Geodesics in Cayley graphs

Andrew Elvey Price

October 10, 2014

Abstract

We will examine the structure of geodesics in Cayley graphs, and more generally the geometry of Cayley graphs, from a number of different perspectives. In some cases understanding the geometry can tell us something about the underlying group. We will consider a number of Cayley graph properties which are actually properties of the underlying group, since they hold for one Cayley graph of a group if and only if they hold for any other Cayley graph of the same group. We will also see some examples of groups which admit multiple Cayley graphs with quite different geometry. Later, we will consider a question of Shapiro: When does a group admit a Cayley graph Γ, such that geodesics in Γ are unique? The conjecture is that only a very simple class of groups, called "basic" groups, have this property. Finally, we consider a related question in the case of Coxeter groups. In answering this new question, we characterise Coxeter groups where all walls are finite.

Acknowledgements

Thanks to my friends and family for their support, in particular Mel, Tian, Adib, Jinghan, Matthew, Dougal, Yi and Thara. Thanks to Murray Elder, Laura Ciobanu, Richard Brak, Arun Ram and Sanming Zhao for their helpful advice, and, of course, a big thanks to my supervisor, Lawrence Reeves.

Contents

Chapter 1

Introduction

Given a group G and a finite generating set S, we can construct the Cayley graph $\Gamma(G, S)$, which is best thought of as a way to visualise the group G. This is useful in two ways - on one hand we can try to understand the group via the Cayley graph, and on the other hand Cayley graphs can simply be a source of many interesting graphs. Another way to understand a group is through the way it acts on spaces, in particular, we will look at actions of groups on trees in some detail.

For most questions in geometric group theory, finite and/or free groups are the trivial examples, so it is natural to ask for a somewhat minimal class of group which contains both. A group is called virtually free if it contains a free subgroup of finite index. This may seem like the smallest interesting class of groups which contains both finite groups and free groups, but an even smaller class of groups, called basic groups, also contains both of the other two classes. A basic group is a group which can be written as a free product of finitely many finite groups and a finitely generated free group. In the second chapter we will discuss four classes, free, finite, basic and virtually free, and give some equivalent formulations for each one from a number of different perspectives, for example the geometry of their Cayley graph, and also the way they act on trees. The class of virtually free groups is certainly more natural than the basic groups, and arguably free groups as well, since properties of free and basic groups tend to be dependent on the generating set whereas properties of finite and virtually free groups tend to be independent of the generating set.

A block of a graph Γ is a maximal two-connected component. In chapters three and four we will use the geometry of the Cayley graph to give a proof that a group G is basic if and only if there exists some generating set S for G such that every block in $\Gamma(G, S)$ is finite. This is equivalent to a result of Haring-Smith [14] which he proved using language theory. The main result in these chapters, though, is that a 2-connected graph B appears as a block in some Cayley graph if and only if there is a group G which acts vertex freely and with finitely many orbits on B. Moreover, if there is such a group G, then there is an integer k and a generating set S for G such that every block in the Cayley graph $\Gamma(G * F_k, S)$ is isomorphic to B.

In chapter five we will discuss a number of properties of Cayley graphs and some implications between them. Many of the questions about which of these properties imply which other properties are still open, but we will prove that some implications do hold and give examples to show that some do not. Many of the properties which we will consider are what we will call *block properties*, since they hold for a Cayley graph Γ if and only if they hold for every block in Γ. As a result, our characterisation of blocks in Cayley graphs could be quite useful for finding examples of Cayley graphs which have some properties and not others, though actually none of the examples which we give for this purpose require any consideration of the blocks in the Cayley graphs. We will also be interested in which of the properties are group properties, that is, which of them are independent of the generating set S. In this case we do use our characterisation of blocks in Cayley graphs to construct some examples.

In chapter six we look at a long standing problem in both geometric group theory and graph theory in general: When does a (Cayley) graph have unique geodesics? In other words, when is it the case that for every pair of vertices u, v in Γ, there is a unique geodesic between u and v? If this happens, Γ is called a geodetic graph. Ore [22] posed this question for general graphs in 1965, whereas the problem for Cayley graphs was posed by Shapiro [26] in 1997, though there seems to be no connection between these two problems in the literature. For the latter problem, the conjecture is that a group G is represented by some geodetic Cayley graph Γ if and only if G is basic. If G is basic then there certainly is such a Cayley graph Γ, however the opposite direction is still an open problem. We will also introduce a generalisation of geodetic graphs which we call almost geodetic graphs. For convenience, we will call a group (almost) geodetic if it admits an (almost) geodetic Cayley graph. We conjecture that a group G embeds into an almost geodetic group if and only if it is virtually free and does not contain any subgroups of the form $\mathbb{Z} \times \mathbb{Z}_k$ for any integer $k \geq 2$. For this question we cannot completely prove either direction, though we do at least prove the forwards direction for virtually free groups. That is, if G is a virtually free, almost geodetic group, then G contains no subgroups of the form $\mathbb{Z} \times \mathbb{Z}_k$.

In the final chapter we will show that the last conjecture is true in the case of virtually free Coxeter groups, and that this is also equivalent to an interesting geometric property of the Cayley graph of a Coxeter group, namely all walls being finite. We do this by completely characterising the Coxeter diagrams of such groups.

1.1 Geometric group theory

We will define two different versions of a Cayley graph. Since they are so similar, these are often used interchangeably, however it will be important that we can easily distinguish between the two constructions.

Definition 1.1. Let G be a group and let S be finite a generating set for G. We define the edge separated Cayley graph $\Gamma_e(G, S)$ to be the labelled, directed graph with vertex set G and such that for each vertex $g \in G$ and each $s \in S$, there is an edge from g to gs, labelled by s.

Definition 1.2. We define the Cayley graph $\Gamma(G, S)$ to be the graph with the same vertex set and with edges between the same pairs of vertices as $\Gamma_e(G, S)$, the only difference being that when there are multiple edges between two vertices in Γ_e we only have a single edge in Γ, so Γ is a simple graph.

The action of a group G on itself by left multiplication extends to an action of G on the Cayley graph $\Gamma(G, S)$. The following proposition shows that we can also define a group from its Cayley graph, rather than the other way around.

Proposition 1.3. *Let S be a finite set and let Γ be a labelled, directed, connected graph, with labels coming from S such that each vertex has exactly one edge with each label pointing into and out of it. If Γ is vertex transitive then Γ is an edge separated Cayley graph of the automorphism group $Aut(\Gamma)$.*

Proof. Since there is only one edge with each label pointing into and out of each vertex, any isomorphism of Γ which fixes a single point must fix all of Γ. Hence, $\mathrm{Aut}(\Gamma)$ acts freely on Γ. Let v_1 be a vertex in Γ. For each element $s \in S$, let v_s be the vertex in Γ which is on the end of the edge labelled s emanating from v_1. We identify each element $s \in S$ with the element of $\mathrm{Aut}(\Gamma)$ which sends v_1 to v_s, so $v_s = sv_1$. More generally, for each element $g \in \mathrm{Aut}(\Gamma)$, let v_g be the vertex gv_1. Since Γ is vertex transitive, every vertex in Γ is labelled v_g for some $g \in \mathrm{Aut}(\Gamma)$. We just need to show three things:

- For each $g \in \mathrm{Aut}(\Gamma)$ and each $s \in S$, there is an edge in Γ from v_g to v_{gs} labelled by s.
- There are no other edges.
- S is a generating set for $\mathrm{Aut}(\Gamma)$.

First, There is an edge labelled by s from v_1 to $v_s = sv_1$, so there is an edge labelled by s from $gv_1 = v_g$ to $gsv_1 = v_{gs}$. For the second part, for each vertex v_g, there is only one edge emanating from it with each label s, so this edge must go to the vertex v_{gs}. Hence there are no other edges.

For the third part, consider an element $g \in \mathrm{Aut}(\Gamma)$. Then we just need to show that g is equal to some product $s_1 s_2 \dots s_n$ where each $s_i \in (S \cup S^{-1})$. Let p be a path in Γ from v_0 to v_g and let $s_1, s_2, \dots, s_n$ be the labels on the edges of p, where we write s^{-1} if p goes backwards along an edge labelled by s. Then, inductively, we can see that $v_{s_1 s_2 \dots s_i}$ is the vertex at the end of the path $s_1 s_2 \dots s_i$. □

1.2 Amalgamated free products and HNN extensions

In geometric group theory it is normal to specify a group by a generating set S and relations R. Although this is very convenient, a serious problem with describing a group in this way is that it can be very difficult to work out what the group really is when given only the generators and relations. For example, even determining in general whether the group specified is trivial an undecidable problem. In this section we describe two standard ways of constructing groups from simpler groups; the main use is that if we construct a group in one of these ways then we really do know a lot about it.

The following definitions and results can be found in [18].

Definition 1.4 (HNN extension)**.** Let G be a group with presentation $\langle S \mid R\rangle$ and let A and B be isomorphic subgroups of G, with $\phi : A \to B$ an isomorphism. We define the HNN extension $G*_\phi$ to be the group with presentation

$$\langle S \cup \{t\} \mid R, t^{-1}at = \phi(a)\forall a \in A\rangle.$$

Note that if S_A is a finite generating set for A and S and R both finite, we have a finite presentation for $G*_\phi$ given by:

$$\langle S \cup \{t\} \mid R, t^{-1}at = \phi(a)\forall a \in S_A\rangle.$$

We will call this a standard presentation for the HNN extension.

We will also define a generalisation of HNN extensions, which we will only use a few times. While the following is more general than a single HNN extension, it is strictly less general than applying multiple HNN extensions to a group, one after the other since, for example, the group

$$\langle a, b, c \mid b^{-1}ab = a^2, c^{-1}bc = b^2\rangle$$

is not a multiple HNN extension of $\mathbb{Z}$. It is, however, a HNN extension of $\langle a, b \mid b^{-1}ab = a^2\rangle$ which is a HNN extension of $\mathbb{Z}$.

Definition 1.5 (Multiple HNN extension)**.** Let G be a group with presentation $\langle S \mid R\rangle$ and let $A_1, A_2, \ldots, A_n, B_1, B_2, \ldots, B_n$ be subgroups of G, with $\phi_i : A_i \to B_i$ an isomorphism for each i. We define the multiple HNN extension $G*_{\phi_1,\ldots,\phi_n}$ to be the group with presentation

$$\langle S \cup \{t_1, \ldots, t_n\} \mid R, t_i^{-1}a_it_i = \phi_i(a_i)\forall a_i \in A_i, i \in \{1, \ldots, n\}\rangle.$$

Note that if each group A_i has a finite generating set S_i and S and R are also finite, then we have a finite presentation for $G*_{\phi_1,\ldots,\phi_n}$ given by:

$$\langle S \cup \{t_1, \ldots, t_n\} \mid R, t_i^{-1}a_it_i = \phi_i(a_i)\forall a_i \in S_i, i \in \{1, \ldots, n\}\rangle.$$

We will call this the standard presentation for the multiple HNN extension.

Proposition 1.6 (Britton's lemma)**.** *Let $G*$ be a multiple HNN extension with presentation given as above. Let w be a word over the alphabet*

$$S \cup S^{-1} \cup \{t_1, \ldots, t_n, t_1^{-1}, \ldots, t_n^{-1}\},$$

such that $\overline{w} = 1$, and w contains at least one letter which is not in $S \cup S^{-1}$. Then there is some subword of w which takes either the form $t_i^{-1}vt_i$, where $\overline{v} \in A_i$, or $t_ivt_i^{-1}$, where $\overline{v} \in B_i$, and such that v only contains letters in $S \cup S^{-1}$.

Proposition 1.7. *Let $G*$ be a multiple HNN extension with a standard presentation, and let Γ be the corresponding Cayley graph. Let $\sim$ be the equivalence relation on $G*$ given by $a \sim b$ if and only if $a^{-1}b \in G$, in other words a and b are in the same left coset of G in $G*$. Then the quotient graph $\Gamma/\sim$ is a tree.*

We will call this the Bass-Serre tree of the (multiple) HNN extension, and we will call each coset a sheet of the Cayley graph Γ.

Proposition 1.8. *Let $G*$ be a multiple HNN extension with presentation as above and let w be a word over the alphabet $S \cup S^{-1} \cup \{t_1, \ldots, t_n, t_1^{-1}, \ldots, t_n^{-1}\}$ with $\overline{w} \in G$. Then (at least) one of the following is true:*

I $|w| \leq 1$.

II The word w decomposes as $w = w_1 w_2$ where w_1 and w_2 are non-empty subwords of w and $\overline{w_1}, \overline{w_2} \in G$.

III There is some j and some word w_1 with $\overline{w_1} \in A_j$ such that $w = t_j^{-1} w_1 t_j$. Note that this is only possible if $w \in B_j$.

IV There is some j and some word w_1 with $\overline{w_1} \in B_j$ such that $w = t_j w_1 t_j^{-1}$. Note that this is only possible if $w \in A_j$.

Proof. Since $\overline{w} \in G$, the path in the Bass-Serre tree T corresponding to w must be a loop l. If l contains any intermediate vertices which are at the same point as the first and last vertices in l, then we can split l into two loops l_1 and l_2. Hence, the corresponding words w_1 and w_2 satisfy condition II. Otherwise, if $|w| > 1$, then the second and second last vertices of l must be the same, so we get either condition III or condition IV, depending on which generator t_j or t_j^{-1} the first edge in l corresponds to. □

Definition 1.9 (Amalgamated free product)**.** Let A and B be groups with presentations $\langle S_A \mid R_A \rangle$ and $\langle S_B \mid R_B \rangle$ respectively. Let H be a group and let $\phi_A : H \to A$ and $\phi_B : H \to B$ be injective group homomorphisms. We define the amalgamated free product $A *_H B$ to be the group with presentation

$$\langle S_A, S_B \mid R_A, R_B, \phi_A(h) = \phi_B(h) \ \forall h \in H \rangle.$$

Note that if S_H is a finite generating set for H and the sets S_A, S_B, R_A, R_B are all finite, then we have a finite presentation for $A *_H B$ given by

$$A *_H B = \langle S_A, S_B \mid R_A, R_B, \phi_A(h) = \phi_B(h) \ \forall h \in S_H \rangle.$$

We will regularly identify H with $\phi_A(H)$ and $\phi_B(H)$, and will even say that $A \cap B = H$. There is a tree related to amalgamated free products too:

Definition 1.10. Given an amalgamated free product $A *_H B$, we construct a graph T as follows: There is one vertex in T for each left coset gA of A in $A *_H B$ and one vertex for each left coset of B in $A *_H B$. Two vertices in T are joined by an edge if the corresponding cosets intersect each other. Note that this can only happen if one of the vertices corresponds to a coset of A and the other corresponds to a coset of B; moreover, in this case the intersection will be a left coset of H.

Proposition 1.11. *the graph T is a tree.*

We will call T the Bass-Serre tree of the amalgamated free product and we will call each coset of A or B in $A *_H B$ a chamber of the amalgamted free product.

Proposition 1.12. *If we have an alternating word* $w = (b_0)a_1b_1a_2b_2 \ldots a_nb_n(a_{n+1})$ *where each* $a_i \in A$ *and each* $b_i \in B$ *and such that* $\overline{w} = 1$ *in* $A *_H B$*, then there is some* $a_i \in \phi_A(H)$ *or some* $b_i \in \phi_B(H)$*.*

Proposition 1.13. *Let* $a_1, a_2, \ldots, a_n \in A \setminus H$ *and let* $b_1, b_2, \ldots, b_n \in B \setminus H$*. Then the alternating product* $a_1b_1 \ldots a_nb_n$ *is not torsion in* $A *_H B$*.*

Proof. If $a_1b_1 \ldots a_nb_n$ was torsion, we would have an alternating product $(a_1b_1 \ldots a_nb_n)^k = 1$, which would contradict the previous proposition. □

1.3 Ends of groups

The theory of ends of groups is covered in detail in [28]. We will simply give a few definitions and basic results from there.

Definition 1.14 (Ends of a graph)**.** Let Γ be a graph, let $x \in \Gamma$ be a point and for $d \in \mathbb{R}_{>0}$, let $B(x, d)$ denote the ball of radius d around x. We define $e(d)$ to be the number of infinite connected components of the space $\Gamma \setminus B(x, d)$. We can easily see that $e(d)$ is non-strictly increasing as a function of d. We define the number of ends of Γ to be the following limit:

$$\lim_{d \to \infty} e(d).$$

Definition 1.15 (Ends of a group)**.** Let G be a group and let S be a finite generating set for G. The number of ends of G is equal to the number of ends of the Cayley graph $\Gamma(G, S)$. This does not depend on the generating set.

We now list basic facts about the ends of finitely generated groups here:

- If G is a finitely generated group, then G has 0,1,2 or ∞ ends.
- A finitely generated group G has 0 ends if and only if it is finite.
- A finitely generated group G has 2 ends if and only if $\mathbb{Z}$ is a subgroup of G with finite index, in other words G is virtually $\mathbb{Z}$.

The following theorem was proven in [29]

Theorem 1.16 (Stalling's theorem)**.** *If* G *is a finitely generated group with more than one end, then* G *splits as an amalgamated free product or HNN extension over a finite subgroup.*

1.4 Language theory

We will now give a very brief introduction to formal language theory. For a more detailed introduction, see [15].

Given a finite set Σ, which we will call an alphabet, Σ^* denotes the set of words over the alphabet S. A language over the alphabet Σ is any subset of Σ^*. Hence, given a group G and a finite generating set S, there are a number of naturally arising languages over the alphabet $(S \cup S^{-1})$. For example, the word problem of (G, S) is the set of all words over the alphabet $(S \cup S^{-1})$ which denote the same group element. We will also consider the language of all geodesics, that is the set of all words w which correspond to geodesics in the Cayley graph $\Gamma(G, S)$.

Formal languages are broken into a number of classes based on how complex a machine needs to be to be able to decide if a word is in the language. At one end, a regular language is one which can be recognised by the simplest possible type of machine, one with only finitely many states. Such a machine is called a (deterministic) finite state automaton. At the other end of the hierarchy, we have the class of decidable languages, which is the set of languages which can be recognised by a Turing machine. We will also mention Context free languages and simple languages. These three main classes, regular, context free and decidable lie in the implication chain:

$$\text{Regular} \Rightarrow \text{Context free} \Rightarrow \text{Decidable},$$

though there are certainly many classes in between context free and decidable languages. The class of simple languages is a subclass of context free, though neither it nor the class of regular languages is contained in the other.

Given one language, there are a number of operations we can apply to it to construct other languages. One which we will refer to is the Kleene star. Given a language L over a finite alphabet Σ, the Kleene star L^* of σ is the set of all words which can be obtained by concatenating finitely many words in L. Note that the empty word is always an element of L^*.

Chapter 2

Free, basic and virtually free groups

The following table shows a number of equivalent characterisations for finite groups, finitely generated free groups, finitely generated virtually free groups and basic groups. In the table, G is a finitely generated group and, unless otherwise stated, S is any finite generating set for G.

Finite	**Free**	**Basic**	**Virtually free**
$\Gamma(G, S)$ is finite.	For some finite generating set S of G, $\Gamma_e(G, S)$ is a tree.	For some finite generating set S of G, $\Gamma(G, S)$ has finite block size [14].	$\Gamma(G, S)$ is quasi-isometric to a tree [17].
The word problem for G w.r.t. S is a regular language.	-	For some finite generating set S of G, the word problem of G w.r.t. S is the Kleene star of some simple language [14].	The word problem of G w.r.t. S is context-free [20].
G is finite.	G is the fundamental group of a finite graph.	G is the fundamental group of a finite graph of finite groups with trivial edge groups.	G is the fundamental group of a finite graph of finite groups [16].
G acts with finite vertex stabilisers on some finite tree.	G acts freely and with finitely many orbits on some tree.	G acts with finitely many orbits, finite vertex stabilisers and trivial edge stabilisers on some tree.	G acts with finitely many orbits and finite vertex stabilisers on some tree [16].

2.1 Virtually free groups

Definition 2.1. A group is called virtually free if it contains a free subgroup of finite index.

In a recent survey [8] Diekert and Weiss discuss a number of equivalent conditions. We list a few of those conditions here. If G is a group and S is a finite generating set for G the following conditions are equivalent:

- G is virtually free.
- The Cayley graph $\Gamma(G, S)$ is quasi-isometric to a tree [17].
- The word problem of G with respect to S is context-free [20].
- G is the fundamental group of a finite graph of finite groups [16].
- G acts with finitely many orbits and finite vertex stabilisers on some tree [16].

Later, when we consider Cayley graphs with unique geodesics, we will particularly utilise the second condition, that G is virtually free if and only if $\Gamma(G, S)$ is quasi isometric to some tree.

2.2 Basic groups

Definition 2.2. A group G is said to be basic if it decomposes as the free product of finitely many finite groups and a finitely generated free group. That is, $G = G_1 * G_2 * \ldots * G_n * F_m$ for some non-negative integers n, m and some finite groups $G_1, \ldots, G_n$.

Perhaps surprisingly, all basic groups are virtually free, and the converse certainly does not hold, for example $\mathbb{Z} \times \mathbb{Z}_2$ is virtually free but not basic. We find many more examples of non-basic, virtually free groups, and give a useful characterisation of such groups in theorem 4.6, namely that a virtually free group G is non-basic if and only if it contains an infinite subgroup which is either a non-trivial amalgamated free product of finite groups or a non-trivial HNN extension of a finite group.

Our main interest with basic groups will be the fact that every basic group G has a generating set S such that geodesics between vertices in $\Gamma(G, S)$ are unique (in other words, $\Gamma(G, S)$ is geodetic), and it appears that no other groups have this property. We will also explore a number of different characterisations of basic groups which are analogues of the some of the characterisations of virtually free groups which we described in the previous section. Unfortunately, unlike their counterparts for virtually free groups, these tend to depend on the generating set.

In [14], Haring-Smith showed that for a finitely generated group G, the following three conditions are equivalent:

- G is basic.

- G can be represented by some Cayley graph Γ such that every vertex is contained in only finitely many simple cycles, in other words, every block in Γ is finite.

- There is some set of generators of G such that the word problem of G with respect to these generators is the Kleene star of a simple language.

In theorems 3.20, 3.21, 4.4 and 4.5 we show two equivalent conditions, namely that G is basic if and only if it acts edge freely and with finitely many orbits on some nontrivial, locally finite tree T, and G is basic if and only if it appears as the fundamental group of some finite graph of finite groups with trivial edge groups. These also provide a new proof of the equivalent condition that G is basic if and only if G can be represented by some Cayley graph in which every block is finite.

We will now prove some closure properties of basic groups.

Proposition 2.3. *If A and B are basic groups then the free product $A*B$ is basic.*

Proof. Let $A = A_1 * A_2 * \ldots * A_n * F_j$ and let $B = B_1 * B_2 * \ldots * B_m * F_k$ where each A_i and each B_i is a finite group. Then

$$\begin{aligned} A * B &= A_1 * A_2 * \ldots * A_n * F_j * B_1 * B_2 * \ldots * B_m * F_k \\ &= A_1 * A_2 * \ldots * A_n * B_1 * B_2 * \ldots * B_m * F_{j+k}, \end{aligned}$$

which is basic. □

Hence the free product of finitely many basic groups is basic. Now we will show that any finitely generated subgroup of a basic group is basic.

Theorem 2.4 (Kurosh's Theorem)**.** *[18] A subgroup H of a free product $\prod_j A_j$ is itself a free product*

$$H = F * \prod_k C_k$$

where each C_k is a conjugate of a subgroup of one of the A_js and F is a free group.

Corollary 2.5. *Any finitely generated subgroup of a basic group is basic.*

Proof. Let H be a finitely generated subgroup of a basic group $G = G_1 * G_2 * \ldots * G_n * F_m$. Then by Kurosh's theorem, H is equal to a free product $F * \prod_k C_k$ where F is free and each C_k is either a subgroup of a finite group G_j or a subgroup of the free group F_m. Therefore, each C_k is either finite or free. Now, we break up all of the free groups C_k, along with F up into a free product of copies of $\mathbb{Z}$ So $H = \prod_k B_k$ where each B_k is finite or isomorphic to $\mathbb{Z}$.

Now let S be a finite generating set for H. Then each element $s \in S$ can be written as a product of finitely many elements of $\bigcup_k B_k$. Let T_s be a set of these elements. So s is in the subgroup of H generated by T_s. Now let $T = \bigcup_{s \in S} T_s$. Then S is a subset of $\langle T \rangle$. But S generates H, so $H = \langle T \rangle$.

So T is a finite subset of $\bigcup_k B_k$ which generates $H = \prod_k B_k$. Since T generates $\prod_k B_k$, T must contain at least one element from each group B_k. So there must only be finitely many values of k. Since each of the groups in the free product is basic, H must also be basic. □

Proposition 2.6. *The only two-ended basic groups are* $\mathbb{Z}$ *and* D_∞

Proof. Let H be a basic group. If H is a free product of two or more non-trivial groups and any of these groups has more than two elements we will get a group with infinitely many ends. So H must be either a single group with no free product, so must be finite or free, or H is the free product of finitely many copies of $\mathbb{Z}_2$, since that is the only group with two elements.

If H is finite then it has 0 ends, not two ends. If H is free then it must be isomorphic to $\mathbb{Z}$, otherwise it would have infinitely many ends. If H is the free product of more than two copies of $\mathbb{Z}_2$ then, again, it will have infinitely many ends.

The final remaining case is that H is the free product $\mathbb{Z}_2 * \mathbb{Z}_2$, which is isomorphic to D_∞ □

Later, when we consider almost geodetic groups, we will prove that if G is a virtually free group which is also almost geodetic, then any 2-ended subgroup of G is basic, so we will be interested in how we can tell if G contains a 2-ended subgroup which is not basic.

Theorem 2.7. *Let G be a virtually free group. Consider an action of G on a tree T which has finitely many orbits and finite vertex stabilisers. G contains a 2-ended subgroup which is not basic if and only if there is an element g of G which fixes infinitely many vertices of T.*

Proof. First assume that G has a 2-ended subgroup H which is not basic. Since H is 2-ended, it is virtually $\mathbb{Z}$. So there is an element $a \in H$ with infinite order, such that $\langle a \rangle$ is a finite index subgroup of H.

Now consider a vertex $v \in T$ which minimises $d(v, av)$, where d is the path metric in T. Note that if a fixes v, then $\langle a \rangle$ fixes v, which is impossible since the vertex stabilisers are all finite. Therefore $d(v, av) \geq 1$.

Now let $v = v_0, v_1, \ldots, v_k = av$ be the path between v and av, where $k = d(v, av)$. Now we consider the infinite path p given by $\ldots, v_{-1}, v_0, v_1, v_2, \ldots$ where for $j > k$ we let $v_j = av_{j-k}$ and for $j < 0$ we let $v_j = a^{-1}v_{j+k}$. Note that this also holds for $j = k$ since $v_k = av_0$. We know that v_j and v_{j+1} are joined by an edge when $0 \leq j < k$, so inductively, for $j \geq k$ we see that $v_j = av_{j-k}$ and $v_{j+1} = av_{j-k+1}$ are joined by an edge because v_{j-k} and v_{j-k+1} are joined by an edge. Similarly, for $j < 0$ the vertices v_j and v_{j+1} are joined by an edge. So p is a path in the tree T.

Now we will show that p is a geodesic. Since T is a tree, we just need to show that $v_j \neq v_{j+2}$ for $j \in \mathbb{Z}$. If $k = 1$, then $v_{j+2} = av_{j+1} = a^2v_j$. But a^2 has infinite order, so it cannot be an element of any vertex stabiliser. Therefore, $v_j \neq v_{j+2}$. If $k \geq 2$ then we have a path of length $k - 2$ between v_{j+2} and v_{j+k} so $d(v_{j+2}, v_{j+k}) \leq k - 2$. Since k is minimal, we know that $d(v_j, av_j) \geq k$. So $d(v_j, v_{j+k}) = d(v_j, av_j) \geq k > k - 2 \geq d(v_{j+2}, v_{j+k})$. So $v_j \neq v_{j+2}$. Therefore, p is a geodesic.

Now we will show that any element $h \in H$ stabilises p.

Recall that H is virtually $\langle a \rangle$. So there is a finite subset $\{b_1, b_2, \ldots, b_n\}$ of H such that any element of H can be written as $a^t b_i$ for some $t \in \mathbb{Z}$ and $i \in$

$\{1, 2, \ldots, n\}$. Now let M be the maximum value of $d(v_j, b_i v_j)$ for $i \in \{1, 2, \ldots, n\}$ and $j \in \{0, \ldots, k-1\}$.

Now, if $h \in H$ and $v_j \in p$, then $h = b_i a^t$ for some $t \in \mathbb{Z}$ and $i \in \{1, 2, \ldots, n\}$. So $hv_j = a^t b_i v_j$. Therefore, $d(hv_j, v_{j+tk}) = d(a^t b_i v_j, a^t v_j) = d(b_i v_j, v_j) \leq M$.

So for any $h \in H$ and $v \in p$, the distance between hv and p is at most M. Therefore, the entire geodesic hp is within a distance of at most M of p. But these paths are in a tree, so hp must be the same path as p, so h stabilises p.

Therefore, H stabilises p, so we can consider the action of H on p. Since p is a tree, if H acts edge freely then H is basic. But H is not basic, so there must be some $h \in H$ which fixes an edge (v_j, v_{j+1}) in p. So $hv_j = v_j$ and $hv_{j+1} = v_{j+1}$. But then, since p is a path, h must fix every vertex of p. So there is an element $g = h$ of G which fixes infinitely many vertices of T.

Now we will prove the opposite direction. Assume that there is an element $g \in G \setminus \{1\}$ which fixes infinitely many vertices of T. We will show that G contains a 2-ended subgroup H which is not basic.

First, $\langle g \rangle$ also fixes the infinitely many vertices. But the vertex stabilisers are all finite, so $\langle g \rangle$ must be finite. Therefore, g has finite order.

Since g fixes infinitely many vertices and there are only finitely many vertex orbits in the action, there must be infinitely many vertices $u_0, u_1, \ldots$ which are fixed by g and are in the same orbit of the action of G. For each $j \in \mathbb{Z}_{>0}$, let $a_j \in G$ such that $a_j u_0 = u_j$. Then $a_j^{-1} g a_j u_0 = a_j^{-1} g u_j = a_j^{-1} u_j = u_0$, so $a_j^{-1} g a_j$ is in the stabiliser of u_0. But the stabiliser of u_0 is finite, so there must only be finitely many different elements $a_j^{-1} g a_j$. So there must be an infinite subsequence $b_1, b_2, \ldots$ of $a_1, a_2, \ldots$ such that $b_j^{-1} g b_j$ is the same for all $j \in \mathbb{Z}_{>0}$.

Let $C_G(g)$ be the centraliser of g. For each j we have

$$b_1 b_j^{-1} g = b_1 b_j^{-1} g b_j b_j^{-1} = b_1 b_1^{-1} g b_1 b_j^{-1} = g b_1 b_j^{-1}$$

So $b_1 b_j^{-1} \in C_G(g)$. Since this is true for each j, and no two terms b_j are the same, $C_G(g)$ is infinite. Also, $C_G(g)$ is a subgroup of G, which is virtually free, so the $C_G(g)$ is virtually free. Therefore, $C_G(g)$ contains some element a of infinite order.

Since g has finite order and g commutes with a, the group $\langle a, g \rangle$ is equal to $\langle a \rangle \times \langle g \rangle$ which is isomorphic to $\mathbb{Z} \times \mathbb{Z}_k$ for some integer $k > 1$. This group is 2-ended and is not equal to either $\mathbb{Z}$ or D_∞ so it is not basic. □

Note that this also proves that any 2-ended non-basic group contains a subgroup of the form $\mathbb{Z} \times \mathbb{Z}_k$ for some integer $k \in \mathbb{Z}_{\geq 2}$, so if we are considering virtually free groups with no non-basic subgroups, we can equivalently consider virtually free groups with no subgroups of the form $\mathbb{Z} \times \mathbb{Z}_k$ where $k \in \mathbb{Z}_{\geq 2}$.

Chapter 3

Blocks

3.1 Blocks in general graphs

We will define blocks in a way similar to Ore [22], p85-87. Note that Ore calls them lobe graphs not blocks. First we define a relation $\sim$ on edges by $a \sim b$ if and only if $a = b$ or there is a simple cycle which contains the edges a and b.

Lemma 3.1. $\sim$ *is an equivalence relation.*

Proof. Clearly $\sim$ is reflexive and symmetric, so we just need to show that $\sim$ is transitive. Let a, b and c be edges such that $a \sim b$ and $b \sim c$. We will show that $a \sim c$.

If either $a = b$ or $b = c$ then the result is trivial, so we will assume that $a \neq b$ and $b \neq c$. Let C_{ab} be a simple cycle which contains edges a and b and let C_{bc} be a simple cycle which contains the edges b and c. If a is contained in C_{bc} then C_{bc} is a cycle containing both a and c so $a \sim c$.

Otherwise, let b_1 and b_2 be the vertices at the ends of b and let a_1 and a_2 be the vertices on the ends of a such that $a_1 a_2 b_2 b_1$ are on cycle C_{ab} in that order. Note that we may have $a_1 = b_1$ or $a_2 = b_2$.

Let l_1 be the path between a_1 and b_1 along C_{ab} which does not contain a and let l_2 be the path between a_2 and b_2 along C_{ab} which does not contain a. So C_{ab} is made up of the paths a, l_1, b, l_2 in that order. Now let x_1 be the point on the intersection of l_1 and C_{bc} which is closest to a_1 along l_1. Now let p_1 be the subpath of l_1 between a_1 and x_1. Then p_1 only intersects C_{bc} at x_1. Define p_2 and x_2 similarly. Now let p be the path in C_{bc} between x_1 and x_2 which contains c. Then pp_1ap_2 forms a cycle which contains both a and c, so $a \sim c$. □

Definition 3.2. A block is an equivalence class of $\sim$

An immediate consequence of this definition is that the blocks of a graph partition its edges. We say that a vertex is contained in a block if it is attached to an edge in that block.

Definition 3.3. A graph Γ is called 2-connected if for any vertex $v \in V(\Gamma)$, the graph Γ remains connected after the removal of v.

Proposition 3.4. *Each block B of a graph Γ is 2-connected.*

Proof. If B only contains one edge then it only contains two vertices so it is certainly 2-connected.

If B contains more than one edge and v is a vertex in B, then for any two vertices $v_1, v_2 \in V(B) \setminus \{v\}$, there are distinct edges $e_1, e_2 \in E(B)$ such that e_1 is attached to v_1 and e_2 is attached to v_2. So there is a cycle C in B containing e_1 and e_2, hence C contains v_1 and v_2. Now, this cycle will form two paths in B between v_1 and v_2 which intersect only at v_1 and v_2. Clearly at least one of these paths does not pass through v, so v_1 and v_2 are connected in $B \setminus \{v\}$. Therefore the graph $B \setminus \{v\}$ is connected. Hence, since this is true for any vertex $v \in B$, the block B is 2-connected. □

Proposition 3.5. *If a graph H is 2-connected, then H has only one block, and this block contains all of H.*

Proof. First, if $e_1 = (u_1, v)$ and $e_2 = (u_2, v)$ are any two edges which meet at a vertex v, then after the removal of v, the graph is still connected, so there is a path in H between u_1 and u_2 which does not pass through v. If we add this to the edges e_1 and e_2, we get a cycle in H which contains both e_1 and e_2. Therefore, any two edges in H which share a vertex are in the same block of H.

Let a and b be any edges in H. Since H is connected, there is a sequence of edges $(a = e_1), e_2, \ldots, e_{k-1}, (e_k = b)$ where each pair of adjacent edges e_i, e_{i+1} shares a vertex. Therefore, each pair e_i and e_{i+1} are in the same block of H, so $a = e_1$ and $b = e_k$ are in the same block of H.

Therefore every edge in H is in the same block, so H has only one block, and it contains all of H. □

Together these two propositions give us the following equivalent definition of a block:

Proposition 3.6. *Let Γ be a graph and let B be a subgraph of Γ. Then B is a block if and only if it is a maximal 2-connected component, that is, B is two connected and contains no other two connected subgraphs of Γ.*

Proof. If B is two connected, then by proposition 3.5, every edge in B is in the same equivalence class of $\sim$, so B is contained in some block B' of Γ, which by proposition 3.4, is 2-connected. Hence if B is a maximal 2-connected component, then $B = B'$.

If B is a block, then by proposition 3.4, B is 2-connected. Moreover, if B is contained in some 2-connected component B', then by proposition 3.5, B' is contained in some block B''. Hence B'' contains B, so $B'' = B$ and hence $B' = B$. Therefore, B is a maximal 2-connected component of Γ. □

Definition 3.7. Let R be a connected graph and let $B(R)$ be the set of blocks of R. We define the block tree $BT(R)$ as follows:

- $V(BT(R)) = V(R) \cup B(R)$
- $E(BT(R)) = \{(v, b) : v \in V(R), b \in B(R), v \in V(b)\}$

Proposition 3.8. *$BT(R)$ is a tree.*

Proof. Suppose there is a simple cycle $v_1b_1v_2b_2\dots v_nb_n$ in $BT(R)$, with b_n connected to v_1. Set $v_{n+1} = v_1$. Then for each i, there are edges (v_i, b_i) and (v_{i+1}, b_i) in $BT(R)$. So $v_i, v_{i+1} \in b_i$. So there is a simple path from v_i to v_{i+1} contained entirely in b_i. Now, if we connect the simple paths for each i, we get a simple cycle, which passes through multiple different blocks. A contradiction.

So there are no simple cycles in $BT(R)$, hence $BT(R)$ is a tree. □

Proposition 3.9. *$BT(R)$ is locally finite if and only if every block in R is finite and R is locally finite.*

Proof. The degree of a vertex b of $BT(R)$ corresponding to a block in R is equal to the number of vertices in that block. So it has finite degree in $BT(R)$ if and only if it has finite size in R.

The degree of a vertex v of $BT(R)$ corresponding to a vertex in R is equal to the number of blocks which that vertex is part of, so it has finite degree in $BT(R)$ if and only if it is part of finitely many blocks in R. If every block has finite size, then each block in which v is contained will contribute finitely many, but at least one, edge to v in R. Hence, v will have finite degree in R if and only if v has finite degree in $BT(R)$.

So, if $BT(R)$ is locally finite, then R has finite block size, so R is locally finite. If R is locally finite and has finite block size then $BT(R)$ is locally finite. □

3.2 Blocks in Cayley graphs

If a group G acts on a graph R, then we obtain a natural action for G on $BT(R)$. In particular, we will be interested in the action of G on $BT(\Gamma(G, S))$, for some finite generating set S of G.

Lemma 3.10. *If G acts vertex freely on some graph R, then G acts edge freely on $BT(R)$.*

Proof. Assume that G acts vertex freely on R. If $g \in G$, and (v, b) is an edge in $BT(R)$ which is fixed by g, then $g(v, b) = (v, b)$. Therefore, since $v \in V(R)$ and $b \in B(R)$ we have $gv = v$ and $gb = b$. Since G acts vertex freely on R, and $gv = v$, g must be the identity in G.

Therefore G acts edge freely on $BT(R)$ □

Proposition 3.11. *If G is a group with finite generating set S, then G acts edge freely on $BT(\Gamma(G, S))$.*

Proof. Since G acts vertex freely on $\Gamma(G, S)$, this follows from the previous lemma. □

Lemma 3.12. *Let G be a group with finite generating set S, and let B be a block in the corresponding Cayley graph $\Gamma(G, S)$. Then B contains a left coset of a subgroup of G which has finite index in B.*

Proof. Let v_0 be the vertex in $\Gamma(G,S)$ corresponding to the identity. Then v_0 is also a vertex in $BT(\Gamma(G,S))$ and its degree is finite. Also, B is a vertex in $BT(\Gamma(G,S))$. Now, if $g \in G$, we will say that $g \in B$ if the vertex corresponding to g is in the block B. Then $g \in B$ if and only if the vertex corresponding to g in $BT(\Gamma(G,S))$ is adjacent to B. Now, the vertex corresponding to g is gv_0, so $g \in B$ if and only if $g^{-1}B$ is adjacent to v_0.

Let $B_1, B_2, \ldots, B_k$ be the vertices adjacent to v_0 in $BT(\Gamma(G,S))$ which are in the orbit of B. Then the sets $\{g \in G : gB_i = B\}$ partition B. Also, each of these sets is a left coset of the stabiliser of B. Therefore, each set $\{g \in G : gB_i = B\}$ is a coset of a subgroup of G which has finite index in B. □

Proposition 3.13. *Every block of a Cayley graph has 0,1,2 or ∞ ends.*

Proof. From the previous lemma, every block B contains a coset of a group which has finite index in B. Therefore, the number of ends of the block must be the same as the number of ends of this group, hence this number must be 0,1,2 or ∞. □

Theorem 3.14. *A 2-connected graph B is isomorphic to a block of some edge separated Cayley graph Γ_e if and only if $Aut(B)$ has a subgroup H which acts freely, and with finitely many orbits on B. Moreover, the stabiliser of the block in Γ_e is isomorphic to H.*

Proof. If B does appear as a block in some Cayley graph $\Gamma(G,S)$, consider the subgroup H of G which stabilises B. Then H is also a subgroup of $\mathrm{Aut}(B)$. Also, since G acts freely on $\Gamma(G,S)$, the subgroup H must act freely on B.

