Hyperbolic groups

an introduction

Aluna Rizzoli

A dissertation presented for the degree of
Master of Mathematics

Mathematics and Statistics
University of St Andrews
United Kingdom
April 22, 2017

Abstract

For every finitely generated group it is possible to construct a metric space (a Cayley graph) on which the group acts by isometries. By analysing the geometry of this particular class of metric spaces we give a definition of hyperbolic groups. We then investigate other characterisations of hyperbolicity that maintain a geometric flavour. We then move to analysing more algebraic properties that are connected to the word problem. It is here that we begin to realise how our starting notion of hyperbolicity (due to Gromov) relates to the word problem. By expanding and sharpening existing results from the literature on the subject we show that an equivalent characterisation of hyperbolicity is in fact the existence of particular type of presentation (a Dehn presentation), which gives a very efficient solution to the word problem.

Contents

1 Introduction

Group presentations were probably first introduced by Walther von Dyck in 1882 and are the starting point of a lot of modern (combinatorial) group theory [vonDyck1882][Mül2017]. In his highly influential 1911 paper [Dehn1911] Max Dehn posed, in a clear manner, questions that in his opinion would determine the direction of (combinatorial) group theory. Looking at groups in terms of generators and relators Dehn phrased the word problem, the conjugacy problem and the isomorphism problem for groups [Dehn1911][Mül2017]. In the translated [Dehn1987] original formulation they were as follows:

- **The identity [word] problem:** *An element of the group is given as a product of generators. One is required to give a method whereby it may be decided in a finite number of steps whether this element is the identity or not.*

- **The transformation [conjugacy] problem:** *Any two elements S and T of the group are given. A method is sought for deciding the question whether S and T can be transformed into each other, i.e. whether there is an element U of the group satisfying the relation $S = UTU^{-1}$.*

- **The isomorphism problem:** *Given two groups, one is to decide whether they are isomorphic or not (and further, whether a given correspondence between the generators of one group and elements of the other is an isomorphism or not).*

As pointed out by Dehn, these three problems have different degrees of difficulty and it is apparent that the first question, i.e. determining if a word is the identity or not, is the easiest of the three and is what we will be concerned with. Notice, for example, that if one is able to determine if two words are conjugates, one can check if a word is the identity by checking if it is a conjugate of the identity.

Dehn provided an answer to the word problem question for surface groups, and gave the algorithm (*Dehn algorithm*) in question. Gromov generalised this to *hyperbolic groups*, a concept that he developed and introduced in 1987 [Gr1987]. Gromov showed that the existence of a *Dehn algorithm* actually characterises his notion of hyperbolicity in groups. The starting definition of hyperbolicity will depend on the 'geometry' of groups. The idea is to construct a (geodesic) metric space on which the group acts by isometries and by doing so gather information about the group. This idea is fundamental in the area of geometric group theory, which gained momentum with Gromov's work [Br1999].

Our main focus throughout the dissertation will be to understand various different characterisations of hyperbolicity and their connection to the word problem. The writer's main aim is to revisit proofs given in the standard literature on the subject and fill in the details when needed. On a side note we also try to optimise the constants that appear in the proofs, something that was not considered in [Alonso et al. 1990]. While doing all of this we fix some issues encountered in [Alonso et al. 1990].

In this dissertation we don't consider the other two problems set by Dehn, the conjugacy and isomorphism problems, for which it turns out hyperbolic groups provide a very nice solution.

In order to do this we start by going through a basic set of definitions such as free groups, group presentations and words in Section 2.
In Section 3 we set the groundwork for the first definition of hyperbolicity, by seeing how it is possible to turn a group into a metric space and considering a particular class of metric spaces.
In Section 4 we give a first characterisation of hyperbolicity (based on *slim triangles*).
In Section 5 we provide different equivalent characterisations, while maintaining a geometric point of view.
In Section 6 we introduce the fundamental concept of Cayley or disc diagrams.
Section 7 is where the hardest results are proven and is by far the longest section of the dissertation. Here we move away from only thinking about triangles and provide different characterisations of hyperbolicity that are strongly related to the word problem. In this section we improve on one result (Theorem 7.12) by [Alonso et al. 1990], while fixing some of the lemmas and propositions leading to it (original work).
In Section 8 we conclude by rounding up the results of the previous sections.

The biggest group of intended readers of this dissertation consists of advanced undergraduate students and first year graduate students who wish to get to know the fundamentals of hyperbolic groups. Due to the presence of some original work which sharpens a couple of standard results, a potential reader of a small part of the dissertation can also be found in a computational group theorist.

The assumed knowledge is in line with what a final year undergraduate student typically knows. Some basic definitions, such as the ones for free groups and group presentations, are restated, but a reader who has never encountered them before is strongly advised to consult some of the cited resources. The hardest part of the background theory is probably the one regarding disc diagrams, to which we dedicate a whole section, mainly stating standard results with only sketches of proofs (see [LynSch2002] for a better understanding). The latex package TikZ ([TikZ2013]) has been used to generate all the pictures that are present throughout the dissertation.

2 Basic definitions and results

In this section we give the most basic definitions on which we are going to build our theory about hyperbolic groups. Most definitions are standard, unless it is specifically pointed out. The reader might want to consult [Rot1995]. The examples are original, unless otherwise noted.

The first notion we introduce is that of a word. In general terms a word is just a sequence of letters. If these letters are taken from some particular setting, such as group elements then we get the following definition.

Definition 2.1. Given a set of elements X of a group G we denote by X^{-1} the set of inverses of the elements in X and we say that a *word* w on X is an expression of the form $w = x_1 x_2 \dots x_n$ where $n \in \mathbb{N}_0$ and $x_i \in X \cup X^{-1}$ for all $1 \leq i \leq n$. We denote by $W(X)$ the set of all such words. If $n = 0$ we say that w is the empty word, which we denote by ϵ.

If we were not told that $w = x_1 x_2 \dots x_n$ is a word, we would think we are just multiplying together group elements $x_1, x_2, \dots, x_n$. We can in fact define equivalence classes of words in G as the following definition points out.

Definition 2.2. Given a word w on a set X in a group G, since w is just a sequence of elements of the group G, we can evaluate w in G. If v is another word which equates to the same element of G we say that $w =_G v$.

Example 2.3. If a, b are the standard generators for the group $\mathbb{Z}^2$, then the set of words on $\{a, b\}$ is $W(\{a, b\}) = \{x_1 \dots x_n : n \in \mathbb{N}_0, x_i \in \{a, b, a^{-1}, b^{-1}\}\}$.

In particular $abb^{-1}a^{-1}, \epsilon \in W(\{a, b\})$ but even though $abb^{-1}a^{-1} =_{\mathbb{Z}^2} \epsilon$ the two words are different in the set of words.

We will often refer to the set of words on some set X without specifying any group. To make sense of this we need to define a set of inverses X^{-1} which is simply a set disjoint from X, but in bijective correspondence to X. Then $W(X)$ is just the set of expressions of the form $w = x_1x_2 \dots x_n$ with $x_i \in X \cup X^{-1}$, as before.

Definition 2.4. Given a word $w = x_1 \dots x_n$ on a set X we say that the word $y = y_1 \dots y_m$ is a *subword* of w if there are $i, j \in \mathbb{N}$ with $i \leq j$ such that $x_i \dots x_j = y$, or if y is the empty word.

Example 2.5. The word $xy^{-1}y$ is a subword of the word $xyxy^{-1}yx^{-1}$ in the set of words on $\{x, y\}$.

When thinking of words as elements of a group it is not efficient to have subwords that are equivalent to the identity. In some cases it is easy to find them, for example when they are just a product of a word and its inverse.

Definition 2.6. Given a word $w = x_1 \dots x_n$ on a set X we say that w is *(freely) reduced* if it contains no subword of the form xx^{-1} or $x^{-1}x$ for some $x \in X$. We say that it is *cyclically reduced* if it is reduced and $x_1 \neq x_n^{-1}$.

Example 2.7. Let $X = \{a, b\}$. The word $ababb^{-1} \in W(X)$ is not freely reduced since it contains bb^{-1} as a subword. The word $ababa^{-1}$ is reduced but is not cyclically reduced.

If we take a word $w = uv$ and we conjugate it by u and reduce it we get the word vu. As we will be interested in words that are equal to the identity it makes sense to consider conjugates, since being equivalent to the identity is preserved under conjugation.

Definition 2.8. A *cyclic conjugate* of a word $w = x_1 \dots x_n$ is any word $x_i \dots x_n x_1 \dots x_{i-1}$ for $1 \leq i \leq n$.

Example 2.9. Let $X = \{a, b, c\}$. The set of cyclic conjugates of the word $abca^{-1}$ is $\{bca^{-1}a, ca^{-1}ab, a^{-1}abc, abca^{-1}\}$.

Definition 2.10. Given a word $w = x_1x_2 \dots x_n$ on a set X, we denote by $|w|$ the length of w, i.e. $|w| = n$. If w is the empty word then $|w| = 0$.

Definition 2.11. Given two words on a set X, $w = x_1 \dots x_n$ and $v = y_1 \dots y_m$, we can define the associative operation of *concatenation* which gives the word $wv = x_1 \dots x_n y_1 \dots y_m$.

Before defining a *group presentation* we need the notion of a *free group*. Different authors like to start with different definitions of free groups. Here we follow the approach found in [Rot1995, Chapter 11] which consists of giving a more abstract definition followed by a way of constructing a more concrete one.

Definition 2.12. Let X be a subset of a group F. Then F is a *free group with free basis* X if for any group H and any set function $\phi : X \to H$ there exists a unique homomorphism $\tilde{\phi} : F \to H$ such that $\tilde{\phi}(x) = \phi(x)$ for all $x \in X$.

It can be shown that up to isomorphism there is a unique free group on a basis X ([Rot1995, Theorem 11.4]). This justifies the following definition.

Definition 2.13. Given a set X we denote by F_X the unique *free group with basis* X up to isomorphism.

The elements of F_X are equivalence classes of words on X where two words are in the same equivalence class if one can get from one to the other by a finite sequence of expansions and contractions (adding or removing a subword of the form xx^{-1} or $x^{-1}x$ for $x \in X$). Multiplication between two equivalence classes is just given by the equivalence class of the concatenation of any two words in the equivalence classes we are multiplying.

An alternative way of looking at the elements of the free group is by choosing a representative for each equivalence class and defining an operation between the chosen representatives. The following two propositions and the fact that the definition is well defined can be easily derived from [Rot1995, Theorem 11.1].

Proposition 2.14. Let F_X be a free group with free basis X. Then there is a unique freely reduced word for each equivalence class.

Definition 2.15. We define an operation of multiplication between two freely reduced words w, u to be equal to the freely reduced word that we get when doing all possible contractions (well defined, see [Rot1995]) on the word obtained when concatenating u on the right of w.

Proposition 2.16. The set of freely reduced words in a free group with free basis X, equipped with the operation defined above, forms a group isomorphic to F_X.

Example 2.17. Let $X = \{a, b\}$. Then in the free group F_X the equivalence class of the word $abb^{-1}baa$ has the word $abaa$ as the unique freely reduced representative. The multiplication with the equivalence class of the word $a^{-1}b$ gives the equivalence class of the word $abb^{-1}baaa^{-1}b$ which has $abab$ as the representative. The word $abab$ is in fact what we obtain after reducing the word $abaaa^{-1}b$.

We are finally ready to define *group presentations*.

Definition 2.18. Given a set X and a set $R \subseteq F_X$ we denote by $\langle X|R\rangle$ the group

$$G = F_X/\langle\langle R\rangle\rangle$$

where $\langle\langle R\rangle\rangle$ denotes the normal closure of R in F_X. We say that G has presentation $\langle X|R\rangle$.

Example 2.19. The presentation $\langle a, b|a^2, b^3, abab\rangle$ is a presentation for the group S_3. Another equivalent presentation for S_3 is $\langle a, b|a^2, b^2, (ab)^3\rangle$. Note that this shows that presentations are not unique.

A presentation for a group can give information about the group itself and the concept is fundamental in combinatorial group theory. One of the reasons presentations are so important is that it is possible to give a presentation for any given group. What is not necessarily easy to do is finding an 'efficient' presentation for a group.

Proposition 2.20. Every group G has a presentation.

Proof. Let $\phi : G \to G$ be the identity set function. Then there is an extension of ϕ to a surjective homomorphism $\tilde{\phi} : F_G \to G$ sending $g \in F_G$ to $g \in G$. If we denote $\ker \tilde{\phi}$ by R we have that by the First Isomorphism Theorem $F_G/R \cong G$. □

The construction given above is clearly not efficient and could be improved by defining ϕ only from a set of generators for G. In fact if the group G is infinite, the construction gives us a presentation of G with infinitely many generators. We will be mostly interested in groups that have a presentation with finitely many generators and finitely many relators. This leads us to the next definition.

Definition 2.21. A group $G = \langle X|R\rangle$ is *finitely presented* by the *finite presentation* $\langle X|R\rangle$ if both X and R are finite. We say that a group is *finitely presentable* if it has a finite presentation.

As seen above presentations are not unique. It also easy to take a presentation and create a new, but equivalent (it presents an isomorphic group), one. In 1908 Tietze formalised this process by defining four different types of transformations that produce equivalent presentations. They are as follows:

Definition 2.22. [Tietze1908][LynSch2002, Chapter 2.2] Let $\langle X|R\rangle = G$ be a presentation (not necessarily finite). We define four types of *Tietze transformations.*

(i) Add a relator. If $w \in WP(G, X)$ transform to the presentation $\langle X|R\cup \{w\}\rangle$.

(ii) Remove a relator. If $r \in R$ and $r \in \langle\langle R \setminus \{r\}\rangle\rangle$ transform to the presentation $\langle X|R \setminus \{r\}\rangle$.

(iii) Add a generator. Transform to a presentation $\langle X \cup \{g\}|R \cup \{wg^{-1}\}\rangle$ where w is a new symbol $\notin X$ and $w \in W(X)$.

(iv) Remove a generator. If $g \in X$ and $wg^{-1} \in R$ with $w \in W(X \setminus \{g\})$ transform to a presentation $\langle X \setminus \{g\}|R'\rangle$ where R' is obtained from R by removing the relator wg^{-1} and substituting w in the place of every subword g of any relator.

It is rather clear that each of these transformations takes a presentation for a group to a presentation for an isomorphic group. In the following example we can see the transformations in action on presentations for the group S_3.

Example 2.23. The numbers above the arrows indicate the type of Tietze transformation used.

$$S_3 = \langle a, b|a^2, b^3, abab\rangle \xrightarrow{\text{(iii)}} \langle a, b, c|a^2, b^3, abab, abc^{-1}\rangle \xrightarrow{\text{(i)}} \langle a, b, c|a^2, b^3, abab, abc^{-1}, acb^{-1}\rangle \xrightarrow{\text{(ii)}}$$

$$\langle a, b, c|a^2, b^3, abab, acb^{-1}\rangle \xrightarrow{\text{(iv)}} \langle a, c|a^2, (ac)^3, aacaac\rangle \xrightarrow{\text{(i)+(ii)}} \langle a, c|a^2, (ac)^3, c^2\rangle.$$

What makes Tietze transformations special is the fact that they are the building blocks of any possible type of transformation which produces equivalent presentations. In fact the following powerful theorem holds.

Theorem 2.24. [Tietze1908] Given two equivalent finite presentations there is a finite sequence of Tietze transformations taking one to the other.

Given two presentations which have disjoint sets of generators and relators we can create a new presentation by merging the two in the following way:

Definition 2.25. Let $G = \langle X|R\rangle$ and $H = \langle Y|S\rangle$, with $X \cap Y = R \cap S = \emptyset$. We define the *free product* of G and H, denoted by $G*H$, to be $\langle X \cup Y|R \cup S\rangle$.

As discussed in the introduction we now give the very important definition of word problem.

Definition 2.26. The *word problem* for a group G generated by a set X is the set of words

$$\{w = x_1 \dots x_n : w =_G 1, n \in \mathbb{N}_0, x_i \in X \cup X^{-1}\}$$

and is denoted by $WP(G, X)$.

Example 2.27. The word problem for $\mathbb{Z}^2 = \langle a, b|[a, b]\rangle$ is the set of words on $\{a, b\}$ such that the sum of the exponents of the letters of type a and the sum of the exponents of the letters of type b is 0. For example $abab^{-1}a^{-1}a^{-1} \in WP(\mathbb{Z}^2, \{a, b\})$ while $aaba^{-1}b^{-1} \notin WP(\mathbb{Z}^2, \{a, b\})$.

We now introduce a particular type of presentation that will provide us with a nice solution to the word problem.

Definition 2.28. A finite presentation $\langle X|R\rangle$ of a group G is called a *Dehn presentation* if there is a finite set of words $u_1, v_1, \dots, u_n, v_n$ on X with $u_i =_G v_i$ and $|u_i| < |v_i|$ for all $1 \leq i \leq n$ such that $R = \{u_1 v_1^{-1}, \dots, u_n v_n^{-1}\}$ and any freely reduced non-empty word w which is the identity in G has v_i as a subword for some i.

We now provide some examples of Dehn presentations.

