mid R \rangle$. For $G(S)$ we are able to show that this rank is some constant for all $S \subseteq \mathbb{N}_0$, in fact we have

$$rank(N^{ab}) = |K^2| + \sum_P l_P$$

where $|K^2|$ is the number of polygons in K and l_P is the length of the boundary of P.

It remains to find some increasing sequence $(b_n)_{n \in \mathbb{N}_0}$ such that the minimal number of relators required to present $G(S)$ is at least $b_{|S|}$ and we are done. This is because we could choose some S large enough such that $b_n > rank(N^{ab})$. Bounding the number of relators across all presentations, even after fixing a generating set, is difficult.

Given the results on taut loops the following idea is natural. Let X be a generating set for a group G and let H be the taut loop spectrum for the corresponding Cayley graph Γ.

If $\langle A \mid R \rangle$ is a presentation for G, is $|R| \geq |H|$?

The set of taut loops gives a presentation for G and many canonical examples have this property, however the following construction breaks it down.

Example 4.21. Let $G(k, l, m)$ the group defined by the presentation

$$\left\langle a_1, a_2, \ldots a_l, t \mid a_1 a_2 \cdots a_l = 1, a_i t^m = t^m a_i^k \right\rangle .$$

We observe that $G(k, l_P, 1)$ is exactly G_P.

The group $G(k, l, m)$ has $l + 1$ relators and taut loop spectrum

$$H = \{l, lk, lk^2, \ldots, lk^n, 2m + k + 1\}$$

where n is the maximal natural number such that $lk^n < 2m + k + 1$.

Computing this taut loop spectrum is not trivial but can be achieved by considering similar arguments to our work on $G(S)$.

The words in our generators that define the taut loops are summarised below.

Word w representing 1 in $G(k, l, m)$	$\lvert w\rvert$
$a_1 a_2 \cdots a_l$	l
$a_1^k a_2^k \cdots a_l^k$	lk
$a_1^{k^2} a_2^{k^2} \cdots a_l^{k^2}$	lk^2
$a_1^{k^3} a_2^{k^3} \cdots a_l^{k^3}$	lk^3
$\vdots$	$\vdots$
$a_1^{k^n} a_2^{k^n} \cdots a_l^{k^n}$	lk^n
$t^{-m} a t^m a^{-2}$	$2m + k + 1$

We have a class of groups for which our idea does not work. However, there may be a way to classify when this does happen and further perhaps one can show that the property does hold in the case of $G(S)$.

While we expect there to be some finite choice of S such that $G(S)$ does have a finite relation gap, we do admit the following evidence against.

Corollary 4.22. *When $S = \emptyset$, $|S| = 1$ and $S = \{0, 1\}$, there is a presentation for $G(S)$ with exactly*

$$|K^2| + \sum_P l_P$$

relators.

These presentations are exactly the aspherical presentations given in Section 1.5.

5 The homological Dehn function of $G(S)$

This section is prompted by (16) where an in depth background Dehn functions and their homological generalisation is given, as well as a computation of the homological Dehn function for the groups in (18).

Theorem 5.11. If K is very-suitable and k is sufficiently large (depending on K) then the homological Dehn function satisfies

$$FA_{G(S)}(n) \simeq e^n$$

for all $S \subseteq \mathbb{N}_0$.

The equivalence above is defined as follows.

Definition 5.1. Let $f, g : [0, \infty) \to [0, \infty)$ be two functions, we say $f \preccurlyeq g$ if there are constants $A, B > 0$ and $C, D, E \geq 0$ such that $f(n) \leq Ag(Bn + C) + Dn + E$ for all $n \geq 0$. We say $f \simeq g$ if $f \preccurlyeq g$ and $g \preccurlyeq f$. We extend this equivalence relation to function $\mathbb{N}_0 \to [0, \infty)$ by assuming they are constant on each interval $[n, n + 1)$.

Dehn functions are defined on group presentations and give a measure of complexity of van Kampen-reduced disk diagrams of a given boundary length. If a group is finitely presented then any two finite presentations will realise Dehn functions that are equivalent up to $\simeq$ and so we can consider the Dehn function as a group invariant (11). If instead a group is type FP_2 (which each $G(S)$ is since they are type FP and cohomological dimension 2), but possibly not finitely presented, then the classical Dehn function may no longer be well defined for a given group. In this section we will see that, for groups of type FP_2, we are able to consider a well defined homological analogue of the Dehn function that can be considered as a group invariant.

5.1 Homological Dehn functions for groups of type FP_2

We will use the following set up which allows us to directly apply results of Section 1.

Definition 5.2. Let $\langle A \mid R \rangle$ be a presentation for a group G and let Γ be the corresponding Cayley graph. A **homological presentation** for a group G is a pair $\langle A || R_0 \rangle$ where $R_0 \subseteq R$ and $X(A, R_0)$, the 2-complex defined by attaching only the 2-cells corresponding to the relators in R_0 to Γ, has trivial first homology group. We say $X(A, R_0)$ is a **homological Cayley complex** in this case.

In a homological Cayley complex there may be loops that are not the boundary of any disk diagram. The natural homological analogue is that, since $H_1(X(A, R_0)) = 0$, for every loop γ there is a **surface diagram** in $\langle A \,||\, R_0 \rangle$ with a single boundary component reading the word read around γ. For concreteness we consider the following definition.