For any two edges e_1, e_2 in B with the same label from S, there is an element $g \in G$ which sends e_1 to e_2. Hence, g must send the block containing e_1, which is B, to the block containing e_2, which is also B. Therefore, $g \in H$ and e_1, e_2 are in the same orbit of the action of H on B. Hence, for each element $s \in S$, every edge in $\Gamma(G,S)$ labelled by s is in the same orbit. Therefore, H is a subgroup of $\mathrm{Aut}(B)$ which acts freely and with finitely many orbits on B.

Now we will do the opposite direction. Suppose B is a 2-connected graph such that there exists a subgroup H of $\mathrm{Aut}(B)$ which acts freely and with finitely many orbits on B. Let $A_1, A_2, \ldots, A_n$ be the edge orbits of the action and let $V_1, V_2 \ldots, V_k$ be the vertex orbits of the action. Since the action is free, no element of H flips any edge, so for each edge orbit, we can choose an edge e in that orbit and make a directed edge e' connecting the same two vertices as e. Now, for each edge d in the orbit of e, let $a \in H$ be the element which sends e to d, then we construct the directed edge $d' = ae'$ to replace d. In this way we create a directed graph B', with edges in the same places as B, such that the action of H preserves the directions.

Now let $S = \{a_1, \ldots, a_n\}$ be an arbitrary finite set with n elements. For each edge e in B', if e is in the orbit A_i, label e with the letter a_i. We will now construct a graph Γ, with all blocks isomorphic to B', which we will prove is an edge separated Cayley graph. Let Λ be a countably infinite set of pairs (f, C) where C is a labelled, directed graph which is isomorphic to B', via the isomorphism $f : B' \to C$. Now we construct Γ as follows:

- We start with the graph consisting of all of the graphs C in Λ, where these graphs are disconnected from each other. We denote one of these graphs "the main connected component"
- at each step we take k pairs $(f_1, C_1), (f_2, C_2), \ldots, (f_k, C_k)$ in Λ, such that exactly one of these graphs C_j is in the main connected component. We then choose a vertex v_i in each graph C_i, which is also in the vertex orbit $f_i(V_i)$, and we do this in such a way that v_j is not in any graphs C in Λ other than C_j. We then identify the k vertices $v_1, v_2, \ldots, v_k$ with each other, and colour this new vertex red. Then the main connected component is now the union of the graphs $C_1, \ldots, C_k$ with the old connected component.

We do the second step in such a way that every vertex which is in one of the blocks is eventually coloured red. In the limit, we obtain a graph which we will denote by Γ. Then the blocks of Γ are all isomorphic to B', and each vertex v in Γ is contained in exactly k blocks $B_1, B_2, \ldots, B_k$, which have corresponding isomorphisms $f_i : B' \to B_i$. Moreover, The vertices given by $f_i^{-1}(v)$ are all in different vertex orbits in the graph B'. We will now prove that Γ is an edge-separated Cayley graph.

Let v be any vertex in Γ.

If two edges (v, w_1') and (v, w_2') in Γ have the same label, then the corresponding edges (v_1, w_1) and (v_2, w_2) in B' have the same label. Hence, there must be some element $h \in H$ which sends one to the other, so $hv_1 = v_2$ and $hw_1 = w_2$. Hence, v_1 and v_2 are in the same orbit, so the two edges (v, w_1') and (v, w_2') must be in the same block of Γ, so $v_1 = v_2$. Since the action of H on B is free, this means that $h = 1$, so the two edges are actually the same. Therefore there is at most one edge with each label in B' pointing out of v.

Now, for each label $a_i \in S$ there is some vertex $u \in B'$ such that an edge labelled a_i points away from u. Also, there is some vertex $f_j^{-1}(v)$ which is in the orbit of u. Therefore, there exists some edge e labelled by a_i which points away from $f_j^{-1}(v)$. Hence, $f_j(e)$ is an edge labelled by a_i which points away from v. So there is exactly one edge with each label in B' pointing out of v. Hence, for each vertex v in Γ, exactly one edge with each label points out of v. Similarly exactly one edge with each label points into v. Hence, exactly one edge with each label points into, and out of each vertex in Γ.

Finally, by construction, Γ is vertex transitive. Therefore, Γ is an edge-separated Cayley graph, as required. □

Corollary 3.15. *A 2-connected simple graph B is isomorphic to a block of some Cayley graph Γ if and only if $Aut(B)$ has a subgroup H which acts vertex freely, and with finitely many orbits on B. Moreover, the stabiliser of the block in Γ is isomorphic to H.*

Proof. First assume that a subgroup H exists. We replace each edge e in B with two directed edges, one in each direction to form a directed graph B'. Then, since H acts vertex freely on B, it acts freely on B'. Therefore, B' is isomorphic to a block in some edge separated Cayley graph Γ_e. Hence, B is a block in the corresponding Cayley graph Γ.

Now assume that B is a block in some Cayley graph $\Gamma(G,S)$. Then the corresponding Block B' in the edge separated Cayley graph $\Gamma_e(G,S)$ has edges between the same pairs of vertices. Hence, there must be a group H which acts freely and with finitely many orbits on B'.

Since B and B' have the same vertices, and two vertices are joined by an edge in B if and only if they are joined by an edge in B', the action of H on $V(B') = V(B)$ extends to an action on B. Moreover, since H acts freely on $V(B') = V(B)$, it acts vertex freely on B. Finally, the action of H on B has finitely many vertex orbits, so it has finitely many orbits. Note however that the action may not be free since there may be elements $h \in H$ which fix mid-edges in B. □

In the next chapter we will use Bass-Serre theory to determine more specifically what groups can contain what blocks. In particular, if B is a 2-connected simple graph and H is a group which acts vertex freely and with finitely many orbits on B, then there exists a finitely generated free group F_k and a generating set S for $H * F_k$ such that every block in $\Gamma(H * F_k, S)$ is isomorphic to B.

Example 3.16. Any 2-connected finite graph appears as a block in some Cayley graph. For example if we take the graph B with vertex set $\{1,2,3,4,5\}$ and edge set $\{(1,2),(1,3),(1,4),(2,3),(3,4),(3,5),(4,5)\}$ then we can add labels a,b,c,d,e,f,g to the edges $(1,2),(1,3),(1,4),(2,3),(3,4),(3,5),(4,5)$ respectively, as shown below.

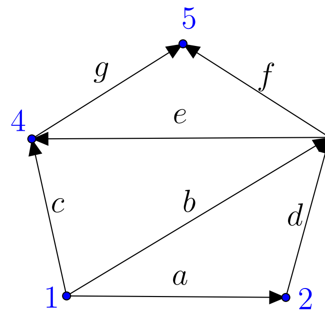

Then every block in the Cayley graph of the group

$$\langle a,b,c,d,e,f,g \mid d = a^{-1}b, e = b^{-1}c, f = b^{-1}cg\rangle$$

is isomorphic to B.

Lemma 3.17. *Let G be a group with finite generating set S and let $s \in S$. Let B_1 and B_2 be blocks in $\Gamma(G,S)$ which both contain an edge labelled by s. Then B_1 is isomorphic to B_2 as a labelled graph.*

Proof. Let e_1 and e_2 be edges in B_1 and B_2 respectively which are both labelled by s. Then there is some element $g \in G$ which sends e_1 to e_2. Therefore, g must send the block containing e_1 to the block containing e_2. So g must send B_1 to B_2. But g is an isomorphism of $\Gamma(G,S)$, so B_1 must be isomorphic to B_2. □

Proposition 3.18. *Let $G_1, G_2, \ldots, G_n$ be groups with finite generating sets $S_1, S_2, \ldots, S_n$ respectively. If B is a block in the Cayley graph $\Gamma(G_1 * \ldots * G_n, S_1 \cup \ldots \cup S_n)$ and B_1 is a block in the Cayley graph $\Gamma(G_i, S_i)$ such that B and B_1 both contain an edge labelled by s, then B is isomorphic to B_1 as a labelled graph.*

Proof. Any simple cycle in $\Gamma(G_1 * \ldots * G_n, S_1 \cup \ldots \cup S_n)$ corresponds to a word w which represents the identity and which contains no proper subwords which also represent the identity. Hence, the letters in w must all be from the same set S_j. Hence, for each block B' in $\Gamma(G_1 * \ldots * G_n, S_1 \cup \ldots \cup S_n)$ the edge labels in B' all come from the same set S_j. In particular, this means that every edge label in B comes from the set S_i.

Let (u, us) be an edge in B and let $(u_1, u_1 s)$ be an edge in B_1. Then for every vertex $v \in B$, there is an edge path from u to v with labels from $S_i \cup S_i^{-1}$, so $u^{-1}v \in G_i$. Define $f_1 : V(B) \to G_i$ by $f_1(v) = u_1 u^{-1} v$ and $f_2 : V(B_1) \to G$ by $f_2(v) = u u_1^{-1} v$.

If two vertices $v_1, v_2 \in V(B)$ are adjacent then the edge joining them is labelled by the element $v_1^{-1}v_2 \in S_i \cup S_i^{-1}$. In this case, $f_1(v_1)^{-1} f_1(v_2) = v_1^{-1} v_2$, so $f(v_1)$ and $f(v_2)$ are joined by an edge in $\Gamma(G_i, S_i)$ with the same label. Hence, f_1 extends to a labelled graph homomorphism $f_1 : B \to \Gamma(G_i, S_i)$. Then, since f_1 is injective, the image of f_1 is isomorphic to B as a labelled graph. In particular, this means that $\mathrm{im}(f_1)$ is 2-connected. Since $\mathrm{im}(f_1)$ contains the edge $(u_1, u_1 s)$, which is in the block B_1, the image $\mathrm{im}(f_1)$ must be contained in B_1.

In exactly the same way we can show that $\mathrm{im}(f_2)$ is contained in B. Finally, from the definitions, f_1 and f_2 are clearly inverse functions, so they must actually be isomorphisms between B and B_1. Hence, B and B_1 are isomorphic as labelled graphs. □

It follows immediately from this that every block in $\Gamma(G_1 * \ldots * G_n, S_1 \cup \ldots \cup S_n)$ is isomorphic to a block in some $\Gamma(G_i, S_i)$.

Theorem 3.19. *Let G be a group and let S be a finite generating set for G. Then we can partition S into sets $S_1, S_2, \ldots, S_n$ such that the following hold:*

- *For each i let G_i be the subgroup of G generated by S_i. Then $G = G_1 * G_2 * \ldots * G_n$.*
- *Any two blocks in the Cayley graph $\Gamma(G_i, S_i)$ are isomorphic to each other as labelled graphs.*
- *Every block in $\Gamma(G, S)$ is isomorphic to a block in one of the graphs $\Gamma(G_i, S_i)$.*

Proof. For $s, t \in S$, we will say that $s \sim t$ if there is a block B in $\Gamma(G, S)$ which contains an edge labelled s and an edge labelled t. First we will prove that $\sim$ is an equivalence relation.

Clearly $\sim$ is reflexive and symmetric, so we just need to show that $\sim$ is transitive. Let $r \sim s$ and $s \sim t$. Then there is a block B_1 which contains an edge labelled by r and an edge labelled by s and there is a block B_2 which contains an edge labelled by s and an edge labelled by t. Since both B_1 and B_2 contain an edge labelled s, by lemma 3.17, B_1 and B_2 are isomorphic as labelled graphs. In particular, this means that B_1 contains an edge labelled t. So $r \sim t$. Therefore, $\sim$ is transitive, so it is an equivalence relation.

Let $S_1, S_2, \ldots, S_n$ be the equivalence classes of $\sim$.

In the following paragraph we will let $\overline{w}$ be the group element in $G_1 * G_2 * \ldots * G_n$ which is given by the word w.

Since $G_1, G_2, \dots, G_n$ are subgroups of G whose union contains all of S, these subgroups generate G. Therefore, there is some normal subgroup N of $G_1 * \dots G_n$ such that G is equal to a quotient group

$$G = G_1 * G_2 * \dots * G_n / N.$$

Suppose that N is not trivial. Let w be a word over the alphabet S with $\overline{w} \in N \setminus \{1\}$ in $G_1 * G_2 * \dots * G_n$. Moreover, let w be the shortest of all such words. Since w represents an element of N, it represents the identity in G. So w forms a cycle in $\Gamma(G, S)$. If w forms a simple cycle in G then the cycle must be contained in a single block B. So every letter in w must be in the same set S_i. But then w is the identity in G_i, so it is the identity in $G_1 * \dots G_n$, a contradiction. So w does not form a simple cycle in $\Gamma(G, S)$. Therefore we can break w up into subwords $w_1 w_2 w_3$ such that $0 < |w_2| < |w|$ and $w_2 =_G 1$. Then $w_1 w_3$ and w_2 are both shorter than w, so $\overline{w_2}, \overline{w_1 w_3} \notin N \setminus \{1\}$. But $\overline{w_2} \in N$, so we must have $\overline{w_2} = 1$ in $G_1 * \dots * G_n$. But then $\overline{w_1 w_3} = \overline{w_1 w_2 w_3} = \overline{w} \in N \setminus \{1\}$, a contradiction. So N is trivial. Therefore,

$$G = G_1 * G_2 * \dots * G_n.$$

Now we will prove that any two blocks in the Cayley graph $\Gamma(G_i, S_i)$ are isomorphic to each other as labelled graphs.

Let B_1 and B_2 be two blocks in $\Gamma(G_i, S_i)$. We will show that B_1 and B_2 are isomorphic as labelled graphs. Let s be a label of an edge in B_1 and let t be a label on an edge in B_2. Then $s, t \in S_i$. Let e be an edge in $\Gamma(G, S)$ with label s and let B be the block containing e. Then by proposition 3.18, B is isomorphic to B_1 as a labelled graph. Since $s, t\ S_i$ we must have $s \sim t$, so t is also an edge label in B. Therefore, B is isomorphic to B_2 as a labelled graph. So B_1 and B_2 are isomorphic as labelled graphs.

Finally, we will show that every block in $\Gamma(G, S)$ is isomorphic to a block in one of the graphs $\Gamma(G_i, S_i)$. Let B be a block in $\Gamma(G, S)$ and let s be an edge label in B. Let $s \in S_i$, let e be an edge in $\Gamma(G_i S_i)$ which is labelled by s and let B_1 be the block containing e. Then B and B_1 are isomorphic as labelled graphs. ☐

3.3 Blocks in basic groups

Previously we stated a result that G is basic if and only if G can be represented by some Cayley graph such that every vertex is contained in only finitely many cycles. Equivalently, every block of that Cayley graph is finite. We now prove the forwards direction of that theorem:

Theorem 3.20. *Let G be a finitely generated group. If G is basic, it can be represented by a Cayley graph Γ such that every block in Γ is finite.*

Proof. Suppose $G = G_1 * \dots * G_n * F_k$ is basic, where each G_i is finite and F_k is free. Let S' be a basis for F_k. We choose the generating set

$$S = (G_1 \cup \dots \cup G_n \cup S') \setminus \{1\}.$$

Then from proposition 3.18, each block in $\Gamma(G, S)$ is isomorphic to a block in some $\Gamma(G_i, G_i \setminus \{1\})$ or a single edge labelled by an element of S'. Therefore, each block is finite. □

Theorem 3.21. *Let G be a finitely generated group. If G can be represented by a Cayley graph Γ such that every block in Γ is finite, then G acts edge freely and with finitely many orbits on some nontrivial, locally finite tree T.*

Proof. Let S be a generating set for G such that every block in $\Gamma(G, S)$ has finite size. Then, by lemma 3.9, the tree $BT(\Gamma(G, S))$ is locally finite. Also, by lemma 3.11, G acts edge freely on $BT(\Gamma(G, S))$. Now, since G acts vertex transitively on $\Gamma(G, S)$, G must act transitively on the vertices of $BT(\Gamma(G, S))$ which correspond to the vertices of $\Gamma(G, S)$. Since each edge in $BT(\Gamma(G, S))$ is connected to exactly one such vertex, and each vertex is connected to finitely many edges, the action has only finitely many orbits. Finally, $\Gamma(G, S)$ has at least one block and at least one vertex, so $BT(\Gamma(G, S))$ is nontrivial. □

Lemma 3.22. *If $a \in G \setminus \{1\}$ has finite order then a fixes exactly one vertex in $BT(\Gamma(G, S))$.*

Proof. If a fixes two vertices $u, v \in BT(\Gamma(G, S))$, then, since $BT(\Gamma(G, S))$ is a tree, it must fix the whole path between u and v. In particular, a must fix some edge in $BT(\Gamma(G, S))$, a contradiction.

Now, suppose for the sake of contradiction that a fixes no vertices in $BT(\Gamma(G, S))$. Consider a vertex $u \in BT(\Gamma(G, S))$ which minimises $d(u, au)$, where d is the path metric in $BT(\Gamma(G, S))$. Note that $d(u, au) \neq 1$, since then one of u and au would be a block in $\Gamma(G, S)$ and the other would be a vertex in $\Gamma(G, S)$. So $d(u, au) \geq 2$.

Now consider a geodesic $(u = u_0), u_1, \ldots, u_{k-1}, (u_k = au)$ between u and au, so $k = d(u, au) \geq 2$. Now we consider the infinite path $u_0, u_1, \ldots, u_k, u_{k+1}, u_{k+2}, \ldots$ where for $j > k$ we let $u_j = au_{j-k}$. Note that this also holds for $j = k$ since $u_k = au_0$. We know that u_j and u_{j+1} are joined by an edge when $0 \leq j < k$, so inductively, for $j \geq k$ we see that $u_j = au_{j-k}$ and $u_{j+1} = au_{j-k+1}$ are joined by an edge because u_{j-k} and u_{j-k+1} are joined by an edge.

Now, each vertex in this path takes the form $a^t u_j$ for some $t \in \mathbb{Z}_{>0}$ and $j \in \{0, 1, \ldots, k-1\}$. Therefore, since a has finite order, there are only finitely many different vertices in the infinite path. So there must be some j such that $u_j = u_{j+2}$. Let j be minimal.

If $j \geq k$, then $u_{j-k} = a^{-1}u_j = a^{-1}u_{j+2} = u_{j-k+2}$, which contradicts the minimality of j. So $j < k$.

But $u_0, u_1, \ldots, u_k$ is a geodesic, so these vertices are all different. Therefore, $j = k - 1$, so $u_{k-1} = u_{k+1}$.

But then $d(u_1, au_1) = d(u_1, u_{k+1}) = d(u_1, u_{k-1}) = k - 2$, which contradicts the minimality of k.

So a fixes exactly one element of $BT(\Gamma(G, S))$. □

Lemma 3.23. *If $a, b \in G \setminus \{1\}$ such that a, b and ab have finite order then a and b fix the same vertex of $BT(\Gamma(G, S))$.*

Proof. If $b = a^{-1}$ then clearly a and b fix the same vertex. Otherwise, $ab \neq 1$, so ab fixes exactly one vertex of $BT(\Gamma(G, S))$.

Let v_a, v_b and v_{ab} be the vertices which are fixed by a, b and ab respectively.

Since $BT(\Gamma(G, S))$ is a tree, the simple paths from v_a to v_b, v_b to v_{ab} and v_{ab} to v_a intersect at some point w. Similarly, the simple paths from v_a to v_b, v_b to bv_{ab} and bv_{ab} to v_a intersect at some point w'.

Let d be the path metric in $BT(\Gamma(G, S))$. Then

$$d(v_b, bv_{ab}) = d(bv_b, v_{ab}) = d(v_b, v_{ab}), \quad \text{and} \quad d(v_a, bv_{ab}) = d(av_a, abv_{ab}) = d(v_a, v_{ab}).$$

Therefore,

$$\begin{aligned}
d(v_a, w) &= \frac{1}{2}(d(v_a, w) + d(v_{ab}, w) + d(v_a, w) + d(v_b, w) - d(v_{ab}, w) - d(v_b, w)) \\
&= \frac{1}{2}(d(v_a, v_{ab}) + d(v_a, v_b) - d(v_{ab}, v_b)) \\
&= \frac{1}{2}(d(v_a, bv_{ab}) + d(v_a, v_b) - d(bv_{ab}, v_b)) \\
&= \frac{1}{2}(d(v_a, w') + d(bv_{ab}, w') + d(v_a, w') + d(v_b, w') - d(bv_{ab}, w') - d(v_b, w')) \\
&= d(v_a, w').
\end{aligned}$$

Moreover, w and w' are both on the simple path between v_a and v_b, so we must have $w = w'$.

But b sends the path between v_b and v_{ab} to the path between v_b and bv_{ab}. Since w is on both paths, it is also the same distance along both paths, so b fixes w. Therefore, $w = v_b$. Similarly, $w = v_a$, so $v_a = v_b$. □

Theorem 3.24. *An amalgamated free product of finite groups is basic if and only if the resultant group is also finite or the amalgamation is trivial.*

Proof. If the resultant group is finite, then clearly it is basic. If the amalgamation is trivial, then the resultant group is just the free product of two finite groups, so it is basic.

Now assume $G = A *_H B$ is infinite, where A and B are finite and H is non-trivial. Let S be any finite generating set for G. Then we just need to show that some block in $\Gamma(G, S)$ is infinite.

Let $h \in H \setminus \{1\}$ be an arbitrary element. Then h has finite order, so it fixes exactly one vertex v of $BT(\Gamma(G, S))$. Now if $a \in A$, then a, h and ah are all in A, so they all have finite orders. Therefore a and h fix the same vertex of $BT(\Gamma(G, S))$, so a fixes v. Therefore, all of A fixes v. Similarly, all of B fixes v.

Now, since G is generated by A and B, v must be fixed by all of G. We know that G acts freely on the vertices of $BT(\Gamma(G, S))$ which correspond to vertices of $\Gamma(G, S)$, so v must correspond to some block B in $\Gamma(G, S)$. Since B is fixed by all of G, it must contain every element of G, hence B is infinite. □

Note that an amalgamated free product $A *_H B$ is finite if and only if $H = A$ or $H = B$, so the previous theorem could be restated as: an amalgamated free product $A *_H B$ of finite groups is basic if and only if H is either A, B or the trivial group.

Theorem 3.25. *A HNN extension of a finite group G is basic if and only if the subgroup used for the HNN extension is trivial.*

Proof. Let $\alpha : H \to K$ be the isomorphism between subgroups of G, so that $G*_\alpha$ is the HNN extension. Let t be the generator which is not in G. If the subgroups H and K used are trivial, then $G*_\alpha$ is equal to $G * \langle t \rangle$, which is also basic.

Now assume that H and K are not trivial and let S be any finite generating set for $G*_\alpha$. Then we just need to prove that $\Gamma(G, S)$ contains an infinite block.

Since G is finite, we know that G must fix exactly one vertex v of $BT(\Gamma(G*_\alpha, S))$. Let $h \in H \setminus \{1\}$ be an arbitrary element. Then $tht^{-1} \in K \subset G$, so tht^{-1} fixes v. Also, $tht^{-1}(tv) = thv = tv$, so tht^{-1} fixes tv. But tht^{-1} can only fix one vertex in $BT(\Gamma(G*_\alpha, S))$, so $tv = v$. Hence, t fixes v. But $G*_\alpha$ is generated by G and t, so v must be fixed by all of $G*_\alpha$.

We know that $G*_\alpha$ acts freely on the vertices of $BT(\Gamma(G*_\alpha, S))$ which correspond to vertices of $\Gamma(G*_\alpha, S)$, so v must correspond to some block B in $\Gamma(G*_\alpha, S)$. Since B is fixed by all of $G*_\alpha$, it must contain every element of $G*_\alpha$, so B is infinite. ∎

Chapter 4

Some applications of Bass-Serre theory

4.1 Introduction to Bass-Serre theory

Bass-Serre theory is a general study involving graphs of groups, the main point being to find decompositions of groups acting on trees via the fundamental graph of groups of the action. The main references are [25] and [2].

Given a group G acting on a tree T, we can construct the quotient graph of groups $G \setminus T$. The main theorem of Bass-Serre theory is that the fundamental group of the graph of groups $G \setminus T$ is isomorphic to G. The process can go backwards as well: Given a graph of groups A with fundamental group G, we can construct the Bass-Serre tree T of A, which comes with an action of G on T. Moreover, this tree satisfies $G \setminus T = A$. This generalises the notion which we mentioned in the introduction of the Bass-Serre tree of amalgamated free products and HNN-extensions.

Definition 4.1. A graph of groups A is a directed connected graph equipped with the following information:

- There is a fixed point free involution of the edges of A given by $e \mapsto \overline{e}$, such that if u and v are vertices in A and e is an edge in A from u to v, then $\overline{e}$ is an edge from v to u.
- For each vertex v in A there is an associated vertex group G_v.
- For each edge e in A there is an associated edge group $G_e = G_{\overline{e}}$.
- For each edge e in A from u to v, there is an associated injective group homomorphism $f_e : G_e \to G_v$. So, since $G_e = G_{\overline{e}}$, there is also a homomorphism $f_{\overline{e}} : G_e \to G_u$.

Since each edge e has an associated reverse edge $\overline{e}$, we have a natural undirected graph K_A associated with A which has the same vertex set as A and has one edge between two vertices u and v in A for each pair of edges $\{e, \overline{e}\}$ in A where e goes from u to v and hence $\overline{e}$ goes from v to u. We will call A a finite graph of groups if K_A is a finite graph.

We will now describe a way to construct the fundamental group of a finite graph of groups. There is an analogous notion for infinite graphs of groups but it will not be necessary for our purposes.

Definition 4.2. We construct the fundamental group of a finite graph of groups A inductively as follows:

If A contains a single vertex v and no edges, then the fundamental group of A is G_v.

If K_A contains a vertex with degree 1, let the vertex be v, let e be the edge in A that goes out from v and let A' be the graph of groups consisting of everything in A except for v and the edges which connect to v (which are just e and $\overline{e}$). Let u be the vertex at the other end of e from v. Let G' be the fundamental group of A'. Then f_e gives an inclusion map of G_e into G_v and $f_{\overline{e}}$ gives an inclusion map of G_e into G_u which is contained in G'. We define the fundamental group G of A to be the amalgamated free product:

$$G = G' *_{G_e} G_v.$$

Finally, if K_A contains no vertices with degree 1 then K_A must not be a tree, so there must be some edge in K_A whose removal does not disconnect the graph. Let e and $\overline{e}$ be the associated edges in A and let u and v be the vertices in A such that e is an edge from u to v.

Let A' be the graph of groups obtained by removing e and $\overline{e}$ from A and let G' be the fundamental group of A'. Then f_e and $f_{\overline{e}}$ give two different inclusion maps of G_e into G. Therefore, we have an isomorphism $\phi : \mathrm{im}(f_e) \to \mathrm{im}(f_{\overline{e}})$ between subgroups of G' given by $\phi = f_{\overline{e}} \circ f_e^{-1}$. We define the fundamental group G of A to be the HNN extension:

$$G = G' *_\phi .$$

Note that there were many choices in how we could have chosen the edges to remove to build up the fundamental group of G, however an important fact in Bass-Serre theory is that the fundamental group obtained is independent, up to isomorphism, of the choices made.

Finally, we will give the definition of a quotient graph of groups of a group action on a tree:

Definition 4.3. Let G be a group which acts on a tree T without edge inversion. The quotient graph of groups $A = G \setminus T$ is given by the following data:

- The underlying graph K_A is the quotient graph of the action.
- We define each vertex group G_v to be the stabiliser of a vertex in the vertex orbit of the action of G on T corresponding to v. Note then that G_v is unique up to isomorphism.
- We define the edge group $G_e = G_{\overline{e}}$ to be the stabiliser of an edge in the edge orbit of the action of G on T corresponding to e. Note then that G_e is unique up to isomorphism.

- If u, v are vertices in A and e is an edge from u to v, then e corresponds to a directed edge e' in T from u' to v', where u', v' and e' are elements of the orbits corresponding to u, v and e respectively. Since G acts without edge inversion, the stabiliser G_e of e' stabilises e' as a directed edge. Therefore, the stabiliser G_e of e' is a subgroup of the stabiliser G_v of v'. This gives us a natural inclusion map $f_e : G_e \to G_v$.

  Note that even if u and v are the same vertex, the maps f_e and $f_{\overline{e}}$ may be different since u' and v' may be different vertices in the same orbit.

A theorem of Dunwoody [9] states that all finitely generated groups are accessible, which, in the language of Bass-Serre theory, means that any finitely generated group G is the fundamental group of a finite graph of groups where each edge group is finite and each vertex group does not split as an amalgamated free product or HNN extension over a finite subgroup. Together with Stallings theorem, this means that any finitely generated group is given by a fundamental group of a finite graph of groups where each vertex group has at most one end. We have already stated the theorem that such a group is virtually free if and only if every vertex group is finite; therefore, a finitely generated group is virtually free if and only if it does not contain a one-ended subgroup.

4.2 Applications of Bass-Serre theory to basic groups

In this section we will finish off the cycle of implications involving basic groups which shows that for a finitely generated group G the following are equivalent:

(a) G is basic.

(b) There is a finite generating set S for G such that every block in $\Gamma(G, S)$ is finite.

(c) G acts edge freely and with finitely many orbits on some non-trivial, locally finite tree T.

(d) G is the fundamental group of a finite graph of finite groups with trivial edge groups.

We proved that (a)⇒(b) and (b)⇒(c) in the previous chapter, so in this section we will show that (c)⇒(d) and (d)⇒(a)

Theorem 4.4 ((c)⇒(d)). *If G acts edge freely and with finitely many orbits on some nontrivial, locally finite tree T, then G is the fundamental group of a finite graph of finite groups with trivial edge groups.*

Proof. Let $Y = G \setminus T$ be the quotient graph of groups of the action of G on T. Then we know that the fundamental group of Y is G, since T is a tree. Since each edge in T has a trivial stabiliser, each edge group in Y is trivial. Since T is non-trivial, each degree is finite and non-zero. Also, since the edge stabilisers are trivial, the vertex stabilisers must all be finite. Therefore, Y is a graph of finite groups. Finally, since the action of G on T has finitely many orbits, Y is a finite graph of groups. Therefore, Y is a finite graph of finite groups with trivial edge groups, and G is the fundamental group of Y. □

Theorem 4.5 ((d)⇒(a))**.** *If G is the fundamental group of a finite graph of finite groups A with trivial edge groups, then G is basic.*

Proof. Since the edge groups are trivial, each amalgamated free product in the definition of the fundamental group will just be a free product with a vertex group and each HNN extension will be a free product with $\mathbb{Z}$. Therefore, the fundamental group will be the free product of the vertex groups and a finite number of copies of $\mathbb{Z}$. Since each vertex group is finite, the fundamental group of A is the free product of finitely many basic groups, hence it is basic. □

So we have finished the cycle of implications involving basic groups.

Given a group which is basic, we can show that it is basic by simply finding a way to write it as a free product of finite groups. If a group G is not basic, it is less obvious how one goes about proving that G is not basic. If G is not virtually free either, then we can simply find a one-ended subgroup to show that it is not virtually free, and hence not basic either. So the remaining case is when we have a virtually free group which is not basic. The following theorem gives us a way to prove that such a group is not basic:

Theorem 4.6. *If G is finitely generated and virtually free, then G is non-basic if and only if it contains an infinite subgroup H which is either a non-trivial amalgamated free product of finite groups or a non-trivial HNN extension of a finite group.*

Proof. From 3.24 and 3.25 we know that such a subgroup H is not basic, so if G contains such a subgroup then it is not basic.

Now assume that G is finitely generated and virtually free but not basic. We just need to show that G contains an infinite subgroup which is either a non-trivial amalgamated free product of finite groups or a non-trivial HNN extension of a finite group.

Since G is virtually free, it is the fundamental group of some finite graph of finite groups J. We will assume that J has the minimal number of edges for any graph of groups with this property. Now, since G is not basic, there must be some non-trivial edge group G_e in J. Consider the subgraph of groups K containing the edge e and each vertex which is attached to e. If the vertices at the ends of e are distinct then the fundamental group H of K is an amalgamated free product of the vertex groups, whereas if both ends of e are attached to the same vertex, then H is a HNN extension of the vertex group. Moreover, since H is the fundamental group of a subgraph of groups of J, it is a subgroup of the fundamental group G of J, so we just need to show that H is infinite.

Suppose for the sake of contradiction that H is finite. Then we can simply remove the edge and identify the vertices at the two ends of it, then make H the vertex group of the new vertex. This creates a new finite graph of finite groups with fewer edges than J which has the same fundamental group as J, contradicting the minimality of J. □

4.3 Characterisation of blocks of Cayley graphs

In this section we will use Bass-Serre theory to finish the characterisation of blocks in Cayley graphs which we started in the previous chapter. Theorem 3.15 already tells us exactly what blocks B can occur in a Cayley graph, but we also want to know what we can say about the group based on these blocks and what if anything one block in a Cayley graph implies about the other blocks in the Cayley graph.

Proposition 4.7. *Let G be a group with finite generating set S such that any two blocks in $\Gamma(G,S)$ are isomorphic as labelled graphs. Let B be a block in $\Gamma(G,S)$, let H be the subgroup of G which stabilises B and let k be the number of vertex orbits of the action of H on B. Then G is isomorphic to $H * F_{k-1}$.*

Proof. Let $Y = G \setminus BT(\Gamma)$ be the quotient graph of groups of the action of G on $BT(\Gamma(G,S))$. Since $BT(\Gamma)$ is a tree, G is the fundamental group of Y.

Let v_0 be a vertex in B. Then v_0 is a vertex in Γ and hence a vertex in $BT(\Gamma)$. If $v_1 \in BT(\Gamma)$ then either v_1 corresponds to a vertex in Γ or a block in Γ. If v_1 corresponds to a vertex in Γ, then it is in the orbit of v_0. If v_1 corresponds to a block in Γ, then the block is isomorphic to B as a labelled graph, so it is in the orbit of B. Therefore, $BT(\Gamma)$ has only two vertex orbits, so Y has only two vertices. Let v be the vertex in Y which corresponds to the orbit of v_0 and let u be the vertex in Y which corresponds to the orbit of B.

Let $u_1, \dots, u_k$ be representatives for the vertex orbits of the action of H on B. We will now show that $(u_1, B), (u_2, B), \dots, (u_k, B)$ are representatives of edge orbits of the action of G on $BT(\Gamma)$. If i, j are integers with $1 \leq i, j \leq k$, and the edges (u_i, B) and (u_j, B) in $BT(\Gamma)$ are in the same orbit of the action of G, then there is some $g \in G$ which satisfies $gu_i = u_j$ and $gB = B$. Since $gB = B$, we must have $g \in H$, so u_i and u_j are in the same orbit of the action of H on B. Hence, $i = j$. In particular, this means that, $(u_1, B), \dots, (u_k, B)$ are in different edge orbits. Next we will show that every edge in $BT(\Gamma)$ is in the same orbit as one of these edges.

Let (u', B') be an edge in $BT(\Gamma)$ where u' is a vertex in Γ and B' is a block in Γ, so u' is a vertex in B'. Then there is some $g \in G$ with $gB' = B$. Then $gu' \in B$, so there exists some $h \in H$ and some integer i with $1 \leq i \leq k$ such that $hgu' = u_i$. Therefore, $hg(u', B') = (u_i, hB) = (u_i, B)$, so (u', B') is in the same orbit as (u_i, B). Therefore the action of G on $BT(\Gamma)$ has exactly k edge orbits, so there are exactly k edges in Y. Note also that each edge e in Y must go between u and v since a representative of the edge orbit corresponding to e joins B, a vertex in the orbit corresponding to u, to a vertex in the orbit corresponding to v.

Since the edge stabilisers of the action of G on $BT(\Gamma(G,S))$ are trivial, the edge groups in Y are trivial. Since the stabiliser of v_0 is trivial, the group G_v is trivial. The group G_u is equal to the stabiliser of B, which is H, so $G_u = H$. Now that we know everything about the graph of groups Y, we just need to compute its fundamental group.

For $i \in \mathbb{Z}$ with $1 \leq i \leq k$, let Y_i be the sub graph of groups of Y which contains u and v, but only contains i of the edges between u and v. We choose these in such a way that for each i, the subgraph Y_{i+1} contains Y_i. The fundamental group G_1 of Y_1 is just the amalgamated free product $H *_{\{1\}} \{1\} = H * \{1\} = H$. If G_i

is the fundamental group of Y_i, then since Y_{i+1} is obtained from Y_i by adding a single edge with trivial edge group, the group G_{i+1} is a HNN extension of G_i where the associated subgroups of the HNN extension are trivial. Hence, $G_{i+1} = G_i * \mathbb{Z}$. Therefore, by induction, we have $G_i = H * F_{i-1}$. So $G = G_k = H * F_{k-1}$. □

Theorem 4.8. *Let B be a simple 2-connected graph and let H be a subgroup of $Aut(B)$ which acts vertex freely and with k orbits on B. Then there is a finite generating set S for the group $H * F_{k-1}$ such that every block in $\Gamma(H * F_{k-1}, S)$ is isomorphic to B.*

Proof. From theorem 3.15 we know that there is a Cayley graph Γ' such that B appears as a block in Γ'. Then from theorem 3.19, we know that there is another Cayley graph Γ such that all blocks in Γ are isomorphic to each other, as labelled graphs, and they are also isomorphic to B. Let G be the associated group and let S be the associated generating set. Then all blocks of $\Gamma(G, S)$ are isomorphic to B. Now by proposition 4.7, we see that G is isomorphic to $H * F_{k-1}$, so every block of $\Gamma(H * F_{k-1}, S)$ is isomorphic to B. □

Example 4.9. Let G be a finitely generated group with generating set S such that $\Gamma(G, S)$ is 2-connected and let H be a subgroup of G with finite index k. Then from the action of G on $\Gamma(G, S)$ we get a restricted action of H on $\Gamma(G, S)$. Since H has index k in G, this action will have k vertex orbits. Moreover, since G acts vertex freely on $\Gamma(G, S)$, the subgroup H also acts vertex freely on $\Gamma(G, S)$. Therefore, there exists a generating set S' for $H * F_{k-1}$ such that every block in the Cayley graph $\Gamma(H * F_{k-1}, S')$ is isomorphic to the Cayley graph $\Gamma(G, S)$.

Example 4.10. Let G be a group and let S be a finite generating set for G. Let $f : S \to \mathbb{Z}_{>0}$ be an arbitrary function. We take the edge separated Cayley graph $\Gamma_e(G, S)$, then for each $a \in S$ we replace each edge Γ_e labelled by a with a chain of $f(a)$ edges to construct a new graph $\Gamma_e^f(G, S)$. Then G acts freely and with finitely many orbits on Γ_e. Moreover, Γ_e^f is 2-connected if and only if Γ_e is 2-connected. Therefore, if Γ_e is 2-connected then there is a positive integer k and a generating set S' for $G * F_k$ such that every block of $\Gamma_e(G * F_k, S')$ is isomorphic to Γ_e^f. Moreover, if 1 is not in the image of f, so that every edge in Γ_e is replaced by at least two edges, then Γ_e^f is a simple graph, so it is isomorphic to every block of $\Gamma(G * F_k, S')$.

Example 4.11. Let G be a group and let S be a finite generating set for G. Let $f : S \to 2\mathbb{Z}_{>0} + 1$ be an arbitrary function which assigns an odd number to each generator in S. We take the Cayley graph $\Gamma(G, S)$ then construct a new graph $\Gamma^f(G, S)$ as follows: For each generator $a \in S$, we replace each edge in Γ labelled by a with a chain of $f(a)$ edges. Clearly G acts with finitely many orbits on Γ^f. Moreover, if $x \in \Gamma^f$ is fixed by some element $g \in G$, then x corresponds to a mid-edge in Γ. Therefore, since the corresponding edge was broken into an odd number of pieces, x is a mid-edge in Γ^f. So G acts vertex freely on Γ^f. Note that Γ^f is 2-connected if and only if Γ is 2-connected. Therefore, if Γ is 2-connected then there is a positive integer k and a generating set S' for $G * F_k$ such that every block of $\Gamma(G * F_k, S')$ is isomorphic to $\Gamma^f(G, S)$.

Chapter 5

More classes of groups and Cayley graphs

In this section we discuss a number of Cayley graph properties from the literature, and also introduce some new related properties. The main questions in this section are of the form "does property K imply property J?" and "is property K independent of generating set?". There are certainly many more properties in the literature, but we will focus on properties which relate to the path metric and geodesics.