Example 2.29. If $G = \langle X\rangle$ is a finite group then we can simply set R to be the set of freely reduced words on X that are the identity in G. Since G is finite we have that R is finite, and we get a Dehn presentation $G = \langle X|R\rangle$.

To create an example for infinite groups we can use free products, thanks to the following proposition. We only give a sketch of proof since free products go beyond the scope of this dissertation.

Proposition 2.30. Let $G = \langle X|R \rangle$ and $H = \langle Y|S \rangle$ be Dehn presentations, with $X \cap Y = R \cap S = \emptyset$. The free product

$$G * H = \langle X \cup Y | R \cup S \rangle$$

is then a Dehn presentation.

Proof. From the normal form theorem for free products ([Rot1995, Theorem 11.52]) we know that a freely reduced word in $G * H$ is equivalent to the identity if and only if the alternating blocks belonging to the original groups are equivalent to the identity in their original groups. Using the fact that we started with two Dehn presentations it is then trivial to conclude that we can find a subword which is more than half than one relator in one of the two original groups, and hence also in the free product. □

Example 2.31. The free product of any two finite groups has a Dehn presentation.

It is more difficult to get something more interesting, i.e. a group which is not finite,a free group, or a free product of finite groups, and which has a Dehn presentation. The example we provide was proven by Dehn himself using arguments from small cancellation theory.

Example 2.32. [Dehn1987] The standard presentation for the fundamental group of surfaces of genus $g \geq 2$, $\left\langle a_1, b_1, \ldots, a_g, b_g | a_1 b_1 a_1^{-1} b_1^{-1} \ldots a_g b_g a_g^{-1} b_g^{-1} \right\rangle$, is a Dehn presentation.

We finally provide an example of a group with no Dehn presentation. The proof is deferred.

Example 2.33. The group $\mathbb{Z}^2$ has no Dehn presentation. We will be able to derive this result easily after our discussion about hyperbolic groups. Furthermore no group which contains a copy of $\mathbb{Z}^2$ as a subgroup has a Dehn presentation [BrHae2011, Corollary 3.10].

The following proposition gives us an idea of how special a group with a Dehn presentation is.

Proposition 2.34. [GhHar1990] Let $\langle X|R \rangle = G$ be a Dehn presentation. Then there exist finitely many conjugacy classes of elements of finite order in G.

Proof. Let C be a conjugacy class of an element of finite order which is not the identity in G. Consider an element $w = x_1 \dots x_k \in C$ of minimal length. Since w is a conjugate of an element of finite order it must also have finite order, i.e. $w^n =_G 1$ for some $n \in \mathbb{N}$. Since $\langle X|R \rangle$ is a Dehn presentation there must be $u, v \in F(X)$ such that $|u| > |v|$, u is a subword of w^n and $u =_G v$. Now if $|u| \leq |w| = k$ we have that there is a cyclic conjugate of w which can be reduced to a shorter word in G, but this is absurd since w is of minimal length in its conjugacy class.

Hence $|w| < |u|$ and if $\rho = \max_{r \in R} |r|$ we have $|w| < \rho$. But there are only finitely many words on X which have length less than ρ. We can therefore conclude that there are only finitely many conjugacy classes of elements of finite order. □

Definition 2.35. We say that a word problem $WP(G, X)$ is *solvable* if there is an algorithm that can determine whether a word on X is in $WP(G, X)$ or not, in finite time.

In general there is no single algorithm to solve word problems. In fact it has been proven that some presentations have undecidable word problems (see [Rot1995, Chapter 12]), i.e. word problems for which there is no algorithm that solves them in finite time. In the case of a Dehn presentation we can however do the following:

Definition 2.36. A Dehn presentation $\langle X|R \rangle = G$ gives an efficient algorithm, called a *Dehn algorithm*, for solving the word problem $WP(G, X)$. Given a word w we look for a subword which equals one of the v_i and if we can find it we substitute it with u_i and repeat the process. A word w will then be equivalent to the identity in G if and only if the algorithm terminates in the empty word.

The most efficient algorithm possible will have computation time depending linearly on the length of the word. In what follows we give a more formal definition of the Dehn algorithm describing how it can be implemented on a Turing machine.

Proposition 2.37. [DomAn1985]. Let $\left\langle X|u_1 v_1^{-1}, \dots, u_n v_n^{-1} \right\rangle = G$ be a Dehn presentation, where we adopt the notation from the definition. Then the word problem $WP(G, X)$ can be solved in a time bounded linearly by the length of the input word.

Proof. Let $w = x_1 \dots x_m \in W(X)$ be an input word. Let $Q = \max_{1 \le i \le n} |u_i|$. We are going to use a Turing machine with two stacks S, T with enough memory to hold Q letters each. The Turing machine will accept w if and only if $w \in WP(G, X)$.

We now give a description of a general step of the algorithm. Assume that i letters from w have been read and that $S = s_1 \dots s_j$ and $T = t_1 \dots t_k$ (when we start we have $i = j = k = 0$ and hence S, T are empty). Furthermore assume that each time we add a letter to a stack we freely reduce if we have the possibility, so that each stack always contains a reduced word. Let $p = \min\{j, Q\}$.

We have the following cases:

(1) The last p letters of S contain a subword equal to some u_l.

(2) The last p letters of S do not contain a subword equal to any u_l. This splits into the following sub-cases:

   (a) $i < n$.

   (b) $i = n; j = 0; k = 0$.

   (c) $i = n$ and either $j \neq 0$ or $k \neq 0$.

First of all we notice that we can easily check if the conditions (1) or (2) are satisfied, in bounded time (irrespective of the input word), since we have finitely many u_l's and Q is fixed.

In case (1) we can substitute u_l with v_l. From the subword $s_{j-p+1} \dots s_j$ we get a shorter word $z_1 \dots z_r$. We then transfer the shortened subword to the stack T so that $T = t_1 \dots t_k z_r \dots z_1$ and $S = s_1 \dots s_{j-p}$.

In case $(2)(a)$ if $k \neq 0$ we transfer the last letter of T, i.e. t_k, to the stack S so that $S = s_1 \dots s_j t_k$ and $T = t_1 \dots t_{k-1}$. If $k = 0$ we read a new input letter onto S so that $S = s_1 \dots s_j x_{i+1}$.
In case $(2)(b)$ we accept w and in case $(2)(c)$ we reject it.

What ends up happening when we adopt this algorithm is that the stack S is filled up until there is a subword equal to some u_l and through the stack T we are able to substitute it with the shorter word v_l and at the same time keep checking if there is a subword equal to some u_l. Since we know that a word in the word problem is either trivial or contains one u_l, we know that this process will stop with empty stacks S, T if and only if the input word is in the word problem. □

The algorithm described in [DomAn1985] runs in linear time, but it does not read the input at constant speed. It is possible to construct an algorithm that runs in real time, as shown in [Holt2001].

Given a presentation $\langle X|R\rangle = G$ we have that by definition an element $w \in WP(G,X)$ must be in $\langle\langle R\rangle\rangle$ and we must therefore be able to write $w = \prod_{i=1}^{N} c_i^{-1} r_i c_i$ for some $r_i \in R \cup R^{-1}$ and $c_i \in F_X$. This leads us to the following definition. Note that one can find the three following definitions and relative discussion in [Gerst1991].

Definition 2.38. Let $\langle X|R\rangle = G$ be a presentation. We define the *area of a relation* to be a function A which given any freely reduced word w in $WP(G,X)$ assigns to it the least number N of conjugates of relators or inverses of relators, denoted by $A(w)$, needed to write

$$w = \prod_{i=1}^{N} c_i^{-1} r_i c_i.$$

In general we don't know anything about how many terms are needed, i.e. how big $A(w)$ is. If we are working with a finite presentation we only have finitely many words up to a certain fixed length. We can therefore define the following useful function.

Definition 2.39. Let $\langle X|R\rangle = G$ be a finite presentation. We define the *Dehn function* of $\langle X|R\rangle$ to be the function $Dehn : \mathbb{N}_0 \to \mathbb{N}_0$ such that $Dehn(n)$ is the maximum of $A(w)$ over all freely reduced words $w \in WP(G,X)$ with $|w| \leq n$.

It is in fact true that a finite presentation for a group has a solution to the word problem if and only if its Dehn function is recursive (in a computational sense, see [Sk1923]). A proof of this powerful result can be found in [Gerst1991] and is beyond the scope of this dissertation. Note that one particular case of solvable word problem occurs when the Dehn function is a polynomial.

Intuitively it makes sense to imagine that if the Dehn function does not have a fast growth then the word problem has an easier solution. In particular the best non-trivial case we can hope for is when the Dehn function is linear. This takes us to the following definition.

Definition 2.40. A finite presentation $\langle X|R\rangle = G$ satisfies a *linear isoperimetric inequality* with constant $K \in \mathbb{R}$ if for any freely reduced $w \in WP(G,X)$ we have $A(w) \leq K|w|$.

As pointed out in [Alonso et al. 1990] we limit ourselves to the finitely generated case, since we could otherwise add every word in the word problem to the set of relators and get a very uninteresting linear isoperimetric inequality. At the end of our discussions we will be able to understand how having a Dehn presentation and satisfying a linear isoperimetric inequality are equivalent conditions. For now we limit ourselves to showing one direction, as the following theorem states. Note that results until the end of this section are standard, but the proofs are original work.

Theorem 2.41. If $\langle X|R\rangle = G$ is a Dehn presentation then $\langle X|R\rangle$ satifies a linear isoperimetric inequality with constant $K = 1$.

Proof. Let w be a freely reduced word in $WP(G, X)$. We prove the theorem by induction on $|w|$. If $|w| = 0$ then the statement is trivial.

Suppose that up to $M \in \mathbb{N}_0$, if $|w| \leq M$, then we can write $w =_{F_X} \prod_{i=1}^{N} c_i^{-1} r_i c_i$ for $c_i \in F_X$, $r_i \in R$ and with $N \leq |w|$.

Suppose w has $|w| = M + 1$. Since $\langle X|R\rangle$ is a Dehn presentation there is a subword v of w and a word u, with $|u| < |v|$ and with $r = uv^{-1} \in R$. Since v is a subword of w, $w = avb$ for some subwords a, b. Thus $w =_{F_X} ar^{-1}ub =_{F_X} ar^{-1}a^{-1}aub =_{F_X} (ar^{-1}a^{-1})aub$.

Since w is equal to the identity in G and $ar^{-1}a^{-1}$ is the conjugate of the inverse of the relator r by a^{-1}, we must have that aub is the identity in G as well. But $|aub| < |w| = M + 1$, so by the inductive assumption we can write aub as a a product of at most M conjugates of relators. Hence we can write w as product of at most $M + 1$ conjugates, as required. □

We now consider a couple of easy, but significant, examples.

Example 2.42. The standard presentation for a free group with finite free basis X, $\langle X|\emptyset\rangle$, satisfies a linear isoperimetric inequality with constant $K = 0$, since the only freely reduced word in the word problem is the empty word.

Example 2.43. The presentation $\langle X \cup \{y\}|y\rangle = G$ is also a presentation for a free group with free basis X. However in this case the presentation satisfies a linear isoperimetric inequality with constant $K = 1$ but not one with $K = 0$ since the word y^n is freely reduced and in $WP(G, X)$, but can't be written with fewer than n conjugates of relators or inverses of relators.

Another important property of linear isoperimetric inequalities is that if we have a presentation which satisfies one and we use a Tietze transformation, we get a presentation which also satisfies a linear isoperimetric inequality. We therefore get a sort of invariance under change of generators and relators. Note however that the two examples above show us that we need to be careful about this. In fact, in general, it must be the case that some linear terms are added to the inequality when changing presentations, since when adding just one generator and relator to the standard presentation of a free group we do not get a linear isoperimetric inequality with the constant being a multiple of the previous one.

The following theorem formalises the invariance of the linear isoperimetric inequality under Tietze transformations.

Theorem 2.44. Let $\langle X|R\rangle = G$ be a finite presentation satisfying a linear isoperimetric inequality (with constant K). Then any single Tietze transformation produces a presentation satisfying a linear isoperimetric inequality.

Proof. We need to consider the different types of Tietze transformations.

(i) It is clear that adding a single relator can only improve the linear isoperimetric constant, and in particular gives a presentation which still satisfies a linear isoperimetric inequality with constant K.

(ii) If we remove a relator $r \in R$ such that $r \in \langle\langle R \setminus \{r\}\rangle\rangle$ we can let Q be the area of r in the new presentation. Now if we have a freely reduced word $w \in WP(G, X)$ we know that it is possible to write it as a product of at most $K|w|$ conjugates of relators or inverses of relators in R. Every time r appears in the product we can substitute it with a product of at most Q relators or inverses of relators in $R \setminus \{r\}$. We therefore see that the new presentation satisfies a linear isoperimetric inequality with constant at KQ.

(iii) If we add a new generator g and a new relator wg^{-1} such that $w \in W(X)$ we consider the new presentation $\langle X'|R'\rangle$ a freely reduced word v on $X' = X \cup \{g\}$ which is equivalent to the identity in the new presentation. If the word v does not contain the generator g as a subword we are done, since we started with a presentation satisfying a linear isoperimetric inequality and we only added a generator and a relator. Otherwise we have

$$v =_{F_{X'}} u_1 g u_2 =_{F_{X'}} u_1 g w^{-1} u_1^{-1} u_1 w u_2 =_{F_{X'}} u_1 (g w^{-1}) u_1^{-1} u_1 w u_2.$$

We are therefore able to push every subword of v equal to the new generator g to the left in the form of a conjugate of the inverse of the relator gw^{-1}. What is left is then a word in $WP(G,X)$ of length at most $|v||w|$. Thus in the worst case scenario we get a word equal to v in $F_{X'}$ starting with $|v|$ conjugates of g and then a subword of length $|v||w|$ in $WP(G,X)$ which we can write as a product of at most $|v||w|K$ conjugates of relators or inverses of relators. Therefore we can write v as a product of at most $|v||w|K+|v|$ conjugates of relators or inverses of relators in R', getting that the new presentation satisfies a linear isoperimetric inequality with constant $K|w|+1$.

(iv) If we remove a generator g which appears in a relator of the form wg^{-1} with $w \in W(X \setminus \{g\})$ we consider a freely reduced word on $X \setminus \{g\}$ equivalent to the identity in the group. This is also a word in $WP(G,X)$ so we can write it as a product of at most K conjugates of relators or inverses of relators in R. If we then substitute w in the place of g for every relator in the product we get what we needed, i.e. a product using only relators of the new presentation. The isoperimetric inequality doesn't thus change.

□

We can compare the results of Theorem 2.44 with the two previous examples. We see that the addition of a generator and a relator to the standard presentation for a free group corresponds to case (*iii*) in our discussion, where we in fact see that the linear isoperimetric constant needs to increment by 1. We can finally get the corollary that we interested in, which allows us not to worry about changes of presentations in relation to linear isoperimetric inequalities. Note that we are not able to give an explicit value for the linear isoperimetric constant of the new presentation, and that when we will try do to it in further discussion, it will be a special case. For a more general discussion about isoperimetric inequalities the reader is directed to [Gerst1991].

Corollary 2.45. Let $\langle X|R\rangle, \langle Y|S\rangle$ be two finite presentations for the same group and suppose that $\langle X|R\rangle$ satisfies a linear isoperimetric inequality. Then also $\langle Y|S\rangle$ satisfies a linear isoperimetric inequality.

Proof. We know that there is a finite sequence of Tietze transformations taking $\langle X|R\rangle$ to $\langle Y|S\rangle$ (Theorem 2.24). Each time we get a presentation satisfying a linear isoperimetric inequality by Theorem 2.44. So in particular $\langle Y|S\rangle$ satisfies a linear isoperimetric inequality. $\square$

3 Geometric definitions

In this section we give the geometric definitions that we are going to need in order to be able to talk about hyperbolic groups. We will look at metric spaces and see how it is possible to turn a group into a metric space, with the group itself acting on it by isometries. Following this we are going to define the notion of geodesic metric spaces and see how the metric space constructed for the group is an example of such a space.

We start by determining how we can give a consistent notion of distance to a group with respect to some finite generating set. Most definitions are taken from [Br1999] and [BrHae2011]. Examples are original unless otherwise stated.

Definition 3.1. [Br1999] Given a group G with finite generating set X we can endow it with a left-invariant metric:

$$d(g,h) = \inf\{|w| : w \in F_X, w =_G g^{-1}h\}.$$

Note that by left-invariant metric we mean that the group acts on itself by isometry when using the action of multiplication on the left of an element. This happens because $w =_G g^{-1}h \Leftrightarrow w =_G (tg)^{-1}th$ for all $t \in G$. Also note that the action of multiplication on the right is not necessarily an isometry, since $w =_G g^{-1}h \not\Leftrightarrow w =_G (gt)^{-1}ht$. The following two examples cover a finite and an infinite group.

Example 3.2. The group $S_3 = \langle \sigma, \rho | \sigma^2, \rho^3, \sigma\rho\sigma\rho \rangle$ has $d(\sigma, \rho) = 2$ since $\sigma^{-1}\rho = \sigma\rho$ is not a generator but is a product of two.

Example 3.3. Given any two elements $x = a^{n_1}b^{m_1}, y = a^{n_2}b^{m_2}$ of the group $\mathbb{Z}^2 = \langle a, b\rangle$, we can compute their distance by $d(x,y) = |n_1 - n_2| + |m_1 - m_2|$.