Definition 5.3. A **surface diagram** D in a homological presentation is cellular decomposition of a surface $S_{g,1}$ with genus g and a single boundary component. Each 2-cell in D reads the boundary of a relator $r \in R_0$ and the word read around the single boundary component ∂D reads a trivial word in $\langle A \mid R \rangle$. Let $|D|$ denote the number of 2-cells in D.

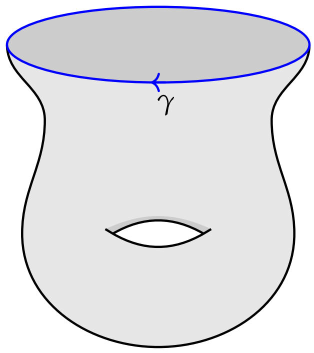

FIGURE 3.48: A surface diagram with boundary γ.

The following proposition allows us to consider finite homological presentations for groups of type FP_2.

Proposition 5.4 (Proposition 2.7 in (16))**.** *Let $G = \langle A | R \rangle$ with A finite. The group G is type FP_2 if and only if there exists a finite homological presentation $\langle A \,||\, R_0 \rangle$ for G.*

Definition 5.5. Let X be a homological Cayley 2-complex. The **homological area** of γ is

$$\mathrm{HArea}_X(\gamma) := \min\{|D| \mid D \text{ is a surface diagram with boundary } \gamma\}.$$

The **homological filling function** of X is then

$$\mathrm{FA}_X(n) := \sup\{\mathrm{HArea}_X(\gamma) \mid \gamma \text{ is a loop in } X^1 \text{ of length } \leq n\}.$$

A **homological Dehn function** of a group G with a finite homological presentation $\langle A \,||\, R_0 \rangle$ is the homological filling function of $X(A, R_0)$.

Proposition 5.6 (Proposition 2.24 in (16))**.** *Up to $\simeq$ equivalence, the homological Dehn function is independent of a homological finite presentation for a group G.*

So, for groups of type FP_2, the homological Dehn function is well defined.

5.2 Homological Dehn function for $G(S)$

The key observation here is that, since $H_1(K) = 0$ and K is finite, for each closed path in K^1 there is a surface diagram made up of polygons in K with the closed path as its boundary. Let M be some constant such that, for every closed path γ in K^1 that does not self-intersect, there is a surface diagram D with $\partial D = \gamma$ and $|D| \leq M$.

Let $\langle A|R \rangle$ be the presentation $\mathcal{P}(S)$ and consider $R_0 \subseteq R$ consisting of only the polygon relators of degree 0 and the t_P-relators.

Proposition 5.7. *The* 2*-complex* $X(A, R_0)$ *is a homological Cayley complex for* G.

Proof. We consider each relation in $R - R_0$ and argue that each can be realised as the boundary of a surface in $X(A, R_0)$.

We first consider the polygon relator P^n with some degree $n > 0$. This can be realised by a disk (surface with genus 0) diagram by attaching n closed t_P-corridors around the polygon relator P^0.

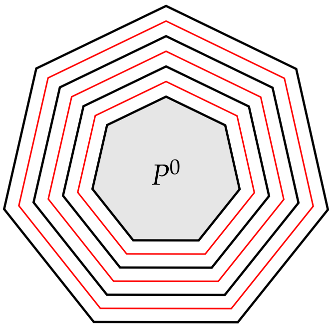

FIGURE 3.49: We realise a surface diagram for P^3 of genus 0.

Next we consider a closed path relator γ^n for some $n \in S$. This corresponds to a closed path γ' in K^1 which is the boundary of some surface made of polygons of K. As such we can construct a surface diagram realising γ^n by taking the surface for γ', subdividing each edge into k^n edges, then replacing every polygon of degree n that appears by a disk as described above. □

We now wish to bound the number of 2-cells in these surface diagrams. We will first bound the number of 2-cells in a disk diagram in $\mathcal{P}(S)$ and then use it to deduce a bound for our surface diagrams in $\langle A||R_0 \rangle$.

Proposition 5.8. *Let* γ *be a loop in* Γ. *There is a disk diagram* D *in* $\mathcal{P}(S)$ *with* $\partial D = \gamma$ *and*

$$|D| \preccurlyeq e^{|\gamma|}.$$

Proof. This follows directly from the first part Theorem 1.27 which gives the existence of a diagram D with $\partial D = \gamma$ such that

$$|D| \leq \left(\frac{1}{2}|\gamma|^2 + |\gamma|\right) C^{\frac{|\gamma|}{2}} \preccurlyeq e^{|\gamma|}$$

where $C = |L^0|d$. $\square$

Next we replace each of the 2-cells from $R - R_0$ by surfaces of a bounded size.

Proposition 5.9. *Let γ be a loop in Γ. There is a surface diagram D in $\langle A|R_0\rangle$ with $\partial D = \gamma$ and*

$$|D| \preccurlyeq e^{|\gamma|}.$$

Proof. Let D' be the disk diagram described in Proposition 5.8. Here we use both parts of Theorem 1.27, the second of which states that there are no K-relators of degree greater than $|\gamma|$.

We replace each closed path relator γ^n by at most M polygon relators of degree n to give a diagram D'' with

$$|D''| \preccurlyeq Me^{|\gamma|} \preccurlyeq e^{|\gamma|}$$

2-cells. We then replace each polygon relator P^n by a copy of P^0 and n closed t_P-corridors which consists of

$$dl_P\frac{k^n - 1}{k-1} \leq dl_P\frac{k^{|\gamma|} - 1}{k-1} \preccurlyeq e^{|\gamma|}$$

2-cells. The resulting surface diagram therefore has

$$|D| \preccurlyeq e^{|\gamma|} \cdot e^{|\gamma|} \preccurlyeq e^{|\gamma|}$$

2-cells. $\square$

It remains to show that, for all $N \in \mathbb{N}$, there is a surface diagram D such that $N \leq e^{|\partial D|} \preccurlyeq |D|$. We will do this via a surface of genus 0.