The properties which we will consider are listed below:

- Hyperbolicity
- The order of the Dehn function. Given a finite presentation for a group, we can define the Dehn function corresponding to that presentation. Two Dehn functions for the same group need not be equal, but they are in some sense equivalent. This means that the order of the Dehn function for a presentation is actually a property of the group. Possible orders include: linear, quadratic, cubic and exponential.
- Almost convexity properties: Almost convexity (AC), minimal almost convexity (MAC), M'AC and $P(2)$. We also introduce the seemingly weaker condition MAC+ and a closely related property MAC$\sim$.
- The falsification by fellow traveller property (FFTP).
- The loop shortening property (LSP) and basepoint loop shortening property (BLSP).
- We also introduce for each $k \in \mathbb{Z}_{\geq 2}$, the k-decomposability property (k-DP). We say that a graph has the bounded decomposability property (BDP) if it is k-decomposable for some $k \in \mathbb{Z}_{\geq 2}$.

Definition 5.1. We call a property K for graphs a *block property* if it is the case that a graph Γ, which has finitely many non-isomorphic blocks $B_1, B_2, \ldots B_n$, has property K if and only if every block of Γ has property K.

Most of the properties which we consider are block properties and as a result, our characterisation of blocks in Cayley graphs is quite useful for finding counter examples to some of the questions about these properties.

Given a graph Γ, we contruct the m-partition Γ^m of Γ by replacing each edge of Γ with a chain of m edges. Then there is a bijection $f : \Gamma \to \Gamma^m$ which multiplies distances by m. Note that while f sends all vertices of Γ to vertices in Γ^m, not all vertices of Γ^m will be images of vertices in Γ. Since Γ^m has essentially the same geometry as Γ, it is natural to ask which properties are invariant under this transformation.

Definition 5.2. We will say that a graph property K is *independent of m-partition* if it is the case that an edge separated Cayley graph Γ_e has the property K if and only if Γ_e^m has property K.

From example 4.10, we know that for any group G along with a finite generating set S, there is a positive integer k and a finite generating set S' for $G * F_k$ such that every block in $\Gamma_e(G * F_k, S')$ is isomorphic to some block in $\Gamma_e^m(G, S)$. So the question is particularly relevant to block properties.

Below we give a table which summarises what we know about the properties which we will consider in this chapter. Note that we also include the properties "finite block size" and "quasi isometric to some tree" which, as we have seen, are the graph properties corresponding to basic and virtually free groups respectively.

Property	Block Property	Independent of m-partition	Group property
Finite block size	yes	yes	no
Quasi isometric to some tree	yes	yes	yes
Hyperbolic	yes	yes	yes
Order of Dehn function	yes	yes	yes
Almost convex	yes	yes	no [32]
P(2)	yes	yes	?
M'AC	yes	iff = MAC$\sim$	no
MAC	yes	iff = MAC$\sim$	no
MAC+	yes	iff = MAC$\sim$	no
MAC$\sim$	yes	yes	no
FFTP	yes	yes	no[21]
LSP	?	yes	no
BLSP	yes	yes	no
k-DP (for fixed k)	yes	yes	no
BDP	yes	yes	?

Now we give an implication diagram, which shows all of the known implications between these properties for Cayley graphs. There are almost certainly many implications which are not shown in this diagram, for example it seems very likely that "regular language of geodesics" implies MAC.

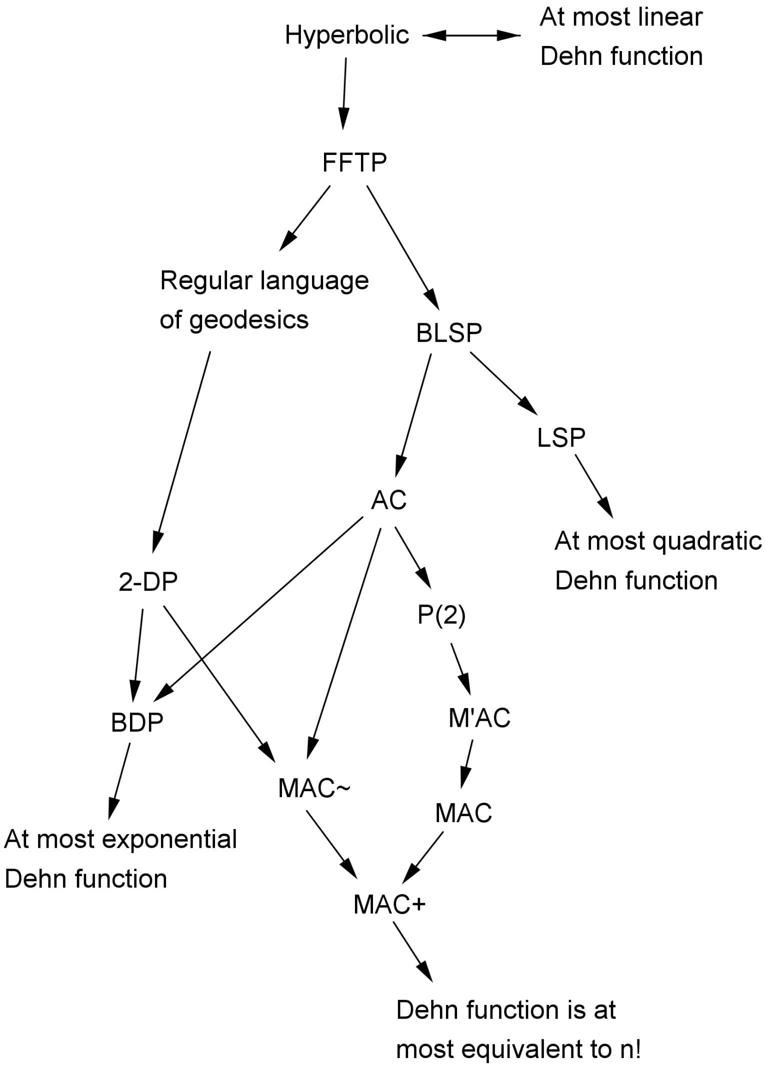

5.1 Hyperbolic groups

Definition 5.3. Let Γ be a graph. We say that Γ is δ-hyperbolic if it has the following property:

If x, y and z are points in Γ, and γ_1, γ_2 and γ_3 are the sides of a geodesic triangle with vertices x, y and z, then any point $u \in \gamma_1$ is within a distance of δ of one of the other two sides. We say that Γ is (Gromov) hyperbolic if it is δ-hyperbolic for some δ.

If G is a group and S is a generating set for G such that $\Gamma(G, S)$ is hyperbolic, then G is called a hyperbolic group. In particular, this property is independent of the generating set. Note that if Γ is a δ-hyperbolic graph, then any two geodesics l_1 and l_2 between the same pair of point forms a geodesic bigon, which is just a geodesic triangle where one side has length 0, therefore, the Hausdorff distance between l_1 and l_2 is at most δ.

Proposition 5.4. *Hyperbolicity is a block property.*

Proof. If a graph Γ is hyerbolic then every block is certainly hyperbolic. Now assume that every block of Γ is δ-hyperbolic. Let γ_1, γ_2 and γ_3 be sides of a geodesic triangle in Γ and let u be a point on γ_1. Then u is either contained in a simple geodesic bigon or a simple geodesic triangle whose sides are subpaths of γ_1, γ_2 and γ_3. Since this geodesic polygon is simple it is contained in a single block B. Therefore, since B is δ-hyperbolic, u is within a distance of δ of one of the sides γ_2 or γ_3. □

Proposition 5.5. *Hyperbolicity is independent of m-partition.*

Proof. If Γ is a graph, then the geodesic triangles in Γ correspond exactly to the geodesic triangles in the m-partition Γ^m. Hence, Γ is δ-hyperbolic if and only if Γ^m is $m\delta$-hyperbolic. □

5.2 The Dehn function

Definition 5.6. Let G be a group with finite presentation $\langle S \mid R\rangle$. Then for any word w over the alphabet $S \cup S^{-1}$ with $w =_G 1$, we can think of w as representing an element of the free group F_S and write w as a product of conjugates of elements of R:

$$w =_{F_S} f_1 r_1 f_1^{-1} f_2 r_2 f_2^{-1} \ldots f_n r_n f_n^{-1},$$

where each $r_i \in R$ and each $f_i \in F_S$. The area $\text{Area}(w)$ of the loop w is the minimum possible value of n in this product.

Definition 5.7. Let G be a group with finite presentation $\langle S \mid R\rangle$. The Dehn function $\text{Dehn} : \mathbb{Z}_{>0} \to \mathbb{R}_{>0}$ of G with respect to the presentation $\langle S \mid R\rangle$, is defined by

$$\text{Dehn}(n) = \max\{\text{Area}(w) | w =_G 1, |w| \leq n\}.$$

While the Dehn function is certainly not independent of the generating set, or even the relators, it is independent up to the following equivalence:

Definition 5.8. Two (non-strictly) increasing functions f and g are called equivalent if there is some constant $c \in \mathbb{R}_{>1}$ such that

$$\frac{1}{c} g\left(\frac{1}{c}x - c\right) - cx - c \leq f(x) \leq cg(cx + c) + cx + c$$

For example, two polynomials are equivalent if and only if they have the same degree, and any two exponential functions are equivalent.

Since the Dehn function is independent of this equivalence, we can partition finitely presentable groups into sets based on the order of their Dehn functions. It is known that a group has linear Dehn function if and only if it is hyperbolic [13], and no group has a Dehn function which lies strictly between linear and quadratic [23]. Surprisingly, however, very little is known about Cayley graphs with merely quadratic isoperimetric inequalities, indeed almost all of the properties which we will consider in this chapter fail for some group with a quadratic Dehn function.

The Dehn function is closely related to the complexity of the word problem. In particular, the Dehn function for a group G is a computable function if and only if the word problem problem for G is decidable.

We will now define the Dehn function for graphs in general, in a way which is equivalent for Cayley graphs:

Definition 5.9. Given a loop directed l in a graph Γ, we define a (c,k)-decomposition of l to be a pair (P,f) where p is a planar graph has at most k inner faces and each one has degree at most c and $f : P \to \Gamma$ is a graph homomorphism which sends the anti-clockwise loop around outside of P to the loop l. Note that this means that the outer face of P has degree $|l|$.

Then, given a constant c_0, we define the area of a loop l in Γ to be the minimum k such that there exists a (c_0,k)-decomposition of l in Γ. We now define the Dehn function for a graph Γ as follows:

$$\mathrm{Dehn}(n) = \max\{\mathrm{Area}(l) | l \text{ is a loop in } \Gamma, |l| \leq n\}.$$

Then, for a Cayley graph $\Gamma(G,S)$, this will be equal to the Dehn function given by the presentation $\langle S \mid R\rangle$, where we let R be the set of words w of length at most c with $w =_G 1$. Note that the area function will only be finite for all loops l if $\langle S \mid R\rangle$ forms a presentation for G, so we must insist that c is sufficiently large for this to be the case. Similarly for graphs in general, we insist the c is sufficiently large that the area function is always finite. Although this is not always possible, we will only define the Dehn function in cases where it is.

Proposition 5.10. *Let Γ be a graph and let c_0 be a constant such that the Dehn function of Γ with respect to c_0 well defined. If l is a loop in Γ which is (c,k)-decomposable, then*

$$Area(l) \leq kDehn(c)$$

Proof. Let (P,f) be a (c,k)-decomposition of l. Then the image of the loop around any inner face F_j of P is a loop l_j with $|l_j| \leq c$. Therefore, l_j is $(c_0, \mathrm{Dehn}(c))$-decomposable, so let (P_j, f_j) be a $(c_0, \mathrm{Dehn}(c))$-decomposition of l_j. Then we can fill in the face F_j with the planar graph P_j, thereby replacing it with at most $\mathrm{Dehn}(c)$ faces, each with degree at most c_0. Now we do this with every inner face of P to obtain a $(c_0, k\mathrm{Dehn}(c))$-decomposition of l.

Note that when we fill in a face F_j of P we may need to identify some vertices on the outside of this face. This is fine as long as the loop l_j around the outside of F_j is simple. If l_j is not simple, then we can simply remove everything in the region enclosed by F_j before proceeding to fill it in. □

We already stated an equivalence between Dehn functions for Cayley graphs which represent the same group, and as one would expect, Dehn functions for graphs satisfy the same equivalence. In fact, they satisfy a stronger equivalence:

Proposition 5.11. *Let Γ be a graph and let c_0 and c_1 be constants with $c_0 \leq c_1$ such that the Dehn function of Γ with respect to both of these constants is well* *defined. Let $Dehn_0$ and $Dehn_1$ be these two Dehn functions. Then there is some constant $t \in \mathbb{R}_{>1}$ such that*

$$Dehn_0(x) \geq Dehn_1(x) \geq \frac{1}{t} Dehn_0(x) \qquad \forall x \in \mathbb{Z}_{>0}$$

Proof. The first inequality is clear since $c_1 \geq c_0$. For the second inequality, let Area_0 denote that area function with respect to c_0 and let Area_1 denote that area function with respect to c_1. Let $t = \text{Dehn}_0(c_1)$. Then we just need to prove that for any loop l we have

$$t\text{Area}_1(l) \geq \text{Area}_0(l).$$

Since l is $(c_1, \text{Area}_1(l))$-decomposable, this follows immediately from proposition 5.10. □

Proposition 5.12. *Let Γ be a graph and let l be a loop in Γ which is contained in a single block B and which is (c, k)-decomposable. Then l is (c, k)-decomposable in B.*

Proof. Let (P, f) be a (c, k)-decomposition of l where the number of edges of P is minimal. We will show that the image of f is contained in B.

Suppose for the sake of contradiction that the image of f is not contained in B. Then the space $f^{-1}(\Gamma \setminus B)$ is non-empty. Let M be a connected component of $f^{-1}(\Gamma \setminus B)$. Then the image of the boundary of M is a single vertex v in Γ. Let P' be the connected component of $P \setminus M$ which contains the outer loop of P. Then the boundary of P' is contained in the boundary of M. Hence, the image of the boundary of P' under f is $\{v\}$. Therefore, we can identify all of the vertices in the boundary of P' to obtain a new (c, k)-decomposition (P', f) of l. But P' has fewer edges than P, a contradiction.

Therefore, the image $f(P)$ is contained in B, so (P, f) is a (c, k)-decomposition of l in B. □

Proposition 5.13. *Let $f : \mathbb{Z}_{>0} \to \mathbb{Z}_{>0}$ be a super additive function (so $f(a+b) \geq f(a)+f(b)$ $\forall a, b \in \mathbb{Z}_{>0}$) and let Γ be a graph with only finitely many non-isomorphic blocks. The Dehn function for Γ is at most equivalent to f if and only if the Dehn function for every block of Γ is at most equivalent to f.*

Proof. Let $c \in \mathbb{Z}$ be a constant such that the Dehn function for Γ and every block of Γ are well defined with respect to c. All Dehn functions in this proof will be using this constant c. First we will Show that for any block B in Γ and any $x \in \mathbb{Z}_{>0}$, we have $\text{Dehn}_\Gamma(x) \geq \text{Dehn}_B(x)$.

So we just need to show that if l is a loop in B, then l is $(c, \text{Dehn}_\Gamma(|l|))$-decomposable in B. By the definition of the Dehn function, we know that l is $(c, \text{Dehn}_\Gamma(|l|))$-decomposable in Γ. Hence, by proposition 5.12, l is (c, k)-decomposable in B. Therefore, if the Dehn function for Γ is at most equivalent to f, then the Dehn function for every block of Γ is at most equivalent to f.

Now assume that the Dehn function for every block of Γ is at most equivalent to f. Then there is some super additive function g such that the Dehn function for every block in Γ is at most equal to g. Any simple loop l in Γ is contained in a single block B, so its area is at most $g(|l|)$. Now if l' is any loop in Γ which is not simple, we can split l it into two non-empty loops l_1, l_2 with $|l_1| + |l_2| = |l|$. Hence,

$$\text{Area}(l) \leq \text{Area}(l_1) + \text{Area}(l_2) \leq \text{Dehn}_\Gamma(|l_1|) + \text{Dehn}_\Gamma(|l_2|).$$

Therefore, since g is super additive, we can show inductively that $\text{Dehn}_\Gamma \leq g$. Hence, the Dehn function on Γ is at most equivalent to f. □

Proposition 5.14. *Let Γ be a graph and let $Dehn_0$ be the Dehn function of Γ with respect to some constant c_0. Let $m \in \mathbb{Z}_{>0}$ and let $Dehn_1$ be the Dehn function of Γ^m with respect to the constant mc_0. Then for $x \in \mathbb{Z}_{>0}$, we have $Dehn_0(x) = Dehn_1(mx)$.*

Proof. Let $f : \Gamma \to \Gamma^m$ be the natural bijection which multiplies lengths by m. If l is a loop in Γ, with $|l| \leq x$ then $|f(l)| \leq mx$, so there is some $(mc_0, \text{Dehn}_1(mx))$-decomposition (P', g') of $f(l)$ in Γ^m. We will assume without loss of generality that the number of edges in P' is minimal amongst all $(mc_0, \text{Dehn}_1(mx))$-decompositions of $f(l)$. Since the number of edges in P' is minimal, there cannot path of length 2 in P' made of an edges e_1, e_2, such that $g'(e_1)$ is the same as the edge $g'(e_2)$ reversed. Hence there is some planar graph P such that P' is isomorphic to P^m. Moreover, $(P, f^{-1} \circ g' \circ f_1)$ is a decomposition of l, where $f_1 : P \to P^m$ is the natural bijection which multiplies lengths by m. Hence we have a $(c_0, \text{Dehn}_1(mx))$-decomposition of l in Γ. Since this works for any loop l with $|l| \leq x$, we have

$$\text{Dehn}_0(x) \leq \text{Dehn}_1(mx).$$

To show the other inequality, let l' a loop which satisfies $\text{Area}_1(l') = \text{Dehn}_1(mx)$ and the length of l' is minimal. Then $|l'| \leq mx$. Since the length of l' is minimal, it must not contain an edge followed immediately by the same edge reversed. Hence, there is a corresponding loop l in Γ such that $l' = f(l)$. Since $|l| \leq x$, there must be some $(c_0, \text{Dehn}_0(x))$-decomposition (P, g) of l in Γ. Then we have a $(mc_0, \text{Dehn}_0(x))$-decomposition $(P^m, f \circ g \circ f_1^{-1})$ of l in Γ, where $f_1 : P \to P^m$ is the natural bijection which multiplies lengths by m. Hence,

$$\text{Dehn}_1(mx) \leq \text{Dehn}_0(x).$$

□

5.3 Almost convexity and related classes

Definition 5.15. Let G be a group with finite generating set S and let $Sph(r)$ and $B(r)$ denote the Sphere and Ball respectively of radius r around 1 in $\Gamma(G, S)$. Let $f : \mathbb{Z}_{>0} \to \mathbb{R}_{>0}$ be a function and let $r_0 \in \mathbb{Z}_{>0}$. The pair (G, S) is said to satisfy the almost convexity condition $\text{AC}_{f,r_0}(k)$ if for any integer $r > r_0$, and any pair of vertices $a, b \in Sph(r)$ which satisfy $d(a, b) \leq k$, there is a path inside the closed ball $\overline{B(r)}$ between a and b which has length at most $f(r)$.

In order to ask whether such properties are block properties, we will need to define them for graphs in general:

Definition 5.16. Let Γ be a graph and let $Sph(v, r)$ and $B(v, r)$ denote the Sphere and Ball respectively of radius r around the point v in Γ. Let $f : \mathbb{Z}_{>0} \to \mathbb{R}_{>0}$ be a function and let $r_0 \in \mathbb{Z}_{>0}$. We say that a graph Γ satisfies the almost convexity condition $\mathrm{AC}_{f,r_0}(k)$ if for any vertex $v \in \Gamma$, any integer $r > r_0$, and any pair of vertices $a, b \in Sph(v, r)$ which satisfy $d(a, b) \leq k$, there is a path in the closed ball $\overline{B(v, r)}$ between a and b which has length at most $f(r)$.

For a Cayley graph $\Gamma(G, S)$, the only difference between these two definitions is that in the second definition we insist that the ball and sphere can be centred at any vertex in Γ, whereas in the first definition they are centred at the identity. Since any Cayley graph $\Gamma(G, S)$ is vertex transitive, these definitions are equivalent. Hence, the pair (G, S) satisfies the almost convexity condition $\mathrm{AC}_{f,r_0}(k)$ if and only if the Cayley graph $\Gamma(G, S)$ satisfies the same almost convexity condition $\mathrm{AC}_{f,r_0}(k)$.

Now we will define a few special cases of this condition. For each one we will define the condition for a general graph, though the standard definition is the analogous definition for a pair (G, S).

Definition 5.17 (Almost Convexity)**.** For a given $k \in \mathbb{Z}_{\geq 1}$, the graph Γ is said to be $\mathrm{AC}(k)$ if there is some $r_0 \in \mathbb{Z}_{>0}$ and some constant function $f : \mathbb{Z}_{>0} \to \mathbb{R}_{>0}$ such that Γ satisfies the almost convexity condition $\mathrm{AC}_{f,r_0}(k)$.

This property was introduced by Cannon in [5], where he also shows that the condition $\mathrm{AC}(k)$ is equivalent for any value $k \geq 2$. So we simply say that Γ is almost convex (AC) if it satisfies $\mathrm{AC}(k)$ for some $k \geq 2$.

The following property was introduce by poenaru in [24].

Definition 5.18 (P(2))**.** The graph Γ is said to satisfy Poenaru's $P(2)$ condition if there is some $r_0 \in \mathbb{Z}_{>0}$ and some sublinear function $f : \mathbb{Z}_{>0} \to \mathbb{R}_{>0}$ such that Γ satisfies the almost convexity condition $\mathrm{AC}_{f,r_0}(2)$

Note that if $f : \mathbb{Z}_{>0} \to \mathbb{R}_{>0}$ is given by $f(r) = 2r$, then any graph Γ trivially satisfies the property $AC_{f,1}$ since any two points u, v in the sphere $Sph(r)$ are joined by a path of length $2r$ which passes through the vertex corresponding to the identity in G. Hence minimal almost convexity is defined by the weakest possible non-trivial restriction on f:

Definition 5.19 (MAC)**.** The graph Γ is said to be minimally almost convex (MAC) if there is some $r_0 \in \mathbb{Z}_{>0}$ such that Γ satisfies the almost convexity condition $\mathrm{AC}_{f,r_0}(2)$ where $f : \mathbb{Z}_{>0} \to \mathbb{R}_{>0}$ is given by $f(r) = 2r - 1$.

Definition 5.20 (M'AC)**.** The graph Γ is said to satisfy the property M'AC if there is some $r_0 \in \mathbb{Z}_{>0}$ such that Γ satisfies the almost convexity condition $\mathrm{AC}_{f,r_0}(2)$ where $f : \mathbb{Z}_{>0} \to \mathbb{R}_{>0}$ is given by $f(r) = 2r - 2$.

Due to the succesive strengthenings of these conditions, we have the implication chain

$$\mathrm{AC} \Rightarrow \mathrm{P(2)} \Rightarrow \mathrm{M'AC} \Rightarrow \mathrm{MAC}$$

The last two conditions M'AC and MAC are discussed in greater detail in [12], where Elder and Hermiller show that M'AC does not imply $P(2)$. It is still an open question, however, whether the properties M'AC and MAC are equivalent.

Proposition 5.21. *Let $r_0 \in \mathbb{Z}_{>0}$ and let $f : \mathbb{Z}_{>0} \to \mathbb{R}_{\geq 2r_0}$ be a non strictly increasing function. The almost convexity condition AC_{f,r_0} is a block property.*

Proof. First, let Γ be a graph such that every block of Γ satisfies AC_{f,r_0}. We will show that Γ also satisfies the condition AC_{f,r_0}. Let $Sph_\Gamma(v,r)$ denote the sphere of radius r in Γ centred at v and let $B_\Gamma(v,r)$ denote the ball of radius r centred at v. Then we need to show that for any $v \in V(\Gamma)$, any integer $r > r_0$, and any pair of vertices $a, b \in Sph_\Gamma(v,r)$ which satisfy $d(a,b) \leq 2$, there is a path in the closed ball $\overline{B_\Gamma(v,r)}$ between a and b which has length at most $f(r)$.

Now let p be a geodesic between a and b, so p has length at most 2. If p contains a vertex which lies inside the open ball $B_\Gamma(v,r)$, then p must have length 2 and must lie entirely inside the ball, so we have a path in $\overline{B_\Gamma(v,r)}$ between a and b which has length $2 \leq f(r)$, so we are done.

Otherwise, p lies entirely outside $B_\Gamma(v,r)$. Let g_1 be a geodesic between v and a and let g_2 be a geodesic between v and b. Then the only point on g_1 which is outside $B_\Gamma(v,r)$ is a and the only point on g_2 which is outside $B_\Gamma(v,r)$ is b. Now let Let u be the point of intersection between g_1 and g_2 which is furthest from v and let g_1' and g_2' be the sections of g_1 and g_2 which are not between u and v. So g_1' is a geodesic between u and a and g_2' is a geodesic between u and b. Now g_1', p and g_2' form a geodesic triangle between the vertices u, a and b. Moreover, these geodesics only intersect each other at these vertices. Therefore, this geodesic triangle forms a simple cycle, so it is contained in a single block B, so $u, a, b \in V(B)$.

Let $r_1 = d(u,a)$. Then $a, b \in Sph_B(u, r_1)$. Now, by our assumption, B satisfies the condition AC_{f,r_0}, so if $r_1 > r_0$ then there is a path l in $\overline{B_B(u,r_1)}$ between a and b of length at most $f(r_1)$. But f is increasing, so $f(r_1) \leq f(r)$. Also, by the triangle inequality, since $d(u,v) = r - r_1$, we have $\overline{B_\Gamma(u,r_1)} \subset \overline{B_\Gamma(v,r)}$. Therefore, l is a path inside $\overline{B_\Gamma(v,r)}$ of length at most $f(r)$. So we are done.

Finally, if $r_1 \leq r_0$, then the union of g_1' and g_2' is a path of length $2r_1 \leq 2r_0 \leq f(r)$ in $\overline{B_\Gamma(v,r)}$. So we have finished this direction of the proof.

Now assume that Γ is a graph which satisfies AC_{f,r_0}. We will show that any block B in Γ satisfies AC_{f,r_0}. We just need to show that for any $v \in V(B)$, any integer $r > r_0$, and any pair of vertices $a, b \in Sph_B(v,r)$ which satisfy $d(a,b) \leq 2$, there is a path in the closed ball $\overline{B_B(v,r)}$ between a and b which has length at most $f(r)$.

Since $a, b \in Sph_B(v,r)$, we have $a, b \in Sph_\Gamma(v,r)$, so by our assumption that Γ satisfies AC_{f,r_0}, there is some simple path l in $\overline{B_\Gamma(v,r)}$ between a and b of length at most $f(r)$. Since l is a simple path between two vertices in the same block B, the entire path l must be contained in B. Therefore, since

$$\overline{B_\Gamma(v,r)} \cap B = \overline{B_B(v,r)},$$

The path l is contained in $\overline{B_B(v,r)}$. So we have a path in $\overline{B_B(v,r)}$ between a and b which has length at most $f(r)$. □

As a simple corollary, the four almost convexity conditions which we have considered so far are all block properties. We will now prove that AC is independent of m-partition. The following argument does not, however, work for any of the other almost convexity properties.

Proposition 5.22. *AC is independent of m-partition.*

Proof. Let $f : \Gamma \to \Gamma^m$ be the natural bijection which multiplies distances by m.

If Γ^m is almost convex, then it is AC($2m$) with respect to some constants $c, r_0 \in \mathbb{Z}_{\geq 1}$. If $v \in \Gamma$, $r \in \mathbb{Z}_{>r_0}$, $a, b \in Sph_\Gamma(v, r)$ and $d(a,b) \leq 2$, then

$$f(a), f(b) \in Sph_{\Gamma^m}(f(v), mr) \text{ and } d(f(a), f(b)) = 2m.$$

Hence, there is a path p inside $\overline{B_{\Gamma^m}(f(v), mr)}$ between $f(a)$ and $f(b)$ of length at most c. Therefore, $f^{-1}(p)$ is a path in $\overline{B_\Gamma(v, r)}$ of length at most $\frac{c}{m} \leq c$. Hence, Γ is almost convex.

Now we will prove the other direction. Assume that Γ is almost convex, then we just need to prove that Γ^m is almost convex. Since Γ is almost convex, it is AC(8) with respect to some constants $c, r_0 \geq 4$. Let $v \in \Gamma^m$, $r \in \mathbb{Z}_{>mr_0}$, $a, b \in Sph_{\Gamma^m}(v, r)$ and $d(a, b) \leq 2$. Then we just need to show that a and b are joined by a path of length at most $m(c+6)$ inside the ball $\overline{B_{\Gamma^m}(v, r)}$.

Let v' be a vertex in Γ such that $d(v', f^{-1}(v)) < 1$. Then

$$\begin{aligned} d(v', f^{-1}(a)) &\geq -d(v', f^{-1}(v)) + d(f^{-1}(v), f^{-1}(a)) \\ &= -d(v', f^{-1}(v)) + \frac{r}{m} \\ &> \frac{r}{m} - 1 \\ &\geq \left\lfloor \frac{r}{m} \right\rfloor - 1 > r_0 - 2 > 1. \end{aligned}$$

Let a' be a vertex in $Sph(v', \lfloor \frac{r}{m} \rfloor - 1)$ which is on a geodesic between $f^{-1}(a)$ and $f^{-1}(v)$ and let b' be defined similarly. Then

$$\begin{aligned} d(a', f^{-1}(a)) &\leq d(f^{-1}(v), f^{-1}(a)) - d(a', f^{-1}(v)) \\ &= \frac{r}{m} - d(a', f^{-1}(v)) \\ &\leq \frac{r}{m} - d(a', v') + d(v', f^{-1}(v)) \\ &< \frac{r}{m} - d(a', v') + 1 \\ &= \frac{r}{m} - \left(\left\lfloor \frac{r}{m} \right\rfloor - 1 \right) + 1 < 3. \end{aligned}$$

Similarly, $d(b', f^{-1}(b)) < 3$. Let p_1 be a geodesic between a' and $f^{-1}(a)$ and let p_2 be a geodesic between b' and $f^{-1}(b)$. Then p_1 and p_2 both have length at most 3. Hence,

$$\begin{aligned} d(a', b') &\leq d(a', f^{-1}(a)) + d(f^{-1}(a), f^{-1}(b)) + d(f^{-1}(b), b') \\ &\leq d(a', f^{-1}(a)) + d(f^{-1}(b), b') + \frac{2}{m} < 6 + \frac{2}{m} \leq 8. \end{aligned}$$

Therefore, there is a path p between a' and b' contained in the ball $\overline{B_\Gamma(v', \lfloor \frac{r}{m} \rfloor - 1)}$ which has length at most c. By the triangle inequality, this is contained in the ball $\overline{B_\Gamma(f(v), \frac{r}{m})}$, so joining p to p_1 and p_2 gives a path p' between a' and b' in $\overline{B_\Gamma(f^{-1}(v), \frac{r}{m})}$ of length at most $c + 6$. Therefore, $f(p')$ is a path between a and b contained in $\overline{B_\Gamma(v, r)}$, which has length at most $m(c+6)$. ∎

We will now define two new properties MAC+ and MAC∼ which are closely related to MAC and M'AC. The property MAC∼ is nice because it is independent of m-partition, however it is not clear whether either of the properties MAC∼ or MAC implies the other. MAC+ is useful because it is implied by both MAC and MAC∼ and still satisfies the important properties of the almost convexity conditions. In particular, a graph Γ which satisfies MAC+ has a Dehn function which has order at most $n!$, though we will delay the proof of this fact until we discuss k-decomposability. Hence MAC+ is in some sense an even more minimal version of minimal almost convexity. It is quite possible, however, that the four properties M'AC, MAC, MAC+ and MAC∼ are all equivalent for Cayley graphs anyway.

Definition 5.23 (MAC∼)**.** Let Γ be a graph. We will say that Γ satisfies MAC∼ if Γ^2 satisfies MAC.

Since MAC is a block property, MAC∼ is also a block property.

Definition 5.24 (MAC+)**.** Let Γ be a graph. We will say that Γ satisfies MAC+ if there is some constant $r_0 \in \mathbb{Z}_{>0}$ such that for any vertex $v \in \Gamma$ and any $r \in \mathbb{Z}$ with $r > r_0$, the following two conditions hold:

- If $a, b \in Sph(v, r)$ satisfy $d(a, b) = 1$, then there is a path in the closed ball $\overline{B(v, r)}$ between a and b which has length at most $2r - 1$.
- If $a, b \in Sph(v, r)$ are joined by a path of length 2 outside $\overline{B(v, r)}$, then there is a path p in Γ between a and b which has length at most $2r - 1$, such that every vertex in p is in the closed ball $\overline{B(v, r)}$.

In the second condition, by outside the ball $\overline{B(v, r)}$, we mean that the path only intersects $\overline{B(v, r)}$ at its endpoints a and b. In particular, this means that a vertex x, which is attached to both a and b by an edge, satisfies $d(x, v) = r + 1$.

There are only two differences between MAC and MAC+: For MAC we would insist that even when a and b are joined by a path of length 2 outside $\overline{B(v, r)}$, the entire path p lies inside $B(v, r)$, rather than just the vertices of p being forced to lie inside Γ. Moreover, such a path exists for a graph which satisfies MAC whenever $d(a, b) = 2$, not just when the path of length 2 between a and b lies outside the ball. Hence we have

$$\text{MAC} \Rightarrow \text{MAC+}.$$

The proof that MAC+ is a block property is essentially the same as the proof that the other almost convexity conditions are block properties, so we will not write it out.

Proposition 5.25. *If Γ is a graph which satisfies MAC∼ then Γ satisfies MAC+*

Proof. Since Γ satisfies MAC∼, the graph Γ^2 satisfies MAC with respect to some constant $r_0 \in \mathbb{Z}_{>0}$. Let $f : \Gamma \to \Gamma^2$ be the natural bijection which multiplies lengths by 2.

Let $v \in \Gamma$ be a vertex and let $r \in \mathbb{Z}$ satisfy $r > r_0$. Then we just need to show the two conditions in the definition of MAC+.

For the first condition, let $a, b \in Sph(v, r)$ satisfy $d(a, b) = 1$. Then we just need to show that there is a path between a and b in the closed ball $\overline{B(v, r)}$ of length at most $2r - 1$. Since $a, b \in Sph(v, r)$, we have $f(a), f(b) \in Sph(f(v), 2r)$ and $d(f(a), f(b)) = 2$. Therefore, there is a path p in $\overline{B(f(v), 2r)}$ between $f(a)$ and $f(b)$ which has length at most $4r - 1$. Therefore, $f^{-1}(p)$ is a path in $\overline{B(v, r)}$ between a and b which has length at most $2r - 1/2$. But $f^{-1}(p)$ is a path between vertices in Γ, so it must have integer length, therefore, $f^{-1}(p)$ has length at most $2r - 1$.

To prove the second condition in the definition of MAC+, let $a, b \in Sph(v, r)$ be joined by a path or length 2 which lies entirely outside the ball $\overline{B(v, r)}$. Then we just need to show that there is a path in Γ between a and b of length at most $2r - 1$, every vertex of which is inside the ball $\overline{B(v, r)}$. Since $a, b \in Sph(v, r)$, we have $f(a), f(b) \in Sph(f(v), 2r)$ and $d(f(a), f(b)) = 4$. Let $f(a), a', x, b', f(b)$ be the vertices in a path between $f(a)$ and $f(b)$ which lies outside $\overline{B(f(v), 2r)}$. Then $a', b' \in Sph_{\Gamma^2}(f(v), 2r + 1)$ and $d(a', b') = 2$, so there is a simple path p in $\overline{B(f(v), 2r + 1)}$ between $f(a)$ and $f(b)$ which has length at most $4r + 1$. Hence, if we cut off the two edges at the ends of $4p$ we get a path p' in $\overline{B(f(v), 2r + 1)}$ between $f(a)$ and $f(b)$ of length at most $4r - 1$. Therefore, $f^{-1}(p')$ is a path in $\overline{B(v, r + 1/2)}$ between a and b which has length at most $2r - 1/2$. But $f^{-1}(p)$ is a simple path between vertices in Γ, so it must be a path in $\overline{B(v, r)}$ of length at most $2r - 1$. ☐

Proposition 5.26. *If Γ is an almost convex graph, then Γ satisfies MAC$\sim$.*

Proof. Since Γ satisfies AC, and AC is independent of m-partition, Γ^2 satisfies AC. Therefore, Γ^2 satisfies MAC, so Γ satisfies MAC$\sim$. ☐

Proposition 5.27. *Let Γ be a bipartite graph and let $m \in \mathbb{Z}_{>0}$. Then Gamma satisfies MAC if and only if Γ^m satisfies MAC.*

Proof. Let $f : \Gamma \to \Gamma^m$ be the natural bijection which multiplies lengths by m. First we will assume that Γ satisfies MAC with respect to some constant r_0 and prove that Γ^m satisfies MAC using the constant $(r_0 + 2)m$.

Let v be a vertex in Γ^m, let $r \in \mathbb{Z}$ with $r > (r_0 + 2)m$ and let $a, b \in Sph_{\Gamma^m}(v, r)$ satisfy $d(a, b) \leq 2$. Then we just need to show that a and b are joined by a path in $\overline{B(f(v), r)}$ of length at most $2r - 1$. Let p_1 be a geodesic in Γ^m between v and a and let p_2 be a geodesic in Γ^m between v and b. Let p' be a path of length 2 between a and b. Let a' be the vertex on $f^{-1}(p_1)$ which is closest along $f^{-1}(p_1)$ to $f^{-1}(a)$ and let b' be defined similarly. There is at most one vertex in Γ^m between a and b on p', so there is at most one vertex in Γ between $f^{-1}(a)$ and $f^{-1}(b)$ on $f^{-1}(p')$. Hence, the distance $d(a', b') \leq 2$.

We will define a new vertex b'' in Γ. If $d(a', b') = 2$, let $b'' = b'$. If $d(a', b') = 1$, let b'' be the vertex on $f^{-1}(p_2)$ which is adjacent to b', such that b'' is closer to v than b'. Then we still have $d(a', b'') = 2$. So either way, b'' is on p_2 and $d(a', b'') = 2$.

Consider a vertex path $(a' = v_0), v_1, \ldots, v_{k-1}, (v_k = b'')$ in Γ which travels along the union of $f^{-1}(p_1)$ and $f^{-1}(p_2)$. Then, since Γ is bipartite the distances $d(v_i, a')$ and $d(v_i, b'')$ change by exactly one each time i is increased by 1. Therefore, $d(v_i, a') - d(v_i, b'')$ changes by 0 or 2 each time i changes by 1. Moreover,

$$d(v_0, a') - d(v_0, b'') = -2 \text{ and } d(v_k, a') - d(v_k, b'') = 2,$$

so there exists some j such that $d(v_j, a') - d(v_j, b'') = 0$.

Let $r' = d(v_j, b'') = d(v_j, a')$. Then $a', b'' \in Sph(v_j, r')$. If v_j is on the geodesic p_2, then

$$\begin{aligned} r' =& d(f^{-1}(v), b'') - d(v_j, f^{-1}(v)) \\ \leq& d(f^{-1}(v), f^{-1}(b)) - d(v_j, f^{-1}(v)) = \frac{r}{m} - d(v_j, f^{-1}(v)). \end{aligned}$$

Also, $d(a', b'') = 2$, so by the triangle inequality,

$$B(v_j, r') \subset B(f^{-1}(v), r/m) = f^{-1}(B(v, r)).$$

Similarly, this is true if v_j is on p_1. Now let q be the shortest path between a' and b'' inside the ball $\overline{B(v_j, r')}$. Let q' be the path made up of q and the section of p_1 between a' and $f^{-1}(a)$ and the section of p_2 between b' and $f^{-1}(b)$. Then q' is a path in $f^{-1}(\overline{B(v, r)})$ between and $f^{-1}(a)$ and $f^{-1}(b)$ and $|q'| < |q| + 2$. Hence, $f(q')$ is a path in $\overline{B(v, r)}$ between and a and b of length strictly less than $m(|q|+2)$, so we just need to show that $m(|q| + 2) \leq 2r$. If $r' \leq r_0$ then, since q has length at most $2r'$, we have

$$m(|q| + 2) \leq m(2r' + 2) \leq m(2r_0 + 2) < m(2r_0 + 4) < 2r.$$

If $r' > r_0$ then since Γ satisfies MAC using the constant r_0, we have $|q| \leq 2r' - 1$. In fact, since Γ is bipartite and $d(a', b'') = 2$, the length $|q|$ must be even, so $|q| \leq 2r' - 2$. Hence,

$$m(|q| + 2) \leq m(2r') \leq m\left(\frac{2r}{m}\right) = 2r.$$

Now we will assume that Γ^m satisfies MAC with respect to some constant $r_0 \geq 2$ and prove that Γ satisfies MAC. Let v be a vertex in Γ, let $r \in \mathbb{Z}$ with $r > r_0$ and let $a, b \in Sph_\Gamma(v, r)$ satisfy $d(a, b) \leq 2$. Then we just need to show that a and b are joined by a path in $\overline{B(v, r)}$ of length at most $2r - 1$. Since Γ is bipartite, $d(a, b)$ must be even, so we must have $d(a, b) = 2$. Let x be a vertex in Γ which is attached to both a and b by an edge. Since Γ is bipartite, we either have $d(v, x) = r - 1$ or $d(v, x) = r + 1$. If $d(v, x) = r - 1$ the the path of length 2 between a and b through x is contained in $\overline{B(v, r)}$, so we are done. Otherwise, $d(v, x) = r + 1$.