We now introduce the notion of a metric graph, which is essentially a graph turned into a metric space by giving a length to each edge and by seeing each edge as isometric to a closed interval. The part that requires carefulness is the fact that we want the points on the edges to be part of the metric space.

Definition 3.4. [Har1969]. A *metric graph* is a connected graph (V, E) where each edge $e \in E$ is isometric to an interval $[0, L_e]$ (via a map ϕ_e), for $L_e \in \mathbb{R}^+$, with the ends of e corresponding to $0, L_e$. For two points x, y on edges e_x, e_y in the graph, where points along the edges are allowed, the distance $d(x, y)$ is defined as

$$min\{d(v_x, v_y)+|\phi_{e_x}(v_x)-\phi_{e_x}(x)|+|\phi_{e_y}(v_y)-\phi_{e_y}(y)| : v_x \in e_x \cap V, v_y \in e_y \cap V\},$$

where $d(v_x, v_y)$ is simply the length of the shortest path between the vertices v_x and v_y with each edge e on the path counting L_e.

Example 3.5. Consider the following simple metric graph on vertices $V = \{v_1, \dots, v_8\}$, where each edge e is labelled by its length L_e. Consider also a point x on the edge $v_5 v_6$ at distance 1.2 from v_5 and the point $y = v_8$.

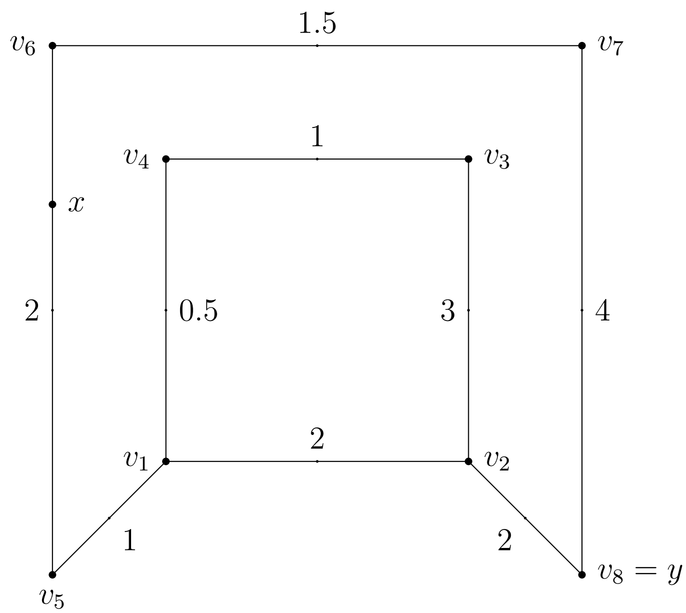

Figure 3.1

In order to find the distance between x and y we first need to find the distances $d(v_5, v_8)$ and $d(v_6, v_8)$ which can be easily seen are 5 and 5.5 respectively. We then check which number is smaller between $d(v_5, v_8) + d(x, v_5)$ and $d(v_6, v_8) + d(x, v_6)$ to determine that $d(x, y) = 6.2$ obtained by going from x to v_5 to v_1 to v_2 to y.

In particular notice that even though $d(x, v_6) < d(x, v_5)$ the shortest path is through v_5.

We are now ready to turn a group into a metric space, extending the previously defined distance between two elements of the group with respect to some generating set.

Definition 3.6. The *Cayley graph* of a group G with respect to a finite generating set X, denoted by $\Gamma(G, X)$, is the metric graph whose vertices are in $1-1$ correspondence with the elements of G and which has a directed edge of length 1 from g to gx, labelled x, for each $g \in G, x \in X$.

Note that we require metric graphs to be connected, so we need to check that our definition of Cayley graph produces a connected graph.

Lemma 3.7. Cayley graphs are connected.

Proof. This follows simply from the fact that for any g, h in a group $G = \langle X \rangle$ we have that $gh^{-1} \in G$ can be written in terms of the generators in X and there is therefore a path between g, h in the Cayley graph. □

Notice that if a generating set is symmetric (i.e. contains all its inverses and does not contain the identity), then for every directed edge there is one going in the opposite direction with inverted label. Also, since in a symmetric set we don't have the identity, the graph does not have any loops.

In general when we talk about a Cayley graph we will think about the undirected graph induced by the directed one, i.e. the graph with an edge between two vertices if and only if there is at least one directed edge between them. We will use the directed graph when thinking about paths and their labels.

Example 3.8. The Cayley graph of a free group F with respect to a free basis X is a tree, since it is a connected graph and there is no reduced word in $W(X)$ of positive length which is equal to the identity, and hence no circuit.

We now define a particular class of metric spaces.

Definition 3.9. A metric space (X, d) is called *geodesic* if for all $x, y \in X$ there is a path of length $d(x, y)$ between x and y, i.e. an isometry $f : [0, d(x, y)] \to X$ such that $f(0) = x$ and $f(d(x, y)) = y$. Any such path is called a *geodesic* and is denoted by $[xy]$.

Note that in general we will use $[xy]$ to denote one particular geodesic, even if there is more than one.

Definition 3.10. We say that the *euclidean plane* $\mathbb{E}^2$ is the metric space given by $\mathbb{R}^2$ equipped with the standard metric.

It is not difficult to find examples of geodesic or non-geodesic metric spaces.

Example 3.11. $\mathbb{E}^2$ is a geodesic metric space since the segment between two points is a geodesic path. If we remove a single point, for example $(0, 0)$, then we don't have a geodesic metric space any more, because there is no path of length 2 between $(-1, 0)$ and $(1, 0)$. Another example of a non-geodesic metric space is the set $\{0, 1\}$ equipped with the discrete metric. Clearly there is no path isometric to $[0, 1]$ between 0 and 1 since there are no points in the middle. In general, any countable metric space with more than one point cannot be geodesic.

Note that in the euclidean plane there is a unique line segment joining two points, and therefore geodesics are unique. In general this is not the case, one example being the plane with taxicab distance.

Example 3.12. The Cayley graph of a group G with respect to a generating set X is a geodesic metric space.

We will often want to consider the distance between two geodesics instead of the distance between two points.

Definition 3.13. Let (X, d) be a geodesic metric space. Let α, β be geodesics in X. We define the distance between α, β, denoted by $d(\alpha, \beta)$ to be

$$d(\alpha, \beta) =_{def} \inf_{x \in \alpha, y \in \beta} d(x, y).$$

The following proposition will turn out to be useful when considering the distance between two geodesics. It basically states that the infimum of pairwise distances is realised, and is original work.

Proposition 3.14. Let (X, d) be a geodesic metric space. Let α, β be geodesics in X. Then we can find points $a \in \alpha$ and $b \in \beta$ such that

$$d(a, b) = d(\alpha, \beta)$$

Proof. First of all we notice that by definition geodesics are continuous images of compact sets, and are therefore compact. In particular α and β are compact sets in (X, d).

It is clear that $d(\alpha, \beta) = \inf_{x\in\alpha} d(x, \beta)$. We therefore consider a point $x \in \alpha$. The distance between x and β is $\inf_{y\in\beta} d(x, y)$. We can therefore construct a sequence $(y_n)_n \in \beta$ such that $d(x, y_n) \to d(x, \beta)$. By compactness of β this sequence has a convergent subsequence and we get a point $y_x \in \beta$ such that $d(x, y_x) = d(x, \beta)$.

We now notice that $d(\alpha, \beta) = \inf_{x\in\alpha} d(x, \beta) = \inf_{x\in\alpha} d(x, y_x)$. Construct a sequence $(x_n)_n$ such that $d(x_n, y_{x_n}) \to d(\alpha, \beta)$. Using compactness two more times on the sequences $(x_n)_n$ and $(y_{x_n})_n$ we get two limit points $a \in \alpha$ and $b \in \beta$ such that $d(a, b) = d(\alpha, \beta)$. □

From now on we will use the above proposition without worrying much about it, so we will freely move from distance between geodesics to pairs of points realising the distance.

We will naturally want to consider more than one geodesic path at once, the following definition agrees with our intuition about polygons in the euclidean plane.

Definition 3.15. A *geodesic n-sided polygon* $x_1x_2 \dots x_n$ in a geodesic metric space (X, d) is the union of geodesics $[x_1x_2], [x_2x_3], \dots, [x_nx_1]$.

We wish to consider geodesic polygons where the sides do not cross more than they are supposed to, hence the following definition.

Definition 3.16. A geodesic triangle xyz in a geodesic metric space is called *non-degenerate* if $x \neq y \neq z \neq x$ and $[xy] \cap [xz] = \{x\}, [yz] \cap [xz] = \{z\}, [xy] \cap [yz] = \{y\}$. It is called *degenerate* otherwise.

In some geodesic metric spaces it is actually not possible to find a non-degenerate triangle.

Example 3.17. Any geodesic triangle in the Cayley graph of a free group with respect to a free basis is degenerate since the Cayley graph of a free group is a tree and if a triangle was non-degenerate it would form a circuit.

4 Hyperbolicity, a first definition

In this section we give a first characterisation of hyperbolic metric spaces and hyperbolic groups. We start from geodesic metric spaces and we define a condition on geodesic triangles that will give a first definition of hyperbolicity (slim triangles). We will then introduce the notion of thin triangles which will turn out to be an equivalent characterisation. Most of the definitions come from [Alonso et al. 1990], unless stated.

Definition 4.1. Let $\delta \in \mathbb{R}_{\geq 0}$. A geodesic triangle xyz in a geodesic metric space (X, d) is called *δ-slim* if for any point $w \in [xy]$ we have that $min(d(w, [xz]), d(w, [yz])) \leq \delta$ and if this holds similarly for all points $w \in [xz], w \in [yz]$.

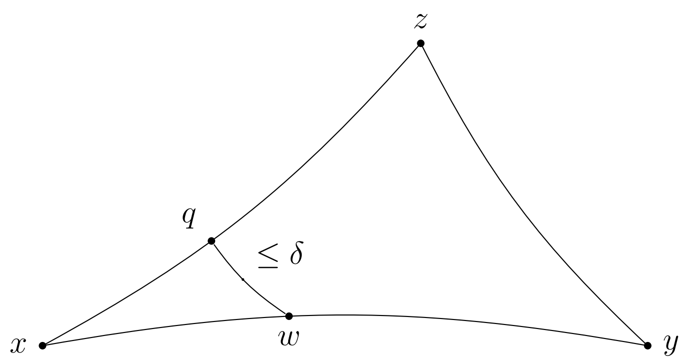

Figure 4.1

We say that *triangles are slim* in X if there is a constant δ such that all geodesic triangles in X are δ-slim.

Here we can see an example of a slim triangle in the euclidean plane and then a more interesting one involving the hyperbolic disc in which triangles are slim.

Example 4.2. An equilateral triangle of side length 1 in the euclidean plane is $\frac{\sqrt{3}}{4}$−slim but not $\frac{1}{3}$−slim since the distance between the middle point of one side and the other two sides is exactly $\frac{\sqrt{3}}{4} > \frac{1}{3}$ (it can be easily seen that this is the limit case). Note that for any euclidean triangle we can find some constant δ such that our particular triangle is δ-slim, but there is no δ such that all triangles are δ-slim because the triangle $(-n, 0)(n, 0)(0, 2n)$ is for example not $\frac{n}{4}$-slim.

Example 4.3. Triangles in the open unit disk in $\mathbb{C}$ equipped with the Poincaré metric are $\log(1 + \sqrt{2})$-slim.

The proof of the example above is omitted, since it would require us to give all the necessary definitions of hyperbolic disc. We are now ready to give a first definition of a hyperbolic metric space, in terms of slim triangles.

Definition 4.4. Let (X, d) be a geodesic metric space. We say that the space X is
(δ-)*hyperbolic* if triangles are (δ-)slim in X.

Example 4.5. From our previous example we get that the hyperbolic disc is a hyperbolic metric space. In general we also have that every bounded geodesic metric space is hyperbolic. This is true because if we pick any point p and a constant δ such that every point is within δ of p we have that every geodesic triangle is definitely 2δ-hyperbolic (the distance between any two points is at most 2δ).

Groups are what we are really interested in, and in order to give a definition of what a hyperbolic group is, we need to go through the Cayley graph. The first problem that arises is that Cayley graphs are defined in terms of a specific generating set. After our discussion we will be able to observe that the definition is in fact independent of the generating set (Corollary 8.2).

Definition 4.6. Let G be a group generated by a finite set X. We say that G is *hyperbolic* with respect to the generating set X if the Cayley graph $\Gamma(G, X)$ is hyperbolic.

Example 4.7. Any finite group is hyperbolic with respect to any generating set since we can always find the maximum distance between all pairs of possible elements and that is a constant δ which makes the Cayley graph δ-slim.

Example 4.8. Any free group is hyperbolic with respect to a finite free basis since the Cayley graph is a tree and, as seen before, any point in a geodesic triangle actually belongs to at least two sides, making the Cayley graph 0-slim.

In what follows we take a slightly different approach from the one in [Alonso et al. 1990], in the sense that definitions for internal points and thin triangles are given without the use of euclidean comparison triangles. The definitions and propositions are original.

Definition 4.9. Given a geodesic triangle xyz in a geodesic metric space (X, d) we define the *internal points* c_x, c_y, c_z of xyz as points such that $c_x \in [yz], c_y \in [xz], c_z \in [xy]$ and $d(y, c_x) = d(y, c_z), d(x, c_z) = d(x, c_y), d(z, c_x) = d(z, c_y)$.

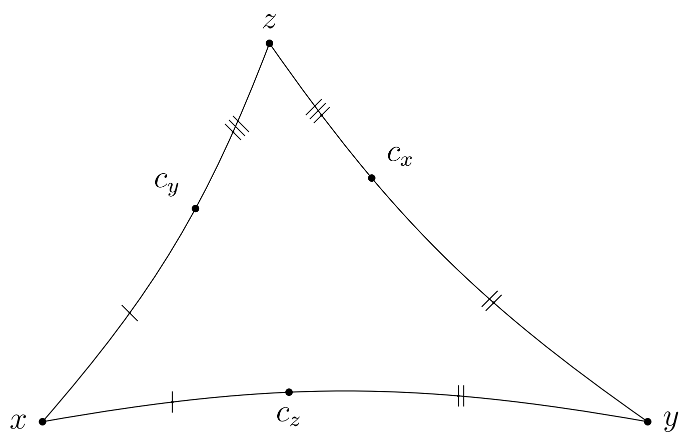

Figure 4.2

It is not clear that such points even exist, but the following proposition shows that they do.

Proposition 4.10. Given a geodesic triangle xyz in a geodesic metric space (X, d) the internal points exist and are unique.

Proof. Let $m = \min\{d(y, x), d(y, z)\}$ and let $t \in [0, m]$.

We consider the unique pair of points $p_t \in [xy], q_t \in [yz]$ both at distance t from y. Let $D_t = d(x, p_t) + d(z, q_t)$. If we increase t continuously from $t = 0$, we continuously decrease D_t. Since $D_0 = d(x, y) + d(y, z) \geq d(x, z)$ and $D_m = |d(y, z) - d(y, x)| \leq d(x, z)$, by the Intermediate Value Theorem we get that there must exist a value $t' \in [0, m]$ such that $D_{t'} = d(x, z)$.

Fix $c_z = p_{t'}, c_x = q_{t'}$ and define c_y as the unique point on $[xz]$ such that $d(x, c_y) = d(x, c_z)$. Then $d(z, c_y) = d(x, z) - d(x, c_y)$ and by choice of t' we have that

$$d(x, z) = d(x, c_z) + d(z, c_x) = d(x, c_y) + d(z, c_x).$$

Hence we have that $d(z, c_y) = d(x, c_y) + d(z, c_x) - d(x, c_y)$ and we get as well that $d(z, c_y) = d(z, c_x)$. Uniqueness follows since the conditions required in the construction are necessary. $\square$

The following proposition gives a more technical characterisation of the internal points and will be useful to shorten some proofs later on.

Proposition 4.11. Given a geodesic triangle xyz in a geodesic metric space (X, d), the points $c_x \in [yz], c_y \in [xz], c_z \in [xy]$ are the internal points of xyz if and only if for any permutation (i, j, k) of the vertices $\{x, y, z\}$ and any points $p \in [ij], q \in [ik]$, we have $d(p, c_k) \leq d(c_i, p)$ and $d(q, c_j) \leq d(c_i, q)$.

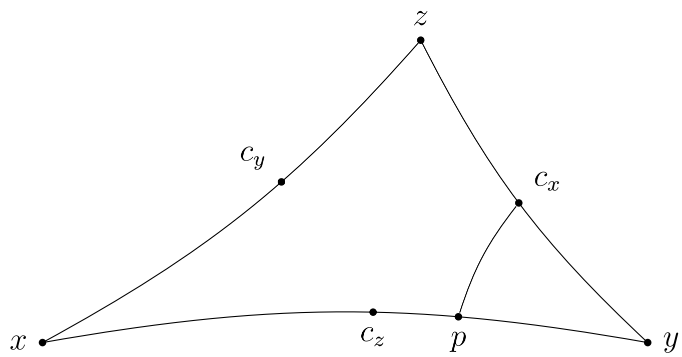

Figure 4.3

Proof. ($\Rightarrow$). It suffices to show that in the case of Figure 4.3 we have $d(c_z, p) \leq d(p, c_x)$. The rest follows similarly.