Proposition 5.10. *For all $n \in \mathbb{N}$, there is a van Kampen-reduced disk diagram in $\mathcal{P}(S)$ with $|\partial D| = 4n + 4$ and $|D| = \frac{2(k^n-1)}{k-1}$.*

Proof. Let T^n be the unique van Kampen-reduced disk diagram where we take n open t_P-corridors and attach them, short side to long side, such that ∂T^n reads

$$at_P^n = t_P^n a^{k^n}$$

as follows.

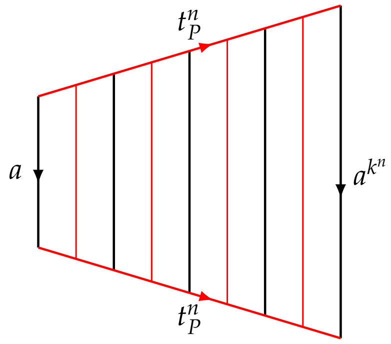

FIGURE 3.50: A van Kampen-reduced disk diagram made up of copies of a single t_P-relator.

We then take two copies of this diagram and attach them via a path of a^{k^n-1} edges labelled a.

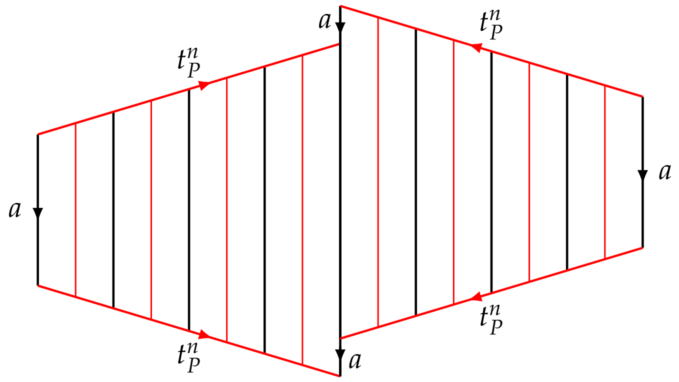

FIGURE 3.51: A van Kampen-reduced disk diagram made up of copies of a single t_P-relator.

This diagram has all the conditions claimed. □

This is the same construction used to prove that the Baumslag-Solitar groups of the form $BS(1,k)$ have exponential Dehn function.

Theorem 5.11. *If K is very-suitable and k is sufficiently large (depending on K) then the homological Dehn function satisfies*

$$FA_{G(S)}(n) \simeq e^n$$

for all $S \subseteq \mathbb{N}_0$.

Proof. Proposition 5.9 provides an upper bound and, since

$$\frac{2(k^n-1)}{k-1} \preccurlyeq e^n \preccurlyeq e^{4n+4},$$

Proposition 5.10 gives an equivalent lower bound. □

In (16) the author asks what homological Dehn functions are attainable when considering groups of type FP. Currently the following question is open.

Open Question 5.12. Are there uncountably many groups of type FP with homological Dehn function $FA_G(n) = n^2$?

It seems that the exponential part of the homological Dehn function in our case is coming from the t_P-relators. If one could find a group G_P such that:

- $G(S)$ is type FP;
- we still have a grasp on disk diagrams in the corresponding presentation;
- and G_P has Dehn function n (or possibly n^2).

If so we may be able to positively answer this question.

6 Acylindrical hyperbolicity of $G(S)$

In this section we will discuss the hyperbolicity of $G(S)$.

Small cancellation groups are heavily associated with hyperbolicity. Yet it is well known that hyperbolic groups must be finitely presented and cannot contain Baumslag-Solitar subgroups; both of of these obstacles arise here. Despite this, given a condition on K, we are able to adhere to a weaker version of hyperbolicity, namely acylindrical hyperbolicity.

Theorem 6.5. If K is suitable and there are polygons P, Q in K such that there are no edges in $\partial P \cap \partial Q$ then, for all $S \subseteq \mathbb{Z}_{\geq 0}$, the group $G(S)$ is acylindrically hyperbolic.

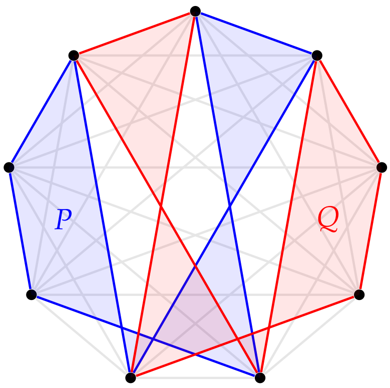

FIGURE 3.52: A pair of polygon boundaries that do not share any edges.

Every suitable complex constructed so far has this property. This induces the following open question.

Open Question 6.1. Is there a suitable complex K such that, for all polygons P, Q in K, $P \cap Q$ contains an edge?

This does not seem an unreasonable property, although it seems less likely the more polygons in K and, since we have found an array of 'small' suitable complexes each without any such pairs of polygons, perhaps it can be proven there is no such construction.

If such a suitable complex does exist it would also be interesting to ask whether $G(S)$ is still acylindrically hyperbolic; it would certainly require an alternative proof.

6.1 Acylindrical hyperbolicity

We begin with the definition of an acylindrical action.