Let a' be the vertex in the image $f((a, x))$ of the edge between a and x such that $d_{\Gamma^m}(a', f(x)) = 1$ and let b' be defined similarly. Then

$$a' = b' \in Sph_{\Gamma^m}(f(v), mr + m - 1)$$

and $d(a', b') = 2$. Therefore, a' and b' are joined by some simple path p' in $\overline{B_{\Gamma^m}(f(v), mr + m - 1)}$ of length at most $2(mr + m - 1) - 1$. This path must pass through $f(a)$ and $f(b)$, so let p be the section of p' which lies between $f(a)$ and $f(b)$. Then p has length at most $2(mr + m - 1) - 1 - 2(m - 1) = 2mr - 1$. Therefore, $f^{-1}(p)$ is a path in $\overline{B_\Gamma(v, r + 1 - \frac{1}{m})}$ between a and b of length at most $2r - \frac{1}{m}$. But $f^{-1}(p)$ is a simple path between vertices of Γ, so it must be contained in $\overline{B_\Gamma(v, r)}$ and have length at most $2r - 1$. □

Proposition 5.28. *MAC$\sim$ is independent of m-partition.*

Proof. Let $m \in \mathbb{Z}_{>0}$. From the previous proposition, since Γ^2 is bipartite, Γ^2 satisfies MAC if and only if Γ^{2m} satisfies MAC. Therefore, Γ satisfies MAC$\sim$ if and only if Γ^m satisfies MAC$\sim$. Therefore, MAC$\sim$ is independent of m-partition. $\square$

Proposition 5.29. *Let Γ be a bipartite graph. If Γ satisfies any of the three properties M'AC, MAC or MAC+, then Γ satisfies all three of these properties.*

Proof. We have already shown that for any graph M'AC $\Rightarrow$ MAC $\Rightarrow$ MAC+, so we just need to show that if Γ satisfies MAC+, then it satisfies M'AC. Assume that Γ is a bipartite graph which satisfies MAC+ with respect to some constant $r_0 \geq 4$. Then we will show that Γ satisfies M'AC with respect to the same constant r_0.

Let $v \in \Gamma$ be a vertex, $r \in \mathbb{Z}_{>r_0}$ and let $a, b \in Sph(v, r)$ satisfy $d(a, b) \leq 2$. Then we just need to show that a and b are joined by a path in $\overline{B(v, r)}$ of length at most $2r - 2$. Note that since Γ is bipartite, the distance $d(a, b)$ must be even, so $d(a, b) = 2$. Moreover, the path of length 2 between a and b lies either entirely inside or entirely outside the ball $\overline{B(v, r)}$. If this path is entirely inside $\overline{B(v, r)}$ then we are done. Now we can assume that a path of length 2 between a and b lies entirely outside $\overline{B(v, r)}$. Then since Γ satisfies MAC+, the vertices a and b must be joined by a path p in Γ of length at most $2r - 1$ such that every vertex in p is in the ball $\overline{B(v, r)}$. Since Γ is bipartite, no edge in p can go between two vertices in the sphere $Sph(v, r)$, so every edge in p is in the ball $\overline{B(v, r)}$. Moreover, the length of p must be even, so $|p| \leq 2r - 2$. Therefore, p is a path in $\overline{B(v, r)}$ between a and b with length at most $2r - 2$. $\square$

Since for any graph Γ, the graph Γ^2 is bipartite, the graph Γ satisfies MAC$\sim$ if and only if Γ^2 satisfies any of the properties M'AC, MAC or MAC+.

Proposition 5.30. *MAC is independent of m-partition if and only if it is equivalent to MAC$\sim$*

Proof. If MAC is equivalent to MAC$\sim$, then it is certainly independent of m-partition. If MAC is independent of m-partition, then a graph Γ satisfies MAC if and only if Γ^2 satisfies MAC. But this happens if and only if Γ^2 satisfies MAC$\sim$. This implies that MAC is equivalent to MAC$\sim$. $\square$

Similarly, the properties M'AC and MAC+ can only be independent of m-partition if they are equivalent to MAC$\sim$.

5.4 The falsification by fellow traveller property (FFTP)

Definition 5.31. Let $k \in \mathbb{R}_{\geq 0}$ and let $u : \mathbb{R}_{>0} \to \Gamma$ and $w : \mathbb{R}_{>0} \to \Gamma$ be paths in a graph Γ, which move at the same constant velocity 1. The two paths are said to k-fellow travel if the following holds: If $t \in \mathbb{R}_{>0}$ and $t \leq m, n$ then $d(u(t), w(t)) \leq k$.

The two paths u and w are said to asynchronously k-fellow travel if there is some bijective increasing function $\phi : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ such that if $t \in \mathbb{R}_{>0}$ and $j \leq m$ then $d(u(j), w(\phi(j))) \leq k$.

We will actually be interested in when finite paths fellow travel, so to define a finite path u of length k, we will use a function $u : \mathbb{R}_{>0} \to \Gamma$, which travels at velocity 1 on $[0, k]$ and u is constant on $\mathbb{R}_{\geq k}$.

Definition 5.32. Let Γ be a graph. Γ has the synchronous falsification by fellow traveller property if there is a constant k such that for any path u in Γ which is not a geodesic, there is a shorter path w in Γ which has the same end points as u such that the paths u and w k-fellow travel. Γ has the asynchronous falsification by fellow traveller property if there is a constant k such that for any path u in Γ which is not a geodesic, there is a shorter path w in Γ which has the same end points as u such that the paths u and w asynchronously k-fellow travel.

It is fairly straightforward to show that the synchronous and asynchronous versions of this property are equivalent, so they are both simply called the falsification by fellow traveller property (FFTP). If a graph has this property, it means that one can check if a word w is a geodesic by simply checking if there are any shorter words which fellow travel with w and which represent the same element. If there is a shorter word, then w is certainly not geodesic, while if there is no such shorter word, then the falsification by fellow traveller property implies that w is geodesic. The falsification by fellow traveller property was introduced in [21], where Neumann and Shapiro show that this property depends on the generating set. Many interesting classes of groups have been shown to satisfy the falsification by fellow traveller property, in particular, abelian groups and hyperbolic groups satisfy FFTP with respect to any generating set.

Proposition 5.33. *The falsification by fellow traveller property is a block property.*

Proof. First, let Γ be a graph which satisfies the (asynchronous) falsification by fellow traveller property using the constant k and let B be a block in Γ. If l is a path in B which is not geodesic then there is a path l' in Γ which asynchronously k-fellow travels with l. Then path l'', made of just the edges of l which are in B, also k-fellow travels with l. Therefore, B satisfies FFTP.

Now for the other direction, let Γ graph with only finitely many non-isomorphic blocks, such that every block B of Γ satisfies FFTP. We will show that Γ also satisfies FFTP.

Let k be the largest constant associated to the (asynchronous) falsification by fellow traveller property in any of the blocks of Γ and let p be a path in Γ which is not a geodesic. We will show that there is a shorter path q in Γ which has the same endpoints as p such that q and p asynchronously k-fellow travel.

Since p is not a geodesic, there is some subpath p' of p which is not a geodesic and such that p' is contained entirely inside a single block B in Γ. Therefore, since B satisfies FFTP with constant k, there is a shorter path q' between the same two endpoints as p' such that p' and q' asynchronously k-fellow travel. Hence, if we let q be the path p but with the subpath p' replaced with q', then q is shorter than p and these two paths asynchronously k-fellow travel. □

Proposition 5.34. *The falsification by fellow traveller property is independent of m-partition.*

Proof. Let m be a positive integer and let Γ be a graph which satisfies the (synchronous) falsification by fellow traveller property using constant k. We will show that Γ^m satisfies FFTP using constant km.

Let $f : \Gamma \to \Gamma^m$ be the natural bijective map which multiplies distances by m.

Given a path w in Γ^m which is not geodesic, the path $f^{-1}(w)$ is also not geodesic. Therefore, there is a shorter path l in Γ with the same endpoints which k-fellow travels with $f^{-1}(w)$. It follows that $f(l)$ must mk-fellow travel with w.

Conversely, assume that Γ^m satisfies FFTP using constant k. We will show that Γ satisfies FFTP using constant k/m.

Given a path w in Γ which is not geodesic, the path $f(w)$ is also not geodesic. Therefore, there is a shorter path l in Γ^m with the same endpoints which k-fellow travels with $f(w)$. It follows that $f^{-1}(l)$ must (k/m)-fellow travel with w. □

In [21], Neumann and Shapiro show the following result. The converse, however, does not hold [11].

Proposition 5.35. *If S is a finite generating set for a group G such that (G, S) has the falsification by fellow traveller property, then the language of all geodesic words in G with respect to the generating set S is regular.*

5.5 Loop shortening properties

In [10] Elder introduced the loop shortening and basepoint loop shortening properties as a natural generalisation of the falsification by fellow traveller property. Where FFTP gives a simple way to check if a word is a geodesic, each of the loop shortening properties gives a somewhat simple way to check if a word represents the identity in the group.

Definition 5.36. Let G be a group with finite generating set S. (G, S) has the (synchronous) loop shortening property (LSP) if there is a constant k such that for any loop $v_0, v_1, \ldots, v_n$ in $\Gamma(G, S)$ with $n \geq 1$, there is a shorter loop $u_0, u_1, \ldots, u_m$ such that $d(u_j, v_j) < k$ for each $j \leq m$, and $d(u_m, v_j) < k$ for $m \leq j \leq n$. In other words, the paths (synchronously) k-fellow travel.

Definition 5.37. Let G be a group with finite generating set S. (G, S) has the (synchronous) basepoint loop shortening property (BLSP) if there is a constant k such that for any loop $v_0, v_1, \ldots, v_n$ in $\Gamma(G, S)$ with $n \geq 1$, there is a shorter loop $(v_0 = u_0), u_1, \ldots, u_m$ such that $d(u_j, v_j) < k$ for each $j \leq m$, and $d(u_m, v_j) < k$ for $m \leq j \leq n$. In other words, the paths (synchronously) k-fellow travel.

Note that the only difference between these two properties is that for the basepoint loop shortening property, the initial loop is around a basepoint which the shorter loop has to pass through, whereas for the loop shortening property no such restriction is imposed. So it is clear that

$$\text{BLSP} \Rightarrow \text{LSP}.$$

It is not known, however, whether the reverse implication holds. Moreover, since a loop is not a geodesic, the basepoint loop shortening property is certainly a generalisation of the falsification by fellow traveller property, so

$$\text{FFTP} \Rightarrow \text{BLSP}.$$

In [10], Elder proved that the synchronous and asynchronous versions of both of these properties are equivalent, and he also proved that the basepoint loop shortening property implies almost convexity.

Proposition 5.38. *the basepoint loop shortening property (BLSP) is a block property.*

Proof. Let Γ be a graph with only finitely many non-isomorphic blocks. If Γ satisfies BLSP, then every block of Γ certainly satisfies BLSP. Hence, we just need to show that if every block of Γ satisfies BLSP, then Γ satisfies BLSP. Assume that every block of Γ satisfies the asynchronous basepoint loop shortening property with respect to the constant k, and let l be a loop in Γ based around some vertex v_0. Then we just need to show that there exists a shorter loop l' in Γ which k-fellow travels with l.

Let $v_0, v_1, \ldots, v_{n-1}, (v_n = v_0)$ be the vertices in l in order. Let $i, j \in \mathbb{Z}$ with $0 \leq i < j \leq n$ such that $v_i = v_j$ and such that $j - i$ is minimal amongst all pairs with these conditions. Note that $v_0 = v_n$, so i and j certainly exist. Since $j - i$ is minimal, the loop p with vertices $v_i, v_{i+1}, \ldots, v_{j-1}, (v_j = v_i)$ is simple, so p is contained in a single block B in Γ. Since B satisfies BLSP with constant k, there exists some shorter loop p' in B which is based around the vertex v_i such that p and p' asynchronously k-fellow travel. Then we can construct a loop l' which is shorter than l by replacing the subpath p of l with p'. Moreover, l and l' asynchronously k-fellow travel. □

On the other hand, the loop shortening property does not seem to be a block property. To see this, consider two large simple loops l_1 and l_2 which lie in different blocks of a graph Γ which share a single common vertex v_0. Then we can link these loops to create another loop l. Then any shorter loop l' which k-fellow travels with l must pass through v_0, so such a loop exists if and only if one of l_1 or l_2 is "basepoint loop shortenable" using the constant k. In fact, for a Cayley graph Γ which is not two connected, Γ enjoys the loop shortening property if and only if Γ enjoys the basepoint loop shortening property, so in some sense the question of whether LSP is a block property is equivalent to the question of whether LSP is equivalent to BLSP.

Now, clearly a graph Γ satisfies LSP with respect to some constant k if and only if the m-partition Γ^m satisfies LSP with respect to the constant mk, hence LSP is independent of m-partition. It is slightly less obvious that BLSP is independent of m-partition, since the basepoint in Γ^m may not correspond to a vertex in Γ. So we will write that more formally:

Proposition 5.39. *The property BLSP is independent of m-partition.*

Proof. Let Γ be a graph and let $f : \Gamma \to \Gamma^m$ be the natural bijection which multiplies lengths by m. If Γ^m satisfies BLSP with respect to some constant k, then for any loop l around a fixed basepoint v_0 in Γ, there is a loop $f(l')$ is Γ^m which k-fellow travels loop $f(l)$ around the basepoint $f(v_0)$. Hence, the loop l' passes through v_0 and k/m-fellow travels with l. Hence, Γ satisfies BLSP with respect to the constant k/m.

Now suppose that Γ satisfies the asynchronous basepoint loop shortening property with respect to some constant k, and let l be a loop in Γ^m around some basepoint v_0. We will prove that there is some shorter loop l' with the same basepoint v_0, such that l' and l $(3km + 2m)$-fellow travel. Let the vertices in l be $v_0, v_1, \ldots, v_{n-1}, (v_n = v_0)$. First note that if there is any j such that $v_j = v_{j+2}$, then we can obtain a shorter loop l' by removing v_{j+1} and the edges that it connects to from l, so we may assume that this does not happen. Let v_j be the the vertex such that $f^{-1}(v_j)$ is a vertex in Γ and j is minimal, so $j < m$. Let p be the loop with the same vertices and edges as $f^{-1}(l)$, but such that p is based around $f^{-1}(v_j)$, so $f(p)$ is a loop which synchronously j-fellow travels with l. Then there is a loop p' in Γ with basepoint $f^{-1}(v_j)$ such that p' is shorter than $f^{-1}(l)$ and p' and $f^{-1}(l)$ k-fellow travel. Repeating this, there are loop p'' and p''' in Γ which are based around $f^{-1}(v_j)$ such that p'' k-fellow travels with p' and p''' k-fellow travels with p'', and $|p'| > |p''| > |p'''|$. Then the loop p''' must $3k$-fellow travel with p, so $f(p''')$ is a loop which $(3km + j)$-fellow travels with l. Let l' be the loop in Γ^m based at v_0 consisting of a geodesic from v_0 to v_j, followed by the loop $f(p''')$ followed by a geodesic for v_j to v_0. Then l' asynchronously j-fellow travels with $f(p''')$. Hence l' asynchronously $(3km + 2j)$-fellow travels with l. Moreover,

$$|l'| = |f(p''')| + 2j = m|p'''| + 2j \leq m|p'| - 2m + 2j < m|p'| < m|p| = |f(p)| = |l|,$$

so l' is a loop which $(3km+2m)$-fellow travels with l and which satisfies $|l'| < |l|$. □

5.6 Convex cycle size and k-decomposability

Definition 5.40. Let Γ be a graph. We say that a loop l in Γ is k-decomposable there is some (c, k)-decomposition of l, where $c < |l|$. We say that Γ has the k-decomposability property if there is some constant c_0 such that every loop l with $|l| > c$ is k-decomposable.

Similarly to the Dehn function, for any fixed k, the property k-DP this is clearly independent of m-partition.

Definition 5.41. We will say that Γ satisfies the bounded decomposability property (BDP) if there is some constant $k \in \mathbb{Z}_{\geq 2}$ such that Γ satisfies the k-decomposability property.

Proposition 5.42. *If Γ is a graph and l is a loop in Γ which is not simple, then l is 2-decomposable.*

Proof. Since l is not simple, l is made up of two loops l_1 and l_2 which are based around the same vertex v. Let P be the planar graph made up of two simple loops

c_1 and c_2 which intersect at a single vertex p, and such that neither loop contains the other, where $|c_1| = |l_1|$ and $|c_2| = |l_2|$. Then there is a graph homomorphism $f : P \to \Gamma$ which sends c_1 to l_1 and c_2 to l_2, so it sends the outer face to l. Therefore, l is 2-decomposable. □

Proposition 5.43. *If $k \geq 2$, then k-decomposability is a block property.*

Proof. First, let Γ be a graph which satisfies the k-decomposability property, using constant c_0. We will show that every block B in Γ also satisfies the property. Let B be a block in Γ and let l be any loop in B with $|l| > c_0$. Then l is k-decomposable, so by proposition 5.12, l is k-decomposable in B. Hence, B satisfies k-DP.

Now assume that Γ is a graph with only finitely many non-isomorphic blocks such that each block B in Γ satisfies the k-decomposability property. Let c_0 be the largest constant associated with the k-decomposability property in any of the blocks B in Γ and let l be any loop in Γ such that $|l| > c_0$. Then we just need to show that l is k-decomposable.

If l is not contained in a single block, then by proposition 5.42, the loop l is 2-decomposable. If l is contained in a single block B, then l is k-decomposable because B has the k-decomposability property. So in both cases, l is k-decomposable. □

We call a subgraph U of a connected graph Γ weakly convex if for any points $x, y \in U$, there is a geodesic between x and y in Γ which is contained entirely inside U. We will be particularly interested in the property 2-decomposability, since as we will see, this is equivalent to having a bounded k on the size of any weakly convex simple cycle.

Proposition 5.44. *a graph Γ has the 2-decomposability property, using the constant c if and only if every weakly convex, simple loop l in Γ satisfies $|l| \leq c$.*

Proof. First assume that Γ satisfies 2-DP, using constant c. We will show that if l is a simple, weakly convex loop then $|l| \leq c$.

If $|l| > c$, then there is some planar graph P with two inner faces, each with degree less than $|l|$ and a graph homomorphism $f : P \to \Gamma$ which sends the outer loop of P onto the loop l. Since l is simple, the outer loop of P is simple, so P is given by two vertices u and v and three paths between them p_1, p_2 and p_3, where p_2 is the middle path. Since the inner faces have degree less than $|l|$, we have

$$|p_1| + |p_2| < |l|, \qquad |p_2| + |p_3| < |l| \;\text{ and }\; |p_1| + |p_3| = |l|.$$

Hence, we have $|p_2| < \min\{|p_1|, |p_3|\}$ so $d_\Gamma(f(u), f(v)) \leq d_P(u, v) = |p_2|$, and this is shorter than either of the paths between $f(u)$ and $f(v)$ in l. Therefore, l is not weakly convex.

Now we will show the other direction. Assume that every weakly convex simple loop l in Γ satisfies $|l| \leq c$, we will show that Γ satisfies 2-DP using the constant c. Then we just need to show that if l is a loop with $|l| > c$ then l is 2-decomposable. Since $|l| > c$, it is not weakly convex, there are vertices $u, v \in l$ which split l into two paths l_1 and l_3 between u and v such that neither of these paths is a geodesic. Let l_2 be a geodesic between u and v. Then we can construct a planar graph P

with two vertices u' and v' joined by three paths p_1, p_2 and p_3 where p_2 is the inner path and $|pi| = |l_i|$ for each i. Then along with a map $f : P \to \Gamma$ which sends u' to u, v' to v and each path p_i to l_i, the graph P forms a 2-decomposition of l. □

Proposition 5.45. *If* Γ *is an almost-convex graph, then* Γ *satisfies the bounded decomposability property (BDP)*

Proof. Let r_0 and c be the constants associated which almost convexity in Γ. We will show that Γ satisfies $(c+1)$-DP. Let l be a loop in Γ which satisfies $|l| > \max\{2r_0+2, c+2\}$. Then we just need to show that l is $(c+1)$-decomposable. First note that if l is not (weakly) convex, then it is 2-decomposable, so it is $(c+1)$-decomposable, so we will assume that l is weakly convex.

Let $r = \left\lfloor \frac{|l|-1}{2} \right\rfloor$, so $r > r_0$.

Let the vertices in l, in order, be $v_0, v_1, \ldots, v_{|l|-1}, (v_{|l|} = v_0)$. Then the vertices $v_r, v_{|l|-r} \in Sph(v_0, r)$, and $d(v_r, v_{|l|-r}) \leq 2$. Hence, there is a path p in the ball $\overline{B(v_0, r)}$ with vertices $(v_r = u_0), u_1, u_2, \ldots, u_{m-1}, (u_m = v_{r+2})$ such that $m \leq c$.

Let P be a planar graph consisting of:

- an outer loop l' containing the vertices $v'_0, v'_1, \ldots, v'_{|l|-1}, (v'_{|l|} = v'_0)$,
- a path with vertices $(v'_{r-1} = u'_0), u'_1, \ldots, u'_{m-1}, (u'_m = v'_{|l|-r})$ and
- for $j = 1, 2, \ldots, m-1$, a path p_j of length $d(v_0, u_j)$ between v'_0 and u'_i.

Let $f : P \to \Gamma$ be a function defined by $f(v'_i) = v_i$, $f(u'_j) = u_j$ and such that f sends each path p_j to a geodesic between v_0 and u_j. Then f is a graph homomorphism, so (f, P) is a decomposition of l. Now, P has $m+1$ inner faces. For each $j = 0, 1, \ldots, m-1$, let F_j be the face which is bordered by the paths p_j, p_{j+1} and the edge between u'_j and u'_{j+1}. Then F_j has degree

$$|F_j| = |p_j| + |p_{j+1}| + 1 = d(v_0, u_j) + d(v_0, u_{j+1}) + 1 \leq 2r < |l|.$$

The other face F, whose border contains the path $u'_0, u'_1, \ldots, u'_m$, has degree at most $m + 2 \leq c + 2 < |l|$. Hence every inner face of P has degree less than $|l|$, so (f, P) is an $(m+1)$-decomposition of l. Note that $m + 1 \leq c + 1$, so l is $(c+1)$-decomposable. □

Lemma 5.46. *the property 2-DP implies MAC+.*

Proof. Let Γ be a graph which satisfies 2-DP using constant c. Then we will show that Γ satisfies MAC+ with the same constant $r_0 = c$. Let $r \in \mathbb{Z}$ satisfy $r > c$ and let $a, b \in Sph(v, r)$. Then we just need to prove the following two statements:

- If $d(a, b) = 1$ then there is a path in the closed ball $\overline{B(v, r)}$ between a and b which has length at most $2r - 1$
- If a and b are joined by a path of length 2 outside $\overline{B(v, r)}$, then there is a path p in Γ between a and b which has length at most $2r-1$, such that every vertex in p is in the closed ball $\overline{B(v, r)}$.

First, assume that $d(a,b) = 1$. Let e be the edge between a and b, let p_1 and p_2 be geodesics with p_1 between v and a and p_2 between v and b. Then, together p_1, p_2 and e form a loop l in Γ of length $2r+1$. Since $2r+1 > c$, the loop l is not weakly convex, so there exists some pair of vertices v_1, v_2 in this l such that neither path between v_1 and v_2 in l is a geodesic. Hence one of these two vertices must be on p_1 and the other must be on p_2, so without loss of generality v_i is on p_i for $i = 1, 2$. Now, since no geodesic between v_1 and v_2 goes along l, we have

$$d(v_1, v_2) \leq d(v, v_1) + d(v, v_2) - 1$$

and

$$d(v_1, v_2) \leq d(a, v_1) + d(b, v_2).$$

Now let p' be a geodesic between v_1 and v_2. Then for any point x on p' we have

$$\begin{aligned} d(v,x) \leq& \frac{1}{2}(d(v,v_1) + d(v_1,x) + d(v,v_2) + d(v_2,x)) \\ =& \frac{1}{2}(d(v,v_1) + d(v_1,v_2) + d(v,v_2)) \\ \leq& \frac{1}{2}(d(v,v_1) + d(a,v_1) + d(b,v_2) + d(v,v_2)) \\ =& \frac{1}{2}(d(v,a) + d(v,b)) = r. \end{aligned}$$

So $x \in \overline{B(v,r)}$.

Now let p be the path constructed from the part of p_1 between a and v_1, the part of p_2 between b and v_2 and p'. Then p lies entirely inside $\overline{B(v,r)}$ and

$$|p| = d(a,v_1) + d(b,v_2) + d(v_1,v_2) \leq d(a,v_1) + d(b,v_2) + d(v,v_1) + d(v,v_2) - 1 = 2r-1.$$

Now for the second case, assume that a and b are joined by a path of length 2 outside $\overline{B(v,r)}$. Let u be a point outside $\overline{B(v,r)}$ which is attached to both a and b by edges. Also, let p_1 be a geodesic between u and v which contains a and let p_2 be a geodesic between u and v which contains b. Then since Γ satisfies 2-DP with constant $c < 2r+2$, there is some v_1 on p_1 and some v_2 on p_2 such that

$$d(v_1,v_2) + 1 \leq \min\{d(v,v_1) + d(v,v_2), d(u,v_1) + d(u,v_2)\}.$$

Now let p' be a geodesic between v_1 and v_2 and let p be the a path between a and b containing the section of p_1 which lies between v_1 and a and the section of p_2 which lies between v_2 and b, joined by p'. Then

$$|p| = d(a,v_1) + d(b,v_2) + d(v_1,v_2) < d(a,v) + d(b,v) = 2r.$$

Also, if x is a vertex in p' then

$$\begin{aligned} d(v,x) \leq& \frac{1}{2}(d(v,v_1) + d(v_1,x) + d(v,v_2) + d(v_2,x)) \\ =& \frac{1}{2}(d(v,v_1) + d(v_1,v_2) + d(v,v_2)) \\ <& \frac{1}{2}(d(v,v_1) + d(u,v_1) + d(u,v_2) + d(v,v_2)) \\ =& r+1. \end{aligned}$$

So $\underline{d(v,x)} \leq r$. Hence, every vertex in p' and hence every vertex in p is contained in $\overline{B(v,r)}$. □

Proposition 5.47. *the property 2-DP implies MAC∼.*

Proof. Let Γ be a graph which satisfies 2-DP. Then, since 2-DP is independent of m-partition, Γ^2 satisfies 2-DP. Therefore Γ^2 satisfies MAC+, so Γ satisfies MAC∼. □

Proposition 5.48. *Let G be a group with finite generating set S. If the language of geodesics in $\Gamma(G,S)$ is regular, then $\Gamma(G,S)$ satisfies 2-DP.*

Proof. Suppose the contrary, that there is some group G with a finite generating set S such that the language of geodesics in $\Gamma(G,S)$ is regular and $\Gamma(G,S)$ does not satisfy 2-DP. Let k be the number of states in the deterministic finite state automaton M for the language of geodesics of $\Gamma(G,S)$. Then, since $\Gamma(G,S)$ does not satisfy 2-DP, there is some weakly convex cycle l in $\Gamma(G,S)$ with $|l| > 2k+2$. Let $w = s_1 s_2 \ldots s_n$ be a word corresponding to l, so $\overline{w} = 1$, and let $m = \lfloor n/2 \rfloor$, so $m > k$. Then since l is weakly convex, every subword of w of length at most m represents a geodesic, but every subword of length greater than m does not represent a geodesic. Therefore, the words

$$s_1 s_2 \ldots s_m, s_2 \ldots s_m, \ldots, s_m$$

are all geodesics. Moreover, since there are more than k geodesics listed, at least two geodesics $s_i \ldots s_m$ and $s_j \ldots s_m$ must correspond to the same state of the finite state automaton M. Hence, for any word w over the alphabet S, the word $s_i \ldots s_m w$ is a geodesic if and only if $s_j \ldots s_m w$ is a geodesic. But $s_i \ldots s_{m+i-1}$ is a geodesic and $s_i \ldots s_{m+i}$ is not a geodesic, which implies that $s_j \ldots s_{m+i-1}$ is a geodesic and $s_j \ldots s_{m+i}$ is not a geodesic. But this is only possible if $j = i$, a contradiction. □

Proposition 5.49. *Let Γ be a graph which satisfies k-DP using the constant c_0. Then the Dehn function function of Γ with respect to c_0 satisfies*

$$Dehn(x) \leq k^x \quad \forall x \in \mathbb{Z}_{>0}$$

Proof. We will prove this by induction on x. If $x \leq c_0$ then any loop of length at most x has area at most 1, so

$$\text{Dehn}(x) \leq 1 \leq k^x.$$

If $x > c_0$, assume that the result is true for $x' = x - 1$.

If l is a loop with $|l| = x$, then since $|l| > c_0$, the loop l is $(|l|-1, k)$-decomposable. Therefore, by proposition 5.10, we have

$$\text{Area}(l) \leq k\text{Dehn}(|l|-1).$$

Moreover, by the inductive hypothesis,

$$\text{Dehn}(|l|-1) = \text{Dehn}(x-1) \leq k^{x-1}.$$

Therefore, Area$(l) \leq k^x$. Since this is true for any loop of length x, and by the inductive hypothesis is also true for any shorter loop, the Dehn function satisfies

$$\text{Dehn}(|l|) \leq k^x.$$

This completes the induction. □

The following proposition shows that graphs which satisfy MAC+ are finitely presentable and have solvable word problems. Note that the functions $n!$ and $n^n = e^{n \log(n)}$ are equivalent, so a Dehn function which is equivalent to $n!$ is really only a small step above exponential.

Proposition 5.50. *Let Γ be a graph which satisfies MAC+. Then the Dehn function of Γ is at most equivalent to $n!$.*

Proof. Let r_0 be the constant associated with MAC+ in Γ. During this proof we will use the constant $2r_0 + 2$ to define the Dehn function, so any loop of this length or less has area 1. First we will show that if l is a loop in Γ which satisfies $|l| > 2r_0 + 2$, then l is $(|l| - 2)$-decomposable. If l is not convex, then it is 2-decomposable. Otherwise, we will consider two cases.

First we consider the case where $|l|$ is even. Let $|l| = 2r + 2$, so $r > r_0$. Let the vertices in l, in order, be $v_0, v_1, \ldots, v_{2r+1}, (v_{2r+2} = v_0)$. Then $v_r, v_{r+2} \in Sph(v_0, r)$, and they are joined by a path of length 2 outside the ball $\overline{B(v_0, r)}$. Hence, since Γ satisfies MAC+, there is a path p with vertices $(v_r = u_0), u_1, \ldots, u_{m-1}, (u_m = v_{r+2})$ such that $m \leq 2r - 1$ and for each i, the vertex $u_i \in \overline{B(v_0, r)}$.

Let P be a planar graph consisting of an outer loop l' which contains the vertices

$$v'_0, v'_1, \ldots, v'_{2r+1}, (v'_{2r+2} = v'_0),$$

an inner path consisting of the vertices

$$(v'_{r-1} = u'_0), u'_1, \ldots, u'_{m-1}, (u'_m = v'_{r+1})$$

and for $j = 1, 2, \ldots, m-1$, a path p_j of length $d(v_0, u_j)$ between v'_0 and u'_i. Let $f : P \to \Gamma$ be a function defined by $f(v'_i) = v_i$, $f(u'_j) = u_j$ and f sends each path p_j bijectively to a geodesic between v_0 and u_j. Then f is a graph homomorphism, so (f, P) is a decomposition of l. Now, P has $m + 1$ inner faces:

- For each $j = 0, 1, \ldots, m-1$, there is a face F_j which is bordered by the paths p_j, p_{j+1} and the edge between u'_j and u'_{j+1}.
- There is also the face F, which is bordered by the loop $v'_{r+1}, v'_r, (v'_{r-1} = u'_0), u'_1, \ldots, u_{m-1}, (u'_m = v'_{r+1})$.

Each F_i has degree $|p_i| + |p_{i+1}| + 1 = d(v_0, u_i) + d(v_0, u_{i+1}) + 1 \leq 2r + 1$. Moreover, the face F has degree $m + 2 \leq 2r + 1$. Hence every inner face of P has degree less than $2r + 2 = |l|$, so (f, P) is an $|l|$-decomposition of l.

Now we consider the case where $|l|$ is odd. Let $|l| = 2r + 1$, so $r > r_0$. Let the vertices in l, in order, be $v_0, v_1, \ldots, v_{2r}, (v_{2r+1} = v_0)$. Then $v_r, v_{r+1} \in Sph(v_0, r)$. Hence, there is a path p in $\overline{B(v_0, r)}$ with vertices $(v_r = u_0), u_1, \ldots, u_{m-1}, (u_m =$

v_{r+1}) such that $m \leq 2r - 1$. Since the entire path is in $\overline{B(v_0, r)}$, for each $j = 0, 1, \ldots, m-1$ we must have $d(v_0, u_j) + d(v_0, u_{j+1}) \leq 2r - 1$, since if $d(v_0, u_j) + d(v_0, u_{j+1}) = 2r$, then the edge (u_j, u_{j+1}) would lie outside the ball $\overline{B(v_0, r)}$.

As in the previous case, we construct a decomposition (f, P) of l, and we define the faces F_j and F in the same way. So the degree of any face F_j is $|p_j| + |p_{j+1}| + 1 = d(v_0, u_j) + d(v_0, u_{j+1}) + 1 \leq 2r$. In this case F only contains one edge (v'_r, v'_{r+1}) in the outer loop of P, so the degree of F is $m + 1 \leq 2r$. Therefore, every inner face of P has degree less than $2r + 1 = |l|$, so (f, P) is an $(m+1)$-decomposition of l. Note that $m + 1 \leq 2r < |l|$, so (f, P) is an $|l|$-decomposition of l.

In both cases l is $|l|$-decomposable. Therefore, for any loop l with $|l| > 2r_0 + 2$, we have

$$\text{Area}(l) \leq |l| \text{Dehn}(|l| - 1).$$

Hence, if $m > 2r_0 + 2$, we have

$$\text{Dehn}(m) \leq m \text{Dehn}(m - 1).$$

Moreover, $\text{Dehn}(2r_0 + 2) = 1 < (2r_0 + 2)!$, so we can prove inductively that for $m \geq 2r_0 + 2$ we have $\text{Dehn}(m) \leq m!$. □

5.7 Examples

5.7.1 Solvable Baumslag-Solitar groups

A Baumslag-Solitar group $BS(m, n)$ is the HNN extension of $\mathbb{Z}$ using the subgroups $m\mathbb{Z}$ and $n\mathbb{Z}$ with the isomorphism $\phi : m\mathbb{Z} \to n\mathbb{Z}$ given by $\phi(mx) = nx$. They have presentation $\langle a, b \mid ab^m a^{-1} = b^n \rangle$. A solvable Baumslag-Solitar group is one where either $m = 1$ or $n = 1$. Note that $BS(m, n)$ is isomorphic to $BS(n, m)$, so we will consider groups of the form $BS(1, n)$.

In [19], Miller and Shapiro showed that no solvable Baumslag Solitar group is almost convex with respect to any generating set. Baumslag-Solitar groups also have exponential Dehn functions, so they do not have any of the loop shortening properties with respect to any generating set. In fact, by increasing n, we can make the base of the exponential in the Dehn function for $BS(1, n)$ as large as we want. Hence, for any k, there is some positive integer K such that $BS(1, n)$ does not satisfy the k-decomposability property for $n \geq K$.

We will show, however, that with the standard generating set, the Baumslag-Solitar group $BS(1, 2) = \langle a, b \mid aba^{-1} = b^2 \rangle$ satisfies the 2-decomposability property (2-DP). In particular, the Cayley graph of $BS(1, 2)$ contains no convex cycles of size greater than 5. In [12], Elder and Hermiller showed that $BS(1, 2)$ satisfies M'AC. In the same paper, they also show that $BS(1, 8)$ is not minimally almost convex (MAC), and we will use the same argument to show that it is not MAC+. Hence, using our characterisation of blocks in Cayley graphs, we will show that the four properties M'AC, MAC, MAC+ and MAC∼ are dependent on the generating set. We will also show, in essentially the same way, that for any given k the k-decomposability property depends on the generating set. We do not know, however, whether the bounded decomposability property (BDP) depends on the generating set.

For the next lemma and proposition, we will think of the Cayley graph Γ of $BS(1,2)$ as being embedded in the plane, where right multiplying by a corresponds to moving up 1 and right multiplying by b corresponds to moving right from the point (x,y) to $(x+2^y, y)$

Lemma 5.51. *If p is a geodesic in $BS(1,2)$, then p does not contain two horizontal edges at the same height unless they are at the maximum height in p.*

Proof. Suppose for the sake of contradiction that p does contain two horizontal edges e_1 and e_2 which are at the same height, and this is not the maximum height in p. Let m be the height of these two edges. If e_1 and e_2 are in opposite directions, then we can remove both of these edges to obtain a shorter path between the same two vertices as p. If e_1 and e_2 are in the same direction, then we can remove both of these edges and add an edge e at height $m+1$ in the same direction as e_1 and e_2 to obtain a shorter path between the same two vertices as p. Either way p is not a geodesic, which is a contradiction. □

Proposition 5.52. *Let w be a word in the Cayley graph Γ of $BS(1,2)$, using the presentation $\langle a, b \mid aba^{-1} = b^2 \rangle$ with $\overline{w} = 1$. If $|w| > 5$ then the cycle formed by w is not convex.*

Proof. Let c be the cycle in the Cayley graph Γ represented by w. Suppose for the sake of contradiction that $|w| > 5$ and c is weakly convex.

Let m be the lowest y value of any point in c and let M be the highest y value. At the lowest point, there must be a path in c given by a word of the form $a^{-1}b^k a$ where k is an integer. Since $a^{-1}b^2 = ba^{-1}$ and $a^{-1}b^{-2} = b^{-1}a^{-1}$, we must have $k = 1$, otherwise c would not be convex. Therefore, each time the c comes down to the lowest point, it goes right or left by one edge labelled b, so it goes right or left by a distance of 2^m. Now, every edge in c which lies above this lowest point will move a multiple of 2^{m+1} in the x direction, therefore there must be an even number of horizontal edges in c at height m. In particular, this means that c goes down to this height at least twice. Let v_1, v_2, v_3, v_4 be vertices at this height, such that (v_1, v_2) and (v_3, v_4) are edges in c with (v_1, v_2) on the left of (v_3, v_4). Without loss of generality we may assume that a geodesic p_1 between v_1 and v_3 in c does not intersect either of these two edges. The path between v_2 and v_4 in c which contains p_1 contains both the edges (v_1, v_2) and (v_3, v_4) so by lemma 5.51, it cannot be geodesic. Hence the other path p_2 between v_2 and v_4 in c is geodesic. So c is made up of the two edges (v_1, v_2) and (v_3, v_4) along with the geodesics p_1 and p_2.

Without loss of generality, a highest point in c is in p_1. Around this point the path must take the form $ab^k a^{-1}$, where k is an integer. Then, since aba^{-1} is not a geodesic, we must have $|k| > 1$. Let $u_1, u_2, \ldots, u_n$ be the vertices in p_1 at height M, in order from left to right. Then $n \geq 3$ and the two x values for each adjacent pair differ by 2^M. Let p' be the directed path in c from u_1 to u_n which does not contain $u_2, \ldots, u_{n-1}$.

Now, at each height below M, there is at most one horizontal step in p_1. Therefore, the horizontal distance travelled by these steps is at most

$$2^{M-1} + 2^{M-2} + \ldots + 2^{m+1} = 2^M - 2^{m+1}.$$

Since the horizontal distance travelled by p_1 at height M is precisely $2^M(n-1)$ from left to right, the total horizontal distance travelled by p_1 is at least

$$2^M(n-1) - (2^M - 2^{m+1}) \geq 2^M + 2^{m+1}.$$

Hence, the x value of $v_3 - v_1$ is at least $2^M + 2^{m+1}$. Therefore, the x value of $v_4 - v_2$ is at least 2^M.