By the triangle inequality we have that $d(p, y) \geq d(c_x, y) - d(p, c_x)$. Also $d(c_z, p) = d(c_z, y) - d(p, y)$. Hence $d(c_z, p) \leq d(c_z, y) + d(p, c_x) - d(c_x, y)$. Now since $d(y, c_x) = d(y, c_z)$ we get that $d(c_z, p) \leq d(p, c_x)$.

($\Leftarrow$). If we pick the point p to be x and $i = z$ we get that $d(x, c_y) \leq d(x, c_z)$. If we instead set $i = y$ we get $d(x, c_y) \geq d(x, c_z)$. Hence $d(x, c_y) = d(x, c_z)$. This is true for all other pairs as well. $\square$

We are now ready to give the definition of thin triangles, which will depend on the internal points.

Definition 4.12. Let (X, d) be a geodesic metric space and let xyz be a geodesic triangle in X. Let c_x, c_y, c_z be the internal points of xyz and let $\delta \in \mathbb{R}_{\geq 0}$.

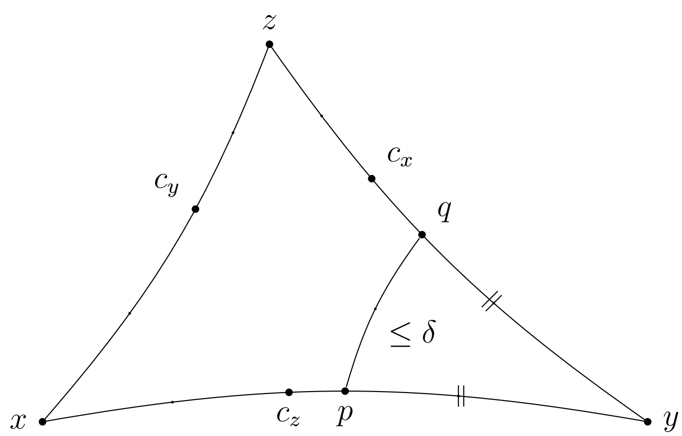

Figure 4.4

We say that the triangle xyz is δ-*thin* if for all permutations (i, j, k) of (x, y, z) we have that for all pairs of points $p \in [kc_i], q \in [kc_j]$ such that $d(p, k) = d(q, k)$, $d(p, q) \leq \delta$. We say that *triangles are thin* in X if there is a constant δ such that all geodesic triangles in X are δ-*thin*.

We give the final definition of the section before showing, in the next section, that triangles are slim if and only if they are thin.

Definition 4.13. Let (X, d) be a geodesic metric space and let xyz be a geodesic triangle in X. We denote by $insize(xyz)$ the maximum of $d(c_x, c_y), d(c_x, c_z)$ and $d(c_y, c_z)$.

5 Equivalence of geometric definitions

In this section our aim is to show that we can give equivalent definitions of what a hyperbolic group is, using thin triangles and insize. We will then want to be able to characterise hyperbolic groups avoiding the use of triangles, and in the next section we are going to extend the set of definitions to include linear isoperimetric inequalities and Dehn presentations.

We are in particular also interested in optimising the constants that appear when proving these equivalences.

Proposition 5.1. Let (X, d) be a geodesic metric space and let $\delta \in \mathbb{R}_{\geq 0}$. If all geodesic triangles in X are δ-slim, then the insize of all geodesic triangles is at most 4δ.

Proof. Let xyz be a geodesic δ-slim triangle with internal points c_x, c_y, c_z. Since xyz is δ-slim, the point c_x is at distance at most δ from $[xy] \cup [xz]$. Without loss of generality assume $d(c_x, [xy]) \leq \delta$. Let p be a point on $[xy]$ such that $d(p, c_x) \leq \delta$.

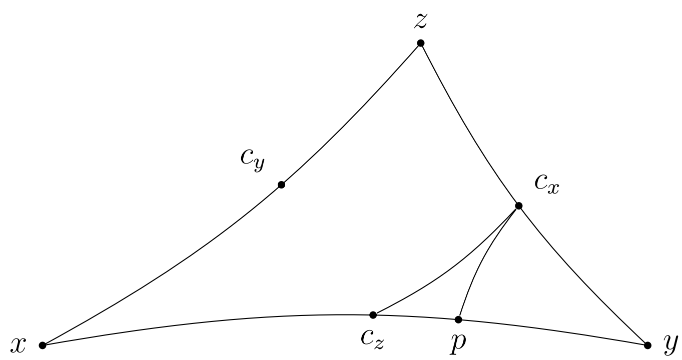

Figure 5.1

By Proposition 4.11 we immediately get that $d(p, c_z) \leq \delta$ and by the triangle inequality $d(c_x, c_z) \leq 2\delta$.

The same reasoning applies to c_y and either $d(c_y, c_x) \leq 2\delta$ or $d(c_y, c_z) \leq 2\delta$. Hence, by the triangle inequality, $\max\{d(c_x, c_y), d(c_x, c_z), d(c_y, c_z)\} \leq 4\delta$. □

We are now going to use Prop 5.1 to go from slim to thin. The following proof can be found in [Alonso et al. 1990]. We are then going to provide an alternative proof that improves the bound from 6δ to 4δ.

Proposition 5.2. [Alonso et al. 1990]. Let (X, d) be a geodesic metric space and let $\delta \in \mathbb{R}_{\geq 0}$. If all geodesic triangles in X are δ-slim then they are also 6δ-thin.

Proof. Let xyz be a geodesic δ-slim triangle with internal points c_x, c_y, c_z. Let $p \in [xc_z], q \in [xc_y]$ be points equally distant from x. We want to show that $d(p, q) \leq 6\delta$. If either $d(p, [xc_y]) \leq \delta$ or $d(q, [xc_z]) \leq \delta$ we can, without loss of generality, consider a point $t \in [xc_y]$ such that $d(p, t) \leq \delta$. Since $d(x, q) = d(x, p)$ we can use the same argument we used to prove Proposition 5.1 to say that $d(q, t) \leq d(p, t) \leq \delta$ and by the triangle inequality $d(q, p) \leq 2\delta$.

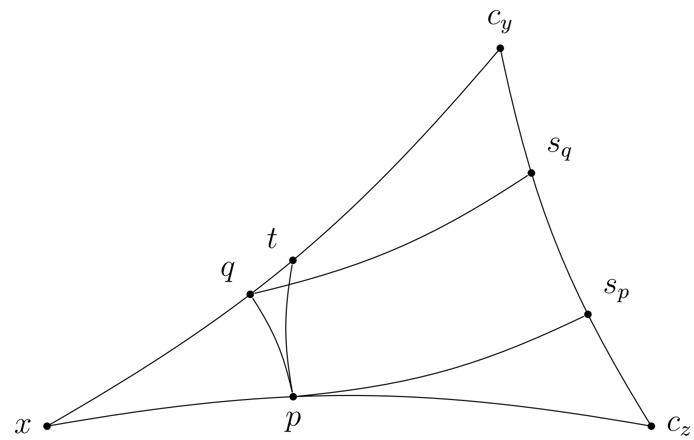

Figure 5.2

Otherwise, since the geodesic triangle xc_zc_y is δ-slim, we can find two points $s_p, s_q \in [c_zc_y]$ such that $d(p, s_p), d(q, s_q) \leq \delta$. By Proposition 5.1 $d(c_z, c_y) \leq 4\delta$ and since $d(s_p, s_q) \leq d(c_z, c_y)$ we get that $d(p, q) \leq 6\delta$, by considering the path qs_qs_ppq of length at most 6δ. This completes the proof since the same reasoning applies to pairs of points on the other segments. $\square$

For the alternative proof we use a construction similar to the one we used to show that internal points exist.

Proposition 5.3. Let (X, d) be a geodesic metric space and let $\delta \in \mathbb{R}_{\geq 0}$. If all geodesic triangles in X are δ-slim then they are also 4δ-thin.

Proof. Let xyz be a geodesic δ-slim triangle with internal points c_x, c_y, c_z. Let $p \in [xc_z], q \in [xc_y]$ be points equally distant from x. We want to show that $d(p, q) \leq 4\delta$. The idea is to simply show that the points p, q are two of the internal points of some geodesic triangle and by Proposition 5.1 we are done.

To do this consider points $a_t \in [qz]$ that vary continuously from $a_0 = z$ to $a_1 = q$ as t goes from 0 to 1 in $[0, 1]$. Each time we consider a geodesic triangle Δ_t on the points x, y, a_t with $[xa_t]$ and $[ya_t]$ subpaths of the original triangle xyz. The length of the sides of Δ_t must also vary continuously and we can therefore say that the internal points of Δ_t move continuously. In particular the internal point of Δ_t that lies on $[xa_t]$ goes from c_y when $t = 0$ to a point on $[xq]$ when $t = 1$. Hence by the Intermediate Value Theorem we have that there is $t' \in [0,1]$ such that the internal point of Δ'_t on $[xa_{t'}]$ is precisely q.

But then we have that q and p are the internal points of $\Delta_{t'}$ and since also $\Delta_{t'}$ is δ-slim, by Proposition 5.1 there is a bound of 4δ on the insize. This implies that in particular $d(q,p) \leq 4\delta$.

□

To conclude the circle of equivalences we prove the following proposition, which follows easily from the definition of thinness.

Proposition 5.4. Let (X,d) be a geodesic metric space and let $\delta \in \mathbb{R}_{\geq 0}$. If all geodesic triangles in X are δ thin then they are also δ-slim.

Proof. Let xyz be a geodesic triangle and consider a point $p \in [xy]$. If c_x, c_y, c_z are the internal points of the triangle xyz, then p either lies on $[xc_z]$ or on $[yc_z]$. Without loss of generality assume that $p \in [xc_z]$. Then pick a point q on $[xc_y]$ such that $d(p,x) = d(q,x)$. Note that this point must exist since $d(x,c_z) = d(x,c_y)$ By definition of thinness we must have that $d(p,q) \leq \delta$. □

We can finally state three equivalent characterisations of hyperbolicity.

Theorem 5.5. Let (X,d) be a geodesic metric space. Then the following are equivalent:

(i) X is hyperbolic.

(ii) There is a bound on the insize of geodesic triangles.

(iii) Triangles are thin in X.

Proof. This follows directly from Propositions 5.1, 5.3 and 5.4. □

6 Disc diagrams

In this section we are going to give the definition of disc diagrams (also called Cayley diagrams) and talk about some of the most important results concerning them. Disc diagrams will be essential in extending our definitions of hyperbolicity since they put together geometric and purely algebraic properties. The theory is completely taken from [LynSch2002, Chapter 5.1].

Definition 6.1. Let $S \subseteq \mathbb{E}^2$. We denote by δS the boundary of S and by $\overline{S}$ the closure.

We start by defining maps, based on which we will be able to define disc diagrams.

Definition 6.2. A *map* D is a finite collection of vertices, edges and regions where a *vertex* is a single point in $\mathbb{E}^2$, an *edge* is a bounded subset of $\mathbb{E}^2$ homeomorphic to $(0,1)$ and a *region* is a bounded set homeomorphic to the open unit disk, with the following properties:

(i) If e is an edge of D, there are vertices a, b in D such that $\overline{e} = e \cup \{a\} \cup \{b\}$.

(ii) If A is a region of D, the boundary δD of A is connected and there are edges $e_1, .., e_n$ in D such that $\delta A = \overline{e_1} \cup \cdots \cup \overline{e_n}$.

It is useful to visualise maps as graphs where some cycles are the boundary of regions.

Example 6.3. If we consider just the x-axis in $\mathbb{E}^2$, the edge $(0,1)$ together with the vertices $\{0\}$ and $\{1\}$ form a map with no regions. The open unit disk together with the edge consisting of the unit circle without the point $(1,0)$ and with the vertex $(1,0)$ form a map with one region.

We now give some definitions that connect the notion of a map to the one of a directed graph. In order to do this we give orientations to the edges and define paths. The conventions are the same as in graph theory.

Definition 6.4. Given an edge e in a map D the vertices a, b such that $\overline{e} = e \cup \{a\} \cup \{b\}$ are the *endpoints* of e. A *closed edge* is an edge together with its endpoints.

Any edge can be *traversed* in 2 directions. Given an oriented edge e from vertex a to b, we denote a by $s(e)$ and b by $t(e)$. We denote the oriented edge from b to a by e^{-1}. We denote the set of oriented edges as $E(D)$.

Definition 6.5. A *path* γ of length $n \in \mathbb{N}_0$ in a map D is a sequence of oriented closed edges $e_1, \ldots, e_n$ such that $t(e_i) = s(e_{i+1})$ for $1 \leq i \leq n-1$. We allow the empty path, which corresponds to $n = 0$. We denote $s(e_1)$ by $s(\gamma)$ and $t(e_n)$ by $t(\gamma)$. We denote the length n by $l(\gamma)$.

A *closed path* or *cycle* is a path $e_1 \ldots e_n$ with $s(e_1) = t(e_n)$ and is *reduced* if it does not contain a successive pair of edges of the form ee^{-1}. A reduced path is *simple* if for all $i \neq j$, $s(e_i) \neq s(e_j)$.

We have said that any edge can be trasversed in two directions, one per directed edge. Now that we have given an orientation to edges we also assign an orientation to the regions. An oriented region will then tell us in which direction we should traverse boundary paths. Given a boundary of a region and a starting vertex on the boundary we can then traverse the entire boundary in one of two directions, by picking one of the two edges on the boundary that are adjacent to the vertex as the starting edge of a closed path on the boundary. For this reason we say that a region can be oriented in one of two directions (clockwise and anti-clockwise).

Definition 6.6. If A is an oriented region in a map D we say that a cycle is a *boundary cycle* of A if it is a cycle of minimal length which includes all the edges of δA and its edges are oriented in accordance to A.

We want the regions to be consistently oriented in the following way (original proof).

Proposition 6.7. Let D be a map in $\mathbb{E}^2$. We can orient each region of D and each component of $\mathbb{E}^2 \setminus D$ in a way that traversing the boundary of every region and every component in the given orientation, each edge is traversed precisely once in each of the two directions.

Proof. We prove the result by induction on the number of regions of D plus the number of components of $\mathbb{E}^2 \setminus D$. The base case occurs when there is only one component of $\mathbb{E}^2 \setminus D$ and no regions. In this case the map D is a graph with no circuits and it is clear that either orientation of $\mathbb{E}^2 \setminus D$ does the job.

Now assume that the result holds up to the case with $N \in \mathbb{N}$ regions and components of $\mathbb{E}^2 \setminus D$ and consider a map D with $N+1$ regions and components of $\mathbb{E}^2 \setminus D$. There must exist an edge $e \in D$ such that e lies on the boundary of either two regions or one region and one component or two components, which we call A, B. We now construct a map D' by removing the edge e and defining $C = A \cup B$ as a component and not a region. We then have reduced the number of regions and components of $\mathbb{E}^2 \setminus D'$ to N.

We can then use the inductive assumption to orient all the regions and the components of D' in a way that traversing the boundary of every region and every component in the given orientation each edge is traversed precisely once in each of the two directions. Now we can use this orientation for the map D, for everything except A, B. If we orient A, B both in the same way as C, i.e. both clockwise or anti-clockwise as C, we have a consistent orientation as required. $\square$

We want to have an definition of boundary cycle that works for a whole map, but in order for this to work we need to give an ordering to the edges with a vertex in common.

Definition 6.8. Let v be a vertex in a map D. We assign a cyclic ordering to the edges of the map D starting at v in the most natural way: two edges are adjacent in the cyclic ordering if and only if they form at least one angle θ at v which is minimal, in the sense that no other pair of edges forms an angle completely contained in θ.

We will be working solely with maps that are connected and simply connected. We are hence interested in the following definition.

Definition 6.9. If a map D is connected and simply connected then we say that a cycle α is a *boundary cycle* of D if it is a cycle of minimal length which contains all the edges of δD and does not cross itself, in the sense that if e_i, e_{i+1} are two consecutive edges in α with common vertex v then e_i^{-1}, e_{i+1} are adjacent in the set of all edges of D starting at v.

We are finally ready to define (disc) diagrams, the fundamental tool that we will use throughout the next section. The idea basically consists of taking a map and assigning names to all oriented edges using elements from a group. It then makes sense to require oppositely oriented edges to have inverse names.

Definition 6.10. A *diagram* over a group G is a map D together with a function $\phi : E(D) \to G$ such that $\phi(e) = \phi(e^{-1})^{-1}$ for all $e \in E(D)$. For a path $\alpha = e_1 \dots e_k$ we define $\phi(\alpha) = \phi(e_1) \dots \phi(e_k)$. If α is a boundary cycle of a region A we say that $\phi(\alpha)$ is a *label* of A.

If γ is a boundary cycle of the diagram D we say that $\phi(\gamma)$ is a *label* of D.

Definition 6.11. Given a diagram D we denote by $A(D)$ the number of regions of D and by $D^{(1)}$ the subdiagram of D consisting of all the edges and vertices of D but no regions. We say that $D^{(1)}$ is the *1-skeleton* of D.