Definition 6.2. A group acts **acylindrically** on a metric space (S, d) if for every $\epsilon > 0$ there are $R, N > 0$ such that, for every two points $x, y \in S$ with $d(x, y) \geq R$, there are at most N elements $g \in G$ with

$$d(x, gx) \leq \epsilon \text{ and } d(y, gy) \leq \epsilon.$$

The intuition behind this definition is that if two points are far apart then they will get moved a large amount, where as points close together will only get moved a small amount. One can think of a rotation around the midpoint between of a pair of points in $\mathbb{R}^2$; if they are close then they will move less than if they are far apart.

Definition 6.3. A group G is said to be **acylindrically hyperbolic** if G acts acylindrically on a hyperbolic metric space.

Hyperbolic groups (and any non-elementary subgroups) are acylindrically hyperbolic, Baumslag-Solitar groups are not. The group $G(S)$ can be seen as a combination of groups of these flavours.

Acylindrically hyperbolic groups have many useful properties, see (20), (22), for a comprehensive background. Since each $G(S)$ is defined via a graph of groups, as shown in Section 3, we will be using (20, Theorem 4.17).

Theorem 6.4. *Let G be the fundamental group of a finite reduced graph of groups such that there are no edges from a vertex to itself. Suppose there are edges e, f and an element $g \in G$ such that*

$$|G_f \cap G_e^g| < \infty,$$

then G is either virtually cyclic or acylindrically hyperbolic.

Here e and f may be the same edge and g is allowed to be trivial.

Theorem 6.5. *If there are polygons P, Q in K such that there are no edges in $\partial P \cap \partial Q$ then, for all $S \subseteq \mathbb{Z}_{\geq 0}$, the group $G(S)$ is acylindrically hyperbolic.*

Proof. Let e and f be the edges corresponding to H_P and H_Q in the graph of groups for $G(S)$ and let g be the trivial element in $G(S)$. Since there are no edges in $\partial P \cap \partial Q$ there are no elements in $H_P \cap H_Q$ and we are able to apply Theorem 6.4, noting that $G(S)$ is certainly not virtually cyclic. □

One can show that $H(S)$ and H_P are also acylindrically hyperbolic (following the idea that they are almost hyperbolic) via (13, Theorem 1.1).

Theorem 6.6 (Gruber)**.** *Let Γ be a graph satisfying $C'(\frac{1}{6})$, then the group corresponding to the graphical presentation of Γ is either virtually cyclic or acylindrically hyperbolic.*

For G_P it is unclear.

7 Automorphisms of $G(S)$

Given an arbitrary group G one often asks about its automorphisms. It is clear that any automorphism of K will induce an automorphism on $G(S)$ for all S. However, as can be seen in Section 2, these automorphisms are not plentiful. In fact we have only found one suitable complex with a non-trivial automorphism: the complex defined in Example 2.9 has a single automorphism of order 2 given by reflecting in the natural line of symmetry.

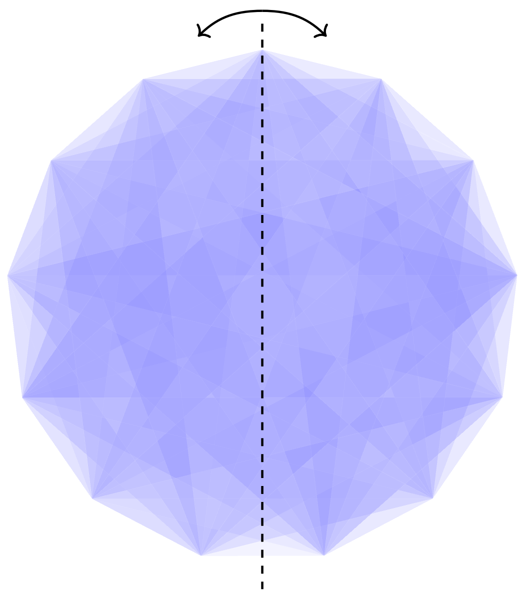

FIGURE 3.53: An automorphism of a suitable complex.

There are subgroups of $G(S)$ that suggest there may be more automorphisms: those isomorphic to Baumslag-Solitar groups, which famously have interesting automorphisms.

7.1 Some non-Hopfian $G(S)$

Baumslag-Solitar groups can be described via the following presentation:

$$BS(m, n) = \left\langle a, t \mid ta^m t^{-1} = a^n \right\rangle .$$

On can show that $BS(2, 3)$ is non-Hopfian.

Definition 7.1. A group G is **non-Hopfian** if there is an epimorphism $\phi : G \to G$ that is not injective.

Example 7.2. Let $G = BS(2, 3)$ then the homomorphism ϕ defined by mapping

$$a \mapsto a^2 \text{ and } t \mapsto t$$

certainly defines an epimorphism. However the element

$$[tat^{-1}, a]$$

is an element of G that lies in the kernel of ϕ. This follows from Britton's Lemma as presented in (15).

In $G(S)$ there are several subgroups isomorphic to $BS(1,k)$. In particular these can be found as subgroups generated by a and t_P where a is some edge in the boundary of the polygon P.

The group $BS(1,k)$ does not share the property of being non-Hopfian with $BS(2,3)$. However, given certain choices of S, the group $G(S)$ can be shown to be non-Hopfian.

Definition 7.3. We define a subset $S \subseteq \mathbb{N}_0$ to be **periodic** if it is of the form

$$S = \{mn + c \mid n \in \mathbb{N}_0\}$$

for some $m, c \in \mathbb{N}_0$. The constant m is called the **periodicity** of S.