Now let M' be the maximum height of any point in p_2. Let $w_1, w_2, \ldots, w_{n'}$ be the vertices in p_2 from left to right at this height. Since at each height below M', there is at most one horizontal step in p_1, the horizontal distance travelled by these steps is at most

$$2^{M'-1} + 2^{M'-2} + \ldots + 2^{m+1} = 2^{M'} - 2^{m+1}.$$

Therefore, the x value of $v_4 - v_2$ is at most $2^{M'}(n'-1) + (2^{M'} - 2^{m+1}) = 2^{M'}n' - 2^{m+1}$. Therefore,

$$2^{M'}n' > 2^{M'}n' - 2^{m+1} \geq 2^M$$

So $n' > 2^{M-M'}$, which means that $n' - 1 \geq 2^{M-M'}$. Now let $r = 2^{M-M'} + 1$. Then $r \leq n'$, so the vertices w_1 and w_r are in p_2 and satisfy

$$w_r - w_1 = (2^M, 0) = u_2 - w_1 = u_3 - u_2.$$

Hence

$$w_1^{-1}u_1 = w_r^{-1}u_2 \text{ and } w_1^{-1}u_2 = w_r^{-1}u_3.$$

Now let d denote the path metric in Γ. Then

$$d(w_1, u_1) = d(w_r, u_2) \text{ and } d(w_1, u_2) = d(w_r, u_3).$$

Since $d(w_1, u_1) = d(w_r, u_2)$, the path in c between u_2 and w_r which contains u_1 and w_1 is not a geodesic. Hence the other path, which contains u_3 is a geodesic. So

$$d(u_3, w_r) = d(u_2, w_r) - 1.$$

But similarly

$$d(u_1, w_1) = d(u_2, w_1) - 1.$$

Which is a contradiction, since $d(u_1, w_1) = d(u_2, w_r)$ and $d(u_3, w_r) = d(u_2, w_1)$. □

We will now go about proving that the standard Cayley graph for the Baumslag Solitar group $BS(1, 8)$ is not MAC+. This proof works in exactly the same way as the proof in [12] that $BS(1, q)$ is not MAC for $q \geq 7$. The following is an equivalent statement of lemma 4.4 from [12].

Lemma 5.53. *Let $q \in \mathbb{Z}_{\geq 7}$, let $BS(1, q)$ be given by the presentation $\langle a, b \mid aba^{-1} = b^q \rangle$ and let d be the path metric in the corresponding Cayley graph. If $m, n \in \mathbb{Z}$ satisfy $d(1, a^m) \leq 2n + 1$, then $m = q^n$ or $m \leq 3q^{n-1}$.*

Using this lemma, the result is straight forward.

Proposition 5.54. *Let $q \in \mathbb{Z}_{\geq 7}$. The Baumslag Solitar group $BS(1, q)$ is not MAC+.*

Proof. Suppose for the sake of contradiction that Γ is MAC+, using constant r_0.

For any $r \in \mathbb{Z}_{>r_0}$, consider the points in the Cayley graph Γ given by $v_1 = a^r b a^{-r} = a^{q^r}$ and $v_2 = b a^r b a^{-r+1} = a^{q^r+1} b$. Then v_1 and v_2 are in the sphere of radius $2r+1$ around the identity and $d(v_1, v_2) = 2$. Since $2r+1 > r_0$, there is some path p in Γ between v_1 and v_2 of length at most $4r+1$, such that every vertex in p is within a distance of $2r+1$ of the identity. If we think of p as going from v_1 to v_2, then consider the last vertex $v = a^t$ in the sheet $\langle b \rangle$. Then, by considering the Bass-Serre tree of Γ, we see that the edge after v must connect to a vertex in the same sheet as v_2, hence $v \neq v_1$ and v is not the identity in Γ.

Therefore, by the lemma 5.53, we have $t \leq 3q^{r-1}$. Hence $q^r - t \geq (q-3)q^{r-1} > 3q^{r-1}$ and $q^r - t \neq q^r$. Therefore, again by lemma 5.53, the distance $d(1, a^{q^r - t}) > 2r+1$. Therefore,

$$d(v_1, v) = d(1, v^{-1} v_1) = d(1, a^{q^r - t}) \geq 2r + 2.$$

Therefore, the length of p is at least

$$d(v_1, v) + d(v, v_2) \geq d(v_1, v) + d(v, v_1) - d(v_1, v_2) \geq 2(2r+2) - 2 = 4r + 2,$$

which is a contradiction, since we defined p to be a path of length at most $4r+1$. □

Proposition 5.55. *Let t be a positive integer. There is a generating set S' for $G = BS(1, 2^t) * F_{t-1}$ such that every block in $\Gamma(G, S')$ is isomorphic to the standard Cayley graph for $BS(1, 2)$.*

Proof. Given the presentation $\langle a, b \mid aba^{-1} = b^2 \rangle$ for $BS(1,2)$, the subgroup $H = \langle a^t, b \rangle$ is isomorphic to $BS(1, 2^t)$ and has index t in $BS(1,2)$. Therefore, we can use example 4.9 to get the desired result. □

We can check that a corresponding presentation for G is

$$\langle a_1, b_1, a_2, b_2, \ldots, a_t, b_t, b_{t+1} \mid a_i b_{i+1} a_i^{-1} = b_i^2 \ \forall i \in \{1, \ldots, t\}, b_1 = b_{t+1} \rangle$$

Proposition 5.56. *All of the Cayley graph properties MAC, M'AC, MAC+, MAC∼ and 2-DP depend on the generating set.*

Proof. Let S be the standard generating set for $G = BS(1,8) * F_2$, coming from the presentation

$$\langle a, b, s, t \mid aba^{-1} = b^8 \rangle$$

and let S' be the generating set from the last lemma, so that every block in $\Gamma(G, S')$ is isomorphic to the standard Cayley graph for $BS(1,2)$. From [12] we know that $BS(1,2)$ satisfies the property M'AC, and from proposition 5.52, we know that $BS(1,2)$ satisfies the property 2-DP. Therefore, $BS(1,2)$ satisfies all of the properties MAC, M'AC, MAC+, MAC∼ and 2-DP, so the Cayley graph $\Gamma(G, S')$ also satisfies all of these properties.

On the other hand, from proposition 5.54, we know that $BS(1,8)$ does not satisfy MAC+. Therefore, since $\Gamma(G, S)$ contains a block which is isomorphic to the standard Cayley graph of $BS(1,8)$, the graph $\Gamma(G, S)$ does not satisfy MAC+. Hence, $\Gamma(G, S)$ satisfies none of the properties MAC, M'AC, MAC+, MAC∼ or 2-DP, so these properties all depend on the generating set. □

Proposition 5.57. *for a fixed k, the property k-DP depends on the generating set.*

Proof. By proposition 5.49, if a graph G satisfies k-DP, then the Dehn function for that graph, with respect to some constant c, is bounded above by an exponential function with base k. Note that while the base of the exponential for the Dehn function may change if we change the generating set for a group, by proposition 5.11 the base will not change if we do not change the graph. Therefore, if a graph Γ satisfies k-DP, then any Dehn function for Γ is at most an exponential function with base k. In particular, this means that if we choose $t \in \mathbb{Z}_{>0}$ which is sufficiently large, then $BS(1, 2^t)$ does not satisfy k-DP. Therefore, using the standard generating set, the group $BS(1, 2^t) * F_{t-1}$ does not satisfy k-DP.

Now let S' be a generating set for $G = BS(1, 2^t) * F_{t-1}$ such that every block in $\Gamma(G, S')$ is isomorphic to the standard Cayley graph for $BS(1, 2)$. Then $\Gamma(G, S')$ satisfies 2-DP, so it satisfies k-DP. Hence the property k-DP depends on the generating set. □

5.7.2 The group $\langle a, b, t \mid ab = ba, a^{10} = tbt^{-1} \rangle$

We consider two generating sets S_1 and S_2 for the group $G = \langle a, b, t \mid ab = ba, t^{-1}at = b^{10} \rangle$ Such that the Cayley graphs $\Gamma(G, S_1)$ and $\Gamma(G, S_2)$ have quite different geometries. In particular, $\Gamma(G, S_1)$ satisfies the basepoint loop shortening property (BSLP). On the other hand, $\Gamma(G, S_2)$ does not have the loop shortening property (LSP), and hence does not satisfy BLSP either. As a result, none of these three properties are independent of generating set. In particular, this answers in the negative the previously open question of whether the loop shortening property is independent of generating set.

The presentations in question are

$$(G, S_1) = \langle a, c, t, d \mid ac = ca, d = a^{10}, d = tct^{-1} \rangle$$

and

$$(G, S_2) = \langle a, b, t \mid ab = ba, a^{10} = tabt^{-1} \rangle.$$

Note that the two groups are the same because we can go between them using basic Tietze transformations as follows:

$$\begin{aligned} &\langle a, c, t, d \mid ac = ca, d = a^{10}, d = tct^{-1} \rangle \\ =&\langle a, c, t \mid ac = ca, a^{10} = tct^{-1} \rangle \\ =&\langle a, b, c, t \mid b = a^{-1}c, ac = ca, a^{10} = tct^{-1} \rangle \\ =&\langle a, b, c, t \mid c = ab, aab = aba, a^{10} = tabt^{-1} \rangle \\ =&\langle a, b, t \mid ab = ba, a^{10} = tabt^{-1} \rangle. \end{aligned}$$

We will first prove that the pair (G, S_1) enjoys the basepoint loop shortening property. In fact, (G, S_1) also satisfies FFTP, though it is not necessary to prove this for our purposes. The fact that (G, S_1) enjoys the basepoint loop shortening property actually follows from a trivial application of a lemma from [10], though before stating the lemma we will need to make some definitions.

Let G' be a group with finite generating set S and relations R. For each $i \in \{1, 2, \ldots, n\}$, let A_i and B_i be subgroups of G' and let $\phi_i : A_i \to B_i$ be an isomorphism. For the next few definitions, consider the multiple HNN extension with the standard presentation:

$$\langle S, s_1, \ldots, s_n \mid R, s_i a_i s_i^{-1} = \phi_i(a_i) \forall a_i \in A_i, \forall i \rangle$$

Definition 5.58. If a group has presentation as given above, we say that the multiple HNN extension is strip equidistant if the following holds:

For each $i \in \{1, \ldots, n\}$ and each $a_i \in A_i$ we have $|\phi_i(a_i)| = |a_i|$ where $|g|$ is the length of g in the group G using the word metric coming from S.

Definition 5.59. Given a graph Γ and a subgraph H of Γ, we say that H is totally geodesic in Γ if every geodesic in Γ between points in H is contained in H. In other words, if $x, y, z \in X$ satisfy $x, z \in Y$ and $d(x, z) = d(x, y) + d(y, z)$, then $y \in Y$.

In [10] Elder shows that if (G', S') is a group and generating set arising from a strip equidistant, multiple HNN extension of a pair (G, S), which satisfies the falsification by fellow traveller property, where the associated subgroups A_i, B_i of G are totally geodesic in $\Gamma(G, S)$, then $\Gamma(G', S')$ satisfies the basepoint loop shortening property.

Now, notice that(G, S_1) is a strip-equidistant HNN extension of $\mathbb{Z}^2$. Moreover, the associated subgroups of this HNN extension are $\langle d \rangle$ and $\langle c \rangle$, which are both totally geodesic in $\langle a, c, d \mid ac = ca, d = a^{10} \rangle$. Finally, since $\mathbb{Z}^2$ satisfies FFTP with respect to any generating set, the pair (G, S_1) satisfies the loop shortening property.

Now we will begin to prove that the pair (G, S_2) does not satisfy the loop shortening property (LSP). Before we get to the result, we need to study the path metric d in the Cayley graph $\Gamma(G, S_2)$. To give an idea as to how the geometry of $\Gamma(G, S_2)$ is quite different to the geometry of $\Gamma(G, S_1)$, consider words for the element a^{1000} in (G, S_2). We can write

$$a^{1000} = ta^{100}b^{100}t^{-1} = t^2a^{10}b^{10}t^{-1}b^{100}t^{-1} = t^3abt^{-1}b^{10}t^{-1}b^{100}t^{-1},$$

and we will eventually prove that the last word in the sequence of equalities is a shortest word for a^{1000}. This is very different from (G, S_1), where the subgroups $\langle a \rangle$ and $\langle a, b \rangle$ both form totally geodesic subspaces of (G, S_1).

We define a function $\phi : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, by

$$\phi(x) = \begin{cases} x, & \text{if } 0 \leq x \leq 1, \\ \dfrac{x+8}{9} + 2\log_{10}(x), & \text{if } x > 1. \end{cases}$$

The point of this is that we will eventually prove that for an integer n, the distance $d(1, a^n) \geq \phi(|n|)$, where d is the path metric in $\Gamma(G, S_2)$, and equality is achieved when n is a power of 10.

Note that since

$$\frac{1+8}{9} + 2\log_{10}(1) = 1,$$

the function ϕ is continuous. Note also that ϕ is strictly increasing on both of the intervals $[0,1)$ and $(1,\infty)$, hence since ϕ is continuous, it is strictly increasing.

Lemma 5.60. *If $x, y \in \mathbb{R}_{\geq 0}$ then $\phi(x+y) \leq \phi(x) + y$*

Proof. For $x > 1$ we have

$$\phi'(x) = \frac{1}{9} + \frac{2}{\log_e(10)x} < \frac{1}{9} + \frac{2}{\log_e(10)} \approx 0.979 < 1.$$

Therefore, if $x \geq 1$ then $\phi(x+y) \leq \phi(x) + y$. If $x < 1$ and $x + y \geq 1$, then

$$\phi(x+y) \leq \phi(1) + x + y - 1 = x + y = \phi(x) + y.$$

Finally, if $x + y < 1$, then $\phi(x+y) = x + y = \phi(x) + y$. □

Lemma 5.61. *If $x, y \in \mathbb{R}_{\geq 0}$ then $\phi(x) + \phi(y) \geq \phi(x+y)$*

Proof. If $y \leq 1$ then

$$\phi(x) + \phi(y) = \phi(x) + y \geq \phi(x+y)$$

and we are done. Similarly the result is true if $x \leq 1$.

If $x, y \geq 1$, then $xy = (x-1)(y-1) + x + y - 1 \geq x + y - 1$. Therefore,

$$\begin{aligned}
\phi(x) + \phi(y) &= \frac{x+8}{9} + 2\log_{10}(x) + \frac{y+8}{9} + 2\log_{10}(y) \\
&= \frac{x+y+16}{9} + 2\log_{10}(xy) \\
&\geq \frac{x+y+16}{9} + 2\log_{10}(x+y-1) \\
&= \frac{x+y+16}{9} + 2\log_{10}(x+y) + 2\log_{10}\left(\frac{x+y-1}{x+y}\right) \\
&= \phi(x+y) + \frac{8}{9} + 2\log_{10}\left(1 - \frac{1}{x+y}\right) \\
&\geq \phi(x+y) + \frac{8}{9} + 2\log_{10}\left(1 - \frac{1}{2}\right) \\
&> \phi(x+y) + 0.287 > \phi(x+y)
\end{aligned}$$

□

Lemma 5.62. *If $x \in \mathbb{R}_{\geq 0}$ then $\phi(10x) \geq \phi(x) + x$, with equality if and only if $x = 0$.*

Proof. If $x > 1$, then we have

$$\phi(10x) = \frac{10x+8}{9} + 2\log_{10}(10x) = x + \frac{x+8}{9} + 2 + \log_{10}(x) = \phi(x) + x + 2.$$

If $1 \geq x > 1/10$ then we have

$$\phi(10x) = \frac{10x+8}{9} + 2\log_{10}(10x) > \frac{10x+8}{9} \geq \frac{10x+8x}{9} = 2x = \phi(x) + x.$$

Finally, if $x \leq 1/10$ then $\phi(10x) = 10x \geq 2x = \phi(x) + x$.

□

Now, recall that d is the path metric in $\Gamma(G, S_2)$.

Lemma 5.63. *If $x, y \in \mathbb{Z}$ then $d(1, a^x b^y) \geq \phi(|x|) + |y|$*

Proof. Suppose for the sake of contradiction that $d(1, a^x b^y) < \phi(|x|) + |y|$ for some $x, y \in \mathbb{Z}$. Let $m, n \in \mathbb{Z}$ satisfy $d(1, a^m b^n) < \phi(|m|) + |n|$ with $d(1, a^m b^n)$ minimal. Let w be a geodesic word over the alphabet $\{a, a^{-1}, b, b^{-1}, t, t^{-1}\}$ with $\overline{w} = a^m b^n$. So

$$|w| = d(1, a^m b^n) < \phi(|m|) + |n|.$$

Since $\overline{w} = a^m b^n \in \langle a, b\rangle$, by proposition 1.8 we have four possible cases:

Case 1: $|w| \leq 1$.

Case 2: The word w decomposes as $w = w_1 w_2$ where w_1 and w_2 are non-empty subwords of w and $\overline{w_1}, \overline{w_2} \in \langle a, b\rangle$.

Case 3: there is some word w_1 with $\overline{w_1} \in \langle a^{10}\rangle$ such that $w = t^{-1} w_1 t$.

Case 4: There is some word w_1 with $\overline{w_1} \in \langle ab\rangle$ such that $w = t w_1 t^{-1}$.

Case 1: Since $|w| \leq 1$, we have $w \in \{1, a, a^{-1}, b, b^{-1}\}$. If $w \in \{a, a^{-1}, b, b^{-1}\}$ then

$$|w| = 1 = |m| + |n| = \phi(|m|) + |n|,$$

which is a contradiction. If $w = 1$, then

$$|w| = 0 = \phi(0) + |0| = \phi(|m|) + |n|,$$

which is also a contradiction.

Case 2: Let $\overline{w_1} = a^{m_1} b^{n_1}$ and let $\overline{w_2} = a^{m_2} b^{n_2}$. Then

$$a^m b^m = \overline{w} = \overline{w_1 w_2} = a^{m_1+m_2} b^{n_1+n_2}.$$

Therefore, $m = m_1 + m_2$ and $n = n_1 + n_2$. By the minimality of $|w|$, we must have

$$|w_1| \geq \phi(|m_1|) + |n_1| \text{ and } |w_2| \geq \phi(|m_2|) + |n_2|.$$

Therefore,

$$\begin{aligned} |w| &= |w_1| + |w_2| \\ &\geq \phi(|m_1|) + |n_1| + \phi(|m_2|) + |n_2| \\ &\geq \phi(|m_1| + |m_2|) + |n_1| + |n_2| \\ &\geq \phi(|m_1 + m_2|) + |n_1 + n_2| \\ &= \phi(|m|) + |n|, \end{aligned}$$

which is a contradiction.

Case 3: Since $\overline{w_1} \in \langle a^{10} \rangle$, we can write $w_1 = a^{10c}$ for some constant c. Then $\overline{w} = \overline{t^{-1}w_1 t} = a^c b^c$. Now, by the minimality of $|w|$, we must have

$$|w_1| \geq \phi(10|c|) = \frac{10|c|+8}{9} + 2\log_{10}(10|c|) = \frac{10|c|+8}{9} + 2\log_{10}(|c|) + 2.$$

Therefore,

$$\begin{aligned}
\frac{10|c|+8}{9} + 2\log_{10}(|c|) + 2 \leq |w_1| & \\
&= |w| - 2 \\
&< \phi(|c|) + |c| - 2 \\
&= \frac{|c|+8}{9} + 2\log_{10}(|c|) + |c| - 2 \\
&= \frac{10|c|+8}{9} + 2\log_{10}(|c|) - 2,
\end{aligned}$$

which is a contradiction.

Case 4: Since $\overline{w_1} \in \langle ab \rangle$, we can write $\overline{w_1} = a^c b^c$ for some constant c. Then $\overline{w} = \overline{t w_1 t^{-1}} = a^{10c}$. Now, by the minimality of $|w|$, we must have

$$|w_1| \geq \phi(|c|) + |c| = \frac{|c|+8}{9} + 2\log_{10}(|c|) + |c| = \frac{10|c|+8}{9} + 2\log_{10}(|c|).$$

Therefore,

$$\begin{aligned}
\frac{10|c|+8}{9} + 2\log_{10}(|c|) \leq |w_1| & \\
&= |w| - 2 \\
&< \phi(10|c|) + |0| - 2 \\
&= \frac{10|c|+8}{9} + 2\log_{10}(10|c|) - 2 \\
&= \frac{10|c|+8}{9} + 2\log_{10}(|c|),
\end{aligned}$$

which is a contradiction. □

Lemma 5.64. *If $x, y, n, p, q \in \mathbb{Z}$ and $n > 0$, then*

$$d(a^x b^y, t^n a^p b^q) \geq |y| + n + \frac{|x|(1 - 10^{-n+1})}{9} + \phi(|p - x \cdot 10^{-n}|) + |q - x \cdot 10^{-n}|$$

Proof. Suppose the contrary. Let $x, y, n, p, q \in \mathbb{Z}$ with $n > 0$ satisfy

$$d(a^x b^y, t^n a^p b^q) < |y| + n + \frac{|x|(1 - 10^{-n+1})}{9} + \phi(|p - x \cdot 10^{-n}|) + |q - x \cdot 10^{-n}|,$$

with $d(a^x b^y, t^n a^p b^q)$ minimal. Let w be a reduced word for the element $a^{-x} b^{-y} t^n a^p b^q$. So

$$|w| < |y| + n + \frac{|x|(1 - 10^{-n+1})}{9} + \phi(|p - x \cdot 10^{-n}|) + |q - x \cdot 10^{-n}|.$$

Then the word $w' = t^{-n}a^xb^yw$ satisfies $\overline{w'} = a^pb^q \in \langle a, b\rangle$, so by proposition 1.8 we have four possible cases:

Case 1: $|w'| \leq 1$.

Case 2: The word w' decomposes as $w' = w_1w_2$ where w_1 and w_2 are non-empty subwords of w' and $\overline{w_1}, \overline{w_2} \in \langle a, b\rangle$.

Case 3: there is some word w_1' with $\overline{w_1'} \in \langle a^{10}\rangle$ such that $w' = t^{-1}w_1t$.

Case 4: There is some word w_1' with $\overline{w_1'} \in \langle ab\rangle$ such that $w' = tw_1t^{-1}$.

Since w' starts with the letter t^{-1}, cases 1 and 4 are clearly impossible, so we are left with cases 2 and 3.

Case 2: First notice that

$$w_1w_2 = w' = t^{-n}a^xb^yw.$$

Since $\overline{w_1} \in \langle a, b\rangle$, the word w_1 must contain just as many occurrences of the letter t as it does of the letter t^{-1}, in particular, since w_1 starts with the letter t^{-1}, it must conain some occurrence of the letter t, so w_1 must contain the entire subword $t^{-n}a^xb^y$. Let $w_1 = t^{-n}a^xb^yw_1'$. Then $w = w_1'w_2$. Recall that w_2 is non-empty, so $|w| > |w_1|$.

Now let $\overline{w_2} = a^{p_1}b^{q_1}$. Then

$$\overline{w_1'} = (\overline{w})(\overline{w_2})^{-1} = a^{-x}b^{-y}t^na^{p-p_1}b^{q-q_1}.$$

Therefore, by the minimality of $|w|$, we must have

$$|w_1'| \geq |y| + n + \frac{|x|(1-10^{-n+1})}{9} + \phi(|p - p_1 - x\cdot 10^{-n}|) + |q - q_1 - x\cdot 10^{-n}|.$$

We also know form the previous lemma that

$$|w_2| \geq \phi(|p_1|) + |q_1|.$$

Finally, recall that

$$|w| < |y| + n + \frac{|x|(1-10^{-n+1})}{9} + \phi(|p - x\cdot 10^{-n}|) + |q - x\cdot 10^{-n}|.$$

Therefore,

$$\begin{aligned} &\phi(|p_1|) + |q_1| \leq |w_2| = |w| - |w_1'| \\ <&\phi(|p - x\cdot 10^{-n}|) + |q - x\cdot 10^{-n}| - \phi(|p - p_1 - x\cdot 10^{-n}|) - |q - q_1 - x\cdot 10^{-n}|. \end{aligned}$$

But then we have

$$\begin{aligned} &\phi(|p - x\cdot 10^{-n}|) + |q - x\cdot 10^{-n}| \\ >&\phi(|p - p_1 - x\cdot 10^{-n}|) + |q - q_1 - x\cdot 10^{-n}| + \phi(|p_1|) + |q_1| \\ \geq&\phi(|p - p_1 - x\cdot 10^{-n}| + |p_1|) + |q - q_1 - x\cdot 10^{-n}| + |q_1| \\ \geq&\phi(|p - x\cdot 10^{-n}|) + |q - x\cdot 10^{-n}|, \end{aligned}$$

which is a contradiction.

Case 3: We know that $\overline{w_1'} \in \langle a^{10} \rangle$, so let $\overline{w_1'} = a^{10c}$. Since $t^{-1}\overline{w_1'}t = \overline{w'} = a^p b^q$, we must have $p = q = c$. Now, we have

$$t^{-n}a^x b^y w = w' = t^{-1}w_1't,$$

so w ends in the letter t. Let $w = w_1 t$. Then $w_1' = t^{-n+1}a^x b^y w_1$. Therefore,

$$\overline{w_1} = a^{-x}b^{-y}t^{n-1}\overline{w_1'} = a^{-x}b^{-y}t^{n-1}a^{10c}.$$

If $n = 1$, then $\overline{w_1} = a^{10c-x}b^{-y}$, so by the previous lemma, we have

$$|w| = |w_1| + 1 \geq \phi(|10c - x|) + |y| + 1.$$

But by our assumption,

$$\begin{aligned}|w| &< |y| + n + \frac{|x|(1 - 10^{-n+1})}{9} + \phi(|p - x \cdot 10^{-n}|) + |q - x \cdot 10^{-n}| \\ &= |y| + 1 + \phi(|c - x \cdot 10^{-1}|) + |c - x \cdot 10^{-1}|.\end{aligned}$$

Putting these together, we get

$$\phi(|c - x \cdot 10^{-1}|) + |c - x \cdot 10^{-1}| > \phi(|10c - x|),$$

which contradicts lemma 5.62. Therefore, $n > 1$.

Since $|w_1| = |w| - 1$, by the minimality of $|w|$, we must have

$$\begin{aligned}|w_1| &\geq |y| + n - 1 + \frac{|x|(1 - 10^{-n+2})}{9} + \phi(|10c - x \cdot 10^{-n+1}|) + |0 - x \cdot 10^{-n+1}| \\ &= |y| + n - 1 + \frac{|x|(1 - 10^{-n+1})}{9} + \phi(|10c - x \cdot 10^{-n+1}|)\end{aligned}$$

Therefore,

$$\begin{aligned}&|y| + n + \frac{|x|(1 - 10^{-n+1})}{9} + \phi(|10c - x \cdot 10^{-n+1}|) \\ <&|y| + n + \frac{|x|(1 - 10^{-n+1})}{9} + \phi(|c - x \cdot 10^{-n}|) + |c - x \cdot 10^{-n}|,\end{aligned}$$

which contradicts lemma 5.62. □

Theorem 5.65. *The pair (G, S_2) does not enjoy the loop shortening property.*

Proof. Let $k \in \mathbb{Z}_{>4}$. We will show that there is a loop l in $\Gamma(G, S_2)$ such that there is no shorter loop l' in $\Gamma(G, S_2)$ which k-fellow travels with l. Hence this will show that $\Gamma(G, S_2)$ does not satisfy the loop shortening property. Let l be the loop given by the word

$$w = t^{2k}abt^{-1}b^{10}t^{-1}b^{100}\dots b^{10^{2k-1}}t^{-1}bt^{2k}a^{-1}b^{-1}t^{-1}b^{-10}t^{-1}b^{-100}\dots b^{-10^{2k-1}}t^{-1}b^{-1}.$$

We can easily check algebraically that $\overline{w} = 1$, so this is indeed a loop. Now let l' be a loop which k-fellow travels with l. Then we just need to show that the length of l' is at least the length of l. First we will prove that there are vertices v_1, v_2, v_3 and v_4 which l' passes through in that order, which take the forms:

$$v_1 = t^k a^x, \quad v_2 = a^{s_2}, \quad v_3 = a^{10^{2k}} b t^k a^y \quad \text{and} \quad v_4 = a^{10^{2k} - s_4},$$

where $x, y, s_2, s_4 \in \mathbb{Z}$ with $|s_2|, |s_4| \leq 10k$.

Let T be the Bass-Serre tree of the HNN-extension and let $f : \Gamma(G, S_2) \to T$ be the associated quotient map. Since l and l' k-fellow travel, there is some vertex v in l' which is within a distance of k of the point a^t. Similarly, there is a vertex u in l' which is within a distance of k of the point $a^{10^{2k}} b t^{2k}$. Let p be a path in Γ made up of a geodesic between a^t and v, a path in one of the directions along l' between v and u and a geodesic between u and $a^{10^{2k}} b t^{2k}$. Now, the vertices in the simple path in T between $f(t^{2k})$ and $f(a^{10^{2k}} b t^{2k})$ are

$$f(t^{2k}), f(t^{2k-1}), \ldots, f(t), (f(1) = f(a^{10^{2k}} b)), f(a^{10^{2k}} bt), \ldots, f(a^{10^{2k}} b t^{2k})$$

Therefore, the image $f(p)$ contains the edge between $f(t^{k+1})$ and $f(t^k)$, so there exists some pair of adjacent vertices v_1, v_1' in p such that $f(v_1) = f(t^k)$ and $f(v_1') = f(t^{k+1})$. Therefore, $v_1 = t^k a^x b^j$ for some i, j and $v_1' = t^{k+1} a^{i'} b^{j'}$. Since these are adjacent, we must have $j = 0$ and $x = 10i' = 10j'$, so $v_1 = t^k a^x$. Finally, since $d(v_1, a^{2k}) \geq k$ and $d(v_1, a^{10^{2k}} b t^{2k}) \geq k$, the vertex v_1 must be on l'. Similarly, there is some vertex $v_3 = a^{10^{2k}} b t^k a^y$ in l'.

Let γ be the path between t^{2k} and $a^{10^{2k}} b t^{2k}$ along l which contains 1. Let γ' be the path along l' between v and u which asynchronously k-fellow travels with γ. Since γ' is a path between u and v, it must contain some pair of adjacent points v_2, v_2' such that v_2 is in the same sheet as 1 and v_2' is in the same sheet as t. Therefore, $v_2 = a^{s_2}$ for some $s \in \mathbb{Z}_{>0}$. Since γ and γ' k-fellow travel, v_2 must be within a distance of k of γ. Therefore, $|s_2| \leq 10k$. Similarly, the path between u and v along the other side of l' passes through a point $v_4 = a^{10^{2k} - s_4}$ where $s_4 \leq 10k$.

Hence l' passes through the vertices

$$v_1 = t^k a^x, v_2 = a^{s_2}, v_3 = a^{10^{2k}} b t^k a^y \text{ and } v_4 = a^{10^{2k} - s_4},$$

in that order, where $x, y, s_2, s_4 \in \mathbb{Z}$ with $|s_2|, |s_4| \leq 10k$.

Therefore, we must have

$$|l'| \geq d(v_1, v_2) + d(v_2, v_3) + d(v_3, v_4) + d(v_4, v_1).$$

now,

$$\begin{aligned} d(v_1, v_2) &= d(a^{s_2}, t^k a^x) \\ &\geq k + \frac{|s_2|(1 - 10^{-k+1})}{9} + \phi(|x - s_2 \cdot 10^{-k}|) + |s_2 \cdot 10^{-k}| \\ &\geq k + \frac{|s_2|(1 - 10^{-k})}{9} + \phi(|x - s_2 \cdot 10^{-k}|), \end{aligned}$$

$$
\begin{aligned}
d(v_2,v_3) &= d(a^{s_2-10^{2k}}b^{-1}, t^k a^y) \\
&\geq |-1| + k + \frac{|(s_2-10^{2k})|(1-10^{-k+1})}{9} \\
&\quad + \phi(|y-(s_2-10^{2k})\cdot 10^{-k}|) + |(s_2-10^{2k})\cdot 10^{-k}| \\
&\geq 1 + k + \frac{|(s_2-10^{2k})|(1-10^{-k})}{9} + \phi(|y-(s_2-10^{2k})\cdot 10^{-k}|),
\end{aligned}
$$

$$
\begin{aligned}
d(v_3,v_4) &= d(a^{-s_4}b^{-1}, t^k a^y) \\
&\geq |-1| + k + \frac{|-s_4|(1-10^{-k+1})}{9} \\
&\quad + \phi(|y+s_4\cdot 10^{-k}|) + |s_4\cdot 10^{-k}| \\
&\geq 1 + k + \frac{|-s_4|(1-10^{-k})}{9} + \phi(|y+s_4\cdot 10^{-k}|), \text{ and}
\end{aligned}
$$

$$
\begin{aligned}
d(v_4,v_1) &= d(a^{10^{2k}-s_4}, t^k a^x) \\
&\geq k + \frac{|10^{2k}-s_4|(1-10^{-k+1})}{9} \\
&\quad + \phi(|x-(10^{2k}-s_4)\cdot 10^{-k}|) + |(10^{2k}-s_4)\cdot 10^{-k}| \\
&\geq k + \frac{|10^{2k}-s_4|(1-10^{-k})}{9} + \phi(|x-(10^{2k}-s_4)\cdot 10^{-k}|)
\end{aligned}
$$

Therefore, we have

$$
\begin{aligned}
|l'| \geq& d(v_1,v_2) + d(v_2,v_3) + d(v_3,v_4) + d(v_4,v_1) \\
\geq& k + \frac{|s_2|(1-10^{-k})}{9} + \phi(|x-s_2\cdot 10^{-k}|) \\
&+ 1 + k + \frac{|(s_2-10^{2k})|(1-10^{-k})}{9} + \phi(|y-(s_2-10^{2k})\cdot 10^{-k}|) \\
&+ 1 + k + \frac{|-s_4|(1-10^{-k})}{9} + \phi(|y+s_4\cdot 10^{-k}|) \\
&+ k + \frac{|10^{2k}-s_4|(1-10^{-k})}{9} + \phi(|x-(10^{2k}-s_4)\cdot 10^{-k}|) \\
=& 2 + 4k + \frac{(|s_2| + |10^{2k}-s_4| + |s_4| + |10^{2k}-s_2|)(1-10^{-k})}{9} \\
&+ \phi(|x-s_2\cdot 10^{-k}|) + \phi(|(10^{2k}-s_4)\cdot 10^{-k}-x|) \\
&+ \phi(|y+s_4\cdot 10^{-k}|) + \phi(|(s_2-10^{2k})\cdot 10^{-k}-y|) \\
\geq& 2 + 4k + \frac{2\cdot 10^{2k}(1-10^{-k})}{9} \\
&+ \phi(|(10^{2k}-s_4-s_2)\cdot 10^{-k}|) + \phi(|(s_2+s_4-10^{2k})\cdot 10^{-k}|) \\
\geq& 2 + 4k + \frac{2(10^{2k}-10^k)}{9} + 2\phi(10^k) - 2\cdot 10^{-k}|s_2+s_4| \\
=& 2 + 4k + \frac{2(10^{2k}-10^k)}{9} + 2\left(\frac{10^k+8}{9} + 2k\right) - 2\cdot 10^{-k}|s_2+s_4|
\end{aligned}
$$

$$
\begin{aligned}
&=2 + 4k + \frac{2(10^{2k} - 10^k)}{9} + 2\left(\frac{10^k + 8}{9} + 2k\right) - 2 \cdot 10^{-k}|s_2 + s_4| \\
&=4 + 8k + \frac{2(10^{2k} - 1)}{9} - 2 \cdot 10^{-k}|s_2 + s_4| \\
&=|l| - 2 \cdot 10^{-k}|s_2 + s_4| \\
&\geq |l| - 20k \cdot 10^{-k} > |l| - 1.
\end{aligned}
$$

Therefore, $|l'| > |l| - 1$. But $|l'|$ and $|l|$ are both integers, so we actually have $|l'| \geq |l|$. □

Hence, we have shown that the pair (G, S_1) enjoys the basepoint loop shortening property (BLSP), and hence also the loop shortening property, whereas (G, S_2) does not satisfy the LSP and hence does not satisfy BLSP either. Therefore, the two properties LSP and BLSP both depend on the generating set.

Chapter 6

Geodetic graphs and groups

6.1 Geodetic graphs

Definition 6.1. A graph Γ is called geodetic if for every pair of vertices v, w in Γ there is a unique geodesic between v and w.

Geodetic graphs were first defined by Ore [22, p. 104]. Since then a large amount of work has been done in an attempt to characterise them. Watkins did this in the case of planar graphs in his doctoral thesis and later he and Stemple simplified the proof [30]. Blokhuis and Brouwer have essentially characterised geodetic graphs of diameter 2 [3].

Theorem 6.2 (Stemple and Watkins [30])**.** *In a planar geodetic graph, every block B can be obtained from either K_2, K_3 or K_4, by replacing each edge with a finite chain of edges. In other words, B is conformal to one of these three graphs.*

Proposition 6.3. *Let Γ be a geodetic graph. If $a_0, a_1, \ldots, a_n$ and $b_0, b_1, \ldots, b_n$ are geodesics in Γ which only intersect at $a_0 = b_0$ and a_n and b_n are joined by an edge, then the cycle c given by vertices $a_0, a_1, \ldots, a_n, b_n, b_{n-1}, \ldots, b_1, (b_0 = a_0)$ is convex.*

Proof. Since the geodesic between a_0 and a_n does not pass through b_1, any path from a_0 to a_n which passes through b_1 has length at least $n+1$. So $d(a_n, b_1) \geq n$. But the path $a_n, b_n, \ldots, b_1$ has length n, so it is a geodesic.

Now we have geodesics between a_n and each of the vertices b_0 and b_1 which go in opposite directions around c. So we can use the same argument as before to show that $b_1, (b_0 = a_0), a_1, \ldots, a_{n-1}$ is a geodesic. Repeating the same argument shows that every path of length n around the outside of c is a geodesic. But any path of length at most n on the cycle is contained in a path of length n, so it is also a geodesic.

Finally, the length of the cycle is $2n+1$, so any two points on c are joined by a path of length at most n on c, so this path is a geodesic. Hence, for any two points on c, the geodesic between them is contained in c, therefore c is convex. $\square$

Definition 6.4. A graph Γ is called l-almost geodetic if for any pair of vertices $v, w \in V(\Gamma)$ with $d(v, w) > 1$, and any pair of geodesics g_1, g_2 between v and w, the geodesics g_1 and g_2 intersect at some point other than v and w.

Proposition 6.5. *A graph* Γ *is geodetic if and only if it is* 0*-almost geodetic.*

Proof. If Γ is geodetic then for any pair of vertices $v, w \in V(\Gamma)$ with $d(v, w) > 0$ (i.e., $v \neq w$) and any pair of geodesics g_1, g_2 between v and w, the geodesics g_1 and g_2 are the same. Therefore, g_1 and g_2 certainly intersect at a point other than v and w.

Now we will assume that Γ is 0-almost geodetic and prove that Γ is geodetic. Suppose for the sake of contradiction that Γ is not geodetic. Let $v, w \in V(\Gamma)$ be distinct vertices such that there are multiple geodesics between v and w and such that $d(v, w)$ is minimal. Let g_1 and g_2 be distinct geodesics between v and w. By the minimality of $d(v, w)$, the geodesics g_1 and g_2 must not intersect at any point other than v and w. But this contradicts the fact that Γ is 0-almost geodetic. □

Proposition 6.6. *A graph* Γ *is* l*-almost geodetic if and only if every block in* Γ *is* l*-almost geodetic.*

Proof. First we will assume that Γ is l-almost geodetic and prove that any block B in Γ is l-almost geodetic. Let $v, w \in V(B)$ with $d(w, v) > l$ and let g_1, g_2 be geodesics in B between v and w. We will show that g_1 and g_2 intersect at some point other than v and w. Since g_1, g_2 are geodesics in B, they are also geodesics in Γ. Therefore, since Γ is l-almost geodetic, g_1 and g_2 intersect at some point other than v and w.

Now we will assume that every block B in Γ is l-almost geodetic and prove that Γ is l-almost geodetic. Let $v, w \in V(\Gamma)$ with $d(w, v) > l$ and let g_1, g_2 be geodesics in Γ between v and w. We will show that g_1 and g_2 intersect at some point other than v and w.

Suppose for the sake of contradiction that g_1 and g_2 only intersect at v and w. Then the union of g_1 and g_2 forms a simple cycle, so this cycle is contained in some block B in Γ. Therefore, g_1 and g_2 are geodesics in B. Now by our assumption, B is l-almost geodetic, so g_1 and g_2 must intersect at some point other than v and w, which is a contradiction. □

Theorem 6.7. *Let* Γ *be a graph which is* l*-almost geodetic. If* l_1 *and* l_2 *are two non-intersecting infinite geodesics in* $\Gamma(G, S)$ *then the Hausdorff distance between* l_1 *and* l_2 *is infinite.*

Proof. Suppose for the sake of contradiction that the Hausdorff distance between l_1 and l_2 is finite and let this distance be k.