We can extend the notion of length of a path to a set of paths.

Definition 6.12. Given a diagram D and a subset $C \subseteq E(D)$, we denote by $l(C)$ the number of edges in C.

The reason we want to use disc diagrams is that they can give us a way of identifying an element of a group with the boundary of a disc diagram and at the same time get information from the regions of the diagram. We can start to see this with the following proposition, which considers elements of a free group.

Proposition 6.13. Given a free group F and a finite sequence $(c_1, \ldots, c_n)$ of non-trivial elements in F, we can construct a connected and simply connected diagram D over F which satisfies the following conditions:

(i) $\phi(e) \neq 1$ for all $e \in E(D)$.

(ii) D has a vertex O on δD and a boundary cycle $\alpha = e_1 \ldots e_m$ beginning and ending at O such that $\phi(\alpha) =_F c_1 \ldots c_n$ and the word $\phi(e_1) \ldots \phi(e_m)$ is freely reduced.

(iii) If $\beta = e_1 \ldots e_k$ is a boundary cycle of a region A in D then $\phi(\beta)$ is a cyclically reduced conjugate of some c_i.

The above proposition can be proven by induction on the length of the sequence, but requires a lot of work to fully appreciate it, so we will not present it here, since it is only tangential to our main concerns. The idea is to start with a vertex O from which we have as many cycles as the length of the sequence of c_i's, each representing the word c_i.

The following proposition links the boundary of a connected and simply connected diagram over a free group to the boundary of the regions that make up the diagram. We will give an outline of the proof.

Proposition 6.14. Given a connected and simply connected diagram D over F with regions $A_1, \ldots, A_n$ and boundary cycle α there exist labels $r_1, \ldots, r_n$ where r_i is a label of A_i, and elements $u_1, \ldots, u_n$ of F such that

$$\phi(\alpha) =_F u_1^{-1} r_1 u_1 \ldots u_n^{-1} r_n u_n.$$

Proof. We use induction on the number of regions. The base case is trivial. Assume that the result holds for diagrams with $N \in \mathbb{N}$ regions and consider a connected and simply connected diagram D with $N+1$ regions. The idea is to construct a connected and simply connected diagram D' with N regions by removing an edge on the boundary cycle of D. We can then use the inductive assumption on D' and by choosing an appropriate u_{N+1} we get the required result. □

By definition the boundary cycle of a given region in a diagram is not unique. It makes sense to define a set of elements of a free group that is closed under cyclic conjugation.

Definition 6.15. We say that a subset R of a free group F is *symmetrized* if every $r \in R$ is cyclically reduced and $r \in R$ implies that all cyclic conjugates of r and r^{-1} are in R.

Example 6.16. We can easily symmetrize any set of relators of a given presentation. For example if we take the standard presentation for $S_3 = \langle a, b | a^2, b^2, (ab)^3 \rangle$ we need to add in the relators $a^{-2}, b^{-2}, (b^{-1}a^{-1})^3, (ba)^3, (a^{-1}b^{-1})^3$ and we have a symmetrized set of relators.

We now start to link diagrams to presentations, and in particular to the word problem. The idea is to use the relators as boundary cycles, so that the boundary cycle of the diagram is an element in the word problem.

Definition 6.17. Let R be a symmetrized subset of a free group F. We say that a diagram D over F is an *R-diagram* if any label of any region of D is in R.

The next proposition is the key result of the section and we are going to use repeatedly.

Proposition 6.18. Let R be a symmetrized subset of a free group F and $w \in F$. Then $w \in \langle\langle R \rangle\rangle$ if and only if there exists a connected and simply connected R-diagram D with label equal to w in F.

Proof. ($\Rightarrow$). The fact that $w \in \langle\langle R \rangle\rangle$ implies that there exist elements $c_1, \ldots, c_n \in F$ such that $w =_F c_1 \ldots c_n$, with $c_i =_F u_i^{-1} r_i u_i$ for elements $u_1, \ldots, u_n \in F$ and $r_i \in R$. Applying Proposition 6.13 we get the required connected and simply connected R-diagram.
($\Leftarrow$). This direction follows straightforwardly from Proposition 6.14 since by definition an R-diagram has all region labels in R. □

We now know that if we are given a finite presentation $\langle X | R \rangle = G$ with a symmetrized set of relators R, and a word $w \in WP(G, X)$, there exists a connected and simply connected R-diagram with label w. We will be interested in finding a diagram with the smallest possible number of regions and for this we give the following definition.

Definition 6.19. An R-diagram with boundary word w is called *minimal* if there is no other R-diagram with the same label and with fewer regions.

Note that if we are given a minimal R-diagram for a word w with a certain number N of regions, we know that it possible to write w as a product of N conjugates of relators, but not $N-1$, so N is the optimal number that we can use in an isoperimetric inequality. We hence see the connection between minimal R-diagrams and isoperimetric inequalities, which will be central in our discussion.

From now on we will assume that the sets of relators in our presentations are symmetrized, in order to make full use of R-diagrams.

7 Hyperbolicity, equivalent characterisations

In this section we prove that we can give an equivalent characterisation of hyperbolicity via the linear isoperimetric inequality and Dehn presentations. The section contains original work. In particular we fix a problem that was found in one of the proofs in [Alonso et al. 1990] (Corollary 7.6 in our text). Furthermore we pay more attention to the constants in the various proofs, reducing the δ in Theorem 7.12 by a factor of more than 7.

We start with a technical lemma about cutting a geodesic triangle which we will be fundamental later on. The lemma is a slightly modified version (the ball centred at w used to be smaller) of a lemma in [Alonso et al. 1990].

Lemma 7.1 (Cutting the triangle)**.** Let $G = \langle X|R\rangle$ be a finitely presented group and let $\Gamma = \Gamma(G, X)$ be a Cayley graph of G. Let $r \in \mathbb{R}, \lambda \in \mathbb{N}$ with $r > 9(\lambda + 1)$.

Let xyz be a non-degenerate geodesic triangle in Γ with $x, y, z \in V(\Gamma)$. Let $w \in [xy] \cap V(\Gamma)$ and assume that $B_{\frac{10}{9}r}(w) \cap [xz] = B_{\frac{10}{9}r}(w) \cap [yz] = \emptyset$. Let w', w'' be points respectively on $[xw]$ and $[wy]$ at distance r from w. Then we can do one of the following.

Case 1 (hexagon or pentagon). We can find points $x', x'', y', y'', z', z'' \in V(\Gamma)$ such that:

(i) $x' \in [xw']$; $y' \in [w''y]$; $y'', z'' \in [yz]$; $z', x'' \in [xz]$;

(ii) $d(x', x'')$, $d(y', y'')$, $d(z', z'') = \lambda$;

(iii) $d([x'y'], [y''z'']) = d([x'y'], [z'x'']) = d([y''z''], [z'x'']) = \lambda$.

Case 2 (quadrilateral). We can find points $x', z'', y', y'' \in V(\Gamma)$ such that:

(i) $x' \in [xw']$; $y' \in [w''y]$; $z'', y'' \in [yz]$ or $z'', y'' \in [xz]$;

(ii) $d(x', z'')$, $d(y', y'') = \lambda$;

(iii) $d([x'y'], [z''y'']) = \lambda$.

Proof. In this proof we try to construct a full hexagon, but in some cases this is not possible and we end up with either a pentagon or a quadrilateral, as in the cases of the statement. We first define y', y''.

For each vertex $p \in V(\Gamma) \cap [w''y]$ consider $d(p, [yz])$. Notice that when going from p to an adjacent vertex on $[w''y]$ this distance can change at most by 1, by the triangle inequality. We have that $d(y, [yz]) = 0$ and $d(w'', [yz]) \geq \frac{10}{9}r - r \geq \lambda + 1$. Hence, since $\lambda \in \mathbb{N}$, if we consider all the vertices that lie on $[w''y]$ we must find that there is a vertex at distance n from $[yz]$ for all $n \in \mathbb{N}_0 : 0 \leq n \leq \lambda$. In particular there must exist $y' \in V(\Gamma)$ on $[w''y]$ such that $d(y', [yz]) = \lambda$. We can take y' to be the closest such point to w''. Let y'' be the corresponding vertex on $[yz]$ so that $d(y', y'') = \lambda$.

Since y' is the closest vertex to w'' on $[w''y]$ with distance λ from $[yz]$, we have that no other vertex, and hence no other point on $[w''y']$, is at distance less than λ from $[y''z]$, i.e. $d([w''y'], [y''z]) = \lambda$.

Now that we have defined the point y' consider the segment $[xw']$. In order to go on constructing a hexagon it would be ideal to have $d([xw'], [zy'']) \geq \lambda$ so that whatever point x' we define on $[xw']$ we are guaranteed that $d([x'y'], [y''z]) = \lambda$. This might however not be true, in which case we can get a quadrilateral.

If $d([xw'], [zy'']) < \lambda$ we can, similarly as above, find vertices $x' \in [xw'], z'' \in [y''z]$ such that $d(x', z'') = \lambda$ and $d([x'w'], [z''y'']) = \lambda$. Hence, since also $d(B_r(w), [y''z'']) \geq \frac{10}{9}r - r > \lambda$ we have $d([x'y'], [y''z'']) \geq \lambda$. In this case we have a quadrilateral $x'y'y''z''$ satisfying *Case 2* (see Figure 7.2).

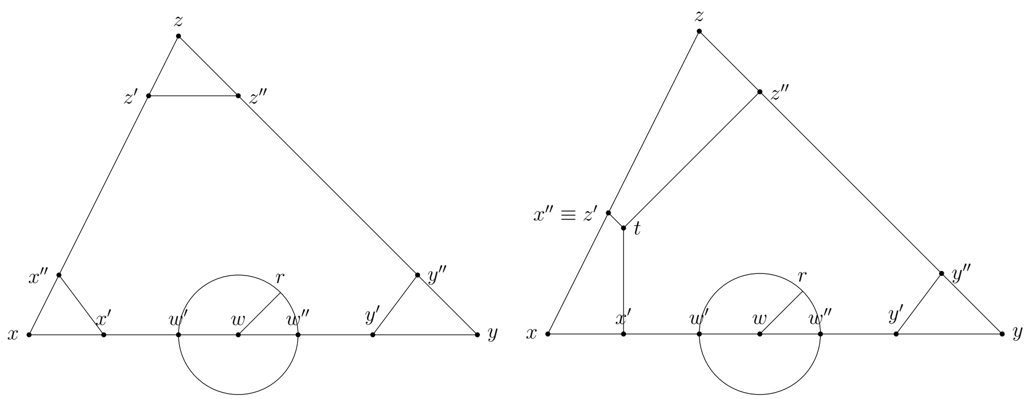

(a) Full hexagon (b) Pentagon

Figure 7.1: Case 1

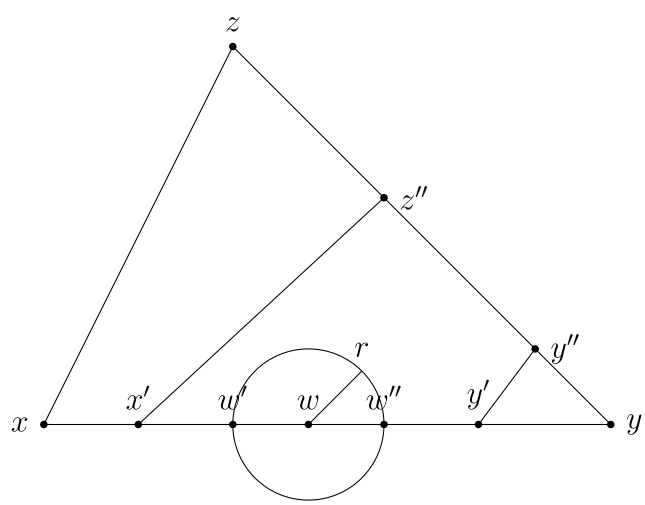

Figure 7.2: Case 2, quadrilateral

Otherwise we immediately have that $d([xy'], [zy'']) \geq \lambda$ and we can go on to find vertices $x' \in [xw'], x'' \in [xz]$ such that $d(x', x'') = \lambda$ and $d([x'w'], [x''z]) = \lambda$.

Now we encounter a similar issue as before. To continue with a hexagon we would like to have $d([w''y'], [x''z]) \geq \lambda$, which is not necessarily true. If $d([w''y'], [x''z]) < \lambda$ we can get another quadrilateral as above. Notice that this quadrilateral would have one side on $[xz]$ instead of $[yz]$. In this case, in order to comply with the statement of the lemma, we need to redefine the points y' and y'' so that $[y'y'']$ is the side of the quadrilateral going from $[w''y]$ to $[xz]$.

Otherwise we have $d([x'y'],[x''z]) = \lambda$ and we continue. Notice that since $r > 9(\lambda+1)$, $d(x'',y'') \geq d(x',y') - d(x',x'') - d(y',y'') \geq 2r - 2\lambda > \lambda$. So if $d(x'',[zy'']) < \lambda$ we can find a vertex $z'' \in [y''z]$ such that $d(x'',z'') = \lambda$ and we can obtain a pentagon $x'y'y''z''x''$ satisfying *Case 1* of the lemma with $z' \equiv x''$.

Otherwise we have $d(x'',[zy'']) \geq \lambda$ and we can take the closest vertex $z' \in [x''z]$ to x'' such that $d(z',[y''z]) = \lambda$. Let $z'' \in [y''z]$ be a vertex such that $d(z',z'') = \lambda$, possibly with $z'' \equiv y''$. We then have vertices satisfying *Case 1* of the lemma (Figure7.1).

Notice that in *Case 1* we are allowing the possibility that either $z'' \equiv y''$ or $z' \equiv x''$. This means that $[z'z''] \cap [y'y'']$ or $[z''z'] \cap [x'x'']$ could be bigger than just one point. In this case (without loss of generality $[z''z'] \cap [x'x''] = [tx'']$) we can get a non-degenerate pentagon $x'y'y''z''t$ with $\lambda < l([x't]) + l([tz'']) < 2\lambda$ (Figure 7.1b). □

The next definition concerns disc diagrams and gives a way of constructing subdiagrams from a given subdiagram.

Definition 7.2. [Alonso et al. 1990]. Let D be a disc diagram. Let E be a subdiagram of D. Then we define $Star_D(E)$ to be the subdiagram of D which has as regions all regions $A \in D$ such that $\overline{A} \cap E \neq \emptyset$ and as edges and vertices all the edges and vertices that are in the boundary of at least one of these regions.

Notice that the *Star* operation is monotone, in the sense that the constructed subdiagram contains the starting one, and that we can get the whole diagram by applying it to a non-empty subdiagram finitely many times.

Example 7.3. In this example we can see the *Star* operation at work. We start with a disc diagram D which is connected, simply connected and has 6 regions.

In the picture on the left we see the disc diagram D and in bold a subdiagram E which has no regions. When we apply the *Star* operation to the subdiagram E we get a subdiagram which consists of the regions whose boundary is intersected by E and the boundaries of these regions. In our case there are two such regions.

The picture on the right indicates $Star_D(E)$, with the 1-skeleton in bold and the regions shaded. Notice that we would have to apply the operation two more times in order to have the full disc diagram D.

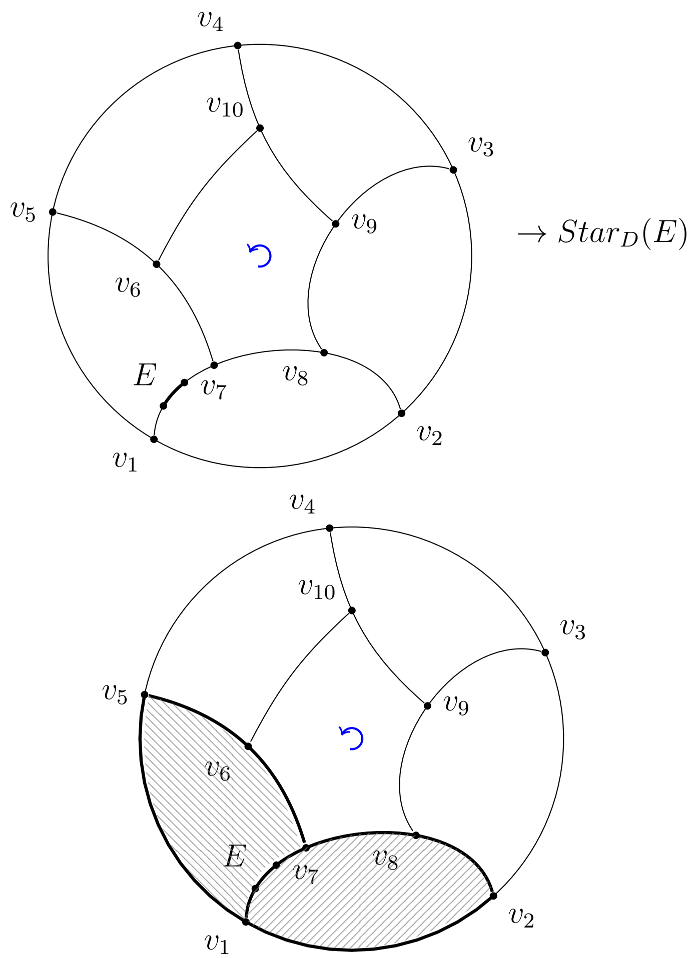

Figure 7.3

If we are given a closed path corresponding to a reduced word in a Cayley graph and we construct a minimal diagram for it, then we know that the label of the diagram will be precisely the word we started with. There is therefore a natural identification of the boundary of the diagram with the closed path. The following lemma (original) concerns the correspondence between the interior of the diagram and the Cayley graph.