Theorem 7.4. *If K is very-suitable and k is sufficiently large (depending on K) then, if S is periodic with non-zero periodicity, $G(S)$ is non-Hopfian.*

Proof. Let m be the periodicity of S and consider the homomorphism ϕ defined by

$$a \mapsto a^{k^m}$$

for each edge a in K and

$$t_P \mapsto t_P$$

for each polygon generator.

We check this is a homomorphism by observing that the periodicity of S is respected by ϕ. This means that every closed path relator is mapped to another closed path relator of higher degree.

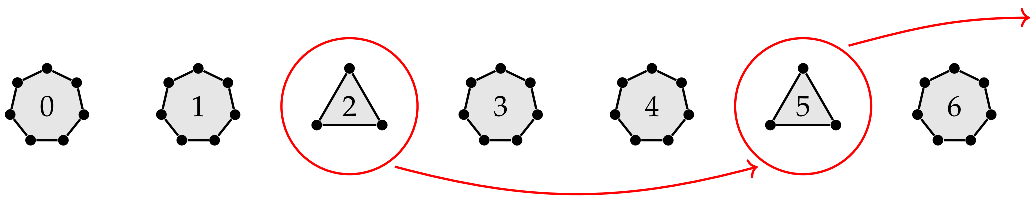

FIGURE 3.54: Here we represent how the periodicity of $S = \{2, 5, 8, 11, ...\}$ is reflected in ϕ.

The surjectivity of ϕ can be seen by considering

$$\phi(t_P^{-m} a t_P^m) = t_P^{-m} a^{k^m} t_P^m = a.$$

This allows us to find a $g \in G(S)$ for each generator h in $\mathcal{P}(S)$ such that $\phi(g) = h$, and thus ϕ is surjective.

Finally we consider injectivity. Consider the word

$$w = t_P^{-m} a t_P^m t_Q^{-m} a^{-1} t_Q^m$$

in $G(S)$ where a is an edge in $\partial P \cap \partial Q$ for some polygons P, Q. We have $\phi(w) = 1$. Now let us suppose that $w = 1$ in $G(S)$. In this case there is a disk diagram D with $\partial D = w$. Since t_P-corridors cannot cross and cannot twist we have the following.

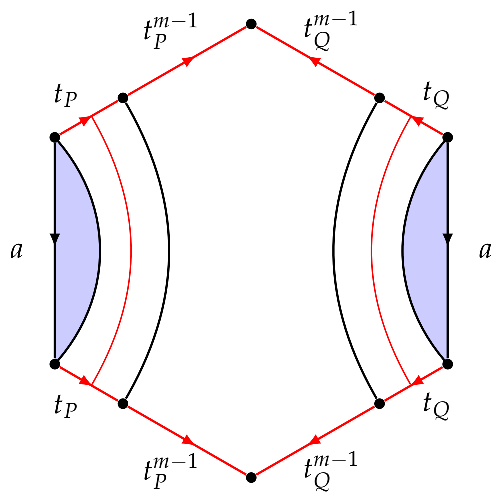

FIGURE 3.55: The boundary of a disk diagram with boundary w.

In which case the blue regions above are leaf regions, but this is a contradiction since the edge a is not long enough to be the visible part of a ladder as in Proposition 1.23. We deduce that w represents a non-trivial element in $G(S)$ that maps to a trivial element via ϕ, and so ϕ is not injective. □

7.2 Other automorphisms

For non-periodic S we cannot consider such a map since it will not induce a homomorphism. It would be interesting to know if there are indeed any other such automorphisms.

Open Question 7.5. Let K be a suitable complex with no automorphisms and let $S \subseteq \mathbb{N}_0$ be non-periodic. Are there any outer automorphisms of $G(S)$?

The work in Section 1 will certainly be useful in attempting to answer this since it allows us to determine whether or not words are trivial.

8 The Σ^1-invariant

In this section we compute the first Bieri-Neumann-Strebel invariant for some $G(S)$. The motivation for this study lies in an attack on the following open question:

Open Question 8.1. Can $\Sigma^1(G)$ take uncountably many values as G ranges over all finitely generated groups?

Since we are working with an uncountable number of groups this question is primed for approach.

We recount some background loosely following (24); for an introduction to the topic use this reference. For an arbitrary group G the invariant is denoted $\Sigma^1(G)$. It is a subset of some sphere $S(G)$, depending on G, and consists of points representing equivalence classes of characters $[\chi]$ which associate to a connected subgraph of a Cayley graph for G.

8.1 Characters and the character sphere

Let G be a finitely generated group.

Definition 8.2. A homomorphism $\chi : G \to \mathbb{R}$ is a **character** of G. Here $\mathbb{R}$ is the additive group of the field of reals.

The set $\mathrm{Hom}(\mathrm{G}, \mathbb{R})$ of all characters of G is a real vector space and, since G is finitely generated, it is finite dimensional. Moreover it can be observed that the dimension of $\mathrm{Hom}(\mathrm{G}, \mathbb{R})$ is equal to the torsion-free rank of the abelianisation

$$G_{ab} = G/[G, G].$$

Definition 8.3. Two characters χ_1 and χ_2 of G will be called **equivalent** if there is a positive real number r with $\chi_1 = r \cdot \chi_2$.

The equivalence class $[\chi]$ of a non-zero character χ can be viewed as a ray from 0 passing through χ.