Let the vertices in l_1 be $a_0, a_1, \ldots$ in that order, let the vertices in l_2 be $b_0, b_1, \ldots$ in that order and let $m = d(a_0, b_0)$.

We will first prove that if $x, y \in \mathbb{Z}_{>0}$, then $|x - y| \leq m + d(a_x, b_y)$. If $x, y \in \mathbb{Z}_{>0}$, we have

$$x = d(a_0, a_x) < (a_0, b_0) + d(b_0, b_y) + d(a_x, b_y) = m + y + d(a_x, b_y),$$

so $x - y \leq m + d(a_x, b_y)$. Similarly, $y - x \leq m + d(a_x, b_y)$. Therefore,

$$|x - y| \leq m + d(a_x, b_y).$$

We will now prove that if $i, j \in \mathbb{Z}_{>0}$ then $d(a_i, b_j) - |i-j|$ is bounded between $-m$ and $m+2k$. If $i, j \in \mathbb{Z}_{>0}$, then since the Hausdorff distance between l_1 and l_2 is equal to k, there is some positive integer s such that $d(a_i, b_s) \leq k$. Now, $|i-s| \leq m + d(a_i, b_s) \leq m+k$ and $|i-j| \leq m + d(a_i, b_j)$. Therefore,

$$-m \leq d(a_i, b_j) - |i-j| \leq d(a_i, b_s) + d(b_s, b_j) - |i-j| \leq k + |s-j| - |i-j| \leq k + |s-i| \leq m+2k,$$

so $d(a_i, b_j) - |i-j|$ is bounded between $-m$ and $m+2k$.

Define $f : \mathbb{Z}^2 \to \mathbb{Z}$ by

$$f(i,j) = d(a_i, b_j) + i - j.$$

Then

$$f(i, j+1) = d(a_i, b_{j+1}) + i - j - 1 \leq d(a_i, b_j) + 1 + i - j - 1 = f(i,j),$$

so f is decreasing in j. Also,

$$f(i+1, j) = d(a_{i+1}, b_j) + i + 1 - j \geq d(a_i, b_j) - 1 + i + 1 - j = f(i,j),$$

so f is increasing in i.

Now, $f(i,j) = d(a_i, b_j) + i - j \geq d(a_i, b_j) - |i-j| \geq -m$. Therefore, for fixed i, f is decreasing in j and bounded below, so $\lim_{j\to\infty} f(i,j)$ exists. For $i \in \mathbb{Z}_{>0}$, let

$$c_i = \lim_{j \to \infty} f(i,j).$$

Note that for $j > i$, we have $f(i,j) = d(a_i, b_j) - |i-j| \leq 2k+m$, so $c_i \leq 2k+m$. Moreover, since f is increasing in i, the sequence $c_1, c_2, \ldots$ must also be increasing. But this sequence is bounded above, so it must converge to some constant c.

Let i_1 be an integer such that $c_{i_1} = c$ and let j_1 be an integer such that $f(i_1, j_1) = c_{i_1} = c$. Now let i_2 be any integer such that

$$i_2 > i_1 + d(a_{i_1}, b_{j_1}) + l.$$

Then since $i_2 > i_1$, we have $c_{i_2} = c$. Let $j_2 > j_1$ be an integer such that $f(i_2, j_2) = c_{i_2} = c$.

Now let g_1 be a geodesic path from a_{i_1} to b_{j_1} and let g_2 be a geodesic path from a_{i_2} to b_{j_2}. Let p_1 be the path along l_1 from a_{i_1} to a_{i_2-l}, let p_1' be the path along l_1 from a_{i_2-l} to a_{i_2} and let p_2 be the path along l_2 from b_{j_1} to b_{j_2}. Since $j_2 > j_1$ we have $f(i_1, j_2) = c$, so

$$d(a_{i_1}, b_{j_2}) = c + j_2 - i_1 = (c + j_1 - i_1) + (j_2 - j_1) = d(a_{i_1}, b_{j_1}) + d(b_{j_1}, b_{j_2}).$$

Also,

$$d(a_{i_1}, b_{j_2}) = c + j_2 - i_1 = (c + j_2 - i_2) + (i_2 - i_1) = d(a_{i_2}, b_{j_2}) + d(a_{i_1}, a_{i_2}).$$

Therefore, g_1 followed by p_2 is a geodesic from i_1 to j_2, and p_1 followed by p_1' followed by g_2 is also a geodesic from i_1 to j_2. Since p_1' is on l_1 and p_2 is on l_2, the path p_1' does not intersect with p_2. Also, the first vertex in p_1' is a_{i_2-l} and

$$d(a_{i_1}, a_{i_2-l}) = i_2 - l - i_1 > d(a_{i_1}, b_{j_1}),$$

whereas the last vertex in g_1 is b_{j_1}. So every vertex in p_1' is further from a_{i_1} than any vertex in g_1. Therefore, p_1' does not intersect with g_1.

Finally, let x be the last point in p_1 which intersects with g_1p_2 and let y be the first point in g_2 which intersects with g_1p_2. Then the sections of $p_1p_1'g_2$ and g_1p_2 between x and y form two non-intersecting geodesics. Moreover, one of these contains p_1', so $d(x, y) > l$. But this contradicts the fact that Γ is l-almost geodetic. □

6.2 Geodetic groups

Definition 6.8. We call a group G geodetic if it has some generating set S such that $\Gamma(G, S)$ is geodetic.

Definition 6.9. We call a group G almost geodetic if it has some generating set S such that $\Gamma(G, S)$ is l-almost geodetic for some l

Example 6.10. The group $\langle a, b, c, d : a^2, b^2, d^2, abab, (bcdc^{-1})^3\rangle$ is not basic as the subgroup $\langle a, b, cdc^{-1}\rangle$ is a non-trivial amalgamated free product. As expected, the Cayley graph with respect to this generating set is not geodetic (for example ab and ba are different geodesics which represent the same group element). It is, however, 12-almost geodetic, so there exist almost geodetic groups which are not basic.

Lemma 6.11. *Let G be a virtually free group with finite generating set S. There exists some constant k such that if p is a path in $\Gamma(G, S)$, between vertices u and v, and x is a vertex which lies on a geodesic between u and v, then the distance between x and p is at most k.*

Proof. Since G is virtually free, the Cayley graph $\Gamma(G, S)$ is quasi isometric to a tree. Let $f : \Gamma(G, S) \to T$ be a quasi-isomorphism and let d denote the path metric in both T and $\Gamma(G, S)$. Then there exists some constant c such that if $v_1, v_2 \in V(\Gamma(G, S))$ then

$$\frac{1}{c}d(v_1, v_2) - c \leq d(f(v_1), f(v_2)) \leq cd(v_1, v_2) + c.$$

We will show that the statement in the lemma is true for $k = 5c^4 + 3c^2$.

Let $(u = x_1), x_2, \ldots, x_{n-1}, (x_n = v)$ be a geodesic in $\Gamma(G, S)$, between vertices u and v, and let p be any path between u and v. Then we just need to show that for $1 < m < n$, the distance between x_m and p is at most $5c^4 + 3c^2$.

Let q be the simple path between $f(u)$ and $f(v)$ in T and let t be the closest point on q to $f(x_m)$. We will first show that $d(f(x_m), t) \leq 5c^3$. If $t = f(x_m)$, then

$$d(f(x_m), t) = d(f(x_m), f(x_m)) \leq cd(x_m, x_m) + c = c.$$

Otherwise, if $t \neq f(x_m)$, we break the tree T into two subgraphs T_1 and T_2 which only intersect at t:

- T_2 is the set of points s in T such that the simple path between s and $f(x_m)$ passes through t.

- $T_1 = (T \setminus T_2) \cup \{t\}$.

Then T_1 contains $f(x_m)$ and T_2 contains $f(u)$ and $f(v)$. Moreover, any path between a point in T_1 and a point in T_2 passes through t.

Let i and j be the maximum and minimum values respectively with $i < m < j$ such that $x_i, x_j \in T_2$. Then $x_{i+1}, x_{j-1} \in T_1$. Therefore,

$$\begin{aligned} d(f(x_i), f(x_j)) &\leq d(f(x_i), t) + d(t, f(x_j)) \\ &\leq d(f(x_i), f(x_{i+1})) + d(f(x_{j-1}), f(x_j)) \\ &\leq cd(x_i, x_{i+1}) + c + cd(x_{j-1}, x_j) + c \\ &= 4c. \end{aligned}$$

Therefore,

$$j - i = d(x_i, x_j) \leq cd(f(x_i), f(x_j)) + c^2 \leq 5c^2.$$

Hence,

$$d(f(x_m), t) \leq d(f(x_m), f(x_i)) \leq cd(x_m, x_i) + c = c(m - i + 1) \leq c(j - i) \leq 5c^3.$$

Now let the vertices in p be $(u = y_1), y_2, \dots, y_{n-1}, (y_n = v)$. Then, since the simple path between y_1 and y_n passes through t, there must be some $h \in \{1, \dots, n-1\}$ such that the simple path between y_h and y_{h+1} passes through t. Hence, the distance

$$d(y_h, t) \leq d(y_h, y_{h+1}) \leq 2c.$$

Therefore,

$$\begin{aligned} d(x_m, y_h) &\leq cd(f(x_m), f(y_h)) + c^2 \\ &\leq cd(f(x_m), t) + cd(t, f(y_h)) + c^2 \\ &\leq 5c^4 + 3c^2 \end{aligned}$$

□

Lemma 6.12. *Let G be a virtually free group and let S be a generating set for G such that $\Gamma(G, S)$ is almost geodetic. If $a \in G$ has infinite order, then there is a bi-infinite geodesic l and a finite constant k such that any point in l is within a distance of k from a vertex corresponding to an element of $\langle a \rangle$. Moreover, there is some t such that l is stabilised by a^t.*

Proof. From lemma 6.11, there is some constant k' such that any point on a geodesic between two points in $\Gamma(G, S)$ is within a distance of k' of any path p between the same two points. Let

$$k = k' + d(1, a).$$

Let $v_1, v_2, \dots, v_n$ be all of the vertices in the ball $B(1, k)$ of radius k around the identity. Now, since G is virtually free, the language of geodesics for the pair (G, S) is regular. Let c be the number of states in a (deterministic) finite state automaton for this language. Let $s \in \mathbb{Z}_{>0}$ such that $d(1, a^s) > cn$ and let p be a geodesic between 1 and a^s.

We construct a path p' between 1 and a^s by joining s copies of a geodesic between 1 and a. Then every point on p' is within a distance of $d(1, a)$ of an element of $\langle a \rangle$. Moreover, every vertex in p is within a distance of k' of a point in p'. Therefore, every vertex in p is within a distance of k of an element of $\langle a \rangle$. Hence each vertex in p can be written as $a^i v_j$ for some $i, j \in \mathbb{Z}$, with $1 \leq j \leq n$.

Let $g_1 g_2 \dots g_m$ be the word corresponding to p. Then $m > cn$, so by the pigeonhole principle, at least $n+1$ of the words $g_1, g_1 g_2, \dots, g_1 g_2 \dots g_m$ correspond to the same state of the finite state automaton. Hence at least 2 both correspond to the same state and are in the same set $\langle a \rangle v_j$. Let these two words be $g_1 \dots g_x$ and $g_1 \dots g_y$, with $y > x$. Since $g_1 \dots g_x$ and $g_1 \dots g_y$ are in the same state W of the finite state automaton, W is an accept state and the word $g_{x+1} \dots g_y$ moves W to W in the finite state automaton. Therefore, the word $g_1 g_2 \dots g_x$ followed by any power of $g_{x+1} \dots g_y$ corresponds to state W, and hence is a geodesic. Hence, we get a bi-infinite geodesic l by repeating the word $g_{x+1} \dots g_y$ infinitely many times in both directions, starting at v_j. Moreover, $\overline{g_{x+1} \dots g_y} = v_j^{-1} a^t v_j$ for some t, so

$$l \subset \langle a^t \rangle p \subset \langle a \rangle p.$$

Hence every vertex in l is within distance k of $\langle a \rangle$. □

Theorem 6.13. *Let G be a virtually free almost geodetic group. Then every two-ended subgroup of G is basic.*

Proof. Let S be a finite generating set for G such that $\Gamma(G, S)$ is almost geodetic. Let v_0 be the vertex corresponding to the identity in $\Gamma(G, S)$.

Suppose for the sake of contradiction that H is a 2-ended subgroup of G which is not basic.

Since H is 2-ended, it must be virtually $\mathbb{Z}$. Let $a \in H$ be a generator of such a subgroup, so H is virtually $\langle a \rangle$. Now, since a has infinite order, we can use lemma 6.12. So there is a bi-infinite geodesic l and a constant k such that any point in l is within a distance of k from some vertex corresponding to an element of $\langle a \rangle$.

Consider the set P of all geodesics p in Γ such that every vertex in p is within a distance of k of a vertex corresponding to an element of H. Then $l \in P$, so P is non-empty. Also, for any $p \in P$ and $h \in H$ we have $hp \in P$.

Now we will show that there is a vertex v in Γ such that every geodesic $p \in P$ passes through v.

First we will show that if $p_1, p_2 \in P$, then the Hausdorff distance between p_1 and p_2 is finite. Let S' be a finite generating set for H and let

$$k' = \max\{k\} \cup \{d(1, s) | s \in S'\}.$$

Then the subspace H' of Γ given by $H' = \{x \in \Gamma | d(x, H) \leq k'\}$ is connected. Moreover, any $p \in P$ is contained in H'. Finally, since H is two-ended, H' is also two-ended. Therefore, if x is a point in p_1, then there is a finite ball B_x centred at x whose removal breaks H' into two infinite pieces, so the ball must intersect p_2. By the definition of H', the group H acts on H' with finitely many orbits. Therefore, the radii of the balls B_x are uniformly bounded. Hence, every point of p_1 is within a uniformly bounded distance of p_2. Conversely, every point in p_2

is within a uniformly bounded distance of p_1. Therefore, the Hausdorff distance between p_1 and p_2 is finite.

Now consider the largest k such that there are vertices u and v in Γ with $d(u,v) = k$ and geodesics p_1 and p_2 between u and v, which only intersect at u and v such that p_1 and p_2 are both subpaths of geodesics in P. Note that k is finite because Γ is almost geodetic. We will show that any geodesic $p \in P$ passes through v.

Suppose that p does not pass through v.

If p does not intersect either p_1 or p_2, then we can extend p_1 in both directions until it does intersect with p, say at u' and v'. Then we get a pair of geodesics between u' and v' with the same condition as p_1 and p_2. But $d(u', v') > k$, which contradicts the maximality of k.

If p does intersect with p_1 or p_2, then without loss of generality, the closest intersection of p to v is with p_1. Let this intersection point be w. Now we extend p_2 beyond v until it intersects p and let the intersection point be v'. Then we have two non-intersecting geodesics between u and v', and $d(u, v') > k$, but this contradicts the maximality of k.

So all paths $p \in P$ pass through v.

Now, if $p \in P$ and $h \in H$, then $h^{-1}p \in P$, so $h^{-1}p$ passes through v. Therefore, p passes through hv.

Now let $p \in P$ be an arbitrary element and let $v_0 = v$. Also, let $\{\ldots, v_{-1}, v_0, v_1, \ldots\}$ be the orbit of v under H, with the order given by the order that the vertices appear along p. Then if $m_1 < m_2 < m_3$ are integers, we have $d(v_{m_1}, v_{m_2}) + d(v_{m_2}, v_{m_3}) = d(v_{m_1}, v_{m_3})$.

Now we can think of $\ldots, v_{-1}, v_0, v_1, \ldots$ as the vertices of a tree T, where for each i, there is an edge between v_i and v_{i+1}. Since $\{\ldots, v_{-1}, v_0, v_1, \ldots\}$ is an orbit of the action of H on Γ, any element h of H permutes the vertices of T. Moreover, two vertices v_i and v_j are adjacent if and only if there is no other vertex v_k such that $d(v_i, v_j) = d(v_i, v_k) + d(v_k, v_j)$, and this condition is preserved by the action of any $h \in H$. Hence, the action of H on $\{\ldots, v_{-1}, v_0, v_1, \ldots\}$ extends to an action on T. Finally, no element $h \in H$ fixes any vertex v_i, so no element fixes all of T. Therefore, H is a subgroup of the automorphism group $\mathrm{Aut}(T) = D_\infty$, so H is basic. □

Theorem 6.14. *Let G be a group with finite generating set S such that $\Gamma(G,S)$ is l-almost geodetic for some l. Suppose also that every block in $\Gamma(G,S)$ has 0 or 2 ends. Then G is basic.*

Proof. By theorem 3.19 We can write G as

$$G = F * \prod_k G_k,$$

where each G_k is the stabiliser of some block in $\Gamma(G,S)$.

Since each block has 0 or 2 ends, each block stabiliser G_k has 0 or 2 ends. If G_k has 0 ends it is finite and therefore basic. If G_k has 2 ends, then since it is a subgroup of G, the previous theorem implies that G_k is basic. Therefore, G is a free product of finitely many basic groups so it is itself basic. □

6.3 Almost geodetic groups which are not virtually free

Definition 6.15. Let Γ be a graph and let $v_0 \in V(\Gamma)$, call this the root vertex. We construct a tree T with the same vertex set as Γ as follows:

For each vertex $v \neq v_0$ in $V(\Gamma)$ at distance k from v_0, there is some edge $(u, v) \in E(\Gamma)$ such that u is at distance $k-1$ from v_0. For each v, choose one such edge (u, v) and add it to $E(T)$.

We call T the shortest path tree of Γ with root v_0.

Proposition 6.16. *If Γ is geodetic then for each vertex $v \in V(\Gamma)$ the shortest path tree of Γ with root v is unique.*

Proof. For each vertex v in Γ there is only one choice of edge (u, v) since otherwise there would be multiple shortest paths from v_0 to v. Therefore the shortest path tree is unique. □

Definition 6.17. For a constant α, we call two vertices u and v in Γ α-converging if there is a point a in Γ, at a distance of at least α from v_0, which may be either a mid edge or a vertex, such that u and v are on different geodesics between v_0 and a which intersect only at v_0 and a.

We simply call u and v converging if they are 1-converging.

Definition 6.18. We call two vertices u and v in Γ ∞-converging if they are α-converging for arbitrarily large α.

Lemma 6.19. *If x and y are ∞-converging vertices in Γ then there is a vertex x' which is adjacent to x, with $d(x', v_0) = d(x, v_0) + 1$ such that x' and y are ∞-converging.*

Proof. Let X be the set of vertices v which are adjacent to x and which satisfy $d(x_1, v_0) = d(x, v_0) + 1$. Then since the degree of x is finite, X is finite.

For any $\alpha > d(v_0, x)$, there is some a with $d(a, v_0) \geq \alpha$ and geodesics l_1 and l_2 between v_0 and a which intersect only at v_0 and a, such that x is on l_1 and y is on l_2. Let x_1 be the vertex adjacent to x on l_1 which lies between x and a. Then x_1 and y are α-converging. Also, $d(x_1, v_0) = d(x, v_0) + 1$, so $x_1 \in X$.

Therefore, for any α there is some vertex $x_1 \in X$ such that x_1 and y are α-converging. But X is finite, so there must be some $x' \in X$ such that x' and y are α-converging for arbitrarily large α. So x' and y are ∞-converging. □

For the remainder of this section, let G be a group with finite generating set S, let T be a shortest path tree of $\Gamma(G, S)$ whose root v_0 is the vertex corresponding to the identity, let d be the path metric in Γ and let d' be the path metric in T.

Lemma 6.20. *Suppose that there is some constant k such that if (u, v) is an edge in Γ, then $d'(u, v) < k$. Then G is virtually free.*

Proof. If x and y vertices in $\Gamma(G, S)$, let $d(x, y) = m$. Then there is a path

$$(x = x_0), x_1, \ldots, x_{m-1}, (x_m = y)$$

in $\Gamma(G, S)$. So $d'(x, y) \leq d'(x_0, x_1) + \ldots + d'(x_{m-1}, x_m) \leq mk$.

So for any vertices x and y in $\Gamma(G, S)$, we have $d'(x, y) \leq kd(x, y)$. Therefore the map $f : V(\Gamma(G, S)) \to V(T)$ which takes vertices to themselves but changes metric from d to d' is a quasi-isometry. Therefore G is quasi isometric to some tree, so it is virtually free. □

Proposition 6.21. *If G is not virtually free, there exist geodesics $a_0, a_1, \ldots$ and $b_0, b_1, \ldots$ which only intersect at $a_0 = b_0 = v_0$ such that for each $i, j > 0$, a_i and b_j are ∞-converging.*

Proof. From the previous lemma, there is no constant k such that $d'(u, v) < k$ for every edge (u, v) in Γ. So for any integer c, there is some edge (u_c, v_c) in Γ such that $d'(u_c, v_c) > 2c$. Let x be the vertex such that $d'(u_c, v_c) = d'(u_c, x) + d'(x, v_c)$, $d'(u_c, v_0) = d'(u_c, x) + d'(x, v_0)$ and $d'(v_0, v_c) = d'(v_0, x) + d'(x, v_c)$.

Now let $g \in G$ satisfy $gv_0 = x$. Let $g^{-1}v_c = v'_c$ and let $g^{-1}u_c = u'_c$. Then $g^{-1}l_u$ and $g^{-1}l_v$ are non-intersecting geodesics from v_0 to v'_c and u'_c respectively.

So for any constant c there is an edge (u'_c, v'_c) in Γ at a distance of at least c from v_0 such that there exist non-intersecting geodesics l'_u and l'_v from v_0 to u'_c and v'_c respectively. Therefore, v_0 and v_0 are ∞-converging.

Now we construct vertices $a_0, a_1, \ldots$ and $b_0, b_1, \ldots$ as follows:

- $a_0 = b_0 = v_0$.
- After choosing a_i and b_i which are ∞-converging, by lemma 6.19 there is some vertex x' which is adjacent to a_i such that $d(v_0, x') = d(v_0, a_i) + 1$ and x' and b_i are ∞-converging. Then, again using lemma 6.19, there is some y' adjacent to b_i such that $d(v_0, y') = d(v_0, b_i) + 1$ and x' and y' are ∞-converging. Let $a_{i+1} = x'$ and let $b_{i+1} = y'$.

Then by construction, $a_0, a_1, \ldots$ and $b_0, b_1, \ldots$ are infinite geodesics and for any i, a_i and b_i are ∞-converging. So if $j < i$, the pairs (a_i, b_j) and (a_j, b_i) are infinity converging. □

Theorem 6.22. *Let G be a group with finite generating set S such that $\Gamma(G, S)$ is l-almost geodetic for some l. If (G, S) has the falsification by fellow traveller property, then G is virtually free.*

Proof. Let (G, S) satisfy the synchronous falsification by fellow traveller property with respect to the constant k.

Suppose for the sake of contradiction that G is not virtually free. Then from proposition 6.21 we know that there are geodesics $(v_0 = a_0), a_1, a_2, \ldots$ and $(v_0 = b_0), b_1, b_2, \ldots$ in $\Gamma(G, S)$ such that (a_i, b_i) is ∞-converging for each i.

From theorem 6.7, we know that the Hausdorff distance between the geodesics $a_1, a_2, \ldots$ and $b_1, b_2, \ldots$ is infinite. In particular this means that $d(a_i, b_i)$ is unbounded. Let r be a positive integer such that $d(a_r, b_r) > l + k$.

Since a_r and b_r are converging, there is a point u in $\Gamma(G, S)$ such that a_r and b_r are on different geodesics between v_0 and u which only intersect at v_0 and u. If u is a vertex of $\Gamma(G, S)$ then, since $\Gamma(G, S)$ is l-almost geodetic, we must have $d(u, v_0) \leq l$, which clearly contradicts the fact that $d(a_r, b_r) > l + k \geq l$.

Therefore, u is a mid edge. So $d(v_0, u) = n + 1/2$ for some positive integer n.

Let $v_0 = a'_0, a'_1, \ldots, a'_n$ and $v_0 = b'_0, b'_1, \ldots, b'_n$ be non-intersecting geodesics where a'_n and b'_n are the vertices on either side of u, $a'_r = a_r$ and $b'_r = b_r$. Then $b'_0, b'_1, \ldots, b'_n, a'_n$ is a path which is not a geodesic. So by the falsification by fellow traveller property, there is a shorter path $(v_0 = b'_0 = c_0), c_1, \ldots, c_{m-1}, (c_m = a'_n)$ such that for each $i \leq m$, we have $d(c_i, b'_i) < k$.

Now, $m \geq d(c_0, c_m) = d(v_0, a'_n) = n$ and $m < n + 1$, so $m = n$. Hence $v_0 = c_0, c_1, \ldots, c_n = a'_n$ is a geodesic. So $c_0, c_1, \ldots, c_n$ and $a'_0, a'_1, \ldots, a'_n$ are geodesics between the same two vertices of $\Gamma(G, S)$. Therefore, since $\Gamma(G, S)$ is l-almost geodetic, $d(a'_r, c_r) \leq l$.

But $d(c_r, b'_r) < k$, so $l + k > d(c_r, b'_r) + d(a'_r, c_r) \geq d(a'_r, b'_r) = d(a_r, b_r) > l + k$, a contradiction.

Therefore, G is virtually free. □

6.4 Geodetic groups which are not virtually free

In this section, let G be a geodetic group which is not virtually free. We wish to reach a contradiction. We certainly fall well short of that, though we will prove that any geodetic Cayley graph Γ for such a group G does not satisfy 2-DP. Hence, the language of geodesics in $\Gamma(G, S)$ is not regular.

Let S be a finite generating set for G such that $\Gamma(G, S)$ is geodetic and let v_0 be the vertex in $\Gamma(G, S)$ corresponding to the identity. From proposition 6.21 we have geodesics $a_0, a_1, \ldots$ and $b_0, b_1, \ldots$ which only intersect at $a_0 = b_0 = v_0$ such that for each $i, j > 0$, a_i and b_j are ∞-converging.

Lemma 6.23. *The path* $\ldots, b_2, b_1, (b_0 = v_0 = a_0), a_1, a_2, \ldots$ *is a bi-infinite geodesic.*

Proof. We just need to prove that for any positive integers i, j, the geodesic between a_i and b_j passes through v_0.

Since a_{i+j} and b_{i+j} are converging in $\Gamma(G, S)$, there is a point u in $\Gamma(G, S)$ such that a_{i+j} and b_{i+j} are on different geodesics between v_0 and u which only intersect at v_0 and u. Let the vertices in these geodesics be $v_0 = a'_0, a'_1, \ldots, a'_n$ and $v_0 = b'_0, b'_1, \ldots, b'_n$ with $a'_{i+j} = a_{i+j}$ and $b'_{i+j} = b_{i+j}$. Note that this also means that $a'_i = a_i$ and $b'_j = b_j$.

Since $\Gamma(G, S)$ is geodetic, u is a mid edge, not a vertex. So b'_n and a'_n are joined by and edge.

Therefore, by proposition 6.3, $a'_0, a'_1, \ldots, a'_n, b'_n, \ldots, b'_1, (b'_0 = a'_0)$ is a convex cycle. So the geodesic between a'_i and b'_j is either $a'_i, a'_{i-1}, \ldots, a'_1, (a'_0 = b'_0), b'_1, \ldots, b'_i$ or $a'_i, a'_{i+1}, \ldots, a'_n, b'_n, b'_{n-1}, \ldots, b'_i$.

The first possibility has length $i + j$ whereas the second has length

$$n - i + n - j + 1 > i + j - i + i + j - j + 1 = i + j + 1.$$

So $a'_i, a'_{i-1}, \ldots, a'_1, a'_0 = b'_0, b'_1, \ldots, b'_i$ must be the geodesic.

Therefore, the geodesic between a_i and b_j passes through v_0 □

Let the geodesic $\ldots, b_2, b_1, b_0 = v_0 = a_0, a_1, a_2, \ldots$ be called l.

Lemma 6.24. *For any positive integer i, there are integers $r, s > i$ and a path γ given by $c_0, c_1 \dots, c_n$ which only intersects l at $c_0 = a_r$ and $c_n = b_s$ and such that the union of γ and l is convex.*

Proof. Since a_i and b_i are ∞-converging, there is a convex cycle which contains the path $a_i, a_{i-1}, \dots, b_i$. Let $(a_{r'} = c'_0), c'_1, \dots, c'_{n'-1}, (c'_{n'} = b_{s'})$ be the section of this cycle which does not travel along l. Then $s', r' > i$. Also $n' > r' + s'$, since $d(c_0, c_{n'}) = d(a_{r'}, b_{s'}) = r' + s'$.

Consider such a convex cycle which minimises the value of $n' - r' - s'$. This gives rise to a path $c_0, \dots, c_n$ which only intersects l at $c_0 = a_r$ and $c_n = b_s$. We will prove that the union Q of l with the path $c_0, \dots, c_n$ is convex.

Let x, y be vertices in the union of l with the path $c_0, \dots, c_n$. Then we just need to show that the geodesic between x and y is contained in this union. If x and y are both on the geodesic l, then the geodesic between x and y is contained in $l \subset Q$. If x and y are both in the convex cycle

$$(a_r = c_0), c_1, \dots, c_{n-1}, (c_n = b_s), b_{s-1}, \dots, a_{r-1}, a_r,$$

then since this cycle is convex, it contains the geodesic between x and y. The final remaining case is where one of the vertices y is on the cycle but not l and the other vertex x is on l but not the cycle. Without loss of generality, assume that $x = b_k$ for some $k > s$. Suppose for sake of contradiction that the geodesic between $x = b_k$ and $y = c_j$ is not contained in Q. We will assume that k is minimal such that this happens for some c_j and j is minimal for this value of k. Then the geodesic between b_k and c_j is not contained in Q, but the geodesic between b_k and c_{j-1} is contained in Q. Clearly then, the geodesic between b_k and c_{j-1} cannot contain c_j, so it must be the path

$$c_{j-1}, \dots, c_1, (c_0 = a_r), a_{r-1}, \dots, b_{k-1}, b_k.$$

Let this geodesic be called p_1 and let the geodesic between b_k and c_j be called p_2. Then, since p_2 is a geodesic, it is shorter than the path

$$b_k, b_{k-1}, \dots, b_{s+1}, (b_s = c_n), c_{n-1}, \dots, c_j.$$

Hence $|p_2| < k - s + n - j$. Now, if p_1 and p_2 intersect at some point other than b_k, it must be a vertex $b_{k'}$ with $k' < k$, but this contradicts the minimality of k. Hence, p_1 and p_2 are geodesics which only intersect at b_k and whose other endpoints are joined by an edge. Therefore, by proposition 6.3, the cycle formed by p_1, p_2 and the edge (c_{j-1}, c_j) is convex. Hence we have a new path γ' which only intersects l at b_k and a_r, such that the union of γ' with the path $b_k, b_{k-1}, \dots, a_{r-1}, a_r$ is convex. But p_2 is a geodesic, so it is shorter than the path $b_k, b_{k-1}, \dots, b_{s+1}, (b_s = c_n), c_{n-1}, \dots, c_j$

$$|\gamma'| - k - r = |p_2| + j - k - r < (k - s + n - j) + (j - k - r) = n - s - r.$$

But this contradicts the minimality of $n - r - s$.

Therefore, the union Q of l and the path $c_0, c_1, \dots, c_n$ is convex. □

Clearly this implies that Γ contains arbitrarily large convex cycles, so it does not satisfy 2-DP.

This seems like a very unusual thing for a geodetic graph to do, though as it turns out, it is not so unusual for Cayley graphs. For example, we saw in chapter 5 that any group with a greater than exponential Dehn function does not satisfy 2-DP. We also give the following example to show that there certainly are fairly simple Cayley graphs with the more specific property described in lemma 6.24.

Example 6.25. The standard Cayley graph of the Baumslag-Solitar group $BS(1,4)$, which has the presentation $\langle a, b \mid ab^4 = ba \rangle$, has the more specific property described in lemma 6.24. We can see this by taking the subgroup generated by a as our bi-infinite geodesic l, so a_i is the vertex corresponding to a^i. Then for any two vertices a_i and a_j, with $j > i$, there is a path $a_i = c_0, c_1, \ldots, c_{3(j-i)+4} = a_j$ corresponding to the word $ba^{j-i}b(a^{-1})^{j-i}b^{-1}a^{j-i}b^{-1}$. we can easily check algebraically that the path only intersects l at its two ends. Moreover, it is easy to check that the union of this path with l is indeed convex.

6.5 Almost geodetic groups which are virtually free

We conjecture that a finitely generated group G embeds into an almost geodetic group G' if and only if G is virtually free and every two-ended subgroup of G is basic. Given that we are unable to prove that geodetic groups are virtually free, we are certainly unable to prove that almost geodetic groups are virtually free. Instead we will consider a more approachable problem: We conjecture that a finitely generated group G embeds into an almost geodetic, virtually free group G' if and only if G is virtually free and every two-ended subgroup of G is basic.

Notice that if $G \leq G'$ and G' is almost geodetic, then every subgroup of G is also a subgroup of G'. Since we know that every two-ended subgroup of an almost geodetic group is basic, this implies that every two-ended subgroup of G is basic. Unfortunately, we are unable to show the other direction in general. In the next chapter, however, we will prove the statement in the case that G is a Coxeter group. As it turns out, this is also equivalent to an interesting geometric property of Coxeter groups, namely that all walls in the Cayley graph are finite.

Chapter 7

A classification of Coxeter groups where all walls are finite

7.1 Introduction to Coxeter groups

For a more thorough introduction to Coxeter groups see [7].

Definition 7.1. A Coxeter group is a group W with presentation

$$\langle s_1, s_2, \ldots, s_n \mid (s_i s_j)^{m_{ij}} \ \forall 1 \leq i, j \leq n \rangle,$$

where each $m_{ii} = 1$ and $m_{ij} = m_{ji}$ for $i \neq j$. In particular, since each $m_{ii} = 1$, each of the generators is an involution. We call the pair $(W, \{s_1, s_2, \ldots, s_n\})$ a Coxeter system.

Definition 7.2. Let (W, S) be a Coxeter system with presentation as above. Whenever $m_{ij} < \infty$, we define c_{ij} to be the alternating product of s_i and s_j of length m_{ij}, starting with s_i. This is sometimes denoted $(s_i s_j)^{m_{ij}/2}$

Note that $(s_i s_j)^{m_{ij}} = c_{ij} c_{ji}^{-1}$, so we can rewrite each relation $(s_i s_j)^{m_{ij}}$ as

$$c_{ij} = c_{ji}.$$

We can easily check algebraically that if m_{ij} is even then $s_i c_{ij} = c_{ij} s_i$ and if m_{ij} is odd then $s_i c_{ij} = c_{ij} s_j$.

Definition 7.3 (Coxeter diagram)**.** Let (W, S) be a Coxeter system with the usual presentation. The Coxeter diagram is a graph with vertex set S and with an edge between each pair s_i, s_j labelled by m_{ij}. As a standard convention, we draw no edge if $m_{ij} = 2$ and we draw an unlabelled edge if $m_{ij} = 3$.

Using this convention, if the Coxeter diagram A of a Coxeter group (W, S) is not connected, then W decomposes as a direct product of the Coxeter groups given by the connected components of A.

Theorem 7.4 (Classification of finite Coxeter groups)**.** *The finite Coxeter groups are direct products of Coxeter groups given by the Coxeter diagrams show below. Except in the case of I_n, the subscript n corresponds to the number of generators, and hence the number of vertices in the diagram.*

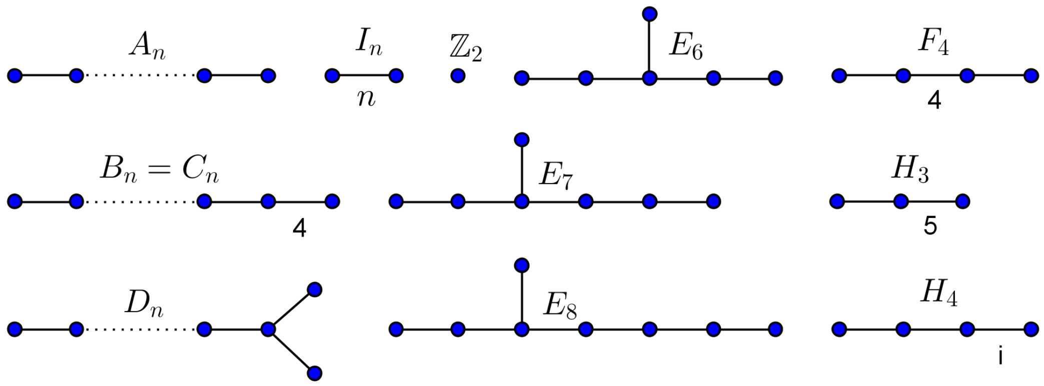

Note that $\mathbb{Z}_2$ is actually the same as A_1, so we don't really need to include it separately.

Definition 7.5. If (W, S) is a Coxeter system and $T \subset S$, we define W_T to be the subgroup of W generated by T. We call this a special subgroup. We will also define $W^T = W_{S \setminus T}$.

The following two theorems are part of theorem 4.1.6 in [7].

Theorem 7.6. *If $T \subset S$ then (W_T, T) is a Coxeter system.*

Theorem 7.7. *If $T_1, T_2 \subset S$ then $W_{T_1} \cap W_{T_2} = W_{T_1 \cap T_2}$.*

Definition 7.8. If $r \in W$ is a conjugate of a generator, we define Ω^r to be the set of midpoints of edges in $\Gamma(W, S)$ which are flipped by the action of r on $\Gamma(W, S)$. We call this the wall corresponding to r.

Theorem 7.9. *If (W, S) is a Coxeter system and Ω^r is a wall in W, then the space $\Gamma(W, S) \setminus \Omega^r$ is made up of two isomorphic connected components which are interchanged by r. In other words, each wall splits the Cayley graph $\Gamma(W, S)$ into two isomorphic pieces.*

Note that if r is conjugate to a generator s_j with $r = w s_j w^{-1}$, then $\Omega^r = w\Omega^{s_j}$. In particular, this means that the wall Ω^r is finite if and only if Ω^{s_j} is finite.

Proposition 7.10. *The centraliser of s_j is finite if and only if Ω^{s_j} is finite.*

Proof. Let $C(s_j)$ denote the centraliser of s_j. By the definition of the wall Ω^{s_j}, we can see that $C(s_j)$ is exactly the stabiliser of Ω^{s_j}. Now, for any two edges e_1, e_2 with the same label which both intersect Ω^{s_j}, there is some element $w \in W$ which sends e_1 to e_2, so w stabilises Ω^{s_j} and hence $w \in C(s_j)$. Therefore, $C(s_j)$ acts with finitely many orbits on Ω^{s_j}. Finally, the stabiliser of any point in Ω^{s_j} contains exactly two elements: 1 and s_j. Therefore, $C(s_j)$ is finite if and only if Ω^{s_j} is finite. □

Definition 7.11. Let (W, S) be a Coxeter system. The nerve of W is an abstract simplicial complex $L(W, S)$ constructed as follows:

- The vertex set of $L(W, S)$ is equal to S.
- A subset $T \subset S$ defines a simplex of L if and only if W_T is finite. We denote this simplex by σ_T.

Theorem 7.12. *([7], theorem 8.7.2) W is one-ended if and only if, for each simplex $\sigma_T \in L$, the punctured nerve $L - \sigma_T$ is connected. In other words, the nerve of the Coxeter system $(W_{S\setminus T}, S \setminus T)$ is connected.*

Proposition 7.13. *If (W, S) is a Coxeter system such that some wall of $\Gamma(W, S)$ is finite, then W is not one-ended.*

Proof. If W is finite then W is 0-ended, so we are done. If W is infinite, let Ω^r be a finite wall in $\Gamma(W, S)$. Then Ω^r cuts $\Gamma(W, S)$ into two isomorphic pieces R_1 and R_2. Since W is infinite, both of these pieces are infinite. Now, since Ω^r is a wall in Γ, it must be finite, so Ω^r is contained in some ball $B(x_0, r)$. Therefore, $\Gamma(W, S) \setminus B(x_0, r)$ contains at least two infinite pieces, so W has at least two ends. □

We will now prove two simple propositions which we will use many times later.