Lemma 7.4. Let $G = \langle X|R \rangle$ be a finitely presented group and let $\Gamma = \Gamma(G, X)$ be a Cayley graph of G. Let α be a closed path in Γ corresponding to a cyclically reduced word $u \in WP(G, X)$. Let D be a minimal R-diagram for the word u. Let $\tilde{h} : \delta D \to \Gamma$ be the natural identification of δD with α. We can extend this to a map $h : D^{(1)} \to \Gamma$ with the following property: any path γ in D has $l(\gamma) \geq d(h(s(\gamma)), h(t(\gamma)))$.

Proof. We first define the extension of the map $\tilde{h}$ on the vertices of the diagram. For any vertex $v \in D^{(1)}$ we take any path from a vertex $v' \in \delta D$ and we multiply $h(v') \in V(\Gamma)$ by the label of this path to obtain a new point $h(v) \in \Gamma$.
We just need to show that this operation is well defined i.e. that choosing different vertices on the boundary or different paths does not change $h(v)$.
If we fix v' we notice that since D is a R-diagram all the paths from v' to v have labels that represent the same element in G.

In the following exemplification (Figure 7.4) we see why it does not matter which path we take from v' to v. We look at a R-diagram with 9 regions. We pick a vertex v' on the boundary and in bold we have a fixed path between v and v'. We start from the picture (a), where we have a second path, which is dotted, between v' and v. In each picture $(b), (c), (d), (e)$ we alter the dotted path making sure that the label stays equivalent in G. We are sure that this happens because we only substitute subpaths of the dotted path with their counterpart on a boundary of a region, which has label equivalent to the identity. Finally in (e) we just need another step to have the dotted path coincide completely with the fixed one.

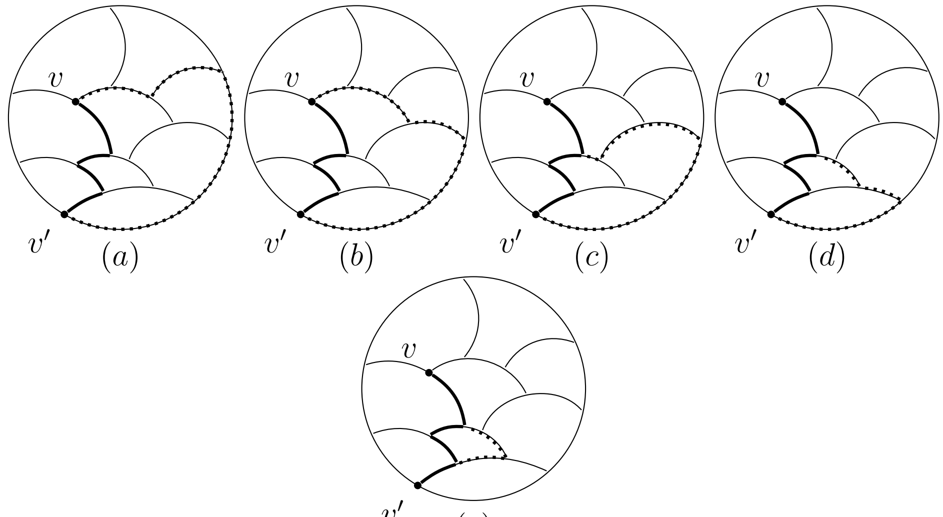

Figure 7.4

If we then have another vertex v'' on the boundary of D we can go from v' to v'' on a path $v'v''$ on δD and then to v via a path $v''v$ so that $h(v')\phi(v'v'')\phi(v''v) = h(v'')\phi(v''v)$, as required. Also, for two vertices a, b in Γ, $d(a, b)$ is the minimal length of any word c such that $ac = b$, so we also have the inequality required.

It is then clear that if there is an edge between two vertices in D we can map it to the edge between the corresponding vertices in Γ, getting the full map h. □

The previous lemma allows us to refer to paths on the boundary of an R-diagram as just paths in the Cayley graph. So if α is one particular circuit of a Cayley graph Γ, corresponding to a cyclically reduced word u in the word problem, and β is a subpath of γ, when we say consider the path β on the boundary of a minimal R-diagram D for u, what we implicitly mean is consider the path on δD which corresponds to the path β in Γ.

We are now going to consider the different circuits that we obtain when cutting a triangle as in Lemma 7.1 and their relative disc diagram. The following lemma will be essential and is original (it was assumed to be true without proof or justification in [Alonso et al. 1990]).

Lemma 7.5. Let $G = \langle X|R\rangle$ be a finitely presented group. Let $\rho = max\{|r| : r \in R\}$. Let $r \in \mathbb{R}, \lambda \in \mathbb{N}$ with $r > 9(\lambda + 1)$. Let xyz be a non-degenerate geodesic triangle in $\Gamma(G,X)$ with $x,y,z \in V(\Gamma(G,X))$. Let $w \in [xy] \cap V(\Gamma(G,X))$ and assume that $B_{\frac{10}{9}r}(w) \cap [xz] = B_{\frac{10}{9}r}(w) \cap [yz] = \emptyset$.

Let α be the circuit in $\Gamma(G,X)$ obtained after cutting xyz as in one of the three cases arising of Lemma 7.1 (so α is either a geodesic hexagon, pentagon or quadrilateral). Assume that $\lambda \geq \rho$. Let u be the cyclically reduced word $u \in WP(G,X)$ corresponding to α. Let D be a minimal R-diagram for u and let θ be one of the paths on δD corresponding to one of the sides of α that lies on the uncut geodesic triangle. For brevity let $\mathcal{A}$ be the finite set of regions in D whose boundary intersects θ.
Then

$$l(\overline{\delta(Star_D(\theta)) \cap int(D)}) \geq l(\theta) - \rho.$$

Proof. We are first going to start from the case where α is a hexagon. We will then outline the proof for the other two cases. Recall from the construction in Lemma 7.1 that we have $d(\theta,\theta') = \lambda$ for all the sides of α that lie on the uncut geodesic triangle and that this distance is realised by following a side of α connecting θ to θ' (these assumptions are all we are going to need).

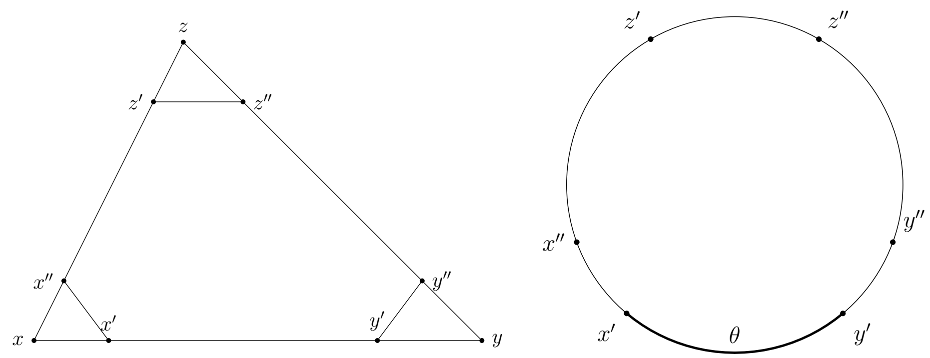

(a) Circuit $\alpha = x'y'y''z''z'x''x'$ (b) Boundary of diagram D

Figure 7.5

In the picture on the left we can see a triangle cut into a hexagon and on the right we have the boundary of a minimal R-diagram for u identified with α. We assume that $\theta = [x'y']$.

Let A be any region in $\mathcal{A}$. First observe that $\delta A \cap [y''z''] = \emptyset$ and $\delta A \cap [z'x''] = \emptyset$ since by construction $d(\theta, [y''z''] \cup [z'x'']) = d([x'y'], [y''z''] \cup [z'x'']) = \lambda$ and the boundary δA has length at most $\rho < \lambda$ (here we are using the fact that if there is a path of certain length in D then there is one equally long or shorter in $\Gamma(G, X)$).

Next we notice that $\delta A \cap [z'z''] = \emptyset$. To see this assume there is a vertex p on δA such that $p \cap [z'z''] \neq \emptyset$. We can then get from θ to $[z'z'']$ on a path of length at most $\frac{\rho}{2}$. The point p splits the path on δD, $[z'z'']$, in two geodesics $[z'p], [pz'']$. Since $d([x'y'], [y''z'']) = \lambda$, we must have that $l([pz'']) \geq \lambda - \frac{\rho}{2}$. Similarly $l([z'p]) \geq \lambda - \frac{\rho}{2}$. Hence $l([z'z'']) = l([z'p]) + l([pz'']) \geq 2\lambda - \rho$. Finally, since $l([z'z'']) = \lambda$, we have $\lambda \geq 2\lambda - \rho$, which implies $\lambda \leq \rho$, absurd.

This shows that the boundary of A can only intersect θ or the neighbouring sides $[x'x''], [y'y'']$, and is otherwise contained in the interior of the diagram. Also, there is at least one region in $\mathcal{A}$ whose boundary intersects $[x'x'']$ and one whose boundary intersects $[y'y'']$.

These observations allow us to say that we can pick two regions $A_x, A_y \in \mathcal{A}$, possibly equal, such that $\delta A_x \cap [x'x''] \neq \emptyset$ and $\delta A_y \cap [y'y''] \neq \emptyset$, and all the regions in $\mathcal{A}$ in between A_x, A_y do not intersect $[x'x''] \cup [y'y'']$. Let $\mathcal{A}^*$ be the subset of $\mathcal{A}$ made up of the regions A_x, A_y and the regions in between. We can see the situation illustrated below.

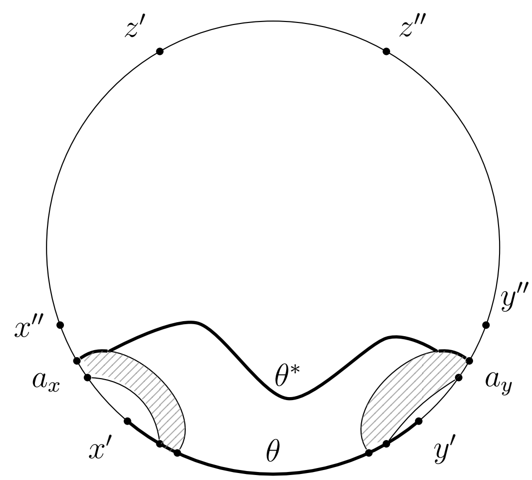

Figure 7.6: Regions A_x, A_y, shaded.

We observe that there is a path θ^* from a vertex $a_x \in [x'x'']$ to a vertex $a_y \in [y'y'']$ which lies completely on the boundary of $\cup_{A \in \mathcal{A}^*} A$. This follows by construction.

By construction we know that we can go from θ to a_x on a path of length less than or equal to $\frac{\rho}{2}$. But since $d([x'y'], [x''z']) = \lambda$, we have that

$d(a_x, x'') \geq \lambda - \frac{\rho}{2}$. This implies that $d(x', a_x) = \lambda - d(a_x, x'') \leq \frac{\rho}{2}$ which means that $l([a_x x']) \leq \frac{\rho}{2}$. Similarly $l([a_y y']) \leq \frac{\rho}{2}$. Hence, since θ is a geodesic and we can go from x' to y' along $[x' a_x]$, θ^* and $[a_y y'']$, we must have $l(\theta^*) \geq l(\theta) - \rho$.

But θ^* is just a subset of $\overline{\delta(Star_D(\theta)) \cap int(D)}$ and the claim is proven. Finally notice that we have not used any special assumption about θ being $[x'y']$, and therefore the exact same arguments work when $\theta = [x''z']$ or $\theta = [y''z'']$

We now look at the case where α is a quadrilateral. Let θ be either $[x'y']$ or $[y''z'']$. In exactly the same way as we did for the hexagon, we find that each region in $\mathcal{A}$ has boundary that can only intersect θ or $[x'z'']$ or $[y'y'']$. We can then construct a path in the interior of D by again picking the last region in $\mathcal{A}$ which has boundary intersecting $[x'z'']$ and the last that has boundary intersecting $[y'y'']$, and conclude as before.

If α is instead a pentagon we let θ be $[x'y']$ (the case $\theta = [y''z'']$ is symmetric). Here we need to be more careful since we have two sides of the pentagon we do not know the exact length of.

First we need to show that it is not possible for a region in $\mathcal{A}$ to also intersect $[tz']$. To see this assume there is such a region. Then it is possible to get from $[x'y']$ to $[tz']$ on a path of length less or equal to $\frac{\rho}{2}$. Consider the whole geodesic $[z'z'']$ (i.e. the geodesic which contains t from z' to z''). Since we know that $d([x'y'], [y''z'']) = \lambda$ and $l([z'z'']) = \lambda$ we must then have that it is possible to get from $[x'y']$ on a path of length less than or equal to $\frac{\rho}{2}$. But this is not possible since by construction the distance between the point z' and $[x'y']$ is λ and $\lambda > \rho$.

Now that we have shown that any region in $\mathcal{A}$ can only intersect $[x't]$ or $[y'y'']$ we still need to check that, as before, no region in $\mathcal{A}$ has boundary which intersects $[x't]$ or $[y'y'']$ too far up. But this can be shown in exactly the same way as before by considering the fact that $[x't]$ is a sub-geodesic of $[x'x'']$, $l([x'x'']) = \lambda$ and $d([x'y'], x'') = \lambda$. $\square$

We now extend the above lemma, in order to be able to deal with repeated iterations of the *Star* operation. All the ideas that are needed are discussed in the proof of the lemma and we will therefore provide only a sketch of the proof of the corollary. In order for it to work we only need to strengthen the assumption about the length λ. The case $n = 1$ will correspond to the original lemma.

Corollary 7.6. Let $G = \langle X | R \rangle$ be a finitely presented group. Let $\rho = max\{|r| : r \in R\}$ and $n \in \mathbb{N}$. Let $r \in \mathbb{R}, \lambda \in \mathbb{N}$ with $r > 9(\lambda + 1)$. Let xyz be a non-degenerate geodesic triangle in $\Gamma(G,X)$ with $x,y,z \in V(\Gamma(G,X))$. Let $w \in [xy] \cap V(\Gamma(G,X))$ and assume that $B_{\frac{10}{9}r}(w) \cap [xz] = B_{\frac{10}{9}r}(w) \cap [yz] = \emptyset$.

Let α be a circuit in $\Gamma(G,X)$ obtained after cutting xyz as in one of the three cases arising of Lemma 7.1 (so α is either a geodesic hexagon, pentagon or quadrilateral). Assume that $\lambda > n\rho$. Let u be the cyclically reduced word $u \in WP(G,X)$ corresponding to α. Let D be a minimal R-diagram for u and let θ be one of the paths on δD corresponding to one of the sides of α that lies on the uncut geodesic triangle.

Then we have that

$$l(\overline{\delta(Star_D^{(n)}(\theta)) \cap int(D)}) \geq l(\theta) - n\rho.$$

Proof. Note that we do not prove this by induction, because each path that we construct is not a geodesic. We simply prove it by noticing that if we apply the *Star* operation n times we can construct a path in the interior of the diagram which has starting and ending points within $\frac{\rho}{2}n$ of the start and end points of θ respectively. This all follows easily by using the same arguments as before and by noticing that we can reach any point on the boundary of $Star^{(n)}(\theta)$ from θ along a path of length at most $\frac{\rho}{2}n$. □

We now use this result to determine a first lower bound on the number of regions of an R-diagram corresponding to our circuit obtained by cutting a triangle. Again, the ideas are basically the same as the ones in [Alonso et al. 1990], with some adjustments.

Lemma 7.7 (Number of regions)**.** Let $G = \langle X|R\rangle$ be a finitely presented group. Let $\rho = max\{|r| : r \in R\}$ and $N \in \mathbb{N}$. Let $r \in \mathbb{R}, \lambda \in \mathbb{N}$ with $r > 9(\lambda+1)$. Let xyz be a non-degenerate geodesic triangle in $\Gamma(G,X)$ with $x,y,z \in V(\Gamma(G,X))$. Let $w \in [xy] \cap V(\Gamma(G,X))$ and assume that $B_{\frac{10}{9}r}(w) \cap [xz] = B_{\frac{10}{9}r}(w) \cap [yz] = \emptyset$.

Let α be a circuit in $\Gamma(G,X)$ obtained after cutting xyz as in one of the three cases arising of Lemma 7.1. Assume that $\lambda > N\rho$. Let u be the cyclically reduced word $u \in WP(G,X)$ corresponding to α. Let D be a minimal R-diagram for u and let θ be one of the paths on δD corresponding to one of the sides of α that lies on the uncut geodesic triangle. Let θ_n be the set (of edges) $\overline{\delta(star_D^{(n)}(\theta)) \cap int(D)}$ for all $n \in \mathbb{N}_0$ with $n \leq N$.
Then the number of regions in $Star_D^{(n)}(\theta)$ is at least $\frac{2l(\theta)n}{\rho} - n^2$ for all $n \leq N$.