Definition 8.4. The set of equivalence classes

$$S(G) = \{[\chi] \mid \chi \in \mathrm{Hom}(\mathrm{G}, \mathbb{R}) \backslash \{0\}\}$$

with the structure inherited from $\mathrm{Hom}(\mathrm{G}, \mathbb{R})$ is the **character sphere** of G.

We observe that the dimension of the sphere $S(G)$ is one less than the dimension of $\mathrm{Hom}(\mathrm{G}, \mathbb{R})$.

8.2 The invariant Σ^1

Choose some generating set X for the group G and let Γ be the Cayley graph of G with respect to X. For every character $\chi : G \to \mathbb{R}$ we define a subset

$$G_\chi = \{g \in G \mid \chi(G) \geq 0\}\,.$$

We view this as a subset of the Cayley graph Γ and define Γ_χ to be the induced subgraph of Γ restricted to the vertices in G_χ. Here we observe the purpose of the equivalence class above: the set of vertices is independent over multiplying a character by a positive constant.

We are now able to introduce the invariant.

Definition 8.5. The set

$$\Sigma^1(G) = \{[\chi] \mid \Gamma_\chi \text{ is connected}\}$$

is the **first Σ invariant** of G.

There are several key statements about the Σ^1 that give it the right to be called an invariant. Namely, it is invariant under isomorphism and change of generating set. These appear to be not obvious, as the graph Γ clearly depends on the choice of generating set, but they can be proved as in (24).

8.3 The Σ^1-criterion

Computing whether or not a subgraph of a Cayley graph is connected is often not an easy task. Here we present a criterion for getting hold of points in $\Sigma^1(G)$.

Theorem 8.6 (Theorem A3.1 in (24))**.** *Let G be a group and let X be a finite generating set for G. For every non-zero character $\chi : G \to \mathbb{R}$ and for every choice of $t \in X \cup X^{-1}$ with $\chi(t) > 0$, the following conditions are equivalent:*

(i) the subgraph Γ_χ is connected;

(ii) for every $y \in X \cup X^{-1}$ there exists a path p_y from t to $y \cdot t$ in Γ that satisfies the inequality $v_\chi(p_y) > v_\chi((1, y))$.

Here v_χ maps a path γ to the lowest value of $\chi(v)$ as v ranges across vertices in γ.

This allows us to compute characters in Σ^1 using diagrams as seen in Figure 3.56.

8.4 Application to $G(S)$

We are now ready to compute $\Sigma^1(G(S))$. We will use the generating set of $\mathcal{P}(S)$, letting a be an edge generator and t_P a polygon generator.

Proposition 8.7. *The character sphere of $S(G)$ has dimension one less than the number of polygons in K.*

Proof. The torsion-free part of G_{ab} is free-abelian on the polygon generators due to the t_P-relators. The dimension of $S(G)$ follows immediately. □

We observe the first interesting fact in that $S(G(S))$ does not depend on S, as such we will just write $S(G)$ for brevity. It remains to show whether or not $\Sigma^1(G(S))$ follows the same pattern.

Theorem 8.8. *The subspace $\Sigma^1(G(S))$ of $S(G)$ is empty.*

Proof. Let t_P be a polygon generator for $G(S)$ and let χ be a character of $G(S)$.

Suppose $\chi(t_P) > 0$, choose some edge a not in ∂P and consider the following diagram.

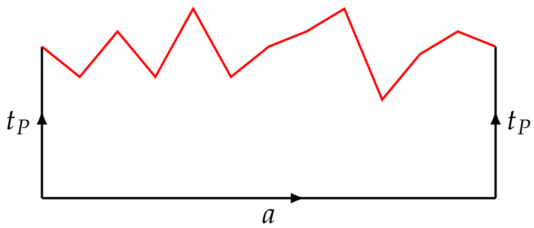

FIGURE 3.56: The edges labelled t_P represent t from Theorem 8.6 and the edge labelled a is our chosen y. The red line represents a potential path from t_P to at_P that does not intersect the path labelled a.

We want information on the red path across the top of the diagram above, we wish to show that every such path must at some point intersect the path of labelled a at the bottom of the diagram. Any such path will combine with the $t_P^{-1}at_P$ path to make a closed path in our Cayley graph and so we are able to apply our work on disk diagrams as in Section 1.

If the two visible edges labelled t_P are the ends of a t_P-corridor then the region enclosed by the edge labelled a and the corridor is a leaf region. This is a contradiction since, as in the proof of Theorem 7.4, the edge a is not long enough to be the visible part of a ladder as in Proposition 1.23. We deduce the diagram has the following form.

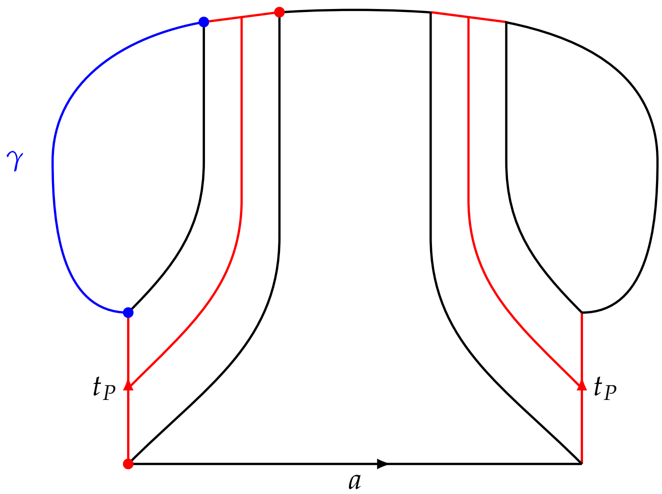

FIGURE 3.57: A disk diagram in $\mathcal{P}(S)$ containing at least two, forced by the three labelled edges, open t_P-corridors.