Proposition 7.14. *If m_{ij} is even, then $c_{ij} \notin W^{\{s_j\}}$*

Proof. By theorem 7.7, we have

$$W^{\{s_j\}} \cap W_{\{s_i,s_j\}} = W_{\{s_i\}} = \{1, s_i\}.$$

Since $c_{ij} \notin \{1, s_i\}$ and $c_{ij} \in W_{\{s_i,s_j\}}$, we have $c_{ij} \notin W^{\{s_j\}}$. □

Proposition 7.15. *If m_{ij} and m_{jk} are odd, then $c_{ij}c_{jk} \notin W^{\{s_j\}}$*

Proof. Suppose for the sake of contradiction that $c_{ij}c_{jk} \in W^{\{s_j\}}$. Then

$$c_{ij}c_{jk} \in W^{\{s_j\}} \cap W_{\{s_i,s_j,s_k\}} = W_{\{s_i,s_k\}}.$$

Also, $s_i c_{ij} c_{jk} = c_{ij} s_j c_{jk} = c_{ij} c_{jk} s_k$. Hence, there is an edge e labelled by s_k between $c_{ij}c_{jk}$ and $s_i c_{ij} c_{jk}$, so e must intersect with the wall Ω^{s_i}. But the Cayley graph of $W_{\{s_i,s_k\}}$ is just a regular $2m_{ik}$-gon, so the only edges in it which intersect with Ω^{s_i} are the edge labelled s_i which connects to the identity, and the opposite edge. Hence e must be the opposite edge, and it must be labelled s_k. Hence $c_{ij}c_{jk} = c_{ik}$ or $c_{ik}s_k$. Notice that $c_{ij}c_{jk}$ is on the same side of the wall Ω^{s_i} as the identity, therefore, we must have

$$c_{ij}c_{jk} = c_{ik}s_k = s_i c_{ik}.$$

Now consider the wall Ω^{s_k}. The path given by $s_i c_{ik}$ only crosses this wall once, so the path $c_{ij}c_{jk}$ must also cross the wall at least once. Let w and ws be subwords from the start of $c_{ij}c_{jk}$ such that $s \in \{s_i, s_j, s_k\}$ and the elements $\overline{w}$ and $\overline{ws}$ are on opposite sides of Ω^{s_k}. Then $\overline{wsw^{-1}} = s_k$. Hence wsw^{-1} contains the letter s_k at least once, so c_{ij} is a subword of w. Let $w = c_{ij}w_1$, so $ws = c_{ij}w_1 s$, and $w_1 s$ is a subword from the start of c_{jk}. Hence, $\overline{s_k c_{jk} w_1 s w_1^{-1} c_{jk} s_k} \in W_{\{s_j,s_k\}}$. But

$$\begin{aligned}\overline{s_k c_{jk} w_1 s w_1^{-1} c_{jk} s_k} &= \overline{s_k c_{jk} c_{ij} w s w^{-1} c_{ij} c_{jk} s_k} \\ &= s_k (c_{ij} c_{jk})^{-1} s_k c_{ij} c_{jk} s_k \\ &= s_k (c_{ik} s_k)^{-1} s_k c_{ik} s_k s_k \\ &= c_{ik} s_k c_{ik} = s_i,\end{aligned}$$

which is a contradiction, since $s_i \notin W_{\{s_j,s_k\}}$. □

7.2 Statement of the theorem

Before we state the theorem, we will need to define the set $\mathscr{X}$ of Coxeter matrices which we will later prove classifies Coxeter groups with finite walls, as well as a number of other equivalent conditions.

Definition 7.16. We define $\mathscr{F}$ to be the set of all Coxeter systems (W, S), where W is finite.

Definition 7.17. For $m, n \in \mathbb{Z}_{>0}$ we define $L_{2m+1,2n+1}$ and L_{2m+1} to be the Coxeter systems given by the Coxeter diagrams shown below:

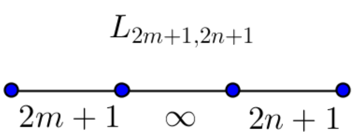

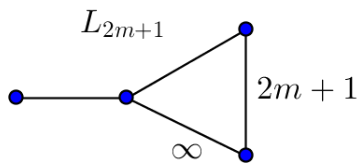

Definition 7.18. We define $\mathscr{X}_1$ to be the smallest set containing $\mathscr{F}$ and every Coxeter system of the form $L_{m,n}$ or L_n, such that if M is an $n \times n$ matrix representing a Coxeter system in $\mathscr{X}_1$ and k, n are positive integers with $1 \leq k \leq n$ then the Coxeter system (W', S') given by the $(n+1) \times (n+1)$ matrix M' defined by:

$$M'_{ij} = \begin{cases} M_{ij} & \text{if } 1 \leq i, j \leq n \\ 2m+1 & \text{if } \{i, j\} = \{k, n+1\} \\ 1 & \text{if } i = j = n+1 \\ \infty & \text{otherwise} \end{cases}$$

is also an element of $\mathscr{X}_1$. We call the Coxeter system (W', S') a $(2m+1)$-extension of (W, S). Note that this operation corresponds to adding a vertex to the Coxeter diagram of (W, S) and joining it to one vertex by an edge labelled $2m + 1$ and every other vertex by an edge labelled ∞.

Definition 7.19. Finally we define $\mathscr{X}$ to be the closure of $\mathscr{X}_1$ under free products. That is, $\mathscr{X}$ is the set of Coxeter systems of the form

$$(W_1 * W_2 * \ldots * W_k, S_1 \cup \ldots \cup S_k),$$

where each pair (W_i, S_i) is a Coxeter system in $\mathscr{X}_1$.

In this chapter we will prove the following theorem:

Theorem 7.20. *Let (W, S) be a Coxeter system. The following statements are equivalent:*

(a) The Coxeter system $(W, S) \in \mathscr{X}$.

(b) W is virtually free and there exists a finite generating set $S' \subset W$ which contains S and a function $f : S' \to 2\mathbb{Z}_{\geq 0} + 1$ such that $\Gamma^f(W, S')$ is almost geodetic. (See example 4.11 for the definition of $\Gamma^f(W, S')$.)

(c) W embeds into a virtually free almost geodetic group G.

(d) W is virtually free and contains no two-ended non-basic subgroups.

(e) W is virtually free and the group $\mathbb{Z} \times \mathbb{Z}_2$ is not a subgroup of W.

(f) Every wall in the Cayley graph $\Gamma(W, S)$ is finite.

We will prove this in a long implication chain of the form:

$$(a) \Rightarrow (b) \Rightarrow (c) \Rightarrow (d) \Rightarrow (e) \Rightarrow (f) \Rightarrow (a)$$

The implications (b) $\Rightarrow$ (c) $\Rightarrow$ (d) $\Rightarrow$ (e) $\Rightarrow$ (f) are fairly straightforward consequences of previous results in this thesis, so we will prove them immediately.

We will prove that (a)$\Rightarrow$ (b) in theorem 7.31 and, finally we will prove that (f)$\Rightarrow$ (a) in theorem 7.51 to complete the chain of implications.

Proposition 7.21. *((b)$\Rightarrow$(c)) Let (W, S) be a Coxeter system with W virtually free and let S' be a generating set for W which contains S. Suppose further that there exists some function $f : S' \to 2\mathbb{Z}_{\geq 0}+1$ such that $\Gamma^f(W, S')$ is almost geodetic. Then W embeds into a virtually free, almost geodetic group G.*

Proof. We know from example 4.11 that there is an integer k and a generating set S'' for $G * F_k$ such that every block in $\Gamma(W * F_k, S'')$ is isomorphic to $\Gamma^f(W, S')$. Therefore, since $\Gamma^f(W, S')$ is almost geodetic, every block in $\Gamma(W * F_k, S'')$ is almost geodetic. Therefore, $\Gamma(W * F_k, S'')$ is an almost geodetic graph, so $W * F_k$ is an almost geodetic group. Therefore, W embeds into the almost geodetic group $W * F_k$.

Finally, since W is virtually free, $W * F_k$ is also virtually free. So W embeds into the almost geodetic, virtually free group $W * F_k$. □

Proposition 7.22. *((c)$\Rightarrow$(d)) Let W be a Coxeter group. If W embeds into a virtually free, almost geodetic group G then W is virtually free and contains no two-ended subgroups.*

Proof. Since G is virtually free and almost geodetic, we know by theorem 6.13 that G does not contain any two-ended non-basic subgroups. Therefore, W contains no two-ended subgroups. Moreover, since W is a subgroup of a virtually free group, W is virtually free. □

Proposition 7.23. *((d)⇒(e)) Let W be a virtually free Coxeter group which contains no two-ended, non-basic subgroups. Then $\mathbb{Z} \times \mathbb{Z}_2$ is not a subgroup of W.*

Proof. The only two-ended basic groups are $\mathbb{Z}$ and D_∞, so $\mathbb{Z} \times \mathbb{Z}_2$ is a two-ended, non-basic group. Therefore, $\mathbb{Z} \times \mathbb{Z}_2$ is not a subgroup of W. □

Proposition 7.24. *((e)⇒(f)) Let (W, S) be a Coxeter system, with W virtually free. If $\mathbb{Z} \times \mathbb{Z}_2$ is not a subgroup of W then every wall in the Cayley graph $\Gamma(W, S)$ is finite.*

Proof. We will prove the contrapositive of this statement:

Let (W, S) be a Coxeter system with W virtually free such that $\Gamma(W, S)$ has a wall which is infinite. We will show that $\mathbb{Z} \times \mathbb{Z}_2$ is a subgroup of W.

Let Ω^r be an infinite wall in $\Gamma(W, S)$ and let r be the corresponding reflection. Since Ω^r is infinite and there are only finitely many labels on the edges, Ω^r must intersect infinitely many edges in $\Gamma(W, S)$ with the same label s_j. Let e be an edge with label s_j which intersects with Ω^r. Since e has label s_j, there is an isomorphism f of $\Gamma(G, S)$ which sends e to the edge labelled s_j which connects to the identity. Therefore, f sends Ω^r to the wall Ω^{s_j}. Moreover, every intersection of Ω^r with an edge labelled s_j is sent to an intersection of Ω^{s_j} with an edge labelled s_j, so Ω^{s_j} must intersect with infinitely many walls labelled Ω^{s_j}.

Let E_H be the set of edges in $\Gamma(G, S)$ with label s_j which intersect Ω^r. Then E_H is infinite. Let H be the set of group elements corresponding to vertices at the ends of edges in E_H. If $h \in H$ then the edge labelled s_j which connects to h is flipped by s_j. Therefore, the reflection s_j sends h to hs_j, so $s_jh = hs_j$. Also, if $g \in G$ and $s_jh = hs_j$ then the action of s_j on $\Gamma(G, S)$ sends g to $s_jg = gs_j$ and gs_j to $s_jgs_j = g$, so s_j flips the edge between g and gs_j (which, of course, is labelled s_j). Therefore, this edge intersects the wall Ω^{s_j}, so $g \in H$. Therefore, H is the centraliser of s_j.

So H is an infinite subgroup of W. Since W is a virtually free group, H is a virtually free group, so H must contain an element g of infinite order. Hence, $\langle g \rangle$ is isomorphic to $\mathbb{Z}$. Also, $g \in H$ which is the centraliser of s_j, therefore, $gs_j = s_jg$. Finally, s_j is in the generating set for W so it has order 2. Therefore, $\langle g, s_j \rangle$ is isomorphic to $\mathbb{Z} \times \mathbb{Z}_2$, so $\mathbb{Z} \times \mathbb{Z}_2$ is a subgroup of W. □

7.3 Coxeter groups in $\mathscr{X}$ embed into almost geodetic groups

In this section we will prove the implication $(a) \Rightarrow (b)$. We will first define the set of Coxeter systems which satisfy condition (b) to be $\mathscr{T}$. Then we just need to prove that $\mathscr{X} \subseteq \mathscr{T}$.

Definition 7.25. Let $\mathscr{T}$ be the set of Coxeter systems (W, S) such that there exists a finite generating set S' for W, such that $S \subset S'$ and a function $f : S' \to 2\mathbb{Z}_{\geq 0} + 1$ such that $\Gamma^f(W, S')$ is almost geodetic.

Lemma 7.26. *If (W, S) is a Coxeter system and W is finite, then $\Gamma(W, S)$ is almost geodetic.*

Proof. Since W is finite, $\Gamma(W, S)$ is a finite graph, so it is certainly almost geodetic. □

Lemma 7.27. *Coxeter systems of the form $L_{2m+1,2n+1}$ are is $\mathscr{T}$.*

Proof. Let $S = a, b, c, d$ and let W be given by the presentation

$$W = \langle a, b, c, d \mid a^2, b^2, c^2, d^2, (ab)^{2m+1}, (cd)^{2n+1}, (ac)^2, (bc)^2, (bd)^2 \rangle.$$

Then $(W, S) = L_{2m+1,2n+1}$.

Now let $S' = \{a, b, c, d, ac, bd\}$ and let $f : S' \to 2\mathbb{Z}_{\geq 0} + 1$ be the constant function with image $\{1\}$. We will show that $\Gamma^f(W, S') = \Gamma(W, S')$ is almost geodetic.

Let d be the path metric in $\Gamma(W, S')$ and let u and v be vertices in $\Gamma(W, S')$ with $d(u, v) > \max\{2m+1, 2n+1\}$. We will show that any two geodesics between u and v intersect at some point between u and v, hence showing that $\Gamma(W, S')$ is almost geodetic.

Recall that W decomposes as an amalgamated free product

$$W = W^{\{a\}} *_{W_{\{b,c\}}} W^{\{d\}},$$

so we can consider the Bass-Serre tree for the amalgamated free product. The chambers of the Cayley graph $\Gamma(W, S')$ are shown below. The edges are highlighted if they are on a geodesic in $\Gamma(W, S')$ from the red vertex at the bottom left of each diagram. The Cayley graph $\Gamma(W, S')$ is made up of these chambers, attached to each other on the squares with edges labelled by b and c.

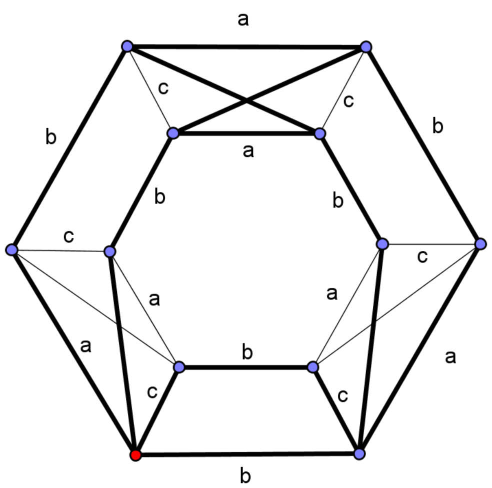

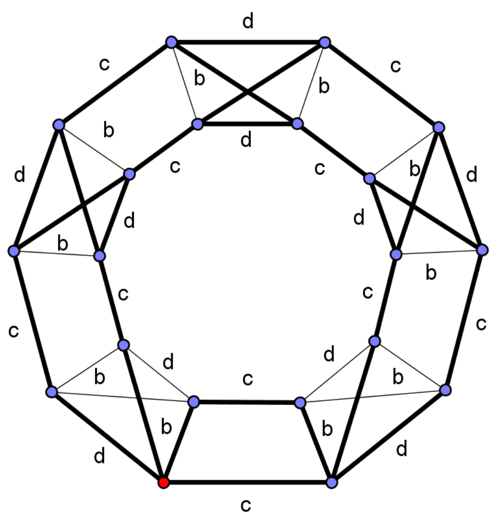

Since $d(u, v) > \max\{2m+1, 2n+1\}$, the vertices u and v cannot be in the same chamber of W. Let e and e' be the edges in the Bass-Serre corresponding to the cosets $uW_{\{b,c\}}$ and $vW_{\{b,c\}}$ respectively. Since u and v are not in the same chamber, the edges e and e' do not meet at a vertex. Let $(e = e_1), e_2, \ldots, e_{k-1}, (e_k = e')$ be the simple path between the two edges e and e', so $k \geq 3$. Then, since each element of S' is in one of the groups $W^{\{a\}}, W^{\{d\}}$, the left coset of $W_{\{b,c\}}$ corresponding to any edge e_i must intersect with any path between u and v.

Let $g_1W_{\{b,c\}}$, $g_2W_{\{b,c\}}$ and $g_3W_{\{b,c\}}$ be the left cosets corresponding to e_1, e_2 and e_3 respectively. Let x and y be vertices in the Bass-Serre tree such that x is contained in both e_1 and e_2 and y is contained in both e_2 and e_3. Then $g_1W_{\{b,c\}}$ and $g_2W_{\{b,c\}}$ are both in the chamber corresponding to x, and similarly $g_2W_{\{b,c\}}$ and $g_3W_{\{b,c\}}$ are both in the chamber corresponding to y. Without loss of generality, the chamber corresponding to x is a left coset of $W^{\{d\}}$ and the chamber corresponding to y is a left coset of $W^{\{a\}}$.

Let g be the vertex in $g_2W_{\{b,c\}}$ which is on the same side of the wall $g_2\Omega^b$ as $g_1W_{\{b,c\}}$, and the same side of the wall $g_2\Omega^c$ as $g_3W_{\{b,c\}}$. We will show that any geodesic in $\Gamma(G, S')$ between u and v passes through g.

Let p be a geodesic between u and v and let u' and v' be vertices in p such that $u' \in g_1W_{\{b,c\}}$ and $v' \in g_3W_{\{b,c\}}$. Then the part of p which lies between u' and v' passes through one of the vertices in $\{g, gc, gb, gcb\}$. Now,

$$d(u', g) = d(u', gc) = d(u', gb) - 1 = d(u', gcb) - 1$$

and

$$d(v', g) = d(v', gc) - 1 = d(v', gb) = d(v', gcb) - 1,$$

so

$$d(u', g)+d(g, v') = d(u', gb)+d(gb, v')-1 = d(u', gc)+d(gc, v')-1 = d(u', gbc)+d(gbc, v')-2.$$

Hence p passes through g. □

Lemma 7.28. *Coxeter systems of the form L_{2m+1} are is $\mathscr{T}$.*

Proof. Let $S = a, b, c, d$ and let W be given by the presentation

$$W = \langle a, b, c, d \mid a^2, b^2, c^2, d^2, (ab)^{2m+1}, (ac)^2, (bc)^2, (cd)^3, (bd)^3\rangle.$$

Then $(W, S) = L_{2m+1}$.

Now let $f : S \to 2\mathbb{Z}_{\geq 0} + 1$ be defined by

$$f(a) = 7,\ f(b) = 3,\ f(c) = 1 \text{ and } f(d) = 7.$$

We will show that $\Gamma^f(W, S)$ is $1000m$-almost geodetic.

Let d denote the path metric in $\Gamma^f(W, S)$. Let u, v be vertices in $\Gamma^f(W, S)$ such that $d(u, v) > 1000m$ and and let p_1, p_2 be geodesics between u anv v. Then we just need to show that p_1 and p_2 intersect at some point between u and v. Since L_{2m+1} splits as an amalgamated free product of $W_{\{b,c,d\}} = A_3$ and $W_{\{a,b,c\}} = I_{2m+1} \times \mathbb{Z}_2$ over the subgroup $W_{\{b,c\}}$, we can consider the Bass-Serre tree T of this amalgamated free product. We will slightly abuse terminology by referring to chambers in the graph $\Gamma^f(W, S)$, where we mean subgraphs of $\Gamma^f(W, S)$ which correspond to chambers in $\Gamma(W, S)$. The chambers of the cayley graph $\Gamma(W, S)$, for the case where $(2m + 1) = 5$ are shown below. The edges are highlighted if they are on a geodesic in $\Gamma^f(W, S)$ from the red vertex at the bottom left of each diagram.

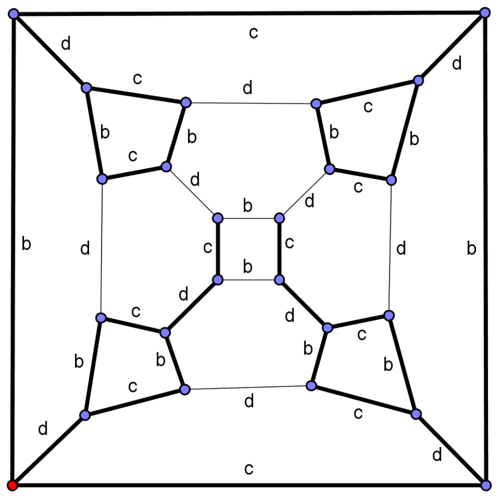

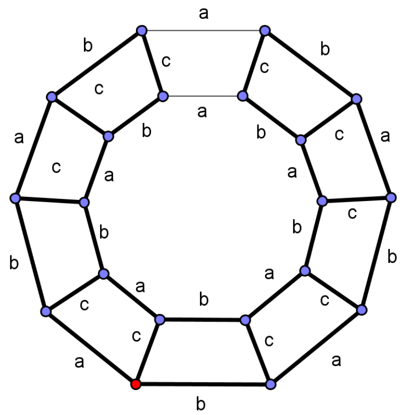

Let u' and v' be the vertices in T corresponding to chambers which contain u and v respectively, such that the distance between u' and v' in T is minimal. Since $d(u,v) > 1000m$, the distance between u' and v' in T is certainly greater than 6. Hence, there are vertices c_1, c_2, c_3, c_4, c_5 in T such that c_1, c_3, c_5 correspond to cosets of $W_{\{b,c,d\}} = A_3$ and c_2, c_4 correspond to cosets of $W_{\{a,b,c\}} = I_{2n+1} \times \mathbb{Z}_2$. Let w_1, x_1, y_1, z_1 be the vertices along p_1 in the cosets of $W_{\{b,c\}}$ corresponding to the edges (c_1, c_2), (c_2, c_3), (c_3, c_4) and (c_4, c_5) in T, such that $d(w_1, x_1)$ and $d(y_1, z_1)$ are minimal. Let w_2, x_2, y_2, z_2 be defined similarly along p_2.

Note that w_1, x_1, w_2, x_2 are all in the chamber corresponding to c_2, which corresponds to a coset of $I_{2m+1} \times \mathbb{Z}_2$. Also, w_1, w_2 are in the same left coset of $W_{\{b,c\}}$ as each other and x_1, x_2 are also in the same left coset of $W_{\{b,c\}}$ as each other. Moreover, since $d(x_1, w_1)$ is minimal, there is a geodesic path given by p_1 between w_1 and x_1 which does not intersect the coset $x_1 W_{\{b,c\}}$ at any point other than x_1, so x_1 must be in the closer of the two left cosets of $W_{\{c\}}$ in $W_{\{b,c\}}$ to $w_1 W_{\{b,c\}}$. Similarly x_2 must also be in this coset, so $x_2 = bx_1$ or $x_2 = x_1$. Similarly $y_2 = by_1$ or $y_2 = y_1$.

Now we just need to consider paths in the chamber L which contains x_1, y_1, x_2, y_2. Note that L corresponds to a coset of $W_{\{b,c,d\}} = A_3$, so there are six cosets of $W_{\{b,c\}}$ in L. Let H_x be the coset $x_1 W_{\{b,c\}} = x_2 W_{\{b,c\}}$ and let H'_x be the coset $x_1 W_{\{c\}} = x_2 W_{\{c\}}$. Let H_y and H'_y be defined similarly. If H_x and H_y are not opposite faces in L, then there is exactly one edge (of length 7) labelled by c which joins them. Hence p_1 and p_2 must both pass through this edge, so p_1 and p_2 intersect at a point between u and v.

If H_x and H_y are opposite faces in L, then H'_x and H'_y contain unique points x_3 and y_3 respectively such that $d(x_3, y_3) = 18$, which is minimal. We can easily see that the only geodesics between x_1 and y_1 pass through these points, so p_1 passes through x_3 and y_3. Similarly p_2 passes through x_3 and y_3, so p_1 and p_2 intersect at a point between u and v. □

Lemma 7.29. *Let (W_1, S_1) be a Coxeter system in $\mathscr{T}$. Let $m \in \mathbb{Z}_{>0}$ and let (W, S) be a $(2m+1)$-extension of (W_1, S_1). Then $(W, S) \in \mathscr{T}$.*

Proof. Let S'_1 be a finite generating set for W_1 which contains S_1 and let $f : S'_1 \to 2\mathbb{Z} + 1$ be a function such that $\Gamma^f(W_1, S'_1)$ is k-almost geodetic. Let s be the generator in S but not S_1 and let $t \in S$ be the generator such that (st) has order $2m+1$. Finally, let $S' = S_1' \cup \{s\}$. Define $f' : S' \to 2\mathbb{Z}+1$ by

$$f'(s') = f(s') \text{ for } s' \in S_1', \text{ and } f'(s) = f(t) + 2.$$

We will show that $\Gamma^{f'}(W, S')$ is $(k + (2m+1)(2f(t)+2))$-almost geodetic. Let u, v be vertices in $\Gamma^{f'}(W, S')$ and let l_1 and l_2 be non-intersecting geodesics between u and v. We will show that $d(u,v) \leq k + (2m+1)(2f(t)+2)$.

Note that W splits as an amalgamated free product of W_1 and $W_{\{s,t\}} = I_{2m+1}$ over the subgroup $W_{\{t\}}$. Hence we can consider the Bass-Serre tree T of the amalgamated free product. Note that the chambers in $\Gamma^{f'}(W, S')$ which correspond to cosets of $W_{\{s,t\}} = I_{2m+1}$ are cycles of size $(2m+1)(2f(t)+2)$. Let u' and v' be the vertices in T corresponding to chambers which contain u and v respectively, such that the distance between u' and v' in T is minimal. Since the geodesics l_1 and l_2 do not intersect, there cannot be any chambers corresponding to $W_{\{s,t\}} = I_{2m+1}$ which lie strictly in between u and v, so any vertices in T which lie strictly between u' and v' correspond to cosets of W. Hence there is at most one vertex x in T, corresponding to a coset of W which lies non-strictly between u' and v'. If there are no such vertices x, then u' and v' must be the same vertex and must correspond to a coset of $W_{\{s,t\}} = I_{2m+1}$, so u and v are in the same finite chamber and

$$d(u,v) \leq \frac{(2m+1)(2f(t)+1)}{2} < k + (2m+1)(2f(t)+2).$$

Hence we may assume that there is exactly one vertex x in T which corresponds to a coset of W_1 and which lies (non-strictly) between u' and v'. Let X be the chamber corresponding to x, so X is k-almost geodetic.

Now we will construct a vertex v_p in X, which satisfies $d(v, v_p) \leq (2m+1)(2f(t)+2)/2$, along with a pair of geodesics l_1' and l_2' between u and v_p which only intersect at u and v_p. If $v' = x$, then v is in X, so we can simply let $v_p = v$.

If $v' \neq x$, then let x_1 and x_2 be the vertices on l_1 and l_2 respectively, such that x_1 and x_2 are in the coset of $W_{\{t\}}$ which corresponds to the edge (v', x) in T. Since l_1 and l_2 do not intersect except at their endpoints, $x_1 \neq x_2$. Hence $x_1 = sx_2$, so the geodesic between x_1 and x_2 is the path p of length $f(t)$ which corresponds to an edge in $\Gamma(W, S')$ labelled by t. Now, by the triangle inequality, $|d(v, x_1) - d(v, x_2)| \leq f(t)$. Moreover, since the chamber containing v, x_1 and x_2 is a cycle of even length, $d(v, x_1) - d(v, x_2)$ and $f(t)$ have the same parity (in particular, they are both odd). Therefore, there is some vertex v_p on the path p such that $d(v, x_1) - d(v, x_2) = d(v_p, x_1) - d(v_p, x_2)$. Now for $i = 1, 2$ we construct the path l_i' between u and v_p by joining the part of the l_i which lies between u and x_i to the geodesic between x_i and v_p. Then since any path between u and v_p passes through one of the two vertices x_1 and x_2, at least one of the paths l_1', l_2' must be a geodesic. But

$$(u, x_2) - d(u, x_1) = d(v, x_1) - d(v, x_2) = d(v_p, x_1) - d(v_p, x_2),$$

so l_1' and l_2' have the same length. Therefore, both l_1' and l_2' are geodesics, so we have constructed a pair of geodesics between u and v_p, which only intersect at u

and v_p. Finally, since v and v_p are in the same cycle of size $(2m+1)(2f(t)+2)$ we must have

$$d(v, v_p) \leq \frac{(2m+1)(2f(t)+2)}{2}.$$

In the same way we construct a vertex u_p in X, which satisfies $d(u, u_p) \leq (2m+1)(2f(t)+2)/2$, along with a pair of geodesics l_1'' and l_2'' in $\Gamma^{f'}(W, S')$ between u_p and v_p, which only intersect at u_p and v_p.

Since l_1'' and l_2'' are geodesics in X which only intersect at their endpoints u_p and v_p, we must have $d(u_p, v_p) \leq k$. Therefore,

$$d(u, v) \leq d(u, u_p) + d(u_p, v_p) + d(v, v_p) \leq k + (2m+1)(2f(t)+2).$$

Hence, $\Gamma^{f'}(W, S')$ is $(k + (2m+1)(2f(t)+2))$-almost geodetic. □

Lemma 7.30. *Let (W_1, S_1) and (W_2, S_2) be two Coxeter systems in $\mathscr{T}$. Then the Coxeter system $(W_1 * W_2, S_1 \cup S_2) \in \mathscr{T}$.*

Proof. Let S_1' be a generating set for W_1 which contains S_1 and let $f_1 : S_1' \to 2\mathbb{Z}+1$ be a function such that $\Gamma^{f_1}(W_1, S_1')$ is almost geodetic. Let f_2 and S_2 be defined similarly. Now define $f : S_1' \cup S_2' \to 2\mathbb{Z}+1$ by $f(s) = f_i(s)$ for $s \in S_i'$. Then every block in $\Gamma(W_1 * W_2, S_1' \cup S_2')$ is isomorphic as a labelled graph to a block in $\Gamma(W_1, S_1')$ or $\Gamma(W_2, S_2')$. Hence, every block in $\Gamma^f(W_1 * W_2, S_1' \cup S_2')$ is isomorphic to a block in $\Gamma^{f_1}(W_1, S_1')$ or $\Gamma^{f_2}(W_2, S_2')$. Hence, $\Gamma^f(W_1 * W_2, S_1' \cup S_2')$ is almost geodetic, so $(W_1 * W_2, S_1 \cup S_2) \in \mathscr{T}$. □

Theorem 7.31. *((a)⇒(b)) If (W, S) is a Coxeter system in $\mathscr{X}$, then W is virtually free and there exists a generating set S' for W which contains S and a function $f : S' \to 2\mathbb{Z}_{\geq 0}+1$ such that $\Gamma^f(W, S')$ is almost geodetic.*

Proof. If (W, S) is a Coxeter system in $\mathscr{F}$, then W is finite, so it is certainly virtually free and almost geodetic. If (W, S) is either $L_{2m+1,2n+1}$ or L_{2m+1} for positive integers m, n, then W decomposes as an amalgamated free product of two finite groups, so W is virtually free. If (W, S) is a $2m+1$-extension of a Coxeter system (W', S'), then W is an amalgamated free product of W' with a finite group over a finite subgroup, so if W' is virtually free, then W is also virtually free. Hence, if $(W, S) \in \mathscr{X}_1$, then W is virtually free. Finally, the free product of finitely many virtually free group is virtually free, so if $(W, S) \in \mathscr{X}$, then (W, S) is virtually free.

Finally, from lemmas 7.27, 7.28, 7.29 and 7.30, any group $(W, S) \in \mathscr{X}$ satisfies $(W, S) \in \mathscr{T}$. So if $(W, S) \in \mathscr{X}$, then there exists a generating set S' for W which contains S and a function $f : S' \to 2\mathbb{Z}_{\geq 0}+1$ such that $\Gamma^f(W, S')$ is almost geodetic. □

7.4 Coxeter groups where all walls are finite are in $\mathscr{X}$

Lemma 7.32. *Let (W, S) be a Coxeter system such that all walls of $\Gamma(W, S)$ are finite. If $T \subset S$, then all walls of $\Gamma(W_T, T)$ are finite.*

Proof. Since $T \subset S$ and W_T is the subgroup of W generated by T, the graph $\Gamma(W_T, T)$ is a subgraph of $\Gamma(W, S)$. Let $f : \Gamma(W_T, T) \to \Gamma(W, S)$ be this embedding. Then if Ω^r is a wall in $\Gamma(W_T, T)$, the image $f(\Omega^r)$ is contained in a wall $f(\Omega^{r'})$ in $\Gamma(W, S)$. Since $f(\Omega^{r'})$ is finite, $f(\Omega^r)$ must also be finite. □

Definition 7.33. Let (W, S) be a Coxeter system. We define the parity graph P_W of (W, S) as follows:

- The vertex set of P is equal to S.
- for each pair $s, t \in S$ with $s \neq t$, we join s and t with a red edge if the order of st is even. We join s and t with a blue edge if the element st has odd order. Finally, if st has infinite order, then s and t are not joined by an edge.

Note that we can obtain this graph from the Coxeter diagram of (W, S) by replacing each edge with an even label by a red edge, replacing each edge with an odd label with a blue edge and, finally, removing each edge labelled by ∞.

Lemma 7.34. *Let (W, S) be a Coxeter system such that all walls of $\Gamma(W, S)$ are finite. If P_W is complete, then W is finite.*

Proof. Let L be the nerve of (W, S). Since all walls of $\Gamma(W, S)$ are finite, W is not one-ended. Therefore, either W is finite or there is some closed simplex in L whose removal will disconnect L.

Since P_W is complete, for every pair of generators $s, t \in S$ the element $st \in W$ has finite order, so $W_{\{s,t\}}$ is a finite dihedral group. Therefore, the 1-skeleton of L is a complete graph. Therefore, if we remove any subset of the vertices of L and all simplices attached to those vertices, the 1-skeleton of the new simplicial complex will still be complete. Therefore W must be finite. □

Lemma 7.35. *Let (W, S) be a Coxeter system such that all walls of $\Gamma(W, S)$ are finite. If c is a simple cycle in P_W of length at least 4 then there are two vertices s_1, s_2 in c which are not adjacent in c but which are adjacent in P_W. In graph theory terminology, the graph P_W is chordal.*

Proof. Suppose that c is a simple cycle in P_W of length at least 4 and any two vertices s_1, s_2 in c which are not adjacent in c are not adjacent in P_W. Let $T \subset S$ be the set of generators which correspond to vertices in the cycle c. Let L be the nerve of (W_T, T). Then the 1-skeleton of L is a cycle of the same length as c. Since this contains no triangles, L must not contain any simplices with dimension greater than 1. So L is a simple cycle. Therefore, if we remove any closed simplex from L, it remains connected, so W_T is one-ended.

But by lemma 7.32, all walls of W_T are finite, so this contradicts lemma 7.13. □

Lemma 7.36. *Let (W, S) be a Coxeter system such that all walls of $\Gamma(W, S)$ are finite. Then there is no simple cycle in P_W which is entirely blue.*

Proof. Suppose the contrary, that there is a cycle c in P_w which is entirely blue.

We will assume that the vertices in c are $s_1, s_2, \ldots, s_k$ in that order. Then $m_{i,i+1}$ is odd for $1 \leq i \leq k-1$ and m_{1k} is also odd. Note that the subgroup of W generated by $s_1, \ldots, s_k$ cannot be a finite group since the Coxeter diagram contains a cycle of edges with labels greater than 2. Therefore, by lemma 7.34, the subgraph of P_W defined by the vertices in c is not complete. So there exist $1 \leq i < j \leq k$ with $m_{ij} = \infty$. Without loss of generality we can assume that $i = 1$. Then $j \neq 2$ or k, so $k \geq 4$.

Let $h = c_{k,1}c_{1,2}c_{2,3}\ldots c_{k-2,k-1}c_{k-1,k}$. We claim that h is in the centraliser of s_k and h has infinite order. Since

$$\begin{aligned}
s_k h &= s_k c_{k,1}c_{1,2}c_{2,3}\ldots c_{k-2,k-1}c_{k-1,k} \\
&= c_{k,1}s_1 c_{1,2}c_{2,3}\ldots c_{k-2,k-1}c_{k-1,k} \\
&= c_{k,1}c_{1,2}s_2 c_{2,3}\ldots c_{k-2,k-1}c_{k-1,k} \\
&= c_{k,1}c_{1,2}c_{2,3}\ldots c_{k-2,k-1}s_{k-1}c_{k-1,k} \\
&= c_{k,1}c_{1,2}c_{2,3}\ldots c_{k-2,k-1}c_{k-1,k}s_k \\
&= hs_k,
\end{aligned}$$

h is in the centraliser of s_k. By proposition 7.15, $c_{k,1}c_{1,2} \notin W^{\{s_1\}}$. Therefore, $c_{k,1}c_{1,2} \in W^{\{s_j\}} \setminus W^{\{s_1,s_j\}}$. Similarly $c_{j-1,j}c_{j,j+1} \in W^{\{s_1\}} \setminus W^{\{s_1,s_j\}}$. Now, since $c_{i,i+1} \in W^{\{s_1,s_j\}}$, if $1 < i < k$ and $i \notin \{j-1, j\}$, then we have

$$c_{2,3}\ldots c_{k-2,k-1}c_{k-1,k} \in W^{\{s_1\}} \setminus W^{\{s_1,s_j\}}.$$

Finally, since the vertices $s_1$and s_j are joined by an edge labelled ∞ in the Coxeter diagram of (W,S), we know that W decomposes as an amalgamated free product $W = W^{\{s_1\}} *_{W^{\{s_1,s_j\}}} W^{\{s_j\}}$. Therefore, by proposition 1.13, we know that $h = (c_{k,1}c_{1,2})(c_{2,3}\ldots c_{k-2,k-1}c_{k-1,k})$ has infinite order. Therefore the centraliser of s_k is infinite, so the wall Ω^{s_k} is infinite. A contradiction. □

Lemma 7.37. *Let (W,S) be a Coxeter system such that all walls of $\Gamma(W,S)$ are finite. Suppose that we have a vertex path $s_{a_1}, s_{a_2}, \ldots, s_{a_k}$ in P_W, such that the two edges at the ends of the path are both red while the other edges are all blue. Then the element*

$$h = c_{a_1,a_2}c_{a_2,a_3}\ldots c_{a_{k-2},a_{k-1}}c_{a_{k-1},a_k}c_{a_{k-2},a_{k-1}}\ldots c_{a_2,a_3}$$

has finite order in W. Note that this is not necessarily a simple path.

Proof. From the colouring of the edges in P_W, we know that $m_{a_i,a_{i+1}}$ is even for $i = 1, k-1$ and odd for $1 < i < k-1$. Now, since

$$\begin{aligned}
s_{a_2}h &= s_{a_2}c_{a_1,a_2}c_{a_2,a_3}\ldots c_{a_{k-2},a_{k-1}}c_{a_{k-1},a_k}c_{a_{k-2},a_{k-1}}\ldots c_{a_2,a_3} \\
&= c_{a_1,a_2}s_{a_2}c_{a_2,a_3}\ldots c_{a_{k-2},a_{k-1}}c_{a_{k-1},a_k}c_{a_{k-2},a_{k-1}}\ldots c_{a_2,a_3} \\
&= c_{a_1,a_2}c_{a_2,a_3}s_{a_3}\ldots c_{a_{k-2},a_{k-1}}c_{a_{k-1},a_k}c_{a_{k-2},a_{k-1}}\ldots c_{a_2,a_3} \\
&= c_{a_1,a_2}c_{a_2,a_3}\ldots c_{a_{k-2},a_{k-1}}s_{a_{k-1}}c_{a_{k-1},a_k}c_{a_{k-2},a_{k-1}}\ldots c_{a_2,a_3} \\
&= c_{a_1,a_2}c_{a_2,a_3}\ldots c_{a_{k-2},a_{k-1}}c_{a_{k-1},a_k}s_{a_{k-1}}c_{a_{k-2},a_{k-1}}\ldots c_{a_2,a_3} \\
&= c_{a_1,a_2}c_{a_2,a_3}\ldots c_{a_{k-2},a_{k-1}}c_{a_{k-1},a_k}c_{a_{k-2},a_{k-1}}\ldots s_{a_3}c_{a_2,a_3} \\
&= c_{a_1,a_2}c_{a_2,a_3}\ldots c_{a_{k-2},a_{k-1}}c_{a_{k-1},a_k}c_{a_{k-2},a_{k-1}}\ldots c_{a_2,a_3}s_{a_2} \\
&= hs_{a_2},
\end{aligned}$$

h is in the centraliser of s_{a_2}. Also, since the wall $\Omega^{s_{a_2}}$ is finite, we know from proposition 7.10 that the centraliser of s_{a_2} is finite. Since h is an element of this centraliser, $\langle h \rangle$ is a subset of this centraliser. Therefore, $\langle h \rangle$ is finite. So h has finite order in W. $\square$

Lemma 7.38. *Let (W, S) be a Coxeter system such that all walls of $\Gamma(W, S)$ are finite. If p is a simple edge path in P_W of length at least 2 between vertices s and t such that the two edges on the ends are red and all of the other edges are blue, then there is an edge between s and t.*

Proof. We will assume that the vertices in the path are $s_1, s_2, \ldots, s_k$ where $s_1 = s$ and $s_k = t$. Since the path contains at least two edges, it contains at least three vertices, so $k \geq 3$. From lemma 7.37 we know that the element $h = c_{1,2}c_{2,3} \ldots c_{k-2,k-1}c_{k-1,k}c_{k-2,k-1} \ldots c_{2,3}$ has finite order in W.

Now suppose for the sake of contradiction that there is no edge between $s = s_1$ and $t = s_k$.