Proof. We prove it by induction on n. The statement is trivially true for $n = 0$. Assume that it holds for $n-1$. We want to show that it holds for $n \leq N$. By inductive hypothesis we have that the number of regions of $Star_D^{(n-1)}(\theta)$ is at least $\frac{2l(\theta)(n-1)}{\rho} - (n-1)^2$.

Consider the set of edges θ_{n-1}. Each one of these edges is on the boundary of exactly two regions, one in $Star_D^{(n-1)}(\theta)$ and one not. Each edge in θ_n is on the boundary of a region in $Star_D^{(n)}(\theta)$ but not in $Star_D^{(n-1)}(\theta)$.
Hence, since each region has boundary length at most ρ, we must have that

$$A(Star_D^{(n)}(\theta)) - A(Star_D^{(n-1)}(\theta)) \geq \frac{l(\theta_{n-1} \cup \theta_n)}{\rho}.$$

But θ_{n-1} can intersect θ_n only in finitely many isolated points. Hence $l(\theta_{n-1} \cup \theta_n) = l(\theta_{n-1}) + l(\theta_n)$ and

$$A(Star_D^{(n)}(\theta)) - A(Star_D^{(n-1)}(\theta)) \geq \frac{l(\theta_{n-1}) + l(\theta_n)}{\rho}.$$

But from Corollary 7.6 we know that $l(\theta_n) \geq l(\theta) - n\rho$ and $l(\theta_{n-1}) \geq l(\theta) - (n-1)\rho$.
We then have that

$$A(Star_D^{(n)}(\theta)) \geq A(Star_D^{(n-1)}(\theta)) + \frac{2l(\theta) - n\rho - (n-1)\rho}{\rho}.$$

which by inductive assumption implies that

$$A(Star_D^{(n)}(\theta)) \geq \frac{2l(\theta)}{\rho} - (n-1)^2 + \frac{2l(\theta)}{\rho} - 2n + 1 = \frac{2l(\theta)n}{\rho} - n^2.$$

□

In the construction of Lemma 7.6 we look at paths between the two sides of length λ in the cut triangle and use a lower bound on their length to determine a lower bound on the number of regions. The geodesic $[x'y']$ has the peculiarity of being the only geodesic of the two or three considered that also intersects the ball centered at w. If we consider the last path, the one that is on the boundary of the last iteration of the *Star* operation on $[x'y']$, it is natural to ask whether we can find more regions that have not been counted. The following lemma starts addressing the question, in a slightly more general setting.

Lemma 7.8. Let $G = \langle X|R \rangle$ be a finitely presented group and let $\Gamma = \Gamma(G, X)$ be a Cayley graph of G. Let $\rho = max\{|r| : r \in R\}$. Let α be a circuit in Γ corresponding to a cyclically reduced word $u \in WP(G, X)$. Let D be a minimal R-diagram for u. Let $L, c \in \mathbb{R}$ with $c > \frac{L}{2}$. Let $w \in \alpha \cap V(\Gamma)$ and let B_c be the ball of radius c centered at w. Let h be the natural identification map from D to Γ and let $\tilde{w}$ be the vertex on δD such that $h(\tilde{w}) = w$. Let $\tilde{\gamma}$ be a path in D with $\tilde{w}$ in a $\frac{L}{2}$ neighbourhood of $\tilde{\gamma}$ and with both $s(\tilde{\gamma})$ and $t(\tilde{\gamma})$ not in a c neighbourhood of $\tilde{w}$.
Then there is a subpath $\tilde{\gamma}'$ of the path $\tilde{\gamma}$, of length at least $2c - \lambda$, such that $h(\tilde{\gamma}') \subseteq B_c$.

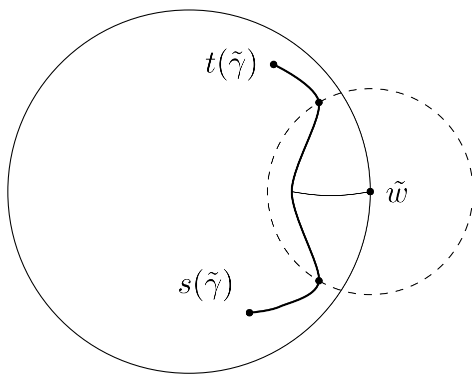

Figure 7.7

Proof. We can go from $\tilde{w}$ to $\tilde{\gamma}$ along a path of length at most $\frac{L}{2}$ since $\tilde{w}$ lies in a $\frac{L}{2}$ neighbourhood of $\tilde{\gamma}$. From there we can go in two directions on subpaths of length $c - \frac{L}{2} > 0$ since we cannot reach $s(\tilde{\gamma})$ and $t(\tilde{\gamma})$ from $\tilde{w}$ on a path of length less than c (see Figure 7.7). But then we have a subpath $\tilde{\gamma}'$ of $\tilde{\gamma}$ which lies in a c neighbourhood of $\tilde{w}$ and hence $h(\tilde{\gamma}') \in B_c$, by property of h. $\square$

Lemma 7.1, that we repeatedly use to cut a geodesic triangle, only works with a non-degenerate triangle xyz with $x, y, z \in V(\Gamma)$ and with a ball (not intersecting $[xz] \cup [yz]$) centred at a vertex $w \in [xy]$. We would like to be able to use the lemma on any geodesic triangle xyz with a ball centred at a point (possibly not a vertex) on $[xy]$ and the following two lemmas (original) show that we are able to do it.

Lemma 7.9 (Non-degenerate triangle)**.** Let (X, d) be geodesic metric space and let xyz be a geodesic triangle in X. Let $r \in \mathbb{R}_{>0}$ and let $w \in [xy]$.
Then if $B_r(w) \cap ([xz] \cup [yz]) = \emptyset$ we can find a non-degenerate geodesic triangle $x'y'z'$ such that each side is a subpath of one of the sides of xyz, $w \in [x'y']$ and $B_r(w) \cap ([x'z'] \cup [y'z']) = \emptyset$.

Proof. Since $B_r(w) \cap ([xz] \cup [yz]) = \emptyset$ we can define the points $x', y' \in [xy]$ such that $x' \in [xz], y' \in [yz]$ and x', y' are the closest such points to w on the two directions from w of the path $[xy]$. Consider $A = [x'z] \cap [y'z] \setminus \{x'\}$. If $A \neq \emptyset$ define z' as the closest point to x' in A. If $A = \emptyset$ define z' as any point on $[yy']$ not equal to x' or y'. The geodesic triangle $x'y'z'$ clearly has $w \in [x'y']$ and since the ball did not intersect the sides $[xz], [yz]$ it does not intersect $[x'z'] \cup [y'z']$. □

Lemma 7.10. Let $G = \langle X|R\rangle$ be a finitely presented group and let $\Gamma = \Gamma(G, X)$ be a Cayley graph of G. Let $\delta \in \mathbb{R}$ with $\delta \geq 1$.
If there exists a geodesic triangle $xyz \in \Gamma$ with a point $w \in [xy]$ such that $B_\delta(w) \cap \{[xz] \cup [yz]\} = \emptyset$ then there exists a non-degenerate geodesic triangle $x'y'z'$ with $x', y', z' \in V(\Gamma)$ and a vertex $w' \in [x'y']$ such that $B_{\delta-1}(w') \cap \{[x'z'] \cup [y'z']\} = \emptyset$.

Proof. Γ is a geodesic metric space. Given the geodesic triangle xyz we can then use Lemma 7.9 to get a non-degenerate geodesic triangle $x''y''z''$ with $w \in [x''y'']$ and $B_\delta(w) \cap \{[x''z''] \cup [y''z'']\} = \emptyset$.
If x'' is not a vertex in Γ it must lie on an edge e which has ends on $[x''y'']$ and $[x''z'']$. We call $x' \in V(\Gamma)$ the end of e which lies in $[x''z'']$. We do the same for y'' and for z'' (in the case of z'' it does not matter which side we shorten).
We then have a non-degenerate triangle $x'y'z'$ with with $x', y', z' \in V(\Gamma)$ and $w \in [x'y']$ such that $B_\delta(w) \cap \{[x'z'] \cup [y'z']\} = \emptyset$. Finally define w' as one vertex on $[x'y']$ on the same edge as w. By the triangle inequality we have $B_{\delta-1}(w') \cap \{[x'z'] \cup [y'z']\} = \emptyset$ as required. □

Before proving the biggest theorem of the section we prove the following small lemma.

Lemma 7.11. Let $G = \langle X|R\rangle$ be a finitely presented group and let $\Gamma = \Gamma(G, X)$ be a Cayley graph of G. Let $\rho = max\{|r| : r \in R\}$.
If $\rho \leq 1$ then Γ is $\frac{1}{4}$-slim.

Proof. We already know from Example 4.8 that if $\rho = 0$ then Γ is 0-slim since every geodesic triangle has every point belonging to two sides. If $\rho = 1$ then Γ is a tree with the addition of a loop on every vertex for each relator. The only possibility of having a geodesic triangle where one point is not on two sides occurs when a vertex is the midpoint of a loop. In this case we get a geodesic triangle which is a geodesic triangle in the tree with the addition of at most three loops. But each loop has length one, hence in the worst case scenario we will need to go on a path of length at most $\frac{1}{4}$ to reach a different side starting from a point on a loop. □

We are finally ready to show that if a group satisfies a linear isoperimetric inequality than it is hyperbolic. The proof essentially follows the steps found in [Alonso et al. 1990], but the value of δ has been improved with the leading term going from $9K^2\rho^3$ to $\frac{5}{4}K^2\rho^3$, effectively reducing the constant δ by a factor of more than 7.

Theorem 7.12 (Isoperimetric $\Rightarrow$ δ-slim)**.** Let $G = \langle X|R\rangle$ be a finitely presented group and let $\Gamma = \Gamma(G, X)$ be a Cayley graph of G. Let $\rho = max\{|r| : r \in R\}$.

If G satisfies a linear isoperimetric inequality with constant $K \geq 1$, then all geodesic triangles in Γ are δ-slim, for $\delta = \frac{5}{4}K^2\rho^3 + 15K\rho^2 + 1$.

Proof. If $\rho \leq 1$ then Lemma 7.11 gives the proof. We can therefore assume that $\rho \geq 2$. Suppose that there is a geodesic triangle xyz which is not δ-slim, for δ as above.

Let $\delta' = \delta - 1$, $r = \frac{9}{10}\delta'$ and $\lambda = \lceil\frac{K\rho^2}{2}\rceil + \rho$. There exists a vertex $w \in [xy]$ such that $B_{\delta'}(w)$ does not intersect $[xz] \cup [yz]$. By Lemma 7.10 we can assume xyz to be a non-degenerate geodesic triangle with $x, y, z \in V(\Gamma)$ and $w \in [xy] \cap V(\Gamma)$ such that $B_{\delta'}(w) \cap \{[xz] \cup [yz]\} = \emptyset$.

We have that

$$r = \frac{9}{10}(\frac{5}{4}K^2\rho^3 + 15K\rho^2) = \frac{9}{8}K^2\rho^3 + \frac{27}{2}K\rho^2$$

and hence, since $\rho \geq 2, K \geq 1$,

$$r > \frac{9}{8}K^2\rho^3 + 9K\rho^2 \geq 9 + \frac{9}{2}K\rho^2 + 9\rho \geq 9(\lambda + 1).$$

Hence, since $r > 9(\lambda + 1)$, we can apply Lemma 7.1 to cut the triangle xyz. In what follows we use the notation from the lemma. Let $u \in WP(G, X)$ be

the cyclically reduced word corresponding to the circuit in Γ constructed by cutting the triangle. Let D be a minimal R-diagram for u.
We now consider the different cases arising from the construction.

If we have a full hexagon, let $\theta_1, \theta_2, \theta_3$ be the paths on δD corresponding respectively to the geodesic paths $[x'y'], [y''z''], [x''z']$. If we have a pentagon or a quadrilateral let θ_1, θ_2 be be the paths on δD corresponding respectively to the geodesic paths $[x'y']$, $[y''z'']$ or $[x''z']$. Let $n = \lfloor \frac{\lambda}{\rho} \rfloor$. For brevity let D_i be the subdiagram $Star_D^{(n)}(\theta_i)$ for $i \in \{1,2,3\}$ in the case of the hexagon or $i \in \{1,2\}$ in the other cases.

For $i \neq j$, by construction $d(\theta_i, \theta_j) \geq \lambda$ which implies that the $\frac{\lambda}{2}$-neighbourhoods of θ_i and θ_j do not intersect each other. Also D_i lies in a $n \cdot \frac{\rho}{2} \leq \frac{\lambda}{2}$ neighbourhood of θ_i and similarly D_j lies in a $\frac{\lambda}{2}$-neighbourhood of θ_j.
Hence $D_i \cap D_j = \emptyset$ for $i \neq j$.
Now, θ_i is a geodesic for $1 \leq i \leq 3$ and $\lambda \geq n\rho$ so by Lemma 7.7 we have that

$$A(D_i) \geq \frac{2l(\theta_i)n}{\rho} - n^2$$

$$\Rightarrow A(D_i) \geq \frac{2l(\theta_i)(\frac{\lambda}{\rho} - 1)}{\rho} - \frac{\lambda^2}{\rho^2}$$

and since $\frac{K\rho^2}{2} + \rho \leq \lambda \leq \frac{K\rho^2}{2} + \rho + 1$,

$$\Rightarrow A(D_i) \geq Kl(\theta_i) - \frac{(\frac{K\rho^2}{2} + \rho + 1)^2}{\rho^2}.$$

Also notice that $r > \frac{\lambda}{2}$, so there is a vertex $\tilde{w}$ on θ_1 such that $h(\tilde{w}) = w$ (h is the identification map from δD to the cut triangle) and the path $\gamma = \overline{\delta(D_1) \cap int(D))}$ satisfies the requirements of Lemma 7.8. We can hence apply this lemma to find a path γ' in the interior of D which is mapped within the ball $B_r(w)$ by h and has length $l(\gamma') \geq 2r - \lambda$. Clearly γ' does not intersect a $\frac{\lambda}{2}$ neighbourhood of either θ_2 or θ_3 and hence it is on the boundary of some regions which are not in D_i for all i. Since the maximum boundary length of a single region is ρ, the number of these regions is at least

$$\frac{2r - \lambda}{\rho} \geq \frac{2r}{\rho} - \frac{K\rho}{2} - 1 - \frac{1}{\rho}.$$

Hence, in the hexagon case

$$A(D) \geq K(l(\theta_1) + l(\theta_2) + l(\theta_3)) + \frac{2r}{\rho} - \frac{3K^2\rho^2}{4} - \frac{7K\rho}{2} - 3K - 4 - \frac{7}{\rho} - \frac{3}{\rho^2}$$

and in the other two cases

$$A(D) \geq K(l(\theta_1) + l(\theta_2)) + \frac{2r}{\rho} - \frac{K^2\rho^2}{2} - \frac{5K\rho}{2} - 2K - 3 - \frac{5}{\rho} - \frac{2}{\rho^2}.$$

In the hexagon case $l(u) = 3\lambda + l(\theta_1) + l(\theta_2) + l(\theta_3)$, in the pentagon case $l(u) = 3\lambda + l(\theta_1) + l(\theta_2)$ and in the quadrilateral case $l(u) = 2\lambda + l(\theta_1) + l(\theta_2)$. Since G satisfies a linear isoperimetric inequality we must have $A(D) \leq Kl(u)$.
But hence, for the hexagon,

$$3K\lambda \geq \frac{2r}{\rho} - \frac{3K^2\rho^2}{4} - \frac{7K\rho}{2} - 3K - 4 - \frac{7}{\rho} - \frac{3}{\rho^2}$$

$$\Rightarrow 3K(\frac{K\rho^2}{2} + \rho + 1) \geq \frac{2r}{\rho} - \frac{3K^2\rho^2}{4} - \frac{7K\rho}{2} - 3K - 4 - \frac{7}{\rho} - \frac{3}{\rho^2}$$

$$\Rightarrow r \leq \frac{9K^2\rho^3}{8} + \frac{13}{4}\rho^2 K + 3K\rho + 2\rho + \frac{7}{2} + \frac{3}{2\rho}$$

and by using again the fact that $\rho \geq 2, K \geq 1$,

$$\Rightarrow r < \frac{9K^2\rho^3}{8} + \frac{27}{2}\rho^2 K = r.$$

This is absurd.

Since we have the same lower bound for the pentagon and the quadrilateral case, and a weaker upper bound for the pentagon, we just need to verify that we get a contradiction in the pentagon case. In fact for the pentagon we get (some steps omitted, same technique as above)

$$r \leq K^2\rho^3 + \frac{11}{4}K\rho^2 + \frac{5}{2}K\rho + \frac{3}{2}\rho + \frac{5}{2} + \frac{1}{\rho}$$

which, by again using the fact that $K \geq 1$ and $\rho \geq 2$ gives us

$$r \leq K^2\rho^3 + \frac{26}{2}\rho^2 K < r.$$

This is absurd and concludes the proof. □

The following corollary basically states the Theorem 7.12 in a more concise way and deals with the case when the linear isoperimetric constant is less than 1.