The blue path γ in the diagram above must start and end with vertices, u and v respectively, such that $\chi(u) = \chi(v)$. This follows from the fact that all edges a are mapped to 0 and, since t_P-corridors cannot cross, for any t_Q in γ there must be a corresponding t_Q^{-1} in γ. In turn this means that the red vertices corresponding to $v \cdot t_P^{-1}$ and our starting vertex take the same values in χ. So we have $v_\chi(p_y) \leq v_\chi((1,y))$ and Γ_χ is not connected.

We next suppose that $\chi(t_P^{-1}) > 0$. The argument follows symmetrically, drawing the t_P edges inversely, and we again find that Γ_χ is not connected.

Therefore the only option is that $\chi(t_P) = 0$. As such, there are no non-zero characters from G to $\mathbb{R}$ and so $\Sigma^1(G(S))$ is empty. $\square$

We observe that $\Sigma^1(G(S))$ therefore does not depend on S and moreover we have not made any progress at answering Question 8.1. We can however deduce that each $G(S)$ does not fibre: there are no maps $\phi : G \to \mathbb{Z}$ such that $\ker \phi$ is finitely generated. In this sense the groups $G(S)$ do not split.

8.5 Altering $G(S)$ for more interesting Σ^1

The main reason that Σ^1 is relatively easy to compute is that $\Sigma(a) = 0$ is forced by the t_P-relators. However this also contributes to the argument that $\Sigma^1(G(S))$ is empty. We can alter the definition of $G(S)$ such that this is not forced and gives a chance of resulting in a more interesting Σ^1.

Consider the sequence

$$a, a^k, a^{k^2}, a^{k^3}, \ldots$$

which is defined by starting with some edge a in the boundary of the polygon P and repeatedly conjugating by t_P. The presentation $\mathcal{P}(S)$ can be seen as having been defined by this sequence, where the K-relators of degree n read words written in the nth term of sequences as the above.

We could consider a different sequence defined by assigning three generators x, y, z to each edge a and using t_P-relators with the following form:

$$t_P x t_P^{-1} = x[y, z^{-1}][y^2, z^{-2}][y^3, z^{-3}]$$

$$t_P y t_P^{-1} = y[z, x^{-1}][z^2, x^{-2}][z^3, x^{-3}]$$

$$t_P z t_P^{-1} = z[x, y^{-1}][x^2, y^{-2}][x^3, y^{-3}]$$

where $[g, h] := ghg^{-1}h^{-1}$ and starting with xyz. So the sequence would start

$$xyz, x[y, z^{-1}][y^2, z^{-2}][y^3, z^{-3}]y[z, x^{-1}][z^2, x^{-2}][z^3, x^{-3}]z[x, y^{-1}][x^2, y^{-2}][x^3, y^{-3}], ...$$

and will grow in length very quickly. In this case we have three generators for each edge in K and the same number of relators as usual. The key point to note here is that these t_P-relators do not force any of the generators a, b, c to map to 0 when considering characters. We choose the long words above such that we are able to maintain all of the small cancellation conditions required to prove our the other results about $G(S)$.

The invariant has not been computed for these groups. We expect $\chi(t_P) = 0$ for all polygon generators t_P but there is potential for the $\chi(a)$ to be non-zero for some edges. We observe that this idea will still fail to answer Question 8.1 because, since K has trivial homology and characters are maps into an abelian group, adding closed path relators of degree n will not change the restrictions on the characters. An answer to Question 8.1 will therefore need a new idea.

We also remark that, using this new construction, there is a graph satisfying $C'(\frac{1}{6})$ that defines G_P. This gives an easier proof that G_P is 2-dimensional and could be used to simplify, and possibly even strengthen, other claims in this thesis.

Appendix A

MAGMA

In Section 2.3 we discuss a construction and a search for which we heavily used MAGMA (4). Here we present, and walkthrough, some of the code used.

1 Computing $H_1(L(q,g))$

We present code for constructing the first homology groups of $L(q,g)$ as in Section 2.3. The blue text shows where we have made a choice for this example: we have chosen to consider $PGL_2(13)$ and its 5th conjugacy class (as listed in MAGMA).

We first choose a prime power q and set up the fundamental group of the 1-skeleton of our complex using a spanning tree S.

```
q := 13;
G := PGL(2, q);

V := GSet(G);
E := [[u, v] : v in V, u in V | u lt v];

F := FreeGroup(#E);
S := [F.edge_index : edge_index in 1 .. #E | 1 in E[edge_index]];
F := quo<F | S>;
```

Next we choose some conjugacy class of G.

```
ConjugacyClasses(G);

class_index := 7;
class := Class(G, ConjugacyClasses(G)[class_index][3]);
```

We then construct a list of polygons corresponding to the conjugacy class chosen above, checking that the polygons have length at least 3, and that we only get each polygon once.

```
P := [Cycle(g, v) : v in V, g in class |
        #Cycle(g, v) gt 2 and
        Cycle(g, v)[2] lt Cycle(g, v)[#Cycle(g, v)] and
        Cycle(g, v)[1] eq Min(Cycle(g, v))];
```

We convert these polygons to words in the free group F.

```
R := [];
for p in P do
    r := F.Index(E, [p[1], p[#p]])^-1;

    for j in [1 .. #p - 1] do
        if p[j] lt p[j+1] then
            r *:= F.Index(E, [p[j], p[j+1]]);
        else
            r *:= F.Index(E, [p[j+1], p[j]])^-1;
        end if;
    end for;
    Append(~R, r);
end for;
```

We are now able to compute the fundamental group and check the first homology.

```
A := quo<F | R>;
AbelianQuotient(A);
```

This set up allows access to subcomplexes by taking some subset of R and computing as above.