From proposition 7.7 we know that $W^{\{s_1,s_k\}} \cap W_{\{s_1,s_2\}} = W_{\{s_2\}} = \{1, s_2\}$, so $c_{1,2} \notin W^{\{s_1,s_k\}}$. Therefore, $c_{1,2} \in W^{\{s_k\}} \setminus W^{\{s_1,s_k\}}$. Similarly $c_{k-1,k} \in W^{\{s_1\}} \setminus W^{\{s_1,s_k\}}$. Now, since $c_{i,i+1} \in W^{\{s_1,s_k\}}$ for $1 < i < k-1$, we have

$$c_{2,3} \ldots c_{k-2,k-1}c_{k-1,k}c_{k-2,k-1} \ldots c_{2,3} \in W^{\{s_1\}} \setminus W^{\{s_1,s_k\}}.$$

Finally, since the vertices $s_1$and s_k are joined by an edge labelled ∞ in the Coxeter diagram of (W, S), we know that W decomposes as an amalgamated free product $W = W^{\{s_1\}} *_{W^{\{s_1,s_k\}}} W^{\{s_k\}}$. Therefore, by proposition 1.13, we know that $h = (c_{1,2})(c_{2,3} \ldots c_{k-2,k-1}c_{k-1,k}c_{k-2,k-1} \ldots c_{2,3})$ has infinite order. A contradiction. $\square$

Lemma 7.39. *Let (W, S) be a Coxeter system such that all walls of $\Gamma(W, S)$ are finite. Let $s_1 s_2 \ldots s_k s_1$ be a cycle in P_W such that the edges (s_1, s_2) and (s_1, s_k) are red and all other edges in the cycle are blue. Then every vertex in the cycle is joined to s_1 by an edge.*

Proof. From lemma 7.37 we know that the element $h = c_{1,2}c_{2,3} \ldots c_{k-1,k}c_{k,1}c_{k-1,k} \ldots c_{2,3}$ has finite order in W.

Now suppose for the sake of contradiction that there is some j with $2 < j < k$ such that s_1 and s_j are not joined by an edge.

From theorem 7.7, we know that $W_{\{s_1,s_2\}} \cap W^{\{s_1\}} = W_{\{s_2\}}$. Clearly $c_{1,2} \notin W_{\{s_2\}}$, so $c_{1,2} \notin W^{\{s_1\}}$. Therefore, $c_{1,2} \in W^{\{s_j\}} \setminus W^{\{s_1,s_j\}}$. Similarly $c_{k,1} \in W^{\{s_j\}} \setminus W^{\{s_1,s_j\}}$. By proposition 7.15, $c_{j-1,j}c_{j,j+1} \notin W^{\{s_j\}}$. Therefore, $c_{j-1,j}c_{j,j+1} \in W^{\{s_1\}} \setminus W^{\{s_1,s_j\}}$. Therefore,

$$c_{2,3} \ldots c_{k-1,k} \in W^{\{s_1\}} \setminus W^{\{s_1,s_j\}}.$$

Finally, since the vertices s_1 and s_j are joined by an edge labelled ∞ in the Coxeter diagram of (W, S), we know that W decomposes as an amalgamated free product $W = W^{\{s_1\}} *_{W^{\{s_1,s_j\}}} W^{\{s_j\}}$. Therefore, by proposition 1.13, we know that $h = (c_{1,2})(c_{2,3} \ldots c_{k-1,k})(c_{k,1})(c_{2,3} \ldots c_{k-1,k})^{-1}$ has infinite order. A contradiction. $\square$

Definition 7.40. We will call a Coxeter system (W, S) reduced if the parity graph P_W of (W, S) is connected, and there is no vertex with degree 1 which is only attached to a blue edge.

Lemma 7.41. *If (W, S) is a reduced Coxeter system such that all walls of $\Gamma(W, S)$ are finite then P_W is 2-connected.*

Proof. Since (W, S) is reduced, P_W is connected.

Suppose that P_W is not 2-connected. Then there is some vertex v in P_W whose removal disconnects P_W. Let B_1' and B_2' be 2-connected components of $P_W \setminus \{v\}$. Let $B_1 = B_1' \cup \{v\}$ and let $B_2 = B_{2'} \cup \{v\}$. Call a vertex in B_1 good if it connects to a red edge in B_1.

We construct a simple path starting at v by travelling along blue edges until we reach a good vertex, a vertex that we have been to before, or a vertex that is only connected to the previous vertex in the path. If we reach a vertex u which is only connected to the previous vertex in the path, then u has degree 1 and the edge it is connected to is blue, contradicting the fact that (W, S) is reduced. If we reach a vertex that we have been to before, it implies that there is a blue cycle, which contradicts lemma 7.36. Therefore, the path must take the form $u_0, u_1, \ldots, u_k$ where $u_0 = v$ and, since we stopped at u_k and not earlier, the only good vertex in the sequence $u_0, \ldots, u_k$ is u_k.

Let u_{k+1} be a vertex in B_1 which is attached to u_k by a red edge. Then u_{k+1} is good, so it is not already in the path. Therefore, $u_0, u_1, \ldots, u_{k+1}$ is a simple path in B_1 such that the last edge is red and all other edges are blue. Similarly we can construct a simple path $v_0, v_1, \ldots, v_m, v_{m+1}$ in B_2 with $v_0 = v$ such that the last edge is red and all other edges are blue. Then attaching these two paths together at v, we get the path

$$v_{m+1}, v_m, \ldots, v_1, v_0 = u_0, u_1, \ldots, u_k, u_{k+1},$$

which has red edges at its ends but all other edges are blue. Therefore, by lemma 7.38, the vertices u_{m+1} and v_{m+1} are joined by an edge. But $u_{k+1} \in B_1'$ and $v_{m+1} \in B_2'$, so this contradicts the fact that B_1' and B_2' are disconnected components of $P_W \setminus \{v\}$. $\square$

Lemma 7.42. *If (W, S) is a reduced Coxeter system such that all walls of $\Gamma(W, S)$ are finite and $|S| \leq 3$ then P_W is complete.*

Proof. From lemma 7.41, we know that P_W is 2-connected. Therefore, since P_W has at most three vertices, it must be complete. $\square$

Lemma 7.43. *If (W, S) is a reduced Coxeter system with $|S| \geq 3$ and $s \in S$, then s is part of some cycle of length 3 in P_W.*

Proof. Since P_W is 2-connected and contains more than two vertices, s must be part of some simple cycle in P_W. Let c be the cycle of minimum length. If c has length at least 4, then by lemma 7.35 there are two vertices u and v in c which are adjacent in P_W but not in c. Then s is on (at least) one of the two arcs l of c between u and v. Hence, the union of l with the edge between u and v is a simple cycle containing s which is shorter than c, contradicting the minimality of c. Therefore, c has length 3. $\square$

Lemma 7.44. *If (W, S) is a reduced Coxeter system and $|S| \geq 4$, then there is some $s \in S$ such that $(W^{\{s\}}, S \setminus \{s\})$ is reduced.*

Proof. Let $t \in S$. If $(W^{\{t\}}, S \setminus \{t\})$ is not reduced, then the graph $P_W \setminus \{t\}$ is either disconnected, or contains a vertex with degree 1. Since P_W is 2-connected, $P_W \setminus \{t\}$ must be connected, so there is some vertex s which has degree 1 in $P_W \setminus \{t\}$. But s must have degree at least 2 in P_W, since P_W is 2-connected. So s has degree exactly 2 in P_W.

By lemma 7.43 there is a cycle of length 3 which contains s. So the two vertices u and v which are attached to s are joined by an edge. If either u or v has degree 2 then removing the other one of these vertices will disconnect the graph. So both u and v have degree greater than 2. Therefore, every vertex in $P_W \setminus \{s\}$ has degree at least 2. Therefore, $(W^{\{s\}}, S \setminus \{s\})$ is reduced. □

Lemma 7.45. *If (W, S) is a reduced Coxeter system and $|S| \geq 3$, then the number of edges in P_W is at least $2|S| - 3$.*

Proof. We will proceed by induction. For $|S| = 3$, by lemma 7.42, the graph P_W is complete, so it has $3 = 2|S| - 3$ edges.

For $|S| \geq 3$, we will assume that the result holds for all reduced Coxeter systems (W', S') where $|S'| < |S|$.

From lemma 7.44, we know that there is some $s \in S$ such that $(W^{\{s\}}, S \setminus \{s\})$ is reduced. Therefore, by the inductive hypothesis, the number of edges in $P_{W^{\{s\}}}$ is at least $2|S \setminus \{s\}| - 3 = 2|S| - 5$. But P_W is 2-connected, so s has degree at least 2. Therefore, the number of edges in P_W is at least $2|S| - 5 + 2 = 2|S| - 3$. □

Lemma 7.46. *If (W, S) is a reduced Coxeter system such that all walls of $\Gamma(W, S)$ are finite and $|S| = 4$ then P_W is complete or P_W is isomorphic to a graph with vertices $\{a, b, c, d\}$ such that all pairs of vertices except a and d are joined by an edge and the edges (a, b) and (b, c) are both red, while the edges (a, c) and (b, d) are both blue. Note that the edge (c, d) may be either red or blue.*

Proof. Let the vertices of P_W be $\{a, b, c, d\}$. If P_W is not complete, then there is some pair of vertices, say a, d, which are not joined by an edge. From lemma 7.45, P_W contains at least 5 edges. Therefore P_W contains exactly 5 edges, and every pair of vertices apart from a, d form an edge. Therefore, a, c, d, b, a is a cycle in P_W. So by lemma 7.36, at least one of the edges in the cycle a, c, d, b, a must be red. So, without loss of generality, assume the edge (a, b) is red. Since a and d are not joined by an edge, by lemma 7.38 the edge (b, d) must not be red, so it is blue. Similarly, at least one of the edges (a, c) and (c, d) must be blue.

If the edge (a, c) is red, then (c, d) must be blue. But then the cycle a, b, d, c, a contradicts lemma 7.39. So (a, c) must be blue.

If the edge (b, c) is blue, then (c, d) must be red, or we would have a blue triangle. But then the path a, b, c, d contradicts lemma 7.38. So (b, c) is red. □

Lemma 7.47. *If (W, S) is Coxeter system such that all walls of $\Gamma(W, S)$ are finite and P_W is isomorphic to a graph of the form described in the previous lemma, then (W, S) is either L_{2m+1} or $L_{2m+1,2n+1}$ for some $m, n \in \mathbb{Z}_{>0}$.*

Proof. Let the vertices in the graph P_W be s_1, s_2, s_3, s_4 such that all pairs of vertices except s_1, s_4 are joined by an edge. Moreover, assume that the edges (s_1, s_2) and (s_2, s_3) are both red, and the edges (s_1, s_3) and (s_2, s_4) are both blue. We will show that if the edge (s_3, s_4) is blue, then $(W, S) = L_{2m+1}$ for some $m \in \mathbb{Z}_{>0}$ and if (s_3, s_4) is red, then $(W, S) = L_{2m+1,2n+1}$ for some $m, n \in \mathbb{Z}_{>0}$.

Case 1: The edge (s_3, s_4) is blue.

We will first prove that $m_{1,2}$ and $m_{2,3}$ are both equal to 2. By lemma 7.34, the group $W^{\{s_4\}}$ is finite. Hence we must have $m_{1,2}, m_{2,3} < 6$. Also, at least one of $m_{1,2}$ and $m_{2,3}$ is equal to 2. Therefore, if either $m_{1,2} > 2$ or $m_{2,3} > 2$, the one which is not equal to 2 must be equal to 4. Moreover, in this case we must have $m_{1,3} = 3$. Therefore, $(m_{1,2}, m_{1,3}, m_{2,3})$ is either $(4, 3, 2)$, $(2, 3, 4)$ or $(2, 2n+1, 2)$ for some $n \in \mathbb{Z}_{>0}$.

Note that $s_2, s_1, s_3, s_4, s_2, s_1$ is a path with red edges at its ends and blue edges everywhere else, so by lemma 7.37, the element

$$h = c_{2,1}c_{1,3}c_{3,4}c_{4,2}c_{2,1}c_{4,2}c_{3,4}c_{1,3}$$

has finite order. Hence, $c_{1,3}c_{2,1}c_{1,3}c_{3,4}c_{4,2}c_{2,1}c_{4,2}c_{3,4}$ also has finite order.

Now, by proposition 7.15, we have $c_{3,4}c_{4,2}, c_{4,2}c_{3,4} \notin W^{\{s_4\}}$ and by proposition 7.14, we have $c_{2,1} \notin W^{\{s_1\}}$. Hence, if $c_{1,3}c_{2,1}c_{1,3} \notin W^{\{s_1\}}$, then by proposition 1.13, the element $c_{1,3}c_{2,1}c_{1,3}c_{3,4}c_{4,2}c_{2,1}c_{4,2}c_{3,4}$ has infinite order in W, a contradiction. So $c_{1,3}c_{2,1}c_{1,3} \in W^{\{s_1\}}$. Therefore,

$$c_{1,3}c_{2,1}c_{1,3} \in W^{\{s_1\}} \cap W_{\{s_1,s_2,s_3\}} = W_{\{s_2,s_3\}}.$$

Now we can just check the Cayley graph of C_3 to see that this is not the case when $(m_{1,2}, m_{1,3}, m_{2,3})$ is either $(4, 3, 2)$ or $(2, 3, 4)$. Therefore, $m_{1,2} = m_{2,3} = 2$.

We will now prove that $m_{2,4}$ and $m_{3,4}$ are both equal to 3. By lemma 7.34, the group $W^{\{s_1\}}$ is finite, hence we must have $m_{2,4}, m_{3,4} < 6$. Also, $m_{2,4}$ and $m_{3,4}$ cannot both be equal to 5. Therefore, $(m_{2,4}, m_{3,4})$ is either $(3, 3)$, $(3, 5)$ or $(5, 3)$.

Note that s_1, s_2, s_4, s_3, s_2 is a path with red edges at its ends and blue edges everywhere else, so by lemma 7.37, the element

$$h = c_{1,2}c_{2,4}c_{4,3}c_{3,2}c_{4,3}c_{2,4}$$

has finite order.

Now, by proposition 7.14, we have $c_{1,2} \notin W^{\{s_1\}}$ hence, if $c_{2,4}c_{4,3}c_{3,2}c_{4,3}c_{2,4} \notin W^{\{s_4\}}$, then by proposition 1.13, the element $c_{1,2}c_{2,4}c_{4,3}c_{3,2}c_{4,3}c_{2,4}$ has infinite order in W, a contradiction. So $c_{2,4}c_{4,3}c_{3,2}c_{4,3}c_{2,4} \in W^{\{s_4\}}$. Therefore,

$$c_{2,4}c_{4,3}c_{3,2}c_{4,3}c_{2,4} \in W^{\{s_4\}} \cap W^{\{s_1\}} = W_{\{s_2,s_3\}}.$$

Now we can just check the Cayley graph of H_3 to see that this is not the case when $(m_{2,3}, m_{2,4}, m_{3,4})$ is either $(2, 3, 5)$ or $(2, 5, 3)$. Therefore, $m_{2,4} = m_{3,4} = 3$. Hence, $(W, S) = L_{2n+1}$, where $2n + 1 = m_{1,3}$.

Case 2: The edge (s_3, s_4) is red.

As in Case 1, $(m_{1,2}, m_{1,3}, m_{2,3})$ is either $(4, 3, 2)$, $(2, 3, 4)$ or $(2, 2n + 1, 2)$.

Note that s_4, s_2, s_1, s_3 is a path with red edges at its ends and blue edges everywhere else, so by lemma 7.37, the element

$$h = c_{4,2}c_{2,1}c_{1,3}c_{2,1}$$

has finite order. Now, by proposition 7.14, we have $c_{4,2} \notin W^{\{s_4\}}$ hence, if $c_{2,1}c_{1,3}c_{2,1} \notin W^{\{s_1\}}$, then by proposition 1.13, the element $c_{4,2}c_{2,1}c_{1,3}c_{2,1}$ has infinite order in W, a contradiction. So $c_{2,1}c_{1,3}c_{2,1} \in W^{\{s_1\}}$. Therefore,

$$c_{2,1}c_{1,3}c_{2,1} \in W^{\{s_1\}} \cap W_{\{s_1,s_2,s_3\}} = W_{\{s_2,s_3\}}.$$

Now we can just check the Cayley graph of C_3 to see that this is not the case when $(m_{1,2}, m_{1,3}, m_{2,3})$ is either $(4, 3, 2)$ or $(2, 3, 4)$. Therefore, $m_{1,2} = m_{2,3} = 2$. Similarly, $m_{3,4} = 2$. Hence, $(W, S) = L_{2m+1,2n+1}$, where $2m + 1 = m_{1,3}$ and $2n + 1 = m_{2,4}$. □

Lemma 7.48. *If* (W, S) *is a reduced Coxeter system such that all walls of* $\Gamma(W, S)$ *are finite and* $|S| = 5$ *then* P_W *is complete.*

Proof. Let $S = \{s_1, s_2, s_3, s_4, s_5\}$. Suppose for the sake of contradiction that P_W is not complete. The structure of the proof is shown in the diagrams below. Grey edges mean there may or may not be an edge and black edges denote edges of undetermined colour.

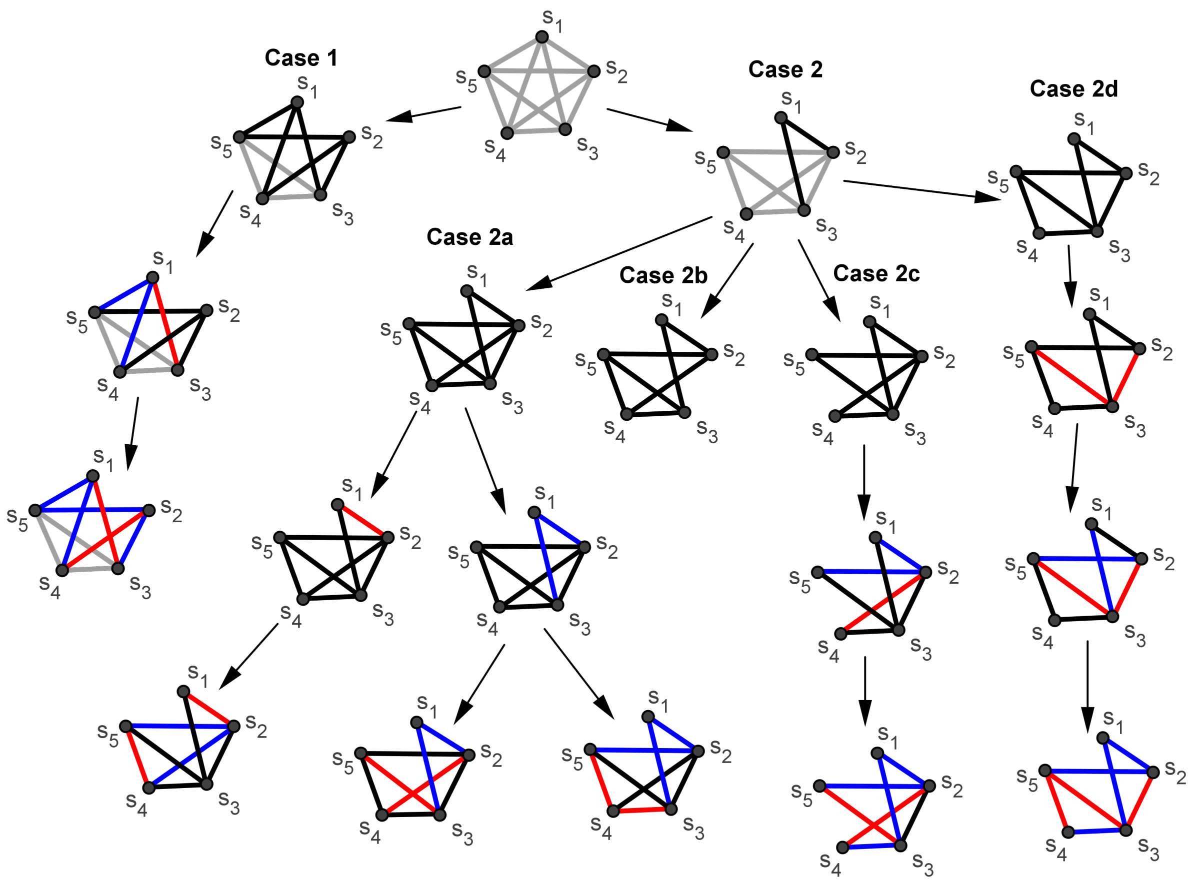

We will consider two cases:

- Case 1: Each vertex in P_W has degree at least three.
- Case 2: There is a vertex in P_W with degree (at most) two.

Case 1:

Since P_W is not complete, there is a pair of vertices which are not joined by an edge. Without loss of generality they are s_1 and s_2. Since s_1 and s_2 have degree at least 3, the pairs (s_1, s_3), (s_1, s_4), (s_1, s_5), (s_2, s_3), (s_2, s_4) and (s_2, s_5) must all be edges in P_W. Also, since s_1, s_3, s_2, s_4, s_1 is a cycle, the Coxeter system $(W^{\{s_5\}}, \{s_1, s_2, s_3, s_4\})$ is reduced. Moreover, since s_1 and s_2 are not joined by an edge, the graph $P_{W^{\{s_5\}}}$ is not complete, so by lemma 7.46, each vertex in $P_{W^{\{s_5\}}}$ is attached to at least one blue edge. Therefore, the edges (s_1, s_3) and (s_1, s_4) are not both red. By the same argument, no pair of edges in P_W involving s_1 are red. So s_1 is attached to at most one red edge in P_W. Similarly, s_2 is attached to at most one red edge in P_W.

Now, since s_1, s_3, s_2, s_4, s_1 forms a cycle in P_W, at least one of the edges in this cycle is red. Without loss of generality, let the edge (s_1, s_3) be red. Then (s_1, s_4) and (s_1, s_5) are both blue. Similarly, at least one of the edges in the cycle s_1, s_4, s_2, s_5, s_1 is red. So either (s_2, s_4) or (s_2, s_5) is red. Without loss of generality we may assume that (s_2, s_4) is red. Then (s_2, s_3) and (s_2, s_5) are both blue. Therefore, $s_1, s_3, s_2, s_5, s_1, s_4, s_2$ is a path with red edges at its ends and blue edges everywhere else. So by lemma 7.37, the element

$$h = c_{1,3}c_{3,2}c_{2,5}c_{5,1}c_{1,4}c_{4,2}c_{1,4}c_{5,1}c_{2,5}c_{3,2}$$

has finite order.

Now, by proposition 7.15, we have $c_{3,2}c_{2,5}, c_{2,5}c_{3,2} \notin W^{\{s_2\}}$ and $c_{5,1}c_{1,4}, c_{1,4}c_{5,1} \notin W^{\{s_1\}}$. Also by proposition 7.14, we have $c_{4,2} \notin W^{\{s_2\}}$ and $c_{1,3} \notin W^{\{s_1\}}$. Therefore,

$$c_{3,2}c_{2,5}, c_{4,2}, c_{2,5}c_{3,2} \in W^{\{s_1\}} \setminus W^{\{s_1,s_2\}}$$

and

$$c_{1,3}, c_{5,1}c_{1,4}, c_{1,4}c_{5,1} \in W^{\{s_2\}} \setminus W^{\{s_1,s_2\}}.$$

Finally, since $m_{1,2} = \infty$, we have

$$W = W^{\{s_1\}} *_{W^{\{s_1,s_2\}}} W^{\{s_2\}}.$$

Therefore, by proposition 1.13, the element $c_{1,3}c_{3,2}c_{2,5}c_{5,1}c_{1,4}c_{4,2}c_{1,4}c_{5,1}c_{2,5}c_{3,2}$ has infinite order in W, a contradiction.

Case 2: There is a vertex in P_W with degree (at most) two.

Without loss of generality we will assume that s_1 has degree at most two. Note that since P_W is 2-connected, s_1 has degree at least two, so it has degree exactly two. Without loss of generality s_1 is joined to s_2 and s_3, but not s_4 or s_5. From lemma 7.45, the graph P_W contains at least 7 edges, so there is at most one pair of vertices in $\{s_2, s_3, s_4, s_5\}$ which are not joined by an edge. The pairs (s_2, s_4), (s_2, s_5), (s_3, s_4) and (s_3, s_5) are all equivalent, so without loss of generality, we can assume that if there is a pair of vertices in $\{s_2, s_3, s_4, s_5\}$ which are not joined by an edge, the pair is either (s_2, s_3), (s_4, s_5) or (s_2, s_4). So we will break this into four cases:

- Case 2a: Every pair of vertices in $\{s_2, s_3, s_4, s_5\}$ are joined by an edge.

- Case 2b: s_2 and s_3 are not joined by an edge.
- Case 2c: s_4 and s_5 are not joined by an edge.
- Case 2d: s_2 and s_4 are not joined by an edge.

Case 2a: Every pair of vertices in $\{s_2, s_3, s_4, s_5\}$ are joined by an edge.

If (s_1, s_2) is red, then by lemma 7.38 using the path s_1, s_2, s_4, the edge (s_2, s_4) is not red. So (s_2, s_4) is blue and similarly (s_2, s_5) is blue. Now by lemma 7.36, the edge (s_4, s_5) is not blue. So (s_4, s_5) is red. But now the path s_1, s_2, s_4, s_5 contradicts lemma 7.38. Therefore, (s_1, s_2) is blue, and similarly (s_1, s_3) is blue.

Now we will consider the case where both of the edges (s_2, s_4) and (s_3, s_5) are red.

By lemma 7.37, the element $c_{4,2}c_{2,1}c_{1,3}c_{3,5}c_{1,3}c_{2,1}$ has finite order in W. Now, by proposition 7.15 we have $c_{2,1}c_{1,3}, c_{1,3}c_{2,1} \notin W^{\{s_1\}}$. Also, by proposition 7.14 we have $c_{4,2} \notin W^{\{s_4\}}$ and $c_{3,5} \notin W^{\{s_5\}}$, so $c_{4,2}, c_{3,5} \notin W_{\{s_2,s_3\}}$. Therefore, $c_{4,2}, c_{3,5} \in W^{\{s_1\}} \setminus W_{\{s_2,s_3\}}$ and $c_{2,1}c_{1,3}, c_{1,3}c_{2,1} \in W^{\{s_4,s_5\}} \setminus W_{\{s_2,s_3\}}$. Now, since $m_{1,4}, m_{1,5} = \infty$, we have $W = W^{\{s_4,s_5\}} *_{W_{\{s_2,s_3\}}} W^{\{s_1\}}$, therefore by proposition 1.13, the element $c_{4,2}c_{2,1}c_{1,3}c_{3,5}c_{1,3}c_{2,1}$ has infinite order, a contradiction.

Therefore, the edges (s_2, s_4) and (s_3, s_5) are not both red. Similarly the edges (s_2, s_5) and (s_3, s_4) are not both red. Without loss of generality we may assume that (s_2, s_4) is blue. Then by lemma 7.36, the edge (s_3, s_4) must not be blue, so it is red. Since (s_2, s_5) and (s_3, s_4) are not both red, (s_2, s_5) must be blue. Therefore, using lemma 7.36 on the triangle s_2, s_4, s_5, the edge (s_4, s_5) must be red. Finally the cycle $s_4, s_5, s_2, s_1, s_3, s_4$ satisfies the conditions in lemma 7.39, so s_4 must be joined to s_1 by an edge, which is a contradiction.

Case 2b: s_2 and s_3 are not joined by an edge.

Then s_1, s_2, s_4, s_3, s_1 is a cycle in P_W, so by lemma 7.35 either (s_1, s_4) is an edge, or (s_2, s_3) is an edge. A contradiction.

Case 2c: s_4 and s_5 are not joined by an edge.

Note that now the three vertices s_1, s_4 and s_5 are symmetric to each other, and no two of them are joined by an edge. By lemma 7.38, no two of the edges (s_2, s_1), (s_2, s_4) and (s_2, s_5) are both red. Similarly at most one of the edges (s_3, s_1), (s_3, s_4) and (s_3, s_5) is red. Now by lemma 7.36, at least one of the edges in the cycle s_2, s_4, s_3, s_1, s_2 is red. Without loss of generality, we will assume that (s_2, s_4) is red. Then the edges (s_2, s_1) and (s_2, s_5) are both blue. Now again by lemma 7.36, at least one of the edges in the cycle s_2, s_5, s_3, s_1, s_2 is red. So either (s_3, s_5) or (s_3, s_1) is red. Without loss of generality, we will assume that (s_3, s_5) is red. Then (s_3, s_1) and (s_3, s_4) are both blue. But now the path s_4, s_2, s_1, s_3, s_5 contradicts lemma 7.38.

Case 2d: s_2 and s_4 are not joined by an edge.

Note that if we swap s_1 with s_4 and also swap s_2 with s_5, then the edges do not change. If we remove s_1 from P_W, to form $P_{W^{\{s_1\}}}$ there are still no vertices of degree 1, so the Coxeter system $(W^{\{s_1\}}, S \setminus \{s_1\})$ must be reduced. Therefore, by lemma 7.46, the graph $P_{W^{\{s_1\}}}$ must be isomorphic to one of the coloured graphs described in that lemma. In particular, this means that the edge (s_3, s_5) is red. Similarly, (s_3, s_2) is red. Now, since $P_{W^{\{s_1\}}}$ is isomorphic to one of the graphs described in lemma 7.46 and the edge (s_3, s_2) is red, both the edges (s_3, s_4) and (s_2, s_5) must be blue. Similarly to (s_3, s_4) being blue, the edge (s_3, s_1) is blue.

Now, by lemma 7.36, the edges in the cycle $s_1, s_2, s_5, s_4, s_3, s_1$ are not all blue. So either (s_1, s_2) or (s_4, s_5) is red. Without loss of generality, (s_4, s_5) is red. Now if (s_1, s_2) is red, then the path s_1, s_2, s_5, s_4 contradicts lemma 7.38. Therefore, (s_1, s_2) is blue.

Now we can use lemma 7.37 on the path $s_4, s_5, s_2, s_1, s_3, s_5$. So the element

$$h = c_{4,5}c_{5,2}c_{2,1}c_{1,3}c_{3,5}c_{1,3}c_{2,1}c_{5,2}$$

has finite order in W. Therefore,

$$c_{5,2}hc_{5,2}^{-1} = c_{5,2}c_{4,5}c_{5,2}c_{2,1}c_{1,3}c_{3,5}c_{1,3}c_{2,1}$$

has finite order in W.

Now, by proposition 7.14, we have $c_{4,5} \notin W^{\{s_4\}}$. However, $c_{5,2} \in W^{\{s_4\}}$. Therefore, $c_{5,2}c_{4,5}c_{5,2} \notin W^{\{s_4\}}$, so $c_{5,2}c_{4,5}c_{5,2} \notin W_{\{s_2,s_3\}}$. Also, $c_{3,5} \notin W^{\{s_5\}}$, so $c_{3,5} \notin W_{\{s_2,s_3\}}$. Therefore,

$$c_{5,2}c_{4,5}c_{5,2}, c_{3,5} \in W^{\{s_1\}} \setminus W_{\{s_2,s_3\}}.$$

Also, by proposition 7.15 we have $c_{2,1}c_{1,3}, c_{1,3}c_{2,1} \notin W^{\{s_1\}}$. Therefore,

$$c_{2,1}c_{1,3}, c_{1,3}c_{2,1} \in W^{\{s_4,s_5\}} \setminus W_{\{s_2,s_3\}}.$$

Finally, since s_1 is not joined to s_4 or s_5 by an edge, both $m_{1,4}$ and $m_{1,5}$ are infinite. So we have

$$W = W^{\{s_1\}} *_{W_{\{s_2,s_3\}}} W^{\{s_4,s_5\}}.$$

So by proposition 1.13, the element $c_{5,2}c_{4,5}c_{5,2}c_{2,1}c_{1,3}c_{3,5}c_{1,3}c_{2,1}$ has infinite order in W. A contradiction. □

Lemma 7.49. *If (W, S) is a reduced Coxeter system such that all walls of $\Gamma(W, S)$ are finite and $|S| \geq 5$, then P_W is complete.*

Proof. We will prove this by induction on $|S|$. For $|S| = 5$ we know that the result is true from lemma 7.48. For the inductive step we will assume that $|S| > 5$ and the result is true for any S' with $|S'| = |S| - 1$.

By lemma 7.44, there is some $s \in S$ such that $(W^{\{s\}}, S \setminus \{s\})$ is reduced. Now, by the inductive hypothesis, $P_{W^{\{s\}}}$ is a complete graph. Now let $t \in S \setminus \{s\}$ be any vertex other than s. We will show that (s, t) is an edge in P_W.

By lemma 7.41, we know that P_W is 2-connected, so s has degree at least 2. Let s_1 and s_2 be vertices in P_W which are both joined to s by an edge. If $t = s_1$ or s_2 we are done. Otherwise, let $s_3 \in S \setminus \{s_1, s_2, t, s\}$ be an arbitrary vertex and let $T = \{s, s_1, s_2, s_3, t\}$. Then, since $P_{W^{\{s\}}}$ is a complete graph, each vertex in P_{W_T} has degree at least 2 and P_{W_T} is connected. Therefore, (W_T, T) is a reduced Coxeter system, so by lemma 7.48, the graph P_{W_T} is complete. Therefore, s and t are joined by an edge. Since this works for any vertex $t \in S \setminus \{s\}$, the vertex s is joined to every other vertex in P_W by an edge. Finally, since $P_{W^{\{s\}}}$ is complete, P_W is also complete. □

Theorem 7.50. *Let (W, S) be a reduced Coxeter system such that all walls in $\Gamma(W, S)$ are finite. Then $(W, S) \in \mathscr{F}$ or (W, S) is given by $L_{2m+1,2n+1}$ or L_{2m+1} for some $m, n \in \mathbb{Z}_{>0}$.*

Proof. From lemmas 7.42, 7.46 and 7.49 we know that P_W is either complete or isomorphic to one of two graphs described in lemma 7.46. If P_W is complete, then by lemma 7.34, the group W is finite, so $(W, S) \in \mathscr{F}$. If P_W is isomorphic to one of the graphs described in lemma 7.46, then by lemma 7.47, the Coxeter system (W, S) is equal to some $L_{2m+1,2n+1}$ or L_{2m+1}. □

Theorem 7.51 ((f)⇒(a))**.** *Let (W, S) be a Coxeter system such that every wall in the Cayley graph $\Gamma(W, S)$ is finite. Then $(W, S) \in \mathscr{X}$.*

Proof. We will prove this by induction on $|S|$.

If $|S| = 1$, then W is isomorphic to $\mathbb{Z}_2$, which is finite. Hence $(W, S) \in \mathscr{X}$.

Now assume that the statement of the theorem is true for all Coxeter systems (W', S') with $|S'| < |S|$. If (W, S) is reduced, then by the previous theorem, $(W, S) \in \mathscr{F}$ or (W, S) is given by some $L_{2m+1,2n+1}$ or L_{2m+1}, hence, $(W, S) \in \mathscr{X}$. If (W, S) is not reduced, then either P_W is disconnected or P_W contains some vertex s which has degree 1 and is only connected to a blue edge.

If P_W is disconnected, then W decomposes as a free product, $W = W_T * W^T$ for some subset $T \subsetneq S$ with $T \neq \phi$. By lemma 7.32, all walls in $\Gamma(W_T, T)$ and $\Gamma(W^T, S \setminus T)$ are finite, so by the inductive hypothesis, we have $(W_T, T), (W^T, S \setminus T) \in \mathscr{X}$. Therefore, $(W, S) \in \mathscr{X}$.

Finally, if P_W contains some vertex s which has degree 1 and is only connected to a blue edge, let this edge be (s, t). Then (W, S) is a $(2m + 1)$-extension of $(W^{\{s\}}, S \setminus \{s\})$, where $2m+1$ is the label on the edge (s, t) in the Coxeter diagram for (W, S). Moreover, by lemma 7.32, all walls in $\Gamma(W^{\{s\}}, S \setminus \{s\})$ are finite, so by the inductive hypothesis, $(W^{\{s\}}, S \setminus \{s\}) \in \mathscr{X}$. Therefore, since $\mathscr{X}$ is closed under $(2m + 1)$-extensions, the Coxeter system $(W, S) \in \mathscr{X}$. □

Bibliography

[1] Avenhaus, Jürgen, Klaus Madlener, and Friedrich Otto. "Groups presented by finite two-monadic Church-Rosser Thue systems." Transactions of the American Mathematical Society 297.2 (1986): 427-443.

[2] Bass, Hyman. "Covering theory for graphs of groups." Journal of pure and applied Algebra 89.1 (1993): 3-47.

[3] Blokhuis, A., and A. E. Brouwer. "Geodetic graphs of diameter two." Geometries and Groups. Springer Netherlands, 1988. 527-533.

[4] Brady, Noel, and Martin R. Bridson. "There is only one gap in the isoperimetric spectrum." Geometric & Functional Analysis GAFA 10.5 (2000): 1053-1070.

[5] Cannon, James W. "Almost convex groups." Geometriae Dedicata 22.2 (1987): 197-210.

[6] Ceccherini-Silberstein, Tullio, et al. "Groups, graphs, languages, automata, games and second-order monadic logic." European Journal of Combinatorics 33.7 (2012): 1330-1368.

[7] Davis, Michael. The geometry and topology of Coxeter groups. Vol. 32. Princeton University Press, 2008.

[8] Diekert, Volker, and Armin Weiß. "Context-Free Groups and Bass-Serre Theory." arXiv preprint arXiv:1307.8297 (2013).

[9] Dunwoody, Martin J. "The accessibility of finitely presented groups." Inventiones mathematicae 81.3 (1985): 449-457.

[10] Elder, Murray J. "The loop shortening property and almost convexity." Geometriae Dedicata 102.1 (2003): 1-17.

[11] Elder, Murray. "Regular geodesic languages and the falsification by fellow traveler property." Algebraic & Geometric Topology 5 (2005): 129-134.

[12] Elder, Murray, and Susan Hermiller. "Minimal almost convexity." Journal of Group Theory 8.2 (2005): 239-266.

[13] Gromov, M. "Hyperbolic groups, essays in group theory, MSRI series vol. 8, S. Gersten, ed." (1988): 54-64.

[14] Haring-Smith, Robert H. "Groups and simple languages." Transactions of the American Mathematical Society 279.1 (1983): 337-356.

[15] Harrison, Michael A. Introduction to formal language theory. Addison-Wesley Longman Publishing Co., Inc., 1978.

[16] Karrass, Abraham, Alfred Pietrowski, and Donald Solitar. "Finite and infinite cyclic extensions of free groups." Journal of the Australian Mathematical Society 16.04 (1973): 458-466.

[17] Kuske, Dietrich, and Markus Lohrey. "Logical aspects of Cayley-graphs: the group case." Annals of Pure and Applied Logic 131.1 (2005): 263-286.

[18] Magnus, Karrass and Solitar. "Combinatorial Group Theory"

[19] Miller, Charles F., and Michael Shapiro. "Solvable BaumslagSolitar groups are not almost convex." Geometriae Dedicata 72.2 (1998): 123-127.

[20] Muller, David E., and Paul E. Schupp. "Groups, the theory of ends, and context-free languages." Journal of Computer and System Sciences 26.3 (1983): 295-310.

[21] Neumann, Walter D., and Michael Shapiro. "Automatic structures, rational growth, and geometrically finite hyperbolic groups." Inventiones mathematicae 120.1 (1995): 259-287.

[22] Ore, Oystein. Theory of graphs. Vol. 38. American Mathematical Soc., 1965.

[23] Papasoglou, P. "On the subquadratic isoperimetric inequality." Geometric Group Theory. Boston: Jones and Bartlett (1995): 149-158.

[24] Ponaru, V. "Almost convex groups, Lipschitz combing, and π_1^∞ for universal covering spaces of closed 3-manifolds." Journal of Differential Geometry 35.1 (1992): 103-130.

[25] Serre, J. P. "Trees, 1980."

[26] Shapiro, Michael. "Pascal's triangles in Abelian and hyperbolic groups." Journal of the Australian Mathematical Society (Series A) 63.02 (1997): 281-288.

[27] Stallings, John. "A remark about the description of free products of groups." Mathematical Proceedings of the Cambridge Philosophical Society. Vol. 62. No. 02. Cambridge University Press, 1966.

[28] Stallings, John R. "On torsion-free groups with infinitely many ends." Annals of Mathematics (1968): 312-334.

[29] Stallings, John R. Group theory and three-dimensional manifolds. Vol. 83. New Haven: Yale University Press, 1971.

[30] Stemple, Joel G., and Mark E. Watkins. "On planar geodetic graphs." Journal of Combinatorial Theory 4.2 (1968): 101-117.

[31] Stemple, Joel G. "Geodetic graphs of diameter two." Journal of Combinatorial Theory, Series B 17.3 (1974): 266-280.

[32] Thiel, Carsten. Zur fast-Konvexitt einiger nilpotenter Gruppen. No. 234. Mathematisches Institut der Universitt Bonn, 1992.