Corollary 7.13. All the geodesic triangles in the Cayley graph $\Gamma(G, X)$ of a finitely presented group $G = \langle X|R \rangle$ which satisfies a linear isoperimetric inequality are δ-slim for some $\delta > 0$.

Proof. If the G satisfies a linear isoperimetric inequality with constant $K < 1$ then it also satifies one with $K \geq 1$. We can then use Theorem 7.12 to get $\delta > 0$ such that all triangles are δ-slim. □

Our next aim is to show the reverse direction of Theorem 7.12, i.e. that if triangles are slim in a Cayley graph, then the group satisfies a linear isoperimetric inequality. In order to do this we start from the equivalence of having slim triangles and thin triangles and we show that the group has a Dehn presentation. This approach can be found in [Can1989] and it differs from the one used by Gromov in [Gr1987].

In order to do this we use the the concept of k-local geodesics, described in the following definition.

Definition 7.14. Let $k \in \mathbb{R}^+$ and let (X, d) be a geodesic metric space. Then a path α in X is called a *k-local geodesic* if any subpath of α of length at most k is a geodesic path.

We want to show that if triangles are thin then k-local geodesics do not lie far away from the true geodesics. If we manage to show this it will then be easy to conclude that there exists a Dehn presentation. We first need the following technical lemma.

Lemma 7.15. Let $\delta \in \mathbb{R}^+$, $k = 4\delta$ and let (X, d) be a geodesic metric space. Let α, β be paths of length $\geq 2\delta$ in X that start and end at the same points g, h, such that α is a k-local geodesic and β is a geodesic. Let s, t be the unique points respectively on α and β at distance 2δ from h. If the geodesic triangles in X are δ-thin then $d(s, r) \leq \delta$.

Proof. We prove the lemma by induction on the length of α, $l(\alpha)$. For the base case let $2\delta \leq l(\alpha) \leq 6\delta$. Now if $l(\alpha) \leq 4\delta = k$, α is a geodesic path and we immediately get that $d(s, r) \leq \delta$, by δ-thinness of the geodesic bigon formed by the paths α, β. Note in fact that we can see a geodesic bigon with two sides of equal length, as a geodesic triangle with one side of length 0 and internal points coinciding with the vertices. Otherwise we can consider a point $p \in \alpha$ at distance 2δ from s. The triangle Δ formed by the subpaths of α in which α is split by the point p, and by the path β, is a geodesic triangle since the two paths on α have length $\leq 4\delta$ and α is k-local geodesic. The side $[gp]$ has length $\leq 2\delta$ so by δ-thinness of Δ we must have $d(s, r) \leq \delta$.

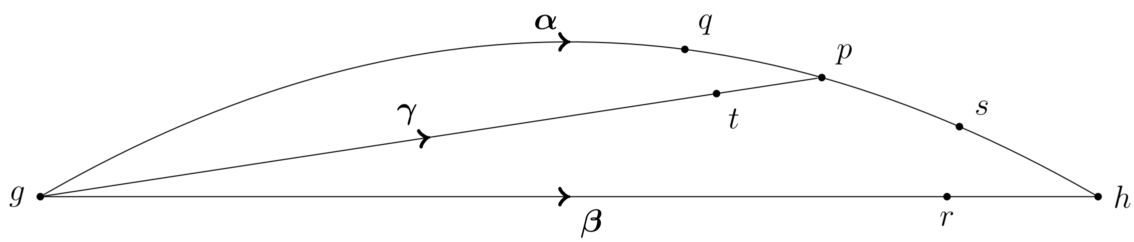

Figure 7.8

Now assume that the statement holds for $l(\alpha) \leq N$, with $N \geq 6\delta$ and consider $N \leq l(\alpha) \leq N + k$. Let p be the point on α at distance 2δ from s, let q be the point on α at distance 2δ from p. Let γ we a geodesic path from g to p. If $l(\gamma) \leq 2\delta$ we conclude as before, so we can let t be the point on γ at distance 2δ from p. By inductive assumption $d(q, t) \leq \delta$. The subpath on α of between q and s has length 4δ and is therefore a geodesic. Hence $d(q, s) = 4\delta$ and since by the triangle inequality $d(q, s) \leq d(q, t) + d(t, s)$ we must have $d(t, s) \geq 3\delta$. But the geodesic triangle formed by γ, β and the path on α between p and h is δ-thin, and since s is not δ close to t it must be within δ of r. $\square$

Now that we have Lemma 7.15 we are able to prove this important result concerning k-local geodesics.

Lemma 7.16. Let $\delta \in \mathbb{R}^+$, $k = 4\delta$ and let (X, d) be a geodesic metric space. Let $g, h \in X$ and assume that the geodesic triangles in X are δ-thin. Then, given any geodesic path β between g and h and any k-local geodesic path α between g, h; α lies in a 3δ neighbourhood of β.

Proof. If there is no point on α at distance at least 2δ from both g, h then there is nothing to prove because α lies in a 2δ neighbourhood of β. Otherwise let z be a point on α such that $d(g, z), d(h, z) \geq 2\delta$. Let x, y be geodesics

$[gz], [zh]$. Let a and b be the 2 points on α at distance 2δ from z and r, s be the points at distance 2δ from z on x, y, as described in Figure 7.9.

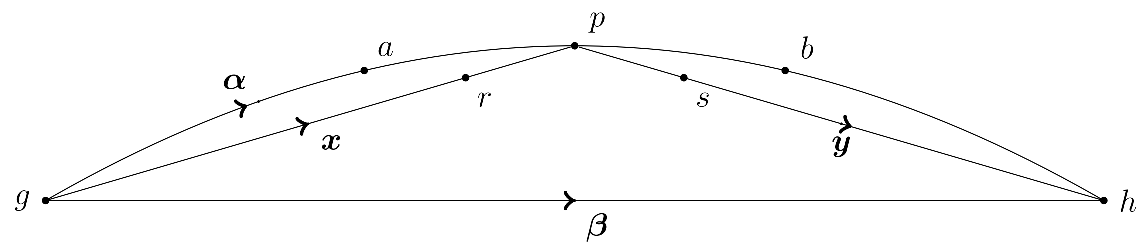

Figure 7.9

By Lemma 7.15 $d(a, r), d(s, b) \leq \delta$. But since α is k-local geodesic $d(a, b) = 4\delta$ and $d(r, s) \geq 2\delta$ by the triangle inequality. The geodesic triangle formed by the geodesic paths x, y, β is δ-thin by assumption. Thus, since $d(r, s) > \delta$, we have $d(r, \beta) \leq \delta$. Therefore $d(z, \beta) \leq d(z, r) + d(r, \beta) \leq 3\delta$. This is then true for all points on α, as required. □

We are finally able to show that if a group has slim triangles than it has a Dehn presentation. Given the results that we have proven regarding k-local geodesics, the idea is simple.

Theorem 7.17. Let $G = \langle X|R \rangle$ be a finitely generated group and let $\Gamma = \Gamma(G, X)$ be a Cayley graph of G. Let $\delta \in \mathbb{R}_{\geq 1}, k = 4\delta$ and assume that the geodesic triangles in Γ are δ-thin. Let $S = \{g \in WP(G, X) : |g| \leq 8\delta - 1\}$. Then $\langle X|S \rangle$ is a Dehn presentation for G.

Proof. It suffices to show that given any freely reduced word $w \in WP(G, X)$ it is possible to find a subword v of w such that $v = u$ for a word u of length $|u| < |v|$ with $uv^{-1} \in S$.

The word w is equivalent to the identity in G and hence corresponds to a closed path α starting and ending at 1_G in Γ. If α is not a k-local geodesic we can find a subpath of α of length at most 4δ which is not geodesic and corresponds to a subword v of w. Letting u be the word corresponding instead to the geodesic path we are done.

We can hence assume that α is a k-local geodesic in Γ. But by Lemma 7.9 we have that α lies in a 3δ neighbourhood of the trivial path at 1_G. This means that the first $3\delta + 1$ letters in w cannot correspond to a geodesic path in Γ. But $3\delta + 1 \leq 4\delta = k$ and α is a k-local geodesic, hence $|w| \leq 3\delta$ and $w \in S$. □

Note that the given Dehn presentation has the same finite generating set we started with. The following corollary follows immediately from the definition of Dehn presentation. We know that a group G having thin triangles implies that G has a Dehn presentation, but hence it also satisfies a linear isoperimetric inequality, for which we can also find the constant, which depends on the given finite presentation for G. Notice that in order to find the constant we use an argument ad hoc since we can't simply apply Theorem 2.44 without knowing the needed Tietze transformations to turn a presentation into the other. Note that we prove the following fact for $\delta \geq 1$, but that if triangles are δ-thin for some $\delta < 1$ they are also δ-thin for any $\delta \geq 1$.

Corollary 7.18. Let $G = \langle X|R\rangle$ be a finitely presented group where the triangles in $\Gamma = \Gamma(G, X)$ are δ-thin for some $\delta \geq 1$. Let $\rho = \max_{r \in R} |r|$. Then $\langle X|R\rangle$ satisfies a linear isoperimetric inequality with constant $K = Dehn(\lfloor 8\delta \rfloor - 1)$.

Proof. By Theorem 7.17 we know that G has a Dehn presentation $\langle X|S\rangle$, where $S = \{g \in WP(G, X) : |g| \leq 8\delta - 1\}$.

Thus, by Theorem 2.41 we have that $\langle X|S\rangle$ satisfies a linear isoperimetric inequality with constant 1. Notice that the set of generators is unchanged between the two presentations.

We know that we can write each relator $s \in S$ as a product of at most $A(s)$ conjugates of relators or inverses of relators in R. Since the set S contains all the words in the word problem of length at most $8\delta - 1$, we know that we can write all of them as a product of at most $Dehn(\lfloor 8\delta \rfloor - 1)$ conjugates of relators or inverses of relators in R. Hence, we have that any freely reduced word in $w \in WP(G, X)$ can be written as a product of at most $|w|Dehn(\lfloor 8\delta \rfloor - 1) = |w|K$ conjugates of relators or inverses of relators in R. □

8 Conclusions

We finally look back at our work, deduce some properties of hyperbolic groups, and state previous results in a more concise way.

Corollary 8.1. Every hyperbolic group is finitely presentable.

Proof. We simply observe that a Dehn presentation is also a finite presentation and by Theorem 7.17 we know that G admits a Dehn presentation. $\square$

We can finally get rid of generating set dependency in the definition of hyperbolic group.

Corollary 8.2. The definition of hyperbolic group is independent of the choice of finite generating set.

Proof. Note that we can assume that the group is finitely presented. We know that a group is hyperbolic if and only if it satisfies a linear isoperimetric inequality, but then from Proposition 2.45 we get that the choice of generating set does not affect hyperbolicity. $\square$

Corollary 8.3. Every hyperbolic group has finitely many conjugacy classes of elements of finite order.

Proof. This follows from the fact the every hyperbolic group has a Dehn presentation (Theorem 7.17) and what we proved before about Dehn presentations in Proposition 2.34. $\square$

We can go back to Example 2.33 and see why $\mathbb{Z}^2$ does not have a Dehn presentation.

Proposition 8.4. The group $\mathbb{Z}^2$ does not have a Dehn presentation.

Proof. We have seen that this is equivalent to being hyperbolic. If we take the standard presentation for $\mathbb{Z}^2$ we see that its relative Cayley graph is basically the tiling of the plane with taxicab distance, which is clearly not a hyperbolic space. $\square$

The above proposition exemplifies perfectly well the core of what we have achieved in our discovery of equivalent characterisations of hyperbolicity. By simply looking at one particular presentation for the group $\mathbb{Z}^2$ we are to determine that it is not possible to find any finite presentation for $\mathbb{Z}^2$ which is a Dehn presentation. By considering the geometry of one Cayley graph we determine a particular algebraic property of the group, related to the word problem. We can say that we were able to determine information about the group by studying its geometric properties. To get a better idea of our accomplishments we now recollect previous results in one single, and fundamental, theorem. The proof is then just a formality.

Theorem 8.5. Let G be a group with finite generating set X. Then the following are equivalent:

(i) G is hyperbolic.

(ii) There is a bound on the insize of geodesic triangles.

(iii) Triangles in $\Gamma(G, X)$ are thin.

(iv) G has a Dehn presentation.

(v) G satisfies a linear isoperimetric inequality.

Proof. $(i) \Rightarrow (ii)$ follows from Proposition 5.1.
$(ii) \Rightarrow (iii)$ follows from Proposition 5.2.
$(iii) \Rightarrow (iv)$ follows from Theorem 7.17.
$(iv) \Rightarrow (v)$ follows from Proposition 2.41.
Finally $(v) \Rightarrow (i)$ follows from Theorem 7.12 and the generating set independence. □

We conclude by summarising what we have done that is new, and talk about what we could have aimed for if there was more time.

In Section 5 we improved the bound when proving that slim implies thin, from 6δ to 4δ (Proposition 5.3). Most important of all, in Section 7 we have restructured the proof of the fact that the existance of a linear isoperimetric inequality implies slimness, fixing several bugs found when working through the proof from [Alonso et al. 1990]. The biggest of these bugs was found in the claim that the *Star* operation applied on a geodesic on the boundary of a minimal R-diagram, produces a path in the interior of the diagram. This is not necessarily true and required a fix, which can be found in Lemma 7.5. Overall we managed to improve the value of δ that makes the Cayley graph of a presentation satisfying a linear isoperimetric inequality δ-slim, by a factor of more than 7, again compared to [Alonso et al. 1990] (Theorem 7.12).

Ideally we would have wanted to get better improvements in terms of bounds on the constants. In particular it would be nice to find a betterbound for the insize, assuming that triangles are δ-slim, and possibly find a different approach for the proof of Theorem 7.12. Furthermore it would be interesting to investigate how far these constants are from the optimal bounds. We would need to start working with more concrete examples in order to produce results in this sense.

References

[Alonso et al. 1990] Alonso, J.M.; Brady, T.; Cooper, D.; Ferlini, V.; Lustig, M.; Mihalik, M.; Shapiro, M.; Short, H., Group theory from a geometrical viewpoint, *Notes on word hyperbolic groups*, World Scientific, Singapore, 1990, pp. 3–63.

[Br1999] M. R. Bridson, *Groups St Andrews 1997 in Bath, I*, Cambridge Univ. Press, Cambridge, 1999.

[BrHae2011] M. R. Bridson; A. Haefliger, *Metric spaces of non-positive curvature*, Springer, Berlin, 2011.

[Can1989] J. W. Cannon, *Lectures at Trieste*, Proceedings of ICPT meeting, Trieste, 1989.

[Dehn1911] M. Dehn, *Uber unendliche diskontinuierliche Gruppen*, Math. Annalen, Vol. 71(1), 1911, pp. 116–144.

[Dehn1987] M. Dehn, *Papers on group theory and topology*, translated and introduced by John Stillwell, Springer Verlag Berlin, Heidelberg, New York, 1987.

[DomAn1985] B. Domanski, and M. Anshel, *The complexity of Dehn's algorithm for word problems in groups*, J. of Algebra **6**, 1985, pp. 543–549.

[vonDyck1882] Walter von Dyck, *Gruppentheoretische Studien*, Math. Annalen, Vol. 20(1), 1882, pp 1–44.

[Gerst1991] S. M. Gersten, *Isoperimetric and isodiametric functions of finite presentations*, Geometric group theory, Vol. 1, London Math. Soc. Lecture Note Ser., 181, Cambridge University Press, Cambridge, 1993, pp. 79–96.

[GhHar1990] E. Ghys, and P. de la Harpe (editors), *Sur les groupes hyperboliques d'apres Mikhael Gromov*, Birkhäser, Progress in Mathematics Series, vol. 83, 1990.

[Gr1987] M. Gromov, *Hyperbolic Groups*, Essays in group theory, MSRI series vol. 8, edited by S. M. Gersten, SpringerVerlag, 1987, pp. 75–263.

[Har1969] F. Harary, *Graph Theory*, Reading, Mass: Addison-Wesley Pub. Co, 1969.

[Holt2001] Holt, D. F., and Rees S., *Solving the Word Problem in Real Time*, Journal of the London Mathematical Society **63**, 2001, pp. 623–639.

[LynSch2002] Lyndon, Roger C., and Paul E. Schupp, *Combinatorial Group Theory*, Springer-Verlag, Berlin, 2002.

[Mül2017] Müller-Stach, S., *Max Dehn, Axel Thue, and the Undecidable*, ArXiv e-prints, arXiv:1703.09750, Mar 2017.

[Rot1995] Joseph J. Rotman, *An Introduction to the Theory of Groups*, Springer-Verlag, Berlin, 1995.

[Sk1923] Thoralf Skolem, *The foundations of elementary arithmetic*, 1923, translated by Jean van Heijenoort in From Frege to Gdel: A Source Book in Mathematical Logic, 1879-1931. Harvard Univ. Press, 1967, pp. 302–333.

[Tietze1908] H. Tietze, *ber die topologischen Invarianten mehrdimensionaler Mannigfaltigkeiten*, Monatsschr. Math. Phys., 19, 1908, pp. 1–118.

[TikZ2013] Till Tantau, *The TikZ and PGF Packages*, http://sourceforge.net/projects/pgf/, 2013-12-20.