2 Searching for pre-suitable subcomplexes

After fixing the number of polygons we want to end up with. We enumerate the subcomplexes and use brute force to check them all until we find a pre-suitable one.

```
X := [1 .. #E - #V + 1];

while true do
    X[#X] +:= 1;

    while #R + 1 in X do
        i := Index(X, #R + 1);
        X[i-1] +:= 1;
        for j in [i .. #X] do
            X[j] := X[j-1] + 1;
        end for;
    end while;

    if IsPerfect(quo<F | [R[x] : x in X]>) then
        X;
        break;
    end if;
end while;
```

Although we have other 'smarter' algorithms, like Algorithm 2.18, for searching this space, in testing this remains the fastest.

For completion we do present an implementation of Algorithm 2.18.

```
X := [];
x := 1;
Append(~X, x);

while true do
    x +:= 1;

    while x gt #R do
        x := X[#X] + 1;
        Prune(~X);
    end while;
```

```
    if IsPerfect(quo<F | [R[y] : y in 1..#R | y notin X cat [x]]>) then
        Append(~X, x);

        if #R + #S eq #E + #X then
            [x : x in 1..#R | x notin X];
            break;
        end if;
    end if;
end while;
```

References

[1] Aschbacher, M. and Segev, Y., 1993. *A fixed point theorem for groups acting on finite 2-dimensional acyclic simplicial complexes.* Proceedings of the London Mathematical Society, 3(2), pp.329-354.

[2] Bestvina, M. and Brady, N., 1997. *Morse theory and finiteness properties of groups.* Inventiones mathematicae, 129(3), pp.445-470.

[3] Bieri, R., 1976. *Homological dimension of discrete groups.* Queen Mary College Mathematics Notes.

[4] Bosma, W., Cannon, J., and Playoust, C., *The Magma algebra system. I.* The user language, J. Symbolic Comput., 24 (1997), 235–265.

[5] Bowditch, B.H., 1998. *Continuously many quasiisometry classes of 2-generator groups.* Commentarii Mathematici Helvetici, 73(2), pp.232-236.

[6] Bridson, M.R., 2002. *The geometry of the word problem.* Invitations to geometry and topology, 7, pp.29-91.

[7] Brown, K.S., 2012. *Cohomology of groups (Vol. 87).* Springer Science & Business Media.

[8] Brown, T. and Leary, I.J., 2020. *Groups of type FP via graphical small cancellation.* arXiv preprint arXiv:2004.04550.

[9] Collins, D.J. and Huebschmann, J., 1982. *Spherical diagrams and identities among relations.* Mathematische Annalen, 261(2), pp.155-183.

[10] Dahmani, F., Guirardel, V. and Osin, D., 2017. *Hyperbolically embedded subgroups and rotating families in groups acting on hyperbolic spaces* (Vol. 245, No. 1156). American Mathematical Society.

[11] Gersten, S.M., 1992. *Dehn functions and l 1-norms of finite presentations.* In Algorithms and classification in combinatorial group theory (pp. 195-224). Springer, New York, NY.

[12] Gromov, M., 2003. *Random walk in random groups* Geometric & Functional Analysis GAFA, 13(1), pp.73-146.

[13] Gruber, D. and Sisto, A., 2018. *Infinitely presented graphical small cancellation groups are acylindrically hyperbolic.* In Annales de l'Institut Fourier (Vol. 68, No. 6, pp. 2501-2552).

[14] Higman, G., 1961. *Subgroups of finitely presented groups.* Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences, 262(1311), pp.455-475.

[15] Kaiser, T., 2020. *A closer look at the non-Hopfianness of $BS(2,3)$.* arXiv preprint arXiv:2005.03396.

[16] Brady, N., Kropholler, R. and Soroko, I., 2020. *Homological Dehn functions of groups of type FP_2.* arXiv preprint arXiv:2012.00730.

[17] Lyndon, R.C. and Schupp, P.E., 2015. *Combinatorial group theory.* Springer.

[18] Leary, I.J., 2018. *Uncountably many groups of type FP.* Proceedings of the London Mathematical Society, 117(2), pp.246-276. Vancouver

[19] McCammond, J.P. and Wise, D.T., 2002. *Fans and ladders in small cancellation theory.* Proceedings of the London Mathematical Society, 84(3), pp.599-644.

[20] Minasyan, A. and Osin, D., 2015. *Acylindrical hyperbolicity of groups acting on trees.* Mathematische Annalen, 362(3-4), pp.1055-1105.

[21] Ollivier, Y., 2003. *On a small cancellation theorem of Gromov*, Bull. Belg. Math. Soc. Simon Stevin, 13, pp.75-89.

[22] Osin, D., 2016. *Acylindrically hyperbolic groups.* Transactions of the American Mathematical Society, 368(2), pp.851-888.

[23] Soroko, I., Kropholler, R. and Leary, I., 2017. *Uncountably many quasi-isometry classes of groups of type FP.*

[24] Strebel, R., 2012. *Notes on the Sigma invariants.* arXiv preprint arXiv:1204.0214.